PROJECT FINAL REPORT

Grant Agreement number: 257528

Project acronym: KHRESMOI

Project title: Knowledge Helper for Medical and Other Information users

Funding Scheme: IP

Period covered: from 1.9.2010 to 31.8.2014

Name of the scientific representative of the project's co-ordinator[1], Title and Organisation:

Mr. Alexandre Cotting
HES-SO
TechnoArk 3
3960 Sierre, Switzerland

Tel: +41 27 606 90 16

Fax: +41 27 606 90 00

E-mail: alexandre.cotting@hevs.ch

Project website address: http://khresmoi.eu

[1] Usually the contact person of the coordinator as specified in Art. 8.1. of the Grant Agreement.

Table of Contents

4.1. FINAL PUBLISHABLE SUMMARY REPORT 3

4.1.1. Executive Summary 3

4.1.2. Project Context and Objectives 4

4.1.3. Main Results 6

4.1.3.1 Khresmoi Prototypes 7
4.1.3.2 Large Scale Data-Driven Image Search and Classification 9
4.1.3.3 Accessible Semantic Search for Linking Multiple Data Sources 10
4.1.3.4 Domain Adaptation for Machine Translation 12
4.1.3.5 Flexible and Adaptive Search Interface 13
4.1.3.6 Integrated System 13
4.1.3.7 Holistic Multi-Component System and User-Centred Evaluation 15

4.1.4. Potential Impact 17

4.1.4.1 Societal Impacts 17
4.1.4.2 Dissemination Activities 18
4.1.4.3 Exploitation 20
4.1.4.4 Impacts on the Consortium 21

4.1.5. Contact Details 24

4.2. USE AND DISSEMINATION OF FOREGROUND 26

4.3. REPORT ON SOCIETAL IMPLICATIONS 26

4.1. Final publishable summary report

4.1.1. Executive Summary

Imagine that a radiologist is examining a computer tomography image and sees an anomaly that he has never come across before. A commonly adopted approach to identifying the anomaly is to page through books to try and find a similar-looking anomaly, or to ask a colleague if she knows what the anomaly is. With Khresmoi, he can do an automated search for similar-looking anomalies in cases in the hospital digital archives. He can then use anonymised reports written by his colleagues attached to the similar cases to guide and support his diagnosis. Based on a text analysis of the returned reports, Khresmoi also launches a search in the medical literature. Even if the returned radiology reports are in German, Khresmoi can access relevant publications in the much larger set of English medical publications. In effect, the radiologist gains immediate access to a huge amount of implicit medical knowledge of his colleagues and relevant information from the medical literature without entering a single search keyword.

Four years ago, such a scenario was not possible. Now that the Khresmoi project has ended, the *Khresmoi Radiology* prototype returns all of the information listed above in less than 4 seconds after the radiologist has marked an anomalous region in an image and pressed the search button.

Khresmoi Radiology is the most exciting result to come out of the Khresmoi project, which ended on the 31st of August 2014 after four years of research and development in the area of medical information search and retrieval. *Khresmoi Radiology* integrates the largest number of technologies developed in Khresmoi into a single system. Khresmoi has developed a data-driven approach to visual similarity search in 3D medical images, taking advantage of the terabytes of medical images stored in hospital archives. As the approach is data-driven, it can be applied anywhere in the human body without the necessity for careful hand-crafting of techniques specific to various organs. Automated semantic annotation of the text of the radiology reports finds mentions of organs and medical conditions, which are used to propose a consensus diagnosis from a group of retrieved radiology reports and automatically construct a query for a search in the medical literature. A key characteristic has been the use of standard medical terminologies in both text and image analysis, allowing straightforward combination of the analysis results for both text and image modalities. Methods for large-scale search and classification of images in the medical literature have also been developed in Khresmoi.

The Khresmoi technology is also used to provide powerful access to medical information through text search. The semantic annotation linking documents to medical terminologies allows the user to easily locate relevant documents even if they only contain synonyms of the words used in the query. But even more powerful queries are possible, such as finding documents that mention medication used to treat diabetes, or documents mentioning diseases that have a dry cough as a symptom. The medicine-specific machine translation allows users to search English documents while entering a query in the language that they are most comfortable with. This currently works for French, German and Czech, but the statistical machine translation approach adopted makes the extension to further languages relatively straightforward. The prototypes specialised on text search have been optimised for two end user groups: *Khresmoi Professional* is designed for medical professionals, while *Khresmoi for Everyone* is easy to use by all. *Khresmoi for Everyone* puts particular emphasis on ensuring the trustworthiness of the websites presented in the search results.

A particular strength of Khresmoi has been the involvement of medical professionals and patients from the design phase to the testing phase, influencing all aspects of the project. In particular, all prototypes have been tested by the actual members of the target user group. *Khresmoi Radiology* has been evaluated by radiologists in hospitals in Austria, Germany, Switzerland and Greece; *Khresmoi Professional* has been tested by medical doctors while they attended symposia in Austria and Germany; while *Khresmoi for Everyone* has been tested by a diverse group of members of the general public in France, Switzerland and the Czech Republic, including patients in a hospital in France.

Around 50 people from 12 organisations worked together over four years to develop this innovative technology and produce new research results, while gaining invaluable experience in areas ranging from system integration to international cooperation. Young researchers have earned their PhD degrees, post-doctoral researchers have taken their first steps toward independent research, and more senior staff have overcome the organisational challenges presented by such a large-scale multinational research and development project. But what happens now that the Khresmoi project is over? Are we going to turn off the servers and disappear? Not if we can prevent it. There are two initiatives to further develop Khresmoi technologies. One has the target of making the Khresmoi medical text annotation, semantic search and machine translation available as commercial-grade web services and to adapt these technologies to patient record processing, while the other initiative deals with bringing the *Khresmoi Radiology* technology to the market. Keep watching the Khresmoi web page for updates on these initiatives.

Khresmoi Prototypes

- Khresmoi for Everyone: http://everyone.khresmoi.eu
- Khresmoi Professional: http://professional.khresmoi.eu
- Khresmoi Radiology: http://radiology.khresmoi.eu

4.1.2. Project Context and Objectives

The Khresmoi project addressed the challenges of searching through large amounts of radiology data, including Magnetic Resonance (MR) and Computed Tomography (CT), in hospital archives, as well as general medical information available on the internet. For the latter, it addressed the issues of trustworthiness and readability levels of the documents. The project consortium, consisting of twelve partners from nine European countries, developed a multilingual multimodal search and access system for health information and documents. The system allows text querying in several languages, in combination with image queries. It returns translated document summaries linked to the original documents. Khresmoi started on the 1st of September 2010 and ran for four years. In summary, the objectives are Khresmoi were:

- Effective automated information extraction from biomedical documents, including improvements using crowd sourcing and active learning, and automated estimation of the level of trust and target user expertise

- Automated analysis and indexing for medical images in 2D (X-Rays), 3D (MR, CT), and 4D (MR with a time component)

- Linking information extracted from unstructured or semi-structured biomedical texts and images to structured information in knowledge bases

- Support of cross-language search, including multilingual queries, and returning machine-translated pertinent excerpts

- Adaptive user interfaces to assist in formulating queries and display search results via ergonomic and interactive visualizations

The research flowed into several open source components, which were integrated into an innovative open architecture for robust and scalable medical information search.

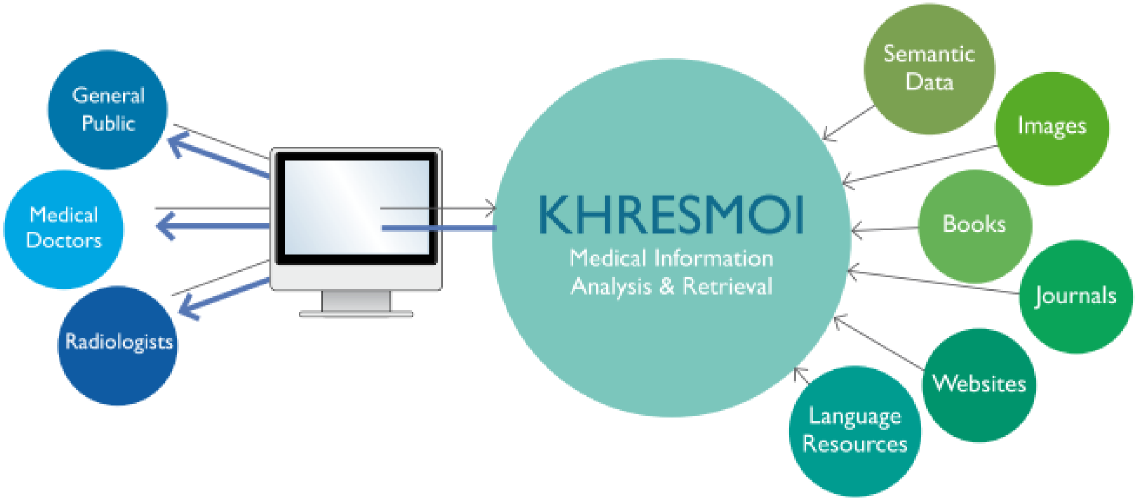

Figure 1: Khresmoi global overview

Khresmoi was evaluated in challenging use cases involving the following target user groups:

- **Members of the general public** want access to reliable and understandable medical information in their own language. At present, web search engines are the most-used tools for finding medical information on the internet, but the web pages returned are of varying quality, with no indication of the reliability of the information.

- **Clinicians and general practitioners** need accurate answers rapidly – a search on PubMed requires on average 30 minutes,[2] while clinicians typically have 5 minutes available.[3]

[2] W. R. Hersh, D. H. Hickam, How Well Do Physicians Use Electronic Information Retrieval Systems? A Framework for Investigation and Systematic Review, Journal of the American Medical Association, Vol 280, No. 15, 1998

[3] A Hoogendam, A. F. H. Stalenhoef, P. F de Vries Robbé, A. J. P. M. Overbeke, Answers to Questions Posed During Daily Patient Care Are More Likely to Be Answered by UpToDate Than PubMed, J Med Internet Res, Volume 10, Number 4, 2008.

- **Radiologists** are drowning in images and need improved automated support for their analysis – at larger hospitals over 100GB of images are produced per day. The huge archives of radiology images available in hospitals (in anonymized form) have a large potential to assist radiologists with diagnosis if search by visual similarity in these archives were possible.

4.1.3. Main Results

The Khresmoi project developed search technologies specifically for the medical domain. These include semantic search, machine translation, image search, search interfaces, and medical knowledge bases. The technologies were integrated into three prototypes each aimed at a different group of end users:

- *Khresmoi for Everyone* is aimed at members of the general public
- *Khresmoi Professional* is aimed at physicians
- *Khresmoi Radiology* has 3D image search features of particular use to radiologists

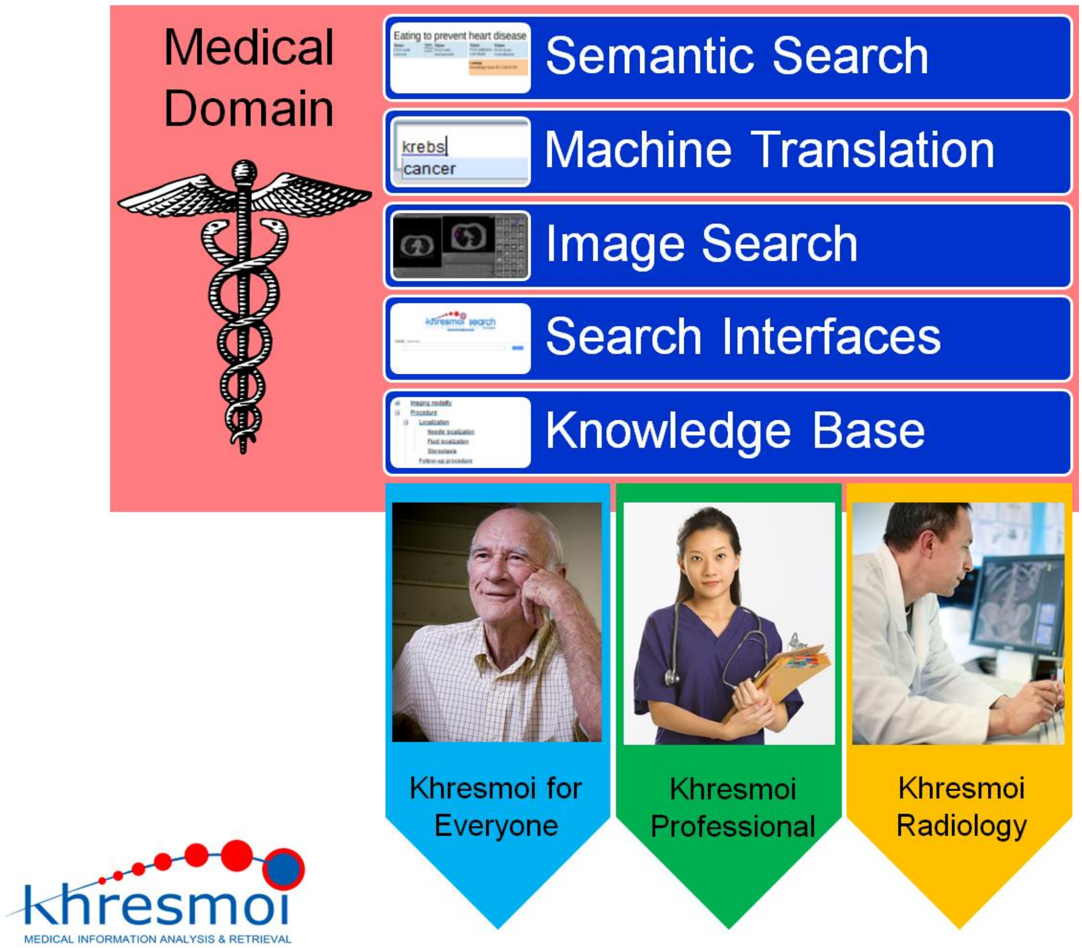

Figure 2: Overview of Khresmoi core achievements

The remainder of this section presents the main results of the Khresmoi project, starting with the prototypes, then presenting some of the components making up the prototypes. Finally, the integration and the evaluation outcomes are presented.

4.1.3.1 Khresmoi Prototypes

There are three Khresmoi prototypes, all based on different combinations of the same basic components. Each prototype meets the requirements of one of the target groups of end users. The three prototypes are:

- **Khresmoi for Everyone:** This prototype presents a straightforward search interface aimed at members of the general public. It also has features specific to the medical domain developed in Khresmoi, such as medicine-specific machine translation and automated estimation of the trustability and readability levels of documents. This prototype is shown in Figure 3. The red or green bar to the left of each result in the result list indicates the estimated readability level, while the scale to the right of each result presents the estimated trustability level of the website. Translation and filtering options are available on the right of the window.

- **Khresmoi Professional:** This prototype, shown in Figure 4, is aimed at medical professionals. The interface is more comprehensive, and allows results to be stored in a personal library, rated and shared with colleagues. Support for medicine-specific machine translation and 2D image search based on visual similarity are also available. Various facets classifying the results are shown on the left of the window.

- **Khresmoi Radiology:** This prototype, shown in Figure 5, makes available the advanced visual search capabilities required by radiologists. It allows search by visual similarity in 3D images (CT, MRI, …) stored in a hospital Picture Archiving and Communication System (PACS), as well as in 2D images in the medical literature. A region of an image can be chosen (on the left in Figure 5), and the system will present the most similar images from the PACS (on the right in Figure 5). Search results and associated radiology reports can be viewed, where the relevant medical terms are highlighted in the radiology reports. Analyses of the texts in the radiology reports accompanying the search results allow the most commonly mentioned pathologies in the radiology reports to be identified, and these are used to automatically create a query to search the medical literature. Machine translation techniques allow the English literature to be searched, even if the radiology reports are in German.

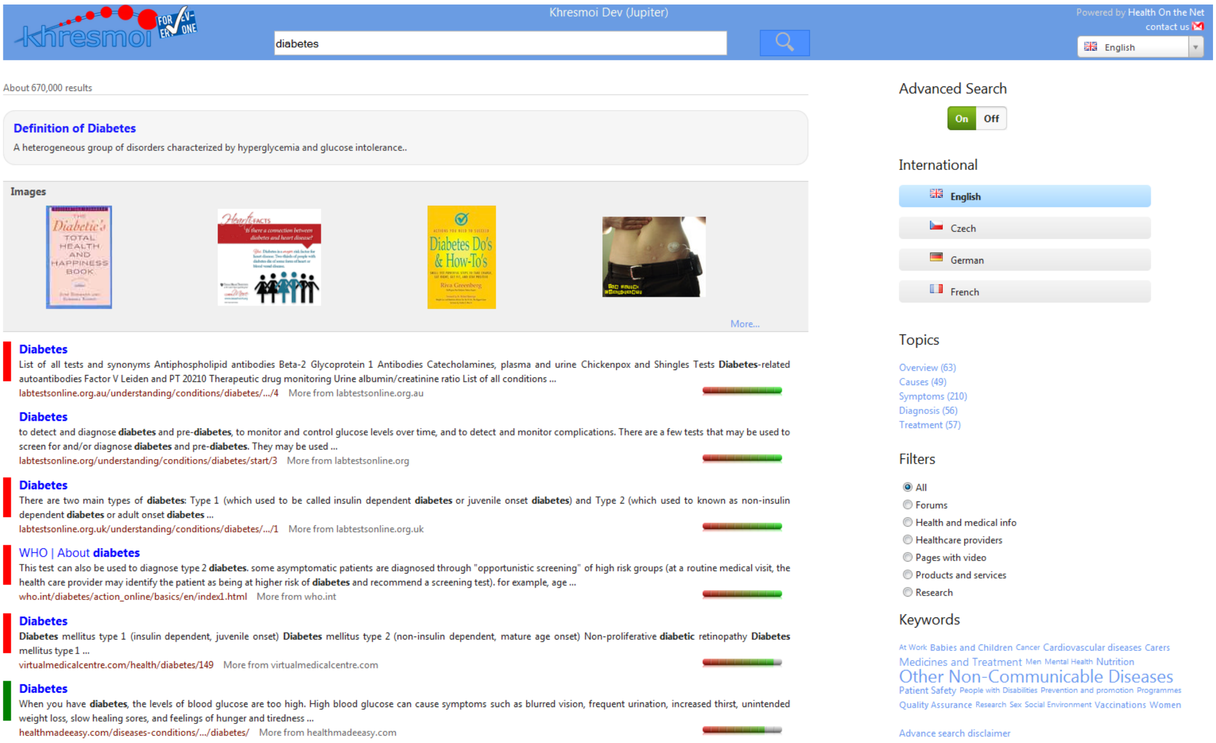

Figure 3: Khresmoi for Everyone

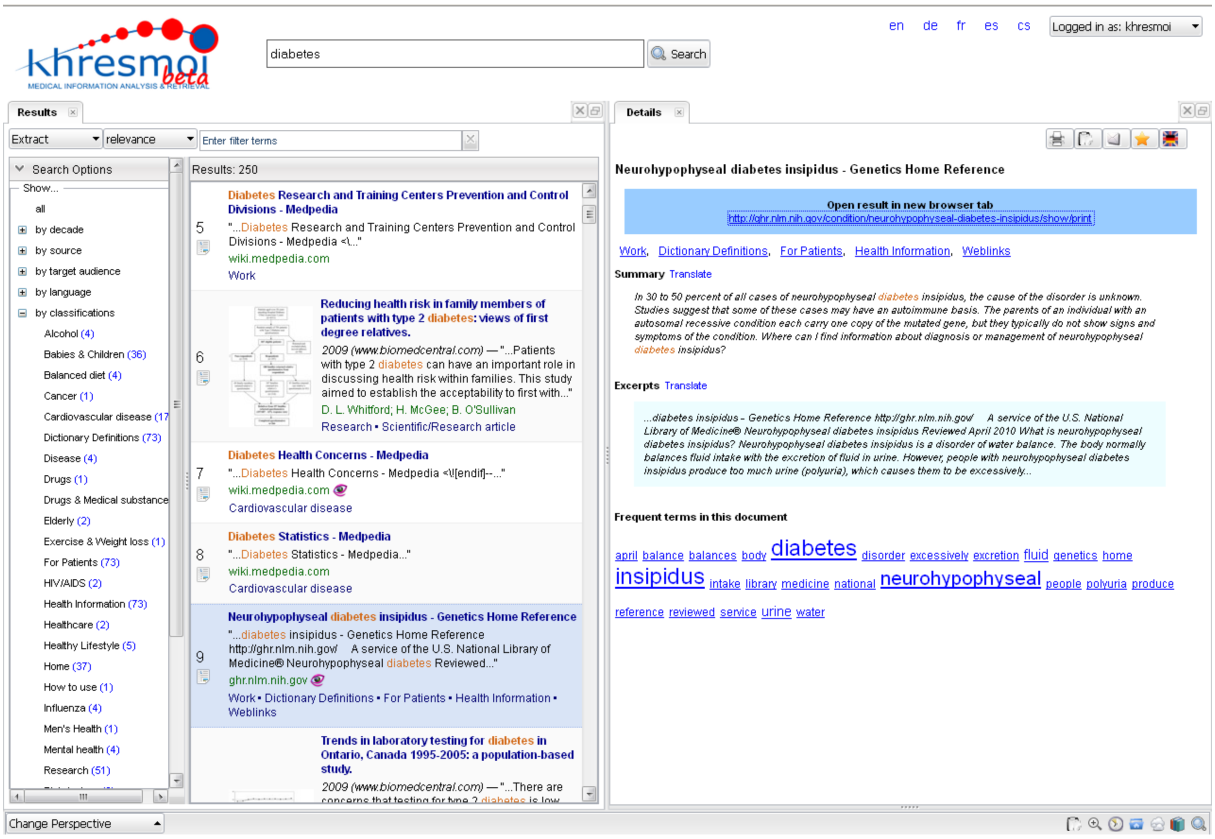

Figure 4: Khresmoi Professional

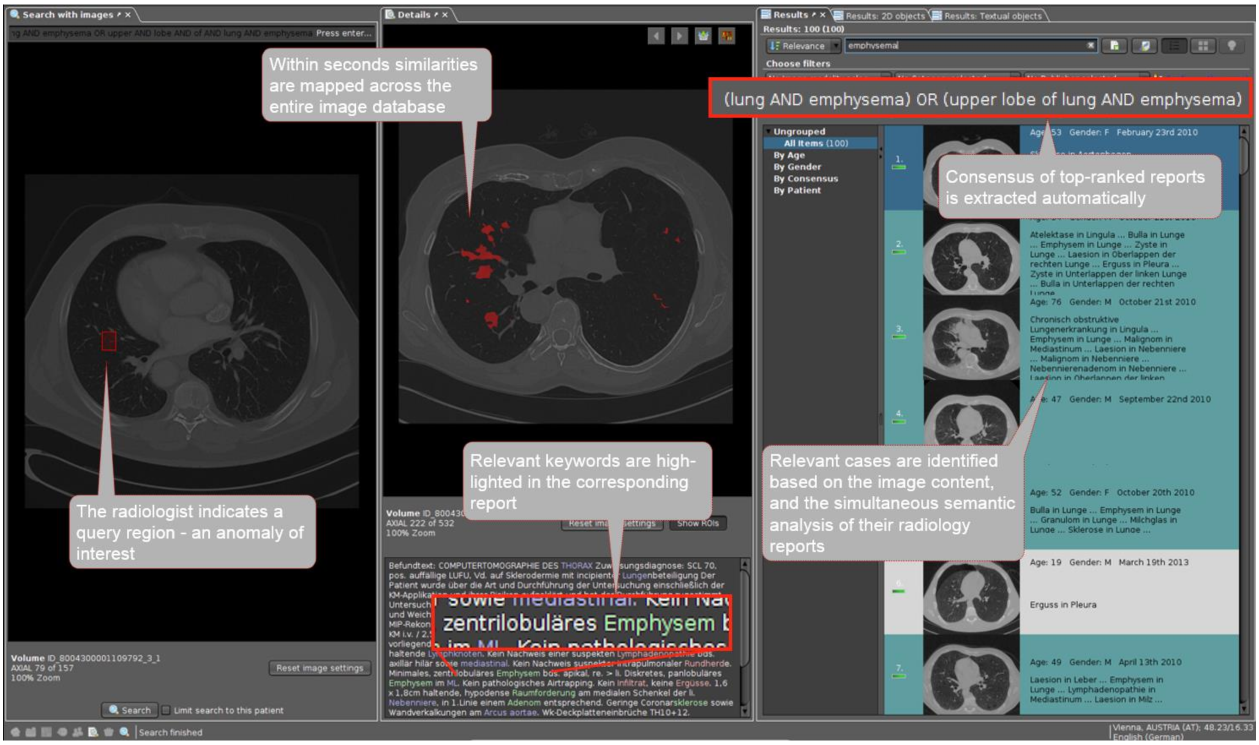

Figure 5: Khresmoi Radiology

4.1.3.2 Large Scale Data-Driven Image Search and Classification

The Khresmoi project has adopted a data-driven approach to image analysis and search in the medical domain. This is possible due to the large amounts of data that were available for processing and analysis in the project, both from a hospital PACS and from the medical literature. Such a data-driven approach is advantageous as it avoids having to manually tune image analysis and search techniques to particular areas of the body – techniques can use machine learning approaches to learn from the sufficiently large number of examples available.

3D Image Search and Analysis

When a user indicates a region of interest in an imaging volume such as a CT, and starts the retrieval, results of similar regions across thousands of cases are now shown within 4 seconds. During this time the visual features of the query region are compared to millions of indexed regions, the most similar regions are identified, and imaging volumes are ranked based on the configuration of those regions. To provide the user with most informative feedback when browsing the search results, image thumbnails that show the relevant portion of the image are rendered. Overall result accuracy is much improved, and the system now accurately identifies similar anomaly patterns across images and patients. The improved accuracy is due to advanced feature extraction and learning methods developed and incorporated into the prototype. The speed of retrieval is due to new indexing algorithms that make the visual information of many millions of image segments comparable within seconds. This is not trivial, since the necessary information cannot be held in the memory, and intelligent query strategies are necessary to ensure speed, and at the same time minimise deviation of distance estimates encoded in the index from the actual distance between examples.

2D Image Search and Analysis

Khresmoi technology also allows images from the medical literature to be searched by visual similarity. The capability to automatically separate compound figures into their constituent sub-figures was an extremely useful addition to image search. To allow for more focused search, for all images the image type or modality was determined automatically and several filters allow the search results to be restricted, for example only to radiology modalities, which account for approximately 20% of the images but are of high interest for our target group, the radiologists. It is also possible to perform keyword and visual search together — this allows images similar to an example that also contains specific keywords in the caption to be found. The speed of the search has been improved both by improving the search algorithms and by using the private cloud infrastructure.

Open Source Outcomes

The outcomes of the 2D image search research and development are implemented in the ParaDISE open source software, with versions available under both the Apache Software Licence 2.0 and the GPL v3 licence. The software can be downloaded from: http://paradise.khresmoi.eu/

4.1.3.3 Accessible Semantic Search for Linking Multiple Data Sources

Semantic Text Annotation and Search

Mimir (from Norse mythology, "The Rememberer"), is a multi-paradigm information management index and repository which can be used to index and search over text, annotations, semantic schemas (ontologies), and semantic meta-data (instance data). Khresmoi created indexes to medical texts that can take search beyond retrieving those documents that match the words of a user's query. Khresmoi uses semantic annotation to find and mark those words and phrases in texts that match complex concepts in the myriad of databases, vocabularies, and ontologies that describe biomedical knowledge. Queries can then be written across both the texts and these knowledge bases. We could, for example, ask to pull back all texts that talk about drugs used in the treatment of malaria. The facts of which drugs treat malaria are retrieved from the knowledge bases, and then the mentions of the individual drugs are retrieved from the text of documents. Mimir allows queries that arbitrarily mix full-text, structural, linguistic and semantic queries and that can scale to gigabytes of text.

A semantic type-ahead interface was developed to ease the entry of semantic queries. Four steps in entering such a query are shown in Figure 6. During the entry of a query, the system queries the knowledge base to obtain query completion suggestions that are coherent with the current state of the query. Step 4 in Figure 6 shows the final query, which requests documents mentioning diseases or syndromes have the symptom of a dry cough. When the query is submitted, the system queries the knowledge base for a list of relevant diseases, and then retrieves documents mentioning these diseases. A list of some of the documents retrieved is shown in Figure 7, with the diseases highlighted in bold. Diseases mentioned include gastroesophageal reflux disease, pleuritic, laryngitis and scleroderma.

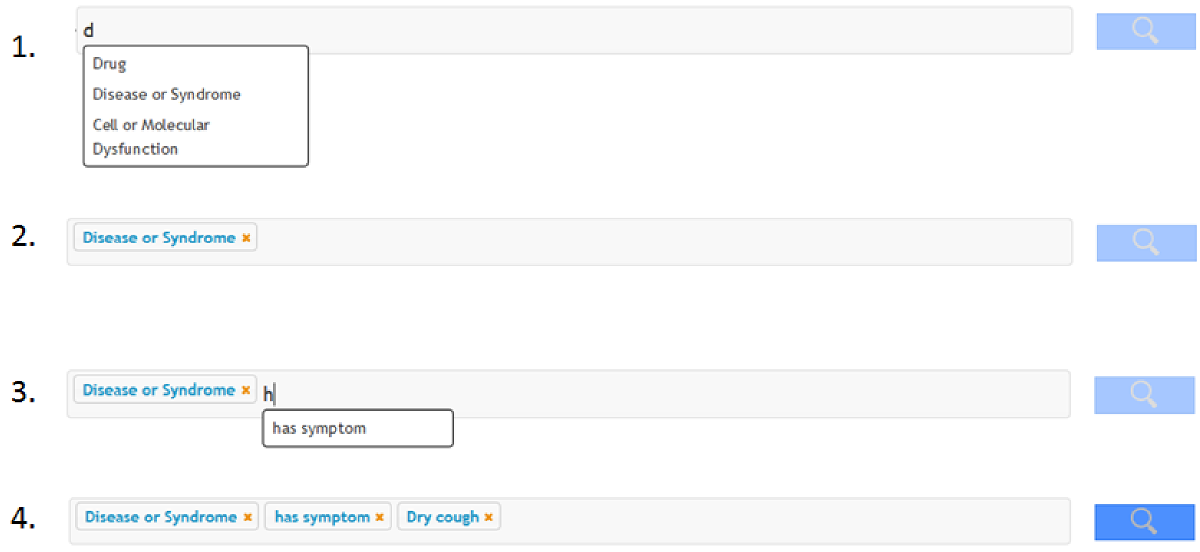

Figure 6: Four steps in entering a semantic query. Query suggestions are obtained from the knowledge base.

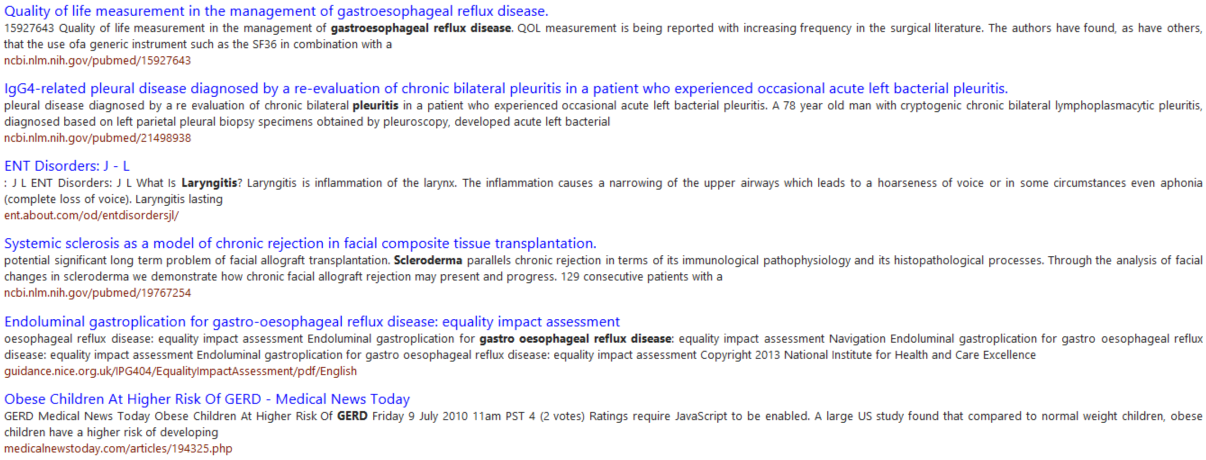

Quality of life measurement in the management of gastroesophageal reflux disease.
15927643 Quality of life measurement in the management of **gastroesophageal reflux disease**. QOL measurement is being reported with increasing frequency in the surgical literature. The authors have found, as have others, that the use ofa generic instrument such as the SF36 in combination with a
ncbi.nlm.nih.gov/pubmed/15927643

IgG4-related pleural disease diagnosed by a re-evaluation of chronic bilateral pleuritis in a patient who experienced occasional acute left bacterial pleuritis.
pleural disease diagnosed by a re evaluation of chronic bilateral **pleuritis** in a patient who experienced occasional acute left bacterial pleuritis. A 78 year old man with cryptogenic chronic bilateral lymphoplasmacytic pleuritis, diagnosed based on left parietal pleural biopsy specimens obtained by pleuroscopy, developed acute left bacterial
ncbi.nlm.nih.gov/pubmed/21498938

ENT Disorders: J - L
: J L ENT Disorders: J L What Is **Laryngitis**? Laryngitis is inflammation of the larynx. The inflammation causes a narrowing of the upper airways which leads to a hoarseness of voice or in some circumstances even aphonia (complete loss of voice). Laryngitis lasting
ent.about.com/od/entdisordersjl/

Systemic sclerosis as a model of chronic rejection in facial composite tissue transplantation.
potential significant long term problem of facial allograft transplantation. **Scleroderma** parallels chronic rejection in terms of its immunological pathophysiology and its histopathological processes. Through the analysis of facial changes in scleroderma we demonstrate how chronic facial allograft rejection may present and progress. 129 consecutive patients with a
ncbi.nlm.nih.gov/pubmed/19767254

Endoluminal gastroplication for gastro-oesophageal reflux disease: equality impact assessment
oesophageal reflux disease: equality impact assessment Endoluminal gastroplication for **gastro oesophageal reflux disease**: equality impact assessment Navigation Endoluminal gastroplication for gastro oesophageal reflux disease: equality impact assessment Endoluminal gastroplication for gastro oesophageal reflux disease: equality impact assessment Copyright 2013 National Institute for Health and Care Excellence
guidance.nice.org.uk/IPG404/EqualityImpactAssessment/pdf/English

Obese Children At Higher Risk Of GERD - Medical News Today
GERD Medical News Today Obese Children At Higher Risk Of **GERD** Friday 9 July 2010 11am PST 4 (2 votes) Ratings require JavaScript to be enabled. A large US study found that compared to normal weight children, obese children have a higher risk of developing
medicalnewstoday.com/articles/194325.php

Figure 7: Results of the query shown in step 4 of the semantic querying process shown in Figure 6.

Semantics and visual information

Aside from improving speed and accuracy, the integration of visual and semantic information proved to open very exciting possibilities for which we have even now only scratched the surface. The semantic information corresponding to the search results from a visual query, such as radiology reports, is mined to generate a summary of the observations made in the majority of reports. The mining algorithms “understand” the meaning of words and their categories based on terminologies such as RadLex. The prototype engine identifies relationships, and is able to determine that a report states that a certain anomaly was observed in a certain anatomical region. This allows for summarising and analysing the entire set of reports retrieved together with the top ranked search results. The structuring of the search result list gains from the information known about individual data. For example, it relates the individual report of a case in the index to the other reports in the search result list. This allows the system to provide simple tagging of results into a consensus set, and a set that might be relevant for differential diagnosis. Based on the consensus diagnosis extracted semantically from the reports, the system queries databases such as PubMed and the educational literature, which can provide representative examples and explanations for observations that might correspond to the query case. This happens before the user has even entered a single keyword. Now the search can be tuned based on textual user input, and the radiologists are enabled to further explore the relevant context of their query.

Knowledge Server

A Large Scale Biomedical Knowledge Server – a semantic warehouse – has been created that makes structured biomedical data available to other components of the Khresmoi system via a set of services. The data layer of the Knowledge Server, referred to as the Knowledge Base, integrates several data sources, including MeSH, RadLex, Drugbank and DBPedia, and includes new links between the data sources. This integration produced more than 1.2 billion facts (RDF statements). The repository contains ontological schema definitions (per data source) and instance data in the form of logical relationships between entities and language resources such as labels and descriptions. The OWLIM technology from Ontotext is used to implement the Knowledge Server.

Open Source Outcomes

The technologies developed in Khresmoi are implemented in the GATE tools, which are available from: http://gate.ac.uk

4.1.3.4 Domain Adaptation for Machine Translation

Machine Translation (MT) provides the cross-lingual capability to search in biomedical documents indexed by Khresmoi. The Khresmoi project made significant advances in adapting Machine Translation to the medical domain for both documents and queries. MT adaptation techniques were employed for the translation proper as well as on MT adaptation to improve cross-lingual information retrieval. The MT service was tuned and evaluated on medical domain-specific test sets carefully created for these specific purposes (and made publically available). The central language used for indexing and searching in Khresmoi is English. The MT service allows 1) to translate non-English user queries to English and 2) to present summaries of search results returned to the user from English to a chosen language. The non-English languages supported by Khresmoi are Czech, German, and French. The MT component is based on the Phrase-based Statistical Machine Translation system Moses and employs MTMonkey, a scalable infrastructure for MT among multiple languages developed within the Khresmoi project and published as open-source. The techniques developed have the potential to be used for adapting other languages to the medical domain. For texts in the medical domain, the Khresmoi Machine Translation system outperformed the best freely available MT services on the web, including Google Translate and Bing Translator, in terms of automatic metrics of translation quality.

Open Source Outcomes

MTMonkey, a distributed infrastructure for Machine Translation web services, has been released as open source under the Apache 2.0 license (https://github.com/ufal/mtmonkey). The MOSES statistical machine translation software (http://www.statmt.org/moses/) has been adapted to machine translation in the medical domain by extensive training on domain-specific texts in English, German, French and Czech.

4.1.3.5 Flexible and Adaptive Search Interface

ezDL (short for "easy access to Digital Libraries") is the basis of the Khresmoi search interface. Originally designed for digital libraries, it has been extended to provide an interactive user interface to a large variety of search systems such as Khresmoi. The system allows for many different clients (such as a Java webstart application, an AJAX application, or an Android app) using common back-end services for user authorization, query conversion, or search suggestions. ezDL is used as the user interface framework for the Khresmoi Professional and Khresmoi Radiology prototypes.

The most noticeable improvements developed as part of the Khresmoi project are the collaborative and organizing functionalities. The ezDL interface allows registered users to keep and organize results beyond a search session within a personal library. Saved results can be sorted or grouped by author, publication year, title and date of addition. In addition, it is possible to apply a filter to the personal library to quickly find a stored document – and of course the documents in the library can be exported and printed. Tags can be used to organize the personal library according to a user's individual needs. For example, a physician could add tags corresponding to a specific case or patient that he is working on. Tags can be used to group the documents in the personal library.

The personal library not only allows for personal information management, but can also be used to collaborate with other users. Documents that have been stored in the personal library (including new documents uploaded by the user) can be shared. To facilitate collaboration and sharing, users can create a personal profile. The privacy setting allows users to control if they want to be found by other users. The search functionality of the interface can be used to search for users based on their name or the description used on their profile. To further support collaboration users can create their own personal contact and sharing lists, or public groups around a specific topic. These allow for easier sharing of documents and discussions.

Another new functionality of ezDL developed in the Khresmoi project is the support for image search by providing positive and negative examples. For search systems that allow similarity search, the "image search" perspective of ezDL, shown in Figure 8, can be used to collect example images. The image examples can be used as positive or negative relevance feedback and allow for easy specification of queries that cannot be expressed through search terms. All previously found images can be used for searching, as can be images from a web browser, or even the local file.

Open Source Outcomes

ezDL is available as open source under the GPL v3 license. It can be downloaded from: http://ezdl.de

4.1.3.6 Integrated System

The three Khresmoi prototypes integrate all technology developed in Khresmoi. The prototypes now run on the Khresmoi Cloud, a private cloud made up of nine servers with one Terabyte of RAM and 28 Terabytes of storage.

The basic architecture of the Khresmoi system is based on common Serviced-Oriented Architecture (SOA) principles. The fundamental part is the logical view that should allow a modular and highly generic structure. For this, a three-tier approach was chosen. This means that the system is decomposed into three different layers which correspond to the following functional blocks: the application, the services, and the persistence of the system.

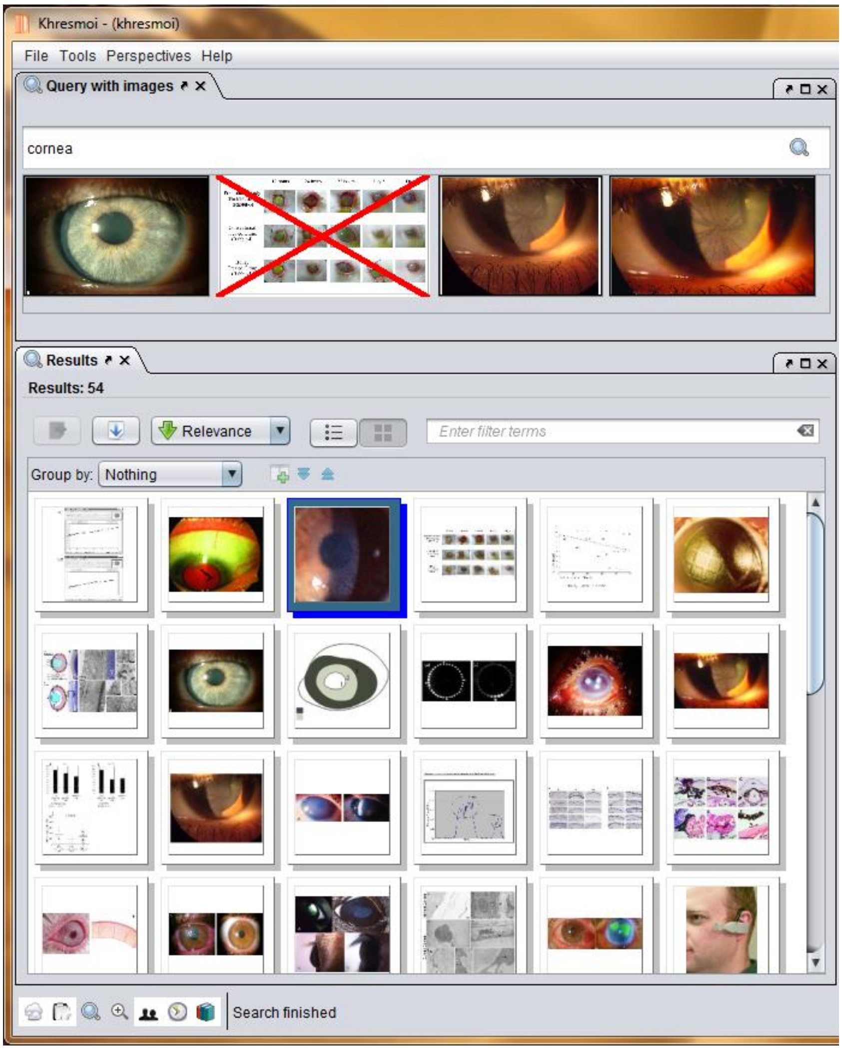

Figure 8: The ezDL image search perspective

The *Application Layer* is the application provided to the end users. According the different use cases defined in the project, different applications can be built to provide adapted user interfaces according to the specific requirements. The Application Layer deals with the configuration of the user interface, and the management of the user interaction to dispatch the events towards the internal system (Service Layer).

The *Service Layer* is the core of the system, as it contains all the main services provided by the system. These services are called Core Services and as they could be numerous and very different, they are grouped by Service Categories. Those services are atomic functions that can be called whenever needed by the system. They are specified with SCA and deployed through the runtime Apache Tuscany.

The last layer is dedicated to the system persistency, the *Persistence Layer*. It has in charge the mechanisms and models to store information. For each kind of information, a repository is required to store the data. Each repository provides a basic API to describe its own CRUD (Create, Read, Update and Delete) functionalities to permit easy access to the data.

Figure 9 shows a diagrammatic view of the components implemented in the integrated Khresmoi system, and how they interact. Each of the three prototypes uses a different combination of components to carry out its tasks. Some components, such as the 3D image analysis and search, is used in only one prototype. Other components, such as the machine translation, is used in all three prototypes.

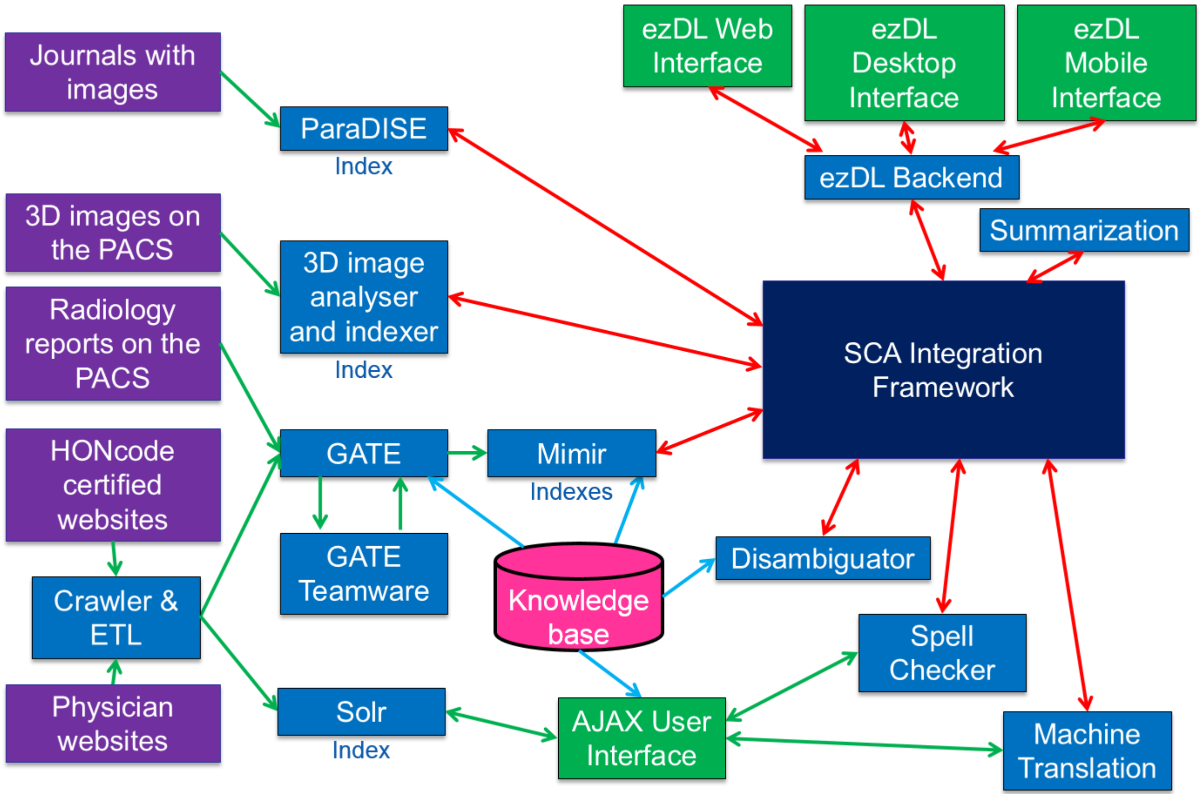

Figure 9: Khresmoi integrated components

4.1.3.7 Holistic Multi-Component System and User-Centred Evaluation

Evaluation Strategy

The creation of an integrated domain-specific search system as has been done in Khresmoi is a complex task requiring modelling of the domain and its users, as well as a specification of the system components required and their interactions. The evaluation of the performance of such a system is challenging, as it involves evaluation of multiple aspects:

- *Computational component-level* evaluations are computational evaluations of the system components taken in isolation;
- *Interactive component-level* evaluations involve an evaluation of components of the user interface and their back-end by end users;
- *Computational system-level* evaluations measure the performance of the full integrated system using a computational approach;
- *Interactive system-level* evaluation involves evaluating the full system by getting end users to perform search tasks on the system in a laboratory-type setting;

In Khresmoi, an evaluation of the system from the point of view of these four aspects was carried out. As a search system is being evaluated, the *performance* is made up of many facets, including: retrieval performance, user satisfaction and efficiency.

A distinguishing characteristic of the Khresmoi project was its implementation of a global coordinated evaluation strategy. An independent evaluation strategy was created near the beginning of the project, which gave recommendations on the evaluations to be carried out in the individual work packages. After the first round of evaluations was complete, a meta-analysis of these results was done, in which the reported results of the evaluations performed were compared to the recommendations in the evaluation strategy. Based on the results of the meta-analysis, an updated evaluation strategy, including approaches to solve the identified shortcomings, was presented. Finally, after the second round of evaluations was complete, a second meta-analysis of the results was done.

End Users

Evaluation of search systems are often not conducted with “real” end users, but with surrogates such as students, who are more readily available than busy professionals. In Khresmoi, we placed a significant emphasis on evaluating the developed prototypes with actual end users. For the evaluation of the final *Khresmoi for Everyone* prototype, 63 members of the general public participated, including patients in a hospital in Paris, France. For encouraging physicians to participate, the technique of conducting the evaluations at booths at medical symposia was adopted (Figure 10), as this allowed access to a larger number of physicians, even though the amount of time that they could spend on doing the evaluation was reduced. Overall, 55 physicians took part in the evaluation of the final *Khresmoi Professional* prototype. Evaluations of the *Khresmoi Radiology* prototype took place in in four hospitals (Medical University of Vienna, Austria; University Hospitals of Geneva, Switzerland; University Hospital of Freiburg, Germany; and University Hospital of Larissa, Greece), with 26 radiologists conducting the evaluations.

Extensive resources for carrying out user-centred evaluations of medical search systems were created in Khresmoi, including the experimental protocols and realistic search tasks for all target groups.

Figure 10: Evaluation of Khresmoi Professional at the STAFAM in Graz, Austria

CLEF eHealth

In 2013 and 2014, members of the Khresmoi consortium were organisers of the CLEF eHealth evaluation lab. The lab is held as part of the Conference and Labs of the Evaluation Forum (CLEF). The first edition of CLEF eHealth, in 2013, included three evaluation tasks: (1) Named entity recognition and normalization of disorders; (2) Normalization of acronyms/abbreviations; and (3) Information retrieval to address questions patients may have when reading clinical reports. Task 3 was managed by members of the Khresmoi consortium, in collaboration with the University of Turku (Finland), CSIRO and NICTA (Australia). The datasets included a document crawl provided by Khresmoi, queries manually built by the nursing group at the University of Turku, and relevance judgements provided by this group. 175 people registered their interest in the lab (64, 56 and 55 respectively for tasks 1, 2 and 3), and 53 teams participated (39, 5 and 9 respectively for tasks 1, 2 and 3). Teams participating included renowned groups from the clinical/medical natural language processing (NLP) and information retrieval (IR) domains. Through the official release of the 2013 task 3 dataset, more teams can use it and investigate new approaches to improve medical IR.

The 2014 edition of the lab also included three evaluation tasks: (1) Visual-Interactive Search and Exploration of eHealth Data; (2) Information extraction from clinical text; (3) User-centred health information retrieval. Again, task 3 was managed by members of the Khresmoi consortium, and a cross language subtask was added. The dataset was created in a similar manner to 2013. 224 people registered their interest in the lab (50, 79 and 55 respectively for tasks 1, 2 and 3), and 53 teams participated (1, 10 and 13 respectively for tasks 1, 2 and 3). The organizers and participants gathered at CLEF 2014 in Sheffield to report results for each task and learn from participants' presentations and posters.

4.1.4. Potential Impact

This section first covers the potential societal impacts of the Khresmoi project, then describes the dissemination activities that have taken place. Plans for exploitation of Khresmoi results are then presented, and finally the impact of the Khresmoi project on the members of the Khresmoi consortium is discussed.

4.1.4.1 Societal Impacts

Extensive studies were carried out during the Khresmoi project on the search behaviour and requirements for all three target groups: members of the general public, physicians in general, and radiologists. These were based on online surveys, interviews with end users, and information gathered during the user-centred evaluations. The results have been made available in public deliverables and in refereed publications. The deliverables covering the results of this work are the most often downloaded among all Khresmoi deliverables.

The Health on the Net Foundation, a partner in the Khresmoi project, certifies medical websites providing reliable information with the HONcode certification. Using technology developed in Khresmoi, HON has been able to improve the efficiency with which the certification, still a largely manual process, is done. The ability to certify websites efficiently is becoming ever more important with the recent sale of the ".health" domain and the concerns about the quality of websites that will use this domain.

4.1.4.2 Dissemination Activities

A total of 153 papers has been published in journals and conferences, based on work done in the Khresmoi project. One quarter of these papers are the result of joint work between two or more partners in the Khresmoi project. The full list of papers published is available online here: http://khresmoi.eu/resources/publications/

The Khresmoi project presented its results at multiple events. The most important events are outlined below.

CeBIT

CeBIT is the biggest computer fair in the world with a large and extremely varied participation from the entire world but in an important part from Germany. In 2013, Khresmoi participated at the CeBIT in a booth together with three other EU projects, while in 2014, Khresmoi participated with its own booth (Figure 11a). One goal of this participation was to present clearly the prototypes to a larger public and get feedback on the prototypes for the preparation of the final Khresmoi prototype evaluations. A second objective was to get commercial contacts and get linked to partners for the Khresmoi technology. Many discussions with companies led to technology exchange and several propositions to distribute the Khresmoi technology if products become available. The Khresmoi torso also helped to clearly brand Khresmoi as a medical project and this attracted interest of many passing persons.

European Data Forum

The European Data Forum 2014 (EDF2014) took place from March 19th to 20th 2014 in Athens, Greece. EDF is the annual meeting-point for data practitioners from industry, research, the public-sector and community initiatives, to discuss the opportunities and challenges of the emerging Data Economy in Europe and took place in the third edition in 2014. The Khresmoi project had a booth at the European Data Forum, where the three prototypes were demonstrated. We were also honoured to be able to present the Khresmoi results to Commissioner Neelie Kroes, Vice President of the European Commission responsible for the Digital Agenda for Europe (Figure 11b).

ICT 2013

Khresmoi had a booth in the exhibition section of the EU ICT 2013 event in Vilnius, Lithuania from November 6th to 8th 2013 (Figure 11c). The ICT is Europe's biggest digital technology and innovation event. Many useful contacts were made with potential adopters of the Khresmoi technology through the extensive discussions that took place at the booth.

World of Health IT

The World Congress of Health IT Conference & Exhibition is the premier forum for the advancement of IT in healthcare in Europe. To address the needs of key stakeholders in the community of eHealth in Europe, The World of Health IT Conference & Exhibition offers professional development sessions, suppliers exhibitions, exchange of best practices, networking sessions and debates and discussions concerning the issues that will shape the future of eHealth.

The Khresmoi project has a booth at the World of Health IT, held from April 2nd to 4th 2014 in Nice, France (Figure 11d). All prototypes were presented at the booth, and the Khresmoi team present also took part in a series of pre-arranged meetings with representatives of various companies attending the event.

European Congress on Radiology

Khresmoi results were presented at a booth at the IMAGINE exhibit of the European Congress on Radiology (ECR), the largest radiology congress in Europe that gathered over 20,000 participants from 102 countries, in 2011, 2012 and 2013. In 2013, Khresmoi participated in the IMAGINE exhibit, with a booth and a prototype demo during the entire congress duration (Figure 11f). An article on Khresmoi was also published in the *ECR Today* congress magazine. The IMAGINE exhibit is significant, since it not only aims at presenting applicable technology to radiologists, but also to communicate work in progress among the medical image analysis community. Both aspects are very valuable for Khresmoi. We could reflect on the applicability of the prototype with radiologists, while at the same time discussing methodological details among peers in the computer science field.

Participation in Medical Symposia

As part conducting the user-centred evaluation of the *Khresmoi Professional* prototype, Khresmoi was demonstrated at various events attended by physicians. This included the STAFAM, the biggest conference for general practitioners in Austria (Figure 10); the Praxis Update Wiesbaden, a medical Continuing Medical Education (CME) conference for practitioners; and multiple events organised by the Association of Physicians in Vienna.

Language Resources and Evaluation Conference

Khresmoi had a booth at the Language Resources and Evaluation Conference (LREC) conference in Reykjavik, Iceland, in 2014 and in Istanbul, Turkey in 2012 (Figure 11e). LREC is the major event on Language Resources and Evaluation for Language Technologies. The LREC conference covers Language Resources and their applications, evaluation methodologies and tools, industrial uses and needs, and requirements coming from the e-society, both with respect to policy issues and to technological and organisational ones. The booth allowed the Khresmoi results in the language technology domain to become known in the language technology community.

Medical Informatics Europe 2012

Khresmoi was present at the Medical Informatics Europe (MIE) Conference in Pisa, Italy from August 26th to 29th 2012. The project had a stand in the Village of the Future (Figure 11g), and a presentation was given in the Village of the Future session on *People and Expectations*. In this session, the scenario of Little Sam was considered. Sam is diagnosed with Cystic Fibrosis (CF) at an early age, and makes use of internet search engines to get information about the disease, and social networks and blogs to get into contact with fellow CF patients. The importance of access to trustable online medical information and the key role that search technology plays in this access was underlined in this session.

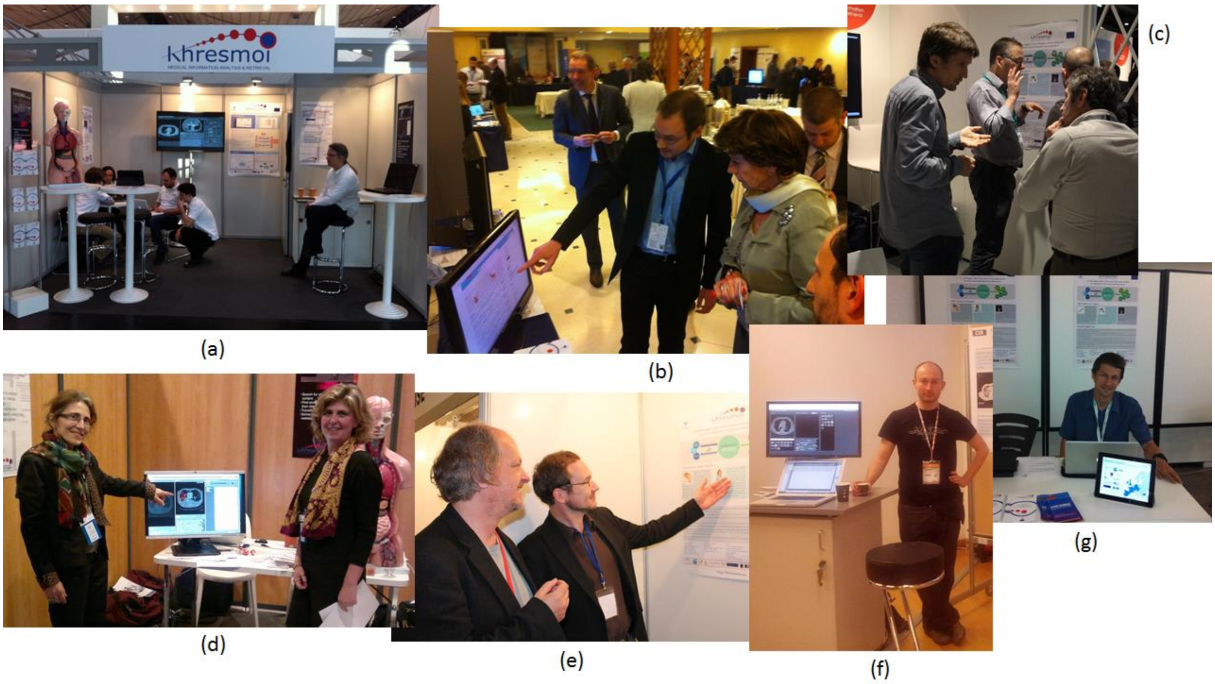

Figure 11: Khresmoi dissemination

4.1.4.3 Exploitation

Key Outcomes

There are two key outcomes of Khresmoi for which avenues of exploitation are currently being investigated:

- **Medical text analysis, retrieval and translation tools:** These tools cover the annotation, indexing, and machine translation of medical texts, as well as the analysis and machine translation of queries to a medical search system. They currently form the basis of many capabilities of all three Khresmoi prototypes. Plans for the exploitation involve providing these tools as commercial web services for use by companies analysing medical texts, and also to extend the tools with the capability to analyse medical records.

- **Radiology analysis and search:** The visual similarity search in 3D radiology images and the semantic linking between these images and the radiology report texts, demonstrated in the *Khresmoi Radiology* prototype, represent the most original outcomes of the Khresmoi project. Plans for the exploitation of these key outcomes are currently in preparation.

Software Outcomes

The software that Khresmoi is built upon has undergone significant advancement through work in Khresmoi. The software is listed below, along with the advances achieved in Khresmoi:

- **GATE** (https://gate.ac.uk/): The General Architecture for Text Engineering (GATE) is used to annotate at word, section and document levels. Through work in Khresmoi, its capabilities for annotating medical documents have been expanded. The use of cycles of human correction to improve the automatic annotation has also been extensively tested.

- **Mimir** (https://gate.ac.uk/mimir/) uses GATE annotations to perform semantic search. The Khresmoi Mimir Interface (KMI) has been developed to allow more user friendly querying of Mimir from Khresmoi. A semantic type-ahead service and corresponding interface has also been developed to allow straightforward semantic querying.
- **ezDL** (http://ezdl.de/) is a framework for interactive search applications. New features have been added, including drop down options for query specification, and automatic translation of non-English query terms if too few results are returned. It has also been made more stable and efficient. Three front-ends are now available for ezDL: the original Java Swing interface, a web interface and a mobile Android interface.
- **ParaDISE** (http://paradise.khresmoi.eu) is a new visual search engine developed in Khresmoi as a successor to the GNU Image Finding Tool (GIFT). It is more scalable than GIFT and contains state-of-the-art image features and visual similarity calculation.
- **MTMonkey** (https://github.com/ufal/mtmonkey) is a distributed infrastructure for Machine Translation web services. It allows a JSON-encoded request for different translation directions to be distributed among multiple MT servers.
- The **OWLIM** semantic repository (http://www.ontotext.com/owlim) has received performance and functionality upgrades, and has also had its medical knowledge base expanded through the addition of new medical vocabularies and new links between the medical vocabularies.

Data Sets

The following datasets have been created in Khresmoi, and are available for further use.

- **Annotated Radiology Images:** Within the VISCERAL project, 3D radiology images have been manually annotated and will be released to the research community by April 2015. Watch the VISCERAL website for more details (http://visceral.eu).
- **Multilingual Corpora from the Medical Domain:** Multilingual datasets for the translation of medical queries and for the translation of summaries of medical documents have been created and are available.
- **Evaluation of Information Retrieval in the Medical Domain:** The data used in the CLEF eHealth retrieval tasks in 2013 and 2014 are available through the ELRA catalogue. The 2014 dataset includes queries in multiple languages.
- **Medical Image Retrieval Evaluation:** The ImageCLEF medical task datasets from 2011, 2012 and 2013 are available by request.

4.1.4.4 Impacts on the Consortium

Around 50 people from 12 organisations worked together over four years on the Khresmoi project, while gaining invaluable experience in areas ranging from system integration to international cooperation. Young researchers have earned their PhD degrees, post-doctoral researchers have taken their first steps toward independent research, and more senior staff have overcome the organisational challenges presented by such a large-scale multinational research and development project. In order to elicit what the impacts on the consortium are, at the final full consortium meeting, we asked consortium members to write their lessons learned in the project on post-it notes, categorised into the following three categories: Scientific and Technical, Organisational, and Personal. An image of the raw collected lessons is shown in Figure 12.

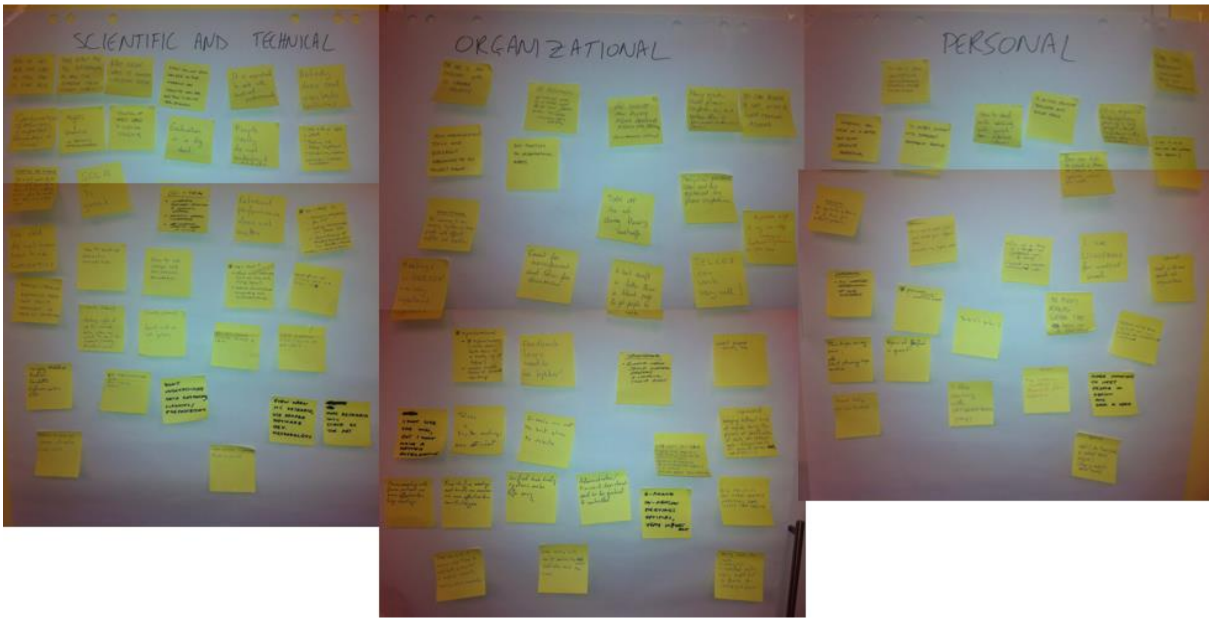

Figure 12: The collected information on the lessons learned.

We provide diagrams summarising the lessons learned (the texts on the post-it notes) in the three categories. In the Scientific and Technical Category, shown in Figure 13, the lessons were clustered into five groups: medical search domain, semantic and language technologies, development, data, and evaluation. The practical side of the development of a large-scale integrated system led to the largest number of lessons learned. As the Khresmoi consortium was a highly multidisciplinary consortium, many of the lessons learned are actually due to the transfer of knowledge between different domains, so that e.g. people working in medical imaging learned about the use of semantics, ontologies and knowledge bases.

In the Organisational category, shown in Figure 14, most of the lessons learned deal with the practicalities of communication in a large and diverse group of people. These lessons are particularly useful for the young researchers who will be running projects in the future.

In the final category, Personal, shown in Figure 15, the main lessons learned were again in the communication and interaction group, with some additional skills also learned.

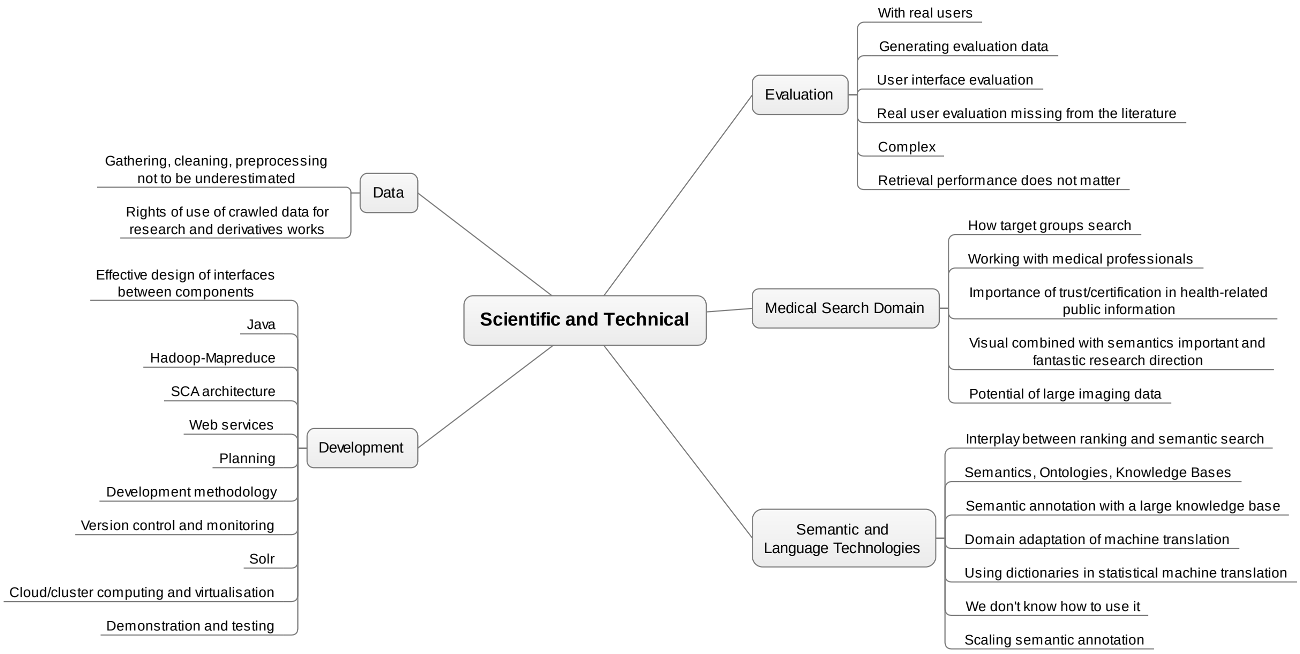

Figure 13: Lessons Learned in the Scientific and Technical category

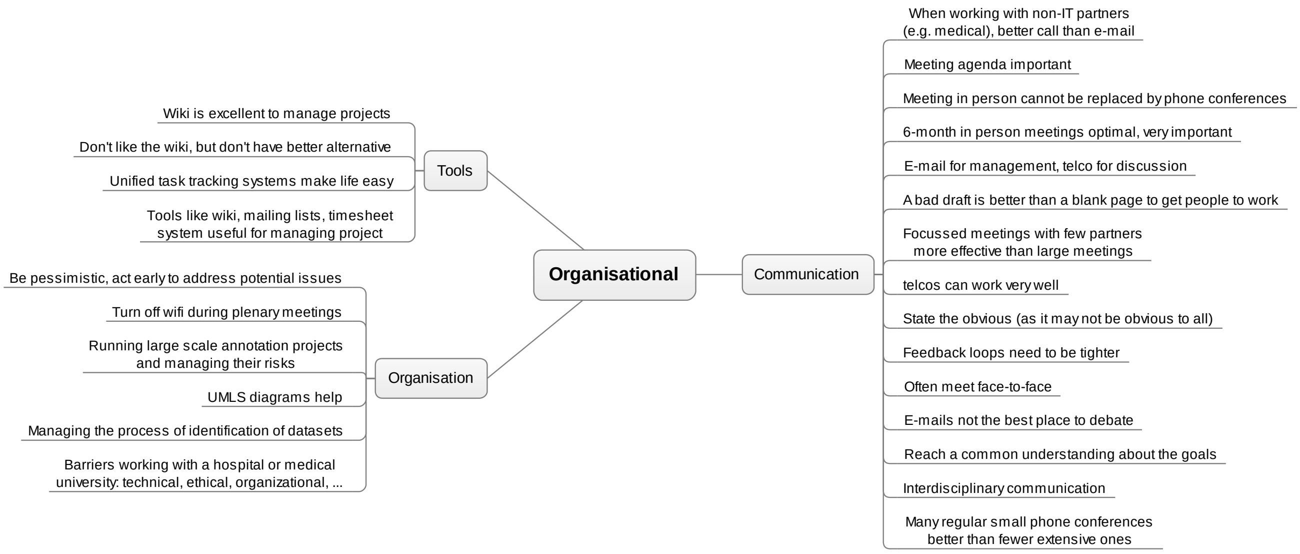

Figure 14: Lessons Learned in the Organisational category

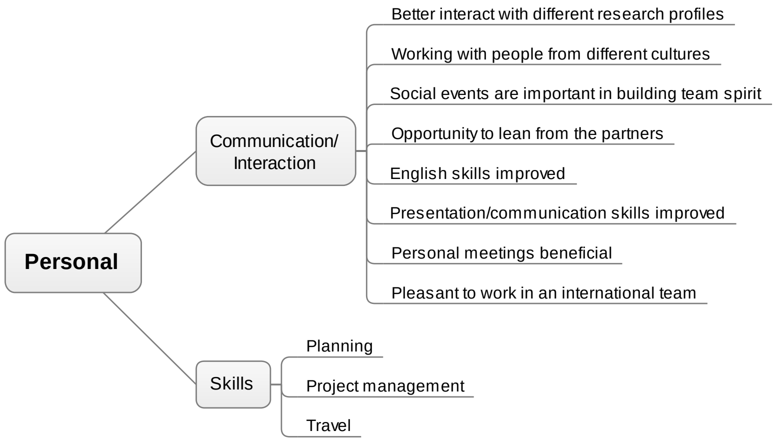

Figure 15: Lessons Learned in the Personal category

4.1.5. Contact Details

Project Website: http://khresmoi.eu

Project Coordinator:

Prof. Henning Müller

University of Applied Sciences Western Switzerland, Sierre (HES-SO)
TechnoArk 3
3960 Sierre, Switzerland

Email: henning.mueller@hevs.ch

Phone: +41 27 606 9036

Scientific Coordinator:

Dr. Allan Hanbury

Vienna University of Technology
Institute of Software Technology and Interactive Systems
Information & Software Engineering Group
Favoritenstrasse 9-11/188
A-1040 Vienna
Austria

Email: hanbury@ifs.tuwien.ac.at

Phone: +43 1 58801 188310

Project logo:

khresmoi

MEDICAL INFORMATION ANALYSIS & RETRIEVAL