

Supporting Collaborative Improvement of Resources in the Khresmoi Health Information System

Lorraine Goeuriot¹, Allan Hanbury², Gareth J. F. Jones¹, Liadh Kelly¹, Sascha Kriewel³,
Ivan Martinez Rodriguez⁴, Henning Müller⁵, Miguel A. Tinte⁴

¹Dublin City University – Dublin, Ireland

lgoeuriot, gjones, lkelly@computing.dcu.ie

²Technische Universitaet Wien – Vienna, Austria

allan.hanbury@tuwien.ac.at

³Universitaet Duisburg-Essen – Duisburg, Germany

sascha.kriewel@uni-due.de

⁴Atos Origin Sociedad Anonima Espanola – Madrid, Spain

ivan.martinez, miguel.tinte@atos.net

⁵University of Applied Sciences Western Switzerland – Sierre, Switzerland

henning.mueller@hevs.ch

Abstract

Since medical knowledge relies on both scientific knowledge and real-life experience, the importance of user contributions to improve resources in health systems cannot be underestimated. We present work from the Khresmoi project which aims to develop a multilingual multimodal search and access system for biomedical information and documents. Khresmoi targets three distinct user classes with differing levels of medical knowledge and information requirements, namely: general public, general practitioners, and, as an example of an area of clinical expertise, radiologists. The Khresmoi system will provide these users with valuable (whose quality has been evaluated and approved) and enriched (meta information from biomedical knowledge bases is added) medical information, selected to fit their medical knowledge and their preferred language. The system will include novel collaborative components of the system are designed to provide means for users to contribute to the system's knowledge by adding or correcting annotations to the documents, as well as a collaborative platform where they will be able to share their own files and both annotate and discuss them.

1. Introduction

Annotation of biomedical data is vital in order to be able to organise and structure the knowledge it contains, and to select and deliver information relevant to the information need of a searcher seeking to address a medical information need from these sources. In this paper, we describe our current work exploring how users (e.g. patients, physicians, etc.) of a medical system can help to improve it by contributing to the quality of its resources and by adding their knowledge to the stored information.

This work is being carried out within the Khresmoi project¹, which aims to develop a multilingual and multimodal search and access system for biomedical information and documents (Hanbury et al., 2011). The Khresmoi project is being targeted at three groups of end users: two groups with general medical interests (general public and general practitioners) and a group of clinicians with specialised expertise (radiologists); all speaking different languages, having different medical knowledge levels and differing levels of knowledge of the languages of the target documents. The system is based on a library of valuable medical documents (images and text) that are enriched using a medical ontology such as UMLS² (Unified Medical Language System) or MeSH³ (Medical Subject Headings)

and knowledge bases such as the LinkedLifeData⁴. The Khresmoi system is being designed to enable our users to correct computed knowledge (meta information and translations), as well as share their experience.

Based on a collection of biomedical documents, including medical 2D images and 3D volumes, automatically annotated with biomedical ontologies, we plan to provide users with the potential to correct errors in these automatic annotations. Since medical knowledge relies both on scientific knowledge and experience, medical literature may not be enough to understand a treatment, a procedure or even the description of a disease. Document meta-information and comments from users can help gathering that knowledge in a single space. For example, a young radiologist will have to check different resources and maybe colleagues to spot an area of interest on an X-ray image. With such a system, he will be able to search for similar images and then use the meta-information/annotations to validate his diagnosis. We will also provide them with tools to share their knowledge through notes and comments on documents. Both the user and the system can benefit from such collaborative tools: improving the quality of data will improve quality of the medical system search, and sharing knowledge and experience helps physicians in their everyday practice. The system will also provide automatic translations of the queries and documents. As automatic translation methods do not give perfect results, we will allow users to correct transla-

¹<http://khresmoi.eu/>

²<http://www.nlm.nih.gov/research/umls/>

³<http://www.nlm.nih.gov/mesh/MBrowser.html>

⁴<http://linkedlifedata.com/>

tion errors as well.

The next section describes related work in medical related collaboration tools. Section 3. provides an overview of the Khresmoi project and its objectives, along with a description of the project's user interface system and resources used. Section 4. describes how users can collaborate to improve the system resources by updating annotations and translations, as well as communicate through comments and discussion threads. Finally Section 5. summarises the paper and outlines the focus of our ongoing work.

2. Related Work

Collaboration by editing digital resources to correct and augment their content is key to obtaining richer information. Knowledge, especially in such specialised domains as medicine, relies on scientific knowledge and experience. However, gathering knowledge from text sources by using automated information extraction methods only produces partially correct scientific knowledge of the data due to errors in the extraction process, and will generally be much less reliable than human-annotation. Web 2.0 technologies enable users to collaborate in the development of content, and an inclination do to this has been observed in the medical domain (Eysenbach, 2008). *Ask Dr Wiki*⁵ and *Medpedia*⁶ are two well-known wikis where physicians can create content, and collaborate on its editing. These wikis must provide complex validation systems in order to guarantee the quality of the information published. The purpose of these websites is mainly to improve online health information. Another online collaborative annotation tool, called *Brat*, provides a user-friendly interface to display and change annotations on text from a web browser. Registered users can view and annotate online files and upload their own files. It has been used for BioNLP extraction tasks and is mainly natural language processing (NLP) focused (Stenetorp et al., 2012). Collaborative projects have also been defined for particular communities of practice, where users sharing patients or interests can discuss cases, information and even manage meetings. For example, the SOMWeb system (Falkman et al., 2008) assists the community of Swedish oral medicine practitioners. Using OWL (Web Ontology Language) to model their data, it allows users to add cases, notes, discussions and manage community aspects.

Medical wikis provide users with a way to gather their knowledge in creating new content, while community of practice collaborative systems are specific software or online systems allowing collaboration in a very specific framework. However, none of these systems provides access to other resources, which is one of the main uses of the Internet. The time practitioners can spend online is rather limited: they spend on average less than 5 minutes to answer a question (Hoogendam et al., 2008). Expecting them to be active on different platforms is unrealistic. A system providing all these services at the same time would be valuable. The Khresmoi system, presented in this paper, is designed to provide a search service on valuable and en-

riched medical documents. The system includes collaborative components intended to enable users to improve resource documents and engage in discussions.

3. Khresmoi System

The Khresmoi project aims to develop a multilingual multimodal search and access system for biomedical information and documents. Khresmoi is adopting a user-centred approach to designing medical information search tools, for which three groups of end users are defined. Two of these are groups with general medical interests: general practitioners and members of the general public. The Khresmoi system is intended to provide them with innovative text search features to interrogate the huge amount of medical information available, including that appearing in journals, websites, Wikipedia and clinical guidelines. These users wish to rapidly find answers to their queries that are suitable for their level of expertise. The other user group that Khresmoi focuses on is radiologists, as an example of clinicians with a specific expertise. For radiologists we plan to provide advanced image search to support them in their work. The Khresmoi system is being developed within a four year project which is now in the first half of year two. During the first year of the project, the requirements of the end users were obtained through surveys and interviews. Following this, the design process for the Khresmoi system has led to a specification of: the characteristics of the target user groups, the types of search tasks that the users would perform, the resources that each user type wishes to access, and the search tools and refinements needed by each user type to carry out their tasks. An interesting result of the survey is the perceived importance of the collaborative aspects of search for medical professionals, who wish to see their peer's opinion on documents and also additional examinations that can increase their confidence in a diagnosis.

3.1. Khresmoi Users and Their Needs

In this section we summarise the surveys carried out within the Khresmoi project to investigate what the different categories of users need from a health information system.

- 385 members of the general public, mostly highly educated and coming from healthcare (not physicians) and IT backgrounds answered the survey. They came from 42 European countries (with the highest numbers of contributors coming from France and Spain). The most researched topics by these users are: general health, chronic diseases and lifestyle. When they were asked what were the most important characteristics of search tools, they mentioned the relevance and trustworthiness of the results.
- 556 physicians and 4 final-year medical students, mostly Internet savvy and with regular patient contact were surveyed. These respondents came mainly from Austria, Switzerland and the United Kingdom. The topics they search on the most are: drugs, treatments and medical education. Currently they mostly use generic search engines (such as Google). Specialist physicians also search for clinical trials and have a preference for medical research databases or society

⁵<http://askdrwiki.com/>

⁶<http://www.medpedia.com/>

websites, whereas general practitioners also search for disease description and tend to use more general health websites.

- 34 radiologists were surveyed, a majority of them young subjects currently with little radiology experience, however several of them had more than 15 years experience. They came mainly from Switzerland and Austria. Image search (search for images matching certain disease or body parts) was mentioned as a common task, but time consuming (often more than 10 minutes) and with 65% success in completing the task. One of the main points of a search tool is then to be able to find good and relevant image results quickly. Subjects would also like to be able to upload an image on a search tool as a query, to find similar images of similar cases.

From these surveys, we can see that the quality of the information, as well as its relevance and trustworthiness are very important criteria for every kind of user. Medical practitioners and radiologists mentioned the need to share information: medical practitioners wanted to have access to a secured community where they can exchange information about cases and share or update their knowledge; and radiologists mentioned that feedback from colleagues on past/current cases was valuable information. Therefore, users express the need of high quality information, as well as interactivity and communication functionalities. More details on these surveys can be found in the public deliverables of the Khresmoi project: (Pletneva and Vargas, 2011) for general public, (Gschwandtner et al., 2011) for medical professionals and (Müller, 2011) for radiologists.

Web2.0 and social media are having an impact on the medical domain, both on the specialist side (Giustini, 2006; Eysenbach, 2008) and on the patient side (Fox, 2011). This change has raised concerns about the quality of information (Denecke and Nejdil, 2009): without any editorial process, how can it be guaranteed? However, Web2.0 is subject to a “socially Darwinian process” (also called *positive network effect*): (Boulos et al., 2006) note with regard to wikis that “because of [the] openness and rapidity that wikipages can be edited, the pages undergo an evolutionary selection process not unlike that which nature subjects to living organisms”.

If the user contributions are done in a controlled and secured way, with an adapted moderation system, the quality of information can still be guaranteed. What we propose here within Khresmoi is to let the users directly contribute to the quality of the information by correcting the metadata (annotation and translation of multilingual content), as well as to be able to freely comment and discuss cases in a secure environment.

3.2. Khresmoi Resources

The Khresmoi system will potentially index a very large number of documents from the biomedical domain. As the collection is a very long process, we gathered datasets for our first prototype in order to observe specific users behaviour. To improve the search, the approach of annotating

the documents with entities important in the medical domain is being adopted, where the entities are taken from a knowledge base of domain ontologies in the medical and life sciences, such as the LinkedLife Data (semantic data integration platform for the biomedical domain).

Datasets used within the project for the first year prototype include 2D and 3D images, as well as text. The 3D image collection consists of: realistic clinical data (medical images and reports from the Vienna University Hospital, constituting over 3 TeraBytes of data) and lung data (medical images and reports collected in the University Hospital of Geneva, corresponding to more than 100 interstitial lung disease cases). These two collections have been anonymized and annotated using RadLex⁷ and MeSH.

The 2D image collection is a collection from ImageCLEF2011 (Kalpathy-Cramer et al., 2011). It contains 231,000 images from the PubMed Central Database and corresponding articles, with articles annotated with MeSH. The text collection gathers MEDLINE⁸ abstracts, UMLS⁹ definitions, a set of Health on the Net¹⁰ classified documents about diabetes. All these documents have been annotated with LinkedLife Data. These datasets have been designed for the first Khresmoi prototype and will be extended as part of the ongoing work of the project.

For the text annotation work during the project, extensive use of manual feedback from professional annotators is made to correct the annotations, and hence allow the system to improve the automated annotation through learning. However, the extensive use of professional annotators is not a sustainable approach, and the system will have to increasingly rely on annotation corrections from the end users. For the cross-lingual search, use of resources for which translated versions of terms are linked to each other is made, such as the MeSH thesaurus¹¹.

4. Collaborative Plans in Khresmoi

In this section, we describe technologies that have been developed within the project and our development plans for the future of the project.

During the first year of the Khresmoi project, a user interface framework based on ezDL technology has been developed. We are currently extending this to implement our plans to create tools to enable users to collaborate. An evaluation phase of these components is planned later in year 2 following their development.

4.1. EzDL System

The user interface of the Khresmoi system is based on ezDL¹², the successor of the Daffodil software (Fuhr et al., 2002) developed at the University of Duisburg-Essen. EzDL is a multi-agent search system for heterogeneous data sources and a tool-set for building search user interfaces to support complex tasks. It allows for simultaneous searches in multiple digital libraries through a unified interface and

⁷<http://www.radlex.org/>

⁸<http://www.ncbi.nlm.nih.gov/pubmed/>

⁹<http://www.nlm.nih.gov/research/umls/>

¹⁰<http://www.hon.ch/>

¹¹<http://www.ncbi.nlm.nih.gov/mesh>

¹²<http://www.ezdl.de/>

query syntax, and presents a merged and enriched view of the results. The tools provided by ezDL allow users to work with the results and can be arranged in customizable perspectives.

EzDL is composed of a server part consisting of a directory and a large number of agents, and clients that contain a selection of loosely-coupled tools which serve as a user interface to the system (see Figure 1).

The server-side agents connect to the search and query support services provided within Khresmoi, handle user authorisation, user profile management, logging, storage of user data and queries, and the caching of documents. Two basic clients are available within Khresmoi: a search desktop written in Java (see Figure 2), as well as a browser application that uses Java Server Faces. Users can either search as guests or obtain a personal account. A personal account allows for a persistent search history spanning multiple search sessions and offers access to a document depository called ‘personal library’, where a user can store found and uploaded documents, as well as favourite queries and authors, and categorise them with personal tags.

An account will also be necessary to contribute to the system’s knowledge by adding or correcting annotations on the documents. Guests and registered users alike can use the search tool with query formulation support which offers spelling corrections and disambiguation of medical terms. The results are presented in a combined list that searchers can group using options like date, type of document (e.g. image or text), category of document (e.g. treatments, symptoms, genetics) or audience (e.g. general public, practitioners or researchers). The search tool also offers filtering, sorting by different criteria, and export options. Documents that have already been inspected, stored, printed or otherwise handled by the user are clearly marked with icons in the result list. The detail tool of ezDL offers a preview of documents from the result list or from the personal library. It shows document metadata (authors, publication date, publication type, journal or conference), annotations of the content and summaries where available. A link to the full document (website, article or media file) is also provided.

4.2. Khresmoi Collaborative Components Development Plan

As mentioned in Section 3.1., surveyed potential users expressed the desire to share knowledge, especially medical professionals and communities of practice. Web2.0 facilitates this knowledge sharing on the web by allowing users to directly contribute information (e.g. Wikipedia or Med-pedia). The Khresmoi system will provide users two ways to share their knowledge:

- correction of existing annotations and translations created by the system;
- creation of comments on Khresmoi documents or on documents uploaded by users, that can target a specific part of the document (region of interest in an image or sentence/paragraph in a text) or the whole document.

These two collaborative approaches will improve Khresmoi resources by adding: explicit knowledge through correc-

tions, and implicit knowledge through comments. While the system can directly benefit from explicit knowledge, both can be useful for users. As mentioned in the surveys (see Section 3.1.), the quality and the relevance of a search result are very important criteria. If users can correct resources on the system that they are also using to get information, they can directly benefit from their input: better translations and annotations improve the quality and relevance of the documents (e.g. though ranking process). We also observed in the user surveys that experience sharing played an important role in physicians and radiologists everyday practice. This system could allow them to do it online, with colleagues that can be in other institutions. For example, a radiologist could give feedback through notes on a radiological image to a general practitioner who needs advices. Physicians can share comments on new clinical trials with other physicians or highlight useful recommendations in a document for patients.

We provide details on these collaborative approaches in Sections 4.2.1. and 4.2.2..

4.2.1. Users Correcting Annotations and Translations

To improve resources in the Khresmoi system users will be able to update and correct errors in such resources while using the system. This can take several forms: direct correction of errors or omissions in annotation or translation, or manual contribution of new knowledge, e.g. translations, or verification or clarification of automatically extracted suggested updates for resources. In addition to supporting users in updating resources in operation, we will also explore methods such as collaborative editing to keep the resources up to date. From a technical perspective we propose the development of a *Collaborative Resources Framework* to support the improvement process. Figure 3 presents an overview of the *Collaborative Resources Framework* as well as the external communication with other components of the Khresmoi system.

We can distinguish two types of processes in the context of collaborative improvement: updating and validating. These processes are aligned with components in the collaborative framework: the *Resource Updater* is responsible for the annotation and translation management; and the *Validator* responsible for managing the life cycle of the user annotations and translations. Both annotations and translations will be by default in a *Pending* state and could change to a *Validated* or a *Refused* state. We next describe these two processes in greater detail.

Resource Updater : This component will manage annotation and translation updates incoming from the ezDL user interface. It consists of two main subcomponents: *Annotation Manager* and *Translation Manager*. The *Annotation Manager* is responsible for implementing the workflows for *New Creation* and *Update Annotation* functionalities as they are offered by ezDL. The Annotation Manager will insert annotations, and updated annotations, into the User Profile database. The Annotation Manager also writes to/reads from an *Annotation State Store*. This store manages the different possible states associated with annotations (*Pending*,

Validated and Refused). The default annotation status will be *Pending*, requiring a user to validate the annotation and change the status to *Validated* or *Refused*. The *Translation Manager* will implement the functionality associated with the *Update Translation* workflow in Figure 3. To fulfil this task, this component will use the *Multilingual Translation Framework* (MTF) provided by our system. The MTF controls the management and storing of translations and user translation updates, hence they will be stored outside of the Collaborative Resources Framework. The *Update Translation* functionality will be provided by the ezDL user interface and the manager will recover the translation from the MTF. Similar to annotations, translations will require user validation. The status associated with user translations will represent their validation status.

Validator : This component will provide the functionalities needed for managing the life cycle associated with annotations and translation. As mentioned previously, when a user adds or updates one concrete document annotation or translation the *Resource Updater* marks as *Pending* the state of the annotation or translation. To support this functionality the Updater will use the *Annotations State Store* for annotations and the MTF for translations. The Validator component will allow users to carry out two types of actions over pending annotations or translations: validate or refuse them. Following user validation, the Validator component commits or discards the annotation/translation as appropriate.

4.2.2. Users adding comments to documents

As we said previously, medical professionals' knowledge is based on scientific knowledge but also relies strongly on their experience. While the scientific knowledge can be more or less similar across persons and available in books and online, experience is rather individual. For this reason it is very important and interesting for practitioners to share this knowledge. Our system aims to provide users with a simple system to share their knowledge and experiences. Registered users will be allowed to share documents from the project library and add comments and discussions on these files. They will also be able to upload their own files (e.g. patients report or x-ray radiography) to the system, which will be anonymous (no patient information) and private (the user will choose people to share the file with). Users can add comments on the whole document or on a region of interest (Figures 4 and 5).

Users' rights fall into 3 categories:

Read : Users will be allowed to read comments from other users. The comments will be accessible for users within the same category (general public, medical practitioners or radiologists), unless the author specifies other categories (e.g. a physician could highlight an interesting paragraph for patients).

Write : Users will be allowed to create/write new comments. Whatever the document is, these users will be




Figure 4: Example of annotations on an image (from <http://en.wikipedia.org/wiki/Radiography>)



Figure 5: Example of annotations on a text (from <http://en.wikipedia.org/wiki/Radiography>)

allowed to add new comments to discuss it or add new knowledge (annotations).

Modify : Users will be able to edit or delete all their comments. They will not be allowed to modify other users comments.

All users, even if they are not registered, will be allowed to read comments written for their category. Registered users will be able to write new comments and modify their own comments. When a new comment is added, the user will have to choose categories of users allowed to read it (e.g. a doctor can write comments for patients). Registered users will be able to edit or delete their own comments.

4.3. Evaluation of the Collaborative Components

Empirical and user-centered evaluation strategies have been developed for the Khresmoi system, which will be conducted in the coming months. The user-centered part of this system evaluation strategy encompasses evaluation of the collaborative components using target user groups. This will entail subjects from each category of user using the system to fulfil predefined scenarios. Feedback gained on the collaborative components through these evaluations will be used to adapt the components to make them more user-friendly and suitable to user practice.

5. Summary and Ongoing Work

In this paper, we have presented a set of collaborative functionalities that will be included in the Khresmoi medical information search system. This system will provide users with a valuable search tool for medical documents that are available in multiple languages, and enriched using medical thesauri. Medical documents can be processed by the system using information extraction tools to include semantic annotations. To do this, a knowledge base of domain ontologies in the medical and life sciences is used. The system will also provide automatic translation of the queries and documents, and provide users with facilities to collaborate to correct these annotations and translations. User collaboration will also be possible through a component which will allow users to add comments and start discussions on documents from the library or their own files. The development of these components is ongoing. These components, along with the system, will be evaluated in the coming months, through both empirical and user-centered evaluations. Patients, medical practitioners and radiologists will partake in the controlled user-centered system evaluations. The system will be improved based on feedback from these evaluations. Following this, the system will be deployed for use by real users. Among other things this will allow us to both assess the quality and value of users' input, and investigate how user input could further contribute to the system. For example, comments and discussions from physicians on a document describing a case might provide rich information that the system could learn to process.

6. Acknowledgements

This research described in this paper has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257528 (KHRESMOI)

7. References

- Maged Kamel Boulos, Inocencio Maramba, and Steve Wheeler. 2006. Wikis, blogs and podcasts: a new generation of web-based tools for virtual collaborative clinical practice and education. *BMC Medical Education*, 6:41.
- Kerstin Denecke and Wolfgang Nejdl. 2009. How valuable is medical social media data? content analysis of the medical web. *Journal of Information Sciences*, 179:1870–1880.
- Gunther Eysenbach. 2008. Medicine 2.0: Social networking, collaboration, participation, apomediation, and openness. *Journal of Medical Internet Research*, 10(3).
- Göran Falkman, Marie Gustafsson, Mats Jontell, and Olof Torgensson. 2008. Somweb: A semantic web-based system for supporting collaboration of distributed medical communities of practice. *Journal of Medical Internet Research*, 10(3).
- Susannah Fox. 2011. Health topics. Technical report, Pew Research Center, February.
- Norbert Fuhr, Claus-Peter Klas, André Schaefer, and Peter Mutschke. 2002. Daffodil: An integrated desktop for supporting high-level search activities in federated digital libraries. In *Research and Advanced Technology for Digital Libraries. 6th European Conference, ECDL 2002*, pages 597–612, Heidelberg. Springer.
- Dean Giustini. 2006. How web 2.0 is changing medicine. *British Medical Journal*, 333(7582):1283–1284.
- Manfred Gschwandtner, Marlene Kritz, and Celia Boyer. 2011. D8.1.2: Requirements of the health professional search. Technical report, Khresmoi Project, August.
- Allan Hanbury, Célia Boyer, Manfred Gschwandtner, and Henning Müller. 2011. Khresmoi: towards a multilingual search and access system for biomedical information. In *Med-e-Tel*, Luxembourg.
- Arjen Hoogendam, Anton F.H. Stalenhoef, Pieter F. de Vries Robbé, and John P.M. Overbeke. 2008. Answers to questions posed during daily patient care are more likely to be answered by uptodate than pubmed. *Journal of Medical Internet Research*, 10(4).
- Jayashree Kalpathy-Cramer, Henning Müller, Steven Bedrick, Ivan Eggel, Alba García Seco de Herrera, and Theodora Tsikrika. 2011. Overview of the CLEF 2011 medical image classification and retrieval tasks. In *CLEF 2011 working notes*.
- Henning Müller. 2011. D9.1: Report on image use behaviour and requirements. Technical report, Khresmoi Project, May.
- Natalia Pletneva and Alejandro Vargas. 2011. D8.1.1. requirements for the general public health search. Technical report, Khresmoi Project, May.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii. 2012. brat: a web-based tool for nlp-assisted text annotation. In *Proceedings of the Demonstrations Session at EACL 2012 (to appear)*.

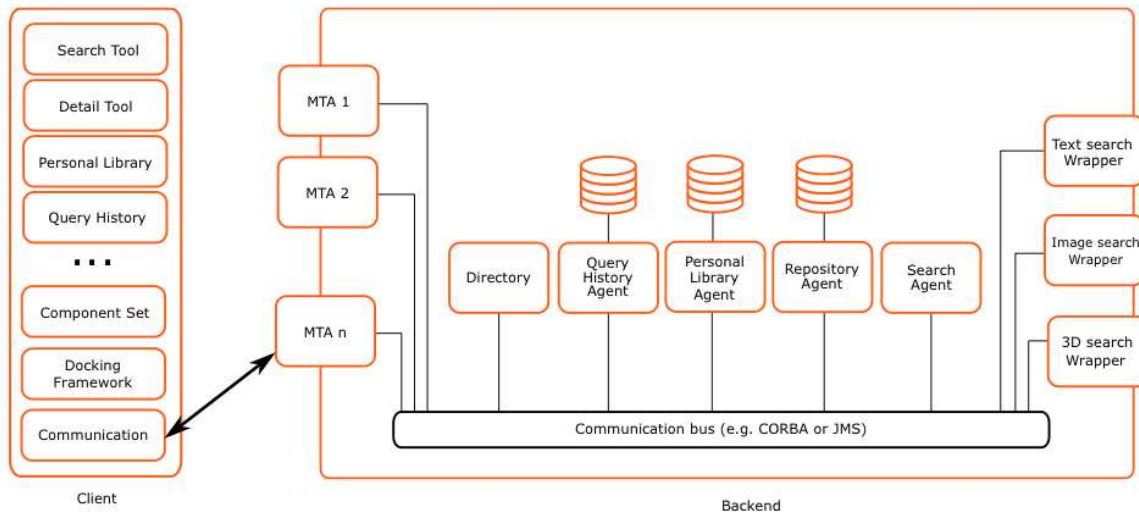


Figure 1: Architecture of ezDL

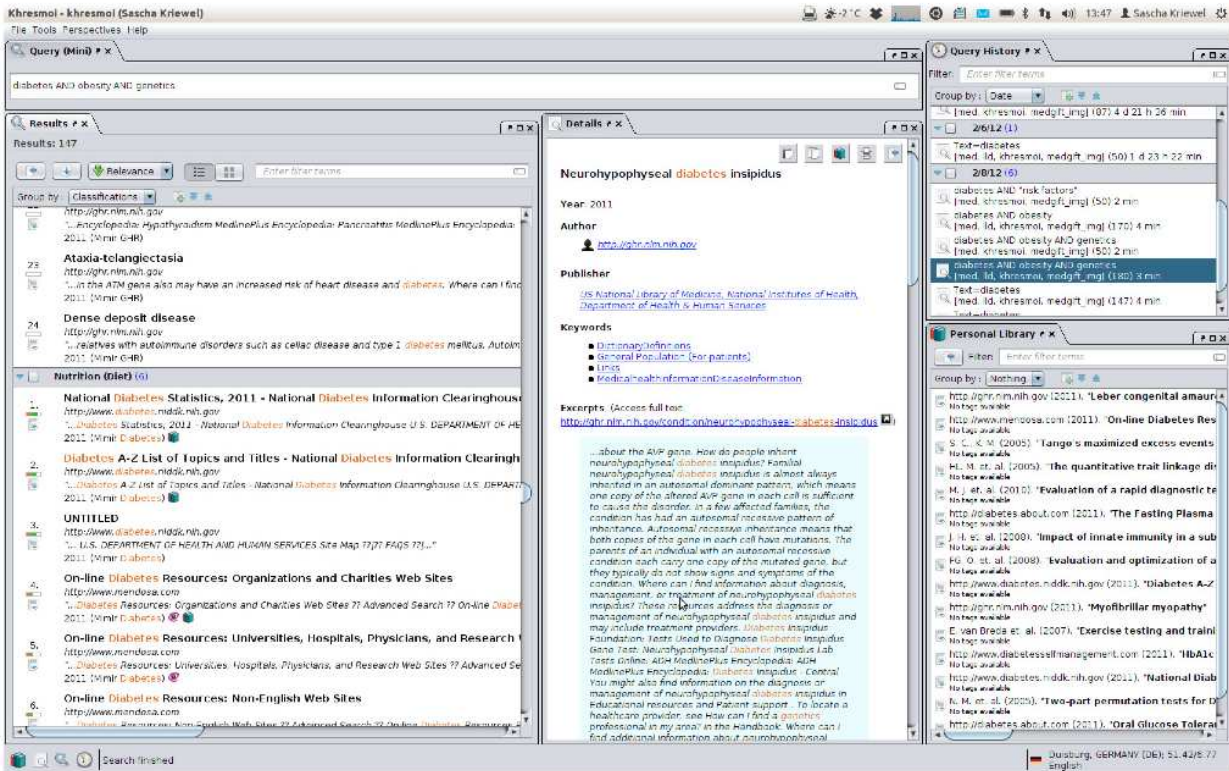


Figure 2: Screenshot of ezDL Interface




Figure 3: Collaborative Resources Framework Architecture