

Search Computing

Business Areas, Research and Socio-Economic Challenges



Media Search Cluster White Paper



European Commission
Information Society and Media

LEGAL NOTICE

By the Commission of the European Communities, Information Society & Media Directorate-General, Future and Emerging Technologies units.

Neither the European Commission nor any person acting on its behalf is responsible for the use which might be made of the information contained in the present publication.

The European Commission is not responsible for the external web sites referred to in the present publication.

The views expressed in this publication are those of the authors and do not necessarily reflect the official European Commission view on the subject.

Luxembourg: Publications Office of the European Union, 2011

ISBN 978-92-79-18514-4

doi:10.2759/52084

© European Union, July 2011

Reproduction is authorised provided the source is acknowledged.

© Cover picture: Alinari 24 ORE, Firenze, Italy

Printed in Belgium

Search Computing: Business Areas, Research and Socio- Economic Challenges

Media Search Cluster White Paper

Search Computing is a fortuitous term used since 2008 by Prof. Stefano Ceri of the Politecnico di Milano, in the EU – ERC project SeCo which pioneered this domain.

Coordinated by the CHORUS+ project co-funded by the European Commission under - 7th Framework Programme (2007-2013) by the –Networked Media and Search Systems Unit of DG INFSO

Table of Contents

1	Executive Summary	5
2	Introduction.....	7
3	Business areas	12
3.1	Searching in Mobile business area.....	12
3.2	Searching in Social business area.....	14
3.3	Searching in Enterprise business area	15
3.4	Searching in Music business area	16
4	Research challenges	18
4.1	Multimodal.....	18
4.2	Affective	19
4.3	Event-based	20
4.4	User experience	20
4.4.1	Visual Search Interfaces.....	21
4.4.2	Augmented reality	22
4.5	Large-scale indexing.....	23
4.6	Network nodes.....	25
4.6.1	Passive Content discovery	25
4.6.2	Content/web-page popularity discovery.....	25
4.6.3	Network topology discovery.....	26
4.7	Real-time	26
4.8	Content diversity.....	27
4.9	Aggregation, mining and Linked Open Data (LOD)	28
4.10	Standardisation	30
5	Socio-economic challenges	33
5.1	Business Models and Policies.....	33
5.2	Search and open innovation	34
5.2.1	Search Ecosystem at European Level	38
5.2.2	Search computing and user generated innovation	38
5.3	Benchmarking: A catalyst to advance the state of the art in multimedia search computing.....	39
5.3.1	What is a benchmark?	39
5.3.2	What are the benefits of benchmarking?.....	39
5.3.3	Benchmarking to reinforce Europe’s competitiveness	41
5.3.4	Challenges in benchmarking.....	42

5.4	Legal and ethical issues	43
5.4.1	EU Legal framework for data protection, privacy and search	43
6	Conclusions.....	46
7	BIBLIOGRAPHY.....	48
8	CONTRIBUTING PROJECTS AND ORGANIZATIONS	55
9	EDITORS AND AUTHORS.....	56
	Editors.....	56
	Authors	56

1 Executive Summary

Search has become an important and necessary component of many diverse ICT applications. A large number of business and application areas depend on the efficiency and availability of search techniques that are capable to process and retrieve heterogeneous and dispersed data. These include: a) the Web, b) mobile devices and applications, c) social networks and social media, and d) enterprise data access and organization.

Search techniques are directly related to a number of research topics and challenges ranging from multimodal analysis and indexing to affective computing and human aspects as well as to various socio-economic challenges including standardization, legal and ethical issues. The area covering all these issues is also known as “Search Computing”.

The objective of this document is to provide an overview of the business areas, the research challenges and the socio-economic aspects related to “Search Computing”. The business areas addressed in this paper include, mobile, enterprise, social networks and music, focusing on their current status, their specific technology and application requirements and the new needs arising as a result of these requirements. Specific technologies and challenges are identified such as device limitations and location-based opportunities in mobile search but also challenges that are applicable across domains such as limitations of text-based approaches and new research directions for content-based search. For example, in social networks, new Social Indexing based approaches can incorporate information about the structure and activity of the users’ social network directly into the search and ranking process.

While significant research has been performed in many related areas, there are still unsolved problems, where new research approaches are required for efficient and human-centric search. Affective-based search and content diversity have appeared more recently as promising research directions for existing search problems. Specific areas where further work is needed include:

- **Multimodal search** enabling queries and interaction independently of the form, which the content is available,
- **Affective-based search** taking into account both the users emotional state and the sentiment contained in multimedia documents
- **Event-based** representation and analysis as an efficient and user-centric approach for the annotation and retrieval of content
- **User experience aspects** including new generations of search interfaces (e.g., visual search interfaces, augmented reality, 3D browsing in virtual places) to enable novel forms of search input and to address the problem of losing the overview of the incredible amount of available content
- **Large-scale indexing** addressing the scalability issues related to the huge volume of content available on the Web by using social connections to implement targeted and thus smarter indexes
- **Content-aware network nodes** extending search in the network with challenges related to “discovery at the network” than “searching in the search engine”

- **Real-time approaches and architectures** enabling to push information to users as fast as it is available, maintaining on the same time the balance between quality, authority, relevance and timeliness of the content
- **Content diversity** in knowledge, providing the ability to identify and exploit the aspects that differentiate a piece of information from another
- **Aggregation, mining and Linked Open Data** enabling data collection and aggregation from social networks and interlinking data chunks under the Linked Open Data principles
- **Standardization** initiatives enabling a complete set of new technologies to face the market and achieve a deep impact in the business of the future

In addition to the research challenges, the socio-economic challenges of “*Search Computing*” are analysed next including business models, search and open innovation, benchmarking, legal and ethical issues (data protection and user privacy). New **business models** are needed in applications where user-generated content plays an important role in order to allow both commercial exploitation and protect user rights. **Open innovation** is the use of purposive inflows and outflows of knowledge to accelerate internal innovation and expand the markets for external use of innovation. **Benchmarks** are valued for their ability to streamline research by eliminating redundancy, enabling direct performance comparison between algorithms, increasing efficiency by sharing resources between research sites and providing a concrete framework in which researchers interact in a productive mixture of competition and collaboration. Benchmarking directly benefits the research community and indirectly produces economic benefits by bringing innovative research closer to market. Finally, social networks pose new challenges to **legal and ethical issues** that should be carefully considered for proper usage and processing of personal data.

In the remaining of this document we provide an overview of the most recent business areas directly related and depending on “*Search Computing*” and go on to identify and describe the research challenges arising from the need for efficient and human-centric search technologies in these areas. The socio-economic challenges of “*Search Computing*” are analysed together with specific application areas and domains. These research and socio-economic challenges can serve as guidelines for defining future research agendas and programmes, such as in the forthcoming European Union’s Common Strategic Framework (Framework Programme 8).

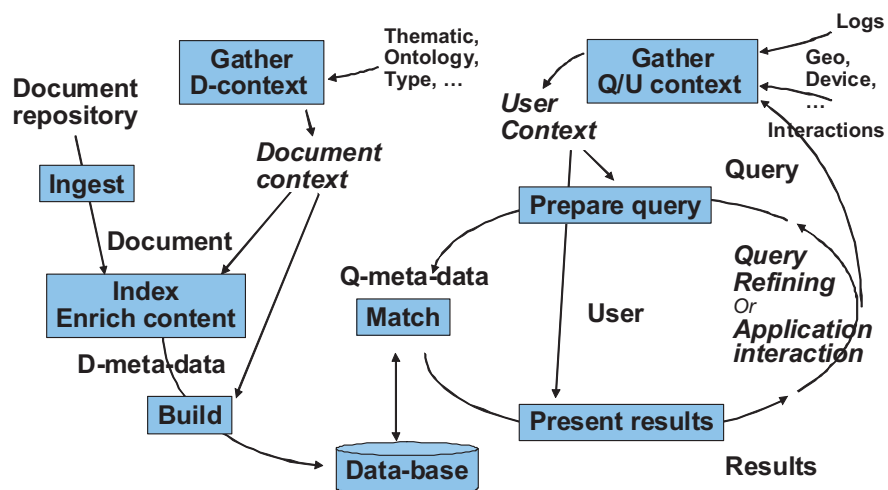
2 Introduction

The rapid advances in the field of information technology have significantly reduced the traditional spatial and temporal obstacles in information exchange. At the same time, with the introduction of affordable and high-quality capture devices, users and professionals are able to generate and share much more content than before. Instant information sharing infrastructures enable users to easily generate and exchange considerable amounts of digital data. Moreover, virtual communities and social networks accommodated by the World Wide Web (WWW) are rapidly turning into a significant source of content generation [Shen06].

However, in this constantly growing digital environment, the fast pace of generating and sharing new digital content poses serious risks on its accessibility. The difficulty of mapping a set of low-level visual features into semantic concepts, generally addressed as bridging the "Semantic Gap" [Smeulders00], makes it difficult for automated systems to interpret and organize digital content in a manner coherent with human cognition, raising the need to discover and adopt intelligent ways to mine, browse, summarize and in general consume digital information [Chang02]. In this context, search engines have been the primer knowledge broker to the abundant availability of information in the Web.

Search engines are one of the very popular 21st century technology inventions. For the vast majority of internet users, a reliable web browser and

a search engine are the preferred entry point to find out what's relevant to them in the World Wide Web. Search engines have become the end-user's knowledge brokers. The engine is what's "in the box": an essential and interactive technology to match one's queries with meaningful results. In big-scale systems, search engines allow entry and can facilitate "connectedness" of web pages based on keywords and content. In a business sense, search engines allow growth and expansion (commercial visibility) to sites that are searchable, as well as for a majority of web services offered. Search plays a key role in the discovery and growth (scalability) of services that are visible on-line.



Functional breakdown of a search engine, CHORUS Final Report 2009

The visible web is but the tip of the submersed iceberg of digital data that is being collected and that could one day be mined or reused. Gigabytes of data flow into the web daily, increasingly due to user-generated and image-based content. For anyone to access, find, or navigate the visible portion of the Internet, requires smarter, faster and more powerful audio-visual search engines. Media search of broadcast material involves database access, storage and very smart indexing & retrieval solutions. Search engines must

evolve to understand better the user's query (semantics).

To get there the best search engines are becoming both multimedia and metadata savvy. Multimedia search engines don't just handle text-based documents, as commercially done today, but also process search in natural language, in music recordings, in photo archives, in streamed video sequences, in live theatre recordings and possibly also multi-avatar virtual performances.

Provided a unified identifier is agreed, and interoperability exists, digital data objects of any kind will be easily modeled, captured, transferred and retrieved by multimedia search platforms. Multimodal and multilingual search will also expand the personalization of services for users (which is expected to be a big winner in terms of revenue and business growth).

Going beyond the internet search domain, the multimedia search technology finds equally important usage within enterprise-defined networks (Enterprise search), or as embedded service component within media rich applications (content enrichment). Search engines allow search of unstructured information which exists as phone call logs, emails, photos, etc, allowing use of materials that are not already organized in a specific database. Indexing programs enable networked PCs to organize knowledge, to understand multimedia as words, images, sounds, etc. Thus analysis of social patterns or relations may be inferred.

It is evident that the advances in Enterprise, Social, Mobile and Music search suggest a fundamental change of the users' needs in the way they search and consume information. Indeed, search is spreading broadly beyond its initial Web search foundation domain. Both socio economics and technology related aspects contribute to the differentiation of several market areas where search is spreading significantly.

In the following we highlight the major differentiating aspects shared by the five search domains: **Web search, Enterprise search, Mobile search, Music search and Social search.**

As already evidenced in the context of **Chorus**¹ published material, search is a set of techniques whose collective goal is to provide relevant answers to an unanticipated and ill formulated query. Because of this fuzzy characteristic, search systems tend to return to their users large quantities of potential results of varying relevance. By contrast, systems which tend to return a single or a small set of perfectly relevant results either fall in the category of data-base systems, or are so constrained by their data-set of indexing requirements that they pose no significant technological problem. In the light of this potentially large result set, it is rather clear that one of the major characteristics of search systems is the ranking method applied to results in order to present to the user the most relevant results first.

Early search engines included Alta Vista but **Web search** has since been dominated by Google, by Yandex (Russia) and by Baidu (China), in large part, thanks to its technology infrastructure and page rank algorithm. This algorithm classifies

¹ <http://www.avmediasearch.eu/>

results according to a relevance factor determined from the Web wide collective interest for a document shown by the presence of links to this document in other pages. Ranking is not the single reason for Google's success in this market (crawling, user interface, performance being others), but we will focus on this factor by contrast to its relevance, or differentiation power on other markets.

Enterprise search for instance cannot rely on a popularity based page rank algorithm to classify the documents returned following a query. Direct application of the Web search technologies to an enterprise intranet falls short of the problem on many aspects (document variety, multiplicity of sources), but in particular, the limited size of the user population and the professional context of the query forces a novel approach to ranking. This has led the actors of this sub-market to put much more significant emphasis on technologies such as semantic analysis and faceted results. One should note, as a supportive argument to the differentiating aspect of this sub-market, that Google is not one of its dominant players, and that Europe holds a strong position with Autonomy² and Exalead³.

Mobile search is yet another growing sub-market in which popularity based ranking comes second to the geographic proximity factor. Another major differentiating factor of this sub-market is the physical size of the device from which search is being conducted. Indeed, Web search can be performed from a GSM telephone, but the limited screen real-estate imposes a significant constraint on the system's ability to present large collection of results.

The most recent emerging sub-market in which search is taking a significant role is that of **Social networks** (we clearly

differentiate here «searching in a social network», which is a classical search problem, from «searching with a social network», which offers a new set of technology and usage opportunities). In this context, ranking is driven by parameters derived from the social network itself. Documents that my friends have recently seen, or have «liked» will rank higher than others. In this sub-market, *Social Indexing* can incorporate information about the structure and activity of the users' social network directly into the search and ranking process.

Similarly, in the music industry an increasing need for **music search** has emerged. Both consumers and professionals call for advanced methods for efficient retrieval of (frequently very specific) musical content from large audio databases. While the interest of private users is foremost to retrieve music matching their own taste, there are a number of scenarios for professional users, ranging from production to recommendation.

We have shown here that ranking is a major differentiating factor in four major search related sub-markets. The Chorus project, through its functional analysis of a search engine, has identified several additional major technological aspects of search which can also become differentiating factors for emerging sub-markets. In the remainder of this document, several such technological domains will be discussed, where progress can either lead to strengthening of existing activities, or emergence of new ones.

But a more general trend is appearing around search: its presence in most if not all applications involving **large quantities of data**. As such, search is disappearing as a specific application, but is becoming one of the key components of major applications. Enterprise search, with the emergence of «Search based

² <http://www.autonomy.com>

³ <http://www.exalead.com>

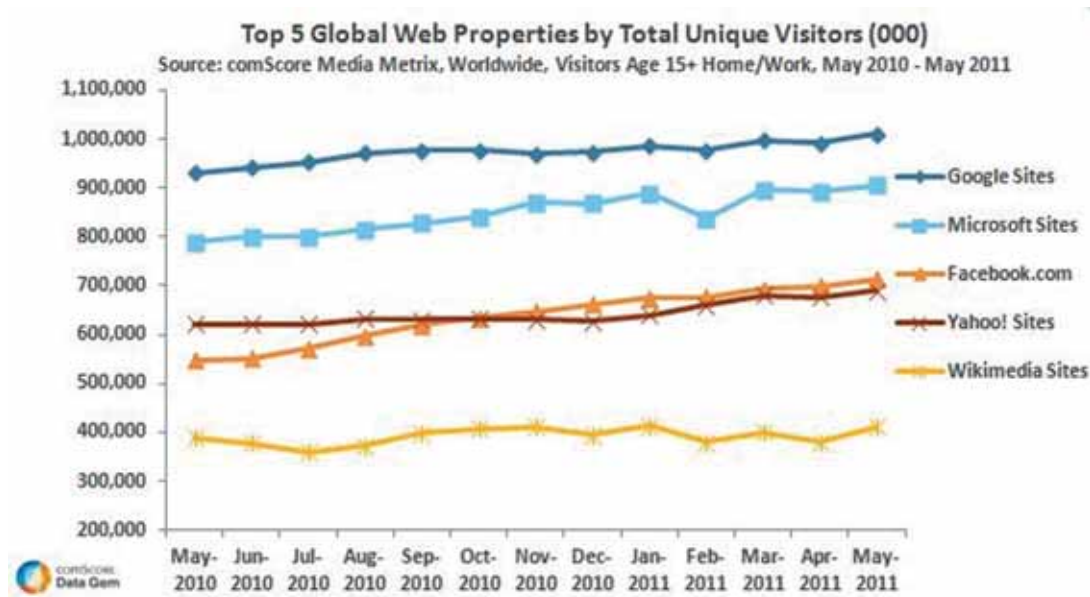
Applications» is clearly one of the sub markets where this trend is highly visible.

In this context the search component is one element of a broader application environment where data or information is not only «found», but also «acted upon» according to the needs of the enterprise. This trend, which is likely to spread across most domains, has led to the choice of the title of this document: “Search Computing”.

Since the internet economy relates to the global economy, and because the ICT sector is fundamental for Europe’s growth, research in the domain of “Search Computing” must be prominent. At present the topic is not mentioned in the Digital Agenda for Europe.

This document is a reflection on what's at stake and why.

The big research challenge is how to progress from today’s commercial text and language-based search engines to multimedia search engines that do more with less. How to achieve this technological breakthrough remains a fundamental and open question.



Europe's competitiveness in this important sector of ICT is hardly visible worldwide, as illustrated by the table below (only Wikipedia, as Europe’s brainchild, makes it as number seven).

Over the next years of experimentation and discovery several EU funded projects will contribute to this effort. These projects are listed at the end of this paper.

WORLD'S TOP WEBSITES (By Traffic Rank)

- 1. Google**
- 2. Facebook**
- 3. YouTube**
- 4. Yahoo!**
- 5. Blogger**
- 6. Baidu**
- 7. Wikipedia**
- 8. Windows Live**
- 9. Twitter**
- 10. QQ.COM (owned by Tencent)**
- 11. MSN**
- 12. Yahoo! Japan**
- 13. LinkedIn**
- 14. Taobao**
- 15. Sina**
- 16. Amazon**
- 17. World Press**
- 18. Yandex**
- 19. eBay**
- 20. Bing**

Source: Alexa Traffic Rankings, Alexa.com June 2011. Search engine providers in red.

The problem of multimedia search engines going beyond text will be even more important if current industry predictions hold true, namely that video traffic over the net will reach 91% of IP consumer traffic and 66% of mobile data traffic by 2014 (CISCO 2010).

3 Business areas

In this section, the specific business areas mentioned in the **Introduction** are described in more detail, focusing on their current status, specific technology and application requirements and the new needs arising as a result of these requirements.

3.1 Searching in Mobile business area

The primal objective of mobile search is to enable people to find either generic web or location-based information and services by entering a word or phrase on their phone.

An example of usage would be a person looking for a local hotel after a tiring journey or taxi company after a night out. With the years, mobile content has changed its media direction towards mobile multimedia. Starting with keyword-based search and going through the step of voice search, now the end user is offered the functionality to capture a photo in his cell phone and find relevant information on the Internet. High-end mobile phones have developed into capable computational devices equipped with high-quality color displays, high resolution digital cameras, and real-time hardware-accelerated 3D graphics. They can also exchange information over broadband data connections, sense location using GPS, and sense the direction using accelerometers and an electronic compass.

According to [Gómez10], in 2006, smartphones accounted only for 6.9% of the total market, while in 2007 the market segment reached 10.6%. The total annual sales of mobile devices reached 1,275 million units in 2008, with 71% of them

sold with data facilities, of which 15% (of total sales) correspond to smartphones. In Europe, 280 million units were sold in 2008, of which 19.3% were smartphones and 65.5% enhanced devices. It is evident that the camera-enhanced, hand-held devices are being spread at a very fast pace. Moreover, according to a leading market research firm eMarketer⁴, by 2011, mobile search is expected to account for around \$715 million. According to a recent study⁵ (April 2011) from Google conducted by the independent market research firm Ipsos OTX, among 5,013 US adult smartphone Internet users at the end of 2010, "71% of smartphone users search because of an ad they've seen either online or offline; 82% of smartphone users notice mobile ads, 74% of smartphone shoppers make a purchase as a result of using their smartphones to help with shopping, and 88% of those who look for local information on their smartphones take action within a day." Earlier this year, Performics predicted that mobile search would soon reach 10 percent of all the search impressions its clients were seeing. In the end of April 2011 the firm said that "mobile impressions accounted for 10.2 percent of all paid search impressions (desktop + mobile)⁶." Although certain obstacles still remain such as roaming charges (now under negotiation), these and other recent studies clearly show signs that mobile search is moving mainstream and gaining momentum. In addition to smartphones and GSMs, tablets are also emerging as the key

⁴ <http://www.emarketer.com/>

⁵

<http://googlemobileads.blogspot.com/2011/04/smartphone-user-study-shows-mobile.html>

⁶<http://blog.performics.com/search/2011/04/mobile-paid-search-impression-share-crosses-10-threshold.html>,<http://searchengineland.com/performics-mobile-impressions-cross-10-percent-threshold-74984>

access device. These popular gadgets will influence the way the web evolves and how users navigate and search for information stored somewhere (e.g. in the cloud).

Despite the many similarities, mobile search is not just a simple shift of PC web search to mobile equipment since it is connected to specialized segments of mobile broadband and mobile content, both of which have been fast-paced evolving recently [Pand10].

Mobile search involves a non-trivial processing/communication trade-off. In fact, even when the terminal is connected through a broadband mobile communication link, the possibility of transmitting huge amounts of media in upstream is usually not granted. On the other side, the on-board processing capabilities are limited by the nature of the processors, the memory, as well as by clock and power limitations. Different architectures can be adopted depending on the most acceptable compromise. In server-based approaches the service is completely provided by a remote server, which acts as repository, data processor and search engine. In this case, the terminal processing unit is almost inactive, while the communication channel has to transmit the full information needed for the search and to receive the results of the search. Client-based approaches use the reverse paradigm, limiting as much as possible the communication and doing all the processing on-board, keeping also a copy of the repository. In this case, the terminal is requested to do all the data processing and almost no transmission. Distributed systems perform part of the processing on-board (e.g., the feature extraction) and part on the server (e.g., the search on the remote repository). In this case there is a balance between communication and processing at the terminal. These three approaches

correspond to different services provided by the server, and therefore different business models for the service provider

Apart from location-based search, one of the most rapidly evolving segments of mobile search is image-based mobile search, which actually extends current location-based applications. Image-based mobile search works like traditional search engines but without having to type any text or go through complicated menus. Instead, users simply turn their phone camera towards the item of interest. Once the system recognizes the user's target it can provide further information (the menu or customer ratings of a restaurant) or services (reserve a table and invite friends for a dinner) [Takacs08], allowing to link between the physical and the digital world [Liu10]. Pointing with a camera provides a natural way of indicating one's interest and enables a new class of augmented reality applications which use the phone camera to initiate search queries about objects in visual proximity to the user. In this context, visual search has been extensively researched in recent years [Philbin07], integrating mobile augmented reality [Wagner08] and outdoor coordinate systems [Takacs08] with visual search technology. Although we are not aware of any figures about the size and dynamics of the image-based mobile search segment of the market, we can reasonably expect that the segment of mobile image search will scale proportionally to mobile search, creating new opportunities and offers. This is also advocated by the fact that major players in the mobile communication and search industry like Nokia and Google, are investing a lot of effort in the mobile image search concept and are aggressively trying to create applications and relationships in order to take advantage of the mobile ad market.

3.2 Searching in Social business area

During the last decade, Information Society witnessed the rapid growth of social networks that emerged as the result of users' willingness to communicate, socialize, collaborate and share content.

The outcome of this massive activity was the generation of a tremendous volume of user-contributed resources that have been made available on the Web, usually along with an indication of their meaning. Indeed, social media sharing networks such as Flickr⁷, Facebook⁸ and YouTube⁹ host billions of images and video, which have been annotated and shared among friends, or published in groups that cover a specific topic of interest.

In this context the notion of "search" takes a radical new shape, since apart from the traditional dimensions for searching such as textual affinity, visual similarity, etc, there is now a bunch of new dimensions available such as "friendshipness" (e.g., facebook's open graph), timestamps, geo-location, tag co-occurrence, etc.

It is evident that the knowledge we can acquire by exploiting the closeness of resources across different dimensions, may help us in overcoming the deficiencies that limit the performance of current search engines and reach closer to "semantic search". Although search mechanisms exist for searching within these very large social media collections, the search functionality largely ignores the social aspect, and solely leverages the textual annotations to make the content

searchable. For example, when a user searches for multi-media content, might very well appreciate relevant photos and videos of a user in her social network above that of an arbitrary user. Apart from such *Social Indexing and Ranking* mechanisms, social media can also be used to improve precision and recall of search results by tag propagation and refinement, for example, of visually similar photos or photos sharing similar tags.

Furthermore, the fact that users annotate and comment on the content in the form of tags, ratings, preferences etc and that these are applied on a daily basis, gives this data source an extremely dynamic nature that reflects events and the evolution of community focus. Although current Web 2.0 applications allow and are based on annotations and feedback by the users, these are not sufficient for extracting this "hidden" knowledge, because they lack clear semantics and it is the combination of visual, textual and social context of social media, which provides the ingredients for a more thorough understanding of media content. Moreover, with the success of the open data movement, the Internet is increasingly becoming a place of deeply interlinked knowledge bases, which further stresses the need for efficient searching with social networks. Therefore, there is a need for scalable and distributed approaches able to handle the mass amount of available data and generate an optimized 'Intelligence' layer that would enable the exploitation of the knowledge hidden within search applications.

⁷ <http://www.flickr.com/>

⁸ <http://www.flickr.com/>

⁹ <http://www.youtube.com/>

3.3 Searching in Enterprise business area

Enterprise search can be understood in two different fashions: a) as a tool deployed on the intranet of a corporation, offering to its employees a service similar to web search, but focused mainly on internal corporate information, b) as a tool embedded into a public service offered by the enterprise such as an e-commerce site or a yellow page service. Early on, those two aspects of Enterprise search were driven by different characteristics of the products, but with the emergence of Search Based Applications, these two domains have now merged into a single one. Search is now being integrated into both inward and outward facing applications, and is in fact disappearing as a stand alone application.

Enterprise search became an identified activity almost at the same time as web search. Developers of AltaVista for instance immediately deployed their solution across Digital Equipment's intranet (circa 1996/1997), and initiated an effort to turn the Web service into a corporate product. Such early attempts immediately brought to light several of the major differences between Web search and Enterprise search:

- **Security:** the corporate space cannot be viewed as a single open space such as the Web. Access control lists and security levels have to be taken into account both for crawling and for querying/results presentation. Implementation of security requirements as an afterthought on web search engines fails to meet the expected security level and must be taken into account at the earliest design stages
- **Diversity:** while the web is (theoretically) a fairly homogeneous space of documents adhering to a limited set of standards (html, doc,

pdf, jpg, ...) the enterprise space is much more diverse, both from the content format point of view (Enterprise products mention over 200 formats), and the repository point of view (file systems, mail archives, data bases of various forms, ...).

- **Precision/Recall:** while users querying the Web expect results in a "best effort" basis, the professional querying his corporate internal search service has much more stringent expectations on the results returned to him and often knows exactly what he is looking for; Thus enterprise search is less about information exploration and more about finding the necessary information (about products, people, sales, etc) in the corporate intranet.
- **Ranking:** by definition, search differs from data-base querying by the potential large amount of returned results due to the imprecise aspect of the query. On the Web, this may mean millions of results, and ranking becomes an essential part of the process, ensuring that the potentially best results are returned at the top of the list. The popularity based ranking of the Web does not apply in the Enterprise space. This has resulted into efforts and innovations on both the indexing side (deeper semantic analysis, context specific analysis, corporate ontologies, ...) and on the result presentation side (faceted results).

The most notable recent development of the Enterprise Search domain is that of **Search Based Applications**, and it is worth exploring somewhat deeper its characteristics.

Initially, search was seen (both on the web and in intranets) as a stand-alone application whose goal was to locate documents within a potentially very large repository. Given the diversity of

repositories present in enterprises (file systems, mail archives, document management systems, product life cycle management systems, data-bases of various kinds), search engines became over time a unified method for accessing all corporate information, irrespective of the origin repository. In this approach, each production silo maintained control over its specific tools optimized for content creation, and the unifying search engine became the privileged solution for read-only access to this information. This read-only characteristic allowed much more efficient solutions (compared to traditional data-bases) and opened the access to corporate information to potentially all employees. The current Search Based Application evolution recognizes the benefit of this approach and develops it further.

Finally, the difference between Web search and Enterprise search has opened an opportunity for competition in a domain in which Google seemed to be the potentially dominant player.

The Gartner Group yearly analysis of this market sector and the publication of their "Magic Quadrant" show all signs of an active and competitive market where Europe plays a significant role with companies such as Autonomy, Fast (now Microsoft) and Exalead. Enterprise search technology and services are key to increasing the competitiveness of the digital economy and as such can be seen as a strategic market for the European Union. Enterprise search solutions are becoming more ubiquitous (providing mobile access to enterprise resources from everywhere), social (using recommendations and guidance from colleagues to find what users are looking for), multimedia (provide capabilities not only to search for text but also images, audio, music, video,...), semantic (understanding better what the user is looking for), and certainly more widely deployed than ever.

3.4 Searching in Music business area

With the tremendous changeover of the music industry and music value chain to digital networked markets an increasing need for music search emerged.

Both consumers and professionals call for advanced methods for efficient retrieval of (frequently very specific) musical content from large audio databases.

While the interest of private users is foremost to retrieve music matching their own taste, there are a number of scenarios for professional users, ranging from production to recommendation. Common to all these scenarios is the limitations of traditional text-based search requests.

Textual search requires exact knowledge of the search criteria (artist, title, genre, etc.), which faces two problems: First, in many scenarios, the exact search criterion is not known (e.g. retrieval of a song

listened to in the radio, retrieval of a song of a specific mood, etc.). Second, in order to enable proper search results through textual search, music needs to be annotated, with an enormous manual effort [Joyce06]. The need for advanced music search and retrieval systems is also witnessed by the recent popularity of smartphone applications in the area of sound-based retrieval, recommendation and discovery.



This photo, is copyright (c) 2011 by photosteve101¹⁰ and made available under a Attribution-NonCommercial-Share Alike 2.0 license¹¹

For more than two decades, Music Information Retrieval (MIR) has been an increasingly popular research domain addressing the advanced needs around organizing, categorizing, searching in and accessing (potentially large) digital music libraries. Research includes both content- and context-based approaches. Content-based approaches perform a deep analysis of the musical content, in order to derive a semantic description of the music, capturing aspects such as loudness, tempo, beat, rhythm, timbre, pitch, harmonics, melody, etc. These descriptors (or “features”) are derived either through advanced signal and spectrum analysis methods (in the case of audio-based files – WAV, MP3, AAC, ...) or notation-based

statistics (in the case of symbolic notation-based formats – MIDI, MusicXML, ...). The MPEG-7 standard (ISO, 2002) includes a definition of audio descriptors. However, the methods for music description are constantly improved, with new tasks and approaches being devised (source separation, instrument recognition, beat detection, chord recognition, etc., to name but a few). Recent approaches also incorporate additional context of music, such as external information about the artist, related artists, lyrics, reviews, origin of artist or genre, tags by users (crowd-based annotations) etc. This context information is frequently retrieved from a multitude of Web resources in order to further increase the search experience.

As a consequence of this variety of approaches, the annual Music Information Retrieval Evaluation eXchange [MIREX11] provides a platform for the comparison and evaluation of state-of-the-art methods for music retrieval algorithms. Research groups and companies also provide frameworks and APIs for music analysis and retrieval e.g. MARSYAS [Tzanetakis02], CLAM [Amatriain06], MIRtoolbox for Matlab¹² or the APIs from MusicBrainz¹³, last.fm¹⁴ or echnonest.com¹⁵. The wealth of approaches in the context of Music Information Retrieval research led to the development of numerous solutions around music search: Identification of music titles is made possible through audio content analysis. The application of machine learning algorithms enables automatic classification of music and allows to detect songs from a particular artist, to recognize the genre of a piece of

¹⁰

<http://www.flickr.com/people/42931449@N07/>

¹¹

<http://creativecommons.org/licenses/by/2.0/>

¹²

<https://www.jyu.fi/hum/laitokset/musiikki/en/research/coe/materials/mirtoolbox>

¹³

http://musicbrainz.org/doc/XML_Web_Service/Version_2

¹⁴ <http://www.last.fm/api>

¹⁵ <http://developer.echonest.com/>

music, or its mood, or to organize an entire music library into a pre-defined genre taxonomy [Lidy05].

Content-based feature analysis also enables music search independent of textual annotations, by providing audio examples as query input: music can now be retrieved by example songs, excerpts of recorded audio, tapped rhythm or hummed melodies. Many of these music search and retrieval applications are now also used in mobile scenarios, as witnessed by popular smartphone applications such as Shazam¹⁶, SoundHound¹⁷, MoodAgent¹⁸, or discover¹⁹. Major work is also being performed in the area of visual search interfaces, through the application of clustering and unsupervised learning techniques that consider audio similarity independently from annotated genre categories. These approaches facilitate the organization of music libraries by automatic clustering of the music. Based on the clustering of music collections, interactive visual user interfaces are created in order to provide a new generation of search interfaces. More details about these approaches are elaborated in Section 4.4 about User Experience through new search Interfaces. [Downie03] provides a review of the Music Information Retrieval research domain. [Orio06] explains and reviews different music processing and music retrieval systems and gives an introduction to scientific MIR evaluation campaigns. A survey on automatic music genre classification is available in [Scaringella06].

4 Research challenges

Apart from the requirements and challenges specific to applications areas and sub-markets, there is a number of research challenges, which are horizontal and in most cases they result from the requirements of various areas.

While significant research has been conducted in most of these areas, there are still unsolved problems, where new research approaches are required. Others like affective-based search and content diversity have appeared more recently as promising research directions for existing search problems.

4.1 Multimodal

Moving beyond traditional text-based retrieval approaches, a lot of research has been conducted on developing methods for content-based multimedia retrieval.

The latter are based on the extraction of low-level features (e.g. colour, texture, shape, etc.) automatically from content.

While there are numerous content-based techniques that achieve retrieval of one single modality, such as 3D objects, images, video or audio, only few are able to retrieve multiple modalities simultaneously.

[Moutzidou10]

Cross-media retrieval, which has been introduced the latest years, comprises all content-based multimedia search methods that take as input a query of one modality to retrieve results of another modality. That is, given as query the image of a dog (2D), to be able to retrieve similar 3D (dogs) objects [Daras09].

¹⁶ <http://www.shazam.com/>

¹⁷ <http://www.soundhound.com/>

¹⁸ <http://www.moodagent.com/>

¹⁹ <http://www.discovermusic.com/>

Moving even further from cross-media retrieval, **multimodal** retrieval allows users to enter multimodal queries and retrieve multiple types of media simultaneously. This is a significant step towards content-based multimedia retrieval, since users will be able to search and retrieve content of any type using a single unified retrieval framework, not a specialised system for each separate media type. Moreover, through multimodal retrieval, users will be able to enter multiple queries simultaneously, thus, retrieve more relevant results. However, this is a highly complicated process, since it requires successful modelling of the low-level feature associations among the different modalities. The main challenges towards realising multimodal search, which could guarantee high-quality search services and improved end-user experience, are given below:

Multimodal content search and retrieval in a unified manner:

Currently, information is perceived, stored and processed in various forms leading to vast amounts of heterogeneous multimodal data (ranging from pure audiovisual data, to fully enriched media information associating also data originating from real world sensors monitoring the environment, etc.). User perception and interpretation of the information is in most cases on a conceptual level, independently of the form this content is available. Assuming the availability of an optimal, user-centric, search and retrieval engine, when users search for content they should be able to:

- express their query in any form most suitable for them;
- retrieve content in various forms providing the user with a complete view of the retrieved information;
- interact with the content using the most suitable modality for the particular user and under the specific context each time.

Sophisticated mechanisms for interaction with content:

Social and collaborative behaviour of users interacting with the content should be exploited at best, which will enable them to better express what they want to retrieve. We need to pay a great research attention on increasing the content utilisation efficiency, measured as the fraction of the relevant delivered content (i.e., content which satisfies their information needs and preferences) over the total amount of delivered content. Towards the direction of delivering relevant multimedia content to users, another barrier that needs to be overcome is that the vast amount of information is not actually annotated.

4.2 Affective

The affective dimension can be a fundamental component in a search operation, although it is currently neglected in search systems due to the difficulty in capturing, representing and exploiting such information.

Being able to extract the sentiment contained in a multimedia document (e.g., the fact that it may inspire positive or negative feelings, or that it refers to happy, sad, funny, annoying, hurting contents) and to match this information with the mood of the user or with a specific request contained in the query, can greatly enhance the search experience of the user.

Affective tagging of media includes a number of challenges such as defining suitable descriptors able to seize the emotional content of a media (e.g., based on the use of colours or other global characteristics of the image, facial expressions, composition of the scene, context, interpretation), defining metrics according to which the media can be matched or ranked on the basis of their affective description, establishing

adequate categories for sentiment clustering and classification.

Capturing and using the feeling of the user along the search/browsing process or in reaction to a search result can be another interesting aspect to be taken into account when designing the interfaces for next generation user-centred search engines [Vrohidis11].

In this case, the ability of analysing the mood of the user through multimodal input channels is a key challenge.

4.3 Event-based

It is commonly accepted that events are a good way to represent what happens in human's lives.

Events involve both the personal sphere (what is occurring to an individual, her family, her friends) and the social sphere (the big happenings that involve thousands or millions of people, such as sports, politics, culture, chronicles, etc.).

Event-based organization of data is already emerged as a key standard in the field of news and reporting (see Events-ML) and is rapidly extending to other professional fields [Glalelis10a].

Its extension to media will have a great potential in a number of application fields, including professional users (management of large media repositories, photo-reporters and freelance, user-generated content) and private users (social networks, thematic networks).

Event-based annotation and organization of media encompasses a number of technologies, including the representation and modeling for describing event-related knowledge, advanced event-related

tagging of media and automatic matching of models and data, development of new search engines and advanced interfaces able to cope with event-based media tagging, presentation tools providing event-structured browsing/navigation capabilities.

4.4 User experience

Facing the incredibly growing dimension of multimedia content not only advanced search techniques have been devised, but also new search interfaces, capable of coping with the content explosion and the need for advanced search methods. As multimedia search is by no means limited to textual queries, new generations of search interfaces have been created, in order to enable other forms of search input. Both in the image and audio domain, query-by-example based search is enabled through computation of similarities between the search items, frequently based either on content analysis or on user relations. Originally, in query-by-example based search systems in the first step the query example was retrieved through a textual search from the content database to be searched in.

More and more, this approach is moving towards "true" query-by-example, especially in the mobile scenarios, where images are searched online by a photo taken with the mobile phone, or music is searched through an audio sample recorded by the mobile phone's microphone.

Besides these new forms of search input – and the appropriate interfaces for them – there is another challenge being addressed: The problem of losing the overview of the incredible amount of available content. For instance, current mobile phones do not have large screens as seen on Apple iPad for example, making the search experience more challenging. The problem is no longer the

availability of content, the major issue is the access to the content – search is about finding the desired content in the shortest time possible – i.e. efficient access. New interfaces have been devised for the search and exploration of multimedia content that take care of these aspects: providing an overview while enabling the best possible way to retrieve the desired content within short time. One such example is 3D browsing in virtual spaces that allows search and retrieval of multimedia content in 3D environment possibly using glass-free screens as seen by Toshiba lately.

4.4.1 Visual Search Interfaces

The following is a review of visual search interfaces mostly in the music search domain – which is a good example of using visual search for another modality (in this case audio).

Frequently, in order to provide for overview even in large(r) media collections, a kind of mapping or projection is performed, from large dimensions to 2D or 3D interfaces. Examples applications are using FastMap algorithm for visualizing audio similarity [Cano02] or force-directed graphs depicting artist relations [VanGulik04]. Other examples are disc- and tree-map-based visualizations based on meta-data [Torrens04]. MusicRainbow is a visual search interface for discovering artists where similar artists are mapped near each other on a circular rainbow and colors encode different types of music [Pampalk06]. Musiccream [Goto05] is an interface, in which pieces of music are represented by discs that stream down on the screen. Disc colors indicate the mood of a piece and reflect similarity between musical pieces. A meta-playlist and sticking function enable visual playlist arrangement.

Among the mapping-based approaches Self-Organizing Maps (SOMs) have become very popular for providing a

topology-preserving mapping from high dimensions on a 2D space [Kohonen01]. By that, unsupervised clustering is performed, while maintaining relations within the data (e.g. preserving acoustic similarities on the 2D space which are derived from audio content analysis, see Section 3.4 on Music Search). SOMs have been applied to organize sound and music collections by acoustic relations [Cosi94], [Feiten94], [Spevak01], [Rauber01], [Pampalk04], [Mörchen05]. The mere mapping of music has been enhanced by visual representations of music clusters of a certain style (“Islands of Music”, [Rauber03]), mnemonic shapes (“Map of Mozart”, [Mayer06]) and a fully interactive interface for exploration of music collections and creation of playlists through visual interaction (“PlaySOM”, [Neumayer05]). Other variations apply SOMs to organize music at the artist level using artist information mined from the web [Knees04] or to enable live clustering of radio stations through analysis of the station’s program content [Lidy06]. The Globe of Music projects music titles on a virtual globe maintaining similarities between the titles [Leitich07]. The Music Thumbnailer visualizes musical pieces as thumbnail images which are created from acoustic analysis [Yoshii08]. Sonarflow creates groups of similar pieces of music into circular objects and provides a hierarchical visual interactive interface for quick access and playback of desired music content [Lidy10].

Visual search interfaces play a very important role on mobile devices, not only because of the limited screen sizes and input possibilities that make textual search tedious, but also because of new ways of interaction through touch interfaces that make search a truly playful user experience.

Recent smartphones provide already a new generation of visual interfaces for multimedia search and discovery,

exhibited by Apps such as cooliris²⁰, discover²¹, sonarflow²², Planetary²³ and others. A more scientific review of visualization approaches in the music domain is provided in [Cooper06].

The key technologies have been built in this area, yet the open challenge is how to combine the technical building blocks to a true user experience. Open (research) questions are how (i.e. by which criteria) to organize a media database to represent it to the user, which criteria to display and which forms of interaction to use in order to retrieve multimedia data within the shortest time possible. This is a topic of combined efforts in software engineering, information visualization, interaction design and usability engineering. Another key issue is to determine and combine the various sources of information in order to create better systems: E.g. do we need to extract more information from the media data, and/or rely on knowledge based systems, wisdom of the crowd, linked data or social networks to gather the information needed to create better retrieval systems? There is already much research effort to this end, yet the proper combination of information sources and interactive representations (visual and other) is an important key challenge to be addressed.

4.4.2 Augmented reality

Augmented reality is one of the most successful ambassadors for ubiquitous computing to date and is based on the idea that location-based data can be overlaid on your view of the real life.

Mobile phones are routinely equipped with GPS receivers, large displays, and quality cameras, which makes the notion of using your mobile phone to augment your own reality with location based information (essentially meta-tagging the real world), highly intriguing [Tish09].

In this context we have been witnessing a growing number of applications that rely on mobile phones to augment our view of the real world. Sekai Camera²⁴ uses positional data to determine the user's location, and, in turn, serving up augmented reality views of the world around him, TwittARound²⁵ provides an augmented reality Twitter view that uses the mobile phone's video camera to show live Tweets pop up based on the user location. Substantial effort has been also allocated on developing general scope augmented reality browsers such as Junaio²⁶ which is an advanced augmented reality browser that allows the augmentation of the camera view captured by the smart phone, with instant source of information about places, events, bargains or surrounding objects, Wikitude²⁷ which combines GPS and compass data with Wikipedia entries and overlays information on the real-time camera view of a smart phone, and Layar²⁸ which provides an open client-server platform that allows content layers to be developed as an equivalent to web pages in normal browsers.

The key research challenges identified in the field of Augmented Reality are primarily concerned with offering more personalized delivery of content and extending the capabilities of the

²⁰ <http://www.cooliris.com/>

²¹ <http://www.discovermusic.com/>

²² <http://www.sonarflow.com/>

²³ <http://planetary.bloom.io/>

²⁴ <http://sekaicamera.com/>

²⁵

<http://thenextweb.com/2009/07/13/twittaround-augmented-reality-twitter-app/>

²⁶ <http://www.junaio.com/>

²⁷ <http://www.wikitude.org>

²⁸ <http://www.layar.com/>

registration mechanism. In the majority of current augmented reality applications the meta-tags overlaid to augment the smart phone's camera view, appear indiscriminately for all users overwhelming them with more information than needed. This has led the researchers to work towards imposing a personalization filter on the displayed meta-tags, based on the user's interests. Similarly, GPS and compass information is currently the registration mechanism adopted from all augmented reality applications, so as to decide when to trigger the meta-tags. However, this registration mechanism is considered insufficient for supporting the full potential of Augmented Reality applications since it is impossible to locate objects without fixed GPS coordinates. To this end, current research efforts are being invested on extending the registration mechanisms to function not only based on the position and orientation of the camera's view, but also from the captured visual content (e.g., a product's brand, underground or railway logo, etc.) [Nikolopoulos11].


4.5 Large-scale indexing

The traditional process of indexing web content is to use various structures, such as lists, trees, inverted indices, etc. [Baeza-Yates99].

The main assumption made by existing indexing schemes is that the information need is the same for all users and thus apply for every user. Other techniques have emerged, especially those based on the construction of a user profile describing as much as possible the information needs of the user, allowing the system to better target search results. These techniques are taking into account the context of information needs, and

create a fundamentally new approach to information retrieval.

The existing techniques index generally the content as a bench of information



This photo, is copyright (c) 2010 by hanspoldoja²⁹ and made available under a Attribution-Noncommercial-Share Alike 2.0 license³⁰

with a use of a similarity measure. However, when the necessary scale grows up to indexing the Web traditional techniques are inadequate.

In order to overcome the existing limitations and tackle the large-scale indexing challenge, current efforts are primarily oriented towards achieving more targeted indexing. The goal of targeted indexing is to exploit the social relations, which may exist between people, objects, resources, etc. either in an explicit or an implicit form. In this way smarter indexes will be developed in the sense that will be targeted to people, groups, contexts, etc [Lajmi09], removing the need to index the Web in its full scale. Towards this end, the notion of social object is being researched as a possible index of social interactions. A social object can be a content that people can annotate or have conversation around and can serve as the means to index people interactions on social platforms and understand how to aggregate such interactions around an object.

²⁹ <http://www.flickr.com/people/hanspoldoja/>
³⁰

<http://creativecommons.org/licenses/by/2.0/>

Thus, one of the key challenges faced by large scale indexing is to enrich over time the “traditional” indexing methods using new forms of indexing based on sociality and achieve large scale indexing at the level of the Web.



Automatic tagging based on visual retrieval content identification with consequent picture to text recognition solution could help content providers to annotate on the fly with the most relevant keywords, their multimedia files.



These photos, are copyright (c) by: Alinari 24 ORE, Firenze, Italy

4.6 Network nodes

In many cases, searching may be extending in the network.

Content-aware nodes may be answer questions such as “where” and “how” in addition to the original “what”, that a search engine normally replies.

In this sub-section we describe some research challenges related to “discovery at the network” more than “searching in the search engine”.

4.6.1 Passive Content discovery

The wide majority of search engines discover content following a technique called active crawling.

When performing active crawling, the search engines’ robots follow all known URLs and references, trying to find and then index new content and services. This function is by definition post-hoc and may often take hours or even days to identify new content. Moreover active crawling can’t face the phenomenon of the “hidden web”; whereby certain portions of the web are hard to crawl because of their lack of linkage to popular web sites, or because they contain dynamically generated content.

By utilizing Content Aware network nodes with Deep Packet Inspection (DPI) capabilities on the content that flows via a content aware node, the node is able to initially discover new URLs and index new content (passive crawling) much faster than active crawling. As a result, it may significantly reduce the discovery time and provide search engines with a major competitive advantage, as it makes possible to find and index new content. Of course, regular active crawling at scheduled time intervals has still to be performed as a follow-up in order to ensure that the content remains

published and to reduce the number of broken or unavailable links returned by the search engine, thereby increasing its effectiveness and reliability.

Apart from content, the same advantages apply also for web services. There are already a number of service search engines (e.g. seekda) that perform active crawling of services (mainly based on WSDL headers). DPI inspection may significantly help in this direction.

4.6.2 Content/web-page popularity discovery

Conventional Search Engines estimate the importance (or relevance) of a particular matching web page typically based on two broad aspects: the content of the web page and the hypertext (or citation) structure of the surrounding web. The conventional search engine analyzes the contents of a particular web page and examines criteria such as the frequency of occurrence of the search terms, the location of the search terms (e.g. the title is more important than the appendix), the font size of the search terms relative to the font size of the surrounding text, the document format (e.g. certain file formats such as word processing files are usually more important than other file formats such as simple web pages), the web location of the document (e.g. documents on major web portals are more important than those on an individual’s web page). Each of these factors plays a role in determining the importance of a web page. Then the Search Engine exploits the hypertext link structure of the World Wide Web by viewing it as a citation index. Pages that are referred to (or linked to) by more pages are likely to be more important than pages that are linked to by fewer pages.

However, conventional search engines are not capable of monitoring how many times particular web pages URLs where actually visited (i.e., the popularity of a web pages) for use in determining the

importance of those web pages, although the actual number of visits to a web page would strongly indicate the importance of the web page. Conventional search engine merely estimate the importance of a particular matching web page based upon the content of the page and hypertext (or citation structure) of the surrounding web. They do not take into consideration the frequency of visits to the web page in estimating the importance of the web page. Furthermore, when propagating scores along the hypertext structure of the web, the score of a page is typically divided equally amongst the destination pages, rather than taking into consideration the relative popularity of the outgoing links from the page. Therefore there is a need for a method and system for monitoring and analyzing the current popularity of pages on a network and also there is a need for monitoring and analyzing the popularity of links between pages in hyperlink network.

Content-aware network nodes equipped with Deep Packet Inspection (DPI) capabilities may calculate the popularity of specific web pages (HTTP) and use such popularity information to rank the web pages retrieved in response to a search.

4.6.3 Network topology discovery

Network nodes with DPI functionality may further help in understanding the network topology by supporting the application of passive network tomography. However, without control plane traffic such as BGP traffic, it is not possible to directly observe topology information.

DPI boxes could on the other hand perform a hash on a packet (including its contents) to create a unique identifier for that packet. Such an identifier can then be sent to a Network Monitoring server (e.g. an extended ALTO server) at the Information Overlay in order to consistently trace the path that the packet is taking through the network as it

travels from source to destination. Initially, when only a few DPI boxes are deployed, only parts of the network will be visible: packets flowing through paths with no DPI boxes are invisible to the system. As more DPI boxes are added and more packets are analyzed in this way, the topology of the network will become clearer

4.7 Real-time

Human intelligence and perception are highly correlated with time awareness and perception. Much discussion is taking place in the web right now about “real-time search engines” that try to push information to users as fast as it is available. Real-time search pertains to the case when freshness of the results is of prime importance (obtaining results which were published a few instants ago, ranking results according to their publication time, etc). The main issue here relates to crawling and access to the publication stream, while the emphasis is given on the freshness of the results rather than on the accuracy, which is also an open research challenge (especially when the posts contain also multimedia). It should be noted that here real-time search primarily relates to the freshness of the results rather than returning the query results in real-time, which refers to the response time of the system.

There are already some attempts on real-time search, such as Google Instant Retrieval³¹ or the known social media search engines such as whostalkin³², Spy³³ or Socialmention³⁴. However, all these approaches are only text-based (keyword based) search engines that provide content mainly coming from twitter, friendfeed or other text-based information sources. Moreover, in the

³¹ <http://www.google.com/instant>

³² <http://www.whostalkin.com/>

³³ <http://spy.appspot.com/>

³⁴ <http://www.socialmention.com/>

case that they also serve multimedia content that is retrieved based on the surrounding text or titles. On the other hand, handling large volumes of multimedia and replying to an increased number of requests dictates the use of efficient mechanisms for caching queries and results as well as storing snapshots of the indexed collections. Since SQL based databases are not efficient for storage of such content a new trend on NoSQL databases emerged. Databases such as Cassandra³⁵ distributed database. Hadoop³⁶ or HyperTable³⁷ are well used.

Even if large data centres and fast indexing algorithms are used from the multimedia sharing web sites it is extremely difficult to cope with the huge volumes of requests without a caching mechanism. Nowadays, most of the large web sites that handle such requests (Facebook, twitter, etc.) use Memcached³⁸, which is a distributed object caching system. Real time search has to balance between quality, authority, relevance and timeliness of the content.

The main research challenges towards realising real-time search are: the ability to find updates, handle real time indexing of the content, execute real time matching and ranking algorithms, and calculate statistics and trends of the multimedia content.

4.8 Content diversity

There is a huge diversity of content and content providers in the Web in terms of formats, languages, topics, opinions. Despite the work done in Information Retrieval (IR), there is still lack of systematic approaches to make this diversity traceable and exploitable.

Current search engines in their ranking strategies tend to give much more importance to Web sites than to their content. In fact, regardless of the query and the topics covered, a few authoritative Websites - such as Wikipedia, Facebook, YouTube - tend to dominate the top positions in the ranked list of results while, because of the overload of information, is it practically impossible to explore the huge amount of documents in the long tail [Madalli11].

Their indexing strategies do not take into account the meaning of the words. For instance, since the word 'java' can mean at least (a) an island, (b) a programming language and (c) a beverage made from coffee beans, documents should be clustered accordingly. These three different topics/concepts can be addressed in different ways in different domains. For instance, Java the island can be discussed in terms of its geographical, geological, historical, political or cultural characteristics, i.e. by differentiating by context or domain [Madalli11]. In this respect, library science provides instead a solid background in the way the notions of topic and domain and formalized and used for knowledge organization (see for instance [Ranganathan65]). The universe of knowledge is carefully analyzed in its fundamental dimensions [Maltese09].

All these problems can be summarized as the impossibility of current tools to deal with diversity in knowledge, i.e. the ability to identify and exploit the aspects that differentiate a piece of information from another. In this respect, the Living

³⁵ <http://cassandra.apache.org/>

³⁶ <http://hadoop.apache.org/>

³⁷ <http://hypertable.org/>

³⁸ <http://memcached.org/>

Knowledge (LK) FET project³⁹ aims to exploit diversity in knowledge and in particular its diversity dimensions, i.e. the axes along which knowledge is differentiated. In Library Science, and more precisely in category based subject indexing systems, there is clear evidence of the existence of this multi-dimensional space of the universe of knowledge, where Topic, Space and Time constitute the three fundamental dimensions.

As first advocated in [Giunchiglia06], diversity in knowledge cannot be avoided, but has to be considered as a key feature with the aim to develop diversity-aware methods and tools for effective design by harnessing, controlling and using the effects of emergent knowledge properties.

This includes indexing, searching, displaying, clustering and aggregating along the diversity axes.

Recently, some attempts have been made in the direction of enhancing user experience. We can mention explorative search approaches (e.g. [Yee03] [White07]), where users can alternate search and browsing often with the support of faceted classifications. It is also worth to mention the approaches to diversification of search results, e.g. the work in [Skoutas10a] [Skoutas10b], and in particular those based on sentiment analysis for controversial topics [Demartini10]. However, even in these approaches there is a lightweight notion of topic and domain and no semantics is used.

4.9 Aggregation, mining and Linked Open Data (LOD)

Data streams coming from ubiquitous sensors are collected and aggregated in existing or newly emerging Social Networks. Aggregated sensor information becomes a source of valuable insights into emerging phenomena and events and enables real-time monitoring and extraction of actionable knowledge⁴⁰. Sensor data is weaved into mainstream social networks and social media platforms enabling open and transparent access to reality mining applications⁴¹. Sensor information becomes extremely commonplace due to widespread availability of smartphones and other mobile sensor-capable devices. Different applications emerge for recording and tracking measurements coming from these sensors⁴². Recorded data are made available through web APIs. Apple and Android phones are already known to collect and transmit location information from mobile phone owners. Several niche applications (e.g. Garmin Connect⁴³, EveryTrail⁴⁴) collect user-contributed sensory information a posteriori. Pachube⁴⁵ attempts to provide a common platform for sharing sensor data on the Web.

The key challenges arising from the need to aggregate and process sensory information can be seen from the following perspectives:

a) **Technology-Push:** The massive adoption of sensor information is expected to lead to massive explosion of data and new requirements for real-time

⁴⁰ <http://www.trendhunter.com/trends/best-photos-of-euro-2008-spain-vs-germany-finale>

⁴¹ <http://www.bbc.co.uk/news/business-13632206>

⁴²

<http://www.readwriteweb.com/hack/2011/06/runkeeper-opens-healthgraph-api.php>

⁴³ <http://connect.garmin.com/>

⁴⁴ <http://www.everytrail.com/>

⁴⁵ <http://www.pachube.com/>

³⁹ <http://livingknowledge-project.eu/>

sharing and processing. Fusion methods that are catered for combining vastly different modalities coming from different sources will be necessary to exploit the full potential of sensor information that is distributed across many sharing platforms.

b) **Application-Pull:** As long as the sensor data collection is not open, users will have limited confidence in making their sensor information available and the research community can do little to tap into the potential of this data.

c) **Deployment:** Several major players (Apple, Google) control the collection and sharing of sensor data. Successful mobile applications and devices will lead to an eco-system of “sensor data operators” that will control the generation, sharing and processing of user contributed sensor data.

Although sensor data processing has been a topic of intense research for several years, it is now gaining new perspective due to the huge scale of sensor data availability, its uncontrolled nature, its distribution mechanism, and the potential to combine the analysis results with information from different modalities (tweets, user contributed photos, videos).

This objective calls for discouraging proprietary and closed sensor data collection practices, encouraging interoperability and transparency to the end user and establishing clear guidelines and terms of data usage.

The dominant trend towards establishing those guidelines is the Linked Open Data (LOD) initiative. The aim of this initiative is to implement the vision of a Web of Data formulated by Tim-Berners Lee in which formerly fragmented data is connected and interlinked with each other based on the so-called Linked Data principles

[Berners-Lee09]. Since a few years from now, the so-called Linked Open Data (LOD) cloud which represents a huge interconnected data set is steadily growing. The key success factor of the LOD movement is the simplicity of its underlying principles. The main tasks that have to be performed in order to publish data as Linked Data are:

(i) to assign consistent URIs to data published [Sauermann08],

(ii) to generate links, and (iii) to publish metadata which allows further exploration and discovery of relevant datasets ([Cyganiak08], [Alexander09]).

The key research challenges that need to be tackled relate to data interlinking, to following the Linked Data dynamics and updates, as well as to the incorporation of trust and provenance information into the LOD cloud, which becomes urgently relevant as more and more user-generated content becomes part of it.

Among these challenges data interlinking is the most critical, aiming to draw connections between different chunks of data such as groups of related text objects (tags, tweets, sms), media objects along with their associated electronic traces (time, place) and user communities with their intrinsic characteristics (structure, dynamics) [Papadopoulos11a]. Current efforts are primarily oriented towards utilizing

a) **Embedded metadata:** Currently, social media platforms and mobile applications embrace vocabularies and standards that aim at describing semantically and/or giving structure to the social media content, creators and their user communities (e.g. FOAF, SIOC, Microformats, RDFa, Linked Data).

b) **Semantic web searches/indexes, like Sindice, Swoogle:** Through such services pieces of related information can be found across platforms which embed RDF

and Microformats and refer to content related to the issue at hand.

c) NLP & Ontology Mapping processes:

They can be used to capture related terms between the content (text or annotated media) and, thus, establish connections between data chunks that describe the same topic.



This photo is copyright (c) by: Alinari 24 ORE, Firenze, Italy

4.10 Standardisation

In the field of multimedia information retrieval, new standardisation initiatives are facing the market. The standards organizations most relevant for “Search Computing” are **ETSI**⁴⁶ and **MPEG**⁴⁷; the following paragraphs provide an overview on these standards.

In the context of Search Computing standardization challenges, ETSI is very present in mobile technologies, network technologies, multimodality, user requirements, standardization of content delivery, etc. ETSI is also well connected to European research and the European legislative framework, i.e. ETSI answers to EC Communications, Recommendations, Directives and Mandates either by

technical specifications or by Harmonized Standards. The goal of ETSI is to respond to the challenges and take up the opportunities offered by new technological developments, as well as to establish intensive links with academia and research bodies, in order to bring research and standardization together. ETSI is one of the key players in standardizing the new convergence of ICT and web technologies on a global scale, enabling the use of ICT in vertical businesses and ensuring quality, interoperability and time to market. ETSI consists of more than 750 members from almost 70 different countries and has links to more than 90 relevant standards organizations and industry fora. ETSI has released more than 30 000 publications, which are global technical standards and European Norms. In the following, some of the ETSI standardization areas will be highlighted, which are relevant for Search Computing.

ETSI is one of the founding partners of the **Third Generation Partnership Project (3GPP)**⁴⁸ in which come together five other standardization organizations from Asia and North America, plus market associations and several hundred individual companies. Established to develop globally applicable specifications for third generation mobile telecommunications (the ITU’s IMT-2000 family), 3GPP is also responsible for the maintenance and evolution of the specifications for the enormously successful Global System for Mobile communication (GSM), which was defined by ETSI, and for transitional technologies, including the General Packet Radio Service (GPRS), Enhanced Data for GSM Evolution (EDGE) and High Speed Packet Access (HSPA). 3GPP’s scope was later extended to develop radio access solutions beyond 3G, and thus encompasses LTE and its evolution towards true 4G technology, LTE-Advanced.

⁴⁶ <http://www.etsi.org>

⁴⁷ <http://www.mpeg.org/>

⁴⁸ <http://www.3gpp.org>

With respect to **Near Field Communications (NFC)** ETSI is looking at the implications of **Machine-to-Machine (M2M)** applications for smart cards. Standard SIMs have been used for specific M2M applications such as metering and device tracking for some time. Other applications may, however, require special functionality and different hardware properties as well as a new form factor. While some ETSI specifications are setup to deal with specific constraints such as data retention, temperature, memory update cycles, vibration resistance and humidity, others describe new form factors, and the physical and logical characteristics of the UICC. The use of confidential applications was developed to allow third-party applications to be loaded on a SIM and executed within a secure and private environment. This will be of particular interest to mobile NFC and M2M application providers who might often not own (or control) the platform onto which their application is loaded.

ETSI works with the **European Broadcasting Union (EBU)**⁴⁹ and the **European Committee for Electrotechnical Standardization (CENELEC)**⁵⁰ in a Joint Technical Committee, **JTC Broadcast**, which defines DVB IPTV standards for technologies on the interface between a managed IP network and retail receivers. JTC Broadcast recently completed the specifications the creation of a new form of media, "Hybrid Broadcast/Broadband" or, more simply, "hybrid broadcasting". Hybrid broadcasting combines the broadcast and broadband delivery of entertainment to the end-user through connected televisions and set-top boxes, offering the advantages and features of both delivery technologies. The JTC also made important revisions of the Multimedia Home Platform (MHP) specifications to facilitate their use by

IPTV (managed IP networks) both in Europe and throughout the rest of the world.

Historically, the standardization of broadcast and telecommunications has followed different paths, to meet differing commercial requirements. Recent developments in the Internet, mobile communications and broadcasting have led to a convergence, in which content delivery has become common ground. Commercial solutions developed by different market players do not currently interoperate across platforms. As a result, content providers face the costly challenge of providing different content formats to the various distribution pipes, whilst customers' buy-in remains below expectations. ETSI is producing standards and interoperability tools to enable content delivery across various distribution platforms, covering all the major elements of media delivery: the networked home, the content/service provider network, the content delivery network (CDN) and the Media Content Distribution (MCD) flow. The goal is the successful overall development of multimedia systems in television and communication, using managed and unmanaged networks to meet present and future market needs. Among other work the Technical Committee on MCD is developing the internal architecture of CDNs, which offer the end-user fast access to media content while optimizing network resources.

It is widely recognized that most deterioration in communications' Quality of Service (QoS) occurs at the network borders due to insufficient standardization of the interfaces. Thus, the **ETSI User Group** developed a Technical Report on end-to-end QoS management at the network interfaces which identifies standardization gaps. ETSI supports EC initiatives on consumer protection which aim to help users of electronic communications services take

⁴⁹ <http://www.ebu.ch/>

⁵⁰ <http://www.cenelec.eu/>

advantage of competition in the market. To gain maximum benefit in terms of choice, price and quality, users need reliable information on the QoS they can expect from the various offers, not just in the actual utilization of the service, but at any stage of the service life cycle.

Only standardization and global cooperation can provide full interoperability of "Search Computing" solutions.

ETSI enforces interoperability through a neutral platform, which offers interoperability testing in all fields of ICT to interested parties. To this end, ETSI Technical Committees carefully track the standard development of other organizations including MPEG Compact Descriptor Visual Search (CDVS) initiative.

The **MPEG Compact Descriptor for Visual Search (CDVS)** initiative is standardizing a low-level descriptor for visual search of images and videos with the aim to enable a complete set of new technologies to acquire a deep impact in the business of the future.

In a recent presentation on the 2nd Workshop on Mobile Visual Search [MobVisSear11], the relevance of the MPEG CDVS standard from a service provider point of view is presented. The buzz on the market perspectives for new technologies as augmented reality and visual search is increasing in the last years. There is an exponential increase in research publications in the academia and a flourish of new companies proposing services for searching of multimedia contents and Augmented Reality. In order to avoid a dispersion of resources and favour interoperability, the MPEG CDVS is proposing the standardization of a low-level descriptor for image. Four different scenarios are envisaged by the committee, ranging from an architecture A, where everything is done in the server and an image is sent to a remote server

for processing; to architecture B, where descriptors are extracted, encoded to reduce bandwidth and then sent to a server for remote processing; to architecture C, where everything is performed in the client. A mix of architecture B and C is generating a new architecture D, where most of the queries are processed on the device, but some others may be further processed on a remote server.

The committee envisages some relevant applications that may benefit from the standard. The first straightforward set of applications enables the mobile visual search of products and services in the Internet. Taking a photo of the product and sending the result to a search server may enable some comparison between different shopping offers or produce a more informed purchase, giving more details about the product or special offers. Mobile augmented reality may improve the quality of tourism, superimposing information about monuments and public events to images in the range of the mobile camera. Shops may inform the potential customers about special deals just pointing the device to the shop window. The TV and IPTV may change from the current static view of a list of channels to a more sophisticated way of organizing channels. The user may interact with the TV clicking to parts of the image relevant to him and in this way proposing to the system what kind of information and advertisement may be appropriate.

All these kind of applications and many others may be enabled by an interoperable solution. As databases that index a compliant set of features are created and proposed to business partners, new business models and companies may face the market. Consortia of European companies representing different aspects of the value chain may propose ecosystems for introducing in the market new services that may range from improving the

efficiency in the retail, to new tourism services, to new social services. The MPEG CDVS is an enabler of all these technologies, but it is just a basic brick that require users that find innovative ways of using it.

The MPEG committee released a roadmap for the image descriptor standardization⁵¹. In the 96th MPEG meeting the databases of testing images and software are made available and the final Call for Proposal was issued. In the 97th MPEG meeting a distractor set will be made available. This set of images would be mixed with the databases of testing images in order to ensure that only correct matches are return from searching algorithms under evaluation. The efficiency of the different proposals will be evaluated according to the specifications given by MPEG⁵². In the 98th MPEG meeting an evaluation of all proposals will be made and the committee will choose the best descriptor. On the 99th MPEG meeting the first Working Draft will be issued. The draft will be modified with contribution in the next meetings to become a Committee Draft (CD) in the 101st MPEG meeting, a Draft Internet Standard (DIS) in the 103rd meeting and a Final Draft Internet Standard (FDIS) in the 105th MPEG meeting. Hence the roadmap for standardization would conclude in July 2013 and there is time for European academia and industry to contribute for its success.

⁵¹

http://mpeg.chiariglione.org/working_documents/explorations/cdvs/cdvs-cfp.zip

⁵²

http://mpeg.chiariglione.org/working_documents/explorations/cdvs/cdvs_framework.zip

5 Socio-economic challenges

5.1 Business Models and Policies

The business models that are currently adopted in “search” can be distinguished between the models of intangible and monetary benefits [Zhao08]. According to the Intangible benefits model, free services are provided to users in exchange of their attention, loyalty and information. Then, the company can “sell” the attention, loyalty and information of users in exchange for money. For instance, advanced search can be used by a company as an attractive application for engaging more customers to their client base. Profit does not derive directly from the use of the searching functionality but from attracting more customers to use a paid service that incorporates this functionality as an additional feature. On the other hand the monetary benefits model comes up in the majority of relationships where a transaction or a subscription process takes place and customers are required to pay in exchanged of services of goods. This model is usually implemented through fixed transaction fees, referral fees, etc. With respect to “search” one aspect of the monetary benefit model is primarily based on advertising. In this case advertisements related to the content of the queries are displayed to the user, in way similar to Google Adds.

Advertising models include several very different business tactics. For instance, there could be off-portal campaigns for certain categories of services such as travel, restaurants, automotive, or consumer electronics to name a few. The traditional strategy consists of simply adding banner ads to search results, usually including a direct response method as well (a link to a microsite, a click-to-call link, or a short code). Another approach is to adopt business models that incorporate the revenue flow from the

application itself, therefore departing from the traditional pay-per-download. There are different business tactics here as well, which include time-based billing for services, event-based billing for specific situations or item-based billing as a function of the results obtained in the search.

If we want to further categorize the existing policies for generating revenue streams based on searching applications we can distinguish between the following:

1. Advertising based: Including a) advertising in general (i.e., like today's internet search), b) Merchandising (i.e., as a way to sell some other product service) or affiliation (i.e., to create opportunities of business for some other site), c) Advertising but based on some product placement (i.e., linked with another product: a TV show, a cinema premiere, etc) and d) User profiling (i.e., selling the user profiles for commercial purposes).
2. Packaging search with some other good or service: Including a) Packaged with the (voice, data) services of the mobile operator, b) Packaged with the mobile handset and c) Packaged with some other product or service not related to ICTs (a flight ticket, a hotel accommodation, a tourist pack, insurance, etc).
3. Premium service: Including a) Premium services (i.e., the basic functionality is free, but the advanced options not), b) Value-added services (i.e., a contract for a pack of services on top of usual ones), c) Pay-as-you-go (impulse purchase) and d) Subscription (monthly/annual fee, etc).
4. Other models: Including a) Business model to be defined at a very late stage when a critical mass of users is achieved (like Twitter today, for example), b) User community maintained by user contributions (like Wikipedia, for example) and c) Not a

commercial service (i.e., a public service).

Based on the above it is evident that current business models are inadequate to cover all different aspects of "Search Computing".

For instance, some of the key challenges that future business models will have to face include content ownership, copyright and licensing especially in cases where user generated content is involved.

Moreover, the fragmentation with respect to regulations and contacts with local businesses for providing mobile services and local advertising and the intellectual property of the technology (who owns what piece of technology and with what right to exploit) are two additional features that will have to be taken into account by the adopted business models and policies.

5.2 Search and open innovation

The notion of Business Ecosystem acquired notoriety from James's Moore's famous book "The death of competition: leadership and strategy in the age of business ecosystems" [MOORE06]. Moore defines a Business Ecosystem as "An economic community supported by a foundation of interacting organizations and individuals – the organisms of the business world. This economic community produces goods and services of value to customers, who are themselves members of the ecosystem."

A Business Ecosystem supersedes the concept of closed innovation, replacing it with the new paradigm of open innovation: "Open innovation is the use of purposive inflows and outflows of knowledge to accelerate internal innovation, and expand the markets for

external use of innovation, respectively. [This paradigm] assumes that firms can and should use external ideas as well as internal ideas, and internal and external paths to market, as they look to advance their technology.” [CHESBROUGH06]

Business ecosystems incorporate one or several value chains models (see Porter 1985), but in such a way that the business system is more than linear mechanical sequences of supplier/buyer relationships. The most relevant asset for a Business ecosystem is the “Community”. For a Business Ecosystem the same actors sometimes compete, sometimes cooperate. Prominent examples of Business Ecosystem are the Silicon Valley or the textile Clusters in Italy or the Champagne region in France.

In the last few years, Business Ecosystems has been discussed also within the Open Source Community (see Cedric Thomas

“Introduction to the OW2 Consortium Business Ecosystems Strategy - OW2 Consortium” [CEDRIC08]). The Open Source Community is a vibrant example of a Business Ecosystem.

With respect to a search paradigm, it is time to start to apply this concept if we are to create value and sustainability. Search computing can start to use the Business Ecosystem paradigm in the different search application domains.

To move forward from a Value Chain to a Business Ecosystem paradigm, involves the employment of a new open source framework and platform, a sustainable community of enterprises and institutions, sometimes collaborating and sometimes competing, but in both cases creating value for end users, their customers, themselves and each other.

A possible search computing Business Ecosystem is shown in figure 1, illustrating three main dimensions:

- B2B value chain with components developers in one hand and Search Applications designers on the other
- B2C value chain with Content providers and end users
- Standard bodies and international policy makers

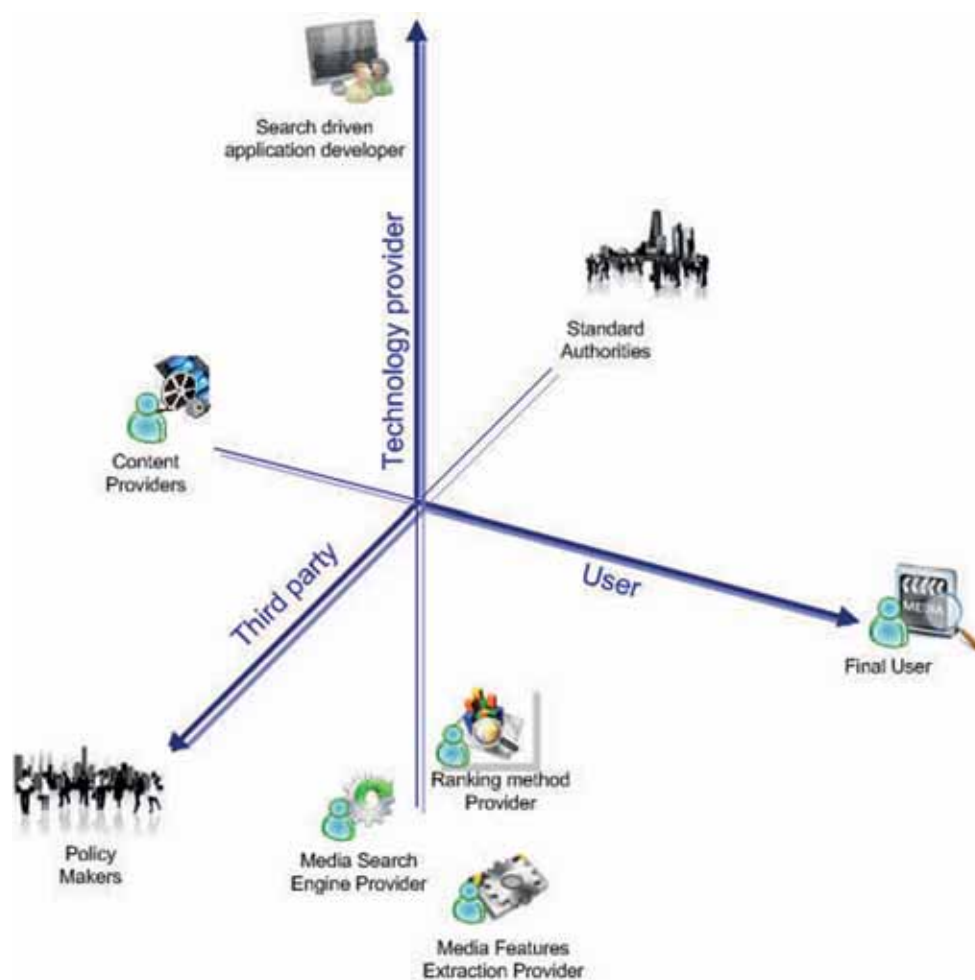


Figure 1 – A possible search computing Business Ecosystem

Search computing can facilitate many broad clusters of stakeholders in its ecosystem, including:

- Content providers (who may be content owners, publishers and/or distributors)
- Search application developers
- Component developers
- Standards bodies or public authorities
- Final users (who may be companies or consumers).

Examples of use cases are presented in the following user scenarios:

Target users	User scenarios
B2B Business to Business	A SW company would like to integrate a new search function inside a big portal (e.g. a system integrator for new Portal of a broadcaster)
	A search company improves the search services for a video archive, providing vertical video and TV search, including third party services for video processing and annotators.
	Two SW companies will work together to integrate in an open source BI application already distributed a new search function
B2B2P Business to Business to Professionals	A movie-maker needs to create a short documentary on a specific topic, e.g. the end of the 2nd World War in Rome region. Search archives of historical records to be integrated into the documentary, requiring an audiovisual search engine.
	A journalist working for a large broadcaster would like to integrate shots and short videos for a news segment on Palestine and Israel.
	A professional journalist would like to find similar Content (an audio interview or a short video)
	A finance professional would like to receive early warning of possible market risks or opportunities, and be prompted by the system to check the relevant video/audio/html news feeds where critical messages were discovered
B2B2C Business to Business to Consumers	A tourist would like to produce a short home made movie after his last Rome vacation (to be uploaded to YouTube) and he would like to search and include some shots from the web of Rome sightseeing.
	A stamp collector would like to find stamps similar to stamps in his collection to compare stamps in the market.
	A French electro-acoustic composer would like to know if in Ireland there are some interest groups on that kind of music and exchange some music files with them.
	An expert on traditional Spanish music would like to exchange information with others who appreciate the same genre.

5.2.1 Search Ecosystem at European Level

To have impact, search business ecosystems require a European level strategy, for the following reasons.

1. First search computing is fully in line with the realization of the vision of the Digital Agenda 2020, which is an European vision of ICT with impacts on the economic and social life of the whole European community.
2. Second, search business ecosystems require an integration of several areas of expertise, approaches, and application domains in spanning the full complexity of the challenges it faces. A European level strategy is needed to guarantee excellence which would not have been possible on a national level.
3. Third, the EU countries have different cultures, and as a consequence the impact and adoption of technologies can be different. This diversity should be used to gather different scientific approaches and business requirements and to exploit extensions and applications that would not be easily obtained at the national level. Global competitiveness requires a European-wide effort that cannot be achieved at the regional or national level only.
4. Fourth, a number of European search projects have already gained significant results relevant to Search Computing and the community that they have gathered so far should be used in future as baseline to start up a search computing ecosystem..

5.2.2 Search computing and user generated innovation

For search computing to become a success the role of users must be central to the collaborative needs elicitation, cooperative solution generation and negotiation of service implementation. This user-centered perspective represents a strategic element in order to achieve a major impact at the social level. Active user communities/social networks should be involved on any strategic decision through moderation tools for large-group dialogues and visual analytics and with tools and services for community and crowd management. This is already visible in some of the larger integrated projects.

Search is not limited to everyday experience of searching the web but it is a “must” feature of business intelligence application for enterprises and archives, where Content is increasingly becoming Multimedia: a combination of text, audio, still images, animation, video and interactivity content into a single form.

Enterprises, news agencies, TV broadcasters, advertising agencies are using digital archives everyday and approx. 50% of applications for such archives requests to include a search facility as a primary interface for end users. A user innovation process should be an open process in which the different actors will cooperate, collaborate and compete in a productive kind of rivalry, providing more, and more valuable, innovation to the entire search computation business ecosystem.

With respect to traditional value-chains, based on the creation of classic “monetary” value, business ecosystems are value chains enriched with “culture” and “non-monetary” values. As explained by the above-cited economics textbooks from an economic point of view, this implies that business relationships are not just “supplier-buyer” and “monetary” relations, but are interactions based on

mutual interests and shared cultural values, resulting in long-term relationships. In practical terms, this means to reinforce the role of users by involving them at a very early stage of the innovation process. Only a very early commitment of the users can ensure this. This vital process is known as “democratizing innovation” [VONHIPPEL05].

This process will turn traditional users into “user-innovators”. Nowadays, innovation is rapidly becoming democratized. Given the advances in computer and communications capabilities, users of products and services – both firms and individual consumers – are increasingly able to innovate for themselves. User-centred innovation processes offer great opportunities to business ecosystems, first of all the possibility to benefit from innovations developed and freely shared by others in the open ecosystem.

However, this emerging process of user-centric, democratized innovation will have to be supported by all the different communities involved: from users and citizens, to government and policy makers, to business and software service developers.

5.3 Benchmarking: A catalyst to advance the state of the art in multimedia search computing

Benchmarks are valued for their ability to streamline research by eliminating redundancy, enabling direct performance comparison between algorithms, increasing efficiency by sharing resources between research sites and providing a concrete framework in which researchers interact in a productive mixture of competition and collaboration.

This contribution discusses the benefits of benchmarking and explains the value of benchmarking for advancing the state of the art in multimedia search computing.

5.3.1 What is a benchmark?

A benchmark is a forum that organises tasks for the research community. Researchers are invited to develop algorithms that address the tasks. In general (according to the so-called Cranfield paradigm [Cleverdon62]), a benchmarking task consists of three parts: 1. A task definition that describes the problem to be solved 2. A data set provided to the benchmark participants 3. Ground truth against which participants’ algorithms are evaluated. Often benchmarks follow a yearly cycle that moves through a series of steps: identifying critical research challenges, developing specific tasks to address these challenges, releasing data, receiving results submitted by participants, evaluating and reporting results, convening participants in a workshop to discuss results and start planning the next year.

5.3.2 What are the benefits of benchmarking?

The benefits of benchmarking fall into two categories: benefits for the research community [Smeaton06], [Tsikrika11] and economic benefits gained by bringing innovative research closer to market [TREC10]. Benchmarking benefits the research community by reducing fragmentation of research agendas and by promoting efficient use of resources. Without benchmarking initiatives, researchers develop solutions working isolated in their labs. Alone, they must identify innovative, high impact problems on which to concentrate their efforts and suffer from a lack of access to the larger, overarching perspective. They must create their own data sets from the ground up in order to evaluate their

results. Individual datasets developed at individual research sites involve not only expending duplicate effort, but, more importantly, block the possibility of cross-site comparison. Comparison and reproducibility are two factors that drive forward the state of the art: if researchers know exactly where they stand with respect to the state of the art they can better direct their efforts to surpass it and more quickly abandon less promising lines of investigation. Benchmarking creates momentum: A researcher can turn to the benchmarking community for support in overcoming practical roadblocks, which are otherwise potentially both discouraging and time consuming. Finally, the mixture of competition and collaboration within a benchmarking initiative provides motivation for the group and makes the payoff of the invested effort immediately and clearly visible, both with and beyond the research community. Sharing data sets not only means that results are better comparable but by combining efforts much larger data sets can be created, improving the significance and generalizability of results.

The economic benefits of benchmarks lie in their ability to coordinate efforts more closely between the research community and the needs of businesses bringing products to markets. In the initial phase of the benchmarking cycle where tasks for the upcoming year are planned, industry can contribute ideas for new tasks. Further, industry can also make datasets available that will allow researchers to focus on real world versions of specific tasks. In some industries such competitions are publicly announced challenges with a prize to win such as a quality improvement in vide recommendation created in [Netfix06]. In this way, algorithms developed within the benchmark are already closely linked to real-world business needs and are better suited for industrialization.

In order to build capacity in business, it is necessary not only to innovate solutions, but also to formulate new problems. A solution to a new problem can provide the basis of a new product with the power to open up new market sectors. The flexibility and efficiency of benchmarking initiatives makes them uniquely suited to quickly address innovative new tasks and to come to a well-supported conclusion about the feasibility of addressing those tasks backed up by a large research community that is participating. Working together with a benchmark, a company is able to explore a new area that would normally be too risky or too expensive for it to tackle itself in addition to the difficulty of having state-of-the-art knowledge in the domain within the company.

Recently, explicit effort has been devoted to measuring the benefits of benchmarking. This effort has been led by large benchmarks sponsored by the US National Institute of Standards and Technology (NIST)⁵³. In particular, the scholarly impact of the TRECvid⁵⁴ Video Retrieval Evaluation has been demonstrated by a bibliographic study [Smeaton06] and a similar study has now been started for the ImageCLEF benchmark [Tsirikika11]. Further, an extensive report is now available on the economic impact of Text REtrieval Conference (TREC) Program [TREC10]. An analysis of the impact of benchmarking data sets also in the long run can be found in [Sanderson10]. Such results surely also extend to benchmarking initiatives that are organized from within Europe.

In remainder of this section, we turn specifically to the European perspective and discuss the importance of benchmarking initiatives for Europe, focusing on two key examples of multimedia benchmarking, MediaEval and ImageCLEF.

⁵³ <http://www.nist.gov/>

⁵⁴ <http://trecvid.nist.gov/>

5.3.3 Benchmarking to reinforce Europe's competitiveness

Europe's ability to compete internationally in ICT is dependent on the success of research initiatives that cross country boundaries. The MediaEval benchmark [Larson11] is a prime example of an initiative that is able to leverage and focus research carried out at individual sites in individual countries, helping to ensure the efficient use of research investment. The MediaEval initiative is sponsored by the PetaMedia Network of Excellence⁵⁵ and the European Institute for Innovation and Technology's ICT Labs⁵⁶, two entities dedicated to fostering innovation and excellence and for bringing research results closer to market. MediaEval is an ongoing initiative and current information can be found on the website. The MediaEval initiative allows for flexible participation – research sites can easily join the initiative and participate in tasks as dictated by their research needs and current capacity. The inclusive and dynamic organization of MediaEval is designed to help dissolve impediments to spontaneous cross-border cooperation and allow Europe to more effectively act as a single, unified research community.

International benchmarking campaigns initiated and coordinated from within Europe make it possible to promote research topics that are critical for Europe within the international research community. It is clearly important for European researchers to take part in international research initiatives whose organization is based abroad. However, in order for Europe to maintain its defining role in international research agendas, it is necessary not only to participate in benchmarks, but also be actively involved in planning and organizing them. One particular aspect is the existence of ICT research areas that are of more central

importance in Europe than they are in Asian or America. Europe can retain its key position by continuing to be the central international initiator of benchmarking efforts in these areas. An example is the importance of developing technologies that make it possible to access multimedia content resources across cultures and independently of language. Here, Europe's own Cross-Language Evaluation Forum (CLEF)⁵⁷ has shown the value of bundling research effort in the area of text and image retrieval. Multilinguality and cultural differences are important in a European context and Europe is clearly a leader in these fields of information retrieval.

MediaEval was originally established within the CLEF campaign and subsequently split off in order to become a stand-alone benchmark focused on multimedia in social contexts. Speech and language issues remain central for MediaEval. Unlike other multimedia initiatives, such as TRECVID, MediaEval is decoupled from investment from governments beyond Europe and in particular from US defense and intelligence spending. Although international cooperation is clearly important, in order to preserve the diversity of the goals pursued within the international research community, it is important to ensure the existence of independent initiatives that are fully free to focus on research areas central to creating more prosperous lives for the citizens and Europe and increased capacity for its businesses.

Similarly to MediaEval, ImageCLEF⁵⁸ [Müller10] started within the CLEF benchmark, but has more focused on classical aspects of multimedia retrieval and on visual information analysis. It has remained part of CLEF despite its orientation towards visual media and language independence. Automatic

⁵⁵ <http://www.petamedia.eu/>

⁵⁶ <http://eit.ictlabs.eu/>

⁵⁷ <http://www.clef-campaign.org/>

⁵⁸ <http://www.imageclef.org/>

annotation of images as well as retrieval combining visual and textual features means including cross-language retrieval has remained an important aspect. ImageCLEF has mainly been a voluntary effort of many people, and the group has evolved and shifted composition over time. Since 2010 a partial support has been provided to ImageCLEF by the Chorus⁵⁹ and Promise⁶⁰ projects. ImageCLEF has run over ten different tasks since its start in 2003 and has raised the number of registered participants from 4 in 2003 to over 110 in 2011.

5.3.4 Challenges in benchmarking

Despite the clear advances of many science domains through benchmarks there are also critical voices stating that benchmarks can block innovation. When a benchmark places too much importance on pure performance and not on novel technologies, participants are motivated to make only slight modifications to existing techniques and not develop something very novel. The critic is clearly valid to the degree to which such effects can be observed in a benchmark. To offer participants alternative incentives and encourage innovation, it is important to stress the collaboration rather than comparison by scores alone. In this way, the goal of “winning” can be transformed into a more subtle, more productive motivator.

Providing alternative incentives, means that benchmarks should continuously change and modify the tasks, so that plenty of room is available for participants to innovate.

MediaEval follows this concept of radically novel tasks such as the current focus on social media. New tasks are developed by discussion and consensus within the community and only offered if group interest reaches the critical mass

necessary for cross-site comparison. ImageCLEF has major changes for each task at least every three years to avoid being technologically locked-in but also maintaining a large body of participants with a long-term focus on specific issues. The impact analysis [Smeaton06] also shows that tasks become cited less usually after three or four years, losing their impact, so it is important to maintain the appropriate balance.

Several issues currently present serious and interesting for benchmarks, for example, the extremely large datasets that may be difficult to distribute to participants. Large data sets also mean that data analysis becomes increasingly sparse and to find clear trends and compare techniques becomes even harder. Another critical point is the division of the yearly circle of benchmarks in phases. This organization structure gives researchers on the one hand a clear frame to focus on, but on the other hand it also blocks researchers to test new tools and techniques and their impact instantly. “Continuous benchmarking”, i.e., ongoing evaluation, would make possible to carry out tests automatically or semi automatically at any time, and could make the adaptation of systems much quicker and more effectively. Component-based evaluation is another important topic for the future as currently mainly entire systems are tested, whereas every retrieval system contains a large number of existing components such as image analysis or text retrieval tools. Measuring the overall performance can hide many important results and allowing to measure performance of components can help understanding the interactions and implications of these components.

In conclusion, we can claim that benchmarking benefits both the research community and industry. Moreover, not only are these benefits potentially high, they are also high-value since the efficiency of benchmarking means that they can be achieved at relatively low

⁵⁹ <http://avmediasearch.com/>

⁶⁰ <http://www.promise-noe.eu/>

costs. European efforts in the organization of benchmarks are necessary in order to secure the Europe's international leadership in its key research area. The MediaEval and ImageCLEF benchmarking initiative make a contribution to this goal and give forum for the research community as well as to industry. The currently existing challenges for benchmarking raise the possibility that concerted effort towards the development of next-generation benchmarking practices and infrastructures will provide payoff in the form of even more effective exploitation of European benchmark potential.

5.4 Legal and ethical issues

Most experts worldwide agree that if personal data and user context were available, search engines' results would be much more targeted and efficient. This approach however, hinders multiple legal and ethical issues that should be carefully considered.

Especially in the case of social networks, proper usage and processing of the personal data should be ensured.

5.4.1 EU Legal framework for data protection, privacy and search

The individual's right of privacy was established with Article 12 of the Universal Declaration of Human Rights⁶¹ in 1948 and with Article 8 of the European Convention for Protection of Human Rights and Fundamental Freedoms⁶² in

⁶¹ Universal Declaration of Human Rights, G.A. res. 217A (III), U.N. Doc A/810 at 71 (1948)

⁶² European Convention for the Protection of Human Rights and Fundamental Freedoms, ETS 5; available from: <http://conventions.coe.int/Treaty/Commun/QueVolezVous.asp?NT=005&CM=8&DF=4/25/2006&CL=ENG>

1950. Since then, there was a broad debate in several European states on the use of electronic data processing by public authorities that led to the initial national data-protection laws.

In 1995 the EU adopted Directive 95/46/EC on the Protection of Individuals with regard to the Processing of Personal Data and on the Free Movement of such Data⁶³ in order to harmonize Member States' laws in providing consistent levels of protections for citizens and ensuring the free flow of personal data within the European Union. It established the basic principles for the collection, storage, and use of personal data that should be respected by governments, businesses and any other organizations or individuals engaged in handling personal data. This Directive (Data Protection Directive) is the reference text, at European level, on the protection of personal data. It sets up a regulatory framework which seeks to strike a balance between a high level of protection for the privacy of individuals and the free movement of personal data within the EU. To do so, the Directive sets strict limits on the collection and use of personal data and demands that each Member State set up an independent national body responsible for the protection of these data.

According to 95/46/EC, personal data can be accessed only by authorised personnel for legally authorised purposes. Personal data stored or transmitted must be protected against accidental or unlawful destruction, accidental loss or alteration, and unauthorised or unlawful storage, processing, access or disclosure, by a security policy with respect to the processing of personal data. In particular, **listening, tapping, storage or other kinds of interception or surveillance of**

⁶³ Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data, OJ L 281, 23.11.1995, p. 31–50

communications and the related traffic data by persons other than users, without the consent of the users concerned, is prohibited except technical storage which is necessary for the conveyance of a communication without prejudice to the principle of confidentiality.

Due to the rapid technical progress of information and communication technologies and the resulting new ways of generating and analyzing personal data extracted from web activities and Search Engines, the existing regulations were increasingly regarded as insufficient. In 1997 the European Union supplemented the 95/46/EC Directive by introducing the Telecommunications Privacy Directive **97/66/EC concerning the Processing of Personal Data and the Protection of Privacy in the Telecommunications Sector**⁶⁴. This directive established specific protections covering telephone, digital television, mobile networks and other telecommunications systems. New advanced digital technologies have been recently introduced in public communications networks in the Community, which give rise to specific requirements concerning the protection of personal data and privacy of the user.

Directive 97/66/EC has been repealed and replaced by the **Directive 2002/58/EC concerning the Processing of Personal Data and the Protection of Privacy in the Electronic Communications Sector**⁶⁵ (Privacy and Electronic Communications Directive, also known as E-Privacy Directive). This Directive forms part of a

new legislative framework for electronic communications networks and services, the main aim of which is to make the electronic communications sector more competitive. This framework consists of one framework Directive and four specific Directives, namely the: “Framework Directive”⁶⁶; “Authorization Directive”⁶⁷; “Access Directive”⁶⁸; “Universal Service Directive”⁶⁹; “Privacy and Electronic Communications Directive”.

The Directive 2002/58/EC has been amended by the Directive **2009/136/EC** amending Directive 2002/22/EC on universal service and users’ rights relating to electronic communications networks and services, Directive 2002/58/EC concerning the processing of personal data and the protection of privacy in the electronic communications sector⁷⁰. This Directive forms part of the telecoms

⁶⁶ Directive 2002/21/EC of the European Parliament and of the Council of 7 March 2002 on a common regulatory framework for electronic communications networks and services, OJ L 108, 24.4.2002, p. 33–50

⁶⁷ Directive 2002/20/EC of the European Parliament and of the Council of 7 March 2002 on the authorization of electronic communications networks and services, OJ L 108, 24.4.2002, p. 21–32

⁶⁸ Directive 2002/19/EC of the European Parliament and of the Council of 7 March 2002 on access to, and interconnection of, electronic communications networks and associated facilities, OJ L 108, 24.4.2002, p. 7–20

⁶⁹ Directive 2002/22/EC of the European Parliament and of the Council of 7 March 2002 on universal service and user’s rights relating to electronic communications networks and services, OJ L 108, 24.4.2002, p. 51–77

⁷⁰ Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009 amending Directive 2002/22/EC on universal service and users’ rights relating to electronic communications networks and services, Directive 2002/58/EC concerning the processing of personal data and the protection of privacy in the electronic communications sector and Regulation (EC) No 2006/2004 on cooperation between national authorities responsible for the enforcement of consumer protection laws (Text with EEA relevance), OJ L 337, 18.12.2009, p. 11–36

⁶⁴ Directive 97/66/EC of the European Parliament and of the Council of 15 December 1997 concerning the processing of personal data and the protection of privacy in the telecommunications sector, OJ L 24, 30.1.1998, p. 1–8

⁶⁵ Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications), OJ L 201, 31.7.2002, p. 37–47

reform package that also includes Regulation (EC) No 1211/2009⁷¹, and Directive 2009/140/EC⁷² and went into force with its publication in the EU's Official Journal (18 December 2009). Its transposition into national legislation in the 27 EU Member States must be ensured by June 2011. The above mentioned amendments are out of the scope of this document, as they have not yet been implemented into national law by Member States. The European Commission will propose in 2011 a new general legal framework for the protection of personal data in the EU covering data processing operations in all sectors and policies of the EU⁷³. Search engine and social networks owners should monitor any changes in this field, and adapt accordingly their practices.

⁷¹ Regulation (EC) No 1211/2009 of the European Parliament and of the Council of 25 November 2009 establishing the Body of European Regulators for Electronic Communications (BEREC) and the Office (Text with EEA relevance), OJ L 337, 18.12.2009, p. 1–10

⁷² Directive 2009/140/EC of the European Parliament and of the Council of 25 November 2009 amending Directives 2002/21/EC on a common regulatory framework for electronic communications networks and services, 2002/19/EC on access to, and interconnection of, electronic communications networks and associated facilities, and 2002/20/EC on the authorization of electronic communications networks and services (Text with EEA relevance), OJ L 337, 18.12.2009, p. 37–69

⁷³

<http://europa.eu/rapid/pressReleasesAction.do?reference=MEMO/10/542>

6 Conclusions

It becomes obvious that search plays an important role in a number of very active sub-markets and business areas. High-end mobile phones have developed into capable computational devices and for many users world-wide are the main device for connecting to the Internet, searching and consuming information. Social Networks are constantly expanding and not only pose new challenges for searching in such vast amount of data but also create new opportunities for enhanced searching through Social Indexing and for data mining in user generated content. Similarly specific “*Search Computing*” challenges have been presented for Enterprise and Music search and other vertical markets.

While a lot of research approaches have been investigated and applied to various aspects of “*Search Computing*”, the problems are far from being solved and new challenges arise as a result to recent developments (e.g. Social Networks). In the Introduction section, it is described in detail as an example how the ranking problem, which could be considered solved in web-based search, is totally different in other sub-markets, with issues unsolved and with new opportunities in Enterprise, Mobile, Music and Social Networks search.

Applicable to these markets but also horizontally to many other sectors, are the enabling technologies and challenges described in Section 4. It seems that a common problem to be solved by all these technologies is the limitations of traditional text-based search requests. Textual search is very restricted in the way queries are formulated by using only one of the available modalities of content, while in many cases, visual, audio, 3D and other types of information is available. Approaches in multimodal search try to overcome these problems. Moreover, existing textual search does not take into account aspects such as affective-based indexing and search and content-aware networks etc, which play an important role in efficient retrieval. Content annotation still remains a problem unsolved in many generic applications, due to the well-known Semantic Gap, and new research in areas such as event-based representation and indexing and content diversity offer new opportunities.

The availability of huge amount of data, which are very dynamically updated, especially in Social Networks and user generated content pose new requirements for research in large-scale and real time processing and architectures. In Mobile area, Augmented Reality applications require new approaches for visual-based search and presentation. In all these novel approaches, user experience and user interfaces are important factors, which should be taken into account in order to deliver technology in user-friendly and acceptable way.

As already mentioned in the Introduction section, search in many applications is not the final or the only component in a service or product. It has become an important step and component in larger systems where retrieved information is further processed in order to extract additional information, to be used then by processes or to be presented to the user. For example, geo-tagged images retrieved for a specific city or area can be further analysed (clustered) in order to extract important locations (landmarks) in this area [Papadopoulos11b]. Instead of presenting the list of images to the user, the processed new information (landmarks) is presented or provided as input to tourism application. Again large-scale and real time processing and architectures play an important role in such applications, together with approaches aggregating various sources of information, for example coming from social networks, web and Linked Open Data. Finally, standardization always plays an important role for large-scale technology adoption, and as new visual-based

technologies are applied to mobile devices, new standards like MPEG *Compact Descriptor for Visual Search* (CDVS) are under development, to give an example.

All these technologies cannot be applicable and transferred to real products without the appropriate socio-economic support and activities. The appropriate business models are always needed for technology transfer and product success. As new technologies arise, new requirements and opportunities are generated that have to be studied and examined. Similarly, a Business Ecosystem is needed to start applying the concept of open innovation in the search applications domain. In the same direction, benchmarking plays an important role for the research community and produces economic benefits by bringing innovative research closer to market. Finally, as more research approaches and products involve user-generated content collection, storing and processing, legal and ethical issues should be carefully considered for proper usage and processing of personal data.

The overall conclusion of this White Paper is that “*Search Computing*” is an active area, with related technologies needed in a number of markets and applications. Many issues are still unsolved and new ones have appeared, formulating a long list of research challenges (summarised in the Executive Summary section).

7 BIBLIOGRAPHY

- [Alexander09] Alexander, K., Cyganiak, R., Hausenblas, M., Zhao, J. (2009). Describing Linked Datasets. Proceedings of the 2nd Workshop on Linked Data on the Web (LDOW2009).
- [Amatriain06] Amatriain, X., Arumí, P., and Garcia, D. (2006). CLAM: A framework for efficient and rapid development of cross-platform audio applications. In Proceedings of ACM Multimedia 2006, Santa Barbara, CA.
- [Baeza-Yates99] R. Baeza-Yates, B. Ribeiro-Neto (1999) Modern Information Retrieval. Addison Wesley Longman Publishing Co. Inc.
- [Berners-Lee09] T. Berners-Lee „Linked Data – Design Issues“, 2009; <http://www.w3.org/DesignIssues/LinkedData>
- [Blandin10] Work, D, Blandin, S, Piccoli, B, Bayen, A (2010). A traffic model for velocity data assimilation. Applied Mathematics Research eXpress.
- [Cano02] Cano, P., Kaltenbrunner, M., Gouyon, F., and Battle, E. (2002). On the use of fastmap for audio retrieval and browsing. In Proceedings of the International Conference on Music Information Retrieval (ISMIR), Paris, France.
- [Cattuto05] Cattuto, C.: Collaborative tagging as a complex system. talk given at international school on semiotic dynamics. In: Language and Complexity. Erice (2005).
- [Cedrico08] Cedric Thomas (2008) “Introduction to the OW2 Consortium Business Ecosystems Strategy” - OW2 Consortium
- [Chang02] S.-F. Chang, “The holy grail of content-based media analysis,” IEEE MultiMedia, vol. 9, no. 2, pp. 6–10, 2002.
- [Chesbrough06] Henry Chesbrough , Wim Vanhaverbeke, Joel West (2006) “Open Innovation: Researching a New Paradigm” – Oxford University Press
- [Cleverdon62] C. W. Cleverdon, Report on the Testing and Analysis of an Investigation into the Comparative Efficiency of Indexing Systems, Cranfield College of Aeronautics, Cranfield, England, 1962.
- [Cooper06] Cooper, M., Foote, J., Pampalk, E., and Tzanetakis, G. (2006). Visualization in audio-based music information retrieval. Computer Music Journal, 30(2):42–62.

- [Cosi94] Cosi, P., De Poli, G., and Lauzzana, G. (1994). Auditory modelling and self-organizing neural networks for timbre classification. *Journal of New Music Research*, 23(1):71–98.
- [Cyganiak08] Cyganiak, R., Delbru, R., Stenzhorn, H., Tummarello, G., Decker, S. (2008): Semantic Sitemaps: Efficient and Flexible Access to Datasets on the Semantic Web. *Proceedings of the 5th European Semantic Web Conference (ESWC2008)*
- [Daras09] P. Daras, A. Axenopoulos, "A 3D Shape Retrieval Framework Supporting Multimodal Queries", SPRINGER, *International Journal of Computer Vision*, DOI 10.1007/s11263-009-0277-2, Jul 2009
- [Demartini10] Demartini, G. and Siersdorfer, S. 2010. Dear search engine: What's your opinion about...?: Sentiment Analysis for Semantic Enrichment of Web Search Results. In the Workshop of WWW on Semantic Search (April, 2010), Raleigh, North Carolina.
- [Downie03] Downie, J. S. (2003). *Annual Review of Information Science and Technology*, volume 37, chapter Music information retrieval, pages 295–340. Information Today, Medford, NJ.
- [Feiten94] Feiten, B. and Günzel, S. (1994). Automatic indexing of a sound database using self-organizing neural nets. *Computer Music Journal*, 18(3):53–65.
- [Giunchiglia06] Giunchiglia, F. 2006. Managing Diversity in Knowledge. Conference on Artificial Intelligence (ECAI), Trento, September 2006. Lecture Notes in Artificial Intelligence. Online presentation, from <http://www.disi.unitn.it/~fausto/knowdive.ppt>
- [Glalelis10] N. Gkalelis, V. Mezaris, I. Kompatsiaris, "A joint content-event model for event-centric multimedia indexing", *Proc. Fourth IEEE International Conference on Semantic Computing (ICSC2010)*, Pittsburgh, PA, USA, September 2010
- [Gómez10] Gómez-Barroso, J.-L., Compañó, R., Feijóo, C., Bacigalupo, M., Westlund, O., Ramos, S., et al. (2010). Prospects of mobile search EUR 24148 EN. Seville: Institute for Prospective Technological Studies. European Commission
- [Goto05] Goto, M. and Goto, T. (2005). Musicream: New music playback interface for streaming, sticking, sorting, and recalling musical pieces. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, London, UK.
- [Joyce06] Joyce, J., (2006). Pandora and the Music Genome Project, *Scientific Computing*, 23 (10): 14, pages 40- 41
- [Knees04] Knees, P., Pampalk, E., and Widmer, G. (2004). Artist classification with web-based data. In *Proceedings of the International Conference*

- on Music Information Retrieval (ISMIR), Barcelona, Spain.
- [Kohonen01] Kohonen, T. (2001). Self-Organizing Maps, volume 30 of Springer Series in Information Sciences. Springer, Berlin, 3rd edition.
- [Lajmi09] S. Lajmi, J. Stan, H. Hacid, E. Egyed-Zsigmond, P. Maret: Extended Social Tags: Identity Tags Meet Social Networks. CSE (4) 2009, pp. 181-187
- [Larson11] Larson, M., Soleymani, M., Serdyukov, P., Rudinac, S., Wartena, C., Murdock, V., Friedland, G., Ordelman, R. and Jones, G.J.F. Automatic Tagging and Geotagging in Video Collections and Communities. ACM International Conference on Multimedia Retrieval (ICMR 2011).
- [Leitich07] Leitich, S. and Topf, M. (2007). Globe of Music - Music library visualization using GeoSOM, In Proceedings of the International Conference on Music Information Retrieval (ISMIR), Vienna, Austria.
- [Lidy05] Lidy, T. and Rauber, A. (2005). Evaluation of feature extractors and psycho-acoustic transformations for music genre classification. In Proceedings of the International Conference on Music Information Retrieval, pages 34–41, London, UK.
- [Lidy06] Lidy, T. and Rauber, A. (2006). Visually profiling radio stations. In Proceedings of the International Conference on Music Information Retrieval, Victoria, Canada.
- [Lidy10] Lidy, T. (2010) Sonarflow: Visual Music Exploration and Discovery. In Proceedings of the International Society for Music Information Retrieval Conference (ISMIR), Utrecht, The Netherlands.
- [Liu10] Xu Liu, Jonathan Hull, Jamey Graham, Jorge Moraleda, Timothee Bailloeuil. 2010. Mobile Visual Search, Linking Printed Documents to Digital Media. In CVPR 2010 Demonstrations.
- [Madalli11] D. P. Madalli, B. Dutta, V. Maltese. Exploiting the Diversity in Search using the Facet Dimensions. At the 1st International Workshop on Eternal Systems (EternalS), 2011 (to appear).
- [Maltese09] Maltese, V., Giunchiglia, F., Denecke, K., Lewis, P., Wallner, C., Baldry, A. and Madalli, D. 2009. On the interdisciplinary foundations of diversity. In the first Living Web Workshop at ISWC (2009).
- [Mayer06] Mayer, R., Lidy, T., and Rauber, A. (2006). The Map of Mozart. In Proceedings of the International Symposium on Music Information Retrieval (ISMIR), Victoria, Canada.
- [MIREX11] Music Information Retrieval Evaluation eXchange (MIREX). Website. http://www.music-ir.org/mirex/wiki/2011:MIREX_Home.
- [MobVisSear11] The 2nd workshop on Mobile Visual Search, Wenjin Hotel, Beijing,

- China, January 22nd 2011, <http://idm.pku.edu.cn/mvs/Program.asp>.
Present:
http://idm.pku.edu.cn/mvs/doc/Addressing%20MVS%20exploitation.%20A%20service%20provider%20perspective%20and%20an%20introduction%20to%20MPEG%20CDVS_Italia_Lab.pdf
- [Moore96] James Moore (1996) "The death of competition: leadership and strategy in the age of business ecosystems" - Harper Business New York
- [Mörchen05] Mörchen, F., Ultsch, A., Nöcker, M., and Stamm, C. (2005). Databionic visualization of music collections according to perceptual distance. In Proceedings of the International Conference on Music Information Retrieval (ISMIR), London, UK.
- [Moutzidou10] A. Moutzidou, A. Dimou, N. Gkalelis, S. Vrochidis, V. Mezaris, I. Kompatsiaris, "ITI-CERTH participation to TRECVID 2010", in Proceedings of TRECVID 2010 Workshop, November 2010, Gaithersburg, MD, USA
- [Müller10] H. Müller, P. Clough, T. Deselaers, B. Caputo, ImageCLEF – experimental evaluation in visual information retrieval, Springer, 2010.
- [Netflix06] The Netflix Prize on video recommendation. <http://www.netflixprize.com/>, 2006.
- [Neumayer05] Neumayer, R., Dittenbach, M., and Rauber, A. (2005). PlaySOM and PocketSOMPlayer – alternative interfaces to large music collections. In Proceedings of the International Conference on Music Information Retrieval (ISMIR), pages 618–623, London, UK.
- [Nikolopoulos11] S. Nikolopoulos, G. Th. Papadopoulos, I. Kompatsiaris and I. Patras, "Evidence driven image interpretation by combining implicit and explicit knowledge in a bayesian network", IEEE Transactions on Systems, Man and Cybernetics – Part B: Cybernetics, 2011, DOI: 10.1109/TSMCB.2011.2147781
- [Orio06] Orio, N. (2006). Music Retrieval: A Tutorial and Review. Foundations and Trends in Information Retrieval. 1(1):1–90.
- [Pampalk04] Pampalk, E., Dixon, S., and Widmer, G. (2004). Exploring music collections by browsing different views. Computer Music Journal, 28(2):49–62.
- [Pampalk06] Pampalk, E. and Goto, M. (2006). MusicRainbow: A New User Interface to Discover Artists Using Audio-based Similarity and Web-based Labeling. Proceedings of the International Symposium on Music Information Retrieval (ISMIR), pages 367-370, Victoria, Canada
- [Pand10] Het Pand, Notes of the Exploring the Future of Mobile Search Expert Workshop 9th June 2010, Ghent, Belgium

- [Papadopoulos11a] S. Papadopoulos, Y. Kompatsiaris, A. Vakali, P. Spyridonos. "Community Detection in Social Media". In Data Mining and Knowledge Discovery, 2011, DOI: 10.1007/s10618-011-0224-z
- [Papadopoulos11b] S. Papadopoulos, C. Zigkolis, Y. Kompatsiaris, A. Vakali. "Cluster-based Landmark and Event Detection on Tagged Photo Collections". In IEEE Multimedia Magazine 18(1), pp. 52-63, 2011
- [Philbin07] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman, "Object retrieval with large vocabularies and fast spatial matching," Proc. of the IEEE Conference on Computer Vision and Pattern Recognition, 2007.
- [Porter85] Porter, Michael E., "Competitive Advantage". 1985, The Free Press. New York.
- [Ranganathan65] Ranganathan, S. R. 1965. The Colon Classification. In the Rutgers Series on Systems for the Intellectual Organization of Information, S. Artandi (etd.), IV. Graduate School of Library Science, Rutgers University, New Brunswick, NJ.
- [Rauber01] Rauber, A. and Frühwirth, M. (2001). Automatically analyzing and organizing music archives. In Proceedings of the European Conference on Research and Advanced Technology for Digital Libraries (ECDL), Darmstadt, Germany. Springer.
- [Rauber03] Rauber, A., Pampalk, E., and Merkl, D. (2003). The SOM-enhanced JukeBox: Organization and visualization of music collections based on perceptual models. Journal of New Music Research, 32(2):193–210.
- [Sanderson10] M. Sanderson. Test Collection Based Evaluation of Information Retrieval Systems. Foundations and Trends in Information Retrieval, Vol. 4, No. 4, (2010) 247–375
- [Sauer mann08] Sauer mann, L., Cyganiak, R. (2008): Cool URIs for the Semantic Web. W3C Interest Group Note. Available at: <http://www.w3.org/TR/cooluris/>
- [Scaringella06] Scaringella, N., Zoia, G., and Mlynek, D. (2006). Automatic genre classification of music content: a survey. Signal Processing Magazine, IEEE, 23(2):133–141.
- [Shen06] Z. Shen, K.-L. Ma, and T. Eliassi-Rad, "Visual analysis of large heterogeneous social networks by semantic and structural abstraction," IEEE Trans. Vis. Comput. Graph., vol. 12, no. 6, pp. 1427–1439, 2006.
- [Skoutas10a] Skoutas, D., Alrifai, M. and Nejdl, W. 2010. Re-ranking web service search results under diverse user preferences. In 4th Int. Workshop on Personalized Access, Profile Management, and Context

Awareness in Databases (PersDB).

- [Skoutas10b] Skoutas, D., Minack, E. and Nejd, W. 2010. Increasing diversity in web search results. In WebSci10: Extending the Frontiers of Society OnLine.
- [Smeaton06] F. Smeaton, P. Over, and W. Kraaij, 2006. Evaluation campaigns and TRECVID. In MIR '06: Proceedings of the 8th ACM International Workshop on Multimedia Information Retrieval, pages 321-330, 2006.
- [Smeulders00] A. W. M. Smeulders, M. Worrying, S. Santini, A. Gupta, and R. Jain, "Content-based image retrieval at the end of the early years," IEEE Trans. Pattern Anal. Mach. Intell., vol. 22, no. 12, pp. 1349-1380, 2000.
- [Spevak01] Spevak, C. and Polfreman, R. (2001). Sound spotting - a frame-based approach. In Proceedings of the Second International Symposium on Music Information Retrieval: ISMIR 2001, pages 35-36, Bloomington, IN, USA.
- [Takacs08] Gabriel Takacs, Vijay Chandrasekhar, Natasha Gelfand, Yingen Xiong, Wei-Chao Chen, Thanos Bimpigiannis, Radek Grzeszczuk, Kari Pulli, and Bernd Girod. 2008. Outdoors augmented reality on mobile phone using loxel-based visual feature organization. In Proceeding of the 1st ACM international conference on Multimedia information retrieval (MIR '08). ACM, New York, NY, USA, 427-434.
- [Thornley11] Thornley, C. V., Johnson, A. C., Smeaton, A. F., Lee, H., 2011. The scholarly impact of TRECVID (2003-2009). J. Am. Soc. Inf. Sci. 62 (4), 613-627.
- [Tish09] Shute, Tish. "Is It 'OMG Finally' for Augmented Reality? Interview with Robert Rice." UgoTrade: Virtual Realities in "World 2.0." Retrieved August 1, 2009.
- [Torrens04] Torrens, M., Hertzog, P., and Arcos, J. L. (2004). Visualizing and exploring personal music libraries. In Proceedings of the International Conference on Music Information Retrieval (ISMIR), Barcelona, Spain.
- [TREC10] National Institute of Standards and Technology, 2010. Economic Impact Assessment of NIST's Text REtrieval Conference (TREC) Program. Final Report, July 2010.
<http://trec.nist.gov/pubs/2010.economic.impact.pdf>
- [Tsikrika11] T. Tsikrika, A. Garcia Seco de Herrera, H. Müller, Assessing the Scholarly Impact of ImageCLEF, Cross Language Evaluation Forum (CLEF 2011), Amsterdam, The Netherlands, 2011.
- [Tzanetakis02] Tzanetakis, G. (2002). Manipulation, Analysis and Retrieval Systems for Audio Signals. PhD thesis, Computer Science Department,

Princeton University.

- [VanGulik04] Van Gulik, R., Vignoli, F., and van de Wetering, H. (2004). Mapping music in the palm of your hand, explore and discover your collection. In Proceedings of the International Conference on Music Information Retrieval (ISMIR), Barcelona, Spain
- [Vonhippel05] Eric Von Hippel (2005) "Democratising Innovation" - MIT Press,
- [Vrohidis11] S. Vrochidis, I. Kompatsiaris and I. Patras, "Utilizing Implicit User Feedback to Improve Interactive Video Retrieval". In Advances in Multimedia, Hindawi, vol. 2011, Article ID 310762, 18 pages, 2011. doi:10.1155/2011/310762
- [Wagner08] D. Wagner, G. Reitmayr, A. Mulloni, T. Drummond, and D. Schmalstieg, "Pose tracking from natural features on mobile phones," Proc. of the 7th IEEE/ACM Int. Symp. on Mixed and Augmented Reality (Sept. 15 - 18, 2008), pp. 125-134.
- [White07] R. W. White, M. Drucker, G. Marchionini, M. Hearst, M.C. Schraefel, 2007. Exploratory search and HCI: designing and evaluating interfaces to support exploratory search interaction. In CHI 2007 Extended Abstracts on Human Factors in Computing Systems, New York: ACM Press
- [Yee03] K.P. Yee, K. Swearingen, K. Li and M. Hearst, 2003. Faceted metadata for image search and browsing. In Proceedings of Conference on Human Factors in Computing Systems, pp. 401-408.
- [Yoshii08] Yoshii, K. and Goto, M. (2008). Music Thumbnailer: Visualizing Musical Pieces in Thumbnail Images Based on Acoustic Features. Proceedings of the International Symposium on Music Information Retrieval (ISMIR), pages 211-216, Philadelphia, USA.
- [Zhao08] H. Zhao, "Emerging business models of the mobile internet market," M.S. thesis, Helsinki University of Technology, Espoo, Finland, 2008.

8 CONTRIBUTING PROJECTS AND ORGANIZATIONS



CHORUS+
<http://www.ist-chorus.org/>

GLOCAL
www.glocal-project.eu



Cubrik
www.cubrikproject.eu



I-SEARCH
<http://www.isearch-project.eu>



WeKnowIt
<http://www.weknowit.eu/>



SocialSensor
www.socialsensor.eu



COAST
<http://www.coast-fp7.eu/>



Living Knowledge
<http://livingknowledge-project.eu/>



PHAROS
www.pharos-audiovisual-search.eu



PetaMedia
<http://www.petamedia.eu>



ETSI
<http://www.etsi.org>



Networked Media and Search Systems Unit of DG INFSO
http://cordis.europa.eu/fp7/ict/netmedia/home_en.html

9 EDITORS AND AUTHORS

Editors

Name	Affiliation	e-mail
Spiros Nikolopoulos	CERTH-ITI	nikolopo@iti.gr
Yiannis Kompatsiaris	CERTH-ITI	ikom@iti.gr

Authors

Name	Affiliation	e-mail
Loretta Anania	European Commission	loretta.anania@ec.europa.eu
Henri Gouraud	INRIA	Henri.Gouraud@LaPoste.net
Petros Daras	CERTH-ITI	daras@iti.gr
Symeon Papadopoulos	CERTH-ITI	papadop@iti.gr
Francesco Denatale	UNITN	denatale@ing.unitn.it
Vincenzo (Enzo) Maltese	UNITN	maltese@disi.unitn.it
Stavri Nikolov	JRC-IPTS	Stavri.NIKOLOV@ec.europa.eu
Thomas Lidy	Vienna University of Technology, Austria	lidy@ifs.tuwien.ac.at
Martha Larson	Delft University of Technology	m.a.larson@tudelft.nl
Henning Müller	University of Applied Sciences Western Switzerland (HES-SO)	Henning.mueller@hevs.ch
Luca Celetto	STMicroelectronics S.r.l.	luca.celetto@st.com
Theodore Zahariadis	Synelixis Solutions Ltd	zahariad@synelixis.com
Menelaos Perdikeas	Synelixis Solutions Ltd	mperdikeas@synelixis.com
Ioannis Koufoudakis	Synelixis Solutions Ltd	ikoufoudakis@synelixis.com
Asterios Platskos	Synelixis Solutions Ltd	aplatskos@synelixis.com
Gaby Lenhart	European Telecommunications Standards Organisation (ETSI)	Gaby.Lenhart@etsi.org
Patrick Guillemin	European Telecommunications Standards Organisation (ETSI)	Patrick.Guillemin@etsi.org
Francesco Saverio Nucci	Engineering SpA	francesco.nucci@eng.it
Claudia Cosoli	Engineering SpA	claudia.cosoli@eng.it
Vincenzo Croce	Engineering SpA	vincenzo.croce@eng.it
Andrea de Polo	Alinari 24 ORE	andrea@alinari.it

European Commission

Search Computing: Business Areas, Research and Socio-Economic Challenges Media Search Cluster White Paper

Luxembourg: Publications Office of the European Union

2011 — 60 pp. — 21 x 29.7 cm

ISBN 978-92-79-18514-4

doi:10.2759/52084

Search Computing

Business Areas, Research and Socio-Economic Challenges



Publications Office

doi:10.2759/52084

ISBN 978-92-79-18514-4



9 789279 185144



European Commission
Information Society and Media

