miKrow: Semantic Intra-enterprise Micro-Knowledge Management System

Víctor Penela, Guillermo Álvaro, Carlos Ruiz, Carmen Córdoba,
Francesco Carbone, Michelangelo Castagnone,
José Manuel Gómez-Pérez, and Jesús Contreras

iSOCO
Avda. Partenón. 16-18, 28042, Madrid, Spain
{vpenela,galvaro,cruiz,ccordoba,fcarbone,mcastagnone,
jmgomez,jcontreras}@isoco.com
http://lab.isoco.net/

Abstract. Knowledge Management systems are one of the key strategies that allow companies to fully tap into their collective knowledge. However, two main entry barriers currently limit the potential of this approach: i) the hurdles employees encounter discouraging them from a strong and active participation (knowledge providing) and ii) the lack of truly evolved intelligent technologies that allow those employees to easily benefiting from the global knowledge provided by them and other users (knowledge consuming). Both needs can sometimes require opposite approaches, tending the current solutions to be not user friendly enough for user participation to be strong or not intelligent enough for them to be useful. In this paper, a lightweight framework for Knowledge Management is proposed based on the combination of two layers that cater to each need: a microblogging layer that simplifies how users interact with the whole system and a semantic powered engine that performs all the intelligent heavy lifting by combining semantic indexing and search of messages and users. Different mechanisms are also presented as extensions that can be plugged-in on demand and help expanding the capabilities of the whole system.

Keywords: enterprise 2.0, knowledge management, social software, web 2.0, microblogging.

1 Introduction

The increasing amount of information generated by enterprises during the last decade has lead to the introduction of the new Knowledge Management (KM) concept, that has grown from a mere accessory to a full discipline that allows companies to grow more efficient and competitive.

Best practices in KM strategies usually attack several key objectives: i) identify, gather and organize the existing knowledge within the enterprise, ii) facilitate the creation of new knowledge, and iii) foster innovation in the company through the reuse and support of workers' abilities. However, in most of the cases,

the potential of these approaches doesn't get properly fulfilled by a fundamental flow in their design: prioritizing backend technologies complexity instead of making them easy to use and attractive enough to really encourage final users. This tends to reduce users' participation leading eventually to a loss of the knowledge that these tools are supposed to capture. In order to improve and extend these solutions, the issues detected are approached from two points of view: the system needs to be made both more attractive, so more users get engaged and actively participate, and smarter, so user interaction is minimized as much as possible making the system more proactive.

For increasing the allure of the system, the Web 2.0 paradigm, and in particular the microblogging approach will be used, where end-user involvement is fostered through lightweight and easy-to-use services and applications. These techniques are increasingly penetrating into the context of enterprise solutions, in a paradigm usually referred to as Enterprise 2.0. In particular, the trend of microblogging (of which Twitter[1] is the most prominent example) based on short messages and the asymmetry of its social connections, has been embraced by a large number of companies as the perfect way of easily allowing its employees to communicate and actively participate in the community, as demonstrated by successful examples like Yammer[2], which has implemented its microblogging enterprise solution into more than 70.000 organizations.

Different strategies are used in order to make the system more intelligent. First and foremost a semantically enriched layer supports KM indexing and search processes on top of the atomic information elements, i.e. status updates. Internally the system uses a domain ontology and thesauri related to the particular enterprise in which it is deployed, which can capture the different concepts relating to the company knowledge, and secondly, by making use of other information sources both internal such as already available internal knowledge information and external such as Linked Data[4] resources. Other techniques are also used for expanding and increasing the system intelligence, such as taking into account user Knowledge Processes [18] that can define what is relevant for employees in a particular context as well as reducing the cold start that a system based on user collaboration usually has.

This paper is structured in three main sections: the State of the Art regarding KM and microblogging is described in 2, the proposed theoretical contribution is shown in 3, several extensions that add new value to the original solution are explained in 4 and finally the implementation details and evaluation results are covered in 5.

2   State of the Art

2.1   Knowledge Management

The value of KM relates directly to the effectiveness[3] with which the managed knowledge enables the members of the organization to deal with today's

---

[1] Twitter: `http://www.twitter.com/`
[2] Yammer: `http://www.yammer.com/`

situations and effectively envision and create their future. Due to the new features of the market like the increasing availability and mobility of skilled workers, ideas sitting on the shelf,..., knowledge is not anymore a static resource of the company. It resides in its employees, suppliers, customers,.... If companies do not use the knowledge they have inside, one of their main resources stale.

In recent years computer science has faced more and more complex problems related to information creation and fruition. Applications in which small groups of users publish static information or perform complex tasks in a closed system are not scalable. In 2004, James Surowiecki introduced the concept of "The Wisdom of Crowds"[17] demonstrating how complex problems can be solved more effectively by groups operating according to specific conditions, than by any individual of the group. The collaborative paradigm leads to the generation of large amounts of content and when a critical mass of documents is reached, information becomes unavailable. Knowledge and information management are not scalable unless formalisms are adopted. Semantic Webs aim is to transform human readable content into machine readable. With this goal languages such as RDF Schema and OWL have been defined.

Computer supported collaborative work[10] research analyzed the introduction of Web 2.0 in corporations: McAfee[11] called "Enterprise 2.0", a paradigm shift in corporations towards the 2.0 philosophy: collaborative work should not be based in the hierarchical structure of the organization but should follow the Web 2.0 principles of open collaboration. This is especially true for innovation processes which can be particularly benefited by the new open innovation paradigm[7]. In a world of widely distributed knowledge, companies do not have to rely entirely on their own research, but should open the innovation to all the employees of the organization, to providers and customers.

Web 2.0 tools do not have formal models that allow the creation of complex systems managing large amounts of data. Nowadays solutions like folksonomies, collaborative tagging and social tagging are adopted for collaborative categorization of contents. In this scenario we have to face the problem of scalability and interoperability[9]: making users free to use any keyword is very powerful but this approach does not consider the natural semantic relations between the tags. Semantic Web can contribute introducing computer-readable representations for simple fragments of meaning. As will be seen, an ontology-based analysis of a plain text provides a semantic contextualization of the content, supports tasks such as finding semantic distance between contents and helps in creating relations between people with shared knowledge and interests.

Different mechanisms for leveraging all this scattered enterprise knowledge have been studied during the last decade, particularly trying to ease the pain of introducing new tools in the already overcrowded worker's desktop by adding a semantic layer on top of current applications. CALO[3] based on the use of cognitive systems and NEPOMUK[4] trying to add the social and semantic aspects

---

[3] CALO is part of the PAL Program: `https://pal.sri.com/`
[4] NEPOMUK Project: `http://nepomuk.semanticdesktop.org/`

to the user's personal desktop are two of the main references of ACTIVE[5], a project that aims to increase productivity of knowledge workers with pro-active and contextualized mechanisms and which technology has been used to improve the proposed solution.

2.2   Semantics in Social Networks

Microblogging is one of the recent social phenomena of Web 2.0, being one of the key concepts that has brought Social Web to more than merely early adopters and tech savvy users. The simplest definition of microblogging, a light version of blogging where messages are restricted to less than a small number of characters, does not make true judgment of the real implications of this apparent constraint. Its simplicity and ubiquitous usage possibilities have made microblogging one of the new standards in social communication. There is a large number of social networks and sites, with more blooming every day, that have some microblogging funcionalities, although currently there are two big players in the field: Twitter and Facebook, with 175 and 600 million users respectively.

One of the main issues microblogging has today is the lack of proper semantics, making building any kind of intelligent system on top of them quite hard. Even though different user initiatives have emerged, such as the use of hashtags to define channels of communication and provide a context for the conversation, its use is mostly related to user consumption of the information, not allowing for any real analysis of the meaning of the related data.

Twitter introduced Annotations[6], as a mechanism to add structured metadata about a tweet. It proposes an open key/value structure as properties of a type entity with recommended types such as "place", "movie" or "review". This low level approach is simplistic in the way that it does not define a formal model, but only a mechanism to add metadata to messages.

Facebook has proposed the Open Graph protocol as a mechanism to add metadata to its network, however the target has been quite the opposite, instead of adding metadata to messages as with Twitter Annotations, the main goal is to improve information linking with external resources by proposing a modified RDFa structure for webpages.

SMOB[12] tries to solve this by proposing the use of semantically-enabled hashtags such as #dbp:Eiffel_Tower in #geo:Paris_France. However this approach puts all the burden of explicitly giving meaning to different elements on the user, which is counterproductive with the idea of microblogging as lightweight communication tools.

This lack of semantics is a stronger constraint in a work environment, where employees need to have both faster and more reliable tools for KM while expecting new tools not to disturb their usual work experience and thus not forcing them into having to perform new tasks. Passant et al.[13] extended their previous approach by trying to solve these issues with a mixture of different user-friendly

---

[5] ACTIVE Project: `http://www.active-project.eu/`
[6] Twitter Annotations: `http://dev.twitter.com/pages/annotations_overview`

Web 2.0 interfaces for users to both provide and consume RDF/OWL annotations. This approach still seems quite hard on common employees, experts in their domain but with no basic knowledge on semantic technologies.

3   Semantic Processing in Knowledge Management

In this section, the theoretical contribution of this paper towards KM is described. We address the benefits of applying the the microblogging approach in 3.1, how the processes involved are enriched by the use of semantic indexing and search in 3.2, and the characteristics of the necessary underlying model in 3.3.

3.1   A Lightweight Approach towards Knowledge Management

The proposed interaction platform is a web application designed following the Web 2.0 principles of participation and usability. Our proposal centers interaction around a simple user interface with a single input option for end-users, where they are able to express what are they doing, or more typically in a work environment, what are they working at. This approach diverges from classical KM solutions which are powerful yet complex, following the simplicity idea behind the microblogging paradigm in order to reduce the general entry barriers for end users.

The purpose of the single input parameter where end-users can write a message is twofold: Firstly, the message is semantically indexed so it can be retrieved later on, as well as the particular user associated to it; secondly, because the content of the message itself is used to query the same index for relevant messages semantically related to it, as well as *experts* associated to those messages.

The semantic functionalities are possible thanks to underlying ontologies able to capture the knowledge of the company[6]. Even though there are already one too many ontologies that try to define the global and generic domain of enterprise relationships and knowledge bases, in terms of final performance, the final model must be as coupled as possible with the particular knowledge of each company. Our solution is to be deployed as a standalone service, with no ties with other deployments in other environments (e.g., in other companies) which further emphasizes the need for domain ontologies to be adapted to the particular needs of each company in order to fully tap into its knowledge needs and sources.

These on-demand domain ontologies will be extended with a set of thesauri in order to cover probable variations such as writing mistakes and commonly accepted alterations, making the whole ontology mapping process suitable for a microblogging environment where users feel less inclined to pursue utter correctness.

3.2   Semantic Indexing and Semantic Search

In the proposed approach, status updates (microposts) are stored in the platform knowledge base along with relevant metadata. The text of such messages is

analyzed and stored in the message index. The set of terms present in users' statuses compose their entries in the experts index. The text of the messages is used to perform a semantic search against the same index as well.

Semantic Indexing. When a user posts a new status message into the system, its content is analyzed and included into a message index (status repository), allowing future retrieval. Similarly, a repository of expert users (experts repository) is populated by relating the relevant terms of the message with the particular author.
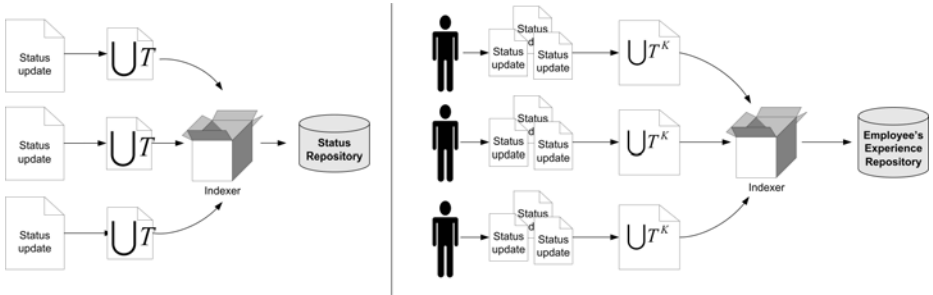
Fig. 1. (a) Message repository creation. (b) Experts repository creation.

Technically, messages that users post to the system are groups of terms T (both key-terms T^K, relevant terms from the ontology domain, and normal terms) $\bigcup T$. The process of indexing each message results in a message repository that contains each document indexed by the different terms it contains, as shown in figure 1(a).

In the case of the update of the semantic repository of experts, which follows the message indexing, each user can be represented by a group of key-terms (only those present in the domain ontology) $\bigcup T^K$. This way, the repository of experts will contain the different users of the systems, that can be retrieved by the key-terms. Figure 1(b) illustrates this experts repository.

Semantic Search. The search process is launched by a user posting a new status update. The new update message is processed by the semantic engine, extracting related concepts from the company knowledge base, modelled as both an ontology and a set of thesauri, and matching them with previously indexed status updates and employees. This is performed seamlessly behind the scenes, i.e., the user is not actively performing a search, but the current status message is used as the search parameter directly.

Two main search approaches are provided by the semantic engine:

- Given the text of a status update, the search on the status index returns semantically related status.

  – Given the text of a status update, the search on the experts index returns
    semantically related people, such as other co-workers with experience on
    related areas.

From a technical point of view, the semantic repository is queried by using
the group of terms $\bigcup T$ of the posted message, as depicted in figure 2(a). This
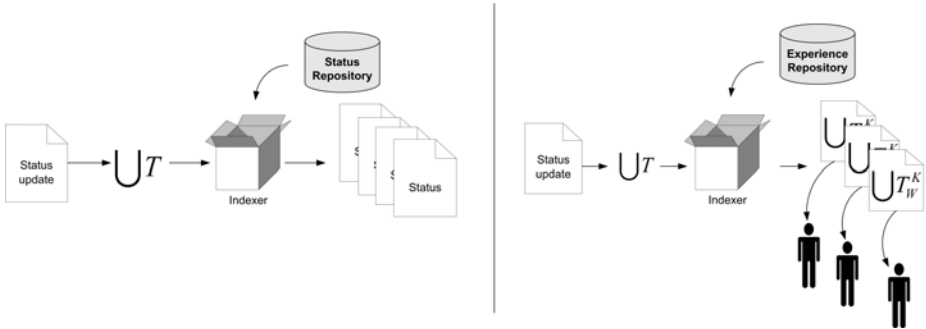search returns messages semantically relevant to the one that the user has just
posted.

Fig. 2. (a) Detection of related statuses. (b) Expert identification.

It is worth noting that, as it will be covered in 3.3, the search process in the
repository is semantic, therefore the relevant messages might contain some of the
exact terms present in the current status message, but also terms semantically
related through the domain ontology.

As it has been stated above, along with the search for relevant messages, the
system is also able to extract experts (identified by the terms present in the
messages they have been writing previously) associated with the current status
message being posted. In this case, the search over the semantic repository of
experts is performed by using the key-terms contained in the posted message
$\bigcup T^K$, as depicted in figure 2(b).

3.3   Knowledge Base Modelling

The precision of the modelled knowledge base, which will be built with collabo-
ration from field experts, is a key performance constrain as the semantic engine
query process is built upon the defined concepts and relationships. Particularly,
the relationships between different elements in both the ontology and the the-
sauri are exploited through techniques based mainly on morphological variations,
orthographic errors and synonyms for the terms defined in the the ontology, in
order to expand the initial queries with different approaches that can extend the
query recall, without threatening the global query precision.

The analysis of the text is not performed on single words: text fragments and
n-grams are considered for ontology matching. A term, composed by one or more
words, in a text can match i) general concepts (e.g. the keyword "product" which

matches the concept "product" in the ontology), ii) semantic relations between concepts, (e.g. the keyword "target" matches the relation *product_has_target*), or iii) instance entities (e.g., the keyword "Sem10 Engine" which matches the instance *Sem10 Engine*, a particular product of the company). This process can produce any number of matches for each term, strongly depending on the size and number of elements of the ontology, how well it covers the business knowledge base of the company and how the message is related to the core elements in it.

Once associations between terms and ontology entities (concepts, attributes and instances) are identified, the semantic search engine builds a query exploiting these relations defined by the ontology and hence in the company knowledge base. Different expansions can be performed depending on the initial input, with each one being weighed accordingly to the relevance of its relation:

– If a synonym of an ontology term is detected, the ontology term is added to the query.
– If a term corresponding to an ontology class is found, subclasses and instances labels are used to expand the query.
– If an instance label is identified, the corresponding class name and sibling instance labels are added to the query.

4   Knowledge Boosting Techniques

As extensions of the mentioned approach, different mechanisms are proposed in this section in order to extend the original solution with value added information that can improve the final user experience as well as to some extent some of the know issues such as the initial cold start and the limitations and lack of proper up-to-date maintenance of the domain ontology.

These proposed "boosting" features are in no way meant to suppress the original index and search engines, but to expand and upgrade the results provided.

4.1   Tackling the Cold Start Problem by Leveraging Existing Knowledge

One of the key issues the usage of a system like this presents is the delay from its formal deployment and the moment the knowledge base is large enough for its results to be really meaningful. This slow starting path could in many cases be long enough for many companies to desist in their investment in this kind of technologies.

Cold start happens when recommendations from a new item that has not been previously rated or classified are required. Since no user information on the item is available, a classical collaborative filtering approach is useless at that point.

This common issue on recommendation system is usually tackled by a large array of techniques ranging from hard and soft clustering of both users and items to other methodologies based on machine learning and probabilistic methods[15].

In order to overcome this issue, the approach proposed is to leverage current resources available in the prepopulated knowledge base to provide simpler

recommendations. That way, even though neither experts or messages will be recommended in the beginning, other secondary elements such as resources, contexts and processes will provide with a first glimpse of the real value of the whole system.

A complementary approach is to query external datasets to perform Named Entity Recognition on the message text, as we cover in the next subsection. This process leverages available datasets in order to provide users with related terms, that even though are not part of the company's knowledge base, could be of some interest. Additionally these relevant terms could also be used as a starting point for the ontology engineering process if seen as having an implicit relevancy as users tend to use them in their conversations.

4.2   Linked Data Consumption

One of the issues of the previous approach is the need of a large ontology that models as close as possible the whole knowledge base of an enterprise, which, depending on the size and the diversity of the company, may differ from difficult to almost impossible (new knowledge concepts being generated almost as fast as they can be modeled). This knowledge curation for the ontology engineering is not only a highly resource consuming process, but also the resulting ontology needs to be maintained and kept up-to-date with new knowledge from the company such as new employees (people), new partners and customers (companies), new business areas (technologies).

As an open approach to tackle this issue the proposed system tries to take advantage of information already available in a structured way via the Linked Data paradigm, providing with an easy and mostly effortless mechanism for adding new knowledge to the system knowledge base. Each new message posted will be processed with NLP methods against the distributed knowledge base that the Linked Data Cloud could be seen as. New concepts or instances extracted from that processing will be added to a temporary knowledge base of terms that could be used to add new information to the system's ontology. These terms would be semiautomatically added to the knowledge via algorithms that weighs the instance usage and the final input of a ontology engineer that decides whether the proposed terms are really valid or is a residue from common used terms with no further meaning to the company.

The main advantage of this approach is that it allows the whole system to adapt and evolve with an organic growth alongside the evolution of the company knowhow. That way, when a new client starts to make business with the company (or even before, when the first contacts are made) some employees will probably start to post messages about it ("Showing our new product to company ACME", "Calling company ACME to arrange a new meeting",...). Querying the Linked Open Data Cloud will automatically detect that this term *ACME* is indeed a company, with a series of properties associated to it (headquarters location, general director and management team, main areas of expertise,...), and would allow for this new knowledge to be easily added to the local knowledge dataset.

4.3   Context-Aware Knowledge Management

In order to extend the relevancy mechanisms proposed, a context-aware[8] approach will extend the current view of messages as the only information element, adding a new layer of external information that could somehow improve the final user experience.

Simple rules will be used for adding a new perspective on top of the previous approach. That way, two employees detected as having a similar level of expertise on a particular topic will be weighed in terms of external data sources such as who is geographically closer (e.g. same office), hierarchically closer (e.g. same department) or available at that particular moment.

For this purpose the original enterprise ontology will be extended by means of an already available context model[5] and the consumption of different services provided by a context-aware infrastructure[14].

4.4   Connecting to Enterprise Information Systems

Even though the global solution is built upon a microblogging environment and obviously focused on lightweight KM, interaction with currently deployed systems in an enterprise environment is a key element in order to ease possible entry barriers as well as leverage already available knowledge information in the company.

As a test use case different levels of information will be extracted from services provided by ACTIVE project[7][16]. ACTIVE aims to increase the productivity of knowledge workers in a pro-active, contextualized, yet easy and unobtrusive way through an integrated knowledge management workspace that reduces information overload by significantly improving the mechanisms through which enterprise information is created, managed, and used. Combining this approach with our microblogging solution will thrive the benefits for workers.

ACTIVE tries to extract information from the whole employee environment, dividing the provided data in three main types of concept:

- Working Context, constructed from a particular set of items (activities, information resources, and people) used to achieve a particular undertaking.
- Resource, seen as placeholder of something that can be used, such as a document or URL.
- Knowledge Process, defined as a loosely defined and structural ramified collection of tasks carried out ky workers as part of their daily activities.

The microblogging tool will extend its classical interface by including links to different instances of each class. These instances will be obtained by consuming ACTIVE services with the detected terms in a particular message as tags for the query and function as interaction channels between both systems, allowing the employee to gather further information and working as a bridge between lightweight KM tool and more resource-intensive platform.

---

[7] ACTIVE Project: `http://www.active-project.eu/`

5   Microblogging as a User Interaction Layer

The theoretical contribution covered in the previous section has been implemented as a prototype, codenamed miKrow, in order to be able to evaluate and validate our ideas. In the following subsections, we address the implementation details and the evaluation performed.

5.1   miKrow Implementation

Figure 3 depicts the Web page of the current implementation of miKrow. After a first version that only included the basic indexing and search mechanisms on top of the microblogging layer, as presented in [1], this new iteration has tried to evolve the initial approach, by adding a general improvement on both the backend and frontend, as well as adding new information boosting techniques that try to improve the final user experience, such as integrating functionalities developed

Fig. 3. miKrow implementation snapshot

inside ACTIVE project, as well as solve some of the issues raised from the first evaluation performed inside iSOCO[8].

miKrow is divided in two main components, a semantic engine that uses Lucene in order to offer search and indexing functionalities, and a microblogging engine, for which Google's Jaiku[9] has been forked and extended to properly include and show the new type of related information that miKrow offers to the final user.

Microblogging Engine. miKrow microblogging capabilities have been built on top of Jaiku, recently open sourced by Google, using basic microblogging functionalities and UI, relying on it for most of the heavy lifting related to low level transactions, persistence management and, in general, for providing with all the basic needs of a simple social network.

Using Jaiku gives the project a good head start, reducing the burden of middleware and infrastructure development, by reusing already production proved Jaiku's software, and thus allowing to extend that effort and focus on adding the semantically enabled layer.

The choice of Jaiku over other possibilities available is based essentially in its condition of having been extensively tested and the feasibility of being deployed in a Cloud Computing infrastructure[2] such as Google App Engine[10], thus reducing both the IT costs and the burden of managing a system that could have an exponential growth.

Semantic Engine. The semantic functionalities are implemented in a three layered architecture: i) ontology and ontology access, ii) keyword to ontology entity, and iii) the semantic indexing and search as the top layer.

The main functionality is the performance of Named Entity Recognition on each new status update, allowing the extraction of some of the real meaning of a message. This process is performed by parsing each message and analyzing different n-gramms with a base ontology that depicts the enterprise knowledge base and several supporting thesauri, that provides a more extended terms and synonym dataset. Each message is then tagged with the entities that have been extracted from that message.

Lucene[11] is used to create both messages and statuses indices. Each index contains terms tokenized using blank space for word delimitation and ontology terms as single tokens (e.g. if the text contains "credit card" and this is a term of the ontology, "credit", "card" and "credit card" are added as tokens to the index). Ontology terms are detected leveraging the keyword to ontology mapping engine, using the OpenRDF framework[12] as an ontology access mechanism to the ontology, and taking into account possible morphological variations, orthographic errors and synonyms.

---

[8] iSOCO: `http://lab.isoco.net/`
[9] Jaiku: `http://www.jaiku.com/`
[10] Google App Engine: `http://code.google.com/appengine/`
[11] Lucene: `http://lucene.apache.org/`
[12] OpenRDF: `http://openrdf.org/`

The original semantic engine is also extended by introducing two main additional functionalities, which main goal is to reduce the usual cold start of this type of services:

- Linked Data entities. External services such as OpenCalais[13] are used to connect the messages posted to external entities in the Linked Data paradigm, allowing the system to propose new entities not included in the enterprise ontology.
- Knowledge resources. ACTIVE technology is used to recommend knowledge resources related with the entities extracted from the user messages, lowering the gap between the lightweight tool and more intensive desktop platforms.

Communication between layers. The communication between both layers, the microblogging engine employed as user interface and the semantic engine that provides the business logic on message and experts recommendation as well as the indexing and search functionalities, is highly decoupled and based on Web Services. This approach provides with a more reliable system, since the microblogging engine will keep providing its basic service even if the semantic engine is down or malfunctioning. That way, even though the user experience will be reduced to a simple microblogging environment, lacking any kind of intelligent analysis and recommendation, users will still be able to check messages by themselves and to update their statuses.

5.2   miKrow Evaluation

A first evaluation of the initial and basic version of miKrow, which was not integrated with existing enterprise information systems, was carried in-house inside iSOCO[1], which has around 100 employees distributed in 4 different cities across Spain. A new evaluation has been made by enabling the new miKrow prototype linked with ACTIVE technologies, in order to assess the "knowledge boosting techniques" as well as the semantic benefits.

From a qualitatively point of view, we extracted the following conclusions from the evaluation process:

- The microblogging paradigm has its own rules and syntax, and therefore reusing a knowledge model without adapting it to the special characteristics of this approach implies a decrease in both precision and recall. On one hand, misleading suggestions are caused by stop-words that should not be considered in a microblogging tool, for instance some initial activity gerunds (e.g., *working*, *preparing*). On the other hand, the particular syntax of microblogging implies new ways of expressing the same things in a simpler form (e.g., *ref* instead of *reference*), and hence the thesauri should capture those.
- Temporal relevance of microposts is not to be disregarded. In some occasions, a message is useful only for a short time span, while in others its validity is much longer. User feedback on the suggestions comes in handy to tackle this issue, if they are able to tag messages as no longer valid, etc.

---

[13] OpenCalais: `http://www.opencalais.com/`

– Informing users about the reasons for the suggestions (both internal to the tool for messages and experts, and external, for documents found in the existing enterprise information systems) is important, as they perceive some sort of *intelligence* in the system, and are significantly more pleased. Also, if the suggestion is not good, they at least know why it has been produced. Again, letting them provide feedback in these occasions will generate a benefitious loop that will enrich the system.

6   Conclusions

This paper has presented the concept of a semantic microblogging tool to be used within an enterprise network as a lightweight KM service. Even though the Web 2.0 philosophy has been used for a while in work environments, in which is usually called the Enterprise 2.0 paradigm, most of the solutions simply apply a new social layer that does not fulfill the particularities of this kind of environments many times becoming more a resource waste than a added-value tool.

The addition of a semantic layer as an indexing and search engine is the proposed solution in terms of extended intelligence and reliability. This semantic engine is in charge of providing employees with related messages and experts on the topics they are talking about. In order to improve the overall performance a set of ontologies and thesauri will be built to fully model each company knowledge base.

Different extensions have been built in order to improve and extend the current solution by adding new sources of information, while providing the user with a single entry point to the application.

References

1. Álvaro, G., Córdoba, C., Penela, V., Castagnone, M., Carbone, F., Gómez-Pérez, J.M., Contreras, J.: mikrow: An intra-enterprise semantic microblogging tool as a micro-knowledge management solution. In: International Conference on Knowledge Management and Information Sharing 2010, KMIS 2010 (2010)
2. Armbrust, M., Fox, A., Griffith, R., Joseph, A., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., et al.: Above the clouds: A berkeley view of cloud computing. EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28 (2009)
3. Bellinger, G.: Systems thinking-an operational perspective of the universe. Systems University on the Net 25 (1996)
4. Berners-Lee, T.: Linked data. International Journal on Semantic Web and Information Systems 4(2) (2006)

5. Cadenas, A., Ruiz, C., Larizgoitia, I., García-Castro, R., Lamsfus, C., Vázquez, I., González, M., Martín, D., Poveda, M.: Context management in mobile environments: a semantic approach. In: 1st Workshop on Context, Information and Ontologies (CIAO 2009), pp. 1–8 (2009)
6. Carbone, F., Contreras, J., Hernández, J.: Enterprise 2.0 and semantic technologies: A technological framework for open innovation support. In: 11th European Conference on Knowledge Management, ECKM 2010 (2010)
7. Chesbrough, H., Vanhaverbeke, W., West, J.: Open Innovation: Researching a new paradigm. Oxford University Press, USA (2006)
8. Dey, A., Abowd, G.: Towards a better understanding of context and context-awareness. In: CHI 2000 Workshop on the What, Who, Where, When, and How of Context-Awareness, pp. 304–307 (2000)
9. Graves, M.: The relationship between web 2.0 and the semantic web. In: European Semantic Technology Conference, ESTC 2007 (2007)
10. Grudin, J.: Computer-supported cooperative work: History and focus. Computer 27(5), 19–26 (1994)
11. McAfee, A.: Enterprise 2.0: The dawn of emergent collaboration. MIT Sloan Management Review 47(3), 21 (2006)
12. Passant, A., Hastrup, T., Bojars, U., Breslin, J.: Microblogging: A semantic and distributed approach. In: Proceedings of the 4th Workshop on Scripting for the Semantic Web (2008)
13. Passant, A., Laublet, P., Breslin, J., Decker, S.: Semslates: Improving enterprise 2.0 information systems thanks to semantic web technologies. In: Proceedings of the 5th International Conference on Collaborative Computing: Networking, Applications and Worksharing (2009)
14. Penela, V., Ruiz, C., Gómez-Pérez, J.M.: What context matters? Towards multidimensional context awareness. In: Augusto, J.C., Corchado, J.M., Novais, P., Analide, C. (eds.) ISAmI 2010. AISC, vol. 72, pp. 113–120. Springer, Heidelberg (2010)
15. Schein, A., Popescul, A., Ungar, L., Pennock, D.: Methods and metrics for cold-start recommendations. In: Proceedings of the 25th ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 253–260 ACM,New York (2002)
16. Simperl, E., Thurlow, I., Warren, P., Dengler, F., Davies, J., Grobelnik, M., Mladenic, D., Gomez-Perez, J.M., Ruiz, C.: Overcoming information overload in the enterprise: the active approach. IEEE Internet Computing 14(6), 39–46 (2010)
17. Surowiecki, J., Silverman, M., et al.: The wisdom of crowds. American Journal of Physics 75, 190 (2007)
18. Warren, P., Kings, N., Thurlow, I., Davies, J., Brger, T., Simperl, E., Ruiz, C., Gómez-Pérez, J., Ermolayev, V., Ghani, R., Tilly, M., Bsser, T., Imtiaz, A.: Improving knowledge worker productivity – the active approach. BT Technology Journal 26, 165–176 (2009)