IST FP7 231507

D3.4 Demonstrator of Presentation Functionality v1.0

File name	PuppyIR_D3.4_PresentationFunctionality_v1.0
Author(s)	Marie-Francine Moens (KUL), Aparna Nurani Venkitasubramanian (KUL), Karl Gyllstrom (KUL)
Work package/task	Task 3.1
Document status	Final
Contractual delivery date	31/12/2011 (adjusted agreement) DoW date was 30/09/2011
Confidentiality	This version is confidential.
Keywords	Aggregated search, presentation of search results
Abstract	This deliverable is a first report on the results for Task 3.1 Multimodal presentation technology: on technologies for aggregated search and presentation of search results. A novel technology for aggregated search is presented. In addition an overview is provided of other presentation approaches developed in the PuppyIR framework.

Table of Contents

Executive Summary

Task 3.1 in WP3 is is concerned with the aggregation of the results of a search and the presentation of these results in a most convenient way. This deliverable reports on a demonstrator for aggregated search. The main idea is that depending on the properties of the search results the broadness or narrowness of a query is computed and depending on this score appropriate search results are presented to the user, which the user can further explore in a navigating and interactive mode. The underlying assumption is that the aggregation and presentation of search results can be largely improved by knowing the distances and closeness in content of the Web pages retrieved. These distances could be computed in many different ways, but in this PuppyIR deliverable we make use of knowledge about the subject categories to which the Web pages belong and their hierarchical taxonomy. More specifically we have made use of the DMOZ kids and teens directory. Based on the location of the search results in such the hierarchical taxonomy of subject categories, a selection of categories or Web pages is made to present to the user. As a side effect the aggregation method contributes to the presentation of a document collection that is hierarchically classified, where we know that large taxonomies cannot be shown to a user and users often get lost when searching in documents classified by their categories. This preliminary deliverable reports on this novel methodology for aggregated search.

In addition, this deliverable also reports on other presentation approaches developed in the framework of the PuppyIR project. In this preliminary deliverable only a short summary of their functionality is given.

In a final version of this deliverable, we will report on the evaluation of the main presentation demonstrator, describe the experimental setup, discuss the results, and update the methodology section accordingly. In addition, additional information on the aggregation aspects of other presentation functionalities of the PuppyIR framework will be given. In this final version the links with the interface development of WP2 will be clarified.

1 Introduction

This document reports on Task 3.1 of the PuppyIR project and describes the multimodal presentation technology for dealing with topical alignment, aggregation and summarization of content. The focus is on retrieval models that aggregate several sources of content. Task 3.1 started from the assumption that search results could be better organized so that children more easily find the information in which they are interested. There is there problem that their search terms might be broad yielding many diverse search results, or their search terms might be too narrow as they have difficulties in expressing a query that is on target on the content in which they are interested. It has been found that facetted search, i.e., a search guided by a taxonomy of subject categories is helpful and helps in navigating to and exploring the relevant information, but hierarchical taxonomies of facets are often very large and it is difficult for users – especially for children - to find their way.

The main goal of the demonstrator described in this deliverable is to develop a novel method for aggregated search that combines the benefits of a regular search based on query terms and a navigational facetted search. More specifically, depending on the properties of the search results the broadness or narrowness of a query is computed and depending on this score appropriate search results are presented to the user, which the user can further explore in a navigating and interactive mode. The underlying assumption is that the aggregation and presentation of search results can be largely improved by knowing the distances and closeness in content of the Web pages retrieved. These distances could be computed in many different ways, but in this PuppyIR deliverable we make use of knowledge about the subject categories to which the Web pages belong and their hierarchical taxonomy. More specifically we have made use of the DMOZ Kids and Teens directory. Based on the location of the search results in such the hierarchical taxonomy of subject categories, a selection of categories or Web pages is made to present to the user. As such the model contributes to the aggregation and summarization of search results. As a side effect our search aggregation method contributes to the presentation of hierarchically classified documents, where we know that large taxonomies cannot be shown to a user and users often get lost when searching in a collection of categorized documents.

This first version of the deliverable is a stepping stone towards its full version that will be delivered at the end of the PuppyIR project (month 36). This preliminary deliverable (version 1.0 of D3.4) reports on this novel methodology for aggregated search, which is described in section 2. In the final version of the deliverable additional subsections on the evaluation setup and discussion of the results will be filled in.

In addition, this deliverable also reports on other presentation approaches developed in the framework of the PuppyIR project. In this preliminary deliverable only a short summary of their functionality is given (Section 3). Also additional information on the aggregation aspects of other presentation functionalities of the PuppyIR framework will be given. In version 2.0 the links with the interface development of WP2 will be clarified.

The Table of Contents in this report already shows the content structure of version 2.0, some (sub)sections are only elaborated in the final version while this versio contains placeholders.
For the full details of the technologies reported in this deliverable we refer to the scientific publications.

2 Aggregated Search Model

2.1 Overview

As discussed in the introduction the goal of the main demonstrator of D3.4 is to select or aggregate results (Web pages or their conceptual summary in the form of a DMOZ category) that will be shown in a navigational interface to the user. We first discuss the rationale of our model in subsecton 2.2. Although not the essential part of the technology subsection 2.3 discusses the interactive interface. This is useful, because it gives the reader a good picture of the interface setting in which the aggregated search technology can operate. Then follows subsection 2.4 on the core of the technology, i.e., the computation of the broadness or narrowness of the query and the consequent selection of Web pages and categories that will be shown, both after firing the query to the system, and during the interaction.

2.2 Rationale

Combining the benefits of a keyword and a facetted search By exploiting the result list of the keyword search and the knowledge of the subject categories to which the Web pages belong and their relationships, the idea is to combine the benefits of both worlds.

Promoting diversity Certain queries are too broad or ambiguous. Examples would be queries like 'Arts' which can have different forms of art or 'Music' which can have various genres. In certain other cases, the query can be about a problem, which can have different dimensions. Since, children's queries can be too broad, it is useful to ensure that different aspects of the query are covered.

Handling queries that are too specific For queries that are too specific, it is useful to recommend other categories that could be of interest. A child who queries 'Common noun' would most likely be interested in 'Proper nouns' and 'Collective nouns' as well. Presenting these categories along with the query's main result would not only help put things in perspective, but also suggest other related topics.

Interactive exploration Search interfaces usually can only show a limited number of search results, even if they are presented in a summarized form, such as a conceptual category. The interactive exploration allows taking into account preferences of the user and the system to reselect search results and categories for further exploration.

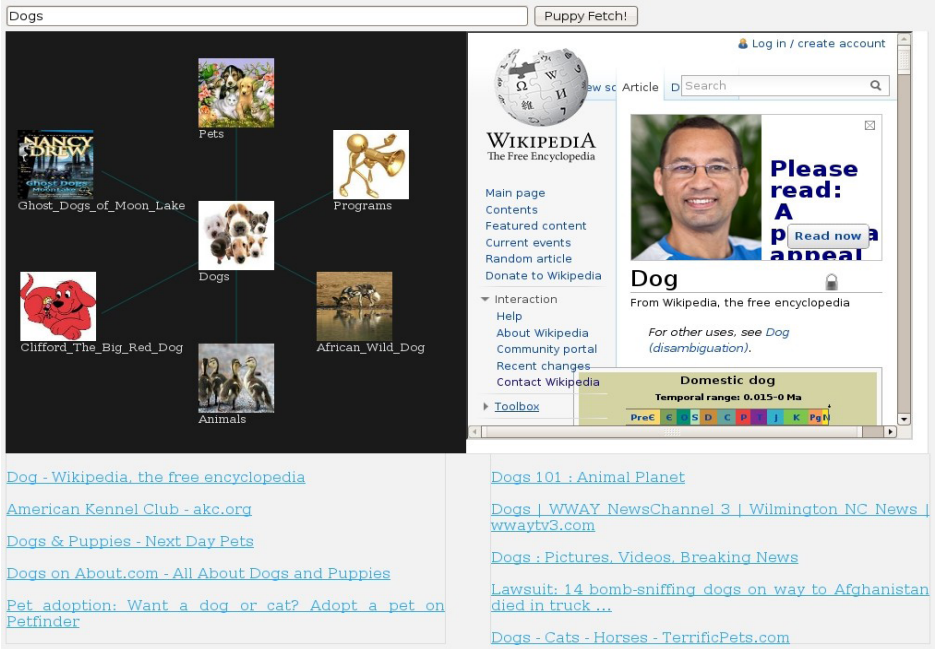

Figure 1: Results for the query 'Dogs'.

2.3 The Interface for Exploratory Search

The interface (see an example in figure 1) that is integrated with the aggregated search displays a network of objects which could be Web pages or categories associated with the query results, apart from the text results. This semantic network is a cluster of nodes, with the query at the centre, surrounded by a selection of nodes that according to the query results are ‘related’. Further, clicking on nodes of the network helps refine queries that are broad or ambiguous, or even too narrow, eventually helping children find the exact information they are looking for.

Any node in the network can be clicked. Clicking a node 'centres' it, and shows results corresponding to the clicked node, and the initial search results. For example, figure 1 shows results for the query 'Dogs', The left pane shows the cluster of nodes with 'Dogs' at the centre, surrounded by the related categories- pets, animal, programs, African wild dog etc. When the node for 'Animals' is clicked, the network shows different objects like mammals, dinosaurs etc., with 'Animals' at the centre, shown in figure 2. Thus, from an initial query, it is possible to refine one's search based on the related nodes displayed.

The selection of the nodes that represent the search results to be shown in the interface and the updating of this selection in an interactive session are discussed in the next section.

Lin et al. [1] have indicated a correlation between the interactivity of a website and the enjoyment it brings about. Likewise, they have also shown a relation between multimedia and enjoyment. In [2], Dania has found that graphics and animation are appealing to children. Additionally, with

children hesitating to read beyond the first page of search results, it would be useful to present a visual summary. So, once the nodes of the network have been identified, an image can be chosen that represents the category or Web page. [3] describes an approach for auto-illustration using statistical associations between images and text which are trained on a large database of images and texts, which can be used to associate the content of a Web with an image. Other approaches are possible such as extracting an image from Wikipedia that corresponds to a certain category, using a category as search term and choosing a representative image, for instance, by choosing the most central image with closest distance in terms of image features to all the other retrieved images, or by choosing the highest ranked image of a search.

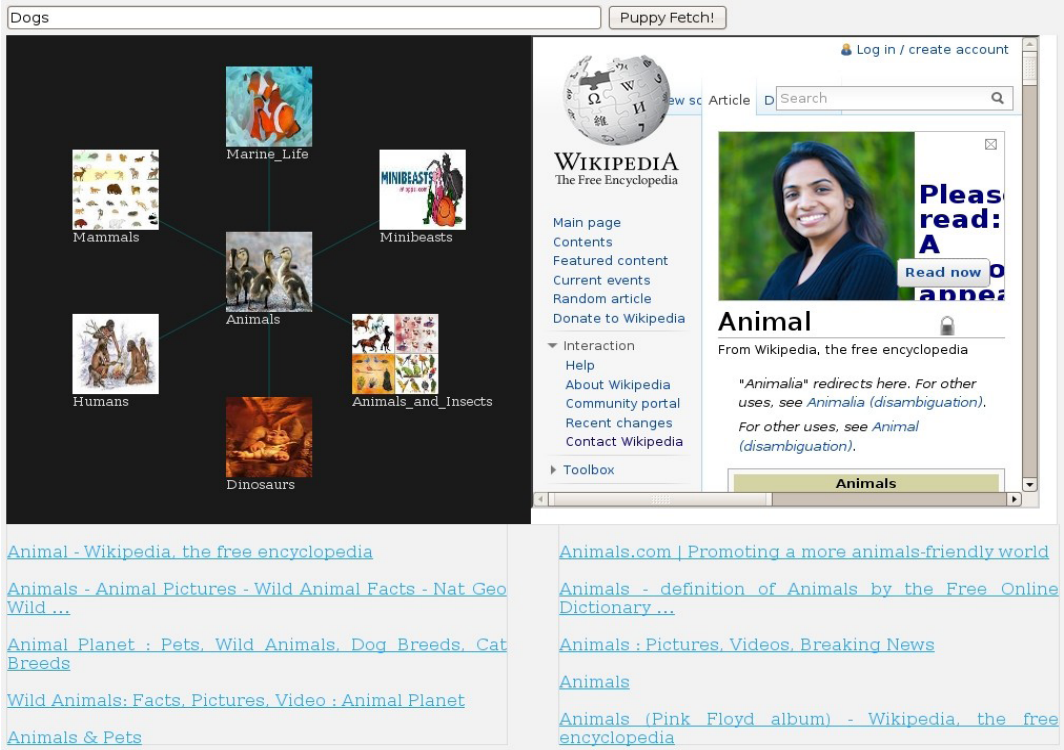

Figure 2: Navigation from 'Dogs' to 'Animals'.

2.4 Selection of Categories and Search Results

General overview

The crux of the technology is the selection of 'related' nodes that will be presented to the user. These can be categories (grouping several Web pages that are the results of a search forming a kind of summary) or individual Web pages. Different queries have to be handled differently. For broad queries, which have results from different diverse domains, we need to present these different facets, allowing the user to zoom-in or focus on one of these. On the other hand, for queries that are too narrow, it would be useful to recommend other items related to the query,

allowing the user to zoom out or widen his/her search and understand the context.

The manner in which this set is chosen largely depends on the nature of the query. For queries that are too broad, various diverse results need to be shown. For the more specific queries, it is useful to recommend other related interesting categories or Web pages. The first step is then to identify the nature of the query. To do this, the DMOZ categories associated with the search results are extracted. Depending on where the search results populate the DMOZ hierarchy, several heuristic techniques can be defined to compute the broadness of the query. Depending on this score several strategies for selecting which nodes (categories or Web pages) to show to the user can be evaluated.

In the interaction with the user, when he or she has selected a category node, the broadness of the query can be recomputed taking into account this additional evidence.

Identifying the categories associated with the query results and estimating their importance

First, the facets of the query's results are identified. In our case, this is done by extracting the DMOZ categories associated with the Web pages that are the result of a search. Then, the 'importance' of each category is estimated by the Discounted Cumulative Gain (DCG) computed over the retrieved documents under this category for the given query. This score gives a measure of the contribution of the particular category to the (high-ranked) results for a particular query.

Computing the query breadth

We assume that broad queries result in diverse search results. This diversity can be computed via the distribution of the category nodes that are activated in the DMOZ hierarchy. Several heuristics here are implemented such as the minimum path length, i.e. the path with minimum length that connects all the retrieved and relevant Web pages, where connection weights can be lower in the base of the hierarchy and become progressively larger as one moves up, and connection weights might be weighted by the DCG of the child category. A longer path indicates a broader query, while a shorter path indicates a narrow one.

Selection of the nodes

As the exploratory interface can only show a limited number of nodes, we need mechanisms to summarize the search results. If several Web pages belong to a subject category, the category will be shown instead of the individual Web pages. But even the number of relevant categories to show can be large, demanding for a kind of summarization over de category nodes. Dense clusters of search results and their representative category are detected while using the DMOZ hierarchy for computing the distances between the Web pages taking into account concepts of informativeness [4] [5]. Representative categories or Web pages possibly weighted by their DCG score are selected and shown in the interface.

Such a presentation of diverse results is useful when the query is broad. When the query is narrow we can use the same summarization of the search results, and when we obtain insufficient nodes to show in the interface, we choose neighboring categories or Web pages given the distances in the DMOZ hierarchy.

Recomputation of the broadness of a query based on the interaction of the user

The above mechanism already provides novel ideas for an aggregated presentation of the search results where by clicking on each category cluster the user can consult the retrieved Web documents belonging to that cluster, and in case of a narrow query can see related pages. If time permits we would like to further expand the model to a fully interactive search setting. When the user has chosen a category in the interactive interface, the documents belonging to that category could be reranked, and the whole selection procedure can be repeated. In such a setting the user should be able to go back to previous configurations of search results.

2.5 Current Implementation of the Demonstrator

We use the Bing search engine to obtain the ranking of relevant Web pages for a query. This set of pages is intersected with the set of Web pages belonging to the DMOZ subtree of Kids and Teens. The Web pages have the DMOZ categories assigned as metadata stored in a database. The category nodes and their relations are also maintained in this database, which is accessed through AJAX calls. The interface network is a 'Hyperbolic tree', built using the Javascript InfoVis toolkit, and is based on a dynamic JSON object. The network allows zooming in and out, enabling one to navigate.

2.6 Experimental Setup (*to be added*)

2.7 Results and Discussion (*to be added*)

3 Survey of Other PuppyIR Presentation Functionality

In this section we describe presentation approaches as supported by other demonstators of PuppyIR. We give a short summary of the demonstrator and in a separate section we discuss their contribution to the Task 3.1 (aggregation and presentation) and evaluation.

3.1 EmSe: Hospital Demonstrator

Summary

EmSe is built using the PuppyIR Framework, which provides a suite of components that can be combined to build child-specific search services. Component types range from interfaces to various search resources (e.g., Bing, YouTube, Twitter) to a collection of information processing components that filter and modify both queries and results to support the user and their search tasks. Fundamentally, EmSe enables searching trusted medial sites as well as the web, over which the following services are built:

1. The Body Browser enables exploration of the dataset via an interactive illustration of a body, where users can zoom to various levels of detail from the entire body to specific organs, which triggers medical Web searches related to the body parts and organs in focus (see Figure 3).
2. The Query Moderation identifies and enhances queries that are unlikely to be child-oriented by applying simple, real-time technology. The focus is not explicitly to remove mature content, but rather to help make results to general queries more child-friendly.
3. Query Suggestions help children explore and query the recommended and related medical sites by providing suggestions that reflect the specific content on these sites. They are generated by extracting meaningful phrases from the anchor text of the recommended resources.
4. Content Simplification helps children understand difficult medical jargon. Instead of a priori rejecting difficult content, as suggested by previous work, requested pages are checked for difficult terms, which the system augments with brief definitions from external sources such as Wiktionary (see Figure 4).

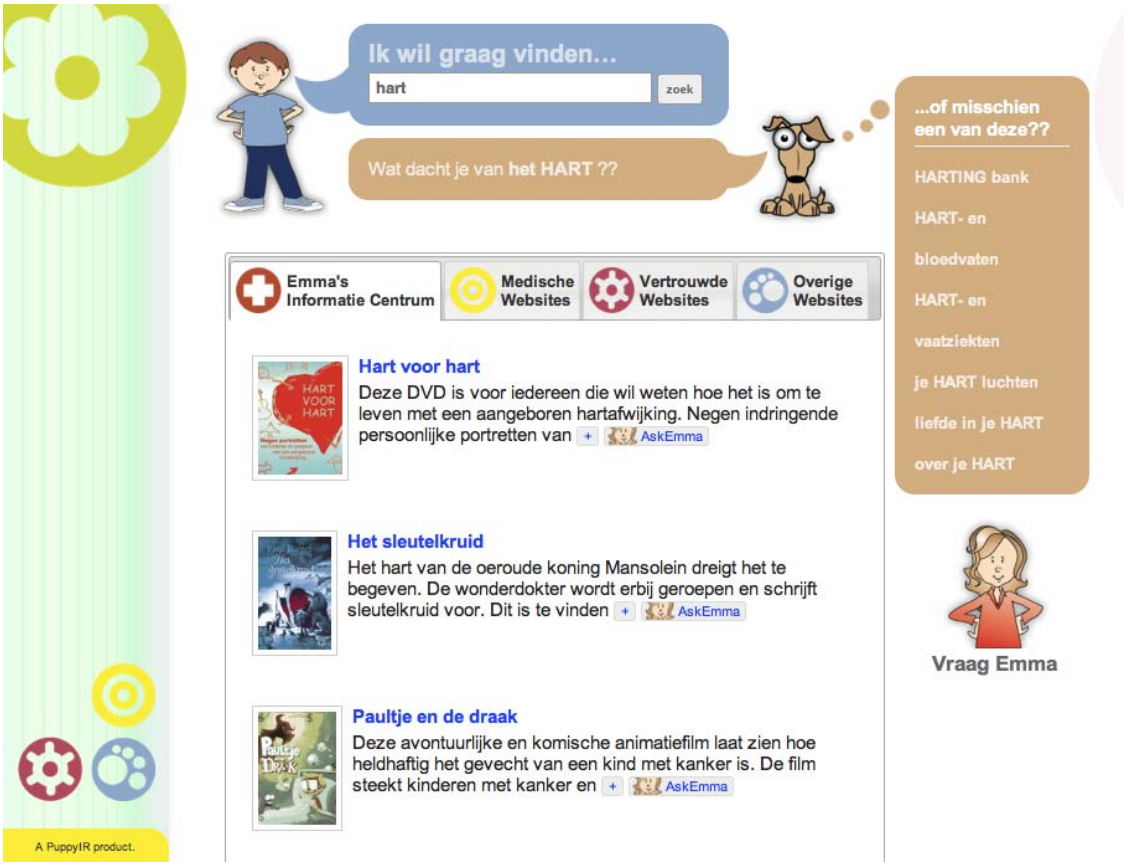

Figure 3: Result presentation interface of EmSe.

The first version is released at the beginning of 2012, where it will be accessible to staff and patients within the hospital (via bedside and other terminals in the hospital), and also out-patients via the web. The planned evaluation will consist of two main stages. The first stage will be to obtain feedback from the staff at the hospital. From this initial evaluation we will then refine the demonstrator to incorporate suggested changes. The next stage will be patient focused, where we will actively solicit feedback from patients, parents and visitors to the hospital. These two evaluations will use questionnaires to illicit feedback from users along with implicit logging to determine how, and how often, the system is being used.

Contribution to the search aggregation and presentation task (to be added in final version)

Evaluation (to be added in final version)

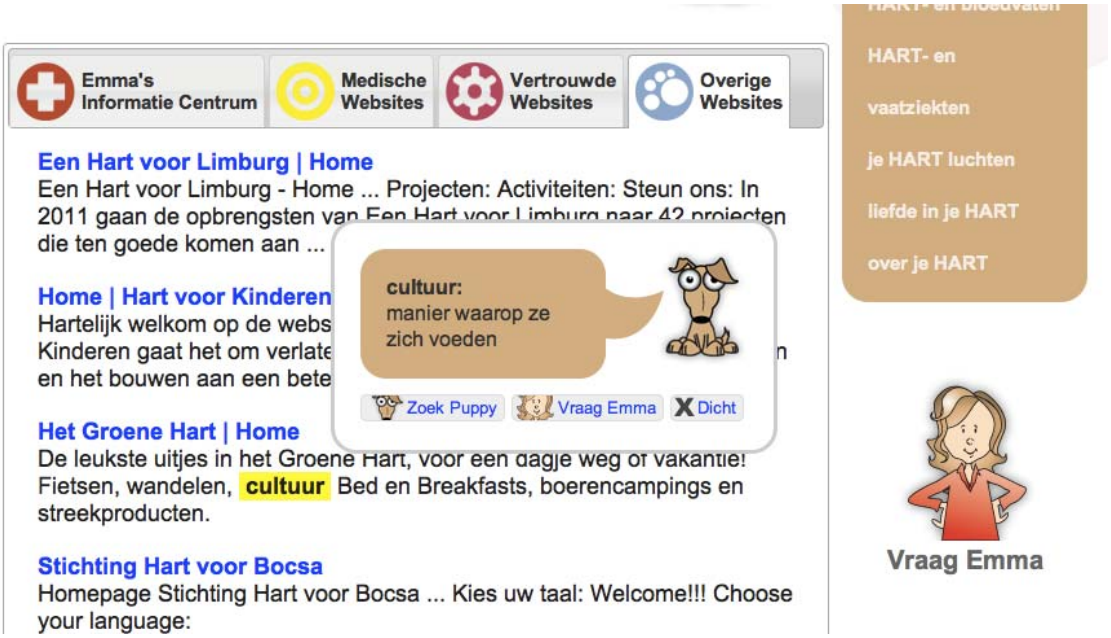

Figure 4: Result presentation interface of EmSe.

3.2 Museon Demonstrator

Summary

In the museum a multi-touch table is used at the beginning and at the end of the museum visit, by groups consisting of two to four children. Often one or two parents are also part of the group. The application at the multi-touch table is used to determine a route through this exhibition, based on the interests of the visitors. The goal is to give the children guidance and to optimize their museum-going experience. The main screen that people see when they arrive at the table has the solar system as a background. When people touch this screen, particles appear with colored backgrounds (see Figure 5).

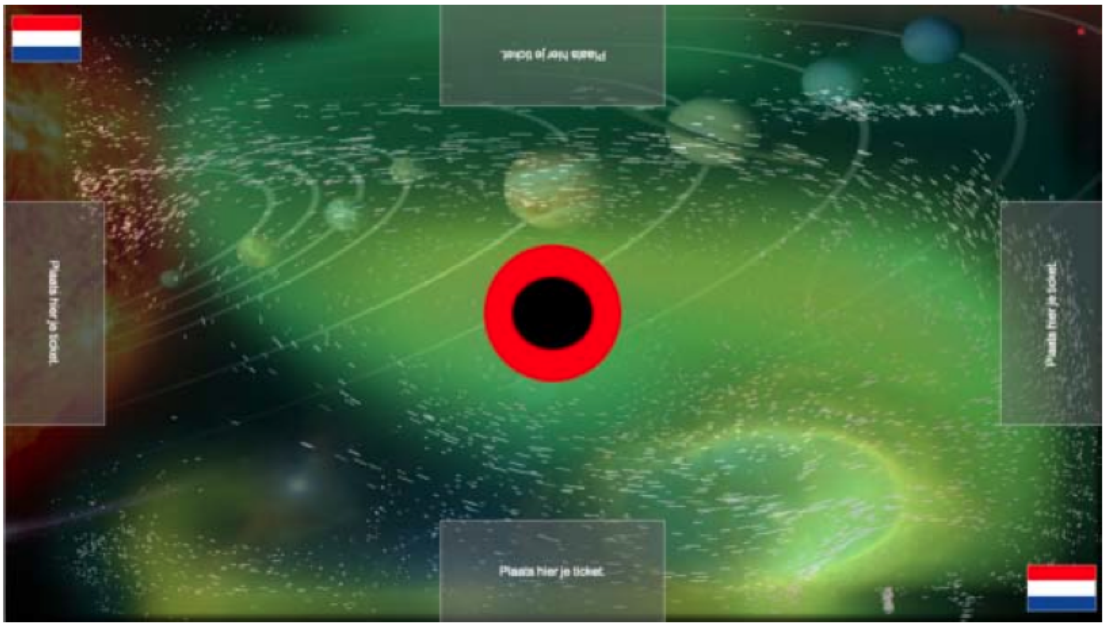

Figure 5: The main screen of the touch table before interaction started.

Interaction with this screen is not really part of the registration procedure but it is meant to be engaging and to encourage exploration by enabling trial and error activity and lightweight interaction. At the middle of each side of the table there are virtual boxes. These boxes can be used to register. Collaboration is facilitated by multiple access points, non-fragmented visibility and space for friends. Children who take part in the experiment get a ticket that fits in these boxes and that has a marker on one side that is recognized by the table and a barcode on the other side that will be used in the quest (see Figure 6). The initial game starts when people put the tickets in the boxes. In the initial game people choose categories of subjects (i.e., parts of the exhibition) they are interested in. There are twelve categories represented by round images. Everybody chooses six of these categories. Here the theme physical and digital representations of our model is relevant. By direct manipulation users choose images and move them to their own area. The relation between this manual action and its effect is easy to understand (isomorph effects) and the action-perception coupling is tight (unification of input and output).

The chosen categories are used to determine a route through the exhibition room of the museum. In the exhibition room many (around 120) touch screens with barcode reader are available, close to the exhibits. The registrations tickets are used here to identify the children and to transfer information from and to the table applications. Based on the results of the initial game the participants receive a personalized quest of twelve questions to be answered at twelve different exhibits. After each good answer, the children choose a virtual object they like. People can help each other whenever they want. They are near to each other and interacting with the other group members. After all members of the group have finished the quest, they go to the multi-touch table again to do the end game.

Figure 6: Putting in the name of the children in the game.

Coming back to the multi-touch table the children use their tickets to login again. From the virtual objects collected during the quest the group chooses twelve different objects. In the end game these objects are in the middle of the table. Twelve boxes with words are positioned at the edge of the table. Collaboration is facilitated by visibility of the words from two sides. The task is to connect the words to the matching virtual objects by drawing lines. The children have limited time for this. After two minutes the connections are checked showing an animation: one by one the virtual objects are highlighted and the connecting lines become green when a connection is correct and red when it is not correct. The animation is meant to be engaging and motivating. After this animation the end score of the group is shown.

Contribution to the search aggregation and presentation task (to be added in final version)

Evaluation (to be added in final version)

3.3 CollAge

Summary

CollAge is a form of aggregated search in which child-friendly media types are presented alongside traditional search results (Figure 7).

Contribution to the search aggregation and presentation task (to be added in final version)

Evaluation (to be added in final version)

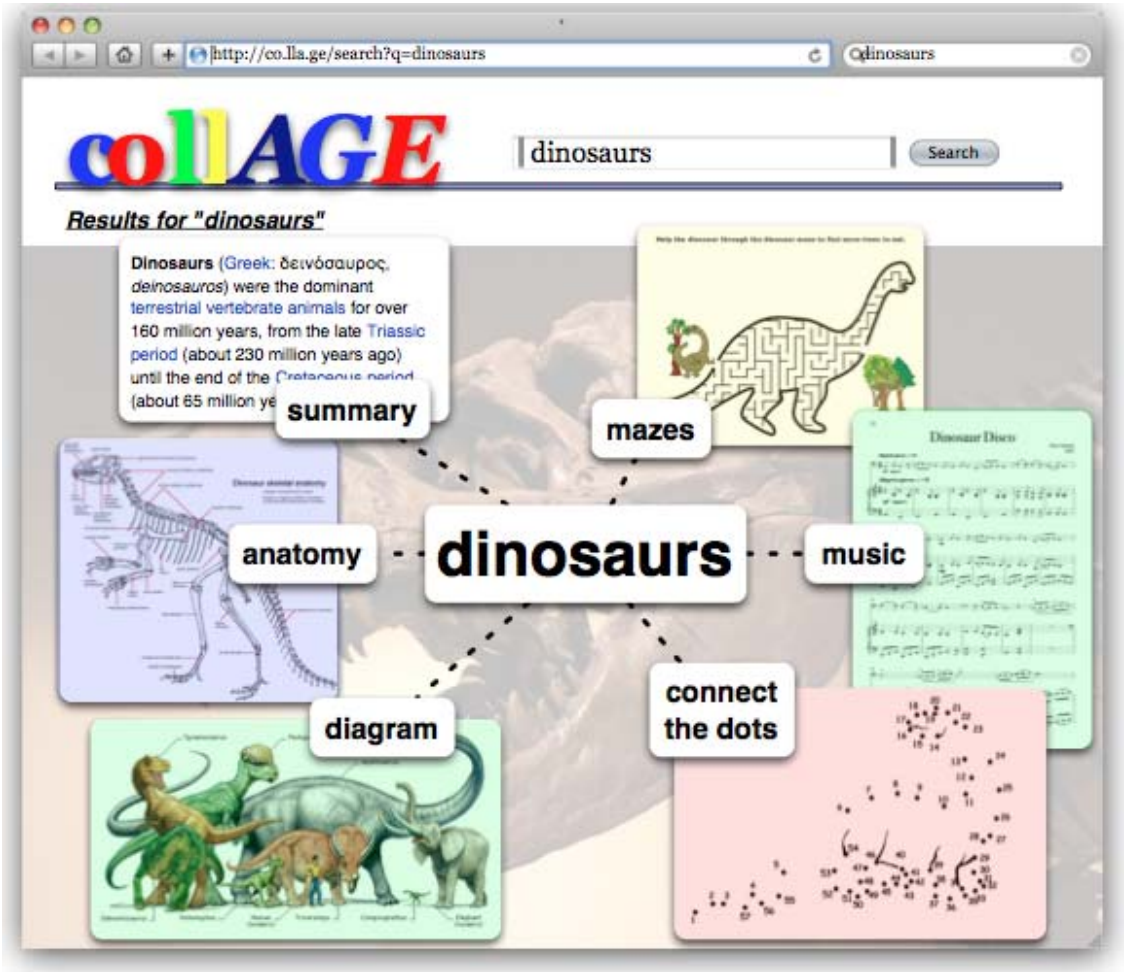

Figure 7: Interface for search results of collAge

3.4 FiFi–Find and Filtering News for Children

Summary

FiFi is a novel news filtering service that was designed specifically for children. A cursory search of news services for children revealed that this was a poorly supported area for children. Existing solutions were either too basic (e.g. CBBC Newsround), or not appropriate for children (Google News).
Figure 8. Search search results in the FiFi-Find interface.

The content categorisation method used in FiFi is based on an information retrieval approach using the Lemur search engine to retrieve recent articles that match a topic request. The articles are converted to JSON and sent to the user interface. The user interface displays topic terms to represent areas of interest. These terms are used to search the indexed collection of entries that have been collected on disk so far. Crucially, as this is a news service and time-order is important, the retrieved entries for a term are ordered not by relevance, but by date. However, the relevance scores are preserved as metadata for the entry so that the user interface can highlight older but more relevant entries.

Figure 8 illustrates the use of FiFi. The set of default topics are shown above the set of discovered topics (1a). In this example, the Science topic has been selected (1b) and the list of document titles filtered for this topic is presented to the child. The amount of space used in presenting a title indicates the relevance of a document, as shown by comparing a somewhat relevant document (2a) and a strongly relevant document (2b). Children are also able to manually define a new topic (3), if needed, and personalise the interface (4), such as the customising the title, colour scheme and style.

When a child clicks on an item in the list, for example (2b), the remainder of the document is presented (not shown). This interaction is used to update the system's knowledge about the child's interests. The history of interaction actions is subsequently used to learn the set of discovered topics (1a), to provide a personalised experience.

Contribution to the search aggregation and presentation task (to be added in final version)

Evaluation (to be added in final version)

3.5 JuSe

Summary

JuSe has been designed to complement existing parent-child educational activities, such as book reading, and has drawn inspiration from educational toys rather than existing search engines. The interface aims to strengthen symbolic associations between words and their written representation. The emphasis is on multiple representations of the same concept, using images, text, and audio. In addition, JuSe can potentially promote Web proficiency among children and provide parents a safe environment to discuss topics concerning Web search.

To address the design concerns associated with browsing and query formulation, the interface combines browsing and picture query-based elements. However, browsing occurs by traversal of keyword categories instead of the document space. Children have a limited vocabulary of content words, mainly concentrated in the high-frequency categories associated with their needs: toys, book and cartoon characters, food, sports, occupations such as doctor or fireman, animals, etc. By identifying these categories from children's vocabulary lists, the JuSe picture dictionary emulates books and television shows used for literacy education. The basic dictionary was compiled by studying lists of children's speech and vocabulary lists of essential English words.

The words used correspond to 'picturable' concepts, mainly nouns, but also simple adjectives. They are arranged in loose categories in a shallow two-level hierarchy. Each category is represented by a word for the parents and an accompanying icon for the children. These categories are condensed into an accordion interface element where only the category titles are shown. When a category title is clicked the corresponding shelf expands and the concept images are revealed. Only one category is shown at a time. Initially the concepts and categories are ordered as they are entered, however, through usage they are reordered so that more commonly used ones are more prominent and easily accessible. As there can be many terms in a category, some scrolling is required, but the need for forward and backward browsing is eliminated by using the accordion interface element.

Querying can be performed by dragging (or alternatively double-clicking) an image corresponding to a concept to the query box on the left (see Figure 9). This type of interaction has a game- like quality and can be performed easily with a mouse or on touch enabled devices. The object identifiers are used to retrieve appropriate query words associated with the image. The query words may not correspond exactly to the description words. For example, it may be advantageous to specify the search terms, e.g. "orange colour" or "orange fruit" instead of just "orange". This feature is included to reduce frustration due to hyponyms.

Contribution to the search aggregation and presentation task (to be added in final version)

Evaluation (to be added in final version)

Figure 9: Search interface of JuSe.

Contribution to the search aggregation and presentation task (to be added in final version)

Evaluation (to be added in final version)

3.6 AmuSe: image and Video Based Browsing and Searching

Summary

AmuSe (see Figure 10) is a multimedia mashup search facility that aims to provide children, between the age of five and twelve, with the opportunity to conduct safe and entertaining searches. Here, mashup is defined to be the merging of two or more sources to create a new service, and so this application merges the results of Google, Bing and Yahoo and present them in a manner considered to be enjoyable for children. As this application is aimed at children the results are moderated to ensure all inappropriate content is removed.

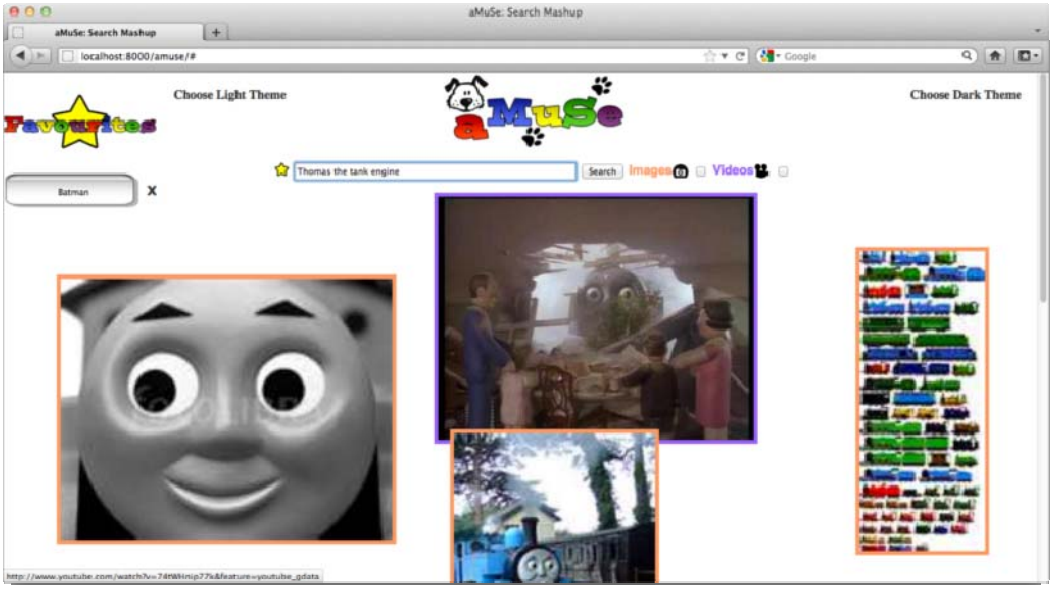

Figure 10: Results interface of AmuSe.

Contribution to the search aggregation and presentation task (to be added in final version)

Evaluation (to be added in final version)

3.7 MaSe: A Customizable Multimedia Mash-up Search Engine for Children

Summary

MaSe (see Figure 11) provides a sandbox environment for high school children to create their own personalised search engine. It is designed as a tutorial in which children can customise, extend and style MaSe to suit their information needs with the use of minimal programming. The application uses PuppyIR as the backend for retrieving and filtering results from various search services and the Django1 Web application framework. Some programming on the part of the child is required to setup search services. Help with the programming is available in the MaSe tutorial as well as an ”answer” file/class. Studies have shown that as a rule, children prefer browsing the web to find information, as opposed to looking for specific answers to a question. MaSe allows

them to create a browsing space in which the desired services can be added and removed as required. From a search query, results of various types are then retrieved from multiple sources, as defined by the services added, leading to new places to browse and explore.

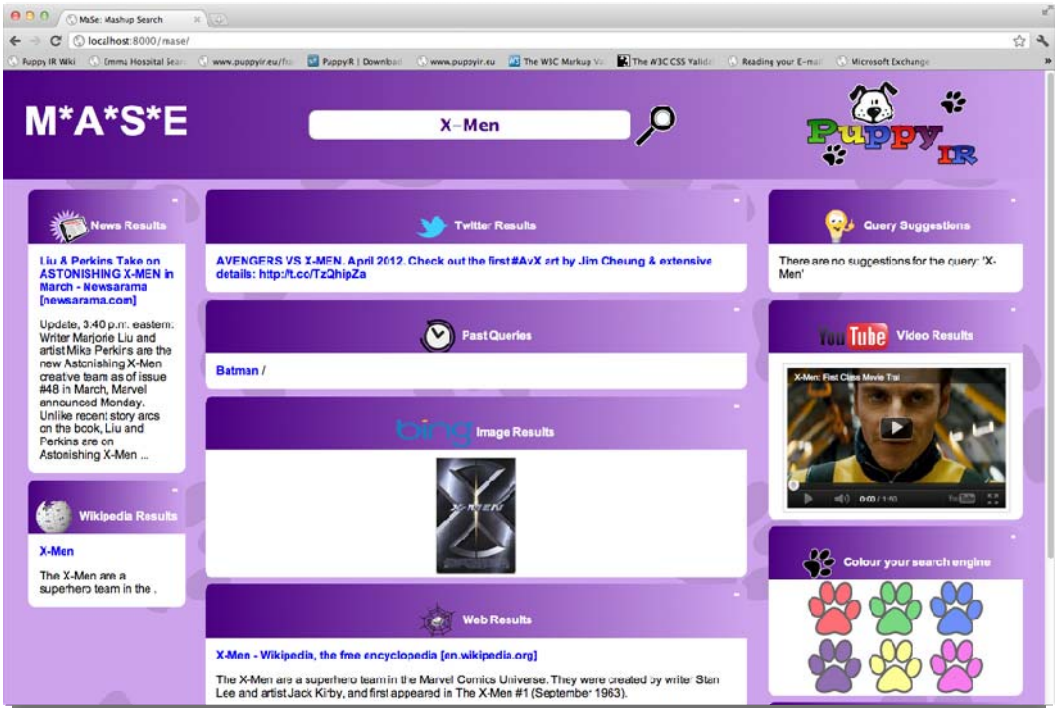

Figure 11: Search results interface of MaSe

Contribution to the search aggregation and presentation task (to be added in final version)

Evaluation (to be added in final version)

4 Conclusion

This first version of Deliverable 3.4 shows work in progress on Task 3.1 of WP3, which focuses on aggregating search results and presenting them to the user in a convenient way. The main demonstrator integrates the benefits of a keyword and a facetted search by exploiting the distribution of predefined categories (here the DMOZ categories of Kids and Teens) assigned to retrieved Web pages. The model defines a way of computing the broadness or narrowness of the search results and presents them as diverse (and possibly summarized through category names) results in case of broad queries and of results expanded with choices of related categories in case of narrow queries. The proposed methods are currently being evaluated and a paper is being prepared to be submitted to the Information Retrieval Facility Conference 2012. The evaluation and the discussion of the results will be part of version 2.0 of D3.4.

Other interesting presentation approaches have been integrated in PuppyIR demonstrators as well. In this version of D3.4 only a summary of this functionality is given. In the final version we will add a detailed description on how the demonstrators contribute to Task 3.1 and what the results of their evaluation are.

References

[1] - Aleck Lin, Shirley Gregor, Michael Ewing, 'Understanding the nature of online emotional experiences: a study of enjoyment as a web experience', Proceedings of the 11th International Conference on Electronic Commerce, ACM, 2009.

[2] - Dania Bilal, 'Children design their interfaces for Web Search Engines: A participatory approach', Proceedings of the 30th Canadian Association of Information Science (CAIS/ACSI) Conference, 2002.

[3] - Kobus Barnard, David Forsyth, 'Learning the semantics of words and pictures', In Proceedings of the International Conference on Computer Vision, 2001.

[4] - Yoëelle S. Maarek, Ronald Faginy, Israel Z. Ben-Shaulz, Dan Pelleg, ‚Ephemeral document clustering for Web applications', IBM technical report, 2000.

[5] - Gong Cheng, Thanh Tran, Yuzhong Qu, RELIN: Relatedness and informativeness-based centrality for entity summarization, Proceedings Semantic Web Conference 2011, 2011.