



info@franz.com  
 http://www.franz.com  
 1-510-452-2000  
 555 12th Street, Suite 1450  
 Oakland, CA 94607 USA

# Allegro CL<sup>®</sup> Success Story

## KPN Research A Portable Webpage Recommender System

KPN Research, the Research Department of the

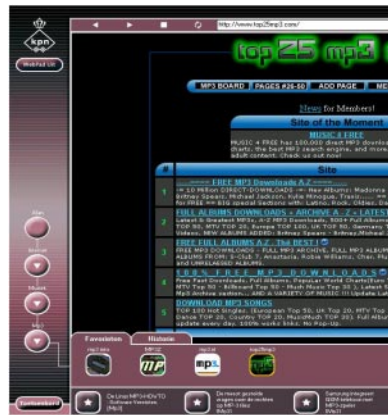
largest Dutch Telecom Operator, is actively designing and testing out new broadband service concepts that make it easier for their customers to use the Internet.

Recently, KPN initiated a pilot project that would give users faster, easier access to web sites that they might not have known about. The project focused on the interface and services that a mobile webpad might have in the Living Room area of the home.

Their goal was to create a bookmark-based browser interface that could be tailored to the information and entertainment needs of every family member, yet was as easy to use as a television remote. It also needed to provide recommendations of interesting web sites based on the users' interests and surfing patterns.

The base technology for this product is a Recommender System. Although quite a bit of literature is available about these systems, information about their strengths, weaknesses and actual field-tested potential, is unavailable or very limited in scope.

The task of the recommendation engine server is to monitor user behaviour and generate recommendations of web pages for each user. This process consists of the following distinct parts: user profiling, content profiling and matching.



**LEGEND:**

- Left Panel:** Shows the current stack for the bookmark tree.
- Bottom Active Tab:** Shows the favorites that belong to this category.
- Bottom Row Recommendation Bar:** Shows the recommendations that belong in this category. Note: The user can add button-icons by selecting a random snapshot from the screen.

**“Given the background of the programmers we might have used Java, but we expected to process millions of URL’s and we have enough experience with Java to know that this would be too slow. C++ was never an option given the short development time granted to us. Lisp was the best solution.”**

*Professor Jans Aasman  
 Manager, Homeservices Group*

The user profiling process builds a user profile for every user of the system. This is constructed by combining explicit topic selection by the user and information gathered by analysing the user's surfing behaviour (by visiting or storing web pages belonging to certain topics, users implicitly show interests in these topics). The implicit data analysis is used to confirm the selections explicitly made by the user and to keep the profile up-to-date.

While user interests are described in terms of topics or categories, so are web pages. Using advanced neural network document classification; web pages are assigned categories describing the page's topic or topics. Apart from the automatic document classification, the user gives input for the classification process as well. A user that places <http://www.cnn.com/weatherforecast> in the folder 'Weather,' unknowingly has classified this web page as being about weather.

Matching user profiles and content profiles is relatively easy. Documents about classical music are recommended to classical music lovers while the latest news in tennis is only recommended to tennis fans.

KPN's web page recommender system was built in Common Lisp using Allegro CL, Allegrostore, Allegroserve and the Clementine neural network engine.

"We had about 6 months to build the recommendation engine with one AI-engineer and two programmers." says Professor Jans Aasman, the Manager of the Homeservices Group. "We also knew beforehand that we would have to experiment with several designs before we would settle for a final one."

"Given the background of the programmers we might have used Java, but we expected to process millions of URL's and we have enough experience with Java to know that this would be too slow. C++ was never an option given the short development time granted to us. Lisp was the best solution." Aasman adds.

KPN's system consists of both server-side and client-side software. Every user has a specialized client-side interface called "The Personal Browser, which allows for

normal web surfing. The interface also contains a 'bookmark-tree' of folders corresponding to the user's interests.

What makes these folders unique, is that they not only contain favorite websites that the user has bookmarked, but recommendations of new websites that the system has found as well. The user's client-side interface is a connection over the web with a recommendation engine server generating the recommendations based on each user's personal interests.

"Our solution is unique." says Alan Verberne, the project leader responsible for the recommendation engine. "First, we combined existing collaborative filtering techniques with advanced document classification software and detailed user monitoring. Secondly, we used hundreds of people to generate and validate an initial bookmark structure for a large number of categories and subcategories. Thirdly, the recommendation engine is integrated in a very user friendly personal browser that has been tested with real users in a number of pilots.

KPN used the AllegroStore object database to store the persistent objects that make up the application. The overall architecture consists of a complex bundle of persistent CLOS classes that reference each other in various ways. "We shudder at the thought of using a relational database for storing these objects." Says Aasman.

AllegroServe was used to prefetch pages in the URL database and to easily program an editorial interface for the system. From this interface a human editor can alter various properties and parameters of the recommendation engine system. "During the development period we encountered some technical difficulties but as always, Franz was quick to provide help and offer solutions." Aasman adds.

Apart from KPN's own usage in a pilot, the software is currently being reviewed by a number of publishers who are interested in using it both for knowledge management purposes within large companies and for distributing their own content (including opt-in advertisements) to consumers.