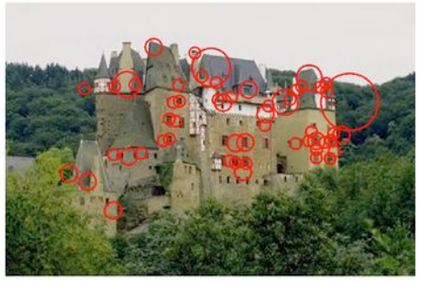

CARTER

The goal of this project is to investigate whether image representations based on local invariant features, and document analysis algorithms such as probabilistic latent semantic analysis, can be successfully adapted and combined for the specific problem of scene categorisation. More precisely, our aim is to distinguish between indoor/outdoor or city/landscape images, as well as (in a later stage) more diverse scene categories. This is interesting in its own right in the context of image retrieval or automatic image annotation, and also helps to provide context information to guide other processes such as object recognition or categorisation.

So far, the intuitive analogy between local invariant features in an image and words in a text document has only been explored at the level of object rather than scene categories. Moreover, it has mostly been limited to a bags-of-keywords representation. Introducing visual equivalents for more evolved text retrieval methods to deal with word stemming, spatial relations between words, synonyms and polysemy is the prime research objective of this project, as well as studying the statistics of the extracted local features to determine to which degree the analogy between local visual features and words really holds in the context of scene classification, or how the local features based description needs to be adapted to make it hold.

Next Generation Information Retrieval

Semantic information recognition and extraction is the major enabler for next generation information retrieval and natural language processing. Yet it is currently only successful in small domains of limited scope. We claim that to move beyond this restriction requires one: (1) to perform integrated semantic extraction incorporating a probabilistic representation of semantic content, and (2) to better employ the broader semantic resources now coming on-line. This project will explore both fundamental research and large scale applications, using the public domain Wikipedia as a driver and a resource. Research will explore the integration of semantic information into the language processing chain. Applications will employ this in broad spectrum named-entity recognition, and in cross-lingual information retrieval using the rich but incomplete data available fron the Wikipedia. Three PASCAL sites will contribute pre-existing software, theory, and skills to the range of tasks involved.

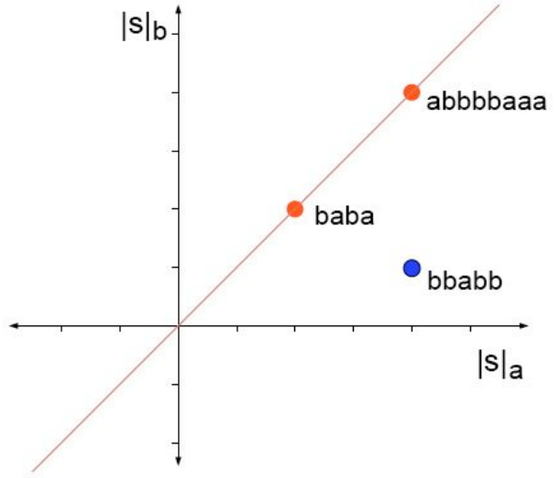

GISK

The purpose of the project is to explore a new family of grammatical inference algorithms, based on the use of string kernels. These algorithms are capable of efficiently learning some languages that are context sensitive, including many linguistically interesting examples of mildly context sensitive languages. The project started on November 1st 2005, and finished at the end of October 2006. It is a collaboration between Royal Holloway and EURISE.

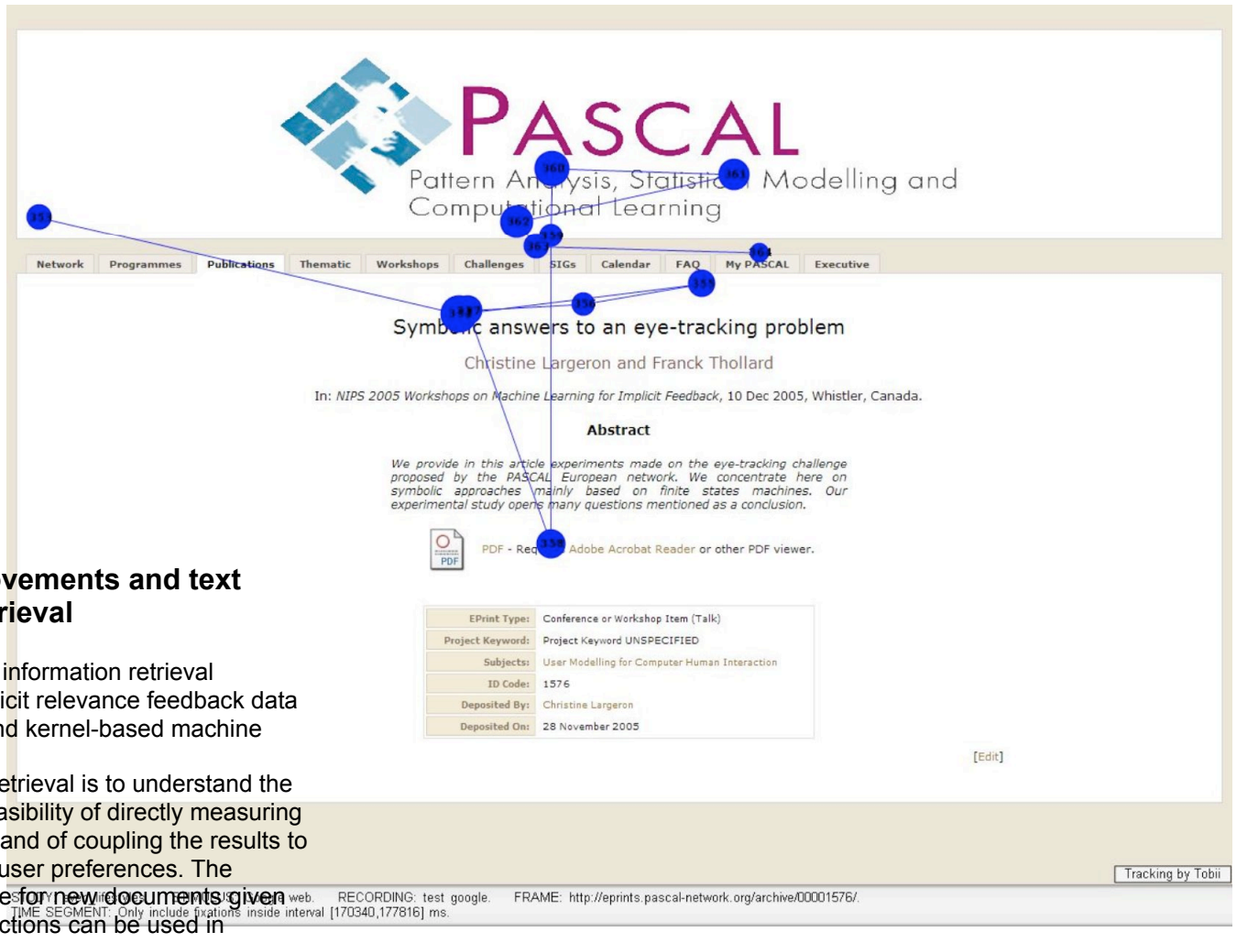

Methods for fusing eye movements and text content for information retrieval

This project develops new kinds of information retrieval systems, by fusing multimodal implicit relevance feedback data with text content using Bayesian and kernel-based machine learning methods.

A long term goal of information retrieval is to understand the "user's intent". We will study the feasibility of directly measuring the interests at the sentence level, and of coupling the results to other relevant sources to estimate user preferences. The concrete task is to predict relevance for new documents given judgments on old ones. Such predictions can be used in information retrieval, and the most relevant documents can even be proactively offered to the user.

The motivation for this project is that by using eye movements we wish to get rid of part of the tedious ranking of retrieved documents, called relevance feedback in standard information retrieval. Moreover, by using the potentially richer relevance feedback signal we want to access more subtle cues of relevance in addition to the usual binary relevance judgments.

The major task in this research is to improve the predictions by combining eye movements with the text content. We aim at combining the relevance feedback to textual content to infer relevant words, concepts, and sentences. We combine two data sources for predicting relevance: eye movements measured during reading and the text content. This is challenging: time series models of very noisy data need to be combined with text models in a task where we typically only have very little data about relevance available.

This novel research task involves dynamic modeling of noisy signals, modeling of large document collections and users' interests, and information retrieval. Multimodal integration and natural language processing are needed to some extent as well. The project also involves a number of interesting challenges from the point of view of applying both kernel and Bayesian methods.