

CONAN: A Framework for contextual recommendations based on augmented domain intelligence

Gbolahan K. Williams, Sarabjot S. Anand
Department of Computer Science, University of Warwick,
United Kingdom
{williamg, s.s.anand}@dcs.warwick.ac.uk

May 30, 2008

General Terms: recommender systems; collaborative filtering; content based filtering; knowledge based systems.

Abstract

It is often the case that users do not know what they are looking for. In such cases (and specific to a particular domain), recommender systems are able to solve this problem either wholly or partially by suggesting items to a user based on previous items consumed by the user (through some sort of domain specific interaction) or items previously consumed by other users of the system.

Technologies such as *collaborative filtering (CF)* and *content based filtering (CBF)* have been used over the years to make this possible, positioning recommender systems as crucial components to e-commerce applications making product discovery much easier for users providing benefits to both application users and businesses. However, many commercially available recommender systems lack the ability to generate recommendations for more complex items - items which are unique and are composed of several complicated related parts.

In order to address this problem, we aim to exploit domain knowledge already available on the web from social websites and other accessible content providers. These portals hold large quantities of data which express numerous individuals interests and opinions about various products and entities in several domains. This user-driven data is yet to be fully exploited by product recommender systems in generating recommendations for complex product items in several domains such as travel, theatre etc. Such social data held by these portals include articles, reviews, blogs, tags, groups and other forms of media content.

Augmenting semantic information derived from this user-driven data unto a fixed or dynamic product domain can alleviate some of the problems encountered when generating recommendations for these complex items. These problems usually arise because of the difficulty of representing the context of the model in a way that can be leveraged by recommender systems and/or the uniqueness of the item (e.g. recommending a *one off theatre event* in the case of a theatre recommender system). A very dominant problem in this field of research is *sparsity*.

In this paper, an approach to solving this problem is described with an example based on a theatre event recommendations system using data provided by LastMinute.com¹. Using empirical studies we aim to investigate how using external domain knowledge can improve the quality recommendations generated by recommender systems.

¹<http://www.lastminute.com>