

Welcome to the Open Relevance Project!

Table of contents

1 What Is the Open Relevance Project?.....	2
2 News.....	2
2.1 14 July 2009 - Wiki now available.....	2
2.2 25 June 2009 - Apache Open Relevance Kickoff.....	2
3 Project History.....	2

1. What Is the Open Relevance Project?

The Open Relevance Project (ORP) is a new Apache Lucene sub-project aimed at making materials for doing relevance testing for Information Retrieval (IR), Machine Learning and Natural Language Processing (NLP) into open source.

Our initial focus is on creating collections, judgments, queries and tools for the Lucene ecosystem of projects (Lucene Java, Solr, Nutch, Mahout, etc.) that can be used to judge relevance in a free, repeatable manner.

2. News

2.1. 14 July 2009 - Wiki now available

The ORP Wiki is now available at <http://cwiki.apache.org/ORP>. Anyone can edit it, so please feel free to help out the ORP!

2.2. 25 June 2009 - Apache Open Relevance Kickoff

The Apache Lucene PMC has officially voted to add the Open Relevance Project (ORP) as a Lucene subproject. ORP's main goal is to build out collections, judgments and queries in an open environment to make it easier for Lucene developers and users to do relevance testing, much like one would get if using [TREC](#) or other evaluation conferences.

3. Project History

For background reading on the genesis of ORP, see:

- <http://wiki.apache.org/lucene-java/OpenRelevance>
- http://mail-archives.apache.org/mod_mbox/lucene-general/200905.mbox/ajax/%3c951F29A3-4461-4B2