

Celebrating Stephen Robertson's Retirement

Djoerd Hiemstra, John Tait, Andrew MacFarlane, Nick Belkin

Stephen Robertson retired from the Microsoft Research Lab. in Cambridge during the summer of 2013 after a long career as one of the most influential, well liked and eminent researchers in Information Retrieval throughout the world.

Stephen began his research career in the late 1960's when he took an M.Sc. in Information Science at City University, before beginning a Ph.D. at University College London, under the renowned Information Scientist B. C. Brookes. Stephen was always more mathematically sophisticated than most Information Scientists having taken a First in mathematics at Cambridge, and this engagement with mathematics characterized his career throughout.

During this period, whilst also working as a Research Assistant at Aslib, the Association of Special Libraries and Information Bureaux, he produced his first papers, the two part: "The Parametric Description of Retrieval Tests" [1], which can still be read with profit by today's researchers. In particular they contain a careful and timeless analysis of issues to do with the presentation of results of performance curves (like precision/recall) and the advocacy of the use of recall/fallout rather than recall/precision.

Stephen left Aslib to study full time with Brookes, on award of a Royal Society Information Science Fellowship, a signal honor. The only other people to have been awarded this Fellowship were Karen Sparck Jones and Keith van Rijsbergen!

After obtaining his Ph.D. in 1976, Stephen moved from UCL back to City where he took up Lectureship in the Department of Information Science, a department with which he has had a continuous association for more than forty years, and indeed continues through his holding of Emeritus Professorial title at City, although in a somewhat different organizational context.

In 1998 he took up a position at the then very new Microsoft Cambridge Research Laboratory, whilst continuing part time at City. Being the modest and charming man he is, Stephen always said he was hired as a playmate for Karen Sparck Jones, the wife of the founding Microsoft Laboratory Director, Roger Needham. We are not so sure: we rather suspect prescience on Roger's part over the emerging importance of search, and the difficulty of hiring his own wife!

In 2008 Stephen took on a third role: returning to University College London as Visiting Professor, a role which he will continue in his retirement.

Stephen Robertson's research at University College London, City University London and Microsoft Research has had a significant and sustained impact on both the theory and practice of search engine technology and information retrieval. His international status and significance for the research and practice in information retrieval was recognized in 2000 when he received the ACM SIGIR Gerard Salton Award. He also received the Tony Kent Strix Award in 1998 from the Chartered Institute of Library and Information Professionals, and many other awards and honors.

Stephen Robertson's work in the 1970's on models for probabilistic term weighting [2] and his probabilistic principle for ranking documents [3] had substantial impact on the theory of information retrieval. Robertson's probabilistic model allows for reasoning about the weighting of search terms, something that the vector space models of that time – that needed so-called *tf.idf* term weighting – were unable to do. The model also allowed re-weighting of search terms based on user feedback, which was hard to obtain at the time, but now available in abundance as user click-

through data of large search engines. His model and term weighting approaches are discussed in all major educational books on information retrieval, and are cited thousands of times in Google Scholar.

Stephen Robertson is one of the main architects, with Stephen Walker, of the experimental search system Okapi, one of the first information retrieval systems that allowed for natural language queries (without boolean operators AND, OR and NOT). Okapi was one of the most successful search systems at the early Text Retrieval Conferences (TREC) that are organized by the US National Institute of Standards and Technology (NIST). The Okapi TREC reports are cited over 2500 times in Google Scholar. At TREC, the Okapi system was not only applied to ad hoc search, but to various other tasks too, including interactive search and information routing and filtering. The TREC adaptive filtering task that was organized by Robertson from 2000 until 2002, researched approaches to filter information from vast streams of documents, a scenario very similar to analysis of Twitter streams today.

With the Okapi team, Robertson aimed at building pragmatic best match term weighting implementations of the probabilistic models that he proposed earlier. One of these approaches, coined BM25 [4], was very successful in research studies, and found its way in many information retrieval systems and products, including open source search systems like Lucene/Solr, Lemur/Indri, Galago, GrapeShot, Xapian, and Terrier. BM25 is used as one of the most important signals in large web search engines, certainly in Microsoft's Bing, and probably in other web search engines too. BM25 is also used in various other Microsoft products such as Microsoft Sharepoint and Microsoft SQL Server. Robertson and his colleagues at Microsoft developed a version of BM25 for fielded search queries specifically for Microsoft products like Sharepoint in 2004 [5]. Arguably, Robertson's BM25 is the most influential information retrieval term weighting and ranking algorithm that exists today. Certainly Bruce Croft, of the University of Massachusetts Amherst, probably the leading academic expert on Information Retrieval in the world, is of the view that the performance of BM25 can only be exceeded by using various forms of machine learning or modeling algorithms on the collection, the query set, click streams and so on.

He continues to publish actively, having had an important paper published as recently as SIGIR 2012.

Stephen Robertson has served as a lively and informative keynote speaker and panel member at many conferences, including ACM SIGIR, ECIR (European Conference on Information Retrieval organized amongst others by the BCS), MIREX (International Conference on Music Information Retrieval), IiiX, (Information Interaction in Context), and NTCIR (NII Testbeds and Community for Information access Research). Robertson has given numerous invited talks, for instance at the S.R. Ranganathan Memorial Lectures, and at UCLA. These are invariably well received.

Stephen has not only been an important intellectual and practical contributor to the emergence of modern Information Retrieval. He has also been an important contributor at a personal level to the development of the IR community.

He is charming and helpful, and always has been. He has an ability to point out what is wrong with someone's thinking in a way which is supportive and encouraging, rather than with any element of point scoring. This is particularly appreciated by his former research students. He deals with all members of the community with patience and equity, whether they are research student at the beginning of their career, or the most senior professor. One of the present authors (Tait) remembers Stephen most firmly, courteously and politely pointing in a program committee meeting that a paper could not be rejected on the basis the ideas had been discussed many years before in Cambridge Coffee Room, unless one could identify where they were written up. Others would have left the

object of the criticism feeling wounded: not Stephen!

Stephen plans to continue to be active in Information Retrieval in his retirement: but “only doing the bits I like”.

I’m sure the whole of the IR community will join the authors in wishing Stephen a long and happy retirement.

- [1] S.E. Robertson, The parametric description of retrieval tests. Part 1: the basic parameters; Part 2: overall measures. *Journal of Documentation* 25, 1-27, 93-107 (1969).
- [2] Stephen Robertson and Karen Sparck-Jones, Relevance weighting of search terms. *Journal of the American Society of Information Science* 27, pp. 129-146, 1976.
- [3] Stephen Robertson, The probability ranking principle in information retrieval, *Journal of Documentation* 29(6), 481-485, 1977.
- [4] Stephen Robertson and Stephen Walker. Some simple effective approximations to the 2-Poisson model for probabilistic weighted retrieval. In *Proceedings of the 17th ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 232-241, 1994. (900 citations in Google Scholar)
- [5] Stephen Robertson, Hugo Zaragoza and Michael Taylor, Simple BM25 extension to multiple weighted fields. In *Proceedings of the ACM Conference on Knowledge and Information Management, CIKM*, pp 42-49, 2004. (385 citations in Google Scholar)