Research Repository UCD

Title	The scholarly impact of TRECVID (2003-2009)
Authors(s)	Thornley, Clare V., Johnson, Andrea C., Smeaton, Alan F., Lee, Hyowon
Publication date	2011-04
Publication information	Thornley, Clare V., Andrea C. Johnson, Alan F. Smeaton, and Hyowon Lee. "The Scholarly Impact of TRECVID (2003-2009)" 62, no. 4 (April, 2011).
Publisher	Wiley
Item record/more information	http://hdl.handle.net/10197/2803
Publisher's statement	This is the authors' version of the following article: "The Scholarly Impact of TRECVID (2003-2009)" published in Journal of the American Society for Information Science and Technology, 62. It is available in its final form at http://dx.doi.org/10.1002/asi.21494 .
Publisher's version (DOI)	10.1002/asi.21494

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

The Scholarly Impact of TREC Vid (2003-2009)

(Pre-Print: Journal of the American Society for Information Science and Technology, 2011)

Clare V. Thornley

School of Information and Library Studies, University College Dublin, Belfield, Dublin 4, Ireland. E-Mail: clare.thornley@ucd.ie

Andrea C. Johnson

School of Information and Library Studies, University College Dublin, Belfield, Dublin 4, Ireland. E-Mail: andrea.johnson@ucd.ie

Alan F. Smeaton

CLARITY: Centre for Sensor Web Technologies, School of Computing, Dublin City University, Glasnevin, Dublin 9, Ireland. E-Mail: alan.smeaton@dcu.ie

Hyowon Lee

CLARITY: Centre for Sensor Web Technologies, School of Computing, Dublin City University, Glasnevin, Dublin 9, Ireland. E-Mail: hyowon.lee@dcu.ie

Abstract

This paper reports on an investigation into the scholarly impact of the TREC Vid (TREC Video Retrieval Evaluation) benchmarking conferences between 2003 and 2009. The contribution of TREC Vid to research in video retrieval is assessed by analyzing publication content to show the development of techniques and approaches over time and by analyzing publication impact through publication numbers and citation analysis. Popular conference and journal venues for TREC Vid publications are identified in terms of number of citations received. For a selection of participants at different career stages, the relative importance of TREC Vid publications in terms of citations *vis a vis* their other publications is investigated. TREC Vid, as an evaluation conference, provides data on which research teams ‘scored’ highly against the evaluation criteria and the relationship between ‘top scoring’ teams at TREC Vid and the ‘top scoring’ papers in terms of citations is analysed. A strong relationship was found between ‘success’ at TREC Vid and ‘success’ at citations both for high scoring and low scoring teams. The implications of the study in terms of the value of TREC Vid as a research activity, and the value of bibliometric analysis as a research evaluation tool, are discussed.

Introduction

In this paper we report on results from a study to investigate the scholarly impact of the annual series of TREC Vid benchmarking conferences. TREC Vid started out as one of several tracks in the larger TREC (Text Retrieval and Evaluation Conference) benchmarking conference series in 2001, and it became a separate activity in 2003. The overall aim of TREC and TREC Vid is to provide access to large scale test collections so that newly developed techniques for content-based operations like search can be tested and compared in an open, metrics-based way and in this way help to progress the field of information retrieval

(IR). TRECVideo uses the same model of evaluation as TREC but the focus is on techniques for digital video whereas TREC focuses on text derived from documents, web pages, blogs, automatic speech recognition, etc. After almost 20 years of activity, the TREC model of evaluation is not without its critics and discussion on its reliability (Robertson, 2008) and, in particular, its validity are widespread in the literature. It does, however, provide the only forum for the large scale testing of new IR techniques.

Our study assesses the scholarly impact of TRECVideo and by this we mean the extent to which the TRECVideo conferences have influenced the development of new thinking and techniques in the field of video retrieval. After 7 years as a standalone benchmarking conference and 2 years as a TREC track it is reasonable to ask whether it has been a successful forum for the development of improved techniques. In a broader sense has TRECVideo been successful in developing new ideas and approaches to the problem of effective video retrieval? In this study investigate these questions using bibliometric tools examining the scientific publications written as result of TRECVideo and their associated citations. This builds upon existing work done by NIST in 2010 on the economic impact study of TREC (Rowe *et al.*, 2010) investigating the economic success of the main TREC activity in developing new IR techniques. It reached some, broadly positive, conclusions regarding the financial return on investment of funding TREC over a 19 year period but it did not examine the scholarly or academic impact of the conference.

What exactly is scholarly impact? How can we measure the effect or impact that TRECVideo has had on video retrieval research both in terms of thinking and practice? How far has the influence of this annual benchmarking conference spread? We can, to an extent, answer some of these questions by examining the number of publications derived from the benchmarking activity in TRECVideo and the number of citations they have received. We defined these as publications that could not have been written if TRECVideo hadn't happened because of their reliance in some way on TRECVideo data and/or the TRECVideo benchmarking process. TRECVideo and TREC in general are different from the majority of conferences because they enable research in the specific way of providing an evaluation and benchmarking process as well as providing a forum for disseminating research. The TRECVideo conference itself is just the final stage of a year-long evaluation process which only participants have access to. The research on new video retrieval techniques could have been done without involvement in the TRECVideo benchmarking. It would be very difficult, however, for researchers to evaluate and compare their results against other possible approaches.

We examine the papers derived from TRECVideo and investigate where they are published and how many citations they each received. We measure the extent to which participation in TRECVideo evaluation has facilitated research which gets through the first hurdle of peer review to get published and secondly, in terms of bibliometrics, receives citations suggesting that other researchers are using the work to progress the field. Success in terms of citations to a work can be seen as wider form of peer response or perhaps use of one's research which although perhaps with less quality control, as reviewers are normally selected for their expertise and knowledge whilst anyone can cite a paper (though in practice is it very likely to be someone else working in the field), does give an indication of the quality and usefulness of the research. Participating in TRECVideo is a means to an end, the ability to evaluate comprehensively one's research approaches, which in turn enable participants to build upon research and to convincingly disseminate their findings to the wider field of computer and information science. We can have reasonable confidence that TRECVideo is successful in its objectives if it makes possible a significant amount of research which is published and then cited. TRECVideo, like the other benchmarking activities TREC, CLEF, MIREX, NTCIR, FIRE, etc., are all open activities whose stated aims are to progress

their respective sub-disciplines within information retrieval. All do this by facilitating scientific investigation into new techniques for analysis, indexing, search, etc., and this investigation is reported through research publications whose impact can be measured through citations. There are other measures of these benchmarks such as numbers of patents, technology transfer or even economic impact as the NIST study from 2010 investigated, but our working premise is that impact is measured by publications and their citations. The purpose of this study is to examine how significant the research that TRECVideo has made possible is through an analysis of the publications which have resulted from TRECVideo. We do this by examining a number of key questions about the publications and citations arising from TRECVideo: how many publications result from TRECVideo; what are they about; how often are they cited; what are the most popular venues in terms of paper and citation numbers; how important are TRECVideo papers to the citation profile of participants; how is TRECVideo 'success' linked to citation 'success'? We address these questions using bibliometric and visualisation techniques. The next section provides a short introduction to TRECVideo to provide a wider context to the study.

TRECVideo: What it is

Information retrieval (IR) research has always had as one of its important components, the systematic and repeatable benchmarking of any new technique for automatic analysis, indexing, retrieval, summarization or other content-based operation. Since the earliest days of information retrieval research (Sparck-Jones, 1981) the pioneers of this field including Luhn, Maron and Kuhns, Salton, Cleverdon, Sparck Jones, van Rijsbergen, Robertson and those who have followed them have recognized the importance of developing effective means of testing the value of new theories, models, and ideas. The development of test collections is one approach, amongst others (see for example the theoretical work in IR by Van Rijsbergen, 2004), of evaluating new potential solutions to the problems of IR.

Up to the end of the 1980s, access to such test collections of data was very limited. Publicly available datasets were small, narrow in scope and over-used and there were no generally available large scale collections of documents, queries, and relevance assessments. In 1991 the National Institute of Standards and Technology (NIST) in the US organized the first Text REtrieval Conference (TREC) with the aim of building such a large scale collection of documents, queries and relevance assessment and allowing uniform evaluation using that dataset and a set of common and shared metrics. The growth rate associated with TREC throughout the 1990s is testimony to its success with increasingly large datasets being made available to the research community and evaluation metrics stabilizing. The IR research community effectively unified around TREC and its tasks so much that TREC started to branch out in terms of the nature of tasks and the variety of (text) data on which benchmarking was taking place. These new data and task types were known as "tracks" and in 2001 TREC launched a new track on video retrieval. This had the usual TREC mode of operation whereby NIST acquired and distributed (video) data to signed-up participants, NIST formulated and distributed search topics which participants executed on the video data using the systems they developed and they then submitted their top-ranked video clips for pooling and manual assessment by NIST personnel. This was then used to calculate performance metrics and at the TREC conference these results and the techniques behind the systems were shared and discussed.

The video track in TREC grew rapidly and in 2003, TRECVideo separated from TREC and became an independent, standalone benchmarking conference. Over the following 7 years TRECVideo operated on a variety of video genres and a range of content-based tasks including automatic detection of video shot boundaries, detection of semantic concepts within

shots, fully automatic, semi-automatic and interactive search for video shots or for known videos, near-duplicate video detection, video summarization, semantic event detection in CCTV and TV news story segmentation. All of these tasks are done in a hugely collaborative and supportive environment with sharing and donation of data and other resources among participants being the default, all in the name of progressing the field of video retrieval.

In this paper we focus on TRECVID during the year 2003 to 2009 inclusive. For the first 4 of those years the video data used was broadcast TV news, initially in the English language but then in 2005 and 2006 also including TV news in Chinese and Arabic. The video was accompanied by speech transcripts derived from automatic speech recognition which was automatically translated into English in the case of Chinese and Arabic. In 2007 TRECVID introduced a new genre of video provided by the Netherlands Institute for Sound and Vision, which consisted of general TV magazine programs. Also in 2007 TRECVID introduced camera rushes video, the raw, unedited video captured by a camera during the recording, and rehearsal of TV shows, provided by BBC. In 2008 TRECVID introduced CCTV camera footage taken in a major international airport.

The task of shot boundary detection ran from 2003 to 2007 at which point progress in the techniques seemed to have reached a plateau. The search task – automatic, manual and interactive – was introduced from the start and continues each year, as does the task of automatically detecting the presence of a set of semantic concepts. Automatic detection of TV news story bounds ran in 2003 and 2004, automatic summarization of BBC rushes video ran in 2007 and 2008, detection of events from surveillance data ran in 2008 and 2009 as did automatic detection of near-duplicate videos.

Participation by research groups in TRECVID increased every year except 2009, peaking in 2008 with nearly 80 groups and dropping to just over 60 in 2009. Participants come from all across the globe, and there is a great geographical spread. Some of the participants are regular and have taken part each year while others have just taken part once or twice; some participants represent larger research groups while others may be just a single PhD student working on a related topic. All, however, take part in the benchmarking in order to test out some new idea or technique they have developed. Each participant is allowed to submit more than one “run” for a task, including search, and participants usually vary some attribute or parameter of their search systems for each of the runs they submit.

Method

A bibliometric study examining both the number of TRECVID publications and the number of times they have been cited gives us both quantitative and qualitative indicators of its scholarly impact. Citations counts are only one indicator of quality as studies show a variety of citing motivations, for a comprehensive review see Bornmann and Daniel (2008), but they do give an indication of the extent to which a publication has made a difference. Recently, there has been a growing recognition of how various data sources and citation metrics may impact on different disciplines and, in particular, the extent to which some established bibliometric tools may disadvantage computer science (Moed & Visser, 2007; Bar-Ilan, 2009). A considerable problem is one of coverage as any bibliometric tool can only accurately measure citations if it has data on all the possible sources of publications and citations in any given discipline. Computer science publications are often conference or technical reports which are not comprehensively included in many of the standard citation analysis tools. Harzing (2010) analyses three different sources for citation analysis, ISI Web of Science, Scopus and Google Scholar across academics working in the Sciences and Social Sciences and Humanities. Computer Science is one of disciplines investigated and over a four year period she found that Computer Science was different from most sciences as

Google Scholar provided five times as many citations as ISI. Another recent study by Freyne *et al.* (2010) also examined the citation scores from Google Scholar and ISI Web of Science for publications in computer science and found that lack of conference coverage by ISI put computer scientists at a disadvantage in evaluations based on ISI.

The table below provides a first pilot of our search results using Scopus and Google Scholar. In 2008, for example, Scopus yielded 130 documents published as a result of TRECVID activity, Google Scholar for the same year has 586.

<u>YEAR</u>	<u>No. of Publications Scopus General</u>	<u>No. of Publications Scopus More</u>	<u>No. of Citation Documents</u>	<u>No. of Citations</u>	<u>H-Index</u>	<u>No. of Publications Google Scholar (Using PoP)</u>	<u>No. of Citations</u>	<u>H-Index</u>
2010	26	33	22	62	1	97	25	2
2009	85	40	62	56	4	401	590	11
2008	130	154	112	114	6	586	2516	19
2007	126	181	195	335	8	411	3655	28
2006	71	241	263	524	11	332	3784	31
2005	34	192	253	998	15	212	2497	28
2004	39	166	214	848	14	171	2195	23
2003	1	145	191	1010	14	58	1180	16
2002	0	3	3	3	2	9	52	4

TABLE 1: Initial pilot results showing comparison of Scopus and Google Scholar.

After consultation with experienced practitioners of citation analysis within computer science we decided, based on the issue of coverage, to use Google Scholar as our main source. ‘Publish or Perish’ (PoP)¹, a software wrapper for Google Scholar, was used to manipulate Google Scholar searches. For each year we searched in PoP using TRECVID as a search term which then provided a list of publications with associated citation counts. TRECVID was selected as a search term as TREC Video Retrieval Evaluation returned TREC papers in addition to TRECVID papers. Whilst Google Scholar has the coverage we needed, its tools for checking duplicates and its ability to deal with large data sets, are not as advanced as its more established alternatives which presented some challenges in terms of manipulating and cleaning the data. Google Scholar will return a maximum of 1000 records, therefore searching across 2003 – 2009 was not possible in a single search. Each year had to be individually searched, with the search criteria being ‘year’ AND ‘TRECVID’.

The initial search results derived from the PoP search were checked and cleaned by an expert in the field of video retrieval. For each year we checked all the PoP search results with the criterion for inclusion being ‘was this publication a direct result of TRECVID activity?’. By this we mean that the paper uses TRECVID data or benchmarking criteria or

¹ <http://www.harzing.com>

describes a technique tried in TRECVID. We excluded papers which just cited or mentioned TRECVID as our aim was to include papers truly derived from TRECVID participation. Clearly there were some papers which included the term TRECVID but were not necessarily about TRECVID. This cleaning eliminated most duplicates and also papers which were only tangentially related to TRECVID but which had been retrieved in the original literature search. PoP does not have the capability of saving data and running further searches from the saved data set. Therefore the duplicates and papers not related to TRECVID were then manually removed from each year of the PoP search results and saved as Excel files for further reference and manipulation. The precision of the data set is, therefore, reasonably reliable but accurately checking the recall is more difficult and, despite our broad search strategies, some papers will have been missed. In terms of the results this means it is most likely that this study slightly under-reports the extent and impact of TRECVID publications. Our aim was to gain a good overview of the impact of the conference over time and thus we aimed for broad coverage of publications with checking to increase precision.

Citation analysis is more commonly used for a single or a small group of authors in which case near complete accuracy is paramount and possible. It is relatively straightforward, for example, to confirm that the results of a citation search are by a particular author as we found in our author analysis, whilst it is much difficult to definitively claim that they are about a particular topic. Indeed if it was easy to establish whether a document is about a particular topic IR would not be the perennially difficult problem it has proved to be. This TRECVID analysis encompasses multiple authors, institutions, and publication types which share a relationship to the topic.

The data set of publications

Our main focus for this study was an investigation into the number and impact of publications written as a result of, or relying upon data from, the TRECVID conferences 2003-2009. The data set 'TRECVID derived papers' includes both the TRECVID conference papers known as workshop papers and published online on the TRECVID website, and papers published in different venues but based on TRECVID work. TRECVID publishes workshop papers describing how well the research groups' techniques did against the evaluation process. These are not refereed, and *most* participants produce a paper, but unfortunately not all. Our initial pilots showed that some TRECVID workshop papers were often highly cited which would suggest that they are used to support papers which have been published in other venues. This gives us an indication that they are having a scholarly impact. Thus, to gain an overall picture of the scholarly impact of TRECVID, it made sense not to exclude them with the 'TRECVID derived publications' of which they consist of, on average, approx 15% of the total. In the first year of TRECVID, they consist of a much greater percentage of approximately 50%, but as the conference matured, more papers were generated for other venues, see Table 2.

After each annual TRECVID conference, many participants publish more detailed descriptions, or further experiments, or comparisons, or overviews, elsewhere, in journals, conferences or workshops. These will in nearly all cases have gone through a competitive peer review process to get published. So we examine all publications that exist because of the TRECVID conferences, either directly as a workshop paper, or indirectly but which could not have been written without the use of TRECVID data in some way. We then examine the citation patterns of these publications. The transition from non-referreed workshop paper, to peer-reviewed publication to receiving citations we envisage as different stages of potential scholarly impact. Peer-review proceeding publications suggests that the reviewers (experts in the field) have agreed that the research is likely to make a useful contribution to the field and

thus is worth publishing. When this research receives citations this suggests that a a broader set of peers, sometimes over along period of time, have indeed found the work an useful contribution to progressing their own work.. Thus publication and citation counts give us some indication of the overall impact of TRECVID related research.

We used this publication and citation data to investigate a number of questions, as outlined earlier, and these are discussed in more detail in the next section. The overall aim of these questions is to provide insights into the quantity and quality of TRECVID's contribution to video retrieval research. We also present some initial investigations into possible factors that may influence the citation rates of different publications to see if these can inform our understandings of bibliometrics as a measurement tool for scientific quality.

Key questions and answers

This section consists of a series of sub-sections, each examining one of the questions we raise about TRECVID, and providing analysis and answers.

How many TRECVID publications are there?

Our study shows that for 2003-9 there were a total of 2,073 TRECVID-derived publications of which 310 were TRECVID workshop papers. As can be seen from the table below, as the conference matures, more publications reach venues outside the conference itself.

<u>YEAR</u>	<u>No. TRECVID papers at CONFERENCE</u>	<u>No. TRECVID derived publication</u>	<u>No. of citations</u>	<u>Cites per paper</u>	<u>H- Index</u>	<u>G- Index</u>
2003	28	64	1,066	16.66	18	30
2004	30	158	2,124	13.44	24	40
2005	37	225	2,537	11.28	28	41
2006	50	361	4,068	11.27	30	52
2007	48	382	3,562	8.97	28	45
2008	64	509	1,691	3.32	16	23
2009	53	374	780	2.09	12	20
Totals	310	2,073	15,828			

TABLE 2: Overview of data 2003-2009.

There is a steady overall increase in publication outputs, which is in line with the increase in participation. The year 2008 produced a particularly high number of publications as is shown in chart of publication trends below. This coincides with the year of greatest participation in TRECVID though one would expect at least a 1-year time lag with the publications from a TRECVID year following at least one year later.

FIG 1: Overview of publication trends (Google Scholar using PoP).

What are TRECVID papers about?

We also used the publication data to analyse and to visualize how the topics treated in TRECVID papers have developed and evolved year-on-year since the start of TRECVID. Similar analysis of topic development using tri-occurrence mapping, in IR as a discipline has been done by Sugimoto & McCain (2010). We examined the titles of all TRECVID-derived papers, the titles and abstracts of the most highly cited papers, and then the titles and abstracts of all TRECVID workshop papers. We were unable to use the abstracts for all the TRECVID-derived papers as the data set is just too large to make retrieving all the abstracts possible within the confines of this study. We thus prioritized a more detailed analysis of highly cited papers and of the actual TRECVID workshop papers.

Using titles of all TRECVID-derived papers

Using the titles of all 2,073 TRECVID-related papers in conferences, journals and workshops we generated word clouds for each year and compared between the years. This helps us analyse how popular sub-topics in TRECVID activities come and go each year and sometimes re-emerge in the later years. Aggregations done were (1) singular and plural terms were put into a singular term and counted together, (2) Capitalised and uncapitalised terms were all turned to uncapitalised with the first letter capitalised and counted together. No semantically similar terms were aggregated.

Year	Most frequently used terms that year	Tasks exercised that year
2003	Shot, Segmentation, Boundary, Features, Framework, Transcript, Browsing	<ol style="list-style-type: none"> 1. Shot boundary determination 2. News story segmentation 3. High-level feature extraction 4. Search
2004	News, Segmentation, News, Semantic, Interactive, Story, Features	<ol style="list-style-type: none"> 1. Shot boundary determination 2. News story segmentation

		<ol style="list-style-type: none"> 3. High-level feature extraction 4. Search
2005	Semantic, Extraction, News, Concept, Annotation, Classification, Learning	<ol style="list-style-type: none"> 1. Shot boundary determination 2. Low-level feature extraction 3. High-level feature extraction 4. Search 5. Rushes exploitation
2006	Semantic, News, Learning, Annotation, Segmentation, Concept	<ol style="list-style-type: none"> 1. Shot boundary determination 2. High-level feature extraction 3. Search 4. Rushes exploitation
2007	Semantic, Learning, Concept, Annotation, Classification, News	<ol style="list-style-type: none"> 1. Shot boundary determination 2. High-level feature extraction 3. Search
2008	Semantic, Annotation, Concept, Summarization, Learning, Rushes, Event	<ol style="list-style-type: none"> 1. Surveillance event detection 2. High-level feature extraction 3. Search 4. Rushes summarization 5. Content-based copy detection
2009	Semantic, Concept, Annotation, Classification, Segmentation, Adaptive	<ol style="list-style-type: none"> 1. Surveillance event detection 2. High-level feature extraction 3. Search 4. Content-based copy detection

TABLE 3: Most frequently used terms in titles 2003-2009.

Firstly, Table 3 shows how in general the topics of “Shot”, “Boundary”, “Segmentation”, “Features” (which have been the main research issues and interests since the research in video information retrieval took off in mid 1990s) were replaced with topics such as “Semantic”, “Concept” and “Learning” over the years. This represents the TRECVID community shifting interest, and incidentally the maturing of the field from a low-level, feature-oriented exploration to a high-level, semantic-oriented one. This change of dominant topics over the years is also in some way steered and guided by the “tasks” introduced in TRECVID each year as seen in the above table.

Using only paper titles might have been limited in terms of finding popular or important terms for specific approaches or techniques used each year as titles tend to pertain only high-level topic-related terms rather than more detailed, technical terms. Acknowledging this, we then extracted full titles and abstracts from the top 10 most frequently-cited papers for each year to see if there are any obvious trends visible (as extracting all abstracts or full-text from all the papers would have been impractical). The result shows that while more technical terms indicating specific approaches or angles appear in the most frequent terms list (e.g. “classifiers”, “SVM”, “categorization”, “tags”, “speech”, “texture” and “edge”), there was no obvious change of usage frequency of these terms over the years. By using 10 most

frequently cited papers of each year, we were capturing terms that were not only most used each year but propagated into the past years as “cited” means cited in the years subsequent to publication.

Thirdly, we were interested in finding out frequency of terms used in each year’s TRECVideo workshop papers to see what topics and methods were mentioned most frequently as a rough indication of their popularity or perceived interests in that year. We simply counted words frequency after removing stop-words such as “and”, “of” and “the”, as well as our own “TRECVideo stop-words” such as “TRECVideo”, “video”, “retrieval”, “search”, “baseline”, “experiment” and “digital”.

We performed term frequency analysis on the titles and abstracts from all 310 TRECVID workshop papers, in order to capture more technical terms that appeared in abstracts and also to represent those that appeared purely within the context of TRECVID participation. Popular aspects or approaches in each year, can, in general be glimpsed in the term analysis with TRECVID workshop papers. Using this result, we generated a simple word cloud visualization taking top 35-40 terms for each year with Wordle², and put each year's cloud next to each other to visually inspect the rising and falling popularities of terms (see Figure 2).

FIG 2: Wordle Visualization of top 35-40 terms used in titles and abstracts from TRECVideo workshop papers 2003-2009.

² <http://www.wordle.net>

As can be seen in Figure 2, low-level technical terms such as “shot”, “boundary”, “colour”, “motion” become smaller as years go on. Also notable in this visualization are the terms “concept” that appeared from 2005 and growing larger each year; the term “SVM” (Support Vector Machine) that appeared from 2005 and grows larger and more or less staying on thereafter; the term “fusion” that started in 2006 and became very popular subsequently; the term “SIFT” (Scale-Invariant Feature Transform) that appeared in 2007 and grows bigger and bigger each year; terms “ASR” (Automatic Speech Recognition) and “text” appear throughout the years indicating that the use of non-visual cues to help video retrieval has been attempted throughout the TRECVID activities.

Finally, we created a bead diagram representing the top 20 terms appearing across the 7 years of TRECVID where the font size of the word represents the importance of that word in that year. This is shown in Figure 3.

FIG 3: Bead plot of the most important 20 words across the 7 years of TRECVID.

Similarly to the trend seen in Figure 2, the bead plot in Figure 3 also shows the diminishing of some topics (“shot”, “boundary”, “ASR”), the growth of others (“concept”, “fusion”, “high-level”, “SVM”, “training”) while other topics remain fairly constant. This is in line with expectations and corresponds to the ending of tasks (shot boundary, use of ASR

in search), and the emergence of new techniques for high-level concept detection based on training support vector machines (SVM).

How often are TRECVID papers cited?

The total number of citations over the time period of the conference was 15,828 and the average cite per paper was 9.58. The citation rates of TRECVID papers are skewed in that a small number of papers receive a very large number of citations with this quickly tailing off. This shows that the citation patterns of TRECVID papers conform to a distribution pattern which has often been observed in other bibliometric studies (Price, 1976). The charts below show citation distribution rates for the year 2007 and also citation trends between 2003-2009. Note that 2007 is chosen as a representative year as it is still relatively recent whilst not being so recent that it is likely to accrue many more citations than it already has.

FIG 4: Citation distribution 2007.

FIG 5: Citation trends 2003-2009.

The TRECVID-derived publications in 2006 have, so far, the largest number of citations. One of the reasons for this is that there is a recommended citation suggested to participants for when TRECVID is referenced in scholarly publications, and the 2006 recommended citation received over 400 citations. This paper skews the peak in 2006 somewhat and if it was removed then the distribution of citations across the years would be more even. We didn't remove it as clearly its citations do indicate that it was an important paper.

Our next two questions investigate some possible reasons why some papers are cited much more than others. Firstly, we investigate whether there is a relationship between success at TRECVID in terms of system performance against the evaluation criteria, and success in terms of citations. Secondly, we examine where highly cited papers are published to see if any particular venues appear to attract more citations. In both cases, particularly the latter, it is problematic to assert causation due to the multiple other factors that can influence citation, but some patterns can be observed.

Does 'good performance' in TRECVID lead to 'high performance' in citations?

Do teams who develop techniques which score highly in TRECVID then go on to produce papers which then also score highly in terms of the number of cites they get, or to put it another way, do people tend to cite papers that describe techniques which were successful at TRECVID more than papers which describe less successful techniques?

We investigated this question by firstly identifying the top performers and the lower performers at TRECVID 2006 using the criteria discussed below. We then analyzed the citation rates of TRECVID derived papers published in 2007 written by those team members, with the assumption that these would have been mainly about work done in 2006.

TRECVID 2006 was the final year of the cycle of using broadcast TV news as the video source before moving on to use video data from the Netherlands Institute of Sound and Vision. The tasks in 2006 were shot boundary detection (SBD), feature detection, and search,

the latter two based on a master shot reference supplied by the organizers. A rushes video summarization task was also on offer as an exploratory task but few groups completed this and there was no formal scoring or feedback to participants in this task. The video used for the SBD, feature and search tasks consisted of 159 hours of TV news from November/December 2005, with news being from TV stations speaking English, Chinese and Arabic, with most of the data being Arabic. Output from an Automatic Speech Recognition (ASR) system was provided for the video, with machine translation into English for the Chinese and Arabic. All this meant that the quality of the text (from ASR or from ASR followed by machine translation) was quite poor in terms of accuracy, forcing participants to focus on visual aspects of content-based retrieval.

In addition to the master shot reference, the MediaMill group at the University of Amsterdam provided the output of 101 automatic feature detectors on the search data, and a group from Columbia University, Carnegie Mellon University and IBM provided the output of manual annotation by 449 features from the LSCOM (large scale concept ontology for multimedia) ontology also on the search data, for all participants to use.

In 2006, 54 participating groups completed one or more of the tasks, broken into 26 who completed SBD, 30 who completed feature detection and 26 who completed at least one form of the search task. Many of these groups went on to publish further details on their TRECVID 2006 work elsewhere, but determining which were the best-performing groups in order to correlate that with subsequent publication and citation is difficult because not all groups did all tasks and even for those who did, they may have performed better in some tasks than in others. This means that a ranking of groups taking part would not only be against the spirit of participation in TRECVID but would also be impossible.

Instead, we have selected from among the 2006 participants, two clusters of threeparticipants each, all of whom have taken part in both the feature detection and the search tasks. We rationale this on the basis that these are the most difficult of the tasks and groups who have completed both these make a serious and large commitment to participation. The first group of three, the high-scorers, all scored highly in each task, and overall were the top performers whereas the second group, the low-scorers, consists of teams who have low-range performances. We then searched for papers written by these teams in our POP results for 2007 and calculated the number of papers for each group (the top scoring and the low scoring groups) and their associated citation counts. We used the measure of cites per paper (see figure 8) to allow for a comparison between the groups despite their different publication numbers.

The results of our analysis show that top scoring teams tended to produce more papers (see figure 6 and figure 7) and also, perhaps more interestingly, that there was a strong connection between high performance in TRECVID and high performance in citations. This was not the only factor in high citation counts, for example, as observed in other studies (Asknes, 2006), review papers were often in the top-cited papers of each year. The key findings are that of the TRECVID 'top scorers' identified, they do nearly twice as well than average in their citation scores, and three times as well as 'low scorers'. TRECVID 'top scorers', however, do not do as well in citation count as some other papers in the top 25 most cited papers so other factors are also at play. Being a 'top scorer' is, however, a good indicator of citation success as is shown in table below for cites per paper comparison of high performing TRECVID team versus the average cites per paper for the entire year. Further detailed analysis of all the top scoring papers, as done recently for ACM published papers by Wainer, de Oliveira & Anido (2010) would provide further insights on these questions but was beyond the current scope of this study.

The table and charts below shows the publication and citation impact for 2007 (based on research done in TRECVID 2006) broken down by ‘all papers’, ‘low scoring teams’ and ‘top scoring teams’.

Breakdown of figures for research teams in 2007.			
	Cite count	Paper count	Mean cite per paper
3 top scoring research teams	990	55	18.0
3 low scoring research teams	48	8	6.0
Other papers	2524	319	7.9
Overall in 2007	3562	382	9.3
Overall without top scoring	2565	327	7.84
Papers of other research teams in top 25 cited papers	1011	15	67.4
Top 25 cited papers overall	1564	25	62.56
Top 3 scoring research teams in the top 25 cited papers	553	10	55.3

TABLE 4: Research teams 2007.

Percentage cites of TRECVID related papers in 2007

FIG 6: Percentage of citations of TRECVID related papers 2007.

Percentage breakdown of TRECVID related papers in 2007

FIG 7: Percentage breakdown of TRECVID related papers 2007.

Average cite count per paper for TRECVID related papers in 2007

FIG 8: average cites of TRECVID related papers 2007.

Teams who performed less well at TRECVID also did less well in terms of citations suggesting that papers discussing techniques which are not successful are not cited as much as papers discussing more successful techniques. This is, perhaps, not altogether unexpected, but it raises some interesting question about the relationships between citations, technological progress and science. Data on techniques which don't perform well at TRECVID still make an important contribution to progress by eliminating certain lines of development. In terms of retrieval performance they may be not be successful but, in terms of science, they form part of the progress. The goal is collective and it is in the best interests of the field if a variety of techniques are tried out some of which, inevitably, will do less well than others.

Where are (highly cited) TRECVID papers published?

Here we examine the publication venues of TRECVID-derived papers to analysis both popular venues (which publish a high number of TRECVID-derived papers) and high impact venues (where TRECVID-derived papers attract a lot of citations). Conference and journal venues 2007-2009 were investigated. We split the publications data set by source, either conference or journal, and then ranked them by citation counts. This also provided an overview of the relative importance of journals and conferences in terms of the publication and citation count of the TRECVID papers, see table and chart below. We see that although in line with other bibliometric studies within computer science, conferences are significantly more important in terms of overall numbers of publications journals are, in fact, more important in terms of citations received on average per paper. This raises some interesting questions about scholarly impact in terms of whether it is better to have a large number of papers widely distributed or a smaller number with higher citations. The response to this is almost certainly different when looking at a topic analysis as in this study as opposed to the publications of a particular individual.

FIG 9: Paper counts and mean cites to TRECVID-derived papers in journals and conferences 2007 – 2009.

Conferences 2007-2009

Figure 10 shows the ranking of top 10 conference venues in 2007, 2008 and 2009, which had highest number of cites to TRECVID-derived papers in each year.

2007	2008	2009	
1. CIVR (565)	1. ACM MM (147)	1. ACM MM (51)	ACM CHI SIGCHI Conference on Human Factors in Computing Sysmtes
2. ACM MM (497)	2. CIVR (130)	2. CIVR (45)	ACM MM ACM International Conference on Multimedia
3. MIR (238)	3. MIR (91)	3. CVPR (41)	BMVC British Machine Vision Conference
4. ICME (114)	4. ECCV (60)	4. ICCV (17)	CBMI International Workshop on Content-based Multimedia Indexing
5. SAMT (75)	5. ICME (47)	5. CLEF (15)	CIVR International Conference on Image and Video Retrieval
6. CVPR (42)	6. BMVC (33)	6. ICME (13)	CLEF Workshop on Cross-language Information Retrieval and Evaluation
7. KDD (17)	7. SIGIR (31)	7. ICML (9)	CVPR IEEE Conference on Computer Vision and Pattern Recognition
8. ACM CHI (16)	8. CVPR (28)	8. WWW (5)	ECCV European Conference on Computer Vision
9. SIGIR (13)	9. ICIP (20)	9. ICDMW (4)	ICCV International Conference on Computer Vision
10.ICCV (13)	10.CBMI (20)	10.WIAMIS (4)	ICDMW IEEE International Conference on Data Minding Workshops
			ICIP IEEE International Conference on Image Processing
			ICME IEEE International Conference on Multimedia and Expo
			ICML International Conference on Machine Learning
			KDD ACM International Conference on Knowledge Discovery and Data Mining
			MIR ACM International Conference on Multimedia Information Retrieval
			SAMT International Conference on Semantic and Digital Media Technologies
			SIGIR SIGIR Conference on Information Retrieval
			WIAMIS International Workshop on Image Analysis for Multimedia Interactive Services
			WWW International World Wide Web Conference

FIG 10: Top 10 conferences by year where TRECVID-derived papers was mostly frequently cited 2007-2009 (number in brackets shows the total number of citations to TRECVID-derived papers in the conference that year).

As the figure shows, between 2007 and 2009, TRECVID papers were consistently being cited especially at the high-profile multimedia/image processing venues such as the International Conference on Image and Video Retrieval (CIVR, rank 1 in 2007, then rank 2 in 2008 and 2009) and the ACM International Conference on Multimedia (ACM MM, rank 2 in 2007, then rank 1 in 2008 and 2009). This shows that TRECVID work is successfully getting through the peer review process and then also receiving citations after publication. Ranking the quality or impact of these conferences through objective means, as opposed to solely examining their importance for TRECVID work, proved difficult as a conference can clearly be important for the image and video retrieval community but still relatively small within the whole field of computer science. In order to provide the most precise overview of conference quality and impact a number of rating systems were identified and each conference was rated using one or more of the systems, dependant on inclusion in the rating system. Not all conferences could be rated in this way, so were possible an h-ndex figure was sought. Conference rating systems included CORE (Computing Research and Education Association of Australasia (<http://www.core.edu.au/>) and the SJR Impact Factor ([SCImago Journal Rankhttp://www.scimagojr.com/](http://www.scimagojr.com/)).

Of the conference venues which consistently published TRECVID work and had a CORE rating, examples include ACM Multimedia, ACM SIGIR and ICCV all with CORE ratings of A+ (the highest). Note these ratings are 2009 figures and while hard-core image/video processing and computer vision conferences such as CVPR and ICCV are seen citing TRECVID papers during these three years, less image/video-centric events such as ACM CHI (International Conference on Human Factors in Computing Systems) and CLEF (Workshop on Cross-Language Information Retrieval and Evaluation) are also seen citing TRECVID papers indicating its impact spilling over to other neighbouring disciplines. It is, of course, difficult to know the relative influence of the quality of the conference venue or the quality of paper in terms of attracting citations but we can confirm that TRECVID papers are appearing at a widespread set of venues.

Journals 2007-2009

Figure 11 shows which journals citing TRECVID-derived papers are published in.

2007	2008	2009
1. IEEE Trans. Multimedia (294)	1. IEEE Trans. Multimedia (100)	1. IEEE Trans. Pattern Analysis and Machine Intelligence (138)
2. Pattern Recognition (191)	2. Proc. the IEEE (72)	2. Computer vision and image understanding (55)
3. IEEE Trans. Circuits and Systems (108)	3. J. Visual Communication and Image Representation (35)	3. IEEE Trans. Multimedia (49)
4. ACM Trans. Multimedia Computing, Communications and Applications (104)	4. Computer vision and image understanding (32)	4. Foundations and Trends in Information Retrieval (26)
5. IEEE Trans. Pattern Analysis and Machine Intelligence (39)	5. IEEE Trans. Systems Man and Cybernetics (27)	5. Pattern Recognition Letters (22)
6. Information Systems (38)	6. Annual Review of Information Science and Technology (21)	6. IEEE Trans. Image Processing (18)
7. IEEE multimedia (37)	7. ACM Trans. Multimedia Computing, Communications, and Applications (16)	7. IEEE Trans. Circuits and Systems (11)
8. Image and Vision Computing (29)	8. IEEE Multimedia (16)	8. Machine Learning (11)
9. Signal Processing: Image Communication (26)	9. J. Electronic Imaging (12)	9. Multimedia tools and applications (10)
10. Multimedia Systems (22)	10. J. Information Science (11)	10. Pattern Recognition (9)

FIG 11: Top 10 journals by year where TRECVID-derived papers are most frequently cited 2007-2009 (number in brackets shows the total number of citations to TRECVID-derived papers in the journal that year).

An important journal for TRECVID papers is IEEE Transactions on Multimedia (rank 1 in 2007 and 2008, then rank 3 in 2009) as, apart from 2009, it is the journal which receives the highest total number of citations for all its TRECVID-related papers. Its JCR impact factor (2009 5 years) is 2.372 putting it just outside upper quartile in JCR rankings so in a similar way to the conferences we can see that TRECVID papers are being published in the top quality computer science journals covering the field. IEEE and ACM Transactions seem to be the most popular journals during these three years where TRECVID-derived papers are cited. Some journals only occur in some years, for example, in 2008 the Annual Review of Information Science and Technology and the Journal of Information Science (ARIST has a JCR impact factor of 3.030), traditionally more information science than computer science publications, are in the top ranking. ARIST 2008 had a paper on ‘Visual image retrieval’ by Peter Enser (2008a) which explains its ranking in that year and, likewise, the Journal of

Information Science had a paper by the same author (Enser, 2008b) on ‘The evolution of image retrieval’. These were current ‘state of the art’ papers reviewing progress and some of their content discussed the role of TRECVID but they are clearly a different kind of paper than one by a participant describing new breakthroughs or techniques.

Impact on Careers

We now look at the impact and influence that TRECVID has had on the careers of individuals by examining the publication and citation patterns of 5 typical TRECVID participants who range from early to late career stage. The total number of participants in TRECVID from 2003 to 2009 is 1,099 but here we select a sample of 5 in order to examine the role that TRECVID publications have played in their publication output and citation counts when compared to their non-TRECVID papers between 2003 and 2009. These were selected as representative by experts in the field and the objective is to provide an indication of the role of TRECVID publications in the overall publication and citation output of some TRECVID participants ranging from junior to senior. It is not a comprehensive analysis of the role of TRECVID in careers overall but an insight into the effect that involvement in international benchmarking events may have on the citation impact of publications. A search was done in POP for the author name and TRECVID as a search term and then a search was done for the author name to include all their publications (i.e. not using TRECVID as a search term but just searching under author). The results were then checked for duplicates and any incorrect authors with similar names which had been included in the search results. We then substracted the TRECVID publications and citations from the non-TRECVID publications and citations to give us set of documents by the TRECVID author but not directly based on their TRECVID work. The data is cumulative so we look at the relative influence of TRECVID papers on their citation scores as their career has progressed. We call these individuals tv1, tv2, tv3, tv4 and tv5 in ascending order of seniority. Figure 12 compares the cites per paper for TRECVID papers and non-TRECVID papers, among the five researchers over the five year period.

FIG 12: Number of cites to TRECVID papers vs. non-TRECVID paper among 5 different researchers in different stages of their career (most junior ‘tv1’ to most senior ‘tv5’).

In Figure 12, tv1 (most junior researcher) naturally has the least number of publications overall and tv5 (most senior researcher) has the highest number of publications overall, and the other three researchers (tv2, tv3 and tv4) are somewhere in between. However, across most of the five researchers, with the exception of tv2, we can see that TRECVID papers receive more citations per paper than their other papers. It is difficult to separate TRECVID publications and non-TRECVID publications completely as clearly all a given researcher's work is inter-related. If a paper isn't about TRECVID for one of these participants, is it likely to be informed by it and *vice versa* for his or her TRECVID papers. This does show that, in most cases, for participants in TRECVID their TRECVID-related work receives more citations per paper than their other work. We cannot, however, say this is definitely because they are about TRECVID as there could be a large number of other reasons for their impact but it does suggest that for most of these researchers their work on TRECVID has had a positive effect on their overall citation count. In terms of bibliometric measures, which are increasingly important in academic promotion and recruitment, this suggests that TRECVID-related work is generally a productive use of their research time.

Discussion and conclusions

TRECVID has resulted in a large number of 'spin off' or derived publications which have received a substantial number of citations in total with some of them being very highly cited. Research carried out at TRECVID has impacted on the field of video research through publication in high quality conference and journals venues and also through being cited by other researchers working in similar or related fields. We can see from the visualizations of TRECVID topics over time how new approaches have been developed through TRECVID. For those involved in TRECVID, their publications relating to the conference have made a significant contribution to their overall research impact. We cannot, of course, know what would have happened to these research ideas or researchers if TRECVID had not taken place, as this would require a control in which TRECVID had not happened.

What does this study tell us about bibliometrics and its reliability and validity as way of measuring scholarly impact? What does it tell us about what scholarly impact actually is? In terms of reliability it reinforces previous work already discussed about problems of coverage for computer science in the established bibliometric tools of Web of Science and Scopus. Publish or Perish, based on Google Scholar, has almost astonishingly better coverage. Despite this, due to the expert checking, we know that it missed some papers. Publish or Perish also has limitations to its 'ease of use' and functionality, particularly for large data sets, compared to its more established rivals. A more detailed paper on methodological issues, describing 'lessons learned' from our chosen methodology as a guide for future related studies, will be published and the data used in our study is available at <http://www.cdvp.dcu.ie/scholarly-impact/>.

In terms of the validity of bibliometrics in general, the main question is whether a high citation rate (quantitative) for a research paper actually tell us something about the quality of that research paper and, by extension, its authors and perhaps their department or institution? Our main contribution to this debate is the investigation into the relationship between TRECVID performance and citation performance. This data strongly suggests that 'success' at TRECVID does lead to 'success' in citation. Thus, one could argue, that citation counts do measure quality if we accept that research quality is about finding solutions to problems that work better than other solutions proposed so far. The set-up in TRECVID is, in one sense, a microcosm of science. In a very limited and finite world, researchers test hypotheses, or at least proposed approaches, against a data set. Some of these turn out to

work well and some do not. For TRECVID and, more importantly, the wider field of video retrieval these 'failures', once confirmed as 'falsified' hypotheses (or more accurately proposed approaches) to use Popper's (1959) terminology, will be important in shaping the research and development trends of the future. Thus in using bibliometrics to measure quality we need to be clear that progress may rely on some researchers not doing too well and coming up against dead ends. They may, during this time, not receive many citations but they may, nevertheless, still make an important contribution. Our understanding of the relationship between citation rates and quality, in terms of what scholarly impact actually means, should include an awareness of this.

Acknowledgements

This material is based upon work supported by Science Foundation Ireland under Grant No. 07/CE/I1147. Thanks to Julia Barrett of UCD Library and Shane McLoughlin of the UCD School of Information and Library Studies for invaluable assistance to this project.

References

- Asknes, D.W. (2006). Citation rates and perceptions of scientific contribution. *Journal of the American Society for Information Science and Technology*, 57 (2), 169-185.
- Bar-Ilan, J. (2009). Which h-index?-a comparison ofWos, Scopus and Google Scholar. *Scientometrics*, 74 (2), 257-271.
- Bornmann, L., & Daniel, H-D. (2008). What do citation counts measure? A review of studies on citing behaviour. *Journal of Documentation*, 64 (1), 45-80.
- Enser,P.G.B. (2008a). Visual Image Retrieval, *Annual Review of Information Science and Technology*, 42,3-42.
- Enser, P.G.B. (2008b).The evolution of visual image retrieval, *The Journal of Information Science*. 34 (4), 531-546.
- Freyne, J., Coyle, L., Smyth, B., Cunningham, P. (2010) Relative Status of Journal and Conference Publications in Computer Science. *Communications. ACM*. 53(11), 124-132.
- Harzing, A-W. (2010).Citation analysis across disciplines: The impact of different data sources and citation metrics. Retrieved January 2011, from http://www.harzing.com/data_metrics_comparison.htm
- Moed, H.F.,&Visser, M.S. (2007). Developing bibliometric indicators of research performance in computer science: an exploratory study. Research report to the council for physical sciences of the Netherlands organisation for scientific research (NWO). CWTS Report 2007-01. Retrieved January 2011, from [http://ict.nwo.nl/files.nsf/pages/NWOA_78NJ63/\\$file/CWTS_Computer_Science_Study.pdf](http://ict.nwo.nl/files.nsf/pages/NWOA_78NJ63/$file/CWTS_Computer_Science_Study.pdf)
- Popper. K.R. (1959).The logic of scientific discovery. London, UK:Hutchinson.

Price, D.D. (1976). A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science*, 27 (5), 292-306.

Risjbergen, C.J. van. (2004). [The Geometry of Information Retrieval](#). Cambridge, UK: Cambridge University Press.

Robertson, S. (2008). On the history of evaluation in IR. *Journal of Information Science*, 34 (4), 439-456.

Rowe, B.R., Wood, D.W., Link, A.N. and Simoni, D.A. (2010) Economic Impact Assessment of NIST's Text REtrieval Conference (TREC) Program. RTI International, Project Number 0211875. Retrieved January 2011, from <http://trec.nist.gov/pubs/2010.economic.impact.pdf>

Spärck-Jones, K. (1981). *Information Retrieval Experiment*. London, UK: Butterworths.

Sugimoto, C.R., & McCain, K. W. (2010). Visualising Changes over time: A history of information retrieval through the lens of descriptor tri-occurrence mapping. *Journal of Information Science*, 36(4), 481-493.

Wainer, J., de Oliveira, H.P., Anido, R. (2010). Patterns of bibliographic references in the ACM published papers. *Information Processing and Management*.
doi:10.1016/j.ipm.2010.07.002