Image Retrieval: Ideas, Influences, and Trends of the New Age - Addendum

Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z. Wang
The Pennsylvania State University

March 13, 2007

1 Publication Trends

We analyze recent publication trends in CBIR and annotation via two exercises, with Google Scholar as aid. The first of these is an analysis of which venues/journals have carried the most CBIR-related work and what the impact is, and which sub-topics generated the most publication count and impact in the last five years. The second one involves generating subtopic-wise time-series capturing trends in publication over the last eleven years.

We query on the phrase "image OR images OR picture OR pictures OR content-based OR indexing OR 'relevance feedback' OR annotation ", year 2000 onwards, for publications in the journals - IEEE T. Pattern Analysis and Machine Intelligence (PAMI), IEEE T. Image Processing (TIP), IEEE T. Circuits and Systems for Video Technology (CSVT), IEEE T. Multimedia (TOM), J. Machine Learning Research (JMLR), International J. Computer Vision (IJCV), Pattern Recognition Letters (PRL), and ACM Computing Surveys (SURV) and conferences - IEEE Computer Vision and Pattern Recognition (CVPR), International Conference on Computer Vision (ICCV), European Conference on Computer Vision (ECCV), IEEE International Conference on Image Processing (ICIP), ACM Multimedia (MM), ACM SIG Information Retrieval (IR), and ACM Human Factors in Computing Systems (CHI). Relevant papers among the top 100 results in each of these searches are used for the study. Google Scholar presents results roughly in decreasing order of citations (again, only rough approximations to the actual numbers). Limiting search to the top few papers translates to reporting statistics on work with noticeable impact. We gathered statistics on two parameters, (1) publishing venue/journal, and (2) sub-topics of interest. These trends are reported in terms of (a) number of papers, and (b) total number of citations. Plots of these scores are presented in Fig. 1 and Fig. 2. Note that the tabulation is not mutually exclusive (i.e. one paper can have contributions in multiple sub-topics such as 'Learning' and 'Region', and hence are counted under both headings), neither is it exhaustive or scientifically precise (Google's citation values may not be accurate). Nevertheless, these plots convey general trends in the relative impact

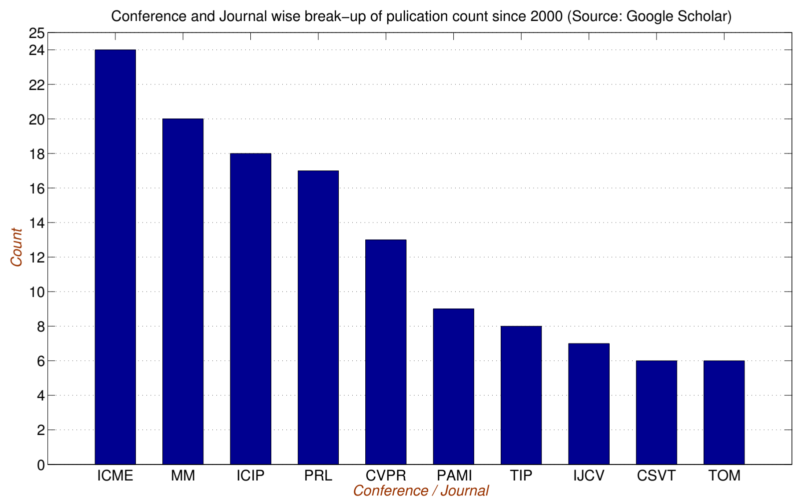

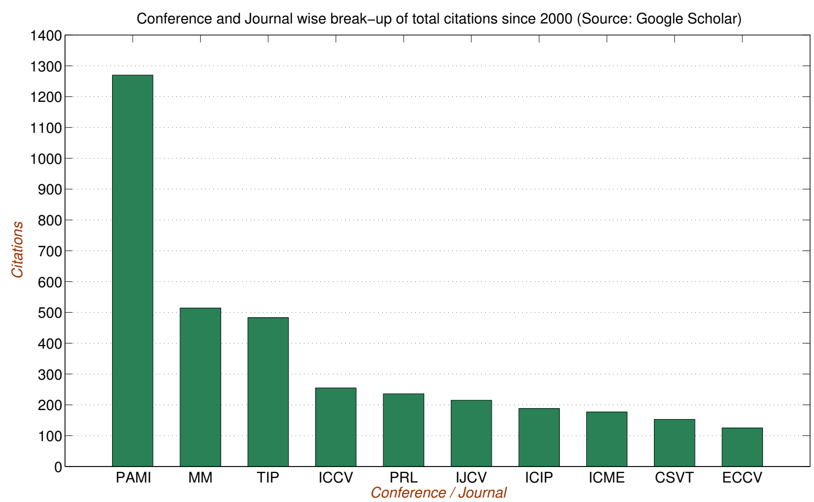

Figure 1: Conference-wise and journal-wise publication statistics on topics closely related to image retrieval, year 2000 onwards. *Top:* Publication counts. *Bottom:* Total citations.

of scholarly work. Readers are advised to use discretion when interpreting these results.

For the second experiment, we query Google Scholar for the phrase "image retrieval" for each year from 1995 to 2005, and note the publication count, say x. We then add a phrase corresponding to a CBIR-related technique, e.g., relevance feedback, and note the publication count again, say y. For each year

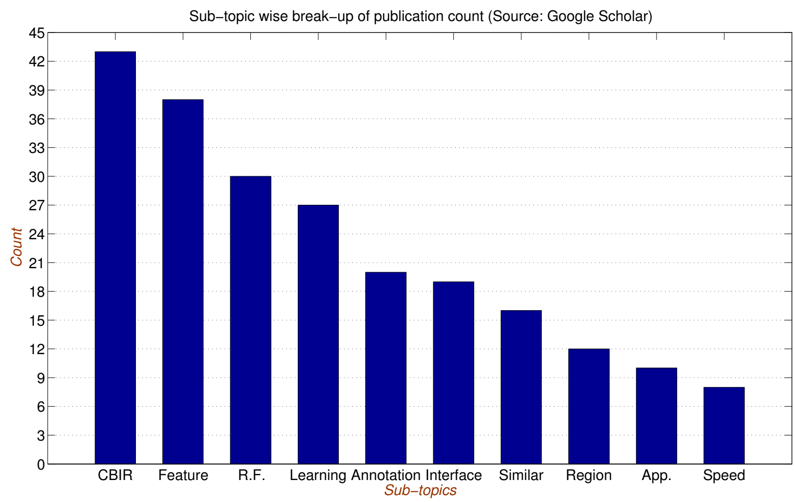

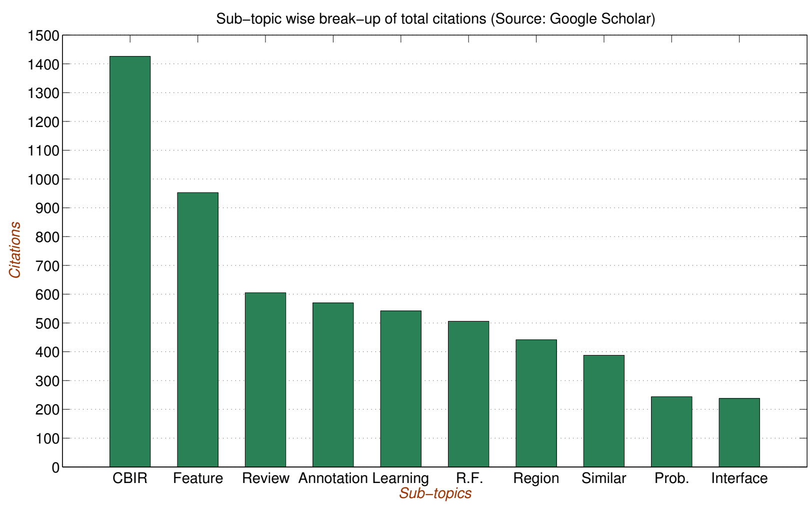

Figure 2: Publication statistics on sub-topics of image retrieval, 2000 onwards. *Top:* Publication Counts. *Bottom:* Total citations. *Abbreviations: Feature* - Feature Extraction, *R.F.* - Relevance Feedback, *Similar* - Image similarity measures, *Region* - Region based approaches, *App.* - Applications, *Prob.* - Probabilistic approaches, *Speed* - Speed and other performance enhancements.

and for each phrase, we take the ratio y/x representing the fraction of relevant publications. The time-series plot for eight such phrases, over the eleven years, can be seen in Fig. 3.

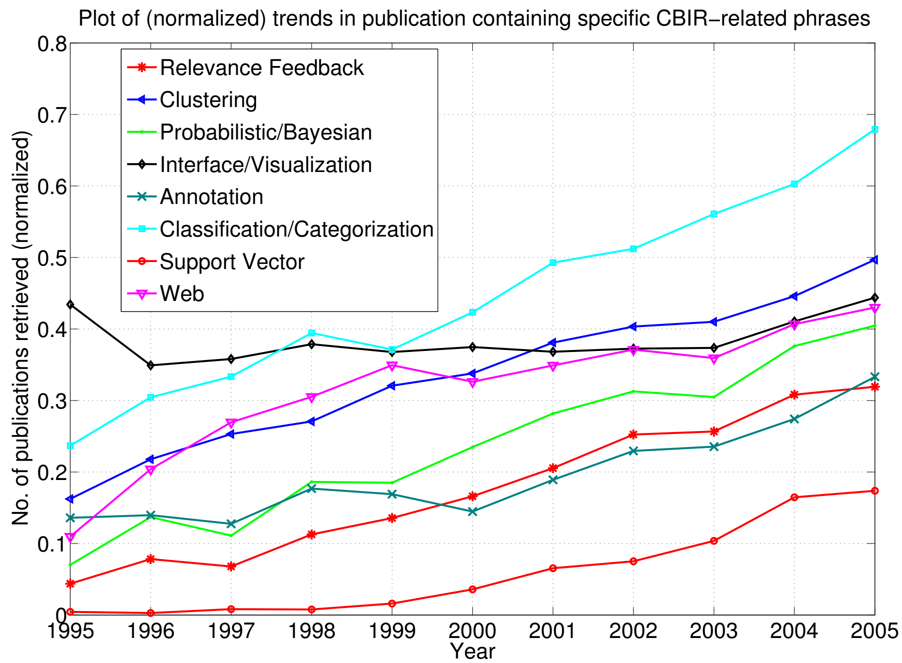

Figure 3: Normalized trends in publications containing "image retrieval" and corresponding phrases, as indexed by Google Scholar. Counts are normalized by the number of papers having "image retrieval" for the particular year.

2 Scientific Impact on Other Research Communities

The list of references in this paper is probably a good way to understand how diverse CBIR as a field is. There are at least 30 different well-known journals or proceedings where CBIR-related publications can be found, spanning at least eight different fields. In order to quantify this impact, we conduct a study. All the CBIR-related papers, cited in this work, are analyzed in the following manner. Let a set of CBIR-related fields be denoted as $\mathbf{F} = \{$*Multimedia (MM), Information Retrieval (IR), Digital Libraries/ World Wide Web (DL), Human-Computer Interaction (HCI), Language Processing (LN), Artificial Intelligence (including ML) (AI), Computer Vision (CV)*$\}$. Note the overlap among these fields, even though we treat them as distinct and non-overlapping for the sake of analysis. For each paper, we note what the core contribution is, including any new technique being introduced. For each such contribution, the core field it is associated with, $a \in \mathbf{F}$, is noted. For example, a paper that proposed a spectral clustering based technique for computing image similarity is counted under both

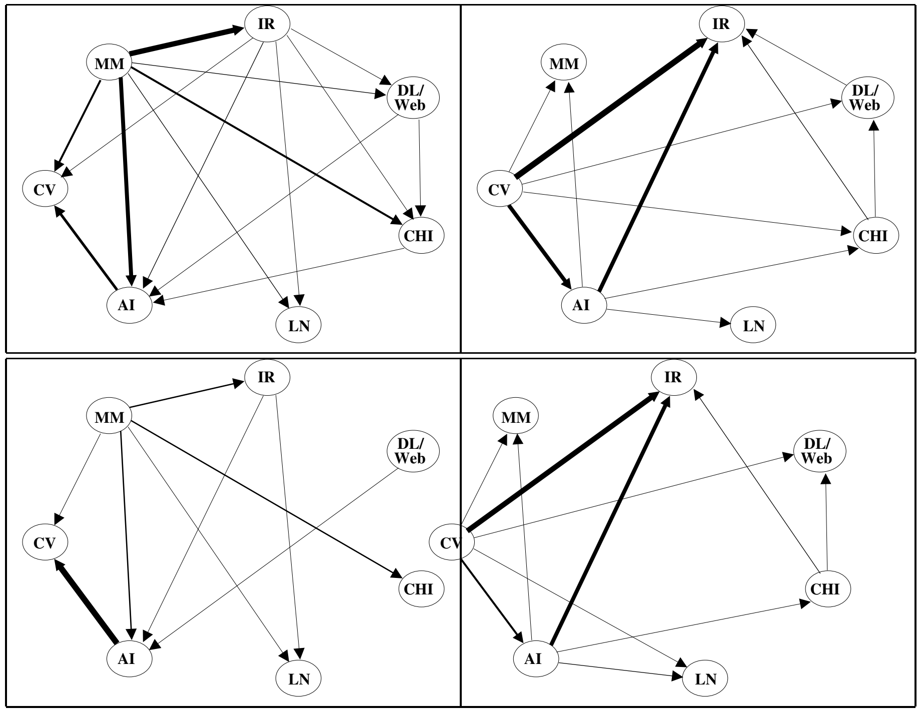

Figure 4: [Acronyms: MM := Multimedia, IR := Information Retrieval, DL := Digital Libraries/ World Wide Web, HCI := Human-Computer Interaction, LN := Language Processing, AI := Artificial Intelligence, and CV := Computer Vision]. Directed graphs representing inter-field impact induced by CBIR-related publications. An edge $a \rightarrow b$ implies publications at venue/journal concerning field b, having content concerning field a. We show oppositely directed edges between pairs of nodes, wherever significant, in the left and right graphs. *Top:* Edge thicknesses represent (relative) **publication count**. *Bottom:* Edge thicknesses represent (relative) **citations** as reported by Google Scholar.

CV and *AI*. Now, given the journal/venue where the paper was published, we note the field $b \in \mathbf{F}$ which it caters to, e.g., ACM SIGIR is counted under *IR* and ACM MIR Workshop is counted under both *IR* and *MM*. Over the 170 papers, we count the publication count and the Google Scholar citations for each $a \rightarrow b$ pair, $a \neq b$. The 7×7 matrices so formed ($|\mathbf{F}| = 7$) for count and citations are represented as directed graphs, as shown in Fig. 4. The thickness represents the publication or citation count, normalized by the maximum in the respective tables. Edges less than 5% of the maximum are not shown.

The basic idea behind constructing such graphs is to analyze how CBIR induces interests of one field of researchers in another field. A few trends are quite clear from the graphs. Most of the *MM*, *CV* and *AI* related work (i.e.

CBIR research whose content falls into these categories) has been published in *IR* venues and received high citations. At the same time, *AI* related work published in *CV* venues has generated considerable impact. We view this as a side-effect of CBIR research resulting in marriage of fields, communities, and ideas. But then again, there is little evidence of any mutual influence or benefits between the CV and CHI communities brought about by CBIR research.