# Measuring and Analyzing the Scholarly Impact of Experimental Evaluation Initiatives

Marco Angelini[1], Nicola Ferro[2], Birger Larsen[3], Henning Müller[4], Giuseppe Santucci[1], Gianmaria Silvello[2], and Theodora Tsikrika[5]

[1] "La Sapienza" University of Rome, Italy
`{angelini,santucci}@dis.uniroma1.it`
[2] University of Padua, Italy
`{ferro,silvello}@dei.unipd.it`
[3] Aalborg University, Denmark
`birger@hum.aau.dk`
[4] University of Applied Sciences Western Switzerland
`henning.mueller@hevs.ch`
[5] Centre for Research and Technology Hellas, Greece
`theodora.tsikrika@acm.org`

**Abstract.** Evaluation initiatives have been widely credited with contributing highly to the development and advancement of information access systems, by providing a sustainable platform for conducting the very demanding activity of comparable experimental evaluation in a large scale.

Measuring the impact of such benchmarking activities is crucial for assessing which of their aspects have been successful, which activities should be continued, enforced or suspended and which research paths should be further pursued in the future.

This work introduces a framework for modeling (or handling) the data produced by evaluation campaigns, a methodology for measuring their scholarly impact, and tools exploiting visual analytics to analyze the outcomes.

## 1 Motivations

Experimental evaluation is a fundamental methodology adopted in *Information Retrieval (IR)* since its inception, which substantially contributed to the scientific advancements of the field. It is based on the Cranfield methodology [3] which makes use of shared experimental collections in order to create comparable experiments and evaluate the performances of different information access systems.

Evaluation activities are very demanding both from the technical and economical point-of-views [4] and to be sustainable and scalable they have been carried out in large-scale evaluation campaigns such as *Text REtrieval Conference (TREC)* in the United States[6], the *Conference and Labs of the Evaluation*

---

[6] `http://trec.nist.gov/`

*Forum (CLEF)* in Europe[7], and the *NII Testbeds and Community for Information access Research (NTCIR)* in Asia[8]. In order to further facilitate their organization and management, each campaign is usually divided into *tracks* (referred to as *labs* in CLEF) and *tasks*. A lab is an area of focus concentrating on a specific evaluation aspect of a particular domain; for instance, CLEF in 2013 was organized into nine labs comprising, for instance, the "Cross Language Image Annotation and Retrieval" (ImageCLEF)[9] concentrating on the experimental evaluation of image classification and retrieval. Each lab may be divided into tasks, each focusing on specific sub-problems concerning the scope of the lab; as an example, ImageCLEF 2013 had four tasks comprising, among the others, the "Photo Annotation and Retrieval" task aimed at studying visual concept detection, annotation and retrieval in the context of diverse collections of photos.

Despite the general agreement about the importance of evaluation campaigns and the data produced by them [2,5], no general and shared methodology for measuring their scientific impact has already been defined. Such a methodology is much needed since measuring the impact of evaluation campaigns is crucial for assessing which of their aspects have been successful, and thus obtain guidance for the development of improved evaluation methodologies and information access systems.

It is our opinion that given that evaluation campaign contribution is mainly indicated by the research that would otherwise not have been possible, it is reasonable to consider that their success can be measured, to some extent, by the scholarly impact of the research they foster [7]. The goal of this work is to introduce the main aspects of a methodology allowing for modeling the experimental data and scientific production related to them, measuring the scholarly impact of evaluation campaigns and analyzing the outcomes by means of visual analytics techniques.

To this purpose, in Section 2 we present the bibliographical area of the *Distributed Information Retrieval Evaluation Campaign Tool (DIRECT)* system [1] which is a comprehensive system allowing for managing the experimental data, providing advanced services over them and defining explicit connections between campaigns and the data produced by them. In Section 3 we outline the three main steps to be followed for measuring the scholarly impact of evaluation campaigns. In Section 4 we show how the results of the impact analysis can be analyzed through visual analytics tools using the outcomes of the study conducted on CLEF campaigns as a use case. Finally, in Section 5 we draw some final remarks.

---

[7] http://www.clef-initiative.eu/

[8] http://research.nii.ac.jp/ntcir/index-en.html

[9] http://www.imageclef.org/

## 2 Modeling Experimental Data and Scientific Production

The necessity of modeling experimental data and designing a software infrastructure to manage and curate them, led to the development of a rather complex system – i.e. DIRECT – covering all the aspects of experimental evaluation. In this paper we focus on the bibliographical area of DIRECT which is responsible for retaining the relationships between the experimental data and the scientific production based on these data. Furthermore, this area models the bibliometrics (e.g. number of citations, h-index and impact factor) that are used to establish the impact of evaluation campaigns.

In Figure 1 we can see the conceptual schema of the bibliographic area. The central entity is `Contribution` which refers to a published piece of writing; a conference or a workshop paper, a journal article, a book, a technical report, a thesis or a manual are examples of contributions.
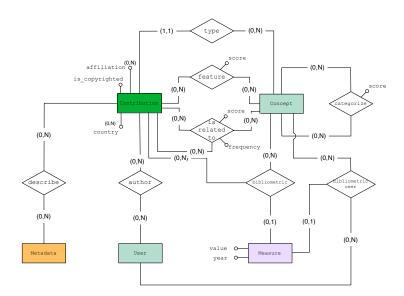


**Fig. 1.** Bibliographical Area relationships

In Figure 1 we can see that `Contribution` is associated to a `Concept` that defines its type; e.g. a `Contribution` can be a generic `Publication`, a `Working Note`, or a `Journal`. In general, `Concept` is defined as an idea or notion, a unit of thought; it is used to define the type of relationships in a semantic environment or to create a vocabulary (e.g. contribution types) and it resembles the idea of concept introduced by *Simple Knowledge Organization System (SKOS)* [8].

Furthermore, each `Contribution` is associated to no, one or more authors (i.e. `User`) via the `Author` relationship and can be described by no, one or more `Metadata` via the `describe` relationship. Similar relationships exist also between

`Contribution` and `Task`, `Track` (i.e. a lab in CLEF) and `Campaign` and allow us to explicitly relate contributions with the experimental data.

The relationship `feature` relates a `Contribution` to a `Concept` which defines its topic; this allows us to determine the topics of a `Contribution` and its relevance for a given topic. As a consequence, a `Contribution` can feature a `Concept` – e.g. "Digital Library" – and given that contributions are related to experimental data, we can conduct topic-oriented analyses on them; for instance, we can calculate how much a task or a campaign are related to the topic "Digital Library". Also the relationship `is related to` is relevant from the scholarly impact point-of-view, because it allows analyses based on the number of citations of a contribution. Indeed, we can say that "Contribution A `cites` Contribution B" where `cites` is a `Concept` relating "Contribution A" with "Contribution B".

Finally, the relationship `bibliometric` relates a `Contribution` to a `Concept` and a `Measure`. This allows us to say that "Contribution A has `impact factor 1.3`"; `impact factor` is defined as a `Concept` and `1.3` as the value of a `Measure`. The relationship `bibliometric user` has the same purpose but oriented to `User` (i.e. author); indeed, through this relationship we can express something like "User A has `h-index 3`", where `h-index` is a `Concept` and `3` is the value of a `Measure`.

## 3 Three Steps for Measuring the Scholarly Impact

Starting from the above described model we can conduct bibliometric studies providing a quantitative and qualitative indication of the scholarly impact of a research activity by examining the number of publications derived from it and the number of citations these publications receive. Such studies can be conducted by following these three main steps: (i) Publication data collection; (ii) Citation data collection; (iii) Data analysis.

So, the first step for assessing the scholarly impact of an evaluation campaign is to identify the publications associated with it and collect them in a dataset so that their citation data can then be obtained and analysed.

The second step involves the selection of citation data sources; the most comprehensive are: Thomson Reuters (formerly ISI) Web of Knowledge[10], Scopus[11] and Google Scholar[12]. Each of these sources follows a different data collection policy that affects both the publications covered and the number of citations found. Once the citation data sources have been selected, the next step is to query them using the publication data as input so as to obtain the citation data.

The last step regards the analyses that can be performed; they can be along several axes, such as the types of publications and the labs and tasks comprising the evaluation campaign while also drilling down the data into time dimension.

---

[10] http://wokinfo.com/
[11] http://www.scopus.com/
[12] http://scholar.google.com/

## 4 Analyzing the Results via Visual Analytics Tools

The three steps depicted above have been applied to the CLEF (2000-2009) Proceedings publications and to the CLEF (2000-2009) Working Notes publications and detailed results are described in [7,6]. For this study, the relationships between experimental data and contributions retained by DIRECT allowed us to calculate the measures determining the impact of evaluation initiatives; for instance, it emerged that three labs – i.e. Adhoc, ImageCLEF, and QA@CLEF – clearly dominate in terms of publication and citation numbers and thus have the higher scholarly impact.
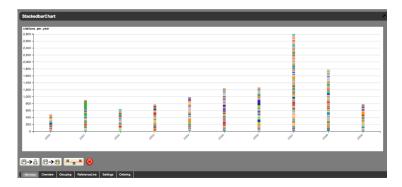


**Fig. 2.** A screen-shot of a part of the interactive visual environment for analysing the results of impact analysis of CLEF.

These conclusions have been drawn thanks to visual analytics tools offered by DIRECT. In Figure 2 we can see a screen-shot of a part of the visual environment developed for conducting impact analysis. This figure reports the stacked bar chart depicting the number of citations for the CLEF labs and tasks over the years (2000-2009). Each color in the bars represents the number of citations received by the tasks belonging to a specific campaign. This environment allows also for selecting specific tasks and comparing their measures, zooming and highlighting specific parts of the graphs and to compare citation numbers with other bibliometrics such as the h-index of authors and the impact factor of publication venues.

By using the analytics possibilities offered by the DIRECT visual environment it is also possible to identify some trends over all labs and tasks; for instance, in many cases there appears to be a peak in their second or third year of operation, followed by a decline [7]. Exceptions include the "Photo Annotation and Retrieval" task of ImageCLEF, which attracted significant interest in its fourth year when it employed a new collection and adopted new evaluation methodologies. Such novel aspects result in renewed interest in labs and tasks, and also appear to strengthen their impact.

## 5 Final Remarks

In this work we present a general framework for modeling the experimental data and their relationships with scientific contributions. This model settles the ground for calculating bibliometrics to be used for assessing the impact of evaluation activities.

We have also introduced the three main steps to be followed for measuring scholarly impact starting from scientific publications. Finally, we have shown how visual and interactive tools can be used for conducting impact analysis.

Future work will focus on the design and development of more advanced visualizations to interact with and explore the scholarly impact data, as well as improving the automation in gathering and cleaning of further bibliographic data in order to carry out deeper analyses.

## Acknowledgements

## References

1. M. Agosti, E. Di Buccio, N. Ferro, I. Masiero, S. Peruzzo, and G. Silvello. DIREC-Tions: Design and Specication of an IR Evaluation Infrastructure. In T. Catarci, P. Forner, D. Hiemstra, A. Peñas, and G. Santucci, editors, *Information Access Evaluation. Multilinguality, Multimodality, and Visual Analytics. Proceedings of the Third International Conference of the CLEF Initiative (CLEF 2012)*, pages 88–99. Lecture Notes in Computer Science (LNCS) 7488, Springer, Heidelberg, Germany, 2012.
2. M. Agosti, G. M. Di Nunzio, and N. Ferro. The Importance of Scientific Data Curation for Evaluation Campaigns. In C. Thanos, F. Borri, and L. Candela, editors, *Digital Libraries: Research and Development. First International DELOS Conference. Revised Selected Papers*, pages 157–166. Lecture Notes in Computer Science (LNCS) 4877, Springer, Heidelberg, Germany, 2007.
3. C. W. Cleverdon. The Cranfield Tests on Index Languages Devices. In K. Spärck Jones and P. Willett, editors, *Readings in Information Retrieval*, pages 47–60. Morgan Kaufmann Publisher, Inc., San Francisco, CA, USA, 1997.
4. B. R. Rowe, D. W. Wood, A. L. Link, and D. A. Simoni. *Economic Impact Assessment of NIST's Text REtrieval Conference (TREC) Program*. RTI Project Number 0211875, RTI International, USA. `http://trec.nist.gov/pubs/2010.economic.impact.pdf`, July 2010.
5. C. V. Thornley, A. C. Johnson, A. F. Smeaton, and H. Lee. The Scholarly Impact of TRECVid (2003–2009). *Journal of the American Society for Information Science and Technology (JASIST)*, 62(4):613–627, April 2011.
6. T. Tsikrika, B. Larsen, G. Bordea, and P. Buitelaar. Deliverable D6.4 – Report on the impact analysis for the CLEF initiative. PROMISE Network of Excellence, EU 7FP, Contract N. 258191. `http://www.promise-noe.eu/documents/10156/9d42701f-7d2f-4450-b6a8-5dec5444a757`, August 2013.

7. T. Tsikrika, B. Larsen, H. Müller, S. Endrullis, and E. Rahm. The Scholarly Impact of CLEF (2000-2009). In P. Forner, H. Müller, R. Paredes, P. Rosso, and B. Stein, editors, *Proc. of Information Access Evaluation. Multilinguality, Multimodality, and Visualization - 4th International Conference of the CLEF Initiative, CLEF 2013*, volume 8138 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2013.
8. W3C. SKOS Simple Knowledge Organization System Reference – W3C Recommendation 18 August 2009. `http://www.w3.org/TR/skos-reference`, August 2009.