Bibliometrics

Werner Nutt

Bibliometrics

- Are there objective (= quantitative) ways to measure
  - the productivity of a researcher?
  - the quality of journals?
  - the productivity of institutions?
- Ideas:
  - Researchers write papers
    - ➔ Analyse the paper output
  - Researchers build upon the work of other researchers, which they cite
    - ➔ Analyse how often papers have been cited

Bibliometric indices

N_p: total number of papers published

N_c: total number of citations received

n_c: mean number of citations per paper (= N_c/N_p)

IF: impact factor, calculated for journals (E. Garfield, 1955)

h-index: Hirsch's h-index (Jorge Hirsch, 2005)

g-index: Response to h-index (Leo Egghe, 2006)

*i*10-index: number of papers with ≥ 10 citations (Google Scholar)

Impact factor of a journal

Published yearly in the *Journal Citation Reports (JCR)*

- by the ISI (= Institute for Scientific Information),
  founded by E. Garfield,
  now part of the *Science and Scholarly Research* division of Thomson-Reuters
- based on two databases with scientific articles, one for science and another one for humanities

Definition of IF_n (= impact factor for year *n*)

- IF_n = average number of citations in year *n*
  for articles published in years *n-1* and *n-2*
- IF_n is published in year *n+1*

Impact factors of some journals in 2013/14

Nature:	42.3/41.5
Science:	31.5/33.6

JACM:	2.939/1.394
ACM TODS	0.750/0.684
VLDB Journal	1.701/1.568
IEEE TOSE	2.292/1.614
TCS	0.516/0.657

JACM	= Journal of the Association of Computing Machinery
ACM TODS	= ACM Transactions on Database Systems
IEEE TOSE	= IEEE Transactions on Software Engineering
TCS	= Theoretical Computer Science

What the IF is supposed to tell us

According to ISI:

- Helps librarians and researchers to decide which journals to subscribe to
- Provides sales arguments for journal publishers
- Supports academic evaluation
  (of researchers, institutions, etc.)

Criticism of IF

- IF can be increased by playing games
  - invite senior researchers
  - reduce types of articles that attract fewer citations
  - publish articles likely to attract citations early in the year

  …
- IFs differ across disciplines
- Articles in a journal with high IF are not necessarily frequently cited
  - Distribution of citations over articles follows a power law
    e.g., 90% of citations in Nature come from 25% of the papers
  - Correlation between paper citation frequency and impact factor is diminishing, since papers are available electronically

    See: Lozano, Larivière, and Gingras. The weakening relationship between the impact factor and papers' citations in the digital age. Journal of the American Society for Information Science and Technology

More criticism of the IF (Lozano et al.)

1) some types of publications within journals, such as letters and commentaries, are used to count citations (the nominator), but do not themselves count as "papers" (the denominator), and hence inflate the journal's IF
2) the IF depends on the number of references, which differs among disciplines and journals
3) the inclusion of journals in the database depends solely on Thomson Reuters, a private company, and not on the fields' practitioners,
4) the exact IF published by Thomson Reuters cannot be replicated using publicly available data,
5) the distribution of citations/paper is not normal, so at the very least the mode or median ought to be used instead of the mean,
6) the 2-year span for papers followed by one year for citations is completely arbitrary and favours high-turnover over long-lasting contributions,
7) journal editors can manipulate and artificially inflate their IFs

More criticism of the IF (Lozano et al.) /2

- It does not make sense to use the IF of a journal as a proxy for paper quality, since today the number of citations of a paper can accessed itself
- Even more troubling, is the 3-step approach of using the IF to infer journal quality, extend it to the papers therein, and then use it to evaluate researchers.

The h-index

Suggested in 2005 by Jorge E. Hirsch, a physicist at UCSD, as a single index tool for determining a researcher's productivity

Definition: For a multiset (bag) of natural numbers M, we say that

$h(M) := \max \{ n \mid \text{there are } n \text{ elements } c \in M \text{ with } c \geq n \}$

Examples:

$M_1 = \{\{100, 99, 80, 8, 4, 4, 1\}\} \rightarrow h(M_1) = 4$

$M_2 = \{\{5, 5, 5, 5, 4, 4\}\} \rightarrow h(M_2) = 4$

In bibliometrics, the multiset M of an author a is the collection of citation counts of a's papers

The h-index: graphically

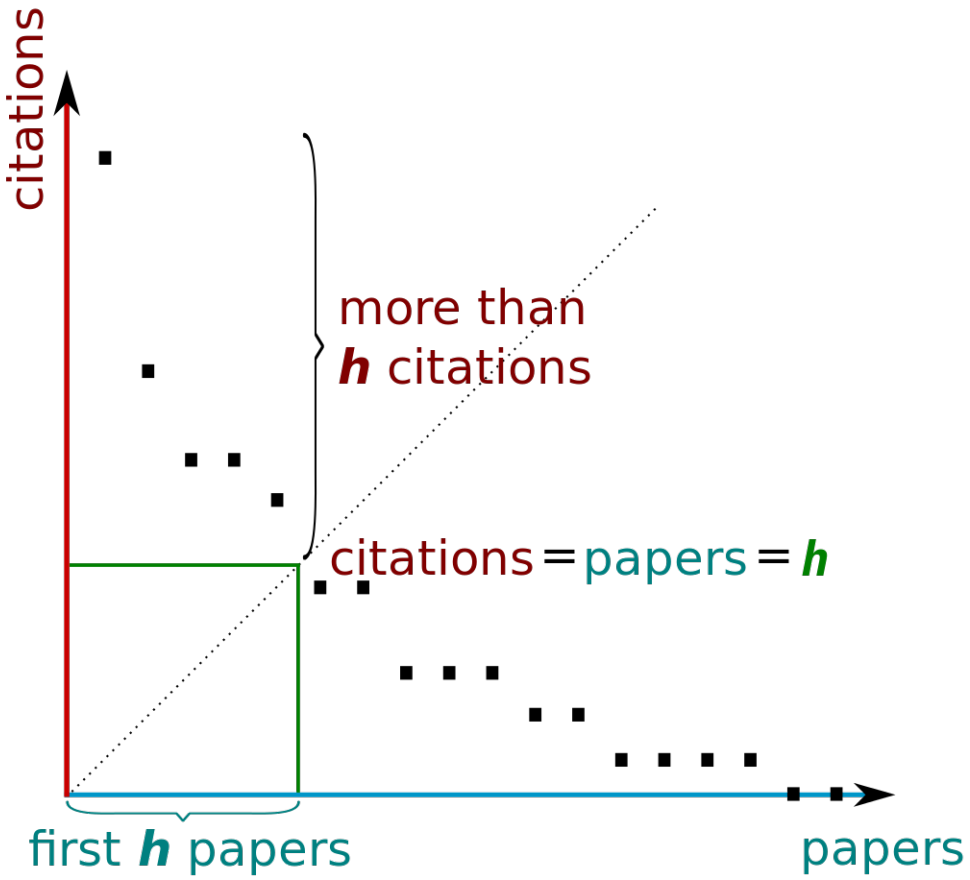

Source: Wikipedia

The h-index: graphically (2)

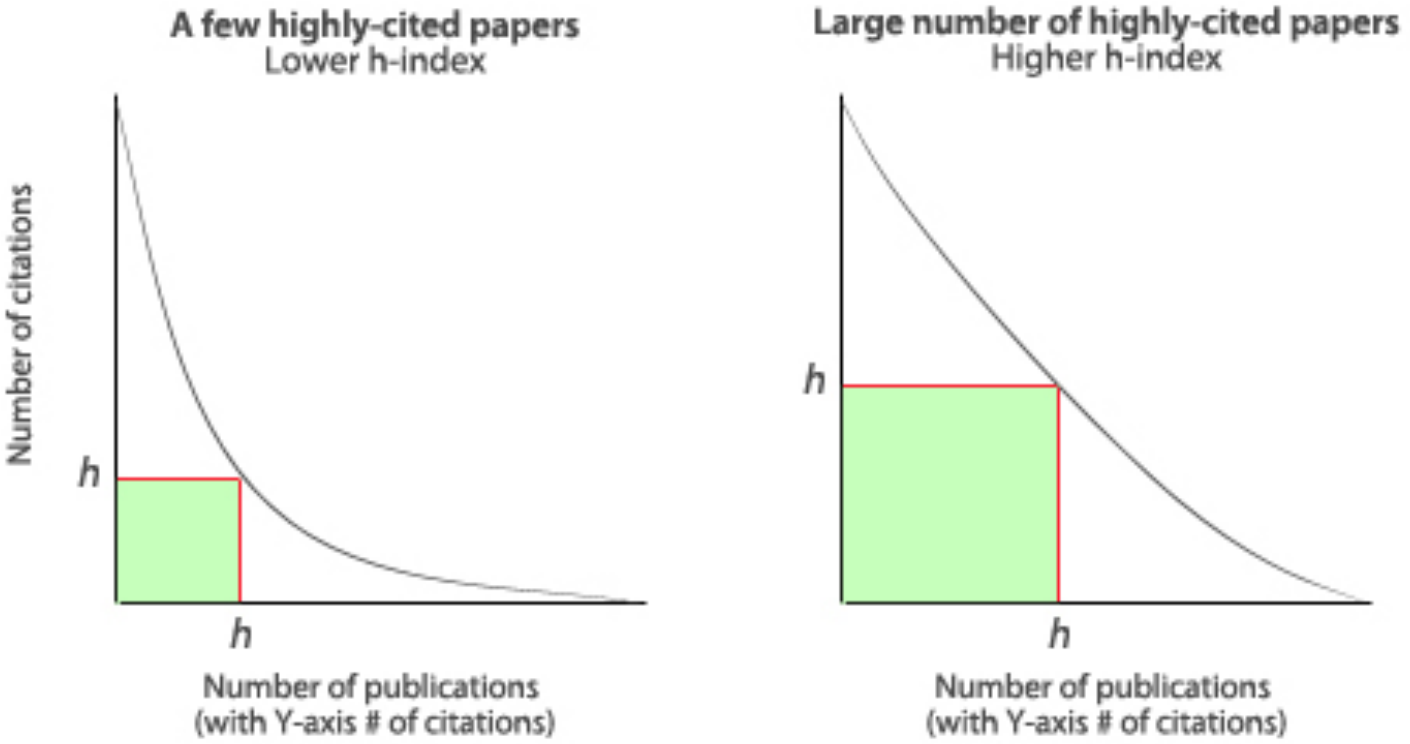

Source: http://www.benchfly.com/blog/h-index-what-it-is-and-how-to-find-yours/

Growth of the h-index

If a researcher publishes at a steady rate, and produces papers of similar quality, then

- each paper collects a constant number, say c, of citations per year
- then, the total number of citations grows as y^2 over the number of years y

➔ the h-index grows linearly with y, i.e., $h \sim m\, y$

- the coefficient m depends
  - on the researcher
  - on the discipline

Advantages of the h-index

- It relies on citations to papers themselves, not the journals
- It is not dramatically skewed by a single well-cited, influential paper (unlike total number of citations would be)
- It is not increased by a large number of poorly cited papers (unlike total number of papers would be)
- It minimizes the politics of publication. A high-impact paper counts regardless of where it was published
- It's good for comparing scientists within a field at similar stages in their careers
- It may be used to compare not just individuals, but also departments, programs or any other group of scientists.

Cited from: http://www.benchfly.com/blog/h-index-what-it-is-and-how-to-find-yours/
Blog by Alan Marnett

Criticism of the h-index

The h-index

- favours papers with many authors
- discards information in author placement on author lists
- does not take into account the context of a citation

(e.g., favorable vs. critical,
fleshing out an introduction vs. result or method
enabling work in current paper)

h-index: Summary

"In summary, I have proposed an easily computable index, h, which gives an estimate of the importance, significance, and broad impact of a scientist's cumulative research contributions. I suggest that this index may provide a useful yardstick with which to compare, in an unbiased way, different individuals competing for the same resource when an important evaluation criterion is scientific achievement."

Hirsch, J. E. (15 November 2005). "An index to quantify an individual's scientific research output". PNAS 102 (46): 16569–16572

The g-index

Idea: Give more credit to researchers with a few (or just one) landmark paper.

Definition: For a multiset (bag) of natural numbers M, we say that

g(M) is the largest number n such that

there are *n* numbers $c_1,...,c_n \in M$ with

$(c_1+...+c_n)/n \geq n$

In other words, the average of the n numbers $c_1,\ldots,c_n$ is at least n.

Equivalently, we can require that n is the largest number such that we find counts $c_1,\ldots,c_n$ with $c_1+...+c_n \geq n^2$

Egghe, Leo (2006) Theory and practise of the g-index, Scientometrics, vol. 69, No 1, pp. 131–152

The g-index: graphically

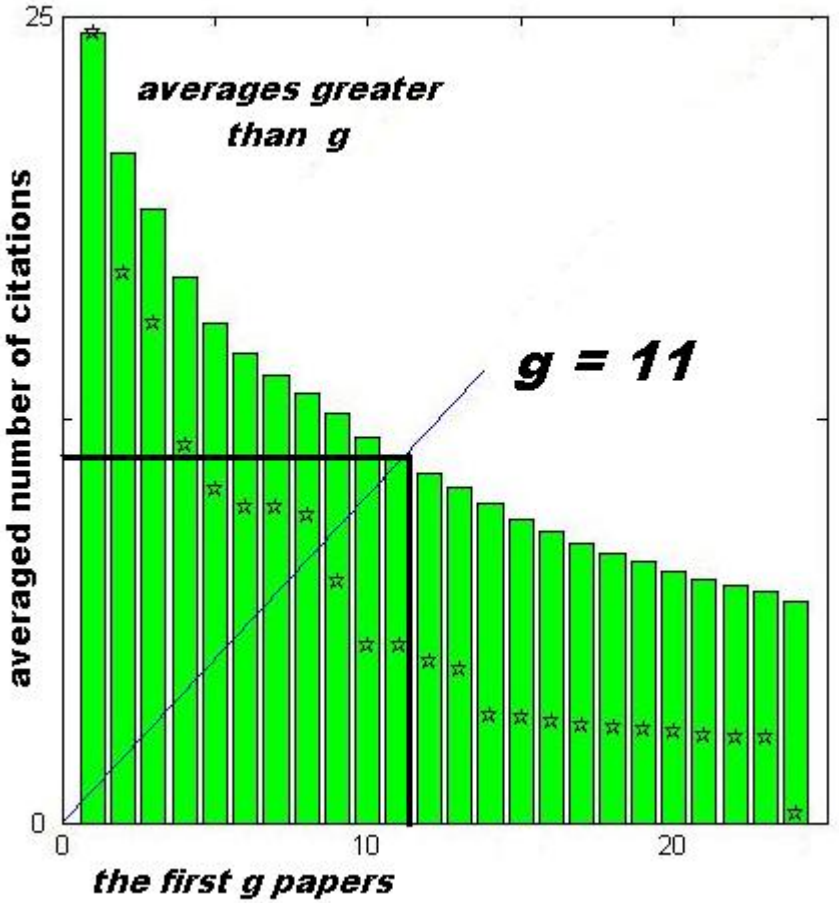

An example of a g-index (the raw citation data, plotted with stars, allows the h-index to also be extracted for comparison).

Source: Wikipedia

Predictive power

Hirsch J. E. (2007). “Does the h-index have predictive power?”.
PNAS 104 (49): 19193–19198. Also: http://arxiv.org/abs/0708.0646

Compared four indices: N_p, N_c, n_c, h

Test on a sample of 50 physicists, each with a 24 years career

- divided career in 2 halves of 12 years
- computed the indices for the first half
- computed the indices for
  - all publications in the entire career (cumulative future performance)
  - only publications in 2nd half of career (exclusive future performance)
- found out:
  - h-index best at predicting itself and other indices (correlation of .91 and .89, resp., for self-prediction)

Hirsch: h-index favors productive authors

For papers with several authors:

- less prolific and junior authors benefit less from the number of citations to a paper

Assume, paper p has N_p citations.

It only helps authors to increase their h, if $h < N_p$

➔ if it helps a productive author (big H),
then its value is H for that author

➔ if it helps a less productive author (little h),
then its value is h for that author

However, other indexes (N_c, g) give the same credit for all authors of a paper

In reality, there is not “the” h-index

The h-index is always computed over some database:

- Google Scholar
- Scopus (http://scopus.com by Elsevier)
- ISI – Web of Knowledge
  (Microsoft Academic Search)

Databases differ wrt to the publications they cover …

Variants of the h-index:

- over the last x years
- w/ or w/o self-citations
- for institutions

Challenges in creating a citation db: …

How one can dupe Google Scholar …

- Invent a new author, e.g., Marco Alberto Pantani-Contador
- Create 6 faked documents by that author, with lots and lots of citations
- Upload the documents on a university website
- Wait until Google Scholar detects them

Outcome: increase of 774 citations in 129 papers, 6 per cited paper

Google Scholar never corrected the fake citations.

Emilio Delgado Lopez-Cozar, Nicolas Robinson-Garcia, Daniel Torres-Salinas. Manipulating Google Scholar Citations and Google Scholar Metrics: simple, easy and temptinghttp://arxiv.org/abs/1212.0638

Ranking of Conferences and Journals

CORE (= Computing Research and Education Association of Australasia) is an association of university departments of CS in Australia and New Zealand

- ranks CS conferences
- ranks CS journals

Based on introduction and reranking requests (with detailed arguments) by researchers

- decision about change are also based on Google Scholar and ArnetMiner

See CORE conference ranking
(http://www.core.edu.au/index.php/conference-rankings)