

Recommended metrics

August 2020
bibliometrics@ucl.ac.uk

The [UCL bibliometrics policy](#) sets out some principles for the use of citation metrics in research assessment at UCL. As part of [the overall guidance](#), this paper sets out some recommendations for appropriate metrics to use.

While no bibliometric indicator is perfect, some are better and more robust than others. Where possible, we always recommend using normalised indicators – these take account of the differences in citation due to the subject area, the age of the paper, and the type of publication, to ensure that citations can be more meaningfully compared. This is a much better approach than simply counting total citations.

Three widely accepted normalised indicators are **category-normalised citations**, the **percentile rank** of papers, and the **share of papers in the top 1%/10% of citations**. All three can be obtained from InCites, and we would recommend using a mix of them to get an overall picture of citation impact – they all measure slightly different things.

If it is not practical to get any of these indicators, and you are confident that all publications are in the same subject area and otherwise comparable, then a quick method of normalisation is "citations per year" – but do be careful that the papers are all directly comparable, and that you are not, for example, comparing research articles to reviews or conference papers.

Category-normalised citations

Category-normalised citation impact (CNCI) calculates the average number of citations to all papers of the same type (articles, reviews, conference papers, etc), in the same publication year, and in the same subject area ("category"). It then divides the number of citations to the target article by this average.

Papers that are multi-disciplinary and assigned to more than one subject area are given a single value, which is an average of their CNCI for each category.

The key weakness of the CNCI is that the "category" is calculated by assuming all papers in a given journal belong to a specific subject area (or multiple different subject areas). This means that it may give confusing results for papers in journals that publish a wide range of topics. In multidisciplinary megajournals, such as *PLOS One* or *Scientific Reports*, it may not be meaningful at all.¹ Further, because normalizations are calculated at the level of the journal, they are usually only useful at that level or higher. Normalisations by discipline do not usually account for variation in approach. E.g., even within a niche research topic, an experimental researcher may have very different publishing and citation practices than a theoretical researcher.

The CNCI is scaled so that the world average is "1", and a CNCI of "1.5" would thus mean that this paper had 1.5x the global average citation for comparable papers. You may find it more helpful to benchmark against the UK average (across all subject areas, approximately 1.4) or the UCL average (across all subject areas, approximately 1.9).²

Percentile rank

The percentile rank calculates where a paper sits, on the basis of its absolute citation count, compared to all others of the same type (articles, reviews, conference papers, etc), in the same

¹ InCites attempts to reassign papers in some multidisciplinary journals, such as *Nature* and *Science*, to more appropriate categories based on the other papers they cite. However, this process is not perfect and does not work for all papers.

² All benchmark figures are calculated for research articles only, published 2013-2018, affiliated to any UK institution or just to UCL, using the November 2019 InCites data.

publication year, and in the same subject area ("category"). A paper with zero citations would be in the hundredth percentile; an average paper would be in the fiftieth percentile, and so on.

As with CNCI, "category" is determined by the journal it was published in, and may give odd results for multidisciplinary research, or papers in multidisciplinary journals. Papers that are assigned to two or more categories are given the best (lowest) percentile rank rather than averaging them.

This is expressed as a whole number, so a value of 1 indicates that a paper is in the first percentile, or in the top 1%. A value of 1.1 indicates that it is just outside the top 1%, and a value of 0.9 indicates that it is just inside it. 0.1 indicates the top 0.1%, and so on.

A value of 100 usually indicates no citations at all, while (due to rounding) an incredibly highly ranked paper may be shown with a value of zero.

The percentile rank applies to single papers; for a group of papers, this is shown as an "average percentile", calculated using a weighted average of all items.³ Again, the lower the value the better. The way this average is calculated tends to dilute the effect of outliers, making it useful to see the overall impact of a set of papers.

As a benchmark, the world average percentile for all papers is around 55 (rather than 50, in part due to cross-disciplinary effects); the UK average is around 48, and the average for UCL is around 41.

Share of papers in top 1%/10%

For large groups of papers, a useful addition to the average percentile rank or CNCI is to identify the share of papers which are among the most highly cited in their field.

These measures can show the amount of work which is very highly rated more effectively than just using an overall average value. For example, the CNCI would look the same for a group of papers which are all quite well cited, or for a group which has many average items and a few really standout ones. Having (for example) 20% of papers in the top 10% band would indicate that papers reach this high citation level about twice as often as the world average.

InCites offers pre-calculated data for papers in the top 10% or top 1% of all comparable papers (again, normalised by category, publication year, and publication type). Where a paper is classed into more than one subject category, it is given the best percentile value available – so a paper which is top 1% in cell biology but outside the top 1% in oncology would be counted as in the top 1% overall.

The global average across all disciplines is (of course) about 10% of papers in the top 10%, and 1% in the top 1%, though the effects of averaging multi-disciplinary papers can mean the exact figure varies a bit by subject area. The UK average is around 16% of papers in the top 10%, and 2.3% of papers in the top 1%, while for UCL the average is around 21% in the top 10%, and 3.9% of papers in the top 1%.

³ It uses the harmonic mean, a form of averaging which gives greater weight to lower values