

# Guidance on comparing metrics over a period of time

August 2020

[bibliometrics@ucl.ac.uk](mailto:bibliometrics@ucl.ac.uk)

## Overview

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 general advice on appropriate ways to interpret and compare citation metrics over a period of time.

One of the most common uses for metrics is for overall reporting and monitoring purposes – how many papers has a unit published over time? Can we tell if the rate of citations has changed?

Making these calculations can be complicated, as some simple approaches like counting the average or total number of citations per year can potentially be quite misleading. Older articles have had more time to gather citations, so some form of normalisation by year is essential.

## Appropriate metrics

The simplest approach is to compare the **number of publications** in each year. This can give a quick summary of relative activity rates. If there have been substantial changes in the size of the group being examined, it may be worth dividing this by the number of active researchers to avoid a misleading conclusion.

The standard **normalised citation impact metrics** discussed in other sections are useful here – it is possible to look at the yearly change in the **category normalised citation impact**, or the **percentage of highly-cited papers** produced by the group in each year.

The normalised metrics produced by InCites are recalculated every time the data is updated, and reflect all citations up to that point. If you found the CNCI for a 2010 paper in 2018, and check it again in 2019 and 2020, you would have three different values – it is likely that they will be reasonably consistent over time, and year-to-year changes will be small, but it does mean that if you are recording data as part of a regular monitoring report, you may need to update the values for earlier years as well as the most recent ones.

If you wish to use citation metrics that do not change over time, one approach is to use the **citations per paper** over a set time frame – usually three or five years since publication. For a five-year timeframe, a paper published in 2010 would count all citations in 2010-15 (year of publication + next 5 full years), a paper published in 2011 would count all citations in 2011-16, and so on. This method is somewhat more complicated to calculate but has the advantage of ensuring all papers are considered equally, and that the values will not change in the future. It has the disadvantage that it cannot be used for more recent papers as they do not have a sufficiently long citation history.

## Where can we find appropriate metrics?

Guidance on downloading the standard normalised metrics from InCites – category normalised citation impact, and percentile ranks – can be found in the specific sections for metrics [for individuals](#), metrics [for UCL departments](#), or metrics [for institutions as a whole](#).

To get **citations over a fixed timeframe**, eg the five years after publication, you can use either Web of Science or Scopus. Run a search to identify the relevant group of papers, and filter it down to a single year (eg “all papers published in 2013”). Then select the “create citation report” function (in Web of Science) or select all items and choose “view citation overview” (in Scopus).

Each of these will give a table for all papers in the search, listing the number of citations each has had in each year since publication. Count up the totals at the top of the columns, or from the graph, for the years you are interested in – the year of publication, plus X many subsequent years. This will allow you to calculate an average number of citations per paper over the period.

In Scopus, you may need to change the years displayed by using the drop-down entries at the top of the overview page. In Web of Science, only a few columns are displayed by default, so you may need to page through to find the relevant ones.

### Interpreting the metrics

The key normalised metrics are the Category Normalised Citation Impact, and the share of documents in the top 1% or 10%

- **Category Normalised Citation Impact (CNCI)**. This is a relative citation ratio, comparing the number of cites this paper has received to the mean of all papers in the same field, same document type, and same publication year. A CNCI of 1 is equal to the world average; the overall CNCI for most UCL publications is around 2, varying by subject area.
- **% of documents in top 1%/10%**. This gives the relative share of papers that exist in the top 1% or top 10% of comparable papers by citation count. The list of comparable papers is defined as those with the same field, same document type, and same publication year.

The CNCI and “top share”, taken together, can give a good sense of where a paper fits in comparison to the rest of the field. We recommend looking at both to help avoid anomalies – for example, a group of papers with a very high CNCI but a more normal-looking share of papers in the top 10% probably indicates that there is one or two very highly cited papers, and the rest are more in line with the field average.

Any normalised metrics are potentially quite volatile for the initial year or two after publication – citation numbers are very low overall, so a small number of citations can easily distort the result. For the first year, they are also strongly skewed by the timing of publication – a January-published paper will have had substantially more time to gather citations than a June-published one. Given these factors, we recommend that citations are not considered until at least one full calendar year after publication. As always, they should also be treated with caution in small fields or for small groups of papers.

The CNCI and percentiles are fully normalised metrics, where subject area, article type, and age are all normalised. However, “fixed-period” metrics, such as the citations over five years after publication, are only normalised by age. They are not normalised by subject area or article type and this means they need to be handled carefully.

These metrics are only meaningful if the papers are broadly comparable in each year – they are all in the same field and they do not, for example, change from being a mixture of reviews and articles to purely reviews. Otherwise, the normalisation will break down.