


altmetrics: a manifesto

NO ONE CAN READ EVERYTHING. We rely on filters to make sense of the scholarly literature, but the narrow, traditional filters are being swamped. However, the growth of new, online scholarly tools allows us to make new filters; these altmetrics reflect the broad, rapid impact of scholarship in this burgeoning ecosystem. We call for more tools and research based on altmetrics.

As the volume of academic literature explodes, scholars rely on filters to select the most relevant and significant sources from the rest. Unfortunately, scholarship's three main filters for importance are failing:



- Peer-review has served scholarship well, but is beginning to show its age. It is slow, encourages conventionality, and fails to hold reviewers accountable. Moreover, given that most papers are eventually published somewhere, peer-review fails to limit the volume of research.
- Citation counting measures are useful, but not sufficient. Metrics like the h-index are even slower than peer-review: a work's first citation **can take years**. Citation measures are narrow; influential work may remain uncited. These metrics are narrow; they neglect impact outside the academy, and also ignore the context and reasons for citation.
- The JIF, which measures journals' average citations per article, is often incorrectly used to assess the impact of individual articles. It's troubling that the exact details of the JIF are a **trade secret**, and that **significant gaming** is **relatively easy**.

Tomorrow's filters: altmetrics

In growing numbers, scholars are moving their everyday work to the web. Online reference managers [Zotero](#) and [Mendeley](#) each claim to store over 40 million articles (making them substantially larger than PubMed); [as many as a third of scholars are on Twitter](#), and a growing number tend scholarly blogs.

These new forms reflect and transmit scholarly impact: that dog-eared (but uncited) article that used to live on a shelf now lives in Mendeley, [CiteULike](#), or Zotero—where we can see and count it. That hallway conversation about a recent finding has moved to blogs and social networks—now, we can listen in. The local genomics dataset has moved to an online repository—now, we can track it. This diverse group of activities forms a composite trace of impact far richer than any available before. We call the elements of this trace altmetrics.

Altmetrics expand our view of what impact looks like, but also of what's making the impact. This matters because expressions of scholarship are becoming more diverse. Articles are increasingly joined by:

- The sharing of “raw science” like datasets, code, and experimental designs
- Semantic publishing or “nanopublication,” where the citeable unit is an argument or passage rather than entire article.
- Widespread self-publishing via blogging, microblogging, and comments or annotations on existing work.

Because altmetrics are themselves diverse, they're great for measuring impact in this diverse scholarly ecosystem. In fact, altmetrics will be essential to sift these new forms, since they're outside the scope of traditional filters. This diversity can also help in measuring the aggregate impact of the research enterprise itself.

Altmetrics are fast, using public APIs to gather data in days or weeks. They're open—not just the data, but the scripts and algorithms that collect and interpret it. Altmetrics look beyond counting and emphasize semantic content like usernames, timestamps, and tags. Altmetrics aren't citations, nor are they webometrics; although these latter approaches are related to altmetrics, they are relatively slow, unstructured, and closed.

How can altmetrics improve existing filters?

With altmetrics, we can crowdsource peer-review. Instead of waiting months for two opinions, an article's impact might be assessed by thousands of conversations and bookmarks in a week. In the short term, this is likely to supplement traditional peer-review, perhaps augmenting rapid review in journals like *PLoS ONE*, *BMC Research Notes*, or *BMJ Open*. In the future, greater participation and better systems for identifying expert contributors may allow peer review to be performed entirely from altmetrics. Unlike the JIF, altmetrics reflect the impact of the article itself, not its venue. Unlike citation metrics, altmetrics will track impact outside the academy, impact of influential but uncited work, and impact from sources that aren't peer-reviewed. Some have suggested altmetrics would be too easy to game; we argue the opposite. The JIF is [appallingly open to manipulation](#); mature altmetrics systems could be more robust, leveraging the diversity of altmetrics and statistical power of big data to algorithmically detect and correct for fraudulent activity. This approach already works for online advertisers, social news sites, Wikipedia, and search engines.



The speed of altmetrics presents the opportunity to create real-time recommendation and collaborative filtering systems: instead of subscribing to dozens of tables-of-contents, a researcher could get a feed of this week's most significant work in her field. This becomes especially powerful when combined with quick "alt-publications" like blogs or preprint servers, shrinking the communication cycle from years to weeks or days. Faster, broader impact metrics could also play a role in funding and promotion decisions.

Road map for altmetrics

Speculation regarding altmetrics ([Taraborelli, 2008](#); [Neylon and Wu, 2009](#); [Priem and Hemminger, 2010](#)) is beginning to yield to empirical investigation and working tools. [Priem and Costello \(2010\)](#) and [Groth and Gurney \(2010\)](#) find citation on Twitter and blogs respectively. [ReaderMeter](#) computes impact

indicators from readership in reference management systems. [Datacite](#) promotes metrics for datasets. Future work must continue along these lines.

Researchers must ask if altmetrics really reflect impact, or just empty buzz. Work should correlate between altmetrics and existing measures, predict citations from altmetrics, and compare altmetrics with expert evaluation. Application designers should continue to build systems to display altmetrics, develop methods to detect and repair gaming, and create metrics for use and [reuse](#) of data. Ultimately, our tools should use the rich semantic data from altmetrics to ask “how and why?” as well as “how many?”

Altmetrics are in their early stages; many questions are unanswered. But given the crisis facing existing filters and the rapid evolution of scholarly communication, the speed, richness, and breadth of altmetrics make them worth investing in.

[Jason Priem](#), University of North Carolina-Chapel Hill ([@jasonpriem](#))

[Dario Taraborelli](#), Wikimedia Foundation ([@readermeter](#))

[Paul Groth](#), VU University Amsterdam ([@pgroth](#))

[Cameron Neylon](#), Science and Technology Facilities Council ([@cameronneylon](#))

How to cite this:

J. Priem, D. Taraborelli, P. Groth, C. Neylon (2010), [Altmetrics: A manifesto](#), 26 October 2010. <http://altmetrics.org/manifesto>

Available at <http://altmetrics.org/manifesto/> (viewed February 16, 2021).

v 1.0 – October 26, 2010

v 1.01 – September 28, 2011: removed dash in alt-metrics

