Obtaining meaningful insights from publication metrics

Tomas James Rees
Director of Innovation, Oxford PharmaGenesis, Oxford, UK

doi: 10.56012/vjht8689

Correspondence to:
Tomas James Rees
tomas.rees@pharmagenesis.com

Abstract

Altmetrics and other article-level metrics offer new opportunities to understand the impact of medical publications and, indeed, clinical trial programmes. For example, we can learn whether the publication has been viewed, shared, engaged with, or cited on different platforms. These metrics have limitations, but new tools and techniques for aggregating and summarising different metrics are making it easier for publication planners to understand the impact of their publications.

So, great news! After weeks, or perhaps months, of effort and work by the authors, publication team, and the writer, your pivotal study has finally been published. Traditionally, this would pretty much mean the end of the story. You move on to the next project.

Nowadays, this is not where the story ends. In the Internet Age, we can gain insight into how the audience has interacted with the publication. We can find out if anyone read or talked about it and even if it has been used to inform clinical practice. This and other information about published articles is critically important for optimising publication planning.

For example, we can learn whether the publication has been communicated effectively. From the perspective of the full clinical programme, we can gather insight into whether the right studies were conducted, the right publications were developed, and which topics attract the greatest interest. We can also look at all of these in comparison to competitor publications. Fundamentally, we can learn whether the effort to develop the publication and publication plan was invested productively or if alternative strategies might be more effective.

Many publication teams still use the journal impact factor (JIF) as a measure of the success of their publications. However, as its name implies, the JIF is a journal-level metric that only provides a rough indication of the impact of individual publications. The JIF is an average score and is highly skewed because most publications receive fewer citations than the mean. Further, the JIF is unidimensional because it is based solely on citations. As a result, it does not take into account the many other forms of impact that can now be assessed.

Altmetrics

It is far better to look at the actual impact of each individual publication using article-level metrics than by using the JIF. Article-level metrics have been transformed in recent years with the development of the so-called "alternative metrics" (altmetrics), which provide an alternative to citations for measuring article-level impact. Altmetrics are a product of the internet, which has provided new avenues for interacting with journal articles. At the most basic level, publishers can track each time an article is viewed or downloaded. There is a lot more that goes on with a publication, however, which was previously entirely opaque. It is now possible to monitor a wide variety of news sources and to be alerted in real time whenever a news article discusses a journal article. Social media platforms make it possible to identify when individuals share or discuss publications. We can also track when individuals save a publication to their reference library or when an article is cited both in the peer-reviewed and in the grey literature, such as in blog posts, on open peer-review sites, and in governmental and non-governmental evidence syntheses, policy documents, and guidelines.

Article-level metrics have been transformed in recent years with the development of the so-called "alternative metrics" (altmetrics), which provide an alternative to citations for measuring article-level impact.

All of these altmetrics have limitations in their coverage. For example, different publishers track article views slightly differently, and they often do not share the data, making it difficult to compare article views across a wide range of publications. Also, social media engagement can only be tracked on platforms that allow machine access to the content, which means that platforms such as LinkedIn are excluded. Similarly, some providers of guidelines (e.g. National Comprehensive Cancer Network) do not allow automated systems to read their reference lists. Only one reference manager (Mendeley) provides anonymised usage data, so the data

from individuals using other reference managers are hidden. Altmetrics providers have to decide which news sources to monitor: do you go broad and inclusive (and potentially pick up a lot of noise) or stick to the more widely read outlets only? And none of these services can capture discussions about a publication when it is not named or referenced in some way.

The value of altmetrics comes from their ability to record different types of interactions with the publication. That means that altmetrics can provide greater insight than simple citation counts, despite limitations in their coverage. For example, the act of posting a link on social media is very different from citing an article in a peer-reviewed publication and is usually done for quite different reasons. People may save an article into their reference manager with the clear intention of citing it later, but it is often simply because they found the article interesting and want to bookmark it. News outlets tend to focus on newsworthy studies that could be of interest to the wider public and are less likely to pick up studies that are scientifically interesting but less immediately relevant to the public.

How to digest the vast amount of data available from altmetrics

Because they represent different actions, correlations between different metrics can vary. This can help us address another issue with altmetrics – how to digest the vast amount of data we now have access to. For example, the company Altmetric.com tracks 21 different metric sources, while their main rival, PlumX, tracks over 40 (although not all of these are relevant to journal articles).

The value of altmetrics comes from their ability to record different types of interactions with the publication.

To try to make sense of all these numbers, Altmetric.com collapses many of the metrics it captures into a single headline number – the Altmetric Attention Score (AAS). To minimise the problem of combining divergent metrics, the AAS simply excludes certain key article-level metrics (reference manager saves and citations in peer-reviewed publications), and to account for the fact that some metrics are more prevalent than others, they are weighted differently; a news article, for example, carries around 30 times the weight of a tweet. Even so, the AAS is almost entirely driven by mentions in news articles and tweets – it really is simply an "attention score".

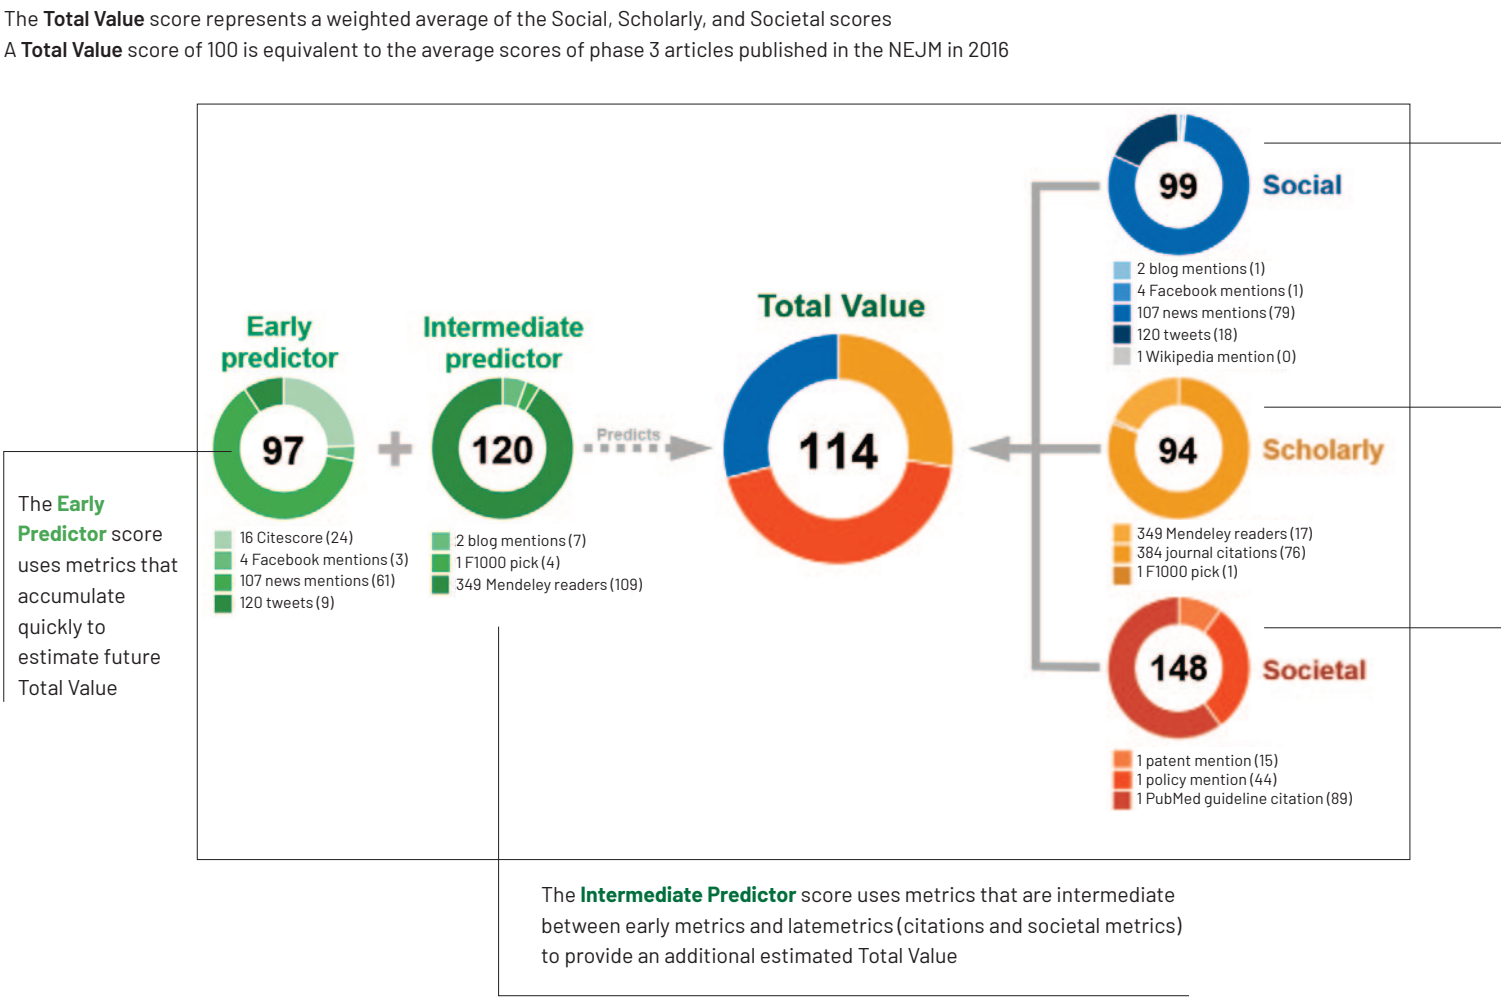

Figure 1. Example of the EMPIRE Index score for a single publication. HCP, healthcare provider; NEJM, New England Journal of Medicine

The EMPIRE Index

Although article-level metrics cannot be reduced to a single dimension, this general approach is helpful in making sense from publication metrics. Using a statistical technique (factor analysis), Avishek Pal and I explored the metrics of nearly 3000 publications of phase 3 studies. We found that, rather than trying to create a single score, the different metrics could be reliably placed in one of three groups, which we have named:

- Social Impact (Twitter, Facebook, news, blog mentions, and Wikipedia citations);
- Scholarly Impact (Mendeley saves, citations in peer-review publications, and citations in the Faculty Opinions service); and
- Societal Impact (citations in guidelines, policy documents and patents).

By appropriately weighting the metrics in these three separate scores, we were able to balance them so that the typical Social Impact would be similar to the typical Scholarly Impact and the typical Societal Impact. As a result, the scoring system, which we called the EMPIRE Index, captures the key different types of interaction that can be measured with article-level metrics in a way that makes it easy to compare across different articles (Figure 1). The EMPIRE Index is fully open for anyone to use and is described more fully in a publication in PLOS One.[1]

The EMPIRE Index can help understand which publications have greater or lesser impact and which can be used as a starting point for deeper insight; metrics can help you understand what and when, but they cannot tell you how or why. Knowing that your publication has been cited is good, but knowing the context in which it has been cited is even better. This applies just as much to social media as it does to citations in peer-reviewed publications. Fortunately, platforms that gather altmetrics also allow you to dive in and see the sources, where these are available.

> As artificial intelligence advances, automating insight-gathering from unstructured text sources will become easier.

The future: artificial intelligence and natural language processing

As artificial intelligence advances, automating insight-gathering from unstructured text sources

The **Social** score represents the impact of the article in public domains such as social media and news

Audience: HCPs , non-specialists, healthcare support staff, patients and other members of the public

The **Scholarly** score represents the impact of the article in academic domains such as journal citations and scholarly reference libraries

Audience: Specialists, experts, scientists, and academics

The **Societal** score represents the impact of the article in treatment guidelines, policy documents, and patents

Audience: Healthcare and policy decision makers, disease management bodies

will become easier. Sentiment analysis is a widely used approach that attempts to classify statements as either emotionally positive or negative. However, sentiment analysis is easily confused by medical discussions (with all their talk of death and pain).[2] Although the lexicons that define the emotional valency of different words can be tweaked, sentiment analysis is fundamentally not suited to understanding publication impact. At the end of the day, we are not really that interested in knowing whether people are happy or sad about the publication.

This is where more advanced natural language processing comes in. By understanding the language used when discussing a publication, we can get rapid insight into the readers' perspectives. We are already seeing the first examples of this in services such Scite – (https://scite.ai), which assesses whether a citation supports or contrasts with the original publication. This is a fast-moving area that has the potential to provide a second transformation in our ability to understand the true impact of publications we develop.

Disclosures and conflicts of interest

The author declares no conflicts of interest.

Data availability statement

For inquiries about data and other supplemental information, please contact the corresponding author.

References

1. Pal A, Rees T. Introducing the EMPIRE Index: a novel, value-based metric framework to measure the impact of medical publications. PLoS One. 2022;17(4):e0265381. doi:10.1371/journal.pone.0265381.
2. He L, Yin T, Zheng K. They may not work! An evaluation of eleven sentiment analysis tools on seven social media datasets. J Biomed Informat. 2022;132:104142. doi:10.1016/j.jbi.2022.104142

Author information

Tomas James Rees has over 20 years of experience in medical communications and publications planning and delivery. He has a longstanding interest in the transformation of scientific publications in the digital environment, including all aspects of open science, increasing reach and engagement across broad stakeholder audiences, and article metrics.