From reactive screening of published articles to proactive assignment of trial registrations

Author: Adam G. Dunn, PhD.[1,2]

Affiliations: 1: Centre for Health Informatics, Australian Institute of Health Innovation, Macquarie University, Sydney, NSW, Australia. 2: Computational Health Informatics Program, Boston Children's Hospital, Boston, MA, United States.

Conflicts of interest: None.

Funding: National Library of Medicine 1R01LM012976-01.

Abstract:

Part of the historical legacy of using published articles as the primary source of information for systematic reviews has been a focus on technologies for improving searching and screening in bibliographic databases. The increased availability of prospective trial registrations and structured results data provide an opportunity for entirely new approaches that may eventually circumvent the need for searching and screening entirely.

The problem

There are known issues associated with relying on evidence extracted from published articles. Articles are slow; around half of all registered trials remain unpublished two years after completion. Articles are incomplete; often missing or changing outcomes in ways that introduce biases. Finally, articles are reported in inconsistent and unstructured ways that make extraction time-consuming for humans and near-impossible for machines.

Despite the rate at which systematic reviews are being published, safety issues in approved drugs are identified too slowly. One in three approved drugs will have a safety issue identified and for half of those, the safety issue is identified more than four years after approval. For type 2 diabetes drugs, more than 23 systematic reviews are published each year but it takes 1.5 years for the results of a published drug trial to be included in a systematic review.

This suggests that much of the global systematic review effort is being spent on the answering the wrong questions (often repeatedly) and mostly using data that are old, incomplete, and biased.

Solutions from computer science

In practice, about 25% of the time spent on a systematic review comes from screening. For computer scientists, the screening task is a recognisable one: given a set of examples that are known to be included or excluded, the aim is to learn which text features distinguish them. Critically, these methods are optimised for perfect recall: to ensure that they do not exclude any article that a human would eventually want to include. There are now dozens of methods designed to help humans screen published trial articles for reviews and most approaches are trained and tested on just a handful of systematic reviews. Unfortunately, these approaches are chasing perfect recall in a far from perfect environment.

Trial registrations offer an alternative approach. One major advantage is the ability to monitor relevant trials as they are registered, completed, and reported. When linked to bibliographic databases and CrossRef, and supported by crowdsourcing, databases like trial2rev now store thousands of examples of systematic reviews linked to the trials they include. This has made it possible to consider other classes of algorithms like collaborative filtering, methods typically used to recommend movies and books to customers.

As a consequence, new methods are now being established for proactively recommending trials to systematic reviews for inclusion in updates as they are registered and completed. Helping us to predict in advance whether a systematic review update is likely to lead to a change in conclusion, these methods should also help reduce waste and target resources at reviews more likely to change policy and practice.