

Doing Things Twice (Or Differently): Strategies to Identify Studies for Targeted Validation

GOPAL P. SARMA  1*

1. School of Medicine, Emory University, Atlanta, GA USA

Abstract

The “reproducibility crisis” has been a highly visible source of scientific controversy and dispute. Here, I propose and review several avenues for identifying and prioritizing research studies for the purpose of targeted validation. Of the various proposals discussed, I identify scientific literature text mining as being a strategy that merits greater attention among those interested in reproducibility. I argue that the tremendous potential of scientific literature text mining for uncovering high-value research studies is a significant and rarely discussed benefit of the transition to a fully open-access publishing model.

I. INTRODUCTION

In recent years, significant attention has been given to problems with reproducibility in many areas of science. Some of these analyses have been theoretical in nature [1] while others have been focused efforts aimed at replicating large numbers of studies in a specific field [2,3].

The issue has become sufficiently high profile that it has been dubbed the “reproducibility crisis,” and is now a major topic of and debate in both the scientific [4–8] and popular press [9–13].

One thing is clear—we simply do not know what the “reproducibility distribution” looks like for the entirety of science. Taking this position as a starting point, how then do we identify and prioritize published results to investigate in greater detail? Should reproducibility initiatives be strictly local and originate from individual scientists themselves, or should there be more global, distributed efforts as well? In this brief note, I examine this and other questions and propose several strategies for identifying key results to be the focus of validation efforts.

*Email: gopal.sarma@emory.edu

II. UNCOVERING “LINCHIPIN” RESULTS

It goes without saying that all scientific results, even the true ones, are not created equally. In order to use resources efficiently, we would ideally identify “linchipin” results, that is, those studies which would carry the highest impact if we had greater certainty in their outcome. For instance, in cases where poorly conducted or fraudulent studies form the basis for guidelines or procedures in medicine, lengthy retractions can have significantly deleterious consequences for the public [14]. To use an ecological metaphor, these studies might be described as “keystone species” of the scientific ecosystem. How can we uncover such results?

i. Reproduction Versus Validation

Although the phrase “reproducibility crisis” has taken root in contemporary discussions, simply re-doing an experiment may not always be the most appropriate course of action. For instance, there might be linchipin theoretical results which simply need greater scrutiny or investigation with alternative methods. The same may be true of certain experimental results, where the highest value would be gained from re-thinking a given experimental design using alternative techniques.

Therefore, to broaden the scope of the discussion

Doing Things Twice (Or Differently)

to include all scientific results, not just experimental ones, as well as approaches other than simply repeating the original study under question, I will use the phrase “validation effort” rather than “reproducibility effort.”

ii. Have Individual Scientists Initiate Validation Efforts

A purely “local” approach would be for validation efforts to be initiated by investigators themselves. For instance, Schooler and several colleagues at UCLA arrived at an agreement whereby each researcher’s lab would attempt to replicate the results of the others prior to publication [15].

However, not all scientists are in a position to create such arrangements. Therefore, formal mechanisms for arranging validation efforts should be encouraged. As an example, this is the principle behind ScienceExchange’s Reproducibility Project¹, a marketplace for scientists to identify researchers from a network of laboratories to validate their research.

iii. Polling Scientists or Crowdsourcing

This approach would be more globally oriented and could be initiated by funding agencies, individual laboratories, or by “open science” projects. The strategy would be to distribute polls to scientists in different disciplines asking them what they believe to be high-value results. Like any poll, a number of practical issues will arise in arriving at reliable data. Questions will likely have to be written by scientists with sufficient experience in a given field to elicit reliable answers and to follow-up on ambiguities in responses. The questions will have to be framed appropriately. The administrators of such polls will have to be mindful of the fact that individuals could use a call for reproducibility as a political tactic for attacking a competitor’s research.

Nonetheless, polling scientists is likely a very straightforward strategy to uncover high-value results. The results of such polls could be used by funding agencies, independent foundations, or in-

dividual scientists themselves in deciding how to allocate resources for validation efforts.

iv. Let Larger Research Agendas be the Focal Points

This strategy would be to focus validation efforts around those results that form the foundation of major research agendas. For example, while soliciting proposals for new programs, funding agencies could ask scientists to submit a bibliography containing results upon which their proposed research relies. Or more directly, researchers could be asked to submit a separate document alongside their grant proposals suggesting experiments that merit additional investigation and which would advance their own research. Scientists could also publish such documents independent of grant applications, perhaps along the lines of a review article.

v. Scientific Literature Text Mining

Scientific literature text mining refers to the use of data analytic techniques to treat the scientific corpus itself as a massive data set for analysis [16]. Although the growth of data science has largely been driven by commercial applications in social media and business intelligence, we are now beginning to see the applications of data science to the scientific literature as well.

At the most basic level, article recommendation by Google Scholar or the many reference managers used by researchers is an example of data science applied to scientific papers. Other examples of contemporary research in scientific literature text mining include fraud detection (applying natural language processing to uncover linguistic signatures of fraudulent research) [17], characterizing the emergence of global scientific trends (using *n*-grams and patent citation networks to model the flow of ideas and technological development) [18], and resource allocation in the biomedical sciences (developing metrics which incorporate disease burden, research literature coverage, and clinical trial coverage to uncover underfunded areas of research) [19].

¹<http://validation.scienceexchange.com/>

One can easily imagine using techniques from machine learning and natural language processing to identify linchpin results, perhaps by examining citations networks, or using entity extraction to model the emergence of new terminology and concepts. The process would not need to be fully automated. We could employ a hybrid approach whereby data analytic techniques allow us to narrow down a corpus of tens of thousands of research papers to a few dozen or a hundred. Subsequently, with the guidance of experts and manual curation, we would arrive at a list of candidate results or studies to be the focus of targeted validation efforts.

III. CONCLUSION

Modern science is witnessing many growing pains, one of which is an increasing concern about the quality of research, as measured by reproducibility, in a number of distinct areas of inquiry. Although this concern has quickly grown to the point of being labeled a “crisis,” the reality is that we simply do not know what the distribution of reproducibility rates looks like for the entirety of science.

Nonetheless, this very uncertainty should be sufficient motivation to put into place procedures and incentives to increase the reliability of published results. There are many approaches to take in addressing this problem, a few of which have been outlined above.

Most of these ideas have been discussed or attempted in some form or another in recent years. The proposal that has received the least attention in the context of the reproducibility crisis is scientific literature text mining. One of the primary roadblocks to open-ended, exploratory data analysis with a large corpus of scientific papers is restrictions on their availability—in other words, closed access publishing models. Therefore, researchers should consider the applications of scientific literature text mining to uncovering linchpin results to be a key motivating factor in encouraging a complete transition to an open access publishing model.

ACKNOWLEDGMENTS

I would like to thank Adam Safron, P. Ravi Sarma, and Daniel Weissman for insightful discussions and feedback on the manuscript. A special thanks to A.S. for suggesting the phrase “keystone species.”

ORCID

Gopal P. Sarma  0000-0002-9413-6202

REFERENCES

- [1] J. P. A. Ioannidis, “Why Most Published Research Findings Are False,” *PLoS Med*, vol. 2, p. e124, 08 2005.
- [2] F. Prinz, T. Schlange, and K. Asadullah, “Believe it or not: how much can we rely on published data on potential drug targets?,” *Nature Reviews Drug Discovery*, vol. 10, no. 712, 2011.
- [3] C. G. Begley and L. M. Ellis, “Drug development: Raise standards for preclinical cancer research,” *Nature*, vol. 483, no. 7391, pp. 531–533, 2012.
- [4] W. Gunn, “Reproducibility: fraud is not the big problem,” *Nature*, vol. 505, no. 7484, pp. 483–483, 2014.
- [5] D. Adam and J. Knight, “Journals under pressure: Publish, and be damned...,” *Nature*, vol. 419, no. 6909, pp. 772–776, 2002.
- [6] E. Check and D. Cyranoski, “Korean scandal will have global fallout,” *Nature*, vol. 438, no. 7071, pp. 1056–1057, 2005.
- [7] R. Horton, “What’s medicine’s 5 sigma?,” *The Lancet*, vol. 385, no. 9976, 2015.
- [8] P. Campbell, ed., *Challenges in Irreproducible Research*, vol. 526, Nature Publishing Group, 2015.
- [9] Editors, “Trouble at the lab,” *The Economist*, 10 2013.
- [10] Neuroskeptic, “Reproducibility Crisis: The Plot Thickens,” *Discover Magazine*, 10 2015.

Doing Things Twice (Or Differently)

- [11] B. Carey, "Science, Now Under Scrutiny Itself," *The New York Times*, 7 2015.
- [12] K. M. Palmer, "Psychology is in a Crisis Over Whether It's in a Crisis," *Wired Magazine*, 3 2016.
- [13] J. S. Flier, "How to Keep Bad Science From Getting Into Print," *The Wall Street Journal*, 3 2016.
- [14] S. Bouri, M. J. Shun-Shin, G. D. Cole, J. Mayet, and D. P. Francis, "Meta-analysis of secure randomised controlled trials of β -blockade to prevent perioperative death in non-cardiac surgery," *Heart*, vol. 100, no. 6, pp. 456–464, 2014.
- [15] J. W. Schooler, "Metascience could rescue the replication crisis," *Nature*, vol. 515, p. 9, 2014.
- [16] G. P. Sarma, "Scientific Literature Text Mining and the Case for Open Access," *The Journal of Open Engineering*, 12 2017.
- [17] D. M. Markowitz and J. T. Hancock, "Linguistic Obfuscation in Fraudulent Science," *Journal of Language and Social Psychology*, p. 0261927X15614605, 2015.
- [18] R. V. Solée, S. Valverde, M. R. Casals, S. A. Kauffman, D. Farmer, and N. Eldredge, "The evolutionary ecology of technological innovations," *Complexity*, vol. 18, no. 4, pp. 15–27, 2013.
- [19] L. Yao, Y. Li, S. Ghosh, J. A. Evans, and A. Rzhetsky, "Health ROI as a measure of misalignment of biomedical needs and resources," *Nature biotechnology*, vol. 33, no. 8, pp. 807–811, 2015.