**Statcheck does not work: All the numbers. Reply to Nuijten et al. (2017)
Updated version, December 2021**

psyarxiv.com/hr6qy

Thomas Schmidt
University of Kaiserslautern, Germany
Experimental Psychology Unit
www.sowi.uni-kl.de/psychologie
thomas.schmidt@sowi.uni-kl.de

Abstract

Statcheck is an R algorithm designed to scan papers automatically for inconsistencies between test statistics and their associated p values (Nuijten et al., 2016). Due to concerns about its reliability and validity, Nuijten et al. (2017) check the output of the program against the sample of 49 papers manually checked by Wicherts et al. (2011), but they do not take into account that the program does only recognize 61 % of all statistical tests. Reconstructing the full table of hits, false alarms, misses, and correct rejections with respect to inconsistently reported tests, I show that Statcheck has poor sensitivity (.52) and poor validity ($\phi = .54$): In 1,120 tests, the program scored 29 hits, committed 19 false alarms, and missed an estimated 27 truly inconsistent tests, while 435 tests went unrecognized. If Statcheck flags a test, it is correct in only 60.4 % of all tests. If a test is truly inconsistent, Statcheck flags it in only 51.8 % of cases. Overall, only an estimated 5.00 % of all tests were actually inconsistent, and the program was clearly biased against flagging them.

Statcheck is an R algorithm designed to scan papers automatically for inconsistencies between test statistics, degrees of freedom, and p values (Nuijten, Hartgerink, van Assen, Epskamp, and Wicherts, 2016). The authors have used the algorithm to automatically scan thousands of scientific psychology papers and post the results on the PubPeer website whenever the program detects some discrepancy, thus flagging papers for reporting inconsistent test statistics. This campaign is strongly debated because the program's reliability and validity are still unclear, and because some PubPeer entries are demonstrably incorrect (including mine, unfortunately: Schmidt, Hauch, & Schmidt, 2015; Seydell-Greenwald & Schmidt, 2012). This prompted the German Psychological Society to criticize the publication of Statcheck results on PubPeer and to demand, without much success, the deletion of all incorrect entries in October 2016:

https://www.dgps.de/uploads/media/Stellungnahme_DGPs_statcheck_v04_eng_251016.pdf

In their new validity analysis of Statcheck, Nuijten et al. (2017) provide many but not all of the figures relevant for evaluating the program's performance in detail. They compare the program's output with the sample of 49 papers manually checked by Wicherts et al. (2011). Those papers contain a total of 1,120 statistical tests. The authors conclude that Statcheck has excellent sensitivity, specificity, and accuracy. In calculating those statistics, however, they do not take into account that Statcheck only recognizes a subset of all statistical tests reported. Because it is a very simple string detection algorithm looking only for reports in one specific APA format, it cannot deal with tests being reported summarily, in prose, in nonstandard order, or in tables. In the reported sample, Statcheck detected only 685 of 1,120 tests (61.2 %). Of course, discounting all the nondetected tests seriously distorts the analysis, especially when it comes to calculating sensitivity. Fortunately, it is not difficult to reconstruct the entire table of hits, false alarms, misses, and correct rejections, assuming that the rate of inconsistent results is the same in the detected and undetected test.

Table 1: The counts on which Nuijten et al.'s (2017) analysis is based. These include only the tests detected by both Statcheck and Wicherts et al. (2011).

	Flagged by Statcheck	Not flagged
Inconsistent	29 Hits (<i>H</i>)	5 Misses (<i>M</i>)
Consistent	19 False Alarms (<i>F</i>)	632 Correct Rejections (<i>CR</i>)

Table 1 gives Nuijten et al.'s (2017) results (the ones for Version 1.2.2 of the program with the one-tailed option toggled on). The 435 results that Statcheck overlooked are missing from this table, so that the rates of misses and correct rejections are underestimated. We can reconstruct the full table under the assumption that the prevalence of inconsistent tests is identical for detected and undetected tests (4.96 %). Rounded to full numbers, we can then estimate the full data (Table 2):

Table 2: Reconstructed counts including the tests not detected by Statcheck.

	Flagged by Statcheck	Not flagged
Inconsistent	29 Hits (<i>H</i>)	27 Misses (<i>M</i>)
Consistent	19 False Alarms (<i>F</i>)	1045 Correct Rejections (<i>CR</i>)

Several features of this table are noteworthy. First of all, inconsistently reported tests are surprisingly rare (5.00 %), giving the table a heavy second line. Second, tests are rarely flagged as inconsistent (4.29 %), giving the table a heavy second column. Third and more importantly, hits are only marginally more prevalent than misses or false alarms. Finally, the overwhelming majority of cases are correct rejections. Many of them (about 40 %) are a result of the program not detecting and thus not flagging the tests. Ironically, this gives Statcheck high accuracy because most tests are in fact consistent: Statcheck correctly classifies 95.89 % of all tests. However, this value is less impressive if we consider that not using Statcheck at all, and just taking all tests at face value, would already yield an accuracy of 95.36 %.

Now we can recalculate Nuijten et al.'s (2017) estimates of specificity and sensitivity. The authors define *specificity* as the proportion of unflagged tests that are classified correctly, or $CR / (CR + F)$. For the reconstructed counts in Table 2, this yields a specificity of 1049/1068, or 98.2 %. This is an excellent value that is just as meaningless as the total accuracy. It merely stems from the fact that the program rarely reports an event that rarely occurs.

More importantly, the authors calculate *sensitivity* as the proportion of truly inconsistent tests that is flagged by the program, or $H / (H + M)$. Because they ignore all the tests that the program failed to detect, they estimate this to be 29/34, or 85.3 %. This value is a gross overestimation. The more realistic value is calculated from the reconstructed table including the undetected test, and is 29/56, or 51.8 %.¹ Another meaningful statistic is the *positive predictive value*, defined as $H / (H + F)$, which is 29/48, or 60.4 %. Similarly poor is the value of the *phi coefficient*, a measure of association in 2x2 contingency tables. Its value is only .538. In summary: If Statcheck does flag a test, it is correct in 60.4 % of all tests. If a test is truly inconsistent, Statcheck flags it in 51.8 % of cases. These are very poor values for any diagnostic procedure, especially one that tries to identify relatively rare events.²

The second part of Nuijten et al.'s (2017) paper deals with my previous commentary on the program (Schmidt, 2017), where I voiced additional concerns about Statcheck's validity. The first one is the problem of corrected p values. Statistical tests make assumptions about the underlying stochastic structure of the data, and it is necessary to correct the tests if those assumptions are violated. Examples include corrections for multiple testing (e.g., Bonferroni corrections, Tukey tests), for making tests post-hoc instead of planning them beforehand (e.g., Scheffé tests), and for violations of stochastic assumptions (e.g., Greenhouse-Geisser or Huynh-Feldt corrections). In their first paper on Statcheck, Nuijten et al. (2016) countered this by performing a string search on two keywords that might signal such corrections: "Bonferroni" and "Huynh-Feldt". Now, their new search reveals that the Huynh-Feldt correction is about 25 times more prevalent than they originally reported, and that the Bonferroni correction is about 30 times more prevalent. They include three additional search terms (out of many others that could also be included), arriving at a prevalence of p -value corrections of 17.9 % of all papers (compared to their previous suggestion of about 0.33 % -- a 54-fold increase). They show that not all of these corrections cause trouble for the program. But the Greenhouse-Geisser and Huynh-Feldt corrections in repeated-measures tests do: Of all tests using one of those corrections, 29.5 % were flagged. These are all false alarms by Statcheck, brought about by the peculiar way these tests are often reported (with corrected p values, but uncorrected degrees of freedom).

The second concern I raised in my review was this: How many tests are only consistent because they *fail* to make the necessary corrections? These would be

¹ For Statcheck's "default mode" (without the option to count in one-sided tests), the authors claim a sensitivity of 100 %. My calculation is 34/76, or 44.7 %.

² *Signal Detection Theory* provides independent estimates of sensitivity and response bias (Macmillan & Creelman, 2005). Its sensitivity index d' is defined as $z(H) - z(F)$ and has a value of 2.15 (standard error: 0.19). The response bias index c is defined as $-2 \cdot [z(H) + z(F)]$ and has a value of 1.02 (standard error: 0.10). Both values are significantly different from zero, indicating that Statcheck performs better than chance but has a conservative bias against flagging tests.

erroneously cleared by the program. They should be counted as additional misses and would further reduce the program's sensitivity and phi coefficient. This problem is not discussed in the new paper.

The authors' conclusion is this:

"...we consider [Statcheck] to show a high level of diagnostic accuracy. Therefore, we recommend the use of statcheck for checking one's own work, for use in peer review, and as a tool to estimate the general prevalence of reporting inconsistencies across a large sample of articles."

I disagree with all of these statements. Statcheck has low validity, misses many inconsistent tests and makes many false alarms. It is not markedly superior to just accepting all tests at face value. It is too inaccurate to make it a useful desktop helper for manuscript preparation, because the time for double-checking the Statcheck output is certainly better spent double-checking the original statistics. The program can be especially misleading for inexperienced users who might think that Statcheck gives them a "clean bill": The program is biased against flagging tests at all, creating a false sense of security, especially when important *p*-value corrections have actually been omitted. It is generally unacceptable as a research tool, and certainly unacceptable for purely automatic scanning of multitudes of papers. The only way to find out about the accuracy of test reports is the meticulous way the authors demonstrated in their new paper and in their earlier work: Actually reading the papers, understanding the type of argument and the context of each test, applying good judgment, and discussing the more difficult cases with others.

References

- Nuijten, M. B., Hartgerink, C. H. J., van Assen, M. A. L. M., Epskamp, S., & Wicherts, J. M. (2016). The prevalence of statistical reporting errors in psychology (1985-2013). *Behavior Research Methods*, 48, 1205-1226. doi:10.3758/s13428-015-0664-2
- Nuijten, M. B., van Assen, M. A. L. M., Hartgerink, C. H. J., Epskamp, S., & Wicherts, J. M. (2017). The validity of the tool "statcheck" in discovering statistical reporting inconsistencies. *PsyArXiv Preprints*. doi:10.17605/OSF.IO/TCXAJ
- Macmillan, N. A., & Creelman, C. D. (2005). *Signal detection theory: A user's guide* (2nd edition). New York: Psychology Press.
- Schmidt, T. (2017). Sources of false positives and false negatives in the STATCHECK algorithm: Reply to Nuijten et al. (2016). *arXiv:1610.01010v7 [q-bio.NC]*
- Schmidt, T., Hauch, V., & Schmidt, F. (2015). Mask-triggered thrust reversal in the negative compatibility effect. *Attention, Perception & Psychophysics*, 77: 2377-2398. doi:10.3758/s13414-015-0923-4
- Seydell-Greenwald, A., & Schmidt, T. (2012). Rapid activation of motor responses by illusory contours. *Journal of Experimental Psychology: Human Perception and Performance*, 38, 1168-1182. <http://dx.doi.org/10.1037/a0028767>
- Wicherts, J. M., Bakker, M., & Molenaar, D. (2011). Willingness to share research data is related to the strength of the evidence and the quality of reporting of statistical results. *PLoS One*, 6, e26828.

Cite as: Schmidt, T. (2017). Statcheck does not work: All the numbers.
Reply to Nuijten et al. (2017). *PsyArXiv*. November 22. psyarxiv.com/hr6qy.

Author note

Correspondence may be sent to the author at thomas.schmidt@sowi.uni-kl.de.