

Writing Good Software Engineering Research Papers: Revisited

Christopher Theisen

Computer Science
NC State University
Raleigh, NC, United States
crtheise@ncsu.edu

Marcel Dunaiski

Computer Science
Stellenbosch University
Matieland, South Africa
marcel@ml.sun.ac.za

Laurie Williams

Computer Science
NC State University
Raleigh, NC, United States
lawilli3@ncsu.edu

Willem Visser

Computer Science
Stellenbosch University
Matieland, South Africa
wvisser@cs.sun.ac.za

Abstract—With the goal of helping software engineering researchers understand how to improve their papers, Mary Shaw presented “Writing Good Software Engineering Research Papers” in 2003. Shaw analyzed the abstracts of the papers submitted to the 2002 International Conference of Software Engineering (ICSE) to determine trends in research question type, contribution type, and validation approach. We revisit Shaw’s work to see how the software engineering research community has evolved since 2002. The goal of this paper is to aid software engineering researchers in understanding trends in research question design, research question type, and validation approach by analyzing the abstracts of the papers submitted to ICSE 2016. We implemented Shaw’s recommendation for replicating her study through the use of multiple coders and the calculation of inter-rater reliability and demonstrate that her approach can be repeated. Our results indicate that reviewers have increased expectations that papers have solid evaluations of the research contribution. Additionally, the 2016 results include at least 17% mining software repository (MSR) papers, a category of papers not seen in 2002. The advent of MSR papers has increased the use of generalization/characterization research questions, the production of empirical report contribution, and validation by evaluation.

Keywords—abstracts; research; guidelines; writing

TALK DESCRIPTION

Mary Shaw published “Writing Good Software Engineering Research Papers” [1] with the goal of helping software engineering researchers understand how to improve their papers. Shaw analyzed the abstracts of the papers submitted to the International Conference on Software Engineering (ICSE) in 2002, and posed three questions for each paper: “What, precisely, was your contribution?” “What is your new result?” and “Why should the reader believe your result?” The clarity of the answers to these questions are, per Shaw, indicative of the authors presenting research results well. In our talk, we reflect on Shaw’s work from 2002 and replicate her study on the abstracts of the papers submitted to the 2016 conference.

We examined the 2016 abstract data and the abstract data from the earlier work to analyze how research question type, contribution types, and validation type have evolved. The abstract data also allowed us to investigate the demographic information of the submissions to ICSE. In our talk, we present several of our findings from our study. We saw an increase of papers classified as “Generalization or Characterization” in 2016 compared to 2002, with 17%

of papers classified as “Generalization or Characterization” versus 6% in 2002. The popularity of mining software repository (MSR) papers in recent years could explain the difference. Evidence of the shift toward MSR papers is found in all types of research questions, types of contributions, and types of validation. Most types of contributions had similar acceptance rates, with “Empirical Report” having a slightly higher rate of acceptance than our other categories.

The program co-chairs observed an emphasis in the review committee discussions on the need for reproducibility. For papers where an artifact is produced, often in the procedure or technique and tool or notation categories, reviewers often want to see that the artifact is available to the general public, such as being available open-source. Review committees had a negative perception if there is not at least mention made that the artifact will be made available if the paper is accepted. These observations also apply to the availability of data that was mined to produce an MSR-type paper.

The validation acceptance rates were the most pronounced of the three types of categorizations performed during this study, indicating that a researcher’s choice of validation technique is the most important factor in determining if a paper is accepted. Simply submitting a novel approach to solving a problem with little backing does not seem to be an approach that will result in a high acceptance rate. The low number of experience-based validations indicates that existing techniques are not often validated by researchers in follow-on work. This could be a combination of the difficulty in getting industrial participants for studies, or a perception that replications or follow-on work is not valued by the ICSE community.

We have submitted the full version of our work to the ACM Transactions on Software Engineering and Methodology (TOSEM). While we are unable to make the full dataset available to researchers, we hope that our presentation will provide valuable information to the software engineering research community.

REFERENCES

- [1] M. Shaw, “Writing Good Software Engineering Research Papers,” *Proceedings of 25th International Conference on Software Engineering (ICSE’03)*, pp. 726-736, 2003.