The Gap between Visualization Research and Visualization Software (VisGap) (2020)
C. Gillmann, M. Krone, G. Reina, T. Wischgoll (Editors)

From Research Topic to Industrial Practice: An Experience Report

P. Gospodnetic[1], M. Rauhut[1] and H. Hagen[2]

[1]Fraunhofer ITWM, Germany
[2]University of Kaiserslautern, Germany

Abstract

Continuous technological progress requires both research and industry to work together. It is a necessity which cannot and should not be avoided. However, due to different interests of the two, it is often accompanied by various challenges. The inability to foresee and overcome the challenges can greatly impact the quality of collaboration results and thus chances of such results being used further by the industry. In this paper, we provide background on the topic and emphasize frequently discussed points. Focus of the work are industry-academia collaborations in applied computer science research such as visualization. For that purpose, a set of requirements is recognized and provided for both industry and research community. Further, we provide an overview of challenges recognized over years of experience from working in industry-academia collaborations. Together, the challenges indicate the gap between the industry and research which is inherently transferred further onto results of collaborative research. Finally, we discuss various possibilities for both industry and research to reduce the gap.

1. Introduction

Product innovation is an important business component which helps companies to maintain their market competitiveness. Continuous industrial development, driven by aim for further growth, often highlights research gaps as well as unanswered questions and fields that need further exploration. In-house research and development (R&D) teams or departments are found in many companies, especially market leaders [Eur19]. However, regardless of the presence of an in-house R&D team, companies will always benefit from cooperation with research institutions in order to allow new technologies to be developed. This makes it possible to greatly extend the the company's knowledge pool on a need-to-solve basis.

Solving a newly-emerged problem or introducing improvements into an existing product is rarely a straightforward process. Typically, it requires a multidisciplinary team of people with various specializations. Well-defined research collaborations between industry and research institutions make it possible to form such work groups and achieve results. Furthermore, it enables bidirectional transfer of knowledge. The research community gets a possibility to gain insight into state-of-the-art industrial practices, while industry gains access to knowledge and methods outside of their specialized area of expertise.

It is important to note that research collaboration can produce new breakthroughs as well as discover deadlocks, i.e. methods which do not work as expected and thus cannot be applied for solving the given problem. Both outcomes are beneficial, but the latter one is, understandably, never preferred and is often disregarded in result presentation. Although it presentation may be beneficial in the future by preventing others to go down the same path, or enabling them to re-evaluate the approach and potentially discover why it did not work.

Research collaboration projects and incentives such as International Research Training Group (IRTG) [DFG20] make it possible to allocate resources and knowledge for testing various hypotheses, enhancing existing products, solving specialized problems, evaluating method performance and more. Those efforts can potentially have a desirable byproduct of discovering new, unexpected, solutions as well as getting inspiration for future research possibilities, which is desirable. While it all sounds flawless in theory, in practice, we do encounter a very frequent, yet, seldom discussed problem - successful transfer of research results into industry and their integration into common workflows.

The problem can be well generalized to any industry-academia collaboration. However, for the purpose of this work, our focus lies on collaborations aiming to solve specific industrial problems using applied computer science research. Visualization research is one such type of research. Aforementioned collaborations are assumed to have an open question in need of solving, together with a dataset defining the area of application. Outcome of the collaboration is one or more methods, implemented in a prototype showing the applicability. Handover of the prototype implementation is, however, not necessarily a required part of the solution.

Over the years of work on research projects, where the main goal is integration of new solutions into industry, we have recog-

DOI: 10.2312/visgap.20201109

nized a set of common challenges. In this paper we want to give our overview of the challenges and propose ways which may help overcome the gap between the research and the industry.

2. Background

Difference between theory (research) and practice (industry) has been discussed in various forms such as papers, books, panels and media, with the latter two being the most common ones. Schofield [Sch13] made a general literature review of industry-academia collaborations. In her work, factors required for successful knowledge transfer between researchers are evaluated, together with cultural and organizational differences between universities and companies. The cultural and organizational differences are further discussed to see how these may build a barrier for knowledge transfer. Finally, the legal aspect of negotiating the intellectual property and how it may block a collaboration is discussed.

2.1. Collaborations in Software Engineering Research

Much of the available work focused on collaborations between industry and academia in the field of software engineering research. Wohlin et al. [WAA*12] describe factors necessary for successful collaboration. Based on surveys conducted in Sweden and Australia, they list success factors and try to find the major differences. Among others, social skills and difference in understanding between different categories of people are found to be the most important. Garousi et al. [GPO16] also focused on challenges in industry-academia collaborations for software engineering research and have collected a detailed list of challenges and best practices mentioned over time. The challenges have been mapped to four stages of a collaboration project lifecycle: problem formulation, planning, operationalization, and transfer and dissemination. Prause et al. [PRD10] also examined the problem from the aspect of software development in highly distributed work groups on research projects. Such groups make use of different software tools for both resource management and development to enhance the quality of cooperation and, thus, the results. In their study, they have emphasized the importance of industrially exploitable research results for achieving and maintaining global competitiveness. Given the focus on software development, several interesting points have been addressed such as lack of automated testing and continuous integration practices, frequent neglect of standardized error handling and bug tracking processes and low rate of code reuse. The study highlights the general attitude that in research projects, software quality and reuse are considered lesser goals because nobody expects flawless operation from a prototype but novelty. Finally, they shape the idea of Research Software Engineering (RSE) for the purpose of highly distributed collaborative projects.

2.2. Research Software Engineering as a Discipline

Since Prause et al. [PRD10] coined the term, Research Software Engineering has received more recognition and has seen a steady rise in the number of such positions on the market. A report by UK Research Software Engineering Network [BCH*17] shows that research software developers tend to be highly educated individuals with research experience, which they continue to use in a way to enhance research software. The report noted that 70% of community have a PhD as a highest level of education and 88% reported their software contributed to a publication. Another interesting point was that 46% of people reported that a loss of a single developer would end the project and 30% percent reported projects would be ended by the loss of two developers, emphasizing how vulnerable research projects are to staff departure. Finally, the report highlights that 78% of projects employing RSEs have no technical handover plan, which would describe the product in enough detail to allow a new RSE group to get up to speed with the software and continue further development.

2.3. Research Reproducibility and Availability

Research availability and reproducibility is currently a hot topic which can be reflected in the growing number of good incentives such as ACM's review and badging process [ACM20], *Papers With Code* community [Pap20] and *Open Science Framework* [OSF20]. ACM's review and badging process is based on the premise that an experimental result is not fully established unless it can be independently reproduced. It is a review process, optional to any of ACM journals and conferences, requiring that any "artifact" (digital object) associated with the work undergoes a formal audit. The paper is further assigned one of 5 badges: Artifacts Evaluated - Functional, Artifacts Evaluated - Reusable, Artifacts Available, Results Replicated and Results Reproduced. *Papers With Code* is a community built around a mission to create a free and open resource providing Machine Learning papers, code and evaluation tables. Open Science Framework, built by Centre for Open Science, is a tool that promotes open, centralized workflows by enabling capture of different aspects and products of the research life cycle. Different aspects include developing a research idea, designing a study, storing and analyzing collected data, and writing and publishing reports or papers [FD17]. Up to now, it has received traction in the life science research community and is working its way to become a part of basic research protocol [SDM19].

2.4. Null Results

Null results, negative results, failed experiments. Different terms are used for what can also be put as a statement - "An experiment has been done, but failed to produce the expected results". It is a valuable experience which no one aims to obtain when conducting research. The main goal is to solve a problem. Efforts to make publishing of null results a common practice are visible in life science research [MH17, MH17], especially in form of on-line commentaries [WUR20], offering a reward for scientists who publish such results. Journals which continuously publish null results in life sciences exist [JAS20, PLO20], but successful journals publishing negative results are hard to find. Many journals which have tried to do the same and failed can also be found [ARJ20], some of them living to have only a single volume, if any. In computer science the results are very similar, with efforts being done by individuals. Journal of Interesting Negative Results in Natural Language Processing and Machine Learning [JIN20] published one volume, containing only one article. Prechelt [Pre97] established the Forum for Negative Results as a special issue of Journal of Universal Computer Science [JUC20], but it is unclear how many articles were published through it. Association for the Advancement of Artificial Intelligence also held their annual Spring Symposium with one of the titles being "What Went Wrong and Why: Lessons from AI Research and Applications".

3. Requirements

Successful execution of quality collaborative research between industry and academia requires from both sides to maintain good cooperation and clearly communicate their requirements and expectations. Throughout academia and industry, there is always the distinction between basic research and applied research. Applied research stands for research that has a direct practical impact in form of products, software, whereas basic research is often defined as an extension of the already existing knowledge. Although the strict separation is hard to be made, we assume that the interests for the industry lie in the applied research in particular. Furthermore, we concentrate mainly on the development of software, algorithms and methods.

3.1. Industrial Requirements

Industry must constantly move forward. Therefore, new technologies and methods should be available to test, evaluate and implement. Beside the technical aspects of implementation, industry must also pay a great deal of attention to legal concerns as well as ensure that new methods satisfy standards, if applicable. Finally all of the above mentioned concerns must be cost-effective.

3.1.1. Method Proposal and Evaluation

Typically, it is the young professionals who bring fresh ideas to the companies. The proposed ideas must be justified before the company decides to pursuit the idea further. For that purpose, the young professionals must ensure their ideas are supported by knowledge tied to the specific area of application. A preliminary research must be done by the industry where experiments would be conducted, ensuring main features and use of the new method, together with evaluation:

- Is the method superior to an existing solution (if it exists)?
- Which costs would the method cut or bring?
- Which overall benefits does the method bring?

As a result, a proof of concept must be made available. It should describe the use-case, emphasize evaluation conclusions and make possible to determine the steps which would be required in order to develop a fully working system, ready for integration into industrial processes.

3.1.2. Method Availability

The above described proposal for method integration into an industry process relies on the amount of knowledge already made available in forms such as research publications. Furthermore, it relies on the previous knowledge of industry professionals conducting the experiments, quality of method presentation in the publication as well as method reproducibility. Every example, source code and dataset comes as a great benefit, making it only more likely that the method will be taken into consideration. Vocabulary used to describe a method in a publication is another interesting point. It can be expected that the vocabulary used in a publication is specific to the field of publication, which may make it rather difficult to read by a person who is not directly familiar with the field. Therefore, publications written in a manner accessible to a wider audience are more likely to be used.

3.1.3. Legal Concerns

Before deciding to use a particular method, or an implementation, it is very important that all the legal aspects of its use are completely clear. Under which licence is the code (or software) made available to the public? Is the method protected by a patent? Who is the right contact person with any potential requests for further use?

3.2. Research Requirements

The goal of applied research is to generate solutions to specific problems, posed by a client, using sound evidence and thinking [Bai18]. Furthermore, every conducted research should make the research findings available for the others to be able take the work further.

Research is a mixture of replicatory and exploratory activities. Replicatory activities seek to confirm previous research findings in order to decide whether they are applicable to the problem in question as well as highlight parts of the method which should be improved. Exploratory activities build further on the foundation achieved by replicatory activities and seek to improve existing methods in order to apply them, or explore them and theorize about specific aspects of the problem. It is often the case that applied research triggers questions about basic research problems, but for the purpose of this paper, basic research is outside of scope of this paper.

3.2.1. Problem Definition

Research problem should be posed in a way to answer the following questions:

- What should the solution be applied to?
- What particular questions should the research answer?
- How should the quality of a solution be evaluated?

Datasets for which the solution is developed must be available in advance, together with additional information describing the process of how it was obtained, what the data represent and a semantic explanation of the desired solution. While it might not always be possible to fulfill the requirement, there should be more than one dataset available to test the possible solution on. Datasets should contain average expected data as well as possible outliers. The presence of the outliers enables robustness enhancement as well as better distinction of what can or cannot be applied and when.

3.2.2. State-of-the-art

If the problem to be solved can be regarded as an enhancement of an existing process, an overview of state-of-the art methods should be given. The overview should elaborate on how this problem has been tackled so far, what are the outcomes of current solutions, and what can or should be improved.

3.2.3. Resources

Time and money resources allocated for the research must take into consideration the time required to get accustomed with the topic and its background, time to learn about state-of-the-art, do applicable method research and testing, solution development, and result presentation. Furthermore, research should be done by a team of people, rather than a single person.

3.2.4. Research Results

Expected research results are one or more prototype solutions to the posed questions. The work put into researching should also result in well documented findings, describing the result and decision process behind the chosen method. In order to support the future development, the obtained knowledge and results should be disseminated to a wider audience in form of reproducible publications.

4. Challenges

Industrial development can be seen as result driven; aiming to obtain advancement in time-cost effective manner. Research development, on the other hand, is knowledge driven; aiming to completely comprehend the problem at hand and find sufficient and theoretically substantiated answers, regardless of how long it may take. The two aspects might be conflicting at times, which can be reflected in the challenges described further in this chapter. Especially when it comes to applied research through collaborations, if the challenges are not overcame or handled with care, the collaboration outcomes might not be completely satisfactory for any of the two sides.

4.1. Problem Statement

Continuous development is based on the premise that most, if not all, of the state-of-the-art methods have a possibility to be enhanced. Choosing which method to enhance, and when, however, is greatly dependent on the impact the method has on the overall process. For that purpose, the most rational thing to do is to focus the research efforts into areas with the biggest impact. The responsibility of highlighting the areas in need for improvement lies exclusively on the individuals working in the industry, as they have the necessary insight and experience to evaluate both performance and impact of the possible enhancement.

Research community obtains access to that kind of information on a need-to-know basis through research collaboration projects. Outside of projects there is almost no way to obtain relevant insight into open problems. University-industry round tables [The20] or seminars like Dagstuhl Seminars [Sch20] are encouraged with intention to help building a bridge between research and industry. While well intended, in reality, meetings like that are often too high-level to make a difference on a level of a single research group.

Such situation has twofold impact on the work of the research community. 1) Research effort, which might be applicable to a real problem, has been done, but has gone unnoticed and, as such, is likely to be forgotten. 2) Research effort has been made to answer a problem which was not even a problem to begin with and is, as such, useless to the industry.

4.2. Research Under Confidentiality Agreements

When a company invests in research, it has to make sure none of the company secrets are revealed during the research. Also, companies have no direct benefit from sharing the research findings with other interested parties (which were not part of the research team). Such practice makes it very difficult to give the obtained knowledge back to the research community and impedes overall research advancements.

Typical problems occurring are:

- Research publishing is not possible due to restrictions in sharing intellectual property obtained during research
- Published research is not reproducible because the dataset and code which was developed to process it cannot be made publicly available
- Verification of research methodology is not possible because data, to which the method might be applicable, will not be made available

4.3. State-of-the-art Solution Availability

Section 4.2 depicts practices which benefit only a narrow range of researchers, while directly hindering advancement of the general knowledge base. It is, however, only one part of the problem. When it comes to software solutions, another problem is the availability of proprietary products. While some companies make their products available free of charge for research purposes, others do not offer it all, or at least do not disclose the practice publicly.

If a research group is working on an enhancement of a particular process, it is very hard to assess the relevance and impact of the solutions against commercially offered tools. Because of that, the chance of developing multiple similar solutions to the same problem is high. Furthermore, the solutions might be sub-optimal when compared to the performance of already existing solutions.

4.4. Software Development Practices

The line between research software development and production software development is hard to define and typically relies on a set of software development best practices which may or may not be enforced during the development process. A study by Prause et al. [PRD10] reported researchers emphasizing that their primary task is to reuse the concepts, not the software. Likewise, they argue that, since they are not expected to deliver complete software solutions to the costumer, software engineering best practices are not applicable for them.

We have observed that such practices frequently result in partial software solutions which have not been designed, rather, their structure grew more complicated as the research progressed. The solutions, as such, are hard to use, mostly undocumented, error prone, hard to maintain and are almost impossible to incorporate into the existing industrial processes. Therefore, in order to actually use the solution, the company must invest further time and money to reverse engineer the implementation and then rebuild it in an efficient way which conforms to software engineering standards and is maintainable.

4.5. Prototyping - Applicability vs. Reproducibility

Bryson et al. [BCS06] define a collaboration between industry and research institutions as the linking of organizations in two or more sectors to jointly achieve an outcome that could not be achieved otherwise. In that sense collaborations can often be seen as a form of feasibility study, where the main goal is to identify concepts capable of advancing current practices.

The research problem is provided to the researchers in form of problem description and a dataset representing the problem. Applicability of the developed solution concepts to the given problem is presented using prototypes. Since the datasets are typically very specialized and industry dependent, the prototype is likely to integrate assumptions which are not easily transferable to a similar dataset of different origin.

Privately funded research collaborations rarely make dissemination of knowledge as one of the project goals. Therefore, if done at all, the dissemination of knowledge obtained during research is done in form of an application paper, where the prototype or datasets are seldom made available. Due to the unavailability of the prototype or the datasets and highly specialized solutions, the concepts are very hard to reproduce.

4.6. Publishing and Transfer of Knowledge

Quality of a research group is measured by number of publications and their impact factor. Such practice frequently drives the so called *publish or perish* attitude, where the group members are encouraged to publish as frequent as they can, regardless of the research having any real industrial value.

Basic research has little ties to the industry and, as such, has no benefit of clearly communicating to an audience outside of their specialized research network. The lack of clear communication to a wider audience renders the research hard to understand by anyone other than the specialized researcher and thus hard to determine circumstances under which the solution would bring benefit if applied.

Furthermore, everyone would like to know which methods work and which do not. However, *null result* papers are rarely published since they are considered as research failure and bring no particular benefit to the research group, beside the insight. Therefore, many research projects are forced to repeat the same method, only to discover for themselves that it is inapplicable.

Finally, as mentioned in Section 4.5, industry has typically little interest in publishing results. Because of that, a lot of successful research projects go by unnoticed by the rest of the research community.

4.7. Research Group

University departments and research groups are specialized in certain algorithms and are often not interested in research of other applicable methods, that differ greatly from their owns. While it can lead to new, unconventional, solutions to a problem, it can also lead to solutions which require an unnecessary amount of overhead in order for them to work, if a solution can be found at all. Therefore, when forming a research collaboration, industry partners should be well informed about the practices and field of work of the research partners.

5. Bridging the Gap

In order to have technological progress, research needs valid questions needing answers, and industry needs high quality answers. Therefore it is necessary that both sides make an effort to find a common ground and thus minimize the encountered challenges.

Industry's best interest is to communicate recognized problems and necessary process improvements to a wide range of researchers. It is understandable that complete and accurate data, needed to describe a problem, cannot be made publicly available without revealing the company's trade secrets. However, providing artificial data describing the problem can make it possible for the research community to depict the idea of the problem and reach out with possible solutions. Therefore, putting an effort into preparation of data similar to the original, without revealing the details, and making the data publicly available can ease the way to completely new solutions, enhancement of current processes and thus profit.

An example of such a situation is as follows: A company manufacturing car parts must have a CAD model of their products. Releasing the model publicly is out of the question because it incorporates precise details about their products, many of which are trade secrets. A research group has recognized that a possible solution exists which would help enhance a vital part of the production line - a system conducting automated quality inspection. The solution, however, requires a CAD model of the product as an input because the solution is in part based on visualization of the object as well as on analysis of geometrical characteristics of the object. In order for the research group to test their solution, they need a 3D model which is realistic to expect in one of such production lines, but none of the publicly available 3D models do not even resemble specialized car parts. When tested on available models, the solution might not look interesting to the industry because the models are not intricate enough.

The situation could be mitigated if the company would make available CAD models of degraded aspect ratios, thus hiding any sensitive details, or models of products which are long out of use and are of no value to the company any longer. Furthermore, artificial or degraded samples would also give more space to the research community to publish the results and new findings about applicable methods.

The round table incentives, mentioned in Section 4.1, would benefit from reinforcing the discussed topics with a set of challenges interesting to the industry. They should be precisely defined, accompanied by an example dataset and provide a possibility of to compare results to used state-of-the-art solutions. The challenges should be made available to a wider audience, especially younger researchers, as it may help them focus their future research towards a certain topic. Also it may bring fresh ideas to the table.

The need to identify high quality research papers in terms of reproducibility is becoming more and more visible with incentives like the ACM badging [ACM20]. Systems like that are much needed and should be encouraged further by both industry and research community as it provides a solid base for method quality evaluation and establishing good practice. Allocation of resources only on replication studies, within the scope of larger government funded research projects is likely to support those effort even more. An interesting and very useful information would also be an evaluation of method applicability from an industry point of view, which is currently unavailable.

It is understandable that the research focus lies on the method development rather than the development of a wholesome software solution. However, every effort to make the research results more available to the industry comes as a great benefit to both industry and research community. An effort to communicate the research results more clearly, give a simplified interpretation of the research methods and provide at least minimal working examples would ensure an easier transfer of knowledge to industry experts. Furthermore, giving more attention to software development in terms of best practices, modularity, and useful API would enable easier integration and testing. Finally, thorough documentation stating the environment in which the solution was developed and describing pipeline logic through steps, inputs and outputs would surely come as a great benefit to the whole research community.

The presence of research software engineers seems to be growing. While computer science and software engineering go hand-in-hand, not every computer scientist is necessarily a good software engineer thus creating the need for research software engineers. Furthermore, every computer science topic comes with a set of software engineering practices which might be specific to it. Knowledge of those practices is very valuable and should be promoted by making topic-dependent research software engineering work a part of every conference program. Also, by establishing research software engineering groups and making them available to researchers would come as a positive step towards availability, quality and reproducibility of research results.

As can be seen from Section 2, publishing null results in form of a traditional scientific publication has been asked for many times, however it is unlikely it will ever become a standard practice, as obtaining results is only half of the way to submitting a final publication. Preparing a publication requires both time and effort which could otherwise be invested into research resulting in a novel, relevant solution, which is far more useful. Yes, lessons learned are needed, but not in a form of a scientific publication. A much better way to do so are experience reports. Online in form of blogs, or live in form of short talks. That way the experience is more likely to be shared, and anyone more interested into it can always reach out to the author to discuss it further. In order to prepare a short commentary or a presentation, far less time and effort are required. Also, it can potentially benefit both industry and research community.

A lack of insight and experience with industry state-of-the-art can cause research findings inferior to existing solutions or answers to questions which are not interesting to the industry. Providing researchers with a possibility to gain industrial experience, or transition of people with industrial experience into research would be beneficial for both parties. Furthermore, when conducting research, researchers should look more into industrial standards and industrial applications. If the research tends to go more in the direction of the basic research, the researchers should take the time to suggest possible scenarios or application which could potentially benefit from the newly found solution. Finally, all approaches have their benefits and drawbacks. The benefits are typically well highlighted within the publications, while the drawbacks tend to be discussed only briefly, if at all. Placing more emphasis on the method's potential drawbacks and vulnerabilities would provide a more wholesome and useful context. An important question to that end is, how would the reviewer board react to such a discussion, given that reviewers tend to reject submissions based on the methods drawbacks and vulnerabilities? Would the proposed practice only give the reviewers more chance to dismiss the approach as irrelevant, or would it be welcomed as a publication written to support research integrity?

Research institutions oriented to bringing academia and industry closer together (e.g. Fraunhofer Society, Hasso Plattner Institute, Helmholtz Association) play an important role in gap minimization, reconciling interests of both parties to mutual benefit. Such institutions have strong connections within both academia and industry, with good insight into what industry needs as well as what academia has to offer. Typically, their staff has a firm academic background with experience in industry applications. As such, institutions of that kind should strongly support scientific journals in form of a reviewing process, with feedback always providing comments on the publication's suggested area of application, as well as a suggestion of other areas which may benefit of the proposed method.

6. Conclusion

Industry needs research and research needs industry, but their interests are frequently conflicting. A silver bullet, solving all the challenges, does not exist and challenges will surely continue to appear. What should be done is to work on identifying current challenges and solving them in a manner which is likely to get accepted by both communities, rather then forcing rigid unnatural solutions which are hard to enforce and bring questionable benefits, along with unknown costs. In Section 5 we recommended a set of actions which could be implemented by all parties in order to make the collaboration more successful and ensure easier transition of research results into good practice accepted by industry. Paying attention to the suggested actions during the planning stage of the collaboration and including them in the overall time schedule is sure to bring benefits. While it might be difficult to focus on all the actions at once, we suggest that first steps be made towards better communication between the two. Industry should communicate current challenges which need to be solved and make datasets available to research. Research community should ensure that all research results are provided openly through a well documented minimal working example and with at least parts of the source code available.

References

[ACM20] ACM, Association for Computing Machinery: Artifact review and badging, 2020. www.acm.org/publications/policies/artifact-review-badging. 2, 6

[ARJ20] ARJ, All Results Journals:, 2020. www.arjournals.com. 2

[Bai18] Baimyrzaeva M.: *Beginners' Guide for Applied Research Process: What Is It, and Why and How to Do It?* Tech. rep., 2018. 3

[BCH*17] Brett A., Croucher M., Haines R., Hettrick S., Hetherington J., Stillwell M., Wyatt C.: *Research Software Engineers: State of the Nation Report 2017*. Tech. rep., 2017. 2

[BCS06] Bryson J., Crosby B., Stone M.: The design and implementation of cross-sector collaborations: Propositions from the literature. *Public Administration Review 66*, SUPPL. 1 (12 2006), 44–55. doi:10.1111/j.1540-6210.2006.00665.x. 5

[DFG20] DFG: International Research Training Group, 2020. www.dfg.de/en/research_funding/programmes/coordinated_programmes/research_training_groups/international_rtg/index.html. 1

[Eur19] European Commission - Joint Research Centre: *EU R&D Scoreboard*. Tech. rep., 2019. 1

[FD17] Foster E. D., Deardorff A.: Open science framework (osf). *Journal of the Medical Library Association 105*, 2 (2017). 2

[GPO16] Garousi V., Petersen K., Ozkan B.: Challenges and best practices in industry-academia collaborations in software engineering: A systematic literature review. *Information and Software Technology 79* (2016), 106 – 127. doi:https://doi.org/10.1016/j.infsof.2016.07.006. 2

[JAS20] JASNH Journal of Articles in Support of Null Hypothesis:, 2020. www.jasnh.com. 2

[JIN20] JINR, Journal of Interesting Negative Results in Natural Language Processing and Machine Learning:, 2020. www.jinr.site.uottawa.ca. 2

[JUC20] JUCS, Journal of Universal Computer Science:, 2020. www.jucs.org/jucs. 3

[MH17] Miller-Halegoua S. M.: Why null results do not mean no results: negative findings have implications for policy, practice, and research. *Translational Behavioral Medicine 7* (2017). 2

[MH17] Mlinarić A., Horvat M., Šupak Smolčić V.: Dealing with the positive publication bias: Why you should really publish your negative results. *Biochemia Medica 27* (2017). 2

[OSF20] OSF, Centre for Open Science: Open science framework, 2020. www.cos.io/our-products/osf. 2

[Pap20] Papers With Code:, 2020. www.paperswithcode.com. 2

[PLO20] PLOS, The Missing Pieces: A Collection of Negative, Null and Inconclusive Results:, 2020. www.collections.plos.org/missing-pieces. 2

[PRD10] Prause C. R., Reiners R., Dencheva S.: Empirical study of tool support in highly distributed research projects. In *2010 5th IEEE International Conference on Global Software Engineering* (2010). doi:10.1109/ICGSE.2010.13. 2, 4

[Pre97] Prechelt L.: Why we need an explicit forum for negative results. *Journal of Universal Computer Science 3*, 9 (1997). 2

[Sch13] Schofield T.: Critical success factors for knowledge transfer collaborations between university and industry. *Society of Research Administrators International 44*, 2 (2013), 38. 2

[Sch20] Schloss Dagstuhl - Leibniz Center for Informatics: Dagstuhl seminars, 2020. www.dagstuhl.de/en/program/dagstuhl-seminars. 4

[SDM19] Sullivan I., DeHaven A., Mellor D.: Open and reproducible research on open science framework. *Current Protocols in Essential Laboratory Techniques 18*, 1 (2019). 2

[The20] The National Academies of Sciences, Engineering and Medicine: Government-University-Industry Research Roundtable, 2020. sites.nationalacademies.org/PGA/guirr/index.htm. 4

[WAA*12] Wohlin C., Aurum A., Angelis L., Phillips L., Dittrich Y., Gorschek T., Grahn H., Henningsson K., Kagstrom S., Low G., Rovegard P., Tomaszewski P., van Toorn C., Winter J.: The success factors powering industry-academia collaboration. *IEEE Software 29*, 2 (2012), 67–73. 2

[WUR20] WUR, The World Univesity Rankings: Scientists offered 1,000 euro to publish null results, 2020. 2