A Procedure and Guidelines for Analyzing Groups of Software Engineering Replications

Adrian Santos, Sira Vegas, Markku Oivo and Natalia Juristo

Abstract—Context: Researchers from different groups and institutions are collaborating on building groups of experiments by means of replication (i.e., conducting groups of replications). Disparate aggregation techniques are being applied to analyze groups of replications. The application of unsuitable techniques to aggregate replication results may undermine the potential of groups of replications to provide in-depth insights from experiment results. **Objectives:** Provide an analysis procedure with a set of embedded guidelines to aggregate software engineering (SE) replication results. **Method:** We compare the characteristics of groups of replications for SE and other mature experimental disciplines such as medicine and pharmacology. In view of their differences, the limitations with regard to the joint data analysis of groups of SE replications and the guidelines provided in mature experimental disciplines to analyze groups of replications, we build an analysis procedure with a set of embedded guidelines specifically tailored to the analysis of groups of SE replications. We apply the proposed analysis procedure to a representative group of SE replications to illustrate its use. **Results:** All the information contained within the raw data should be leveraged during the aggregation of replication results. The analysis procedure that we propose encourages the use of stratified individual participant data and aggregated data in tandem to analyze groups of SE replications. **Conclusion:** The aggregation techniques used to analyze groups of replications should be justified in research articles. This will increase the reliability and transparency of joint results. The proposed guidelines should ease this endeavor.

Index Terms—Replication, Statistical Analysis, Aggregated Data, Individual Participant Data, Narrative Synthesis.

1 INTRODUCTION

Experiments are commonplace in SE [1]–[3]. Still, two main shortcomings usually have an impact on their suitability for evaluating the effectiveness of SE technologies [3]: (1) sample sizes are usually small [4], and (2) results are only generalizable to the configuration of the experimental settings [5].

With the aim of increasing the reliability and generalizability of individual experiment results, SE researchers are working on building groups of experiments by means of replication (conducting groups of replications) [6]–[10]. By collaborating with each other (e.g., sharing experimental material, and assisting each other during the design, execution and/or analysis phase of experiments, etc.), researchers are able to increase the sample size, as well as evaluate the effects of the treatments under different settings. This should increase the reliability of results and their generalizability to different contexts and populations [11].

Groups of replications provide some advantages for evaluating the effectiveness of SE treatments [12]–[16]: (1) access to raw data provides for the use of consistent pre-processing and analysis techniques to analyze each experiment, thereby increasing the reliability of joint conclusions; (2) researchers conducting groups of replications may limit the changes made across the replications in order to increase the internal validity of joint conclusions; (3) joint conclusions are not affected by the detrimental effects

of publication bias, as groups of replications do not rely on already published results; (4) consistent measurement instruments can be used across replications in order to measure participant characteristics with identical methods and scales and, possibly, stratify the results according to such characteristics.

According to a systematic mapping study (SMS) that we undertook [17], five techniques are being applied to aggregate groups of SE replications (listed from most to least used): narrative synthesis, aggregated data, mega-trial or stratified individual participant data, and aggregation of p -values. According to the literature of mature experimental disciplines like medicine and pharmacology [18], [19], some aggregation techniques are more suitable than others depending on the characteristics of the group of replications. We observed similar findings when applying the aggregation techniques to analyze a stereotypical group of SE replications [20] (a small group of replications with small and dissimilar sample sizes, opportunistic participant recruitment, different types of subjects, identical experimental designs and response variable operationalizations, and heterogeneous results [17]). The applied aggregation technique definitely had a big impact on the reliability of joint conclusions [20].

In view of this, the aim of this study is to answer the following main **research question**:

- How should groups of SE replications be aggregated?

To answer this question, we performed a literature review in mature experimental disciplines (medicine and pharmacology) to learn about the techniques applied to aggregate replication results. Along the way, we also noticed

A. Santos and M. Oivo are with the M3S (M-Group), ITEE University of Oulu, P.O. Box 3000, 90014, Oulu, Finland, e-mail: {adrian.santos.parrilla,markku.oivo}@oulu.fi
S. Vegas and N. Juristo are with the Escuela Técnica Superior de Ingenieros Informáticos, Universidad Politécnica de Madrid, Campus de Montegancedo s/n, 28660 Boadilla del Monte, Madrid, Spain, e-mail: {svegas,natalia}@fi.upm.es

some (more profound than expected) differences between groups of replications in SE and medicine. This led to a subsequent literature review of studies with similar circumstances to SE in the fields of medicine, social research, educational research and econometrics. In view of the results, we tailored a *procedure*, with a set of embedded *guidelines*, to facilitate the aggregation of results in groups of SE replications. We apply the proposed procedure to analyze a group of replications in order to illustrate its use.

This article extends our prior work (i.e., [17], [20]) in several ways: (1) by *identifying differences* between the characteristics of groups of replications in the fields of SE and medicine and how such differences may impact the aggregation techniques used to analyze groups of SE replications; (2) by *proposing a step-by-step procedure* to analyze groups of SE replications that takes into account such differences and the typical limitations regarding joint data analysis of groups of SE replications; (3) by *providing a hands-on tutorial* with mathematical formulae and R code snippets to analyze a stereotypical group of SE replications; (4) by *providing a discussion and further pointers* to references indicating how to analyze groups of replications with different experimental designs.

The **take-away messages** of this research are:

- *Random-effects models should be preferred over fixed-effects models*, especially because many variables may impact SE experiment results, changes are frequent across SE replications, and heterogeneous results are commonplace. Differences between groups of replications in medicine and SE make it inappropriate to directly apply medical guidelines to the analysis of groups of SE replications. In particular, the application of fixed-effects models and traditional statistical thresholds (e.g., the traditional p -value of 0.05) in order to detect heterogeneity and moderators does not appear to provide guarantees in SE.
- *Avoid narrative synthesis* [21], *aggregation of p -values* [22] and *mega-trial individual participant data (IPD-MT)* [23], and *use aggregated data (AD)* [24] and *stratified individual participant data (IPD-S)* [25] in tandem instead. AD and IPD-S appear to be the most suitable techniques for analyzing groups of SE replications. AD provides intuitive visualizations to convey joint results and straightforward statistics to quantify heterogeneity. IPD-S increases the interpretability of joint results and offers greater statistical flexibility.
- *Strive to identify both experiment-level moderators¹ and participant-level moderators*. AD and IPD-S appear to be good at identifying experiment-level moderators. IPD-S appears to be preferable for identifying participant-level moderators.
- *Use the following four-step procedure to analyze a stereotypical group of SE replications*: (1) describe the characteristics of the participants using appropriate descriptive statistics and visualizations; (2) use consistent statistical techniques to pre-process, describe and analyze the data of each replication; (3) select suitable aggregation techniques to provide

joint conclusions; and (4) conduct exploratory analyses to identify experiment-level moderators² and participant-level moderators³.

The paper has been organized as follows. In Section 2 we provide the background of this study. In Section 3 we outline the research method that we followed to elaborate the proposed analysis procedure, and present the stereotypical group of replications that we will use to illustrate its application. In Section 4 we show the differences between groups of replications in medicine and SE. Then, in Section 5 we outline the most common limitations with regard to joint data analysis of groups of SE replications. In Section 6 we provide an overview of the four-step analysis procedure that we propose to analyze groups of SE replications. In Sections 7, 8, 9 and 10 we detail each of the steps of the procedure that we apply to the illustrative group of replications. We outline the threats to validity of this study in Section 11, and provide further pointers for the analysis of groups of SE replications with different experimental designs in Section 12. We relate our research to other SE research in Section 13. Section 14 states our conclusions.

2 BACKGROUND

2.1 Replication

The relevance of replication has been widely acknowledged in SE [27], [28]. Replication has been coupled in SE with the concept of applying a similar experimental procedure to the one applied in a previous baseline experiment on a different sample of participants to generate new raw data [29], [30].

However, two different concepts need to be set apart: replication and reproducibility—or reproduction of results. Researchers who want to reproduce results apply the same analysis procedure followed by the original experimenters to the original raw data with the aim of getting the same results [31]. Therefore, reproduction has to do with re-analysis of raw data from the baseline experiment [31]. Replications generate new raw data—and results—that can be later combined with the outcomes of other replications to provide joint conclusions. This research focuses on replication.

Different types of replications can be conducted. According to Gomez et al. [32], replication types vary along a continuum: from exact replications, following exactly the same experimental configurations as their baseline experiments, to conceptual replications, where the only thing that the replications have in common are the baseline experiment research questions and objectives. Somewhere in between these two extremes lie other replication types, where different elements of the baseline experiment configurations remain unchanged [33].

Laboratory packages were proposed to ease replication across research groups and institutions [34]. Laboratory packages contain relevant information needed to replicate an experiment [35]. With a laboratory package, an external group of researchers can reproduce the settings of a baseline

2. Characteristics of the experiments that may be impacting the results, such as the programming language or experimental session length of the experiments.

3. Characteristics of the participants that may be impacting the results, such as participant programming or Java experience.

1. Variables that cause an effect to differ across contexts [26].

experiment, and gather new raw data from a different sample of participants. In addition to sharing laboratory packages, experimenters conducting groups of replications may also collaborate with each other through face-to-face or Internet meetings to plan, design, execute and/or analyze their experiments [36]. This close collaboration may increase the chances of getting similar results across the replications—as experimenter interaction is expected to assure more similar experimental procedures. This should increase the reliability of joint conclusions [22].

Despite even the hardest efforts to conduct exact replications, conflicting results may still pile up [22], [37]. It is then that first-hand knowledge of experiment configurations and participant characteristics plays a central role. If such information is known, it is easier to hypothesize on the variables that may be behind divergent results. In turn, this is useful for hypothesizing on experiment-level or participant-level moderators that may be influencing the results. It is the above flexibility that leads this research to focus on groups of replications.

2.2 Groups of Replications

We conducted a SMS with the aim of learning what aggregation techniques are being used to analyze groups of SE replications [17]. We identified a total of 39 groups of replications that share certain characteristics:

- They are either conducted by *individual researchers* or by *groups of researchers working in close collaboration* across one or multiple research groups, universities and/or institutions. As such, researchers have access to the raw data of all the replications.
- They are formed *opportunistically*. In other words, a priori plans are not typically set for building groups of replications; each replication comes into being individually without a defined protocol at the inception of the group. As a consequence, replications are aggregated—generally after having being published individually—to either increase the reliability of the findings or to elicit moderators (e.g., assessing how the technologies perform for different types of subjects).
- Most groups of replications are composed of *three to five replications* evaluating the performance of a *binary treatment* (e.g., Method A vs. Method B) on a *continuous outcome of interest* (e.g., productivity measured as LOC per hour). Replications are usually *small*⁴, have *dissimilar sample sizes*, evaluate the performance of the treatments for *different types of subjects* (e.g., professionals vs. students), have *identical experimental designs and response variable operationalizations*⁵, and provide *heterogeneous results*.

2.3 Aggregation Techniques

Five aggregation techniques have been used in groups of SE replications [17]: narrative synthesis, AD, IPD-MT,

4. Out of consistency with other SE authors [38], small sample size refers throughout this article to experiments involving fewer than 30 subjects.

5. Therefore, internal and construct threats to validity are not mitigated or new ones cannot be identified.

IPD-S and aggregation of p -values. Thirty-five percent of the groups of replications use more than one aggregation technique. However, they usually serve different purposes: one for providing joint conclusions, and a different one for eliciting moderators. Thus, the groups of SE replications never compare the results achieved with different aggregation techniques for the same objective.

In the following, we review the aggregation techniques used in groups of SE replications starting with the most, and ending with the least, popular.⁶ To do this, we rely on the results of the SMS that we undertook [17].

Narrative synthesis was used to analyze 46% (18 out of 39) of the groups of replications [17]. In narrative synthesis (also known as semi-quantitative aggregation [22]), replication results (in either p -value or effect size terms) are combined textually to provide a summary of results. For example, it is common in SE to analyze each replication individually using a t -test or a Wilcoxon test and to then provide a textual summary of results of the replications as follows: “...while the results are statistically significant/large/negative in experiments X, Y and Z, they are not in experiment M. This difference of results could have been caused by H, N or K moderator variable...”.

AD—commonly known as meta-analysis of effect sizes in SE [24]—was used to analyze 38% (15 out of 39) of the groups of replications [17]. In AD, all replication effect sizes are first computed from summary statistics, such as means, variances or sample sizes—or from experiment statistical test results—and then combined by means of a meta-analysis model [22]. Two different types of meta-analysis models can be fitted: fixed-effects models or random-effects models [22]. Fixed-effects models assume that all the experiments estimate a common population effect size and, thus, differences across experiment results arise due to the natural variation of results (i.e., due to the different participant samples in the experiments). On the other hand, random-effects models assume that differences across experiment results arise not just from the natural variation of results, but also from a real heterogeneity of effects. In other words, random-effects models estimate a distribution of population effect sizes rather than a common population effect size.

IPD-MT was used to analyze 33% (13 out of 39) of the groups of replications [17]. In IPD-MT, the raw data of all the experiments are analyzed jointly as if the raw data came from one big experiment. Since IPD-MT depends on the availability of raw data, researchers typically first analyze each replication individually [23] to later perform IPD-MT by pooling and analyzing the raw data of all the replications applying the same statistical test that was used for performing the individual analyses.

IPD-S was used to analyze 15% (6 out of 39) of the groups of replications [17]. In IPD-S all experiment raw data are analyzed jointly by acknowledging the experiment where the raw data come from. As in AD, two types of IPD-S models can be fitted: fixed-effects models and random-effects models. Commonly used fixed-effects models are linear regression models (e.g., ANOVA) with two factors:

6. The respective percentage use of the aggregation techniques sum more than 100% because 16 groups of replications used more than one aggregation technique.

Experiment and Treatment [19]. Commonly used random-effects models are linear mixed models with two factors: Experiment and Treatment [19].

Aggregation of p -values was used to analyze 7% (3 out of 39) of the groups of replications [17]. In aggregation of p -values, *one-sided* p -values from all replications are pooled together by means of a statistical model such as Fisher's or Stouffer's method [22]. Note that p -values can be either available directly (have been previously reported) or computed by researchers (raw data available from each replication is first analyzed to calculate the p -values).

3 RESEARCH METHOD

We began our research by studying the recommendations and guidelines provided in medicine and pharmacology to analyze—and report—groups of replications (i.e., multicenter clinical trials [39]). We resorted to the medical and pharmacological literature because of their longstanding experimental tradition and because SE researchers have previously looked to these disciplines for advice on how to analyze individual experiments [5], [40], how to conduct systematic literature reviews [24] or how to conceptualize new research paradigms, such as evidence-based software engineering [41] and so on.

Particularly, we began studying the recommendations and guidelines promoted by the Cochrane Association [42], the American Food and Drug Administration [43], the guidelines for analyzing multicenter clinical trials (MCTs) provided by the International Conference on Harmonization [44], the PRISMA-IPD statement of the EQUATOR Network framework [45], and the CONSORT statement for reporting randomized controlled trials [46]. These guidelines for analyzing MCTs are mature and have been widely used. Some have been in use for over 20 years [44] and others have been referenced thousands of times [46].

The above guidelines contain a number of terms (e.g., multicenter, treatment-by-center interaction, etc.), and concepts (e.g., individual participant data, aggregated data, etc.) that we used to drive a subsequent literature review. Throughout the literature review, we came across numerous references to the statistical techniques that can be used to analyze MCTs [15], [25], [47]–[49], meta-analysis [19], [22], [50], hierarchical linear models [18], [51]–[53], and study protocols [54].

After studying the above references, we identified some differences between MCTs and groups of SE replications. According to the meta-analysis and hierarchical linear models literature that we examined, these differences had statistical consequences with respect to results aggregation. This led to another literature review where we discovered studies evaluating data under circumstances more typical of SE: small number of replications with small and unbalanced sample sizes. We found a number of studies from medicine [55], [56], social research [57], educational research [58] and econometrics [59] studying exactly this. In Section 4 we outline the differences between groups of SE replications and MCTs, and the statistical consequences of such differences.

After this second literature review, where we learned about the statistical consequences of the differences between groups of SE replications and MCTs, we revisited the groups

of replications that we identified during the SMS reported in [17]. We compiled a list of four common limitations with regard to joint data analysis in groups of SE replications. We outline these limitations in Section 5.

We developed guidelines to tackle these limitations. We created a four-step analysis procedure with an identical structure to those commonly followed in medicine [42]–[46]. We embedded the guidelines within Steps 3 and 4 of the analysis procedure that we propose (i.e., providing either joint conclusions, or moderator effects, respectively). Before going any further, however, we should clarify that we do not aim to propose an all-encompassing cookbook procedure to analyze all groups of SE replications. Our procedure can be seen as a set of minimum criteria needed to analyze a stereotypical group of SE replications (i.e., with the characteristics outlined in Section 2). In Section 6 we discuss the steps of the proposed analysis procedure, the objectives of each step, the medical guidelines recommending the respective steps, and how we adapted each step to SE by acknowledging differences between MCTs and groups of SE replications, and their common limitations with regard to joint data analysis.

Finally, we outline each of the steps of the proposed procedure. We apply the procedure to analyze a representative group of SE replications to illustrate the procedure.

The chosen group of replications focuses on test-driven development (TDD). One *research question* drives the group of replications: How does TDD affect quality compared to Iterative-Test-Last (ITL)?

The main *independent variable* across all replications is the development approach, with TDD and ITL as treatments. ITL is defined as the reverse-order approach of TDD (following Erdogmus et al. [60]).

All experiments have an identical *experimental design*: an AB repeated-measures design [5] (where subjects first apply ITL, and then TDD at a later date). The *dependent variable* within the group of replications is quality. We measured quality as the percentage of test cases that successfully pass from a battery of test cases that we built for measuring participant solutions. Specifically, we measured quality as follows:

$$QLTY = \frac{\#Test\ Cases(Pass)}{\#Test\ Cases(All)} * 100\%$$

A total of *four replications* were run: three at a multinational online security products company (i.e., F-Secure H, F-Secure K and F-Secure O), and one at UPV, a Spanish university. Six, 11, 7 and 33 subjects participated in each replication, respectively.

The characteristics of this group of replications are typical in SE: a small number of replications (i.e., four replications, the median number of replications within SE groups) evaluating the performance of a binary treatment (i.e., ITL vs. TDD) on a continuous⁷ outcome of interest, with identical response variables and experimental designs, small and dissimilar sample sizes, different types of subjects (i.e., professionals and students) with common knowledge (i.e.,

7. Although the data are measured on a percentage scale (i.e., 0 to 100%), the approach taken throughout this study is to consider the data as continuous as the total number of test cases is large (i.e., greater than 30 [61]).

week-long training on slicing, unit testing, ITL and TDD) and development culture (test last), using the same environment (i.e., Java, JUnit, Eclipse) and showing heterogeneous results.

4 DIFFERENCES BETWEEN GROUPS OF SE REPLICATIONS AND MCTS

After studying the guidelines and recommendations on analyzing MCTs in medicine and pharmacology, we were skeptical about their direct application for analyzing groups of SE replications due to some relevant differences between MCTs and groups of SE replications.

For example, MCTs use controlled experiments (participants are randomly assigned to treatments), while quasi-experiments (where assignment to treatments is non-random) are common in SE. Quasi-experimental designs usually create less compelling support for counterfactual inferences [62]. Quasi-experimental control groups may differ from the treatment condition in many systematic (non-random) ways other than the presence of the treatment. Many of these differences could be alternative explanations for the observed effect and should be ruled out by researchers in order to get a more valid estimate of the treatment effect. By contrast, with random assignment, researchers do not have to think about all these alternative explanations.

Additionally, MCTs tend to have detailed protocols specifying the experimental settings under which all the experiments are to be run and the set of procedures that are to be strictly adhered to during the execution of the experiments [19], [44], [63]. On the contrary, groups of SE replications are usually created ad hoc [17]. In fact, changes are usually made opportunistically across the replications which are then aggregated to either provide joint results or to investigate moderators. Unfortunately, the changes typically made across SE replications may result in an unexpectedly large variation of results (i.e., statistical heterogeneity of results [22]). In practical terms, if statistical heterogeneity materializes, then this is taken as evidence that the treatments may be performing differently across the experiments [22]. It may be misleading in this case to apply fixed-effects models to aggregate the results [19], [22]—as is typically the case in medicine [19], [43], [64]—. This is because, unlike random-effects methods, fixed-effects models provide a common effect rather than a distribution of effects as a joint conclusion [19], [22]. The joint conclusions of fixed-effects model may be especially misleading if results reverse across the experiments [19], [22]—as the averaged effect may not, ultimately, be representative for all the experiments.

We put the absence of protocols in groups of SE replications down to experimental research in SE being less mature than in medicine. Therefore, we regard this as a temporary difference since we expect SE researchers to be convinced by the advantages of developing standardized protocols (e.g., increased internal validity of results [39], [65]) and adopt them when conducting groups of replications.

Besides, stringent and random selection criteria are typically set to recruit the participants in MCT experiments designed to assess the efficacy of new treatments [44], [46],

[63], like specific blood pressure parameters, lack of comorbid conditions, etc. [44], [46], [63]. This ensures consistent results across sites and helps to minimize the risk of confounding effects impacting results [19]. Contrariwise, SE replications rarely set stringent selection criteria for recruiting participants. Instead, participants are usually recruited using convenience sampling. Unfortunately, the different characteristics of the participants across the experiments may result in statistical heterogeneity. Once again, this is an obstacle to the application of fixed-effects models for analyzing groups of SE replications. We think that there are two grounds for the absence of strict recruiting criteria in groups of SE replications. First, SE experimental research is less mature and has not yet developed standardized measurement instruments to classify—and include/exclude participants—in SE experiments [66]. Second, there are differences between the domains of SE and medicine, where SE researchers rarely have the luxury of dismissing participants or an ample array of potential participants. Since we do not expect this to change in the short term, we consider this difference as permanent.

Also, MCTs commonly undertake a planning phase where both participant-level and experiment-level sample sizes are calculated [43], [46], [63]. Participant-level sample sizes define how many subjects are needed, whereas experiment-level sample sizes define how many replications are needed if it is only plausible to allocate X subjects to each experiment. This ensures balanced sample sizes across the experiments and proper statistical power for detecting true population effect sizes. On the contrary, sample size estimation phases are rarely undertaken in SE (considering that only one group of SE replications [67] provided any sample size requirements calculation [17]). Instead, a small number of replications with small and dissimilar (i.e., convenient) sample sizes are usually run and then aggregated. This sample size estimation phase is feasible within the broader population to which medicine interventions apply as opposed to SE experiments where the population is more restricted. This more contrived sampling frame may prevent groups of SE replications from satisfying statistical power requirements. This places several limitations on the use of fixed-effects models to analyze groups of SE replications. First, fixed-effects models fit many parameters. For example, parameter estimates may be potentially biased in ANOVA models including an experiment factor, where a different parameter is fitted for each experiment [19], due to experiment-level sample size limitations. Second, groups of SE replications tend to have dissimilar sample sizes. This may prevent fixed-effects models achieving the statistical power to detect true treatment effects [55], [68]. We think that the failure in groups of SE replications to pay attention to sample size calculations may be due both to SE experimental research being less mature than medicine and to different participant recruitment opportunities between the SE and medicine domains. Therefore, it is regarded as a permanent difference.

Finally, the small sample sizes and number of replications in groups of SE replications also impact the detectability of moderators [19], [47], [69]. In particular, larger sample sizes are usually required to detect moderators than to detect treatment effects [69]. Therefore, it may not be feasible

TABLE 1
Differences between MCTs and groups of SE replications and statistical consequences.

Multicenter Control Trials	groups of SE replications	Statistical consequence
✓ Identical experimental configurations	✗ Opportunistic changes across replications	- Risk of heterogeneity
✓ Rigid & random participant selection criteria	✗ Convenience sampling	- Risk of heterogeneity
✓ Balanced & adequate sample sizes	✗ Unbalanced & small sample sizes	- Low precision & power of fixed effects
✓ Appropriate overall sample size	✗ Small overall sample size	- Inability to detect moderators

in groups of SE replications to get p -values lower than 0.05 in order to claim that there are statistically significant moderator effects. This is especially worrying in the case of experiment-level moderators, as, in most cases, only a few data points are available for moderator detection. It may not be feasible to identify moderators in groups of SE replications unless statistical significance thresholds are adapted (e.g., by increasing them from 0.05 to 0.10 [19]). Unfortunately, this comes at the cost of a larger proportion of statistical errors [19], [70]. In our opinion, the inability of groups of SE replications to detect moderators may be due both to SE experimental research being less mature than medicine (as, after all, moderators could be identified if sample size calculations were made [71]) and differences between SE and medicine (again, in terms of resources). Therefore, it is regarded as a permanent difference.

Table 1 summarizes the differences between MCTs and groups of SE experiments, and the statistical consequences of such differences for joint data analysis.

5 LIMITATIONS OF GROUPS OF SE REPLICATIONS

We designed an analysis procedure that is identical to the steps followed in medicine and pharmacology to analyze and report MCTs [42]–[46], [54]. We adapted this procedure to groups of SE replications taking into account their typical characteristics in order to overcome common limitations with regard to joint data analysis. After revisiting the groups of SE replications that we identified in our SMS [17], we came up with a list of four major limitations regarding joint data analysis practices. In the following, they are reviewed one by one.

Limitation 1: Fifty-three percent of the groups of SE replications use either narrative synthesis or aggregation of p -values to aggregate replication results [17]. Even though we agree with the use of narrative synthesis and aggregation of p -values when the raw data and summary statistics are unavailable or when response variables are incompatible [21], [72], we are skeptical about their use when the raw data are available and the replications have identical designs and response variables [17]. In the last analysis, access to the raw data may offer the possibility of providing more informative joint conclusions than just a textual summary of results (narrative synthesis) or a joint p -value (aggregation of p -values).

Limitation 2: Thirty-three percent of the groups of SE replications were analyzed by means of IPD-MT [17]. This technique may provide misleading results if participants are more similar within replications than between replications (e.g., when the replications are either run with professionals or with students), or if sample sizes are unbalanced across

the treatments and/or replications (e.g., if the replications have different sample sizes and there are missing data).

Limitation 3: Thirty-eight percent of the groups of SE replications were analyzed by means of AD with standardized effect sizes (such as Cohen’s d or the Pearson correlation) [17]. Even though AD with standardized effect sizes can be used to aggregate experiment results in systematic literature reviews (as access to summary statistics or to standardized effect sizes may be guaranteed), we question its use *alone* when the raw data are available and replications have identical response variable operationalizations⁸. This is because standardized effect sizes overlook the response variable scales, and, thus, may affect the interpretability of joint conclusions. For instance, how relevant is a joint Cohen’s d of 0.3? On the contrary, if the replications have identical response variable operationalizations and access to the raw data is guaranteed—as is typically the case in groups of SE replications [17]—, it may be possible to apply IPD-S [13]–[15] and, thus, interpret results in natural units. This practice has already been applied in SE [26], [75], [76] and can lead to more informative joint conclusions.

Limitation 4: SE researchers rarely acknowledge the limitations of the exploratory analyses that they undertake for identifying moderators. Besides, they usually identify moderators textually (e.g., “as the results are ‘statistically significant/positive’ in Experiment 1 and not in Experiment 2, this difference between the results could be due to moderator variable X ” [17]).

6 PROCEDURE FOR ANALYZING GROUPS OF REPLICATIONS

We propose the adoption of a four-step procedure to analyze groups of SE replications.

Step 1. Describe the participants. We propose to start by describing the participants of the replications. The objectives of this step are not only to *describe the population* to which the results should be generalized, but also to *suggest plausible sources of heterogeneity* that may arise when providing joint conclusions [44], [46]. This step can be further broken down into two main activities:

- As typical in MCTs, we propose to start by providing *summary statistics* to describe the characteristics of the participants [44], [46].
- We adapt this step to SE by rounding out the summary statistics with a *profile plot* [77] showing the

8. Groups of SE replications seldom justify the selected aggregation technique [17]. Nevertheless, there appears to be evidence that AD and stratified IPD tend to provide similar results with regard to the provision of joint conclusions [14], [73]. The identification of participant-level moderators, however, is a different matter, where stratified IPD comes out on top [74].

characteristics of the participants averaged across the replications (see Figure 1, Section 7).

Step 2. Analyze individual replications. We propose to pre-process, describe and analyze the data of each replication with consistent statistical techniques. The objectives of this step are to *provide descriptive statistics* to ease the incorporation of results into prospective studies (e.g., by facilitating the recalculation of effect sizes [22]), *identify patterns across replication results*, and ensure that *statistical heterogeneity is not introduced by the different methods* used to analyze the replications [78]. This step can be further broken down into three main activities:

- As in MCTs [44]–[46], [54], we propose to provide *summary statistics and visualizations* (e.g., box plots or violin plots) to describe the data of each replication and use consistent pre-processing steps to remove outliers or replace missing data [79], [80].
- We adapt this step to SE by rounding out the summary statistics and box plots with a *profile plot* [77] showing the mean of the treatments across the replications (see Figure 3).
- We analyze each replication with *consistent analyses* (e.g., *t*-test, ANOVA, etc. [44], [46], [54]). This ensures consistency of results across the replications and eases the integration of results in later phases.

Step 3. Aggregate the results. Following analysis guidelines for MCTs [44]–[46], the results of the individually analyzed replications are aggregated to arrive at *joint conclusions*. The objective of this step is to *increase the reliability of joint conclusions*. We adapt this step to SE by proposing three guidelines, each specifically tailored to address Limitations 1-3 discussed in Section 5:

- Guideline 1 draws upon arguments from groups of data analysis experts in mature experimental disciplines [43]–[45], [63] and the latest recommendations provided by statistical reformers and associations [37], [81], [82] to suggest avoiding the use of narrative synthesis and aggregation of *p*-values to provide joint conclusions.
- Guideline 2 recommends avoiding IPD-MT by default. Identical advice has been already provided in mature experimental disciplines such as medicine and pharmacology [45], [49], [69], [83], [84].
- Guideline 3 draws on arguments from various resources regarding linear mixed models [18], [51], references comparing the performance of IPD-S and AD [13], and articles comparing the performance of various IPD-S models [85] to encourage the use of both AD and IPD-S *in tandem* to analyze groups of SE replications. A similar recommendation has already been provided in mature experimental disciplines such as medicine and pharmacology [73], [86].

We also adapt this step to SE by proposing the use of random-effects models (rather than the fixed-effects models typically used in MCTs [19], [43]) in the two activities into which this step is divided: *apply AD* and *apply IPD-S*. Similar advice to this has also been given under similar circumstances in other disciplines such as the social sciences [22], [87], [88].

Step 4. Conduct exploratory analyses. As in MCTs [45], [46], [54], [86], [89], exploratory analyses should be conducted after providing joint conclusions. The objective of this step is to *identify experiment-level and participant-level moderators* that may be behind the statistical heterogeneity commonly present in groups of SE replications. We adapt this step to SE by developing three new guidelines to address Limitation 4 discussed in Section 5. To do this, we rely on the recommendations provided in references on data analysis in the social sciences, biology, and medicine:

- Guideline 4 provides guidance on how to identify experiment-level moderators by means of AD and IPD-S [12], [47], [70], [83], [89], [90].
- Guideline 5 provides guidance on how to identify participant-level moderators by means of IPD-S [47], [70], [74], [83], [89].
- Guideline 6 outlines the limitations of exploratory analyses [19], [37], [45], [46], [70], [86], [89].

We also adapt the procedure for identifying moderators to SE by suggesting an increase in the statistical significance threshold from 0.05 to 0.10—at the greater risk of a larger proportion of statistical errors [19]. We also suggest that less attention be paid to *p*-values, evaluating instead the relevance of moderator effect sizes—and their respective 95% CIs. These recommendations account for the typically low number of replications and small sample sizes of groups of SE replications, which are an obstacle to moderator detection [17]. The latter two adaptations are used again in the first two activities of this step: *identify experiment-level moderators* and *identify participant-level moderators* and acknowledged in the last activity of this step, *acknowledge limitations of exploratory analyses*.

Table 2 summarizes the steps of the procedure that we propose for analyzing groups of SE replications, the objectives of each step, and their associated activities. Table 3 links each of the proposed steps with the guidelines from medicine recommending the respective step, and the adaptation that we made for SE.

In the next four sections we outline the four steps of the procedure that we propose for analyzing groups of SE replications. For illustrative purposes, we apply the respective step to analyze the stereotypical group of replications described in Section 3. As Step 1 (i.e., describe the participants) and Step 2 (i.e., analyze individual replications) require no further explanation, we merely apply the steps to analyze the illustrative group of replications in Sections 7 and 8. As Step 3 (i.e., aggregate the results) and Step 4 (i.e., conduct exploratory analysis) embed a set of guidelines that require further explanation, we first develop the guidelines and then go on to illustrate their application to the group of replications in Sections 9 and 10.

7 STEP 1: DESCRIBE THE PARTICIPANTS

Activity 1.1. Provide summary statistics. The provision of summary statistics of the characteristics of the participants offers information about the population to which the results are to be generalized.

Example. Table 4 shows the means and standard deviations of participant programming, Java, unit testing and

TABLE 2
Procedure for analyzing groups of SE replications and objectives.

Step	Objectives	Activity
1. Describe participants	- Inform about the population under assessment - Hypothesize on possible sources of heterogeneity	1.1. Provide summary statistics 1.2. Provide profile plot
2. Analyze individual replications	- Ease incorporation of results into prospective studies - Identify patterns across replication results - Avoid heterogeneity due to different analysis procedures	2.1. Provide summary statistics and visualizations 2.2. Provide profile plot 2.3. Perform consistent individual analyses
3. Aggregate results	- Maximize informativeness of joint conclusions	3.1. Apply AD 3.2. Apply IPD-S
4. Conduct exploratory analyses	- Identify experiment-level moderators - Identify participant-level moderators	4.1. Identify experiment-level moderators 4.2. Identify participant-level moderators 4.3. Acknowledge limitations of exploratory analyses

TABLE 3
Mapping of proposed procedure steps to references in medicine and adaptations made for SE.

Step	Recommended in...	Adaptation to SE	Adapted from...
1. Describe participants	[44], [46]	✓ Provide profile plot	[77]
2. Analyze individual replications	[44]–[46], [54]	✓ Provide profile plot	[77]
3. Aggregate results	[43]–[46], [49]	✓ Avoid narrative synthesis & aggregation of p-values ✓ Avoid IPD-MT ✓ Use AD & IPD-S in tandem ✓ Use random-effects models	[37], [63], [81], [82] [69], [83], [84] [13], [18], [51], [73], [85], [86] [19]
4. Conduct exploratory analyses	[45]–[47], [54], [86], [89]	✓ Use AD & IPD-S to assess experiment-level moderators ✓ Use IPD-S to identify participant-level moderators ✓ Acknowledge limitations of exploratory analyses ✓ Increase statistical threshold ✓ Evaluate effect size and 95% CI	[12], [70], [83], [90] [70], [74], [83] [19], [37], [70] [17], [19] [17]

JUnit experience (measured by self-assessment as inexperienced, novice, intermediate and expert) across the replications. We do not find any clear patterns for averaged participant-level characteristics (e.g., averaged experience levels for F-Secure O and F-Secure H participants alternate between programming and Java) across replications.

TABLE 4
Descriptive statistics for participant characteristics.

Experiment	Prog.	Java	Unit	JUnit
F-Secure H	3.67 (0.52)	2.33 (1.21)	2.17 (0.98)	2.17 (1.17)
F-Secure K	2.91 (0.70)	1.82 (0.87)	1.64 (0.5)	1.27 (0.47)
F-Secure O	3.29 (0.76)	2.71 (1.11)	2.71 (0.76)	2 (0.82)
UPV	2.36 (0.57)	1.88 (0.60)	1.04 (0.20)	1 (0)

Activity 2.1. Provide profile plot. The provision of a profile plot showing the mean experience of the participants across replications may enhance the understandability of the summary statistics, help to convey the variability of participant characteristics across replications, ease the identification of patterns in the characteristics of the participants across experiments, and help to identify potential sources of heterogeneity.

Example. Figure 1 shows the profile plot of the illustrative group of replications. Figure 1 indicates that there is an observable decreasing trend in averaged participant experience across replications: participants have relatively more experience with programming and Java than with unit testing or JUnit across all the replications. There appears to be a noticeable heterogeneity of averaged participant experience. This may result in statistical heterogeneity when providing joint conclusions.

Fig. 1. Profile plot for subject experience.

Summary of example. A heterogeneous group of developers participates within the group of replications, with the most senior developers at F-Secure H and F-Secure O, and the junior at UPV.

8 STEP 2: ANALYZE INDIVIDUAL REPLICATIONS

Activity 2.1. Provide summary statistics and visualizations. The summary statistics and box plots provide information about the distribution of the data, facilitate the incorporation of results into prospective studies, and minimize the heterogeneity of results due to the application of different analysis techniques.

Example. Table 5 shows the descriptive statistics for QLTY with ITL and TDD across all replications. The respective box plots—and violin plots—are shown in Figure 2.

As Figure 2 shows, TDD appears to outperform ITL in all replications. However, while the difference in performance

TABLE 5
Descriptive statistics for QLTY: ITL vs. TDD.

Experiment	Treat.	N	Mean	Corr	SD	Median
F-Secure H	ITL	6	30.71	0.59	36.58	24.16
	TDD	6	40.23		33.43	35.34
F-Secure K	ITL	11	22.17	0.42	20.44	17.98
	TDD	11	35.42		35.40	22.41
F-Secure O	ITL	7	16.05	0.52	20.81	7.87
	TDD	7	68.97		31.53	81.03
UPV	ITL	31	33.38	0.47	39.79	6.74
	TDD	29	77.16		21.04	83.93

Fig. 2. Box plot and violin plot: ITL vs. TDD.

between ITL and TDD is small for F-Secure H and F-Secure K, the difference appears to be larger for F-Secure O and UPV. Noticeable—and similar—correlations appear to have materialized among the QLTY scores of the participants across replications (i.e., correlations around 0.5). This will result in greater statistical power (i.e., smaller effect size variances) when analyzing each replication and calculating their respective effect sizes [37], [91]. Finally, some data points (at the bottom of the distributions) for F-Secure O or UPV may be considered outliers. However, due to the already small sample sizes of the replications and missing data for UPV (two participants have data for TDD only, and another two have none), we do not remove any potential outlier from the data analysis.

Activity 2.2. Provide profile plot. A profile plot to complement the descriptive statistics provides a bird's eye-view of the data and helps identify patterns in results.

Example. A profile plot showing the mean QLTY score per treatment across replications is provided in Figure 3.

Fig. 3. Profile plot: ITL vs. TDD.

As Figure 3 shows, TDD appears to outperform ITL across all replications as observed in the violin plot above. The extent to which TDD outperforms ITL varies widely

across replications (see the different slopes of the lines). The different slopes indicate that *there may be heterogeneity in the group of replications*. Besides, there is no apparent pattern between ITL and TDD mean QLTY scores: the larger improvements with TDD over ITL (see the lines with the highest slopes) are achieved in the replications with the lowest and highest ITL mean scores (i.e., UPV and F-Secure O replications, respectively). By chance, they are the replications with the most novice and senior developers, respectively. Thus, *we cannot hypothesize, in principle, on any moderator that may be impacting results*.

Activity 2.3. Perform consistent individual analyses.

An analysis of the replications with consistent statistical methods ensures that differences across experiment results are not due to the use of different analysis procedures.

Example. Since all replications have an identical AB repeated-measures experimental design [5], we analyze each of them with a dependent *t*-test [92]. Table 6 shows the results of the individual *t*-tests performed.

TABLE 6
Individual analyses: ITL vs TDD.

Experiment	Estimate	95% CI	<i>p</i> -value
F-Secure H	9.52	(-19.58, 38.62)	0.483
F-Secure K	13.26	(-7.26, 33.77)	0.193
F-Secure O	52.91	(30.44, 75.39)	<0.001
UPV	42.31	(29.02, 55.62)	<0.001

As Table 6 shows, TDD outperforms ITL in all replications. Besides, the difference in performance between TDD and ITL is large and statistically significant for F-Secure O and UPV but not for F-Secure H and F-Secure K.

Summary of example. TDD outperforms ITL in all replications. However, the extent to which TDD outperforms ITL seems largely dependent upon site.

9 STEP 3: AGGREGATE THE RESULTS

In Sections 9.1, 9.2 and 9.3, we outline the three guidelines that we propose to overcome the most common limitations of groups of SE replications when providing joint conclusions. The description includes its application to the illustrative group of replications.

9.1 Avoid narrative synthesis and aggregation of *p*-values

9.1.1 Perils of narrative synthesis and aggregation of *p*-values

Although *p*-values are commonly used to evaluate the statistical significance of results, numerous criticisms have been made with respect to their inappropriate use across the sciences [93], [94]. The dichotomization of evidence possibly arising as a result of the indiscriminate use of statistical thresholds (such as 0.05 [94]), and the inability of *p*-values to convey the relevance of results (because they confound sample size and effect size [93]) are just two well-known criticisms of *p*-values [93]. As effect sizes and 95% CIs can also be used to assess the statistical significance of results (i.e., if the 95% CI of the effect size does not cross 0, then results are statistically significant), some authors have

suggested that effect sizes and 95% CIs should be used instead of p -values [37], [81], [82].

Bearing this in mind, neither narrative synthesis nor aggregation of p -values appear to be suitable for providing joint conclusions in groups of SE replications: narrative synthesis yields neither an effect size nor a p -value (it merely provides a textual summary of results), whereas aggregation of p -values provides a joint p -value but not an effect size.

Narrative synthesis and aggregation of p -values have another shortcoming: narrative synthesis weights each replication subjectively, while aggregation of p -values weights each replication within the joint conclusion identically [22]. Both types of weighting may be undesirable in groups of replications with different sample sizes. For example, larger (in principle, more precise) replications may have a greater weight than small replications within the joint conclusion, industrial (in principle, more representative) replications may have a greater weight than academic replications, and higher quality experiments may have a greater weight within the joint conclusion, etc.

Finally, narrative synthesis has another relevant shortcoming when providing joint conclusions. Very often non-significant results lead to a joint statistically significant result [22]. However, the joint conclusion of narrative synthesis would be non-significant if there are more non-significant results than significant results (as non-significant is the winner [22]). Narrative synthesis has been known since the 1980s to have low statistical power [22], [95].

9.1.2 Application to the illustrative group of replications

Narrative synthesis is simply applied by providing a textual summary of results of the replications (i.e., their effect sizes and p -values [22]).

To apply narrative synthesis to the illustrative group of replications, the procedure is as follows: "...even though TDD outperformed ITL in all the replications, the extent of such outperformance was largely dependent upon site. Besides, the difference in performance between TDD and ITL was statistically significant only for F-Secure O and UPV. Thus, conflicting results materialized in terms of statistical significance: two replications provided non-significant results, while two others provided significant results. As an identical number of replications point in opposite directions—i.e., non-significant vs. significant, *no final claims* can be made about the statistical significance of results. More replications are needed to argue the statistical significance, and practical relevance of results."

Aggregation of p -values procedures typically involve [19], [22]: (1) the individual analysis of each replication with a one-sided statistical test; and (2) the later combination of the resulting p -values by a statistical technique like Fisher's method.

We first analyzed each replication independently by means of a *one-sided* dependent t -test [92]. Then, we used Fisher's method [22] to pool together the p -values of all the replications. The result is a statistically significant difference between TDD and ITL as a joint conclusion ($\chi^2=47.13$; $df=8$; $p < 0.001$). Thus, the difference in performance between TDD and ITL is statistically significant in at least one replication. However, this was already known before aggregating

the results (as F-Secure O and UPV's results were already statistically significant).

Summary of example. Neither narrative synthesis nor aggregation of p -values provide informative joint conclusions. Narrative synthesis fails to provide a joint effect size or p -value and is not able to provide final claims since there are two significant results versus two non-significant results in our example. Aggregation of p -values fails because it provides a joint conclusion that was already known before results aggregation.

9.1.3 Guideline 1: Avoid narrative synthesis and aggregation of p -values

Avoid narrative synthesis and aggregation of p -values to provide joint conclusions.

What impact may this guideline have on the findings of joint analyses of groups of SE replications? More informative joint conclusions could have been obtained for 53% of the groups of replications (i.e., groups that applied either narrative synthesis or aggregation of p -values, see Section 5). Not applying weak aggregation techniques should enhance the findings of groups of SE replications.

9.2 Avoid IPD-MT

9.2.1 Perils of IPD-MT

IPD-MT should be avoided on two grounds. First, it may be *underpowered* compared to an identical IPD-S model including a factor accounting for the experiment [55], [96]. In other words, IPD-MT may provide a statistically non-significant result when it should be statistically significant. Second, IPD-MT may provide *biased results* [69], [97] when data are unbalanced across treatments and replications (which may be the case in groups of replications with missing data and different sample sizes) and subjects are more similar within, than between, replications (which may be the case when either professionals or students participate in the replications). Here we illustrate the perils of IPD-MT with an intuitive extreme example where it provides a biased result. Like Kraemer's example to illustrate the perils of IPD-MT [69], we produce our example by means of simulation [37].

Particularly, let us simulate two hypothetical replications comparing the performance of two technologies (e.g., Technology A vs. Technology B) on a continuous outcome of interest (e.g., quality). For simplicity's sake, let us suppose that the replications have an identical experimental design: an AB between-subjects design (i.e., a design where each participant is assigned to either Technology A or B). It is straightforward to simulate a group of replications with such characteristics using random draws from the data distributions that simulate the quality scores achieved with Technologies A and B across the replications [37]. Each random draw will represent the quality score achieved by a hypothetical (i.e., simulated) participant. SE data may follow a myriad of data distributions [38]. For illustrative purposes, we simulate the performance of Technologies A and B with normal distributions [37], although many other data distributions could have been used and have obtained the same results [37]. Table 7 shows the normal distributions that we use to simulate the performance of Technologies A and B across the replications, and the sample sizes of each

of the groups (i.e., the number of participants assigned to either Technology A or B) across the replications.

TABLE 7
Example: two simulated replications.

		Technology A	Technology B
Exp. 1	QLTY	$\mathcal{N}(20, 10^2)$	$\mathcal{N}(30, 10^2)$
	Sample Size	90	10
Exp. 2	QLTY	$\mathcal{N}(60, 10^2)$	$\mathcal{N}(70, 10^2)$
	Sample Size	10	90

As Table 7 shows, we aim to simulate two highly unbalanced replications (i.e., 90 subjects assigned to Technology A, and 10 to Technology B in Experiment 1, and vice versa in Experiment 2). Additionally, we aim to simulate a circumstance where the mean difference in performance between Technologies B and A is expected to be around 10 in both replications (i.e., 30-20 in Experiment 1 and 70-60 in Experiment 2), and the participants are more similar within, than between, replications (as they achieve either much larger or much smaller scores with either Technology A or B depending upon the replication in which they participate). These are the exact circumstances under which IPD-MT provides biased results.

As the difference in performance between Technologies B and A in both replications is around 10, we would expect the difference in performance to be similar for joint results.

We analyzed the data with both IPD-S and IPD-MT (i.e., ANOVA models that did not include Experiment as a factor, respectively). IPD-S provides an estimate close to the expected ($M = 9.55$). IPD-MT provides an estimate that deviates from the expected ($M = 41.17$). This is because IPD-MT does not take into account the experiment that is the source of the data and, instead, assumes that all the data come from a single “big” experiment. As such, IPD-MT is unaware that most subjects contributing towards the mean quality score with Technology A (90/100) come from Experiment 1 (with mean scores of 20), whereas most subjects contributing towards Technology B (90/100) come from Experiment 2 (with mean scores of 70). As a result, the unbalance of subjects across the treatments and the dissimilarities of participant scores across the replications distort IPD-MT results (i.e., by providing a much larger difference of results than expected). The larger the unbalance across treatments and/or sample sizes across the replications, the more biased IPD-MT results will be.

9.2.2 Application to the illustrative group of replications

To **apply IPD-MT**, the raw data of all replications are pooled together as if they come from one big experiment and then the same statistical model as used to analyze each experiment individually is applied [49], [69], [83].

We apply a dependent t -test to analyze the raw data of all the replications together. IPD-MT provides a joint estimate equal to $M = 33.67$ and a p -value < 0.001 . Thus, the difference in performance between TDD and ITL is large—insofar as QLTY ranges from 0 to 100—and statistically significant.

To **apply IPD-S**, raw data from all replications are pooled together, and then two factors —Treatment and

Experiment— are used [49], [69], [83]. In other words, a linear regression model (e.g., ANOVA) that takes into account the source of the raw data is fitted.

We applied an ANOVA with Treatment and Experiment as factors to analyze the raw data of the illustrative group of replications. IPD-S provides a joint estimate equal to $M = 34$ and a p -value < 0.001 . Thus, the difference in performance between TDD and ITL is large—insofar as QLTY ranges from 0 to 100—and statistically significant.

So, both IPD-MT and IPD-S provide similar results in the illustrative group of replications (i.e., a large and statistically significant result). This is because data are perfectly balanced within replications (as the replications are AB repeated-measures designs), missing data have a relatively low impact on results (as just a few subjects have missing data for UPV), and subject scores are similar across the replications—note that the mean scores for ILT are clustered around 25 (see Figure 3, Section 8). However, this cannot be guaranteed in all circumstances.

9.2.3 Guideline 2: Avoid IPD-MT

Avoid IPD-MT due to its potential to provide biased or underpowered results.

What impact may this guideline have on the findings of joint analyses of groups of SE replications? Thirty-three percent of the groups of replications (i.e., where IPD-MT was applied to provide joint conclusions, see Section 5) could have arrived at less biased and less underpowered joint conclusions. Not applying IPD-MT should enhance the findings of groups of SE replications.

9.3 Use AD and IPD-S in tandem for joint analysis

9.3.1 Benefits of using AD plus IPD-S

AD and IPD-S are complementary in some respects. While AD provides certain advantages for analyzing groups of SE replications, IPD-S provides others. **Some of the advantages of AD over IPD-S are:**

- AD can be used to analyze groups of replications with *different response variables* (e.g., by computing standardized effect sizes such as Cohen’s d [22]). On the contrary, IPD-S can only be applied whenever identical response variable scales are used [13], [19]. Thus, if response variables change across the replications, AD may be the only available option.
- AD provides *intuitive visual summaries of results* (i.e., forest plots) that have been commonly used in SE to synthesize the findings of experiments gathered by means of systematic literature reviews [24]. On the contrary, less standardized visualizations are available for IPD-S (e.g., 95% CI plots, error bars, etc. [37]). Thus, the appeal and familiarity of forest plots in SE is a plus for AD over IPD-S.
- AD is useful for *interpreting the heterogeneity of results with straightforward statistics and tests* (e.g., the I^2 statistic and the Q -test [22]). Besides, rules of thumb are also available for interpreting the I^2 statistic (i.e., 25%, 50% and 75% for small, medium, and large heterogeneity, respectively [22]). On the contrary, IPD-S may require either: (1) the standard deviation of

results to be contrasted against the joint result; or (2) fixed-effects models (such as ANOVA) to be used with Treatment by Experiment interaction terms to be able to claim that results are heterogeneous when the interaction is statistically significant [19]. However, we do not encourage the latter procedure in groups of SE replications, which are typically small. Therefore, this method of heterogeneity detection would be underpowered [19].

On the other hand, IPD-S has some advantages over AD:

- IPD-S can simultaneously assess the difference in performance between the treatment and the control group (like AD), as well as the performance of the control group in order to weight their relative size in natural units. For example, if the difference in performance between the means of the treatment group and the control group is equal to 20, and the mean performance of the control group is equal to 20, then the treatment doubles the performance of the control. On the contrary, AD commonly relies on standardized effect sizes (e.g., Cohen's d [22]) to convey the difference in performance between the treatment and control. This may affect the interpretability of results: how relevant is a Cohen's d of 0.3?
- IPD-S can simultaneously assess the effect of multiple factors and their interactions on results (e.g., the effects of the treatments, the tasks, and their interaction in ANOVA models [5]). On the contrary, AD is commonly used to perform pairwise comparisons between treatments (e.g., Treatment A vs. Treatment B [22]). Thus, IPD-S is more flexible than AD for analyzing groups of replications when multiple factors are of interest or the results depend upon interaction terms (e.g., when the effects of the treatment reverse depending upon the task being developed).
- Some IPD-S models such as LMMs can be used to analyze groups of replications with missing data—provided that the data can be assumed as missing at random [18]. On the contrary, the calculation of effect sizes—and their respective variances—using AD rests on the assumption of complete observations (otherwise, it would not be possible to compute the variances of some effect sizes for repeated-measures designs [22]). If there are missing data, researchers performing AD to calculate effect sizes and their respective variances may have to either exclude participants with missing data or rely on advanced imputation techniques for their inclusion (see Section 12). Thus, if there are drop-outs or protocol deviators across the replications, IPD-S models such as LMMs may come in handy [98].

Therefore, the application of both techniques in tandem takes advantage of the strengths of each one. Regarding the type of model to be used, both AD and IPD-S are statistical procedures that deliver a weighted average of experiment results as a joint conclusion [19], [22]. The weight—or contribution—of each experiment towards the joint conclusion is proportional to either the sample size of

the experiment—if a fixed-effects model is used—or to the sample size of the experiment and the statistical heterogeneity of results (i.e., the variation of results that cannot be explained by natural variation)—if a random-effects model is used [18], [22].⁹ Besides, if results are more heterogeneous, the weights of all the experiments within the joint conclusion will be more alike. Intuitively, as each experiment may be estimating a potentially different effect size when there is a large heterogeneity of results, smaller experiments are still informative about the distribution of effect sizes (as their effect sizes are also feasible). In turn, both small and large experiments tend to be regarded as being more equally informative in random-effects models (even though larger experiments have a slightly greater weight within the joint conclusion [22]).

As the heterogeneity of results is commonplace in SE experiments [99]–[103], many factors may have an impact on SE experiment results [11], and experimental changes, opportunistic recruitment of participants and different types of subjects (e.g., professionals vs. students) are typical in groups of SE replications (see Section 2), we recommend relying by default on random-effects models to provide joint conclusions.

Specifically, if using AD we suggest the use of random-effects meta-analysis models [22]. If using IPD-S, we suggest the use of linear mixed models (LMMs) [18].

9.3.2 Application to the illustrative group of replications

Activity 3.1. Apply AD. The application of AD requires calculating the effect sizes—and corresponding variances—of the replications from their summary statistics. They are then pooled using a random-effects meta-analysis model [22].

Example. First, we calculate Cohen's d s—and corresponding variances—of the replications from their summary statistics (i.e., sample sizes, means, standard deviations, and correlations between ITL and TDD [22]). As four subjects at UPV have missing data, their data have to be discarded—as they did not provide complete observations for the correlation between ITL and TDD [22]. Then, we pool together all Cohen's d s using a random-effects meta-analysis model [22]. Figure 4 shows the forest plot of the meta-analysis.

Fig. 4. Forest plot: ITL vs. TDD.

As Figure 4 shows, TDD outperforms ITL in all replications. Additionally, the joint effect size ($M = 0.90$) is large—according to rules of thumb [22]—and statistically significant (as the 95% CI does not cross 0). Besides, there is, according to rules of thumb, a medium ($I^2 = 67.8\%$)

9. Assuming a common variance across all replications.

heterogeneity of results [22]. Thus, moderators should be identified to explain the detected heterogeneity of results.

Activity 3.2. Apply IPD-S. IPD-S is straightforward to apply. It is sufficient to fit a LMM with two factors: Treatment and Experiment, considering Treatment as a random effect across the experiments [18], [19].

Example. To analyze the illustrative group of replications with IPD-S, we pool the raw data of all the replications together and then analyze them using a LMM. Table 8 shows the results of the LMM.

TABLE 8
LMM results.

Factor	Estimate	95% CI	p-value
ITL	27.44	(13.08, 41.79)	<0.001
TDD	56.27	(22.12, 90.42)	<0.001
M_{Diff}	28.83	(9.72, 47.93)	0.004
sd_{Diff}	16.09		

As Table 8 shows, the difference in performance between TDD and ITL is relevant ($M_{Diff} = 28.83$) and statistically significant ($p=0.004$). Looking at the difference in performance between TDD and ITL (i.e., M_{Diff}) and the effect of the control approach (i.e., ITL), we reach the conclusion that TDD doubles the performance of ITL (i.e., 56.27/27.44). Unlike AD, participants with missing data have been included to provide joint conclusions [18]. Finally, the standard deviation of the differences between TDD and ITL across the replications (i.e., sd_{Diff}) is relatively large compared with the overall difference (i.e., M_{Diff}): 16.09/28.83=0.56. This suggests that there is heterogeneity. Thus, moderator effects should be identified to explain the observed heterogeneity of results.

Summary of example. AD indicated that there is a large—and statistically significant—joint effect size with medium heterogeneity. Also, AD was able to visualize that TDD outperformed ITL across all the replications. IPD-S showed that TDD doubled the performance of ITL. It also meant that we could include missing data when providing joint conclusions and confirm the statistical significance of results observed with AD.

9.3.3 Guideline 3: Use AD and IPD-S

- Use *AD and IPD-S in tandem* to provide joint results. Use AD because of its intuitive visualizations (i.e., forest plots) and straightforward heterogeneity statistics. Use IPD-S because of its ability to convey joint results in natural units, and its flexibility for analyzing replications with missing data.
- Use *random-effects models* by default to provide joint conclusions. Particularly, use LMMs [18] for IPD-S, and random-effects meta-analysis models [22] for AD.

What impact may this guideline have on the findings of joint analyses of groups of SE replications? Adherence to this guideline may have potentially resulted in more intuitive joint conclusions for 38% of the groups of replications (i.e., groups that only applied AD with standardized effect sizes to provide joint conclusions, see Section 5).

10 STEP 4: CONDUCT EXPLORATORY ANALYSES

In Sections 10.3, 10.1 and 10.2, we outline the three guidelines that we propose to overcome the most common limitations of groups of SE replications for identifying moderators. The description includes its application to the illustrative group of replications.

10.1 Use AD plus IPD in tandem to identify experiment-level moderators

The identification of experiment-level moderators increases knowledge of software development. We suggest applying both AD and IPD-S in tandem to identify experiment-level moderators. The benefits of using AD plus IPD in tandem have already been discussed in Section 9.3.1. Therefore, we will not repeat our arguments here.

10.1.1 Application to the illustrative group of replications

Activity 4.1. Identify experiment-level moderators. To identify experiment-level moderators with AD, perform either a sub-group meta-analysis for *categorical* moderators or a meta-regression for *continuous* moderators. To identify experiment-level moderators with IPD-S, fit LMMs with interaction terms [18], [19].

Example. We run an AD sub-group meta-analysis to assess the effect of the type of subject (i.e., professionals vs. students) on results. Figure 5 shows the forest plot for the sub-group meta-analysis that we performed. Table 9 shows the result of the sub-group meta-analysis.

Fig. 5. Forest plot: professionals vs. students.

TABLE 9
Sub-group meta-analysis: professionals vs. students.

Group	N	Estimate	95% CI	I^2
Professionals	3	0.77	(-0.13, 1.68)	70.8%
Students	1	1.24	(0.72, 1.76)	0
Difference	-	0.47	(-0.58, 1.52)	-

As Table 9 shows, both professionals ($M = 0.77$) and students ($M = 1.24$) perform better with TDD than with ITL. However, the difference in performance between students and professionals ($M = 0.47$) is relevant—medium according to rules of thumb [22]. In other words, students appear to benefit more from TDD than professionals. However, despite the relevance of the moderator effect, there was a wide 95% CI. To be precise, the 95% CI ranges from a *medium* and *negative* effect (i.e., -0.58) to a *large* and *positive* effect (1.52). This results in a non-statistically significant

moderator effect due to the small number of replications analyzed. Thus, more replications are needed to increase the precision of experiment-level moderator effects.

Finally, we complement the results of AD with the findings of IPD-S. To do this, we run a LMM with interaction terms [18], [19]. Table 10 shows the results of the LMM that we performed.

TABLE 10
LMM experiment-level moderators: professionals vs. students.

Interaction	Estimate	95% CI	p-value
Type:Students	16.32	(-37.16, 69.55)	0.545

As Table 10 shows, the difference in performance between students and professionals with TDD appears to be large ($M = 16.32$)—at least compared with the difference in performance between TDD and ITL in the main analysis ($M = 28.83$). In view of this, students appear to perform around 60% (i.e. $16.32/28.83$) better than professionals using TDD. However, this should be further substantiated with more replications as the 95% CI ranges from negative to positive results (95% CI = $(-37.16, 69.55)$). Again, the group of replications is too small to detect experiment-level moderators.

Summary of example. Students appear to benefit more than professionals from TDD. But four replications with 6, 11, 7 and 33 subjects are not enough to detect the effect.

10.1.2 Guideline 4: Use AD and IPD-S to Identify Experiment-Level Moderators

Use AD and IPD-S in tandem to assess experiment-level moderators.

What benefit may this guideline have on joint analysis practices for groups of SE replications? Forty-two percent of the groups of replications (i.e., groups that adopted a textual approach to eliciting experiment-level moderators, see Section 5) could have achieved more transparent moderator effects in groups of SE replications. Using AD plus IPD in tandem to identify experiment-level moderators should enhance the findings of groups of SE replications.

10.2 Use IPD to identify participant-level moderators

New knowledge is also gained by identifying participant-level moderators.

10.2.1 Benefits of using IPD to identify participant-level moderators

IPD-S is better than AD at identifying participant-level moderators [47], [104]. This is because AD *may be underpowered* if the averaged participant characteristics do not vary much across the replications [15], [104] and subject to *ecological bias* when identifying participant-level moderators (i.e., the average effect may not be representative of the effect on the population) [15], [47], [105]. This may result in misleading conclusions. Thus, as is already common practice in medicine [47], [74], we recommend relying by default on IPD-S models to identify participant-level moderators.

10.2.2 Application to the illustrative group of replications

Activity 4.2. Identify participant-level moderators. To identify participant-level moderators with IPD-S, it is sufficient to fit LMMs with interaction terms [47], [74]. As Fisher et al. [47], [74] noted, special attention should be paid to separating the variance of moderator effects within and between experiments.

Example. We ran a series of LMMs with interaction terms to assess the effect of participant experience with programming, Java, unit testing or JUnit on results. Table 11 shows the results of the LMMs. Figure 6 shows the regression plot for the moderator effects.

TABLE 11
LMM participant-level moderators: participant experience.

Interaction	Estimate	95% CI	p-value
Programming	15.76	(0.49, 31.04)	0.04
Java	3.85	(-8.13, 15.83)	0.52
Unit testing	11.79	(-5.95, 29.54)	0.18
JUnit	11.07	(-6.26, 28.41)	0.20

Fig. 6. LMM interactions: participant-level moderators.

As Figure 6 shows, the more experienced participants are with programming, Java, unit testing or JUnit, the more TDD outperforms ITL. In other words, TDD appears to perform better for more experienced developers and vice versa. Additionally, as Table 11 shows, the significance level for programming experience was lower than 0.1—the significance threshold that we recommended for identifying moderators (see Section 10). Thus, in principle, *programming experience may be moderating the effects of TDD on quality*. However, this information should be regarded with caution because there were inflated Type I error rates (due to the execution of a total of five exploratory analyses) and a wide 95% CI interval (varying from 0.49, an almost negligible effect, to 31.04, a large effect). Thus, the results may be spurious. Finally, participant experience with programming may also be confounded with other variables (e.g., the age of the developers, their experience with Java, etc.).

Summary of example. Participant experience with programming appears to moderate TDD effects.

10.2.3 Guideline 5: Use IPD-S to Identify Participant-Level Moderators

Use IPD-S to identify participant-level moderators.

What benefit may this guideline have on the findings of joint analyses of groups of SE replications? Seventy-five percent of the groups of replications (i.e., groups that adopted

a textual approach to eliciting participant-level moderators or did not elicit participant-level moderators, see Section 5) could have detected more transparent moderator effects in groups of SE replications. The use of IPD to identify participant-level moderators should enhance the findings of groups of SE replications.

10.3 Acknowledge limitations of exploratory analyses

The limitations of exploratory analyses should be acknowledged.

10.3.1 Limitations of exploratory analyses

Exploratory analyses have three limitations:

- *They are unable to provide cause-effect relationships* because replications are *designed* exclusively to study the effects of the treatments on the response variables. Thus, it is impossible to establish the cause-effect relationships of other variables (e.g., moderators) on the response variables. In SE terms, it is risky to claim that the programming language is the reason for the different results detected using different programming languages across two replications. Note that other variables than the programming language, such as different participant characteristics, for example, may be the real cause for this difference in the results. If any cause-effect claims are to be made about moderators, experiments assessing such questions have to be undertaken beforehand. For instance, a new experiment where participants are randomly assigned to either one or other programming language could serve to identify whether the programming language is the cause of the effects on the results.
- *They increase the risk of committing statistical errors.* Many statistical analyses are typically run to identify moderators (e.g., one per moderator [15]), which inflates the Type I error rates. In other words, spurious statistical significant results may emerge out of multiple testing merely by chance. Such inflated Type I error rates may need to be corrected with *multiple comparison* correction procedures (such as the Bonferroni correction [70]). However, this is troublesome in groups of SE replications: on top of their already small sample sizes and the small number of replications (and, thus, poor moderator detectability), even more demanding statistical thresholds are set for identifying moderators. For instance, according to the Bonferroni correction, a statistical threshold of $0.05/3$ may be needed to detect moderators and assess three different moderators—one per analysis—. Because of the limitations of groups of SE replications in this regard, we recommend either: (1) setting a statistical threshold of 0.1 to identify moderators—despite the heightened probability of committing statistical errors, or (2) focusing more on the *magnitude* and *sign* of moderator effects—and their corresponding 95% CIs—rather than on their statistical significance.
- *Moderator effects can be confounded if multiple simultaneous changes are made across replications.* For example, if

both the programming language and the unit testing tool change simultaneously across two replications, it may be misleading to claim that differences across replication results are solely due to the programming language. The difference in the results may be due to the programming language, the IDE, or a mixture of both. Another typical case of confounding is to claim that differences across replication results are induced by the subject type when different types of subjects are evaluated across two replications (e.g., professionals vs. students) and the replications provide different results. In particular, other variables may also be behind the difference in results such as age (e.g., professionals may be older than students), motivation (e.g., students whose grades are at stake may be more motivated than professionals), treatment conformance (e.g., professionals may deviate from the procedure more than students [106]), some threats to validity may materialize in some replications and not in others (e.g., drop-outs, missing data, fatigue), etc.

Researchers running joint analyses of groups of replications should acknowledge the limitations of exploratory analyses in their papers.

10.3.2 Application to the illustrative group of replications

Activity 4.3. Acknowledge limitations of exploratory analyses. To acknowledge the limitations of exploratory analyses, it suffices to check through and adapt the list of limitations that we outlined in Section 10.3.1 to the group of replications—and moderators—that are to be investigated.

Example. We plan to investigate the effect of *one* experiment-level moderator (i.e., type of subject, students vs. professionals) and *four* participant-level moderators (i.e., experience with programming, Java, unit testing, and JUnit) on results. We acknowledge that none of the above moderators may be the real reason behind the detected heterogeneity of results, and other confounding variables may also be responsible for the detected difference in results. For instance, students appeared to be more motivated than professionals, students adhered more closely to the TDD process, professionals were older than students, etc. Also, we acknowledge that the chance of achieving spurious statistically significant results increases because we intend to run five data analyses (i.e., one per moderator, following Fisher's approach [47]) [70]. This may invalidate the conclusions reached. Thus, we will only use exploratory analyses as a way of motivating further research and never to draw definite conclusions [89].

Summary of example. The group of replications is too small for making definite claims about the results regarding subject type. Additionally, results for programming experience may be spurious due to the wide 95% CI that materialized. Finally, programming experience may be confounded with other variables (e.g., age of the participants, experience with Java, etc.).

10.3.3 Guideline 6: Acknowledge Limitations of Exploratory Analyses

Separate exploratory analyses from the main analysis and acknowledge their limitations.

What benefit may this guideline have on the findings of joint analyses of groups of SE replications? Thirty-eight percent of the groups of replications that did not perform exploratory analyses (see Section 5), and 88% of the groups of replications that did not acknowledge the limitations of exploratory analyses could have identified informative moderator effects in groups of SE replications. The performance of exploratory analyses and acknowledgment of their limitations should enhance the findings of groups of SE replications.

11 THREATS TO VALIDITY

We focused on the aggregation of quantitative results. What about qualitative results (e.g., text transcripts, etc.)? Throughout this article we focused exclusively on the aggregation of quantitative results into joint conclusions. We acknowledge that this limits the applicability of our guidelines. However, we decided to focus on quantitative results because SE experiments are usually coupled with the acquisition and analysis of this type of results [2], [3], [5] and most of the groups of replications uncovered by our SMS only aggregated quantitative results [17]. For an overview of the methods that can be used for aggregating qualitative—or qualitative and quantitative results—into joint conclusions, we refer interested readers to Cruzes and Dyba [107] and Kitchenham et al. [3]. They discuss meta-ethnography, narrative synthesis, qualitative cross-case analysis, thematic analysis, meta-summary, vote counting, grounded theory, content analysis, case survey, qualitative comparison analysis, aggregated synthesis, realist synthesis, meta-synthesis and meta-study. The aggregation of qualitative and quantitative results is discussed in [21], [108], [109].

We used only one analysis procedure. Are there any limitations to the way in which we tailored it to SE? Like [19], [22], we acknowledge that it is unfeasible to provide definite guidance on aggregation techniques and statistical models for use across the board—independently of the characteristics of the data or the intentions of the analyst (e.g., to provide joint results or to identify moderators)—. However, we tried our best to provide tailored statistical advice for aggregating the results of groups of SE replications considering their common characteristics and the limitations on joint data analysis based on the guidelines typically followed in medicine and pharmacology to analyze MCTs [43]–[45], [63] and our understanding of the statistical methods that can be used in circumstances that are typical of groups of SE replications—at least according to well-known references in medicine, pharmacology, and the social sciences [15], [19], [57]–[59]. We acknowledge that this is only a first approximation towards tailoring a definite analysis procedure and that further research is needed in order to provide more evidence on the suitability of the analysis procedure for analyzing groups of SE replications.

We gathered the references on data analysis opportunistically. Might not this introduce bias? Unfortunately, we could not systematically gather references on the topic of how to analyze MCTs—or groups of replications with characteristics typical in SE—: the number of articles retrieved from online databases with the terms “aggregation” and “experiments”

was unmanageable. Thus, we acknowledge that our guidelines may be open to bias. However, we made every effort to consult reliable resources on the topic of aggregation of experiment results from both mature experimental disciplines, such as medicine and pharmacology, and other areas, such as social research, education and econometrics. We also strove to embed guidelines providing the results not only of single aggregation techniques but also of different aggregation techniques applied in tandem (see the recommendation to use IPD-S and AD in tandem). This should reduce the potential bias that may have been introduced due to the non-systematic selection of references on data analysis.

We selected random-effects models over fixed-effects models. Are there not any limitations to this advice? Contrary to medicine, where fixed-effects models are encouraged for use by default [19], [43], [64], we recommend the use of random-effects models instead [19], [22], [87]. This is a controversial recommendation: some authors from other disciplines suggest that a minimum of five [49], ten [58], [110], fifteen or even more experiments [57], [58], [111] are needed to obtain reliable *variance* parameter estimates in random-effects models. However, we recommend the use of random-effects models by default as SE is commonly concerned with providing joint results (i.e., differences between *means*) rather than making inferences on variance parameters. Additionally, random-effects models tend to provide more conservative results than fixed-effects models (i.e., 95% CIs tend to be wider with random-effects models [19], [50], [65]), and random-effects models produce identical results to fixed-effects models when there is no heterogeneity [22]. Finally, sensitivity analyses assessing the robustness of results to the specification of fixed-effects models—rather than random-effects models—can also be run in groups of SE replications [22]. We refer interested readers to Thabane et al. [112].

We did not check the statistical assumptions of the statistical tests used (t-tests, LMMs). Is this not a limitation? As usual after analyzing the data of individual experiments [5], it is also necessary to check the statistical assumptions of the statistical tests used after aggregating the results [18], [51]. For example, if LMMs are fitted to analyze the data, then the normality assumption needs to be checked [19], [51]. We acknowledge that SE data may be not normal and concede that there are more advanced statistical methods for analyzing non-normal data (see Section 12). However, we resorted to *t*-tests and LMMs in this article because they are robust to departures from normality [113], [114], especially with larger sample sizes—as is the case when the raw data of all the replications are pooled together— [115]. In any case, as model diagnostics procedures are standard across the disciplines, we refer interested readers to specialized literature on the topic [18], [19], [51], [116].

We use only one illustrative group of replications. Is the analysis procedure applicable in other cases? For reasons of space, we illustrated the application of the analysis procedure to only one group of replications. We admit that this places some limitations on the generalizability of the analysis procedure. However, we tried to select what is, according to the results of a previous SMS addressing this issue, a representative group of SE replications [17]. Accordingly, we expect our guidelines to be applicable to the analysis of a large per-

centage of groups of SE replications. Additionally, we also provide further references in Section 12 indicating how to analyze groups of replications with different experimental designs and data characteristics.

12 ALTERNATIVE EXPERIMENTAL DESIGNS

As typical in groups of SE replications, this article analyzes a group of replications where all the replications have an identical experimental design: an AB *within-subjects* design (i.e., a design where the participants apply first Treatment A and then Treatment B in a later session [5]). It is straightforward, based on our procedure, to analyze groups of replications where all the experiments have an identical AB *between-subjects* design (i.e., a design where the participants are randomly assigned to either Treatment A or B): (1) for IPD-S, it is sufficient to remove the repeated-measures structure at participant level [19]; (2) for AD, it is sufficient to adapt the effect size variance formulae to the between-subjects design [22]. Besides, as groups of replications with AB between-subjects designs are commonplace in medicine, there is no shortage of references indicating how to analyze such designs with both IPD-S [15], [19], [47] and AD [22].

Whenever groups of replications contain a mixture of AB between-subjects designs and AB within-subjects designs, researchers can still use LMMs with the IPD-S approach to provide joint conclusions—as LMMs can account for replications with missing data, provided that they are missing at random (e.g., when it is possible to tell from the experimental design whether the subject data will or will not be missing in a specified experimental session [97], [117], [118]). AD can also be used for analyzing mixtures of AB between-subjects designs and AB within-subjects designs [91]: it is sufficient to calculate a consistent pooled standard deviation for standardizing Cohen’s *d* (typically the average standard deviations of Treatments A and B) and then select the appropriate variance formulae depending upon the experimental design (i.e., a within-subjects or between-subjects design [91]). As experiments with different experimental designs may be estimating different true effect sizes [91], however, exploratory analyses investigating the difference of results across experimental designs [91], [112] (e.g., by means of a sub-group meta-analysis [22]) should be used.

Throughout this article, we analyzed experiments whose response variable was measured on a *continuous scale*, which is typical in SE [17]. Also, we relied on the statistical tests that we ran being robust to departures from normality [113], [114]. Still, researchers may question the reliability of their inferences if data largely depart from the normality assumption [38]. If this is the case, researchers may resort to data transformation (e.g., Box-Cox transformations) to make the normality assumption hold and then apply IPD-S or AD to provide joint conclusions [70]. Researchers may also resort to more advanced statistical techniques such as bootstrapping to provide inferences with both IPD-S [119], [120] and AD [121]. Finally, researchers applying AD can also use non-parametric effect sizes such as Cliff’s delta to provide joint conclusions [38]. Eventually, if data do not meet the homogeneity of variances assumption, researchers can opt for generalized least squares (GLS) models under

the IPD-S umbrella (because they can accommodate different variance terms per treatment [122]) or Glass’s delta—rather than Cohen’s *d*—under the AD umbrella [123].

Also, response variables may be measured on a non-continuous scale (e.g., binary, count, etc. [70]). In this case, generalized linear mixed models (GLMMs) may be more appropriate than LMMs for analyzing the data with IPD-S [70]. We refer interested readers to Zuur et al. [122], [124] for an accessible introduction to the topic using the R programming language. Pautz et al. [125] and Fritz et al. [126] provide illustrative examples of how to calculate effect sizes—to be later combined with AD—for non-continuous response variables.

In multilevel data structures typical of groups of SE replications—where subjects are nested within replications and subjects can be measured several times (once or more per treatment) throughout the experiment—, data are typically correlated within *clusters* (i.e., clusters of both replications and participants [19]). This correlation needs to be taken into account when analyzing the data (i.e., by including the clustering units as either fixed factors or as random factors with the selection of appropriate variance-covariance matrices [19]). Throughout this article, we relied on random factors—for both participants and replications—and the variance-covariance matrix used by the *lme* R function when fitting LMMs: the *unstructured* variance-covariance matrix [127]. With this variance-covariance matrix, both the treatment and control (e.g., ITL in the illustrative group of replications) effects are assumed to be correlated across the replications [127]. If this does not hold, however, other variance-covariance matrices may be more suitable for analyzing the group of replications (e.g., by assuming independent treatment and control approach effects). For an accessible introduction to the topic, we refer interested readers to Finch et al. [127] and Zuur et al. [122].

Finally, missing data can materialize in SE experiments due to drop-outs or protocol deviators. Under these circumstances, researchers may resort to LMMs under the IPD-S umbrella to aggregate results (LMMs can be used to analyze groups of replications with missing data as long as data are missing at random [98]). They may also rely on imputation methods (e.g., single imputation, multiple imputation, etc. [79]) to analyze the data with other IPD-S models that do not permit the inclusion of missing data (e.g., repeated-measures ANOVA) or AD. Whichever procedure is finally selected for handling missing data, sensitivity analyses should be conducted to ensure that the results are robust to the missing data procedure specification [112]. We refer interested readers to Little et al. [80] and Schafer and Graham [79].

13 RELATED WORK

To the best of our knowledge, no previous attempts have been made in SE to provide guidelines for analyzing groups of SE replications—when researchers have access to the raw data and first-hand knowledge of the settings and participant characteristics. However, some previous articles already discussed the suitability of various research synthesis methods for combining *published results*. For instance, in the late 1990s, Pickard et al. [128] outlined the advantages

and disadvantages of meta-analysis of effect sizes [22], Fisher's method (i.e., an aggregation of p -values technique [22]) and vote-counting (i.e., a form of narrative synthesis procedure [22]) for aggregating the results of a series of case studies. However, Pickard et al. [128] aggregated the results of case studies whose raw data were accessible—because they argued that the effect size that they used (i.e., the Pearson correlation coefficient [22]) needed to be computed from the raw data to guarantee the consistency of results across the studies [128]. Ultimately, Pickard et al. also acknowledged that meta-analysis could be performed as long as study reports provided appropriate summary statistics to backcalculate the necessary effect sizes [128].

Ever since, meta-analysis has been tightly coupled in SE with the concept of synthesizing the results of already published studies (i.e., typically with standardized effect sizes such as Cohen's d and AD) [3], [101], [129]–[131]. To do this, researchers should backcalculate appropriate effect sizes from study reports, and, if the studies are not very dissimilar, use either fixed-effects models or random-effects models for combination [3], [100], [101], [130]. However, disparate advice with regard to the use of meta-analysis can be found in the SE literature. For example, while Pickard et al. [128] acknowledge that meta-analysis is only appropriate when the studies are homogeneous enough—or when the heterogeneity across the studies can be clearly attributed to certain conditions [128]—, Miller et al. [129] indicate that identical studies (e.g., replications using identical materials) may result in “strong correlations” affecting the reliability of the joint conclusions. This view has also been backed up by others in the SE community [28]. At the same time, and given the commonly heterogeneous results reported in the literature and the myriad variables that typically change across SE studies, Miller et al. [129] finally conclude that “...the heterogeneity of current empirical results is a major limitation to our ability to apply meta-analytic procedures...”.

Due to the limitations of meta-analysis and the particularities of SE studies, other SE researchers have proposed the use of other aggregation techniques for synthesizing already published empirical study results [3], [130], [132]. Briefly, such techniques commonly involve some sort of vote-counting technique (e.g., counting positive vs. negative results, small vs. large results, etc.), or the application of different aggregation techniques (e.g., meta-analysis, vote-counting, etc.) depending upon the characteristics of the studies being aggregated (number of available studies, number of changes made across the studies, etc. [130]).

Similar concerns about the limitations of meta-analysis have been also raised in other disciplines over the years [89], [133], [134]. The overall consensus nowadays seems to be that meta-analysis of effect sizes (i.e., AD) should be preferred over narrative synthesis or vote counting techniques—at least for aggregating quantitative results [12], [16], [22], [134]. Also, the meta-analysis of raw data (i.e., IPD-S) outperforms AD in some circumstances (e.g., in terms of statistical flexibility or for identifying participant-level moderators [13], [63], [74]). Still, the debate about the limitations of meta-analysis techniques and research synthesis is ongoing [134].

14 CONCLUSIONS

Researchers from different groups and institutions are collaborating on the construction of groups of replications in SE. Applying unsuitable aggregation techniques to analyze groups of SE replications may undermine their potential to provide in-depth insights from experiment results.

We learned about the recommendations and guidelines used to analyze and report groups of replications in mature experimental disciplines such as medicine and pharmacology [43], [44], [63]. Unfortunately, such guidelines could not be directly imported for the analysis of groups of SE replications because of the noticeable differences between groups of replications in SE and medicine that we came across (i.e., in terms of the number of changes made across the replications, participant heterogeneity, statistical power, etc.).

We designed an analysis procedure with a set of embedded guidelines to analyze the stereotypical group of SE replications [17]. To do this, we adopted the same basic structure typically followed in medicine and pharmacology for analyzing groups of replications. However, we adapted the steps to the characteristics of groups of SE replications, and their common limitations with regard to joint data analysis. The analysis procedure that we propose outlines a minimum set of steps that may potentially increase the informativeness of joint conclusions and moderator effects. It all boils down to providing appropriate descriptive statistics and visualizations to ease the interpretation and incorporation of results into prospective studies, as well as taking advantage of the raw data to provide joint conclusions and identify moderators. AD and IPD-S—random-effects models—are crucial for this purpose. Table 12 shows a summary of how to use the aggregation techniques proposed in our procedure.

To wrap up, we encourage SE researchers analyzing groups of replications to justify their aggregation techniques and, more importantly, to transparently report the statistical models and the raw data that they used to provide joint conclusions and identify moderators. With the aim of easing the application of the analysis procedure and with a view to reproducibility, the supplementary material of this article includes the step-by-step commented R code and raw data—with the associated R notebook—that led to the results reported throughout this article. In addition, we offer a more technical tutorial including R code snippets, dataset descriptions, and mathematical formulae to complement the understanding of the R code that we provide. All the supplementary material is also available at figshare (URL: <https://doi.org/10.6084/m9.figshare.7583909.v7>). We hope this encourages others to give the analysis procedure a go.

ACKNOWLEDGMENTS

This work was partially funded by Spanish Ministry of Science, Innovation and Universities research grant PGC2018-097265-B-I00.

TABLE 12
Summary of aggregation techniques to be used.

	AD	IPD-S
Recommended for	- Results aggregation - Experiment-level moderators	- Results aggregation - Experiment-level moderators - Participant-level moderators
When can be used	- RV metrics can be different - Complete observations - Effect sizes (or raw data to calculate them) are available	- RV metrics are identical - Allows missing data - Raw data are available
How should be used	- Fit random effects models	- Use Linear Mixed Models (LMMs) - Fit random effects models
How should be used for moderator analyses	Use either: - Sub-group meta-analysis for categorical moderators - Meta-regression for continuous moderators	- Fit LMMs with interaction terms - Increase statistical significance threshold to 0.1 - Pay less attention to p-values and focus on effect sizes and 95% CIs - Run one analysis per moderator

SUPPLEMENTARY MATERIAL

Supplementary material 1: Raw data (XLSX 13 kb).
Supplementary material 2: Characteristics (XLSX 9 kb).
Supplementary material 3: R code (R 17 kb).
Supplementary material 4: R notebook (Rmd 20 kb).
Supplementary material 5: R Tutorial (PDF 316 kb).

REFERENCES

- [1] D. I. Sjøberg, J. E. Hannay, O. Hansen, V. B. Kampenes, A. Karahasanovic, N.-K. Liborg, and A. C. Rekdal, "A survey of controlled experiments in software engineering," *Software Engineering, IEEE Transactions on*, vol. 31, no. 9, pp. 733–753, 2005.
- [2] K.-J. Stol and B. Fitzgerald, "The abc of software engineering research," *ACM Trans. Softw. Eng. Methodol.*, vol. 27, no. 3, pp. 11:1–11:51, Sep. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3241743>
- [3] B. A. Kitchenham, D. Budgen, and P. Brereton, *Evidence-based software engineering and systematic reviews*. CRC Press, 2015, vol. 4.
- [4] T. Dybå, V. B. Kampenes, and D. I. Sjøberg, "A systematic review of statistical power in software engineering experiments," *Information and Software Technology*, vol. 48, no. 8, pp. 745–755, 2006.
- [5] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in software engineering*. Springer Science & Business Media, 2012.
- [6] L. Muñoz, J.-N. Mazón, and J. Trujillo, "A family of experiments to validate measures for uml activity diagrams of etl processes in data warehouses," *Information and Software Technology*, vol. 52, no. 11, pp. 1188–1203, 2010.
- [7] G. Canfora, F. García, M. Piattini, F. Ruiz, and C. A. Visaggio, "A family of experiments to validate metrics for software process models," *Journal of Systems and Software*, vol. 77, no. 2, pp. 113–129, 2005.
- [8] S. Mouchawrab, L. C. Briand, Y. Labiche, and M. Di Penta, "Assessing, comparing, and combining state machine-based testing and structural testing: a series of experiments," *IEEE Transactions on Software Engineering*, vol. 37, no. 2, pp. 161–187, 2011.
- [9] T. Kosar, M. Mernik, and J. C. Carver, "Program comprehension of domain-specific and general-purpose languages: comparison using a family of experiments," *Empirical software engineering*, vol. 17, no. 3, pp. 276–304, 2012.
- [10] S. Abrahao, C. Gravino, E. Insfran, G. Scanniello, and G. Tortora, "Assessing the effectiveness of sequence diagrams in the comprehension of functional requirements: Results from a family of five experiments," *Software Engineering, IEEE Transactions on*, vol. 39, no. 3, pp. 327–342, 2013.
- [11] V. R. Basili, F. Shull, and F. Lanubile, "Building knowledge through families of experiments," *Software Engineering, IEEE Transactions on*, vol. 25, no. 4, pp. 456–473, 1999.
- [12] H. Cooper and E. A. Patall, "The relative benefits of meta-analysis conducted with individual participant data versus aggregated data," *Psychological methods*, vol. 14, no. 2, p. 165, 2009.
- [13] L. A. Stewart and J. F. Tierney, "To ipd or not to ipd? advantages and disadvantages of systematic reviews using individual patient data," *Evaluation & the health professions*, vol. 25, no. 1, pp. 76–97, 2002.
- [14] G. H. Lyman and N. M. Kuderer, "The strengths and limitations of meta-analyses based on aggregate data," *BMC medical research methodology*, vol. 5, no. 1, p. 14, 2005.
- [15] T. Debray, K. G. Moons, G. Valkenhoef, O. Efthimiou, N. Hummel, R. H. Groenwold, and J. B. Reitsma, "Get real in individual participant data (ipd) meta-analysis: a review of the methodology," *Research synthesis methods*, vol. 6, no. 4, pp. 293–309, 2015.
- [16] G. Biondi-Zoccai, *Umbrella reviews: Evidence synthesis with overviews of reviews and meta-epidemiologic studies*. Springer, 2016.
- [17] A. Santos, O. Gómez, and N. Juristo, "Analyzing families of experiments in se: a systematic mapping study," *IEEE Transactions on Software Engineering*, vol. ?, no. ?, pp. ?–?, 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8430603>
- [18] H. Brown and R. Prescott, *Applied mixed models in medicine*. John Wiley & Sons, 2014.
- [19] A. Whitehead, *Meta-analysis of controlled clinical trials*. John Wiley & Sons, 2002, vol. 7.
- [20] A. Santos and N. Juristo, "Comparing techniques for aggregating interrelated replications in software engineering," in *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*. ACM, 2018, p. 8.
- [21] J. Popay, H. Roberts, A. Sowden, M. Petticrew, L. Arai, M. Rodgers, N. Britten, K. Roen, and S. Duffy, "Guidance on the conduct of narrative synthesis in systematic reviews," *A product from the ESRC methods programme Version*, vol. 1, p. b92, 2006.
- [22] M. Borenstein, L. V. Hedges, J. P. Higgins, and H. R. Rothstein, *Introduction to Meta-Analysis*. John Wiley & Sons, 2011.
- [23] A. Field, J. Miles, and Z. Field, *Discovering statistics using R*. Sage publications, 2012.
- [24] B. Kitchenham, "Procedures for performing systematic reviews," *Keele, UK, Keele University*, vol. 33, no. 2004, pp. 1–26, 2004.
- [25] M. C. Simmonds, J. P. Higgins, L. A. Stewart, J. F. Tierney, M. J. Clarke, and S. G. Thompson, "Meta-analysis of individual patient data from randomized trials: a review of methods used in practice," *Clinical Trials*, vol. 2, no. 3, pp. 209–217, 2005.
- [26] J. L. Krein, L. Prechelt, N. Juristo, A. Nanthaamornphong, J. C. Carver, S. Vegas, C. D. Knutson, K. D. Seppi, and D. L. Eggett, "A multi-site joint replication of a design patterns experiment using moderator variables to generalize across contexts," *IEEE Transactions on Software Engineering*, vol. 42, no. 4, pp. 302–321, 2016.
- [27] F. J. Shull, J. C. Carver, S. Vegas, and N. Juristo, "The role of replications in empirical software engineering," *Empirical Software Engineering*, vol. 13, no. 2, pp. 211–218, 2008.
- [28] B. Kitchenham, "The role of replications in empirical software engineering a word of warning," *Empirical Software Engineering*, vol. 13, no. 2, pp. 219–221, 2008.

- [29] F. Q. Da Silva, M. Suassuna, A. C. C. França, A. M. Grubb, T. B. Gouveia, C. V. Monteiro, and I. E. dos Santos, "Replication of empirical studies in software engineering research: a systematic mapping study," *Empirical Software Engineering*, vol. 19, no. 3, pp. 501–557, 2014.
- [30] R. M. Bezerra, F. Q. da Silva, A. M. Santana, C. V. Magalhaes, and R. E. Santos, "Replication of empirical studies in software engineering: An update of a systematic mapping study," in *Proceedings of the 2015 9th International Symposium on Empirical Software Engineering and Measurement (ESEM)*. IEEE, 2015, pp. 1–4.
- [31] S. Amann, S. Beyer, K. Kevic, and H. Gall, "Software mining studies: Goals, approaches, artifacts, and replicability," in *Software Engineering*. Springer, 2013, pp. 121–158.
- [32] O. S. Gómez, N. Juristo, and S. Vegas, "Understanding replication of experiments in software engineering: A classification," *Information and Software Technology*, vol. 56, no. 8, pp. 1033–1048, 2014.
- [33] M. T. Baldassarre, J. Carver, O. Dieste, and N. Juristo, "Replication types: Towards a shared taxonomy," in *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*. ACM, 2014, p. 18.
- [34] F. Shull, M. G. Mendonça, V. Basili, J. Carver, J. C. Maldonado, S. Fabbri, G. H. Travassos, and M. C. Ferreira, "Knowledge-sharing issues in experimental software engineering," *Empirical Software Engineering*, vol. 9, no. 1-2, pp. 111–137, 2004.
- [35] M. Solari, S. Vegas, and N. Juristo, "Content and structure of laboratory packages for software engineering experiments," *Information and Software Technology*, 2017.
- [36] N. Juristo, S. Vegas, M. Solari, S. Abraho, and I. Ramos, "A process for managing interaction between experimenters to get useful similar replications," *Information and Software Technology*, vol. 55, pp. 215–225, 2013.
- [37] G. Cumming, *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. Routledge, 2013.
- [38] B. Kitchenham, L. Madeyski, D. Budgen, J. Keung, P. Breerton, S. Charters, S. Gibbs, and A. Pohthong, "Robust statistical methods for empirical software engineering," *Empirical Software Engineering*, pp. 1–52, 2016.
- [39] L. M. Friedman, C. Furberg, D. L. DeMets, D. Reboussin, and C. B. Granger, *Fundamentals of clinical trials*. Springer, 1998, vol. 3.
- [40] N. Juristo and A. M. Moreno, *Basics of software engineering experimentation*. Springer Science & Business Media, 2001.
- [41] B. A. Kitchenham, T. Dyba, and M. Jorgensen, "Evidence-based software engineering," in *Proceedings of the 26th international conference on software engineering*. IEEE Computer Society, 2004, pp. 273–281.
- [42] J. P. Higgins, S. Green et al., *Cochrane handbook for systematic reviews of interventions*. Wiley Online Library, 2008, vol. 5.
- [43] C. Anello, R. T. O'Neill, and S. Dubey, "Multicentre trials: a us regulatory perspective," *Statistical Methods in Medical Research*, vol. 14, no. 3, pp. 303–318, 2005.
- [44] J. A. Lewis, "Statistical principles for clinical trials (ich e9): an introductory note on an international guideline," *Statistics in medicine*, vol. 18, no. 15, pp. 1903–1942, 1999.
- [45] L. A. Stewart, M. Clarke, M. Rovers, R. D. Riley, M. Simmonds, G. Stewart, and J. F. Tierney, "Preferred reporting items for a systematic review and meta-analysis of individual participant data: the prisma-ipd statement," *Jama*, vol. 313, no. 16, pp. 1657–1665, 2015.
- [46] K. F. Schulz, D. G. Altman, and D. Moher, "Consort 2010 statement: updated guidelines for reporting parallel group randomised trials," *BMC medicine*, vol. 8, no. 1, p. 18, 2010.
- [47] D. Fisher, A. Copas, J. Tierney, and M. Parmar, "A critical review of methods for the assessment of patient-level interactions in individual participant data meta-analysis of randomized trials, and guidance for practitioners," *Journal of clinical epidemiology*, vol. 64, no. 9, pp. 949–967, 2011.
- [48] T. Pincus, C. Miles, R. Froud, M. Underwood, D. Carnes, and S. J. Taylor, "Methodological criteria for the assessment of moderators in systematic reviews of randomised controlled trials: a consensus study," *BMC medical research methodology*, vol. 11, no. 1, p. 14, 2011.
- [49] D. J. Feaster, S. Mikulich-Gilbertson, and A. M. Brincks, "Modeling site effects in the design and analysis of multi-site trials," *The American journal of drug and alcohol abuse*, vol. 37, no. 5, pp. 383–391, 2011.
- [50] D.-G. D. Chen and K. E. Peace, *Applied meta-analysis with R*. CRC Press, 2013.
- [51] J. J. Hox, M. Moerbeek, and R. van de Schoot, *Multilevel analysis: Techniques and applications*. Routledge, 2010.
- [52] W. H. Finch, J. E. Bolin, and K. Kelley, *Multilevel modeling using R*. Crc Press, 2014.
- [53] D. A. Luke, *Multilevel modeling*. Sage, 2004, vol. 143.
- [54] A. de Zoete, M. de Boer, M. van Tulder, S. Rubinstein, M. Underwood, J. Hayden, J. Kalter, and R. Ostelo, "Rational and design of an individual participant data meta-analysis of spinal manipulative therapy for chronic low back pain protocol," *Systematic reviews*, vol. 6, no. 1, p. 21, 2017.
- [55] R. Chu, L. Thabane, J. Ma, A. Holbrook, E. Pullenayegum, and P. J. Devereaux, "Comparing methods to estimate treatment effects on a continuous outcome in multicentre randomized controlled trials: a simulation study," *BMC medical research methodology*, vol. 11, no. 1, p. 21, 2011.
- [56] R. M. Pickering and M. Weatherall, "The analysis of continuous outcomes in multi-centre trials with small centre sizes," *Statistics in medicine*, vol. 26, no. 30, pp. 5445–5456, 2007.
- [57] C. J. Maas and J. J. Hox, "Sufficient sample sizes for multilevel modeling," *Methodology*, vol. 1, no. 3, pp. 86–92, 2005.
- [58] D. M. McNeish and L. M. Stapleton, "The effect of small sample size on two-level model estimates: A review and illustration," *Educational Psychology Review*, vol. 28, no. 2, pp. 295–314, 2016.
- [59] A. Bell and K. Jones, "Explaining fixed effects: Random effects modeling of time-series cross-sectional and panel data," *Political Science Research and Methods*, vol. 3, no. 1, pp. 133–153, 2015.
- [60] H. Erdogmus, M. Morisio, and M. Torchiano, "On the effectiveness of the test-first approach to programming," *IEEE Transactions on Software Engineering*, vol. 31, no. 3, pp. 226–237, 2005.
- [61] M. J. Crawley, *The R book*. John Wiley & Sons, 2012.
- [62] W. R. Shadish, T. D. Cook, and D. T. Campbell, *Experimental and quasi-experimental designs for generalized causal inference*. Wadsworth Cengage learning, 2002.
- [63] L. Bero and D. Rennie, "The cochrane collaboration: preparing, maintaining, and disseminating systematic reviews of the effects of health care," *Jama*, vol. 274, no. 24, pp. 1935–1938, 1995.
- [64] A. Phillips and V. Haudiquet, "Ich e9 guideline statistical principles for clinical trials: a case study," *Statistics in medicine*, vol. 22, no. 1, pp. 1–11, 2003.
- [65] D. B. Petitti, *Meta-analysis, decision analysis, and cost-effectiveness analysis: methods for quantitative synthesis in medicine*. OUP USA, 2000, no. 31.
- [66] D. Falessi, N. Juristo, C. Wohlin, B. Turhan, J. Münch, A. Jedlitschka, and M. Oivo, "Empirical software engineering experts on the use of students and professionals in experiments," *Empirical Software Engineering*, pp. 1–38, 2017.
- [67] O. Laitenberger, K. El Emam, and T. G. Harbich, "An internally replicated quasi-experimental comparison of checklist and perspective based reading of code documents," *IEEE Transactions on Software Engineering*, vol. 27, no. 5, pp. 387–421, 2001.
- [68] A. R. Localio, J. A. Berlin, T. R. Ten Have, and S. E. Kimmel, "Adjustments for center in multicenter studies: an overview," *Annals of internal medicine*, vol. 135, no. 2, pp. 112–123, 2001.
- [69] H. C. Kraemer, "Pitfalls of multisite randomized clinical trials of efficacy and effectiveness," *Schizophrenia Bulletin*, vol. 26, no. 3, pp. 533–541, 2000.
- [70] G. P. Quinn and M. J. Keough, *Experimental design and data analysis for biologists*. Cambridge University Press, 2002.
- [71] J. Ensor, D. L. Burke, K. I. Snell, K. Hemming, and R. D. Riley, "Simulation-based power calculations for planning a two-stage individual participant data meta-analysis," *BMC medical research methodology*, vol. 18, no. 1, p. 41, 2018.
- [72] M. Rodgers, A. Sowden, M. Petticrew, L. Arai, H. Roberts, N. Britten, and J. Popay, "Testing methodological guidance on the conduct of narrative synthesis in systematic reviews: effectiveness of interventions to promote smoke alarm ownership and function," *Evaluation*, vol. 15, no. 1, pp. 49–73, 2009.
- [73] C. T. Smith, J. Oyee, M. Marcucci, M. Rovers, A. Iorio, R. Riley, P. Williamson, and M. Clarke, "Individual participant data meta-analyses compared with meta-analyses based on aggregate data," *Trials*, vol. 12, no. 1, p. A57, 2011.
- [74] D. J. Fisher, J. R. Carpenter, T. P. Morris, S. C. Freeman, and J. F. Tierney, "Meta-analytical methods to identify who benefits most from treatments: daft, deluded, or deft approach?" *bmj*, vol. 356, p. j573, 2017.

- [75] P. Runeson, A. Stefik, A. Andrews, S. Gronblom, I. Porres, and S. Siebert, "A comparative analysis of three replicated experiments comparing inspection and unit testing," in *Replication in Empirical Software Engineering Research (RESER), 2011 Second International Workshop on*. IEEE, 2011, pp. 35–42.
- [76] F. Ricca, G. Scanniello, M. Torchiano, G. Reggio, and E. Astesiano, "Assessing the effect of screen mockups on the comprehension of functional requirements," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 24, no. 1, p. 1, 2014.
- [77] P. Alasuutari, L. Bickman, and J. Brannen, *The SAGE handbook of social research methods*. Sage, 2008.
- [78] R. Silberzahn, E. L. Uhlmann, D. P. Martin, P. Anselmi, F. Aust, E. Awtrey, Š. Bahník, F. Bai, C. Bannard, E. Bonnier *et al.*, "Many analysts, one data set: Making transparent how variations in analytic choices affect results," *Advances in Methods and Practices in Psychological Science*, vol. 1, no. 3, pp. 337–356, 2018.
- [79] J. L. Schafer and J. W. Graham, "Missing data: our view of the state of the art," *Psychological methods*, vol. 7, no. 2, p. 147, 2002.
- [80] R. J. Little, R. D'agostino, M. L. Cohen, K. Dickersin, S. S. Emerson, J. T. Farrar, C. Frangakis, J. W. Hogan, G. Molenberghs, S. A. Murphy *et al.*, "The prevention and treatment of missing data in clinical trials," *New England Journal of Medicine*, vol. 367, no. 14, pp. 1355–1360, 2012.
- [81] R. L. Wasserstein and N. A. Lazar, "The asa's statement on p-values: context, process, and purpose," *Am Stat*, vol. 70, no. 2, pp. 129–133, 2016.
- [82] R. McElreath, "Statistical rethinking. texts in statistical science," 2015.
- [83] G. Abo-Zaid, B. Guo, J. J. Deeks, T. P. Debray, E. W. Steyerberg, K. G. Moons, and R. D. Riley, "Individual participant data meta-analyses should not ignore clustering," *Journal of clinical epidemiology*, vol. 66, no. 8, pp. 865–873, 2013.
- [84] B. C. Kahan and T. P. Morris, "Analysis of multicentre trials with continuous outcomes: when and how should we account for centre effects?" *Statistics in medicine*, vol. 32, no. 7, pp. 1136–1149, 2013.
- [85] D. L. Burke, J. Ensor, and R. D. Riley, "Meta-analysis using individual participant data: one-stage and two-stage approaches, and why they may differ," *Statistics in medicine*, vol. 36, no. 5, pp. 855–875, 2017.
- [86] J. F. Tierney, C. Vale, R. Riley, C. T. Smith, L. Stewart, M. Clarke, and M. Rovers, "Individual participant data (ipd) meta-analyses of randomised controlled trials: guidance on their use," *PLoS Med*, vol. 12, no. 7, p. e1001855, 2015.
- [87] T. Greco, A. Zangrillo, G. Biondi-Zoccai, and G. Landoni, "Meta-analysis: pitfalls and hints," *Heart, lung and vessels*, vol. 5, no. 4, p. 219, 2013.
- [88] T. S. Clark and D. A. Linzer, "Should i use fixed or random effects?" *Political Science Research and Methods*, vol. 3, no. 2, pp. 399–408, 2015.
- [89] J. Lau, J. P. Ioannidis, and C. H. Schmid, "Summing up evidence: one answer is not always enough," *The lancet*, vol. 351, no. 9096, pp. 123–127, 1998.
- [90] J. Higgins, A. Whitehead, R. M. Turner, R. Z. Omar, and S. G. Thompson, "Meta-analysis of continuous outcome data from individual patients," *Statistics in medicine*, vol. 20, no. 15, pp. 2219–2241, 2001.
- [91] S. B. Morris and R. P. DeShon, "Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs," *Psychological methods*, vol. 7, no. 1, p. 105, 2002.
- [92] A. Field, *Discovering statistics using IBM SPSS statistics*. Sage, 2013.
- [93] R. S. Nickerson, "Null hypothesis significance testing: a review of an old and continuing controversy," *Psychological methods*, vol. 5, no. 2, p. 241, 2000.
- [94] J. Cohen, "The earth is round (p. 05)." *American psychologist*, vol. 49, no. 12, p. 997, 1994.
- [95] L. V. Hedges and I. Olkin, "Vote-counting methods in research synthesis," *Psychological bulletin*, vol. 88, no. 2, p. 359, 1980.
- [96] B. C. Kahan and T. P. Morris, "Assessing potential sources of clustering in individually randomised trials," *BMC medical research methodology*, vol. 13, no. 1, p. 58, 2013.
- [97] O.-M. Kwok, A. T. Underhill, J. W. Berry, W. Luo, T. R. Elliott, and M. Yoon, "Analyzing longitudinal data with multilevel models: An example with individuals living with lower extremity intra-articular fractures," *Rehabilitation psychology*, vol. 53, no. 3, p. 370, 2008.
- [98] J. Twisk, M. de Boer, W. de Vente, and M. Heymans, "Multiple imputation of missing values was not necessary before performing a longitudinal mixed-model analysis," *Journal of clinical epidemiology*, vol. 66, no. 9, pp. 1022–1028, 2013.
- [99] N. Juristo and O. S. Gómez, "Replication of software engineering experiments," in *Empirical Software Engineering and Verification: International Summer Schools, LASER 2008-2010, Elba Island, Italy, Revised Tutorial Lectures*, B. Meyer and M. Nordio, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 60–88. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-25231-0_2
- [100] D. I. Sjöberg, T. Dyba, and M. Jorgensen, "The future of empirical methods in software engineering research," in *Future of Software Engineering, 2007. FOSE'07*. IEEE, 2007, pp. 358–378.
- [101] W. Hayes, "Research synthesis in software engineering: a case for meta-analysis," in *Software Metrics Symposium, 1999. Proceedings. Sixth International*. IEEE, 1999, pp. 143–151.
- [102] J. Miller, "Can results from software engineering experiments be safely combined?" in *Software Metrics Symposium, 1999. Proceedings. Sixth International*. IEEE, 1999, pp. 152–158.
- [103] J. E. Hannay, T. Dybå, E. Arisholm, and D. I. Sjöberg, "The effectiveness of pair programming: A meta-analysis," *Information and Software Technology*, vol. 51, no. 7, pp. 1110–1122, 2009.
- [104] P. C. Lambert, A. J. Sutton, K. R. Abrams, and D. R. Jones, "A comparison of summary patient-level covariates in meta-regression with individual patient data meta-analysis," *Journal of clinical epidemiology*, vol. 55, no. 1, pp. 86–94, 2002.
- [105] J. A. Berlin, J. Santanna, C. H. Schmid, L. A. Szczech, and H. I. Feldman, "Individual patient-versus group-level data meta-regressions for the investigation of treatment effect modifiers: ecological bias rears its ugly head," *Statistics in medicine*, vol. 21, no. 3, pp. 371–387, 2002.
- [106] O. Dieste, C. E. R. Fonseca, G. Raura, and P. Rodríguez, "Professionals are not superman: failures beyond motivation in software experiments," in *Proceedings of the 5th International Workshop on Conducting Empirical Studies in Industry*. IEEE Press, 2017, pp. 27–32.
- [107] D. S. Cruzes and T. Dybå, "Research synthesis in software engineering: A tertiary study," *Information and Software Technology*, vol. 53, no. 5, pp. 440–455, 2011.
- [108] M. DixonWoods, R. Fitzpatrick, and K. Roberts, "Including qualitative research in systematic reviews: opportunities and problems," *Journal of evaluation in clinical practice*, vol. 7, no. 2, pp. 125–133, 2001.
- [109] J. Thomas, A. Harden, A. Oakley, S. Oliver, R. Sutcliffe, Katyand Rees, G. Brunton, and J. Kavanagh, "Integrating qualitative research with trials in systematic reviews," *BMJ*, vol. 328, no. 7446, pp. 1010–1012, 2004.
- [110] T. A. Snijders, "Multilevel analysis," in *International Encyclopedia of Statistical Science*. Springer, 2011, pp. 879–882.
- [111] C. Duncan, K. Jones, and G. Moon, "Context, composition and heterogeneity: using multilevel models in health research," *Social science & medicine*, vol. 46, no. 1, pp. 97–117, 1998.
- [112] L. Thabane, L. Mbuagbaw, S. Zhang, Z. Samaan, M. Marcucci, C. Ye, M. Thabane, L. Giangregorio, B. Dennis, D. Kosa *et al.*, "A tutorial on sensitivity analyses in clinical trials: the what, why, when and how," *BMC medical research methodology*, vol. 13, no. 1, p. 92, 2013.
- [113] M. W. Fagerland, "t-tests, non-parametric tests, and large studies—a paradox of statistical practice?" *BMC Medical Research Methodology*, vol. 12, no. 1, p. 78, 2012.
- [114] C. E. McCulloch and J. M. Neuhaus, "Misspecifying the shape of a random effects distribution: why getting it wrong may not matter," *Statistical science*, pp. 388–402, 2011.
- [115] T. Lumley, P. Diehr, S. Emerson, and L. Chen, "The importance of the normality assumption in large public health data sets," *Annual review of public health*, vol. 23, no. 1, pp. 151–169, 2002.
- [116] C. E. McCulloch and J. M. Neuhaus, *Generalized linear mixed models*. Wiley Online Library, 2001.
- [117] L. Hoffman and M. J. Rovine, "Multilevel models for the experimental psychologist: Foundations and illustrative examples," *Behavior Research Methods*, vol. 39, no. 1, pp. 101–117, 2007.
- [118] C. J. Maas and T. A. Snijders, "The multilevel approach to repeated measures for complete and incomplete data," *Quality and Quantity*, vol. 37, no. 1, pp. 71–89, 2003.
- [119] S. Ren, H. Lai, W. Tong, M. Aminzadeh, X. Hou, and S. Lai, "Nonparametric bootstrapping for hierarchical data," *Journal of Applied Statistics*, vol. 37, no. 9, pp. 1487–1498, 2010.

- [120] C. A. Field and A. H. Welsh, "Bootstrapping clustered data," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 69, no. 3, pp. 369–390, 2007.
- [121] S. Nakagawa and I. C. Cuthill, "Effect size, confidence interval and statistical significance: a practical guide for biologists," *Biological reviews*, vol. 82, no. 4, pp. 591–605, 2007.
- [122] A. Zuur, E. Ieno, N. Walker, A. Saveliev, and G. Smith, "Mixed effects models and extensions in ecology with r. gail m, krickeberg k, samet jm, tsiatis a, wong w, editors," *New York, NY: Spring Science and Business Media*, 2009.
- [123] N. Pautz, B. Olivier, and F. Steyn, "The use of parametric effect sizes in single study musculoskeletal physiotherapy research: A practical primer," *Physical Therapy in Sport*, vol. 32, pp. 87–97, 2018.
- [124] A. F. Zuur, J. M. Hilbe, and E. N. Ieno, *A Beginner's Guide to GLM and GLMM with R: A Frequentist and Bayesian Perspective for Ecologists*. Highland Statistics Limited, 2013.
- [125] N. Pautz, B. Olivier, and F. Steyn, "The use of nonparametric effect sizes in single study musculoskeletal physiotherapy research: A practical primer," *Physical Therapy in Sport*, vol. 33, pp. 117–124, 2018.
- [126] C. O. Fritz, P. E. Morris, and J. J. Richler, "Effect size estimates: current use, calculations, and interpretation," *Journal of experimental psychology: General*, vol. 141, no. 1, p. 2, 2012.
- [127] W. H. Finch, J. E. Bolin, and K. Kelley, *Multilevel modeling using R*. Crc Press, 2016.
- [128] L. M. Pickard, B. A. Kitchenham, and P. W. Jones, "Combining empirical results in software engineering," *Information and software technology*, vol. 40, no. 14, pp. 811–821, 1998.
- [129] J. Miller, "Applying meta-analytical procedures to software engineering experiments," *Journal of Systems and Software*, vol. 54, no. 1, pp. 29–39, 2000.
- [130] E. Fernández, "Aggregation process with multiple evidence levels for experimental studies in software engineering," in *Proceedings 2nd International Doctoral Symposium on Empirical Software Engineering*. Pag. Citeseer, 2007, pp. 75–81.
- [131] M. Shepperd, N. Ajenka, and S. Counsell, "The role and value of replication in empirical software engineering results," *Information and Software Technology*, vol. 99, pp. 120–132, 2018.
- [132] B. K. Olorisade, S. Vegas, and N. Juristo, "Determining the effectiveness of three software evaluation techniques through informal aggregation," *Information and Software Technology*, vol. 55, no. 9, pp. 1590–1601, 2013.
- [133] J. Ioannidis, N. Patsopoulos, and H. Rothstein, "Research methodology: reasons or excuses for avoiding meta-analysis in forest plots," *BMJ: British Medical Journal*, vol. 336, no. 7658, pp. 1413–1415, 2008.
- [134] J. Gurevitch, J. Koricheva, S. Nakagawa, and G. Stewart, "Meta-analysis and the science of research synthesis," *Nature*, vol. 555, no. 7695, p. 175, 2018.

Sira Vegas has been associate professor of software engineering with the School of Computer Engineering at the Technical University of Madrid, Spain, since 2008. Sira belongs to the review board of IEEE Transactions on Software Engineering, and is a regular reviewer of the Empirical Software Engineering Journal. She was program chair for the International Symposium on Empirical Software Engineering and Measurement in 2007.

Markku Oivo is professor and head of the M3S research unit at the University of Oulu, Finland. He held visiting positions at the University of Maryland (1990-91), Schlumberger Ltd. (Paris 1994-95), Fraunhofer IESE (1999-2000), University of Bolzano (2014-15), and Universidad Politécnica de Madrid (2015).

Natalia Juristo has been full professor of software engineering with the School of Computer Engineering at the Technical University of Madrid, Spain, since 1997. She was awarded a FiDiPro (Finland Distinguished Professor Program) professorship at the University of Oulu, from January 2013 until June 2018. Natalia belongs to the editorial board of EMSE and STVR. In 2009, Natalia was awarded an honorary doctorate by Blekinge Institute of Technology in Sweden.

Adrian Santos received his MSc in Software and Systems and MSc in Software Project Management from the Technical University of Madrid, Spain, and his MSc in IT Auditing, Security and Government from the Autonomous University of Madrid, Spain. He is a PhD student at the University of Oulu, Finland. He is a member of the American Statistical Association (ASA) and the International Society for Bayesian Analysis (ISBA).