Psychometrics in Behavioral Software Engineering: A Methodological Introduction with Guidelines

DANIEL GRAZIOTIN, University of Stuttgart, Institute of Software Engineering, Germany

PER LENBERG and ROBERT FELDT, Chalmers and the University of Gothenburg, Dept. of Computer Science and Engineering, Sweden

STEFAN WAGNER, University of Stuttgart, Institute of Software Engineering, Germany

A meaningful and deep understanding of the human aspects of software engineering (SE) requires psychological constructs to be considered. Psychology theory can facilitate the systematic and sound development as well as the adoption of instruments (e.g., psychological tests, questionnaires) to assess these constructs. In particular, to ensure high quality, the psychometric properties of instruments need evaluation. In this article, we provide an introduction to psychometric theory for the evaluation of measurement instruments for SE researchers. We present guidelines that enable using existing instruments and developing new ones adequately. We conducted a comprehensive review of the psychology literature framed by the Standards for Educational and Psychological Testing. We detail activities used when operationalizing new psychological constructs, such as item pooling, item review, pilot testing, item analysis, factor analysis, statistical property of items, reliability, validity, and fairness in testing and test bias. We provide an openly available example of a psychometric evaluation based on our guideline. We hope to encourage a culture change in SE research towards the adoption of established methods from psychology. To improve the quality of behavioral research in SE, studies focusing on introducing, validating, and then using psychometric instruments need to be more common.

CCS Concepts: • **Software and its engineering**; • **Human-centered computing** → **Empirical studies in collaborative and social computing**; • **General and reference** → *Reference works*; **Surveys and overviews**; *General literature*; **Cross-computing tools and techniques**; **Empirical studies**; **Measurement**; *Evaluation*; **Validation**; • **Applied computing** → **Psychology**;

Additional Key Words and Phrases: Empirical software engineering, psychology, behavioral software engineering, methodology, questionnaire design

ACM Reference format:

Daniel Graziotin, Per Lenberg, Robert Feldt, and Stefan Wagner. 2021. Psychometrics in Behavioral Software Engineering: A Methodological Introduction with Guidelines. *ACM Trans. Softw. Eng. Methodol.* 31, 1, Article 7 (September 2021), 36 pages.

<https://doi.org/10.1145/3469888>

We acknowledge the support of Swedish Armed Forces, Swedish Defense Materiel Administration and Swedish Governmental Agency for Innovation Systems (VINNOVA) in the project number 2017-04874. The study was also partly funded by Marianne and Marcus Wallenberg Foundation (project number 2017.0071) as well as the Alexander von Humboldt (AvH) Foundation.

Authors' addresses: D. Graziotin and S. Wagner, University of Stuttgart, Institute of Software Engineering, Universitätsstraße 38, 70569 Stuttgart, Germany, Stuttgart, Germany; emails: {daniel.graziotin, stefan.wagner}@iste.uni-stuttgart.de; P. Lenberg and R. Feldt, Chalmers and the University of Gothenburg, Dept. of Computer Science and Engineering, SE-412 96, Gothenburg, Sweden; emails: {per.lenberg, robert.feldt}@chalmers.se.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).

1049-331X/2021/09-ART7 \$15.00

<https://doi.org/10.1145/3469888>

1 INTRODUCTION

Software is developed for people, by people. For decades, we have recognized that, no matter the size and importance of the technical side of software engineering, it is humans that ultimately drive the underlying processes and produce the desired artifacts [143]. Software engineers are knowledge workers and have knowledge as their main capital [130]. They need to construct, retrieve, model, aggregate, and present knowledge in all their analytic and creative daily activities [119]. Operations related to knowledge are cognitive in nature, and cognition is influenced by characteristics of human behavior, including personality, affect, and motivation [62]. It is no wonder that industry and academia have explored psychological aspects of software development and the assessment of psychological constructs at the individual, team, and organization level [14, 16, 42, 57, 80, 136, 141].

Psychological assessment is the gathering of psychology-related data towards an evaluation that is accomplished through the use of tools such as tests, interviews, case studies, behavioral observation, and other procedures [21]. With the present article, we are primarily focused on psychological tests, i.e., quantitative research,¹ as behavioral software engineering has turned much attention to employing theory and measurement instruments from psychology [14, 30, 42, 54, 80].

Psychological tests are instruments (e.g., questionnaires) used to measure unobserved constructs [21]. We can not assess these constructs directly, like when measuring the source lines of code or specific properties of a UML diagram. Hence, the underlying variables are called latent variables [98]. Examples of such unobserved constructs include attitude, mood, happiness, job satisfaction, commitment, motivation, intelligence, abilities, cognitive skills, and performance. We need to create good measurement instruments² for the assessment of such constructs. For ensuring a systematic and sound development of psychological tests and their interpretation, the field of psychometrics was born [98, 121].

Psychometrics is the development of measurement instruments and the assessment on whether these instruments are reliable and valid forms of measurement [48]. Psychometrics is also the branch of psychology and education that is devoted to the construction of valid and reliable psychological tests [121].

The origin of psychometrics go back more than a hundred years [98, 121]. The establishment of testing as a mathematically oriented sub-discipline of psychology are to be found in the 1930s, with three milestone publication events [68]: (1) *Psychometrika* volume 1, issue 1, in March 1936, a venue “devoted to the development of psychology as a quantitative rational science” (subheading of the journal). (2) Guilford’s first edition of *Psychometric Methods* [58]. (3) Thurstone’s article *Psychology as a quantitative rational science* in the journal *Science* [133] to address the establishing of the *Psychometric Society* in 1936.

Thurstone’s paper [133] served as a plea to recognize quantitative roots of psychology, by which rationalizing problems in psychology with mathematics language and methods from statistics would allow the development of psychological science. Proper development and validation of tests has the potential to result in better decisions on individuals, while, on the opposite end, improper development and validation of tests might result in invalid interpretation of results, economic loss, and even harm of individuals [3]. The individual level also affects teams and, ultimately, organizations, and even when these are in focus, measuring constructs that concern them is typically built up from measurements on the individual level (e.g., organizational readiness for change, at the organization level, is computed from aggregated responses of individuals who work at the organization [126]). Personality assessment is a classic example of psychological testing in personnel

¹We are interested in qualitative research as well, and we value it the same as quantitative research in behavioral software engineering, as we state in Section 13.

²In this article, we use the terms psychological test, measurement instrument, and questionnaire interchangeably.

selection, which has been employed across domains, including companies related to information technology [32, 150]. Another example is assessment of job satisfaction of workers, which is an evaluative judgment one makes about one's job or job situation [144]. A company would, arguably, attempt to foster job satisfaction of individuals, for example by employing agile methodologies instead of traditional processes [136]. Similarly, IT-related companies are interested in fostering motivation of software engineers, which is related to factors that energize, channel, and sustain human behavior over time [46]. Stress of software developers can also be measured by psychological tests coupled with physiological measurements [101] and biometrics [39]. Finally, to bring examples related to cognition and behavior, research in software engineering has recently turned attention to understanding and addressing cognitive biases while developing software [18, 151] and identifying and reducing gender bias in development teams [141].

Improper development, administration, and handling of psychological tests could harm the company by hiring a non-desirable person, and it could harm the interviewee because of missed opportunities.

We believe that solid theoretical and methodological foundations should be the first step when designing a measurement instrument. The reality, however, is that not all tests are well developed in psychology [3]. Software engineering research, especially when studying psychological constructs quantitatively, is far from adopting rigorous and validated research artifacts.

1.1 Misconceptions with Psychological Tests in Software Engineering Research

Already in 2007, McDonald and Edwards [90] subtitled their paper “Examining the use and abuse of personality tests in software engineering.” The authors anticipated the issue that we attempt to address in the present manuscript, that is, the “the lack of progress in this [personality research in software engineering] field is due in part to the inappropriate use of psychological tests, frequently coupled with basic misunderstandings of personality theory by those who use them” (p. 67). While their focus was on personality, their concerns are held in a broader way to other psychometric instruments and constructs.

Instances of mishandle³ can, for example, be observed in the papers found by a systematic literature review of personality research in software engineering by Cruz et al. [30]. We noted in the results of Cruz et al. [30] that 48% of the personality studies in software engineering have employed the **Myers-Brigg Type Indicator (MBTI)** questionnaire, which has been shown to possess low to no reliability and validity properties [106] up to the point of being called a “little more than an elaborate Chinese fortune cookie” [64]. Feldt and Magazinius [41] similarly pointed out deficiencies of MBTI and proposed and used an alternative (IPIP) with more empirical support in the psychological literature.

Feldt et al. [42] have argued in favor of systematic studies on human aspects of software engineering; more specifically, to adopt measurement instruments coming from psychology and related fields. Graziotin et al. [52] have echoed the call seven years after but found that research on the affect of software developers had been threatened by a misunderstanding of related constructs and how to assess them. In particular, the authors noted that peers in software engineering tend to confuse affect-related psychological constructs such as emotions and moods with related, yet different, constructs such as motivation, commitment, and well-being.

Lenberg et al. [80] have conducted a systematic literature review of studies on human aspects in software development and engineering that made use of behavioral science, calling the field behavioral software engineering. Among their results, they found that software engineering research is

³We believe that direct accusations bring no value to our contribution and are counterproductive to our advancement of knowledge, so we discuss resources that point to specific issues rather than specific papers or authors.

threatened by several knowledge gaps when performing behavioral research, and that there have been very few collaborations between software engineering and behavioral science researchers.

Graziotin et al. [54], meanwhile, extended their prior observations on affect to a broader view of software engineering research with a psychological perspective. The work offered what we can consider the sentiment for the present article, that is, brief guidelines to select a theoretical framework and validated measurement instruments from psychology. Graziotin et al. [54] called the field “psychoempirical software engineering” but later agreed with Lenberg et al. [80] to unify the vision under “behavioral software engineering.” Hence, the present collaboration.

Our previous studies have also reported that, when a validated test from psychology is adopted by software engineering researchers, its items are often modified, causing the destruction of its psychometric reliability and validity properties. This includes a thorough evaluation of the psychometric properties of candidate instruments. Gren and Goldman [57] have argued in favor of “useful statistical methods for human factors research in software engineering” (paper title), which include underused methods such as Test-Retest, Cronbach’s α , and exploratory factor analysis—all of which are covered in this article. Gren [56] has also offered a psychological test theory lens for characterizing validity and reliability in behavioral software engineering research, further enforcing our view that software engineering research that investigates any psychological construct quantitatively should maintain fair psychometric properties. We agree with Gren [56] that we should “change the culture in software engineering research from seeing tool-constructing as the holy grail of research and instead value [psychometric] validation studies higher” (p. 3).

A mea culpa works better than a *j’accuse* in further building our case, so we bring a negative example from one of our previous studies. As reported in a very recent work by Ralph et al. [108] (which we appreciate in the next paragraph), “there is no widespread consensus about how to measure developers’ productivity or the main antecedents thereof. Many researchers use simple, unvalidated productivity scales” (p. 6). In one of the earliest works by the first author of the present article [53], we compared the affect triggered by a software development task with the self-assessed productivity of individual programmers. While we were very careful to select a validated measurement instrument of emotions and to highlight how self-assessment of productivity converges to objective assessment of productivity, we used a single Likert item scale to represent productivity. This choice was to reduce as much as possible the items of the measurement instrument, which had to be used every 10 minutes, for a total of nine times for each participant. While the results of the study are not invalidated by this choice, the productivity scale itself was not validated, making the results less valuable from a psychometric perspective and, thus, our interpretation of its results. The study was also (successfully) independently replicated twice, and both replications suffer from the same unfortunate choice.

We wish to refrain from being overly negative. The field of software engineering does have positive cases—excluding those from the present authors—that we can showcase here. For example, Fagerholm and Pagels [40] developed a questionnaire on lean and agile values and applied psychometric approaches to inspect the structure of value dimensions. Fagerholm [38] has also embodied psychometric approaches in their PhD dissertation by analyzing the validity of the constructs they studied. A more recent example is by Ralph et al. [108], who analyzed through a questionnaire the effects of the COVID-19 pandemic on developers’ well-being and productivity. The authors constructed their measurement instrument by incorporating psychometrically validated scales on constructs such as perceived productivity, disaster preparedness, fear and resilience, ergonomics, and organizational support. Furthermore, they employed confirmatory factor analysis (which we touch upon in the present article) to verify that the included items do indeed cluster and converge into the factors that are claimed to converge to.

While positive cases do exist, we notice that they are fairly recent and that we can do better than that. We want to synthesize knowledge from psychology fields to software engineering research towards better quantitative studies of behavioral aspects.

Filling the knowledge gap: introduction and guidelines to psychometric evaluation for behavioral software engineering research.

Overall, we argue that one thing that is missing is an introduction to the field of psychometrics for behavioral software engineering researchers. Such an introduction can help improve the understanding of the available measurement instruments and, also, the development of new tests, allowing researchers as well as practitioners to explore the human component in the software construction process more accurately.

1.2 Objective

Our overall objective is to address the lack of understanding and use of psychometrics in behavioral software engineering research including its limitations.

We also hope to increase software engineering researchers' awareness and respect of theories and tools developed in established fields of the behavioral science towards stronger methodological foundations of behavioral software engineering research.

1.3 Contribution

With this article, we contribute to the behavioral software engineering body of knowledge with a set of guidelines that enable a better understanding of psychological constructs in research activities when we interpret them through measurement instruments. This improvement in research quality is achieved by either (1) reusing psychometrically validated measurement instruments, as well as understanding why and how they are validated, or, if no such questionnaires exist, (2) developing new psychometrically validated questionnaires that are better suited for the software engineering domain.

Our contribution is enabled by offering one theoretical deliverable and one companion, practice-oriented deliverable.

- (1) We offer a review and synthesis of psychometric guidelines in form of several textbooks, review papers, as well as empirical studies, packaged in a style that is familiar to the software engineering researchers, including concrete examples and how to execute each activity. The guidelines enable evaluation of existing measurement instruments as well as developing new ones.
- (2) We offer a *hands-on* counterpart to our review by providing a fully reproducible implementation of our guidelines as R Markdown.

The **Standards for Educational and Psychological Testing (SEPT)**, American Educational Research Association et al. [3]) is a set of gold standards in psychological testing jointly developed by the **American Psychological Association (APA)**, **National Council on Measurement in Education (NCME)**, and the **American Educational Research Association (AERA)**. The book defines areas and standards that should be met when developing, validating, and administering psychological tests. We adopted SEPT as a framework to guide the article construction for ensuring that the standards are met and that the various other references are framed in the correct context.

Additionally, we organized the scoping of the article by comparing related work from the fields of psychology research. While the present article is not a systematic literature review or a mapping study—the discipline is so broad that entire textbooks have been written on it—we systematically framed its construction to ensure that all important topics were covered.

Fig. 1. Steps for developing a psychological test. Phases with dark background are uncommon in software engineering research and are covered in the present article.

1.4 Scope

Several authors, e.g., Crocker [28], Rust [121], Singh et al. [129], have proposed different phases for the psychometric development and evaluation of measurement instruments. Through our review, we identified 14 phases that we summarize visually in Figure 1 and outline as follows:

- (1) Identification of the primary purpose for which the test scores will be employed.
- (2) Identification of constructs, traits, and behaviors that are reflected by the purpose of the instrument.
- (3) Development of a test specification, delineation of the items proportion that should focus on each type of constructs, traits, and behaviors of the test.
- (4) Construction of an initial pool of items.
- (5) Review of the items.
- (6) Conduction of a pilot test with the revised items.
- (7) Execution of an item analysis to possibly reduce the number of items.
- (8) Evaluation of an exploratory factor analysis to possibly reduce and group items into components or factors.
- (9) Execution of a field test of the items with a larger, representative sample.
- (10) Determination of statistical properties of item scores.
- (11) Design and execution of reliability studies.
- (12) Design and execution of validity studies.
- (13) Evaluation of fairness in testing and test bias.
- (14) Development of guidelines for administering, scoring, and interpreting test scores.

We focus mainly on those with a dark background in Figure 1, as they are the most challenging and usually not covered in software engineering research.

As represented by the dashed lines, the process is linear only idealistically, when everything goes according to plans. Realistically, the construction of tests following a psychometric approach is failure-prone and iterative for correcting issues.

Furthermore, not all steps have to necessarily be followed when developing a measurement instrument. We are illustrating a wide range of possibilities, some of which are often brought by in future validation studies.

As a final note, the present article, as well as any psychometric construction of measurement instruments, is not a checklist. A psychometric evaluation does not include all elements reported in this article, as many facets of psychometrics are influenced by the research questions, study design, and data at hand. Yet, a proper psychometric evaluation requires a consideration of all elements reported in the present article.

1.5 Structure

After a brief introduction to the key concepts of psychometrics (Section 2) that are required to understand the rest of the article, we focus on test construction in psychometrics and the phases highlighted in Figure 1, namely, pool items (Section 3), item review (Section 4), pilot test (Section 5), item analysis (Section 6), factor analysis (Section 7), statistical properties (Section 8), reliability (Section 9), validity (Section 10), and fairness in testing and test bias (Section 11). We close our guidelines with two opposing sources for inspiration: (1) A comprehensive list for further reading (Section 12) to deepen what we are able to merely surface in this article. (2) A review of limitations of psychometrics and their critique (Section 13).

Finally, we provide a hands-on running example (Section 14) of a psychometric evaluation. We provide R code and generated datasets openly [51] following open science principles in software engineering [43].

2 CONCEPTS

This section provides an overview of basic terms and concepts from psychometrics that will enable an understanding of all remaining sections. In particular, we clarify on psychometric models and test types (and types of testing), as these will sometimes have an influence on the statistical methods and lens to adopt when designing and evaluating a measurement instrument.

2.1 Building Blocks

The fundamental idea behind psychological testing is that what is being assessed is not a physical object, such as height and weight. Rather, we are attempting to assess a construct, that is, a hypothetical entity [3] constructed by humans to represent concepts referring to various, concrete entities that are perceived in the moment, such as behaviors, experience, and attitudes [138].

If we assess the job satisfaction of a software developer, then we are not directly measuring the satisfaction of the individual. Instead, we compare the developer's score with other developers' scores or a set of established norms for job satisfaction. When comparing the satisfaction scores between developers, we are limited to seeing how the scores differentiate between satisfied and unsatisfied developers according to the knowledge and ideas we have about satisfied and dissatisfied individuals.

There are two common models of psychometrics, namely, functionalist and trait [121, 148]. Functionalist psychometrics often occurs in educational and occupational tests; it deals with how the design of a test is determined by its application and not about the constructs being measured [55, 121]. For functionalist design, a good test is one that is able to distinguish between individuals who perform well and individuals who perform less well on a job or in school activities. This is also called local criterion-based validity (explained in Section 10). The functionalist paradigm can be applied to most cases where a performance assessment or an evaluation is required.

Trait psychometrics attempts to address notions such as human intelligence, personality, and affect scientifically [25, 94]. The classic trait approach was based on the notion that, for example,

intelligence is related to biological individual differences, and trait psychometric tests aimed to measure traits that would represent biological differences among people [94].

Both schools of thought have several aspects in common, including test construction and validation methods, and they are linked by the theory of true scores [61]. The theory of true scores, or latent trait theory, is governed by formulas of the form:

$$X = T + E, \quad (1)$$

where X is the observed score, T is the true score, and E is the error. There are three assumptions with the theory of true scores: (1) all errors E are random and normally distributed, (2) true scores X are uncorrelated with the errors, and (3) different measures of the observed score T on the same participants are independent from each other. Besides all issues that come with the three assumptions, the theory has been criticized with the major point being that there is arguably no such thing as a true score, and that all that tests can measure are abstractions of psychological constructs [85].

Elaborations and re-interpretations of the theory of true scores have been proposed, among which is the statistical true score [15]. The statistical true score defines the true score T as the score we would obtain by averaging an infinite number of measures X from the same individual. With an infinite number of measures the random errors E cancel each other, leaving with the true score T . The statistical form of the theory of true scores should not be completely new to readers of software engineering, as most quantitative methods that are in use in our field nowadays are based on it. The statistical interpretation of the theory of true scores applies both to trait and functional psychometrics. A difference lies in generalization. Functional tests can only be specific to a certain context, while trait tests attempt to generalize to an overall construct present in a group of individuals.

2.2 Test Types and Types of Testing

Items on psychological tests can be knowledge-based or person-based [4, 121]. Knowledge-based tests assess whether an individual performs well regarding the knowledge of certain information, including possessing skills favoring performance or quality in knowledge-based tasks. For a software engineering example, debugging skills would be assessed by a knowledge-based test. Person-based tests, however, assess typical performance, or how a person is represented, with respect to a construct. Examples of constructs in person-based tests include personality, mood, and attitudes. The personalities of programmers in pair programming settings would be assessed by a person-based test.

Knowledge-based tests are usually uni-dimensional as they gravitate towards the notion of possessing or not possessing certain knowledge. We can also easily rank individuals on their scores and state who ranks better. Person-based tests are usually multi-dimensional and do not allow direct ranking of individuals without some assumptions. For example, a developer could score high on extroversion. A high score on extroversion does not make a developer with a lower extroversion score a “worse” developer in any way because of the lower score.

A second distinction to be made is between criterion-referenced and norm-referenced testing [49]. Criterion-referenced tests are constructed with reference to performance on *a priori*-defined values for establishing excellence [6, 49].

Continuing with the example on debugging skills, a criterion-referenced test would assess, with a score from 0 to 10, whether a developer is able to open a debugging tool and use its 10 basic functionalities. A score of 10 out of 10 would mean that the developer is able to debug software.

Norm-referenced tests lack *a priori*-defined scores. What constitutes a high score is in relation to how everyone else scores. A test for assessing the happiness of software developers will return

scores for each participant. The test itself will have a theoretical range, say, -10 for strong unhappiness and $+10$ for strong happiness. When a developer scores $+4$ on our happiness scale, all we can say is that the developer is happy rather than unhappy. If we know that software developers score -3 on average, with a standard deviation of 1, then we know that the developer is quite a happy one. The development and evolution of norm-referenced testing attempts, in addition to developing valid and reliable instruments to establish norms, values for populations and sub-populations of individuals. That is, norm-referenced tests allow us to compare scores with respect to what is considered *normal* [7, 49].

3 POOL ITEMS

The first steps in developing a measurement instrument is, of course, developing an initial set of items. Coaley [20] summarizes a planning and designing process in a way that reminds us of the **Goal-Question-Metrics (GQM)** model [11]:

(1) Set clear aims—on defining the purpose of the measurement instrument, which constructs we are targeting, and what is the intended target population. (2) Define the attribute(s)—on moving in the empirical lenses from the construct, or the object, to its attributes being measured. In this step, it is advised to perform a comprehensive literature review of the theoretical concepts being assessed. (3) Write a plan—on drafting a specification of the measurement instrument as if it was completed already, including the test content, target group and population, kinds and number of needed participants, administration instructions, time constraints, and how scores should be interpreted. (4) Writing items—on designing and constructing a pool of items related to all previously defined steps.

Sources for generating items could be experts in the domain or fields, interviewing potential respondents, and prior work [110] paired with constant comparison or update of research questions [100].

We do not want to spend too much space about writing items because of existing good literature: on the questionnaire construction process [9, 100, 110], on question effects and question-wording effects [71, 125, 128], and, for our field, on experiences and reviews on conducting surveys in software engineering research [19, 66, 74, 93, 140].

We will rather focus on selecting items from the produced pool. While some guidelines suggest to generate a double amount of items than those that are likely required [75], we note in Section 7 that factors require three to five associated answers to possess meaningful variance properties. Thus, a better strategy would be to develop 6 to 10 items per envisioned factor, or sub-construct.

4 ITEM REVIEW

When developing a new measurement instrument, we are likely (and encouraged by the previous section) to create more items than are really needed.

Item review and item analysis are a series of methods to reduce the number of items of a measurement instrument and keep the best performing ones [59, 75]. This is a two-step process, as shown in Figure 2. First, it requires a review by experts; then, a pilot study and statistical calculations. During the first step (*item review*), experts in the domain of knowledge evaluate items one-by-one and argue for their presence in the test [75, 124]. During the second step (*item analysis*), the developers of the measurement instrument calculate item facility and item discrimination based on a pilot study that uses the tentative set of items.

During item review, experts in the domain of knowledge discuss candidate items and argue in favor or against them, as it happens when discussing inclusion and exclusion of publications in systematic literature reviews [73].

Fig. 2. Phases for item review and item analysis.

Given its usage in the medical domain [8], we recommend using the Delphi technique [31] with domain experts to identify best candidates after initial qualitative probes to identify sub-constructs and items. To assess the degree of agreement among raters, we recommend using inter-rater reliability measures such as Cronbach's α [29] and Krippendorff's α [34].

After reaching an agreement on the items to be included, a pilot study is required for an analysis of the items.

5 PILOT TEST

A pilot test is necessary to probe the effectiveness of the developed items. After item review, we are left with a set of items that have been qualitatively evaluated by experts in the field. These items are too many, and we have no idea yet on how they would group together or contribute to the total variance related to the constructs we intend to represent. Hence, a pilot test.

Pilot tests allow to discover early issues related to questionnaire items, for example on wording problems, lack of clarity, confusing steps, or even discrimination of and between respondents [110].

There are no clear guidelines on how big the sample size for pilot tests should be. First criterion is rather on representativeness of the pilot sample with respect to the target group or population [132]. Johanson and Brooks [67] have identified suggestions from the literature that sees $n = 30$ participants as the lower limit to gather meaningful data and to reason for the test construction. Johanson and Brooks [67] have conducted a cost-benefit analysis to identify the point at which adding participants to a pilot sample size would incur in notably lesser effects in estimating population parameters. Analysis of sample size changes for a range of Pearson's correlations, range of proportions, and confidence intervals for reliability coefficients. All results confirm the indicated value of $n = 30$ as reasonable minimum recommendation for a pilot study for a psychological test.

A further possible novelty that we would like to report to the software engineering research audience is brought by Collins's [22] overview of cognitive methods to pretest survey instruments. The authors argue, in their review of the literature, that the tradition of survey pre-testing has mostly dealt with quantitative issues, e.g., standardizing data collection procedures such as question wording, worrying about clustering, and variance. All this, while valuable, should be coupled with techniques that go beyond the insofar assumption that respondents all understand the questions being asked in the same way, and that they are willing to respond to the same set of questions the same. Many of these techniques are about qualitative data rather than quantitative data. Cognitive methods enable exploring the process by which respondents think and act with when responding (or not) to questions, and which factors influence these answers.

Cognitive methods include think-aloud and probing in cognitive interviewing, paraphrasing, card sorts, vignettes, confidence ratings, and response latency timing. In absence of consensus on the best process [63], we will briefly summarize cognitive interviewing because it is an established method [35, 92, 146] and it brings qualitative methods to psychometrics.

Cognitive interviewing is a qualitative method that studies the question-response process by participants when answering questions on psychological tests [92]. The aim is to use theories from the cognitive sciences to understand how participants perceive and interpret items and to identify issues that may arise when distributing the instrument [35]. Cognitive interviewing requires a small sample of participants purposely selected for in-depth interviews on how and why they answered the question as they did [92].

The cognitive interview process is characterized by two verbal report methods: think-aloud and verbal probing [146]. In think-aloud, the participant is asked to (preferably) verbally explicating what they are thinking from the point of reading an item to the point of scoring the item. In verbal probing, the interviewer intervenes with specific questions and probes [22]. The former is respondent-driven, the latter is interviewer-driven. Examples of think-aloud questions are of open-ended nature and inviting for more thoughts, e.g.: “What are you thinking while reading this item?”; “What makes you think that?”; or “I noticed you hesitate for a while there, what made you hesitate?”. Examples of verbal-probing (from Collins [22], Table 2) include comprehension (“What does the term X mean to you?”), retrieval (“How did you calculate your answer?”), confidence judgment (“How well do you remember this?”), and response (“How did you feel about answering this question?”).

6 ITEM ANALYSIS

While qualitative methods are useful for identifying issues, we can apply quantitative techniques to reduce items. Item analysis refers to several statistical methods for the selection of items of a psychological test [59, 75]. Item analysis is often the only method for item reduction when a single construct should be evaluated, whereas multiple constructs and sub-constructs require further steps to be taken. Two of the most known techniques can be found in item facility and item discrimination. This section will explain both techniques according to the test type.

6.0.1 Item Facility. Item facility for an item is a measure of tendency in answering one item with the same score. This has different meanings according to the type of test.

Knowledge-based Test. Item facility for an item, also known as item difficulty, is defined as the ratio of the number of participants who provided a right answer over the number of all participants to a test [75, 121]. The value of item facility ranges from 0 (all respondents are wrong) to 1 (all respondents are right). In other words, item facility is the probability of obtaining the right answer for the item [75]. Of interest for test construction is the variance of an item. An item variance is the calculated variance of a set of item scores, which is a set of zeros and ones. That is, the item variance for an item with facility 0 is 0 (all are wrong), and an item variance for an item with facility 1 is 1 (all are right). Both these extremes would render the item rather void, as all individuals scoring the same on an item would not tell us anything interesting on the individuals. What usually happens is that some individuals will get the answer right, some will get the answer wrong. For such cases, we can compute the variance for an item by using Equation (2):

$$var_i = if_i * (1 - if_i), \quad (2)$$

where i is the item and if_i is the item facility for i . The highest possible value for var_i is 0.25, and this is the case when items are neither very easy or very difficult. When var_i has small values, for example 0.047, it means that most respondents tend to reply the same way for that item, making it

either extremely easy or extremely difficult. A value of var_i near 0.00 does not warrant automatic exclusion of item i , but the value should solicit a review.

Person-based Tests. Facility suggests that participants to all tests can be either right or wrong on an answer. What about trait measurement, where participants are not exactly right or wrong but are assessed in terms of a psychological construct? We can calculate item facility for these cases as well. The issue relies only in the naming, because item facility was developed for knowledge-based tests first. Some scholars prefer to use the term item endorsement or item location [112] to better reflect how calculations can be done on traits.

For trait measurement, it is common to have Likert items [84]. Item facility for Likert items can be calculated with the mean value of the item. If a Likert item maps to the values 0 (strongly disagree) to 5 (strongly agree), then the extreme values for the item will be 0 and 5 instead of 0 and 1. An item with average value of 4.8 with variance of 0.09 is a candidate for deletion, whereas an item with average value of 2.72 with variance of 3.02 is deemed interesting, because we prefer items that reflect on diversity of participants rather than having most participants score the same. Items such as those worded negatively should be reversed in their scoring prior to item analysis, so all items have comparable values.

6.0.2 Item Discrimination. Item discrimination is a technique to discover items that behave oddly with respect to what we expect participants to score the item. The meaning of item discrimination differs according to the type of test.

Knowledge-based Test. Item discrimination reflects items that behave oddly, in the sense that individuals that tend to score very high (or very low) on a test as a whole tend to be wrong (right) on the same item [89]. Such an item would possess a negative (positive) discrimination. Ideally, a test should have items with zero discrimination [75].

From a statistical point of view, if an item is uncorrelated with the overall test score, then it is almost certainly uncorrelated with the other items and making very little contribution to the overall variance of the test [121, 129]. Therefore, we calculate item discrimination by comparing the correlation coefficient of an item score and the overall test score. If the computed correlation coefficient is 0 or below, then we should consider removing the item.

Person-based Tests. What holds for knowledge-based tests holds to a wide extent with person-based tests. Instead of assessing how well an item behaves with respect to the test score, we instead assess how an item is in fact measuring the overall trait in question. By calculating the correlation coefficient of an item and the overall test score for a specific trait, we will have an initial estimate of how well an item represents the trait in question. If the computed correlation coefficient is 0 or below, then we should consider removing the item.

The Case of Norm-referenced Tests. The variance of an item, here calculated using the classic definition from statistics, is interesting also in the context of norm-referenced testing. Item facility also applies with norm-referenced testing, as the purpose of the test is to spread out individuals' scores as much as possible on a continuum. A larger spread is due to a larger variance, and we are interested in including items that make a contribution to the variance [75, 121]. Furthermore, if an item has a high correlation to other items and has a large variance, then it derives that the item makes high contribution to the total variance of a test and it will be kept in the pool of items.

The Case of Criterion-referenced Tests. Item analysis is often seen as applicable to norm-referenced testing exclusively [121]. With criterion-referenced testing, it is still possible to calculate item facility and item discrimination, and these can be conducted, for example, before and after teaching and formative activities (and this might include workshops of, say, Scrum methods,

Table 1. Example Correlation Matrix for Items
a, b, c, d, e

items	a	b	c	d	e
a	1.00				
b	0.73	1.00			
c	0.03	0.20	1.00		
d	0.89	0.84	-0.18	1.00	
e	-0.04	0.46	0.12	0.04	1.00

at IT companies). A difference in item facility before and after the teaching activities would indicate that the item is a valid measure of the skill taught. This would turn the measure of item facility into a measure of item discrimination as well.

6.0.3 Limitations of Item Analysis. Item analysis, while valuable and still in use today, is part of the **classical test theory (CTT)**, which assumes that an individual's observed score is the same as a true score plus an error score [135] (see Equation (1)). Modern replacements for CTT have been proposed, and the most prominent one is **item response theory (IRT)** [36]. IRT models build upon a function (called item response function, IRF, or **item characteristics curve, ICC**) that defines the probability of being right or wrong on an item [2]. IRT is outside the scope of the present article, as CTT is still in place to this day [105] and explaining IRT requires a publication on its own.

Item analysis, as presented in this section, assumes that there is a single test score, meaning that a single construct is being measured. Whenever multiple constructs or a construct of multiple factors are being measured, item analysis requires to be accompanied by factor analysis [129].

7 FACTOR ANALYSIS

Factor analysis is one of the most widely employed psychometric tools [75, 77, 129] and it can be applied to any dataset where the number of participants is higher than the number of item scores under observation. Factor analysis is for understanding which test items “go together” to form factors in the data that ideally should correspond to the constructs that we are aiming to assess [121]. At the same time, factor analysis allows to reduce the dimensionality of the problem space (i.e., reducing factors and/or associated items) and explaining the variance in the observed variables compared to underlying latent factors [77]. In case we intend to assess a single construct, factor analysis helps in identifying those items that (best) represent the construct we are interested in, so we can exclude the other items.

Factor analysis techniques are based on the notion that those constructs that we observe through our measurement instruments can be reduced to fewer latent variables that are unobservable but share a common variance [152] (see Section 2).

Factor analysis starts with computed correlation coefficients as its first building block. A way to summarize correlation coefficients is through a correlation matrix, which is a $n \times n$ matrix of n items that displays the correlation coefficient between all n items of the tentative test. Table 1 provides an example correlation matrix for five items a, b, c, d , and e . Given that two items (a and b) correlate with each other the same, no matter their order ($cor(a, b) = cor(b, a)$), and that a single item a correlates with itself with a perfect correlation coefficient ($cor(a, a) = 1.00$), the matrix displays only its lower triangle, omitting the repeated upper part, with the value of 1.00 as its diagonal [120, 152].

A high correlation among certain items, in our case a with b and d , and b with d , indicate these items might belong to the same factor. This approach, however, lacks part of the story. Questions such as “how do our candidate factor explain the total variance of the measurement instrument,” “to which candidate factor does an item belong *more*?” and “how are factors related to each other?” are better answered by further analysis.

There are two main factor analysis techniques, namely, **Exploratory Factor Analysis (EFA)** and **Confirmatory Factor Analysis (CFA)** [77, 112, 129]. EFA attempts to uncover patterns and clusters of items by exploring predictions, while CFA attempts to confirm hypotheses on existing factors [152].

7.1 Exploratory Factor Analysis

EFA is a family of analysis techniques aimed to reduce the number of items by retaining the items that are most relevant to certain factors [77]. Strictly speaking, when developing a measurement instrument, after item analysis, it is desirable to observe whether the measurements for the items tend to *cluster*. These clusters are likely to represent different factors that might or might not pertain to the construct being measured [75, 121]. EFA provides tools to group and select items from a correlation matrix.

EFA operates on Equation (3) for a measure x_1 [129, 152]:

$$X_1 = w_{11}F_1 + w_{21}F_2 + \dots + w_{n1}F_n + w_1U_1 + e_1, \quad (3)$$

where F_s are those factors grouping the items being analyzed, U_s are factors that are unique to each measure, w_s are loading of each item on respective factors, and e_s are random measurement errors. Factor loadings are, in practice, weights that provide us with an idea of how much an item contributes to a factor [152].

From the equation, we derive that the variance of the constructs being measured is explained by three parts: (1) the common factors, also known as *communality* of a variable [37, 129, 152] (2), the influence of factors that are *unique* to that measure, and (3) *random error*, or $V = V_h + V_c + V_e$.

Estimates for *communalities* of an item are often referred to as h^2 . h^2 is the calculated proportion of variance that is free of error and is thus shared with other variables in a correlation matrix [37, 129, 152]. Several techniques calculate the communality of a variable by summing the squared loadings of each variable associated with a variable.

Estimates for the *unique variance*, denoted as u^2 , is the proportion of variance that is associated with communalities, that is, $u^2 = 1 - h^2$ [37, 129, 152]. Determining a value for u^2 for an item allows us to find how much specific variance can be attributed to that variable.

Last, the *random error* that is associated with an item is the last component of the total variance. Random error is also often called the unreliability of variance [37, 152].

Unique factors are never correlated with common factors, but common factors may or may not be correlated with each other [152].

EFA encompasses three phases [77, 112, 121, 129], described in Figure 3. First, we have to select a fitting procedure to estimate the factor loadings and unique variances of the model. Then, we need to define and extract a number of factors. Finally, we need to rotate the factors to be able to properly interpret the produced factor matrices.

Many statistical programs allow to either perform all these phases separately or to perform more than one at the same time. It is not an easy task to assign a methodology to one of the three categories below. The reader is advised that some textbooks avoid our classification of phases and simply revert to a more practical set of questions, e.g., “how to calculate factors?” and “how many factors should we retain?” We also note that recent studies have formulated Bayesian versions of these classical exploratory factor analysis techniques and claimed several benefits [23].

Fig. 3. Phases for exploratory factor analysis.

7.1.1 Factor Loading. The most common technique for estimating the loadings and variance is to use the standard statistical technique of **principal component analysis (PCA)** [102]. PCA assumes that the communalities for the measures are equal to 1.0. That is, all the variance for a measure is explained only by the factors that are derived by the methodology, and hence there is no error variance. PCA operates on the correlation matrix, mostly on its eigenvalues, to extract factors that correlate with each other. The eigenvalue of a factor represents the variance of the variables accounted for by that factor. The lower the eigenvalue, the least the factor contributes to the variance explanation in the variables [97]. Factor weights are computed to extract the maximum possible variance, with successive factoring continuing until there is no further meaningful variance left. PCA is not a factor analysis method *strictu sensu*, as factor analysis assumes a presence of error variance rather than being able to explain all variance. Some advocates prefer to state that its output should be referred to as a series of components rather than factors. While more simplistic than other techniques to estimate factor loading, performing a PCA is still encouraged as a first step in EFA before performing the actual factor analysis [120, 121, 152].

Among the proper factor analysis techniques that exist, we are interested in a widely recommended technique for estimating loadings and variance named **principal axis factoring (PAF)** [75, 120, 145]. PAF does not operate under the assumption that the communalities are equal to 1.0, so the diagonal of the correlation matrix (e.g., the one in Table 1) is substituted with estimations of communalities, h^2 . PAF estimates the communalities using different techniques (e.g., the squared multiple correlation between a measure and the rest of measures) and a covariance matrix of the items. Factors are estimated one at a time up until there is a large enough variance accounted for in the correlation matrix. Under PAF, the ordering of the factors determines their importance in terms of fitting, e.g., the first factor accounts for as much variance as possible, followed by the second factor, and so on. Russell [120] provides a detailed description of the underlying statistical operations of PAF, which we will omit for the sake of brevity. Most statistical software provides functions to implement PAF.

7.1.2 Factor Extraction. Both PCA and PAF result in values assigned to candidate factors. Therefore, there has to be a strategy to extract meaningful factors. The bad news here is that there is no unique way, let alone a single proper way, to extract factors [114, 120, 129]. More than one strategy should be adopted in parallel to allow a comparison of results, ending with a sense-making analysis to take an ultimate decision [37]. Several factor extraction techniques exist, which are mentioned in the cited references in the present section. We provide here those that are used most

widely as well as those that are easier to apply and understand. Our running example (Section 14) enables an easier visualization of the concepts below.

Perhaps the simplest strategy to extract factors is **Kaiser's** eigenvalue-greater-than-one (K1) rule [70]. The rule simply states that factors with eigenvalue *higher* than 1.0 should be retained. Kaiser's rule is quite easy to apply but it is highly controversial [26, 37, 114, 120]. First, the rule was originally designed for PCA and not for PAF or other factor analysis methods, which might make it unsuitable for methodologies that provide estimation for commonalities as diagonals for correlation matrices [26]. Second, the cut-off value for 1.0 might discriminate for factors that are just above or just below 1.0 [26]. Third, computer simulations found that K1 tends to overestimate the number of factors [37]. Yet, K1 is still the default option for some statistical software suites, making it an unfortunate *de facto* main method for factor extraction [26].

Cattell [17] scree test, also based on eigenvalues, foresees the plot of the eigenvalues extracted from either the correlation matrix or the reduced correlation matrix (thus making it suitable for both PCA and PAF) against the factor they are associated with, in descending order. One then inspects the curved line for a break in the values (or an elbow) up to when a substantial drop in the eigenvalues cannot be observed anymore. The break is a point at which the shape of the curve becomes horizontal. The strategy is then to keep all factors before the breaking point. The three major criticisms of this approach is that it is subjective [26], that more than one scree might exist [134], and that data often does not offer a discernible scree and a conceptual analysis of the candidate factor is thus always required [121].

Revelle and Rocklin [117] proposed the **Very Simple Structure (VSS)** method for factor extraction that is based on assessing how the original correlation matrix can be reproduced by a simplified pattern matrix, for which only the highest loading of each item is retained (everything else set to zero) [26]. The VSS criterion to assess how well the pattern matrix performs is a number from 0 to 1, making it a goodness-of-fit measure almost of a confirmatory nature rather than an exploratory one [116]. The VSS criterion is gathered for solutions involving a number of factors that goes from 1 to a user-specified value. The strategy ends with selecting the number of factors that provide the highest VSS criterion.

Finally, the method of **parallel analysis (PA)**, introduced by Horn [65], was found to be very robust for factor extraction [26, 37]. PA starts with the K1 concept that only factors with eigenvalue larger than 1.0 should be kept. Horn [65] has argued that the K1 rule was developed with population statistics and was thus unsuitable when sampling data. Sampling errors would then cause some components from uncorrelated variables to have eigenvalues higher than one in the population [26]. PA takes into account the proportion of variance that results from sampling rather than being able to access to the population. The way it achieves this is a constant comparison of the solution with randomly generated data [116]. PA generates a large number of matrices from random data in parallel with the real data. Factors are retained as long as they are greater than the mean eigenvalue generated from the random matrices.

7.1.3 Factor Rotation. The last step is to rotate the factors in the dimensional space for improving our interpretation of the results [33, 75, 77].

An unrotated output, that is, the one that often results after factor extraction, maximizes the variance accounted for by the first factor, followed by the second factor, the third factor, and so on. That is, most items would load on the first factors and many of them would load on more than one factor in a substantial way.

Rotating factors build on the concept that there are a number of "factor solutions" that are mathematically equivalent to the solution found after factor extraction. By performing a rotation of the factors, we retain our solution but allow an easier interpretation. We rotate the factors to seek

a *simple structure*. A simple structure is a loadings pattern such that each item loads strongly on one factor only and weakly on other factors. If the reader is interested in mathematical foundations of factor rotation, then two deep overviews of factor rotation are offered by Browne [10], Darton [33].

There are two families of rotations, namely, orthogonal and oblique [120]. Orthogonal rotations force the assumption of independency between the factors, whereas oblique rotations allows the factors to correlate with each others. Which methodology to use is influenced by the statistics software; for example, R psych package [115] provides “varimax,” “quartimax,” “bentlerT,” “equamax,” “varimin,” “geominT,” and “bifactor” for orthogonal rotations and “promax,” “oblimin,” “simplicimax,” “bentlerQ,” “geominQ,” “biquartimin,” and “cluster” for oblique rotations. Several rotation methodologies are summarized by Abdi [1], Browne [10], Russell [120].

Perhaps the most known and employed [33] orthogonal rotation method is the Varimax rotation [69]. Varimax maximizes the variance (hence the name) of the squared loadings of a factor on all variables. Each factor will tend to have either large or small loadings of any particular variable. While this solution makes it rather easy to identify factors and their loading on items, the independency condition of orthogonal rotation techniques is hard to achieve. The assumption of independency of factors, especially in the context of behavioral research, belittles the value of orthogonal rotation techniques, to the point that “we see little justification for using orthogonal rotation as a general approach to achieving solutions with simple structure” [37, p. 283].

Oblique rotation is preferred for behavioral software engineering studies, as it is sensible to assume that behavioral, cognitive, and affective factors are separated by soft walls of independence (e.g., motivation and job satisfaction) [37, 120, 121]. If any, one would have to first conduct an investigation using oblique rotation and observe if the solution shows little to no correlation between factors and, in that case, switch to orthogonal rotation [37]. The two most employed and recommended oblique rotation techniques are Direct Oblimin (and its slight variation Direct Quartimin) and Promax, both of which perform well [37].

Fabrigar et al. [37], Russell [120] recommended to use a Promax rotation because it provides the best of both approaches. A Promax rotation first performs an orthogonal rotation (a Varimax rotation) to maximize the variance of the loadings of a factor on all the variables [120]. Then, Promax relaxes the constraint that the factors are independent from each other, turning the rotation to oblique. The advantage of this technique is that it will reveal whether factors really are uncorrelated with each other [120].

7.1.4 Further Recommendations. There have been several recommendations regarding the required sample size, number of measures per factor, number of factors to retain, and interpretation of loadings [120, 129, 152].

The recommended overall sample size as reported by Yong and Pearce [152] is at least 300 participants, with each variable that is subject to factor analysis with at least 5 to 10 observations. This recommendation has, however, low empirical validation. As reported by Russell [120], a Monte Carlo study by MacCallum et al. [88] analyzed how different sample sizes and communalities of the variables were able to reproduce the population factor loadings. They found that with item communalities higher or equal to 0.60, results were very consistent with sample sizes as low as 60 cases. Communality levels around 0.50 required samples of 100 to 200 cases. In their review, Russell [120] also found that 39% of EFA studies involved samples of 100 or fewer cases.

On the number of measures (items) per factor, Yong and Pearce [152] report that for something to be labeled as a factor it should have at least three variables, especially in cases when factors receive a rotation treatment, where only a high correlation (coefficient higher than 0.70) with each other, and mostly uncorrelation with other items would make them worthy of consideration.

Generally speaking, the correlation coefficient for an item to belong to a factor should be 0.30 or higher [131]. Russell [120] identifies that prior work has requested at least three items per factor; however, four or more items per factor was found to be a better holistic way to ensure an adequate identification of the factors. In his review, he identified 25% studies with three or fewer measures per factor.

We reported on the number of factors to retain in Section 7.1.2, so we will not repeat ourselves here. There is not a recommended number, and one would follow (possibly more than) one extraction method to identify the best number of factors according to the case. Tabachnick et al. [131] add that cases with missing values should be deleted to prevent overestimation of factors. Russell [120] wrote something that is worthy of mentioning for the uninitiated behavioral software engineering researchers, that is, even when constructing a new measurement instrument there is already an expectation of possible factors in the mind of the researcher. The reason is that items are developed following an investigation of prior work and/or empirical data (see Section 6). That number is a good starting point to base ourselves when conducting EFA.

Yong and Pearce [152] spend some further explanations on interpretation of loadings when they are produced by a statistical software. There should be few item crossloadings (i.e., split loadings, when an item loads at 0.32 or higher with two or more factors) so each factor defines a distinct cluster of interrelated variables. There are exceptions to this that require an analysis of the items. Sometimes it is useful to retain an item that crossloads, with the assumption that it is the latent nature of the variable. Furthermore, Tabachnick et al. [131] report that, with an alpha level of 0.01, a rotated factor loading with a meaningful sample size would need to have a value of at least 0.32 for loadings, as this would correspond to approximately 10% of the overlapping variance.

As a final note, the reader might now be left wondering, why conducting item analysis to reduce items if factor analysis is available? Kline [75] argues that the so many phases of factor analysis foresee assumptions of data distribution and decisions on techniques, that much could go wrong with it. Furthermore, items selected by item analysis are highly correlated to items selected by factor analysis [99], making item analysis a cheaper, effective *initial* method for item reduction.

7.2 Confirmatory Factor Analysis

Contrary to exploratory factor analysis, **confirmatory factor analysis (CFA)** is for confirming *a priori* hypotheses and assessing how well a hypothesized factor structure fits the obtained data [120]. A hypothesized factor could be derived from existing literature as well as data from a previous study to explore the factor structure.

Once the data is obtained to compare to the hypothesized factor structure, a goodness-of-fit test should be conducted. CFA requires statistical modeling that is outside the scope of this article and the estimation of the goodness-of-fit in CFA is a long-lasting debate, as “there are perhaps as many fit statistics as there are psychometricians” ([113, p. 31]). Russell [120], Rust [121], Singh et al. [129] provide several techniques for estimating the goodness-of-fit in CFA, e.g., Chi-squared test, **root mean square residual (RMSR)**, **root mean square error of approximation (RMSEA)**, and the standardized RMSR. Statistical software implement these techniques, including R psych package [116]. A widely employed technique for CFA is to be found in **structural equation modeling (SEM)**, which is a family of models to fit networks of constructs [72]. MacCallum and Austin [87] provided a comprehensive review of SEM techniques in the psychological science including their applications and pitfalls.

Conducting both EFA and CFA is very expensive. When designing and validating a measurement instrument, and when obtaining a large enough sample of participants, it is common to split the sample for conducting EFA on a part of it and CFA on the remaining part [129]. Most authors,

however, prefer conducting EFA only [120] and rely on future independent studies towards a better psychometrics evaluation of a tool. This is also why statistical tools, e.g., R psych package [115] provide estimates of fits for EFA as well as convenience tools to adapt the data to CFA packages, e.g., R sem [45].

We refer the reader to a prior work of ours in the behavioral software engineering domain [81] where we conducted a CFA and describe its application. We also note that, like for EFA, there have also been Bayesian methods for CFA proposed [5, 78].

8 STATISTICAL PROPERTIES OF ITEMS

Assessing characteristics and performance of individuals poses several challenges when interpreting the resulting scores. One of them is that a raw score is not meaningful without understanding the test standardization characteristics [75]. For example, a score of 38 on a debugging performance test is meaningless without knowing that 38 means to be able to open a debugger only. Furthermore, the interpretation of the results vary wildly when knowing that on average, developers score 400 on the test itself compared to if they score 42. The former issue is related to criterion-referenced standardization, the latter to norm-referenced standardization [75, 121].

Criterion-referenced tests assess what an individual with any score is expected to be able to do or know. Norm-referenced standardization enables to compare an individual's score to the ordered ranking of a population (also see Section 2). We concentrate on norm-referenced standardization, as criterion-referenced standardization is unique to a test criteria.

A first step to norm-reference a test is to order the results of all participants and rank an individual's score. Measures such as median and percentiles are useful for achieving the ranking and compare. When we can treat our data as interval scales and have it approximately following a normal distribution, we can also use the mean and the standard deviation. The standard deviation is useful for telling us how much an individual's score is above or below the mean score. Instead of reporting that an individual's score is, e.g., 13 above the mean score, it is more interesting to know that the score is 1.7 standard deviations above the mean score. Hence, we norm-standardize scores using different approaches. The remaining of this section is modeled after Rust [121] text and augmented by further explanations from other sources.

8.1 Standardization to Z Scores

Whenever a sample approximates a normal distribution, we know that a score above average is in the upper 50%, and by following the three sigma empirical rule [107], we know that a score greater than one or two standard deviations from the mean is in the top 68% and 95%, respectively. For expressing an individual's score in terms of how distant it is from the mean score, we transform the value to its Z score (also called standard score) using the formula in Equation (4):

$$Z_{score}(x_{pc}) = \frac{x_{pc} - \bar{x}_{pc}}{s_{pc}}, \quad (4)$$

where x_{pc} is a participant's score, $\bar{x}_{pc}$ is the mean of all participants' scores, and s_{pc} is the standard deviation for all participants' scores.

The ideal case would be to use the population mean and standard deviation. In software engineering research, we lack studies estimating population characteristics (an example of norm studies was provided by Reference [50]), so we should either aggregate the results of some studies or gather more samples.

An important note is that transforming scores into a Z scores does not make the scores normally distributed. This would require a normalization procedure, explained below.

Table 2. Conversion from Z Scores to Stanine and Sten Scores

Stanine score	Z score	Sten score	Z score
1	<-1.75	1	<-2.00
2	-1.75 to -1.25	2	-2.00 to -1.50
3	-1.25 to -0.75	3	-1.50 to -1.00
4	-0.75 to -0.25	4	-1.00 to -0.50
5	-0.25 to +0.25	5	-0.50 to -0.00
6	+0.25 to +0.75	6	+0.00 to +0.50
7	+0.75 to +1.25	7	+0.50 to +1.00
8	+1.25 to +1.75	8	+1.00 to +1.50
9	>+1.75	9	+1.50 to +2.00
		10	>+2.00

8.2 Standardization to T Scores

Z scores typically range between -3.00 and $+3.00$. The range is not always suitable for its application. A software developer could, for example, object to a Z score of -0.89 , which, at first glance, might suggest to be low (value) or negative (sign).

A T score, not to be confused with t-statistics of the Student's t-test, is a standard Z score that is scaled and shifted so it has a mean of 50 and a standard deviation of 10. T scores thus typically range between 20 and 80.

$$T_{score}(Z_{pc}) = Z_{pc} * 10 + 50 \quad (5)$$

For transforming a Z score into a T score, we use Equation (5).

The software developer in the previous example would have a T score of 41.1 from a Z score of -0.89 .

8.3 Standardization to Stanine and Sten Scores

Stanine and sten scores respond to the need of transforming a score to a scale from 1 to 9 (stanine) or 10 (sten) with a mean of 5 (stanine) or 5.5 (sten) and a standard deviation of 2. These scores purposely lose precision by keeping only decimal values.

The conversion to stanine and sten scores follows the rules in Table 2.

The advantage of stanine and sten scores lies in their imprecision. If our non performing developer with a Z score of -0.89 was compared with other two developers having scores of -0.72 and -0.94 , then how meaningful would be such a tiny difference in scores? Their stanine scores are 3, 4, and 4, respectively. Their sten score would be 4. Stanine and sten scores provide clear cut-off points for easier comparisons.

There is an important difference between stanine and sten scores, besides their range. A stanine score of 5 represents an average score in a sample. An average sten score can not be obtained, because the value of 5.5 does not belong to its possible values. A score of 4 represents the low average band (which ranges from 4.5 to 5.5, which is one standard deviation below the mean), and a score of 6 represents the high average band (which ranges from 5.5 to 6.5, which is one standard deviation above the mean).

8.4 Normalization

The standardization techniques that we presented in the previous section carry the assumption that the sample and population approximate the normal distribution. For all other cases, it is possible to normalize the data. Examples include algebraic transformation, e.g., square-rooting or log

Fig. 4. Reliability in psychometric theory.

transformation, as well as graphical transformation. See introductory statistical texts for more detailed explanations of and a broad set of such transformations.

9 RELIABILITY

Reliability can be seen either in terms of precision, which is the consistency of test scores across replications of the testing procedure (reliability/precision), or as a coefficient, which is the correlation between scores on two equivalent forms of the test (reliability/coefficients) [3]. For evaluating the precision of a measurement instrument, it would be ideal to have as many independent replications of the testing procedure as possible on the same very large sample. Scores are expected to generalize over alternative tests forms, contexts, raters, and over time. The reliability/precision of a measurement instrument is then assessed through the range of differences of the obtained scores. The reliability/precision of an instrument should be assessed with as many sub-groups of a population as possible.

The reliability/coefficients of a measurement instrument, which we will simply call reliability from this point on, is the most common way to refer to the reliability of a test [3]. There are three categories of reliability coefficients, namely, alternate-form (derived by administering alternative forms of test), test-retest (same test on different times), and internal-consistency (relationship among scores derived from individual test items during a single session).

We adhere to Nunnally [98], Rust [121] classification of reliability and provide a brief overview of reliability facets in psychometric theory in Figure 4.

Several factors, as defined in the SEPT [3], affect the reliability of a measurement instrument, especially adding or removing items, changing wording or intervals of items, causing variations in the constructs to be measured (e.g., using a measurement instrument for happiness to assess job satisfaction of developers), and ignoring variations in the intended populations for which the test was originally developed.

We now introduce the most widely employed techniques for establishing the reliability of a test.

9.1 Test-retest Reliability

Test-retest reliability, also known as test stability, is assessed when administering the measurement instrument twice to the same sample within a short interval of time. The paired set of scores for each participant is then compared with a correlation coefficient such as Pearson product-moment correlation coefficient or Spearman's rank-order correlation. A correlation coefficient of 1.00, while

rare, would indicate a perfect test-retest reliability, whereas a correlation coefficient of 0.00 would indicate no test-retest reliability at all. A negative score is no good news either, and it is automatically considered as a value of 0.00.

9.2 Parallel Forms Reliability

Test-retest reliability is not suitable for certain tests, such as those assessing knowledge or performance in general. Participants either face a learning or motivation effect from the first test session or simply improve (or worsen) their skills between sessions. For such cases, the parallel forms method is more suitable. The technique requires a systematic development of two versions of the same measurement instrument, namely, two parallel tests, that are assessing the same construct but using different wording or items. Parallel tests for assessing debugging skills would feature the same sections and amount of items, e.g., arithmetic, logic, and syntax errors. The two tests would need different source code snippets that are, however, very similar. A trivial example would be to test for unwanted assignments inside conditions in different places and with different syntax (e.g., using `if (n = foo())` in version one and `if (x = y + 2)` in version two). As with test-retest reliability, each participant faces both tests and a correlation coefficient can be computed.

9.3 Split-half Reliability

Split-half reliability is a widely adopted and more convenient alternative to parallel forms reliability. Under this technique, a measurement instrument is split into two half-size versions. The split should be as random as possible, e.g., splitting by taking odd and even numbered items. Participants face both halves of the test and, again, a correlation coefficient can be computed. The obtained coefficient, however, is not a measure of reliability yet. The reliability of the whole measurement instrument is computed with the Spearman-Brown formula in Equation (6):

$$R_{instrument} = 2 * \frac{r_{half}}{1 + r_{half}}, \quad (6)$$

where r_{half} is the correlation of the split tests. This formula shows that the more discriminating items a test has, the higher will be its reliability.

9.4 Inter-rater Reliability

Inter-rater reliability is perhaps the most common reliability that is found in software engineering studies. Qualitative studies or systematic literature reviews and mapping studies often have different raters for evaluating the same items. The sets of rates can be assessed using a correlation coefficient. Cohen's k is widely used in the literature for inter-rater coefficient of two entities together with Fleiss' k for the inter-rater coefficient of more than two entities. Cases have also been made for using Krippendorff's α [34].

9.5 Standard Error of Measurement

The standard error of measurement is used for generating confidence intervals around reported scores. The score is strictly related to the reliability coefficient [121], as shown in Equation (7):

$$E_{measurement} = var_{test} * (1 - r_{test}), \quad (7)$$

where var_{test} is the variance of the test scores, and r_{test} is the calculated reliability coefficient of the test. The standard error of measurement also provides an idea of how errors are distributed around observed scores. The standard error of measurement is maximized—and becomes equal to the standard deviation of the observed scores—when a test is completely unreliable. The standard error of measurement is minimized to zero when a test is perfectly reliable.

Fig. 5. Validity in psychometric theory.

If the assumption that errors are distributed normally is met, one can calculate the 95% confidence interval by using the z curve value of 1.96 to construct the interval $x_{obs} \pm E_{measurement} * 1.96$. Confidence intervals could also be used to compare participants' scores. Should one participant score fall below or above the interval, their results would differ significantly from the normality of scores.

10 VALIDITY

Validity in psychometrics is defined as “The degree to which evidence and theory support the interpretation of test scores for proposed uses of tests” [3]. Psychometric validity is therefore a different (but related) concept than the one of study validity that software engineers are used to dealing with [41, 104, 127, 147]. Validation in psychometric research is related to the interpretation of the test scores. For validating a test, we should gather relevant evidence for providing a sound scientific basis for the interpretation of the proposed scores.

Kline [75], Rust [121] have summarized six major facets of validity in the context of psychometric tests, which we summarize and represent in Figure 5 and describe below, augmented with references to material that offers additional explanations.

10.1 Face Validity

Face validity concerns how the items of a measurement instrument are accepted by respondents. For example, software developers expect the wording of certain items to be targeted to them instead of say, children. Similarly, if a test presents itself to be about a certain construct, such as debugging expertise, then it could cause face validity issues if it contained a personality assessment.

10.2 Content Validity

Content validity (sometimes called criterion validity or domain-referenced validity) concerns the extent to which a measurement instrument reflects the purpose for which the instrument is being developed. If a test was developed under the specifications of job satisfaction, but measured developers' motivation instead, then it would present issues of content validity. Content validity is evaluated qualitatively most of the times, because the form of deviation matters more than the degree of deviation, but there are proposals for its quantitative estimation [149].

10.3 Predictive Validity

Predictive validity is a statistical validity defined as the correlation between the score of a measurement instrument and a score of the degree of success of the selected field. For example, the

degree of success of debugging performance capability is expected to be higher with a higher programming experience. Computing a score for predictive validity is as simple as calculating a correlation value (such as Pearson or Spearman). According to the acceptance criterion for predictive validity, a score higher than 0.5 could be considered an acceptable predictive validity for the items. We would then feel justified in including programming experience as an item to represent the construct of debugging performance capability.

10.4 Concurrent Validity

Concurrent validity is a statistical validity that is defined as the correlation of a new measurement instrument and existing measurement instruments for the same construct. A measurement instrument tailored to the personality of software developers ought to correlate with existing personality measurement instruments. While concurrent validity is a common measure for test validity in psychology, it is a weak criterion, as the old measurement instrument itself could have a low validity. Nevertheless, concurrent validity is important for detecting low validity issues in measurement instruments.

10.5 Construct Validity

Construct validity is a major validity criterion in psychometric tests. As constructs are not directly measurable, we observe the relationship between the test and the phenomena that the test attempts to represent. For example, a test that identifies highly communicative team members should have a high correlation with...observations of highly communicative people who have been labelled as such. The nature of construct validity is that it is cumulative over the number of available studies [123].

10.6 Differential Validity

Differential validity assesses how a measurement instrument correlates with measures from which it should differ, and how a measurement instrument correlates with measures from which it should not differ. In particular, Campbell and Fiske [13] have differentiated between two aspects of differential validity, namely, convergent and discriminant validity. Rust [121] mentions a straightforward example of both. A test of mathematics reasoning should correlate positively with a test of numerical reasoning (convergent validity). However, the mathematics test should not strongly correlate positively with a test of reading comprehension, because the two constructs are supposed to be different (discriminant validity). In case of a low discriminant validity, there should be an investigation of whether the correlation is a result of a wider underlying trait, say, general intelligence. Differential validity is overall empirically demonstrated by a discrepancy between convergent validity and discriminant validity.

11 FAIRNESS IN TESTING AND TEST BIAS

Fairness is “the quality of treating people equally or in a way that is right or reasonable” (Cambridge [12]). A test is fair when it reflects the same constructs for all participants, and its scores have the same meaning for all individuals of the population [3]. A fair test does neither advantage nor disadvantage any participant through characteristics that are not related to the constructs under observation. From a participant point of view, an unfair test brings a wrong decision based on the test results. An example of a test that requires fairness is an attitude or skills assessment when interviewing candidates for hire in an information technology company.

The SEPT [3] reports on several facets of fairness. Individuals should have the opportunity to maximize how they perform with respect to the constructs being assessed. Similarly, for a measurement instrument that assesses traits of participants, the test should maximize how it assesses

Fig. 6. Bias in psychometric theory.

that the constructs being measured are present among individuals. This fairness comes from how the test is administered, which should be as standardized as possible. Research articles should describe the environment for the experimental settings, how the participants were instructed, which time limits were given, and so on. Fairness also comes, however, from participants themselves. Participants should be able to access the constructs as being measured without being advantaged or disadvantaged by individual characteristics. This is an issue of accessibility to a test and is also part of limiting item, test, and measurement bias.

We provide an overview of bias in psychometric theory in Figure 6.

Rust [121] provides an overview of item, test, and measurement bias, which we supplement with related work. It almost feels unnecessary to state that a measurement instrument should be free from bias regarding age, sex, gender, and race. These cases are indeed covered by legislation to ensure fairness. In general, there are three forms of bias in tests, namely, item bias, intrinsic test bias, and extrinsic test bias [121].

11.1 Item Bias

Item bias, also known as differential item functioning, refers to bias born out of individual items of the measurement instrument. A straightforward example would be to test a (non-UK) European developer about coding snippets dealing with imperial system units. A more common item bias is about the wording of items. Even among native speakers, the use of idioms such as double negatives can cause confusion. Asking a developer to mark a coding snippet that *is free from logic and syntax errors* is clearer than asking to mark code that *does not possess neither logic nor syntax errors*.

A systematic identification of item bias that goes beyond carefully checking an instrument is to carry out an item analysis with all possible groups of potential participants, for example men and women or speakers of English on different levels. A comparison of the facility values (the proportion of correct answers) of each item can reveal potential item bias. For instruments that assess traits and characteristics of a group instead of function or skills, a strategy is to follow a checklist of questions that researchers and pilot participants can answer [60].

Differential item functioning (DIF) is a statistical characteristic of an item that shows potential unfairness of the item among different groups that should provide same test results otherwise [103]. A presence of DIF does not necessarily indicate bias but unexpected behavior on an item [3]. This is why, after the detection of DIF, it is important to review the root causes of the differences. Whenever DIF happens for many items of a test, a test construct or final score is potentially unfair among different groups that should provide same test results otherwise. This situation is called **differential test functioning (DTF)** [118]. There are three main techniques for identifying DIF, namely, Mantel-Haenszel approach, **item response theory (IRT)** methods, and logistic regression [153].

11.2 Intrinsic Test Bias

Intrinsic test bias occurs when there are differences in the mean scores of two groups that are due to the characteristics of the test itself rather than difference between the groups in the constructs being measured. Measurement invariance is the desired property upon lack of which intrinsic test bias occurs. If a test for assessing the knowledge of software quality is developed in English and then administered to individuals who are not fluent in English, then the measure for the construct of software quality knowledge would be contaminated by a measure of English proficiency. Differential content validity (see Section 10.2) is the most severe form of intrinsic test bias, as it causes lower test scores in different groups. If a measurement instrument for debugging skills has been designed by keeping in mind American software testers, then any participant that is not an American software tester will likely perform worse on the test to different degrees. Rust [121] reports various statistical model proposals over the past 50 years to detect intrinsic test bias, which, however, present various issues including the introduction of more unfairness near cut-off points or for certain groups of individuals. There is not a recommended way to detect intrinsic test bias other than perform item bias analysis paired with sensitivity analysis.

11.3 Extrinsic Test Bias

Extrinsic test bias occurs whenever unfair decision happens based on a non-biased test. These issues usually belong to tests about demographics dealing with social, economical, and political issues, so they are unlikely to belong to measurement instruments developed for the software engineering domain.

12 FURTHER READING

The present article only scratches the surface of psychometric theory and practice, and it is its aim to be broad rather than deep. We collect, in this section, what we consider to be good next steps for a better understanding and expansion of the concepts that we have presented.

The books written by Coaley [20], Kline [75], Nunnally [98], Rust [121] provide an overall overview of psychometric theory, cover all topics mentioned in the present article, and more. In particular, we invite to compare how they present measurement theory and their views and classifications of validity and reliability. A natural followup is The SEPT, [3], which proposes standards that should be met in psychological testing.

While our summary breaks down fundamental concepts and presents them in an introductory way for researchers of behavioral software engineering, our writing can not honor enough the guidelines and recommendations for factor analysis offered by Fabrigar et al. [37], Russell [120], Singh et al. [129], Yong and Pearce [152]. To those, we add the work of Zumbo [153], which explored, through data simulations, the conditions that yield reliable exploratory factor analysis with samples below 50, which is unfortunately a condition we often live with in software engineering research. Furthermore, we wish to point the reader to alternatives to factor analysis, especially

for **confirmatory factor analysis (CFA)**. Flora and Curran [44] analyzed benefits when using **Robust Weighted Least Squares (Robust WLS)** regression. With a Monte Carlo simulation, they have shown that robust WLS provided accurate test statistics, parameter estimates, and standard errors even when the assumptions of CFA were not met. Bayesian alternatives for CFA were proposed in the early '80s [78] and later expanded to cover the exploratory phase as well; see, for example, the works by Conti et al. [24], Lu et al. [86], Muthén and Asparouhov [95].

In the above sections, we have pointed to several papers that can provide a modern, Bayesian statistical view of many psychometric analysis procedures. We also note that a more general treatment and overview can be found in Levy and Mislevy [82]. While it is important for a software engineering researcher who wants to use and develop psychometric instruments to know the key concepts and techniques of the more classical, typically frequentist, psychometric techniques one can then switch to a Bayesian view either on philosophical or for practical reasons (a simpler more unified treatment, for one).

Within the software engineering domain, Gren [56] has offered an alternative lens on validity and reliability of software engineering studies, also based on psychology, that we advise to read. Ralph and Tempero [109] have offered a deep overview of construct validity in software engineering through a psychological lens.

We do want to note that one aspect of the method we have used, which can be seen as a limitation, is that there is much current discussion about the statistical methods that are and/or should be applied in behavioral and social science, including psychology [122, 139], as well as in applied sciences in general [142]. This has also affected software engineering and, for example, a recent paper argued for transitioning to Bayesian statistical analysis in empirical software engineering [47]. However, it is too early to base guidelines on proposals in this ongoing, scholarly discussion, since there is not yet a clear consensus. Thus, since we base our review on the current and more established literature, it is likely that future work will need to consider more powerful and up-to-date statistical methods for the creation and assessment of psychometric instruments. Thus, we foresee future updates to this article that extend it by using such, more recent analysis methods.

13 LIMITATIONS OF PSYCHOMETRICS

Our article has insofar showed appreciation of psychometrics and how they could be adopted in the software engineering research fields. We should balance this with reasons for not adopting psychometrics, which we also recommend as a further reading.⁴

Psychometrics are not universally accepted as a perfect tool for assessing psychological constructs. Fronts of discontents of psychometrics have emerged, especially in recent times and in the medical, social, and education fields [96, 124, 138].

The widest critique to psychometric-based testing is that it might not be the best tool for assessing individuals. On the technical side, there has been some evidence that alternative evaluation systems that matched individuals to a standardized set of holistic and realistic vignettes improved discrimination of the individual's performance and facilitated the identification of severe, medical issues of some participants [111]. The technical side of critique is, however, limited, as psychometric theory has a long history of robust statistical methods. What causes the widest concern with

⁴Following principles of open science, we deposited revisions of the present article on arXiv as we wrote it, submitted for peer review, and revised during peer review. One of the good aspects of depositing manuscripts on arXiv is to attract feedback. A previous version of the present article, deposited on arXiv with identifier [2005.09959v2](https://arxiv.org/abs/2005.09959v2), was used as a basis by Lewis [83] to offer a critique of our stance on psychometrics as well as a broader warning on relying on psychometrics in absolute. The present article takes Lewis's feedback into account, especially on the wording and on the limitations of psychometrics. We recommend reading Lewis's paper, as it offers a view on psychometrics that, on several points, is opposed to ours.

psychometrics is the central argument, summarized by Schoenherr and Hamstra [124], that following a psychometric approach might cause to focus too narrowly on characteristics of individuals in terms of dimensions, features, and competences, while at the same time missing on context, uniqueness of individuals, and team perspective. In other words, psychometric-based assessments potentially neglect a richness of information, much of which comes from direct experience with individuals and, in particular, qualitative data [124]. This missed opportunity may be essential for taking decisions such as related to allocation of resources, promotions, and development or prioritization of skills.

Uher [137] has argued that discussions for and against psychometrics so far have swayed on the “qualitative vs. quantitative data” debate, which links to the issue of clashing worldviews in research. Lewis [83] has expressed skepticism that psychometric methods can be applied in software engineering in a successful way, one of the reasons being that an “empirical study that is purely empirical cannot succeed” (p. 40). This pushback can go to an opposite extreme, with Michell [91] defining psychometrics as *pathological science* because of a lack of testing of the hypothesis that some psychological attributes are quantitative at all and psychometric science is based on accepting this hypothesis into its core. We hope to meet the reader’s mercy if we do not cover these topics in a concluding section,⁵ especially because we do not see any winning idea regarding epistemological stances. In short, as concisely summarized by Schoenherr and Hamstra [124], adopting psychometric approaches can be met by a clash between postpositivism and constructivism, or, the debate concerning the adequacy of mapping constructs to numbers and the assumptions that such a process implies. We agree with Lewis [83] that the idea of psychometrics, as well as the one to add numbers to people, is seductive. We do not see this issue confined to psychometrics, though. Our stance is that, should one decide to embark in a fully empirical and quantitative approach to assess psychological constructs in software engineering, the psychometric approach is an effective tool. There is little debate on the effectiveness of psychometric tests as tools. Schoenherr and Hamstra [124], indeed, argue that “we should not seek to establish a post-psychometric era” (p. 720) and that we should rather focus on the lack of understanding of the theoretical context in which measurement instruments tools have been developed as well as the several, sometimes unresolved, debates within the psychometrics discourse. Our article is a start with this.

On the issue of assigning numbers to individuals, we agree that quantitative studies are not all there is. We want to emphasize the importance of qualitative studies the same way as psychology is rediscovering them. Thurstone [133], in the plea to render psychological science quantitative, mathematical, and robust, was clear in specifying that “the mathematical and rigorous treatment of a psychological problem is in no sense a substitute for the other forms of exploratory and descriptive types of experimentation” (p. 228). This quote should remind us that quantitative methods, while important, are all but one side of a multifaceted coin that we should supply with a mixed-method mindset that includes qualitative studies and other disciplines. As anticipated in the opening of this article, we are interested in qualitative research as well. We have offered our proposals for qualitative behavioral software engineering [79]. In Lenberg et al. [79], we suggest future research in software engineering will benefit from a broader set of methods from qualitative psychology, such as interpretive phenomenological analysis, narrative analysis, and discourse analysis. We invite researchers in qualitative studies of software engineering to emphasize reflexivity on how our thinking came to be and how pre-existing understanding is constantly revised in the light of new insights. Finally, we encourage the adoption of qualitative guidelines and criteria to enhance the quality of qualitative studies.

⁵We refer the reader to the Creswell and Creswell [27] textbook for an introduction to ontology and epistemology, the work of Koltko-Rivera [76] for the psychology of worldviews, and the studies by Petersen and Gencel [104], Siegmund et al. [127] for worldviews and their influence in software engineering research.

To sum up, we see discontent with psychometrics about clashing worldviews, ontological and epistemological issues, and on how psychometric tests are used in decision-making processes as razor-blade cut-off methods for taking complex decisions.

There are also limitations brought by following psychometric approaches that are related to practical challenges. The field is enormous—even our introductory paper has passed 150 referenced sources—with many theoretical and practical contributions that require a wide understanding of the concepts at hand, several of which are taken as implicit. Uher [138], indeed, argues that much issues with psychometrics arise from psychological jargon and underlying, codified conceptual fallacies that bring misconceptions. We believe that basic introductory contributions such as ours help coping with this issue, and we welcome more. Uher [138] provides a comprehensive introduction on terms such as psyche, behavior, constructs, operationalization, variables, attributes, and more, which we recommend reading. Furthermore, we highlight the considerable amount of effort it takes for developing a psychometrically validate test. Gren [56] has argued that “Spending an entire Ph.D. candidacy on the validation of one single measurement of a construct should be, not only approved, but encouraged.” We agree on that enabling a Ph.D. student to develop a measurement instrument for software engineering constructs following a psychometric approach should be allowed, but we want to emphasize here the “entire Ph.D.” part. It might truly take years to develop and validate a measurement instrument. We highly recommend searching the literature for existing, psychometrically validated tests before developing one.

All in all, we emphasize our view that we should be systematic and rigorous regardless of world-view, ontological and epistemological stance, and preferred methods. Psychometrics is one of several ways of quantitative research, which we embrace but not blindly believe in. Qualitative research is as important as quantitative research and they complement each other. Qualitative research is essential to discover rich data that captures the uniqueness of individuals as such, rather than pooling people as unique data points. Psychometric methods are an excellent tool to enhance our way to select, develop, understand, and test metrics. Psychometrics should, however, not be justification to build a metric for the sake of it. Proper application of psychometrics does also not promote the validity of results in absolutism. We remind that validity in psychometrics is defined as “The degree to which evidence and theory support the interpretation of test scores for proposed uses of tests” [3] for which we now highlight the word **interpretation**.

14 RUNNING EXAMPLE OF PSYCHOMETRIC EVALUATION

We believe that a methodology description is best complemented by a concrete example of its application. In an appendix, we thus provide a complete scenario of the development of a fictitious measurement instrument and the establishment of its psychometric properties with the R programming language. The evaluation follows the same structure as the present article for ease of understanding. In the spirit of open science in software engineering [43], we provide the running example as a replication package as well [51]. We wrote the example using R Markdown, making it fully repeatable, as well as the generated dataset, and instructions for replication with newly generated data.

15 CONCLUSION

The adoption and development of validated and reliable measurement instruments in software engineering research, whenever humans are to be evaluated, should benefit from psychology and statistics theory; we need not and should not “reinvent the wheel.” This article provides a brief introduction to the evaluation of psychometric tests. Our guidelines will contribute to a better development of new tests as well as a justified decision-making process when selecting existing tests.

After providing basic building blocks and concepts of psychometric theory, we introduced item pooling, item review, pilot testing, item analysis, factor analysis, statistical property of items, reliability, validity, and fairness in testing and test bias. In an appendix, we also provided a running example of an implementation of a psychometric evaluation and shared both its data and source code (scripts) openly to promote self-study and a basis for further exploration. We followed textbooks, method papers, and society standards for ensuring a coverage of all important steps, but we could only offer a brief introduction and invite the reader to explore our referenced material further. Each of these steps is a universe of its own, with dozens of published artifacts related to them.

A proper psychometric evaluation requires a consideration of all elements reported in the present article. The development of a measurement instrument, however, does not have to execute every single step we summarize, and not even necessarily in the same order. We are illustrating a wide range of possibilities, some of which are often brought by in future validation studies. Critical thinking is required throughout the whole process, which involves choices and tradeoffs.

Adding the steps described in this article will increase the time required for developing measurement instruments. However, the return on investment will be considerable. Psychometric analysis and refinement of measurement instruments can improve their reliability and validity. The software engineering community must value psychometric studies more. This, however, requires a cultural change that we hope to champion with this article.

ACKNOWLEDGMENTS

We are thankful for the comments of three anonymous reviewers and the associate editor, who helped the manuscript mature and grow in value during the peer review process. We are thankful to Marvin Wyrich and Kai Mindermann for offering feedback to early versions of this article. We are grateful for Lewis's comments in the open, which we treasured when revising the manuscript. Our thanks to Katharina Plett for proofreading the article.

REFERENCES

- [1] Hervé Abdi. 2003. Factor rotations in factor analyses. In *Encyclopedia of Social Sciences Research Methods*, A. Lewis-Beck M., Bryman and Futing T. (Eds.). SAGE, Thousand Oaks, CA, 792–795.
- [2] Arnold Alphen, Ruud Halfens, Arie Hasman, and Tjaart Imbos. 1994. Likert or Rasch? Nothing is more applicable than good theory. *J. Adv. Nurs.* 20, 1 (1994), 196–201.
- [3] American Educational Research Association, American Psychological Association, National Council on Measurement in Education, and Joint Committee on Standards for Educational and Psychological Testing (U.S.). 2014. *Standards for Educational and Psychological Testing*. American Educational Research Association, Washington, DC.
- [4] Anne Anastasi and Susana Urbina. 1997. *Psychological Testing*. Prentice Hall, Upper Saddle River, NJ.
- [5] Asim Ansari, Kamel Jedidi, and Laurette Dube. 2002. Heterogeneous factor analysis models: A Bayesian approach. *Psychometrika* 67, 1 (2002), 49–77.
- [6] Michael Berger. 2013. Criterion-referenced testing. In *Encyclopedia of Autism Spectrum Disorders*. Springer New York, 823–823. DOI: https://doi.org/10.1007/978-1-4419-1698-3_146
- [7] Michael Berger. 2013. Norm-referenced testing. In *Encyclopedia of Autism Spectrum Disorders*. Springer New York, 2063–2064. DOI: https://doi.org/10.1007/978-1-4419-1698-3_451
- [8] R. Boulkedid, H. Abdoul, M. Loustau, O. Sibony, and C. Alberti. 2011. Using and reporting the delphi method for selecting healthcare quality indicators: A systematic review. *PLoS One* 6 (2011), e20476. DOI: <https://doi.org/10.1371/journal.pone.0020476>
- [9] Norman M. Bradburn, Seymour Sudman, and Brian Wansink. 2004. *Asking Questions*. John Wiley & Sons, Hoboken, NJ.
- [10] Michael W. Browne. 2001. An overview of analytic rotation in exploratory factor analysis. *Multivar. Behav. Res.* 36, 1 (2001), 111–150.
- [11] Victor R. Basili, Gianluigi Caldiera, and H. Dieter Rombach. 1994. The goal question metric approach. *Encyc. Softw. Eng.* 1, 1 (1994), 528–532.

- [12] English Dictionary Cambridge. 2018. Fairness. *Cambridge Eng. Dict.* 1, 1 (2018), 1. Retrieved from <https://dictionary.cambridge.org/dictionary/english/fairness>.
- [13] Donald T. Campbell and Donald W. Fiske. 1959. Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychol. Bull.* 56, 2 (1959), 81–105.
- [14] Luiz Fernando Capretz. 2003. Personality types in software engineering. *Int. J. Hum.-comput. Stud.* 58, 2 (2003), 207–214.
- [15] Rudolf Carnap. 1962. *Logical Foundations of Probability*. University of Chicago Press, Chicago.
- [16] Jeffrey C. Carver, Henry Muccini, Birgit Penzenstadler, Rafael Prikladnicki, Alexander Serebrenik, and Thomas Zimmermann. 2021. Behavioral science and diversity in software engineering. *IEEE Softw.* 38, 2 (2021), 107–112. DOI: <https://doi.org/10.1109/ms.2020.3042683>
- [17] Raymond B. Cattell. 1966. The scree test for the number of factors. *Multivar. Behav. Res.* 1, 2 (1966), 245–276.
- [18] Souti Chattopadhyay, Nicholas Nelson, Audrey Au, Natalia Morales, Christopher Sanchez, Rahul Pandita, and Anita Sarma. 2020. A tale from the trenches. In *ACM/IEEE 42nd International Conference on Software Engineering*. ACM, New York, NY, 654–665. DOI: <https://doi.org/10.1145/3377811.3380330>
- [19] Marcus Ciolkowski, Oliver Laitenberger, Sira Vegas, and Stefan Biffl. 2003. Practical experiences in the design and conduct of surveys in empirical software engineering. In *Empirical Methods and Studies in Software Engineering*, Gerhard Goos, Juris Hartmanis, Jan van Leeuwen, Reidar Conradi, and Alf Inge Wang (Eds.). Springer Berlin, 104–128. DOI: https://doi.org/10.1007/978-3-540-45143-3_7
- [20] Keith Coaley. 2014. *An Introduction to Psychological Assessment & Psychometrics*. SAGE, Los Angeles, CA.
- [21] Ronald Jay Cohen, Mark E. Swerdlik, and Suzanne M. Phillips. 1995. *Psychological Testing and Assessment: An Introduction to Tests and Measurement*. Mayfield Pub Co., California City, CA.
- [22] D. Collins. 2003. Pretesting survey instruments: An overview of cognitive methods. *Qual. Life Res.* 12, 3 (2003), 229–238.
- [23] Gabriella Conti, Sylvia Frühwirth-Schnatter, James J. Heckman, and Rémi Piatek. 2014. Bayesian exploratory factor analysis. *J. Economet.* 183, 1 (2014), 31–57.
- [24] G. Conti, S. Frühwirth-Schnatter, J. J. Heckman, and R. Piatek. 2014. Bayesian exploratory factor analysis. *J. Econ.* 183, 1 (2014), 31–57.
- [25] Paul T. Costa Jr and Robert R. McCrae. 1992. Comes of age. In *Nebraska Symposium on Motivation: Psychology and Aging*. University of Nebraska Press, Lincoln, NE.
- [26] Matthew Gordon Rau Courtney. 2013. Determining the number of factors to retain in EFA: Using the SPSS r-menu v2.0 to make more judicious estimations. *Practic. Assess. Res. Eval.* 18, 8 (2013), 1–14.
- [27] John W. Creswell and J. David Creswell. 2018. *Research Design*. SAGE Publications, Incorporated, Thousand Oaks, CA.
- [28] Linda Crocker. 2008. *Introduction to Classical and Modern Test Theory*. Cengage Learning, Mason, OH.
- [29] Lee J. Cronbach. 1951. Coefficient alpha and the internal structure of tests. *Psychometrika* 16, 3 (1951), 297–334.
- [30] Shirley Cruz, Fabio Q. B. da Silva, and Luiz Fernando Capretz. 2015. Forty years of research on personality in software engineering: A mapping study. *Comput. Hum. Behav.* 46 (2015), 94–113. DOI: <https://doi.org/10.1016/j.chb.2014.12.008>
- [31] Norman Dalkey and Olaf Helmer. 1963. An experimental application of the DELPHI method to the use of experts. *Manag. Sci.* 9, 3 (1963), 458–467. DOI: <https://doi.org/10.1287/mnsc.9.3.458>
- [32] David P. Darcy and Meng Ma. 2005. Exploring individual characteristics and programming performance: Implications for programmer selection. In *38th Hawaii International Conference on System Sciences*. IEEE, Piscataway, NJ.
- [33] R. A. Darton. 1980. Rotation in factor analysis. *Statistician* 29, 3 (1980), 167.
- [34] Knut De Swert. 2012. *Calculating Inter-coder Reliability in Media Content Analysis Using Krippendorff's Alpha*. Technical Report. Center for Politics and Communication, University of Amsterdam.
- [35] J. Drennan. 2003. Cognitive interviewing: Verbal data in the design and pretesting of questionnaires. *J. Adv. Nurs.* 42, 1 (2003), 57–63. DOI: <https://doi.org/10.1046/j.1365-2648.2003.02579.x>
- [36] Susan E. Embretson and Steven P. Reise. 2013. *Item Response Theory*. Psychology Press, Hove, East Sussex, UK.
- [37] Leandre R. Fabrigar, Duane T. Wegener, Robert C. MacCallum, and Erin J. Strahan. 1999. Evaluating the use of exploratory factor analysis in psychological research. *Psychol. Meth.* 4, 3 (1999), 272.
- [38] Fabian Fagerholm. 2015. *Software Developer Experience: Case Studies in Lean-Agile and Open Source Environments*. Ph.D. Dissertation. Department of Computer Science, University of Helsinki, Helsinki, Finland.
- [39] Fabian Fagerholm and Thomas Fritz. 2020. Biometric measurement in software engineering. In *Contemporary Empirical Methods in Software Engineering*, Michael Felderer and Guilherme Horta Travassos (Eds.). Springer International Publishing, Cham, Switzerland, 151–172. DOI: https://doi.org/10.1007/978-3-030-32489-6_6
- [40] Fabian Fagerholm and Max Pagels. 2014. Examining the structure of lean and agile values among software developers. In *Lecture Notes in Business Information Processing: Agile Processes in Software Engineering and Extreme Programming*. Springer International Publishing, Cham, 218–233.

- [41] Robert Feldt and Ana Magazinius. 2010. Validity threats in empirical software engineering research—An initial survey. In *22nd International Conference on Software Engineering & Knowledge Engineering (SEKE'2010)*... Retrieved from http://www.cse.chalmers.se/~feldt/publications/feldt_2010_validity_threats_in_ese_initial_survey.pdf.
- [42] Robert Feldt, Richard Torkar, Lefteris Angelis, and Maria Samuelsson. 2008. Towards individualized software engineering: Empirical studies should collect psychometrics. In *International Workshop on Cooperative and Human Aspects of Software Engineering*. ACM Press, New York, NY, 49–52.
- [43] Daniel Méndez Fernández, Daniel Graziotin, Stefan Wagner, and Heidi Seibold. 2020. Open science in software engineering. In *Contemporary Empirical Methods in Software Engineering*, Michael Felderer and Guilherme Horta Travassos (Eds.). Springer International Publishing, Cham, Switzerland, 479–504. DOI: https://doi.org/10.1007/978-3-030-32489-6_17
- [44] David B. Flora and Patrick J. Curran. 2004. An empirical evaluation of alternative methods of estimation for confirmatory factor analysis with ordinal data. *Psychol. Meth.* 9, 4 (2004), 466.
- [45] John Fox, Zhenghua Nie, and Jarrett Byrnes. 2017. *sem: Structural Equation Models*. Technical Report. The Comprehensive R Archive Network.
- [46] Cesar Franca, Fabio Q. B. da Silva, and Helen Sharp. 2020. Motivation and satisfaction of software engineers. *IEEE Trans. Softw. Eng.* 46, 2 (2020), 118–140. DOI: <https://doi.org/10.1109/tse.2018.2842201>
- [47] Carlo Alberto Furia, Robert Feldt, and Richard Torkar. 2019. Bayesian data analysis in empirical software engineering research. *IEEE Trans. Softw. Eng.* 1, 1 (2019), 1–26. DOI: <https://doi.org/10.1109/TSE.2019.2935974>
- [48] Annie T. Ginty. 2013. Psychometric properties. In *Encyclopedia of Behavioral Medicine*, Marc D. Gellman and J. Rick Turner (Eds.). Springer New York, NY, 1563–1564. DOI: https://doi.org/10.1007/978-1-4419-1005-9_480
- [49] Robert Glaser. 1963. Instructional technology and the measurement of Learning outcomes: Some questions. *Amer. Psychol.* 18, 8 (1963), 519–521. DOI: <https://doi.org/10.1037/h0049294>
- [50] Daniel Graziotin, Fabian Fagerholm, Xiaofeng Wang, and Pekka Abrahamsson. 2017. On the unhappiness of software developers. In *21st International Conference on Evaluation and Assessment in Software Engineering*. ACM Press, New York, NY, 324–333.
- [51] Daniel Graziotin, Per Lenberg, Robert Feldt, and Stefan Wagner. 2021. Behavioral software engineering—Example of psychometric evaluation with R. DOI: <https://doi.org/10.5281/zenodo.3799603>
- [52] Daniel Graziotin, Xiaofeng Wang, and Pekka Abrahamsson. 2015. The affect of software developers: Common misconceptions and measurements. In *IEEE/ACM 8th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE'15)*. IEEE, Piscataway, NJ, 123–124.
- [53] Daniel Graziotin, Xiaofeng Wang, and Pekka Abrahamsson. 2015. Do feelings matter? On the correlation of affects and the self-assessed productivity in software engineering. *J. Softw.: Evol. Proc.* 27, 7 (2015), 467–487.
- [54] Daniel Graziotin, Xiaofeng Wang, and Pekka Abrahamsson. 2015. Understanding the affect of developers: Theoretical background and guidelines for psychoempirical software engineering. In *7th International Workshop on Social Software Engineering (SSE'15)*. ACM, New York, NY, 25–32. DOI: <https://doi.org/10.1145/2804381.2804386>
- [55] Christopher D. Green. 2009. Darwinian theory, functionalism, and the first American psychological revolution. *Amer. Psychol.* 64, 2 (2009), 75–83. DOI: <https://doi.org/10.1037/a0013338>
- [56] Lucas Gren. 2018. Standards of validity and the validity of standards in behavioral software engineering research. ACM Press, New York, NY. DOI: <https://doi.org/10.1145/3239235.3267437>
- [57] Lucas Gren and Alfredo Goldman. 2016. Useful statistical methods for human factors research in software engineering: A discussion on validation with quantitative data. In *9th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE'16)*. Association for Computing Machinery, New York, NY, 121–124. DOI: <https://doi.org/10.1145/2897586.2897588>
- [58] Joy Paul Guilford. 1954. *Psychometric Methods* (1st ed.). McGraw-Hill Book Co., New York, London.
- [59] Joy Paul Guilford. 1954. *Psychometric Methods* (2nd ed.). McGraw-Hill Book Co., New York, London.
- [60] Ronald K. Hambleton and Jane Rodgers. 1995. *Item Bias Review*. Technical Report. ERIC Clearinghouse on Assessment and Evaluation, the Catholic University of America, Department of Education.
- [61] Ronald K. Hambleton, Hariharan Swaminathan, and H. Jane Rogers. 1991. *Fundamentals of Item Response Theory*. Sage Publications, Newbury Park, CA.
- [62] Ernest R. Hilgard. 1980. The trilogy of mind: Cognition, affection, and conation. *J. Hist. Behav. Sci.* 16, 2 (1980), 107–117.
- [63] Charlotte Emma Hilton. 2017. The importance of pretesting questionnaires: A field research example of cognitive pretesting the exercise referral quality of life scale (ER-QLS). *Int. J. Soc. Res. Methodol.* 20, 1 (2017), 21–34. DOI: <https://doi.org/10.1080/13645579.2015.1091640>
- [64] Robert Hogan. 2007. *Personality and the Fate of Organizations*. Erlbaum, Mahwah, NJ.
- [65] John L. Horn. 1965. A rationale and test for the number of factors in factor analysis. *Psychometrika* 30 (1965), 179–185.

- [66] Junzhong Ji, Jingyue Li, Reidar Conradi, Chunnian Liu, Jianqiang Ma, and Weibing Chen. 2008. Some lessons learned in conducting software engineering surveys in china. In *2nd ACM-IEEE International Symposium on Empirical Software Engineering and Measurement*. ACM, Piscataway, NJ, 168–177.
- [67] George A. Johanson and Gordon P. Brooks. 2010. Initial scale development: Sample size for pilot studies. *Educ. Psychol. Measur.* 70, 3 (2010), 394–400. DOI : <https://doi.org/10.1177/0013164409355692>
- [68] Lyle V. Jones and David Thissen. 2006. A history and overview of psychometrics. In *Handbook of Statistics: Psychometrics*. Elsevier, Amsterdam Boston, 1–27. DOI : [https://doi.org/10.1016/s0169-7161\(06\)26001-2](https://doi.org/10.1016/s0169-7161(06)26001-2)
- [69] Henry F. Kaiser. 1958. The varimax criterion for analytic rotation in factor analysis. *Psychometrika* 23, 3 (1958), 187–200.
- [70] Henry F. Kaiser. 1960. The application of electronic computers to factor analysis. *Educ. Psychol. Measur.* 20, 1 (1960), 141–151.
- [71] Graham Kalton and Howard Schuman. 1982. The effect of the question on survey responses: A review. *J. Royal Statist. Soc.: Series A (Gen.)* 145, 1 (1982), 42–57.
- [72] David Kaplan. 2008. *Structural Equation Modeling: Foundations and Extensions*. Vol. 10. SAGE, Los Angeles, CA.
- [73] B. A. Kitchenham. 2007. *Guidelines for Performing Systematic Literature Reviews in Software Engineering*. Technical Report. Keele University and University of Durham Keele and Durham, UK.
- [74] Barbara A. Kitchenham and Shari L. Pfleeger. 2008. Personal opinion surveys. In *Guide to Advanced Empirical Software Engineering*, Forrest Shull, Janice Singer, and Dag I. K. Sjøberg (Eds.). Springer London, 63–92. DOI : https://doi.org/10.1007/978-1-84800-044-5_3
- [75] Paul Kline. 2015. *A Handbook of Test Construction (Psychology Revivals): Introduction to Psychometric Design*. Routledge, London New York.
- [76] Mark E. Koltko-Rivera. 2004. The psychology of worldviews. *Rev. Gen. Psychol.* 8, 1 (2004), 3–58. DOI : <https://doi.org/10.1021/bi00696a014>
- [77] G. J. Kootstra. 2006. *Exploratory Factor Analysis: Theory and Application*. Technical Report. University of Groningen.
- [78] Sik-Yum Lee. 1981. A Bayesian approach to confirmatory factor analysis. *Psychometrika* 46, 2 (1981), 153–160.
- [79] Per Lenberg, Robert Feldt, Lars Göran Wallgren Tengberg, Inga Tidefors, and Daniel Graziotin. 2017. Behavioral software engineering—Guidelines for qualitative studies. Retrieved from <https://arxiv.org/abs/1712.08341>.
- [80] Per Lenberg, Robert Feldt, and Lars Göran Wallgren. 2015. Behavioral software engineering: A definition and systematic literature review. *J. Syst. Softw.* 107 (2015), 15–37. DOI : <https://doi.org/10.1016/j.jss.2015.04.084>
- [81] Per Lenberg, Lars Göran Wallgren Tengberg, and Robert Feldt. 2017. An initial analysis of software engineers' attitudes towards organizational change. *Empir. Softw. Eng.* 22, 4 (2017), 2179–2205.
- [82] Roy Levy and Robert J. Mislevy. 2016. *Bayesian Psychometric Modeling*. CRC Press, Boca Raton, FL.
- [83] Clayton Lewis. 2020. On personality testing and software engineering. In *31st Workshop of the Psychology of Programming Interest Group (PPIG'20)*. PPIG, 39–41.
- [84] Rensis Likert. 1932. A technique for the measurement of attitudes. *Arch. Psychol.* 22, 40 (1932), 1–55.
- [85] Jane Loewinger. 1957. Objective tests as instruments of psychological theory. *Psychol. Rep.* 3 (1957), 635–694. DOI : <https://doi.org/10.2466/pr0.1957.3.3.635>
- [86] Z. H. Lu, S. M. Chow, and E. Loken. 2016. Bayesian factor analysis as a variable-selection problem: Alternative priors and consequences. *Multivar. Behav. Res.* 51, 4 (2016), 519–539.
- [87] R. C. MacCallum and J. T. Austin. 2000. Applications of structural equation modeling in psychological research. *Annu. Rev. Psychol.* 51 (2000), 201–226.
- [88] Robert C. MacCallum, Keith F. Widaman, Shaobo Zhang, and Sehee Hong. 1999. Sample size in factor analysis. *Psychol. Meth.* 4, 1 (1999), 84–99.
- [89] Geoffrey N. Masters. 1988. Item discrimination: When more is worse. *J. Educ. Measur.* 25, 1 (1988), 15–29. DOI : <https://doi.org/10.1111/j.1745-3984.1988.tb00288.x>
- [90] Sharon McDonald and Helen M. Edwards. 2007. Who should test whom. *Commun. ACM* 50, 1 (2007), 66–71.
- [91] Joel Michell. 2000. Normal science, pathological science and psychometrics. *Theor. Psychol.* 10, 5 (2000), 639–667.
- [92] Kristen Miller, Valerie Chepp, Stephanie Willson, and Jose-Luis Padilla. 2014. *Cognitive Interviewing Methodology*. John Wiley & Sons, Hoboken, NJ.
- [93] Jefferson Seide Molléri, Kai Petersen, and Emilia Mendes. 2016. Survey guidelines in software engineering: An annotated review. In *10th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*. IEEE, Piscataway, NJ.
- [94] W. P. Morgan. 1980. The trait psychology controversy. *Res. Q. Exerc. Sport* 51, 1 (1980), 50–76. DOI : <https://doi.org/10.1080/02701367.1980.10609275>
- [95] Bengt Muthén and Tihomir Asparouhov. 2012. Bayesian structural equation modeling: A more flexible representation of substantive theory. *Psychol. Meth.* 17, 3 (2012), 313–335.

- [96] Geoff Norman. 2002. Research in medical education: Three decades of progress. *BMJ* 324, 7353 (2002), 1560–1562. DOI : <https://doi.org/10.1136/bmj.324.7353.1560>
- [97] Megan Norris and Luc Lecavalier. 2010. Evaluating the use of exploratory factor analysis in developmental disability psychological research. *J. Autis. Devel. Disord.* 40, 1 (01 Jan. 2010), 8–20. DOI : <https://doi.org/10.1007/s10803-009-0816-2>
- [98] Jum Nunnally. 1994. *Psychometric Theory*. McGraw-Hill, New York.
- [99] Jum C. Nunnally. 1978. An overview of psychological measurement. In *Clinical Diagnosis of Mental Disorders*. Springer US, 97–146. DOI : https://doi.org/10.1007/978-1-4684-2490-4_4
- [100] A. N. Oppenheim. 1992. *Questionnaire Design, Interviewing and Attitude Measurement*. Pinter Pub Ltd, London, UK.
- [101] Jan-Peter Ostberg, Daniel Graziotin, Stefan Wagner, and Birgit Derntl. 2017. Towards the assessment of stress and emotional responses of a salutogenesis-enhanced software tool using psychophysiological measurements. In *Towards the Assessment of Stress and Emotional Responses of a Salutogenesis-Enhanced Software Tool Using Psychophysiological Measurements*. IEEE, Piscataway, NJ, 22–25. DOI : <https://doi.org/10.1109/semotion.2017.4>
- [102] Karl Pearson. 1901. LIII. On lines and planes of closest fit to systems of points in space. *London, Edinburgh, Dublin Philos. Mag. J. Sci.* 2, 11 (Nov. 1901), 559–572. DOI : <https://doi.org/10.1080/14786440109462720>
- [103] Michael Perrone. 2006. Differential item functioning and item bias: Critical considerations in test fairness. *Teach. Coll., Columbia Univ. Work. Pap. TESOL Appl. Ling.* 6 (2006), 1–3.
- [104] Kai Petersen and Cigdem Gencel. 2013. Worldviews, research methods, and their relationship to validity in empirical software engineering research. In *Joint Conference of the 23rd International Workshop on Software Measurement and the 8th International Conference on Software Process and Product Measurement (IWSM-MENSURA'13)*. IEEE, Piscataway, NJ, 81–89. DOI : <https://doi.org/10.1109/iwsm-mensura.2013.22>
- [105] Jennifer Petrillo, Stefan J. Cano, Lori D. McLeod, and Cheryl D. Coon. 2015. Using classical test theory, item response theory, and Rasch measurement theory to evaluate patient-reported outcome measures: A comparison of worked examples. *Val. Health* 18, 1 (2015), 25–34. DOI : <https://doi.org/10.1016/j.jval.2014.10.005>
- [106] David J. Pittenger. 1993. Measuring the MBTL... and coming up short. *J. Career Plan. Employ.* 54, 1 (1993), 48–52.
- [107] Friedrich Pukelsheim. 1994. The three sigma rule. *Amer. Statist.* 48, 2 (1994), 88–91.
- [108] Paul Ralph, Sebastian Baltes, Gianisa Adisaputri, Richard Torkar, Vladimir Kovalenko, Marcos Kalinowski, Nicole Novielli, Shin Yoo, Xavier Devroey, Xin Tan, Minghui Zhou, Burak Turhan, Rashina Hoda, Hideaki Hata, Gregorio Robles, Amin Milani Fard, and Rana Alkadhi. 2020. Pandemic programming. *Empir. Softw. Eng.* 25, 6 (Sept. 2020), 4927–4961. DOI : <https://doi.org/10.1007/s10664-020-09875-y>
- [109] Paul Ralph and Ewan Tempero. 2018. Construct validity in software engineering research and software metrics. In *22nd International Conference on Evaluation and Assessment in Software Engineering (EASE'18)*. Association for Computing Machinery, New York, NY, 13–23. DOI : <https://doi.org/10.1145/3210459.3210461>
- [110] J. Rattray and M. C. Jones. 2007. Essential elements of questionnaire design and development. *J. Clin. Nurs.* 16, 2 (2007), 234–243. DOI : <https://doi.org/10.1111/j.1365-2702.2006.01573.x>
- [111] Glenn Regehr, Marion Bogo, Cheryl Regehr, and Roxanne Power. 2007. Can we build a better mousetrap? Improving the measures of practice performance in the field practicum. *J. Soc. Work Educ.ation* 43, 2 (2007), 327–344. DOI : <https://doi.org/10.5175/jswe.2007.200600607>
- [112] William Revelle. 2009. *An Introduction to Psychometric Theory with Applications in R*. Retrieved from personality-project.org.
- [113] William Revelle. 2018. *An Introduction to the Psych Package: Part II Scale Construction and Psychometrics*. Technical Report. The Comprehensive R Archive Network.
- [114] William Revelle. 2018. *Using the Psych Package to Generate and Test Structural Models*. Technical Report. The Comprehensive R Archive Network.
- [115] William Revelle. 2019. *psych: Procedures for Psychological, Psychometric, and Personality Research*. Northwestern University, Evanston, IL. Retrieved from <https://cran.r-project.org/package=psych>. R package version 1.9.12.
- [116] William Revelle. 2020. *How To: Use the Psych Package for Factor Analysis and Data Reduction*. Technical Report. The Comprehensive R Archive Network.
- [117] William Revelle and Thomas Rocklin. 1979. Very simple structure: An alternative procedure for estimating the optimal number of interpretable factors. *Multivar. Behav. Res.* 14, 4 (Oct. 1979), 403–414. DOI : https://doi.org/10.1207/s15327906mbr1404_2
- [118] Judith Runnels. 2013. Measuring differential item and test functioning across academic disciplines. *Lang. Test. Asia* 3, 1 (2013), 9.
- [119] Ioana Rus, Mikael Lindvall, and S. Sinha. 2002. Knowledge management in software engineering. *IEEE Softw.* 19, 3 (2002), 26–38.
- [120] Daniel W. Russell. 2016. In search of underlying dimensions: The use (and abuse) of factor analysis in personality and social psychology bulletin. *Personal. Soc. Psychol. Bull.* 28, 12 (2016), 1629–1646. DOI : <https://doi.org/10.1177/014616702237645>

- [121] John Rust. 2009. *Modern Psychometrics : the Science of Psychological Assessment*. Routledge, Hove, East Sussex New York.
- [122] Daniel J. Schad, Michael Betancourt, and Shravan Vasishth. 2021. Toward a principled Bayesian workflow in cognitive science. *Psychol. Meth.* 26, 1 (Feb. 2021), 103–126. DOI : <https://doi.org/10.1037/met0000275>
- [123] Frank L. Schmidt. 1992. What do data really mean? Research findings, meta-analysis, and cumulative knowledge in psychology. *Amer. Psychol.* 47, 10 (1992), 1173–1181. DOI : <https://doi.org/10.1037/0003-066x.47.10.1173>
- [124] J. R. Schoenherr and S. J. Hamstra. 2016. Psychometrics and its discontents: An historical perspective on the discourse of the measurement tradition. *Adv. Health Sci. Educ. Theor. Pract.* 21, 3 (2016), 719–729. DOI : <https://doi.org/10.1007/s10459-015-9623-z>
- [125] Norbert Schwarz and Daphna Oyserman. 2016. Asking questions about behavior: Cognition, communication, and questionnaire construction. *Amer. J. Eval.* 22, 2 (2016), 127–160. DOI : <https://doi.org/10.1177/109821400102200202>
- [126] C. M. Shea, S. R. Jacobs, D. A. Esserman, K. Bruce, and B. J. Weiner. 2014. Organizational readiness for implementing change: A psychometric assessment of a new measure. *Implement. Sci.* 9 (2014), 7.
- [127] Janet Siegmund, Norbert Siegmund, and Sven Apel. 2015. Views on internal and external validity in empirical software engineering. In *IEEE/ACM 37th IEEE International Conference on Software Engineering (ICSE)*. IEEE, Piscataway, NJ, 9–19. DOI : <https://doi.org/10.1109/icse.2015.24>
- [128] Lee Sigelman. 1981. Question-order effects on presidential popularity. *Pub. Opin. Quart.* 45, 2 (1981), 199. DOI : <https://doi.org/10.1086/268650>
- [129] Kamlesh Singh, Mohita Junnarkar, Jasleen Kaur, Kamlesh Singh, Mohita Junnarkar, and Jasleen Kaur. 2016. Norms for test construction. In *Measures of Positive Psychology*. Springer India, New Delhi, 17–34.
- [130] Juani Swart and Nicholas Kinnie. 2003. Sharing knowledge in knowledge-intensive firms. *Hum. Resour. Manag. J.* 13, 2 (2003), 60–75.
- [131] Barbara G. Tabachnick, Linda S. Fidell, and Jodie B. Ullman. 2007. *Using Multivariate Statistics*. Vol. 5. Pearson Boston, MA.
- [132] Lehana Thabane, Jinhui Ma, Rong Chu, Ji Cheng, Afisi Ismaila, Lorena P. Rios, Reid Robson, Marroon Thabane, Lora Giangregorio, and Charles H. Goldsmith. 2010. A tutorial on pilot studies: The what, why, and how. *BMC Med. Res. Methodol.* 10, 1 (2010). DOI : <https://doi.org/10.1186/1471-2288-10-1>
- [133] Louis Leon Thurstone. 1937. Psychology as a quantitative rational science. *Science* 85, 2201 (1937), 227–232. DOI : <https://doi.org/10.2307/1662685>
- [134] Howard E. Tinsley and Diane J. Tinsley. 1987. Uses of factor analysis in counseling psychology research. *J. Counsel. Psychol.* 34, 4 (1987), 414–424.
- [135] Ross E. Traub. 2005. Classical test theory in historical perspective. *Educ. Measur.: Iss. Pract.* 16, 4 (Oct. 2005), 8–14. DOI : <https://doi.org/10.1111/j.1745-3992.1997.tb00603.x>
- [136] John F. Tripp, Cindy Riemenschneider, and Jason B. Thatcher. 2016. Job satisfaction in agile development teams: Agile development as work redesign. *J. Assoc. Inform. Syst.* 17, 4 (2016), 267.
- [137] J. Uher. 2018. Quantitative data from rating scales: An epistemological and methodological enquiry. *Front. Psychol.* 9 (2018), 2599. DOI : <https://doi.org/10.3389/fpsyg.2018.02599>
- [138] Jana Uher. 2021. Quantitative psychology under scrutiny: Measurement requires not result-dependent but traceable data generation. *Personal. Individ. Differ.* 170 (2021), 110205. DOI : <https://doi.org/10.1016/j.paid.2020.110205>
- [139] Eric-Jan Wagenmakers, Maarten Marsman, Tahira Jamil, Alexander Ly, Josine Verhagen, Jonathon Love, Ravi Selker, Quentin F. Gronau, Martin Šmíra, Sacha Epskamp, et al. 2018. Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychon. Bull. Rev.* 25, 1 (2018), 35–57.
- [140] Stefan Wagner, Daniel Mendez, Michael Felderer, Daniel Graziotin, and Marcos Kalinowski. 2020. Challenges in survey research. In *Contemporary Empirical Methods in Software Engineering*, Michael Felderer and Guilherme Horta Travassos (Eds.). Springer International Publishing, Cham, 93–125. DOI : https://doi.org/10.1007/978-3-030-32489-6_4. Available <https://arxiv.org/abs/1908.05899>
- [141] Yi Wang and Min Zhang. 2020. Reducing implicit gender biases in software development: Does intergroup contact theory work. In *28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ACM, New York, NY. DOI : <https://doi.org/10.1145/3368089.3409762>
- [142] Ronald L. Wasserstein, Allen L. Schirm, and Nicole A. Lazar. 2019. Moving to a world beyond “ $p < 0.05$.” *The American Statistician* 73, sup 1 (2019), 1–19. DOI : <https://doi.org/10.1080/00031305.2019.1583913>
- [143] Gerald M. Weinberg. 1971. *The Psychology of Computer Programming*. Dorset House Publishing, New York.
- [144] Howard M. Weiss. 2002. Deconstructing job satisfaction: Separating evaluations, beliefs and affective experiences. *Hum. Resour. Manag. Rev.* 12, 2 (2002), 173–194.
- [145] Keith F. Widaman. 1993. Common factor analysis versus principal component analysis: Differential bias in representing model parameters. *Multivar. Behav. Res.* 28, 3 (1993), 263–311. DOI : https://doi.org/10.1207/s15327906mbr2803_1

- [146] Gordon B. Willis. 2004. Cognitive interviewing revisited: A useful technique, in theory. In *Methods for Testing and Evaluating Survey Questionnaires: Wiley Series in Survey Methodology*. John Wiley & Sons, Inc., Hoboken, NJ, 23–43. DOI : <https://doi.org/10.1002/0471654728.ch2>
- [147] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in Software Engineering*. Springer Berlin.
- [148] D. Wood, M. H. Gardner, and P. D. Harms. 2015. How functionalist and process approaches to behavior can explain trait covariation. *Psychol. Rev.* 122, 1 (2015), 84–111. DOI : <https://doi.org/10.1037/a0038423>
- [149] C. A. Wynd, B. Schmidt, and M. A. Schaefer. 2003. Two quantitative approaches for estimating content validity. *West. J. Nurs. Res.* 25, 5 (2003), 508–518. DOI : <https://doi.org/10.1177/0193945903252998>
- [150] Marvin Wyrich, Daniel Graziotin, and Stefan Wagner. 2019. A theory on individual characteristics of successful coding challenge solvers. *PeerJ Comput. Sci.* 5 (2019), e173.
- [151] Marvin Wyrich, Andreas Preikschat, Daniel Graziotin, and Stefan Wagner. 2021. The mind is a powerful place: How showing code comprehensibility metrics influences code understanding. In *IEEE/ACM 43rd International Conference on Software Engineering (ICSE'21)*. IEEE, Piscataway, NJ, 512–523. DOI : <https://doi.org/10.1109/icse43902.2021.00055>
- [152] An Gie Yong and Sean Pearce. 2013. A beginners guide to factor analysis: Focusing on exploratory factor analysis. *Tutor. Quantit. Meth. Psychol.* 9, 2 (2013), 79–94.
- [153] Bruno D. Zumbo. 2007. Three generations of DIF analyses: Considering where it has been, where it is now, and where it is going. *Lang. Assess. Quart.* 4, 2 (2007), 223–233.

Received July 2020; revised April 2021; accepted June 2021