

The Software Infrastructure Attitude Scale (SIAS): A Questionnaire Instrument for Measuring Professionals' Attitudes Toward Technical and Sociotechnical Infrastructure

Miikka Kuuttila
Dalhousie University
Halifax, Canada
University of Helsinki,
Finland miikka.kuuttila@
helsinki.fi

Paul Ralph
Dalhousie University
Halifax, Canada
paulralph@dal.ca

Huilian Sophie Qiu
Northwestern University
Evanston, USA
huilian_qiu@
alumni.brown.edu

Ronnie de Souza Santos
University of Calgary
Calgary, Canada
ronnie.desouzasantos
@ucalgary.ca

Morakot
Choetkiertikul
Mahidol University
Nakhon Pathom, Thailand
morakot.cho@mahidol.ac.th

Amin Milani Fard
New York Institute of
Technology
Vancouver, Canada
amilanif@nyit.edu

Rana Alkadhi
King Saud University
Riyadh, Saudi Arabia
ralkadi@ksu.edu.sa

Xavier Devroey
NADI, University of Namur
Namur, Belgium
xavier.devroey@
unamur.be

Gregorio Robles
U. Rey Juan Carlos
Madrid, Spain
grex@gsyc.urjc.es

Hideaki Hata
Shinshu University
Nagano, Japan
hata@shinshu-u.ac.jp

Sebastian Baltes
Heidelberg University
Heidelberg, Germany
sebastian.baltes@
uni-heidelberg.de

Vladimir Kovalenko
JetBrains, Amsterdam
The Netherlands
vladimir.kovalenko@
jetbrains.com

Shalini Chakraborty
University of Bayreuth
Bayreuth, Germany
shalini.chakraborty
@uni-bayreuth.de

Eray Tuzun
Bilkent University
Ankara, Turkey
eraytuzun@
cs.bilkent.edu.tr

Hera Arif
Dalhousie University
Halifax, Canada
hera.arif@dal.ca

Gianisa Adisaputri
Dalhousie University
Halifax, Canada
gianisa@dal.ca

Kelly Garcés
Universidad de los Andes
Bogotá, Colombia
kj.garces971@
uniandes.edu.co

Anielle S. L. Andrade
PUCRS
Porto Alegre, Brazil
anielle.lisboa@gmail.com

Eyram Amedzor
Dalhousie University
Halifax, Canada
eyramm@dal.ca

Bimpe Ayoola
Dalhousie University
Halifax, Canada
Bimpe.Ayoola@dal.ca

Keisha Gaspard-Chickoree
Dalhousie University
Halifax, Canada
Keisha.Gaspard
-Chickoree@dal.ca

Arazoo Hoseyni
Dalhousie University
Halifax, Canada
a.hoseyni@dal.ca

Abstract

Context: Recent software engineering (SE) research has highlighted the need for sociotechnical research, implying a demand for

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICSE '26, Rio de Janeiro, Brazil

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM

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

customized psychometric scales. **Objective:** We define the concepts of technical and sociotechnical infrastructure in software engineering, and develop and validate a psychometric scale that measures attitudes toward them. **Method:** Grounded in theories of infrastructure, attitudes, and prior work on psychometric measurement, we defined the target constructs and generated scale items. The items were reviewed and refined by domain experts. The scale was administered to 225 software professionals and evaluated using a split sample. We conducted an exploratory factor analysis (EFA) on one half of the sample to uncover the underlying factor structure and performed a confirmatory factor analysis (CFA) on the other half to validate the structure. Further analyses with the whole sample

assessed face, criterion-related, and discriminant validity. **Results:** EFA supported a two-factor structure (technical and sociotechnical infrastructure), accounting for 65% of the total variance with strong loadings. CFA confirmed excellent model fit. Face and content validity were supported by the item content reflecting cognitive, affective, and behavioral components. Both subscales were correlated with job satisfaction, perceived autonomy, and feedback from the job itself, supporting convergent validity. Regression analysis supported criterion-related validity, while the Heterotrait–Monotrait ratio of correlations (HTMT), the Fornell–Larcker criterion, and model comparison all supported discriminant validity. **Discussion:** The resulting scale is a valid instrument for measuring attitudes toward technical and sociotechnical infrastructure in software engineering research. Our work contributes to ongoing efforts to integrate psychological measurement rigor into empirical and behavioral software engineering research. The scale can also be used by companies and organizations to assess their employees' attitudes toward the infrastructure they provide.

CCS Concepts

• **Social and professional topics** → **Project and people management**; • **Human-centered computing** → *Empirical studies in collaborative and social computing*; • **General and reference** → *Metrics*; *Empirical studies*.

Keywords

Scale, Psychometrics, Attitudes, Infrastructure, Exploratory Factor Analysis, Confirmatory Factor Analysis, Validity

ACM Reference Format:

Miikka Kuuttila, Paul Ralph, Huilian Sophie Qiu, Ronnie de Souza Santos, Morakot Choetkiertikul, Amin Milani Fard, Rana Alkadhi, Xavier Devroey, Gregorio Robles, Hideaki Hata, Sebastian Baltes, Vladimir Kovalenko, Shalini Chakraborty, Eray Tuzun, Hera Arif, Gianisa Adisaputri, Kelly Garcés, Anielle S. L. Andrade, Eyram Amedzor, Bimpe Ayoola, Keisha Gaspard-Chickoree, and Arazoo Hoseyni. 2025. The Software Infrastructure Attitude Scale (SIAS): A Questionnaire Instrument for Measuring Professionals' Attitudes Toward Technical and Sociotechnical Infrastructure. In *2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE '26), April 12–18, 2026, Rio de Janeiro, Brazil*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

Software engineering (SE) researchers are increasingly aware not only of the importance of psychological phenomena (e.g., beliefs, attitudes, intentions) in determining project outcomes, which is central to *behavioral software engineering* [46], but also of the difficulty of measuring these phenomena [57]. While the importance of scale creation has long been understood in psychology [36], scale creation has historically been very limited in SE. More recently, however, researchers working in software engineering have awoken to this fact. Graziotin et al. [23] call for greater attention to creating and validating psychometric instruments for SE research, including more rigorous and theoretically grounded tools for measuring psychological and behavioral constructs.

In modern software development, professionals rely heavily on the infrastructure supporting their work. This includes software tools, but also processes and norms that shape a work environment.

Infrastructure for software development captures technical infrastructure (e.g., CI/CD pipelines, IDEs) as well as sociotechnical infrastructure (e.g., communication tools, documentation standards). However, we currently lack reliable ways to measure the perceptions of these infrastructures, especially in survey-based research. One solution is to assess software professionals' attitudes toward their infrastructure, as attitudes are thought to be good predictors of behavior [3, 21]. Thus, attitudes can act as proxies for infrastructure quality in studies looking at professional beliefs (e.g., job satisfaction) and behavior (e.g., hiring and quitting). We therefore created and validated a psychometric scale to measure developers' attitudes toward technical and sociotechnical infrastructure, guided by the following research question:

Research Question: *Can software professionals attitudes toward technical and sociotechnical infrastructure be validly measured using a psychometric scale?*

Such an instrument is relevant for both software engineering research and practice. Researchers can use the instrument to consider infrastructure as a factor in their studies with developers. Practitioners can use the instrument in longitudinal developer surveys that many companies and organizations regularly conduct [11].

2 Background

This section reviews: (1) theories and definitions related to infrastructure; (2) theories and definitions of attitudes; (3) seminal theories from social psychology that explain phenomena using attitudes and their application to software engineering research; (4) measurement and validity-related work in software engineering.

2.1 Software development infrastructure

Modern software development work is characterized by multitudinous tools, technologies, systems, practices, methods, policies, standards, rituals, and ways of working. In their book *Software engineering at Google*, Winters, Manshreck, and Wright [74] devote whole chapters to examples including version control, build systems, and continuous integration. Such collections of systems have been called “work oriented infrastructures”, due to sharing the same general characteristics as traditional infrastructure [29].

In information systems and human-computer interaction research, studies focusing on infrastructure have long been common. Particularly, the seminal works by Star [62, 63] introduce the concept of sociotechnical infrastructure and demonstrate the importance of studying it. While originally introduced to HCI and IS research, these ideas are increasingly applicable to modern software development environments, which depend on layered, intertwined technical and social systems. The popularity of the concept is demonstrated by the reference counts of these articles, as well as the recent systematic review by Lyu et al. [47], which goes over 190 primary studies noting that the main themes investigated in HCI research are growing infrastructure, appropriating infrastructure, and coping with infrastructure.

Why, then, is studying infrastructure important? The economic theory of infrastructure-led development [2] posits that public infrastructure drives long term economic growth due to network effects and complementarities, but that the benefits from the investment are non-linear. Such investment in software infrastructure is

also done with public money¹. Agenor [2] uses the example of roads and electricity being needed on one hand to produce commodities in rural areas, but also to transport them to urban areas. Examples for the current time could include how investments into mobile and financial infrastructure enable increased usefulness of them in allowing for network effects from externalities such as application stores, mobile payments, or e-commerce platforms. These complementarities amplify the value of the original infrastructure by enabling new services and innovations that rely on it, increasing and reinforcing adoption and participation. In our view, analogous network effects can be seen in infrastructure used in software development. For example, the investment in version control systems, continuous integration pipelines, or shared libraries allow for the use of them across projects, teams and companies. This in turn allows more developers to use them, rewarding standardizing and documentation efforts and reducing barriers to entry. Complementarities arise with additional tools such as deployment automation, collaborative platforms or the recent trend of integrating LLMs into software development. All of these additions have enhanced the utility of the core infrastructure.

However, it has been argued that infrastructure should be divided into groups when investigated, e.g. an archaeological typology of infrastructure divides infrastructure into static (e.g. storehouse), circulatory (e.g. highway), bounding (e.g. palisade) and signaling (e.g. lighthouses) infrastructure [72]. One way of doing this is by dividing infrastructure according to the function it performs [8]. However, as we are interested in infrastructure that is used for software development, we argue that these kinds of typologies do not work as same systems can have multiple purposes simultaneously. As an example, a version control system simultaneously allows for asynchronous work on a complex system with internal dependencies, documenting changes, branching and merging features, conflict resolution, and release planning, to name a few. This division would be thus blurry and ambiguous.

To address these challenges, we draw from systems modeling literature, which offers a productive way to conceptualize infrastructure based on its dependence on human actors and institutions. Ottens et al. [55] define technical systems as “systems that perform their function without either actors or social institutions” or “systems in which some actors perform subfunctions but social institutions play no role” and consequently sociotechnical systems as “engineering systems that need actors and some social/institutional infrastructure to be in place in order to perform their function”. Thus, in an analogous way, we divide technical and sociotechnical infrastructure along these lines. To us, software development systems that professionals use alone are part of technical infrastructure, whereas systems that require multiple actors to perform their function are sociotechnical infrastructure.

2.2 What are attitudes?

Attitudes are “overall evaluations of objects”, which consist of “affective information (e.g., feelings towards an object), cognitive information (e.g., beliefs associated with an object) and behavioral information (e.g., past experiences with an object)” [28]. This tripartite model of attitudes has held since the '60s [54]. However,

¹<https://www.sovereign.tech/programs/fund>

attitudes are more stable over time than core affect [28], understood as “neurophysiological state that is consciously accessible as a simple, nonreflective feeling that is an integral blend of hedonic (pleasure–displeasure) and arousal (activation–deactivation) values” [61]. This relative stability of attitudes makes them attractive and potentially useful for investigating enduring psychological states toward objects or systems, especially given that much existing research in software engineering has already focused more on transient affective states than on more stable attitudes [25, 52].

2.3 Why measure attitudes?

Attitudes are part of several seminal theories in psychology, which are used to explain a multitude of outcomes. These outcomes include, but are not limited to, volitional behavior (actions under conscious control) [21], behavior with limited control [3], behavior in regard to perceived chances of success and values [4], acceptance of technology [12, 67], cognitive consistency [32], attitude change [40], identity and self [7, 39], emotional regulation [20, 30], and perceived values and motivation [6, 64].

Three seminal theories referenced above are used to explain human behavior, including the Theory of Reasoned Action (TRA) [21], the Theory of Planned Behavior (TPB) [3], and Expectancy-Value Theory [4, 17]. Both TRA and TPB assume that behavior follows from intentions, and intentions from attitudes and social norms. Expectancy-Value Theory [4, 17] posits that behavior is influenced by beliefs about whether engaging with a task or system (e.g., using a tool or following a process) will lead to desired outcomes (expectancy), and by the value placed on those outcomes. From the point of view of these theories, attitudes toward technical and sociotechnical infrastructure are thus shaped by developers’ beliefs about their usefulness and the importance they attach to outcomes like productivity and collaboration.

Theories such as Cognitive Dissonance Theory [20] and the Elaboration Likelihood Model [56] explain attitude change through inconsistencies or persuasive communication. For example, developers may dislike documentation or testing practices but adjust their views over time. Our scale enables tracking such changes.

Theories like Symbolic Interactionism [7] and Social Cognitive Theory (SCT) [6] highlight how attitudes are shaped through social roles, observation, and perceived efficacy. For example, a developer placed in a mentoring role may adopt more positive views toward documentation tools. Similarly, Value-Belief-Norm Theory (VBN) [64] explains how attitudes can underpin moral obligation, motivating developers to act in alignment with shared values.

Thus, measuring developers’ attitudes toward technical and sociotechnical infrastructure would be valuable in survey research for a multitude of reasons. The scale could offer insight into how well work environments support, for example, productivity, collaboration and job satisfaction. This is done by introducing these constructs into a set of related constructs that are often called a nomological networks, i.e. “interlocking system of laws which constitute a theory” [10].

2.4 Scales in software engineering research

A recent book chapter has called attention to measurement issues in software engineering research [57]. This is further emphasized

with calls for software development specific scale development [23]. However, rather than using scales that are developed specifically for software professionals in survey research, it is common to adapt items from scales that have been validated with more general samples (e.g., [16, 24]). While some scale development papers exist [49], they have stopped after expert evaluation and do not generally administer the instrument to practitioners, nor do they rigorously and quantitatively assess the validity of the developed measures. To our knowledge, the present work's highly novel contribution to the field of software engineering is in that it documents the development and evaluation of a scale with psychometric properties.

3 Method

We followed Trochim's [66] guidelines for developing Likert scales, with reference to relevant sections of Graziotin et al.'s [23] guidelines. Trochim [66] recommends a four-step approach: (1) Define the construct being measured, (2) create a pool of possible items to measure the construct, (3) panel review and scoring of items by experts and (4) select a subset of items based on these evaluations. In addition, we (5) administered the scale for a larger audience and assessed reliability and validity. Figure 1 visualizes this process and shows the number of possible questionnaire items after each step.

We provide a comprehensive replication package including all of the items and their scoring at each stage (see Data Availability). In the following subsections, we elaborate on each step.

3.1 Defining the constructs

Both Trochim's [66] and Graziotin's [23] guidelines start by defining what we want to measure. To measure software engineers' attitudes toward systems that they use at work, we developed a construct that integrates two complementary dimensions: technical infrastructure and sociotechnical infrastructure.

These dimensions are built on the perspectives offered in Section 2.1. Technical infrastructure includes the tangible and primarily automated systems that support software professionals day-to-day development work, such as computer hardware and software, programming languages, frameworks, IDEs, version control systems, sketching and testing tools. These are all tools and systems that are often, but not always, used by a single actor to perform their work duties in software development work. In contrast, sociotechnical infrastructure encompasses systems that need multiple actors and institutional infrastructure to perform their function and allow individual software professionals to coordinate, collaborate and manage software development work. In sociotechnical infrastructure, we include tools, processes, roles, norms, management, and collaborative practices that are used to communicate and organize software development work. The tools included, but not necessarily limited to, are communication and information sharing tools, documentation, coding standards, pair programming and code ownership policies, peer code review practices, sprints, user story planning, and team meetings. This distinction divides software development work into two categories: (1) technical work, where the individual primarily interacts with technical systems to produce technical artifacts, and (2) sociotechnical work, where the individual engages with collaborative systems and practices to coordinate and manage work with others. Thus, we define:




Figure 1: PRISMA type diagram showing the steps in creating the scale.

Definition 1 (Technical Infrastructure). The tools and technical systems individuals use to perform their work, such as computer hardware, software, programming languages, frameworks, version control systems, sketching tools, testing tools, networks, security mechanisms, and other technical systems.

Definition 2 (Sociotechnical Infrastructure). The tools, practices, and processes that are used to communicate and organize work, such as management processes, communication tools, information sharing tools, documentation, coding standards, pair programming, code ownership policies, peer code review practices, sprints, user stories, and team meetings.

Viewing software development infrastructure through these two dimensions allows researchers to investigate how technical and sociotechnical elements interrelate with each other and allows for

analyzing how these dimensions affect factors such as productivity, well-being and team dynamics. These two dimensions also allow for a holistic view beyond individual tools or practices. In our view, this holistic approach allows for measurement that has less construct overlap and contextually confounding factors that can be hard to control for [71]. For instance, a single treasured tool could affect measurements of tools that are often used alongside it in ways reminiscent of the halo effect [51]. Thus, by dividing the concepts this way with the still broad and holistic concepts of technical and sociotechnical infrastructure, software engineering research can capture software professionals' attitudes and experiences.

3.2 Generating items

Trochim calls this step "Generating the Items" [66], while Graziotin et al. [23] call it "Pool Items". We gave eight domain experts preliminary definitions of the two dimensions above and asked them to generate items. By domain expert, we mean individuals with theoretical knowledge and practical experience in the subject, as is common in psychological literature [31]. The eight individuals included a professor, postdoctoral researcher, doctoral students, and a master's student, all of whom had years of experience in SE research, teaching, and practice. The group of domain experts consisted of individuals with diverse backgrounds, including participants originally from North America, South America, the Caribbean, Africa, the Middle East, and the Nordics. Experts were asked to generate 10 items per dimension, though they could stop at any time. They collectively generated 57 items for technical infrastructure and 71 items for sociotechnical infrastructure. After these preliminary items were collected, the first and second authors revised them for clarity. We have marked the original items and the modifications in our replication package.

To supplement these pools of items, we generated 50 additional items on the 7th of July 2024 for each type of infrastructure using ChatGPT 4o with the default parameters accessed via a browser, and with various prompting techniques. We used true random numbers from random.org to choose examples for the prompts that require examples (One-Shot, Few-Shot, and Few-Shot chain of thought) from the list of items generated by domain experts. This process, together with the exact prompts, is documented in our replication package. We did not use LLMs in any other part of this research.

3.3 Panel review and expert scoring

Trochim calls this step "Rating the items" and recommends it *before* pilot testing [66], while Graziotin et al. call it "Item analysis" and recommend it *after* pilot testing [23]. We followed Trochim to limit the workload for the experts providing the favorability ratings (examples below).

The first two authors pruned and de-duplicated the items in a face-to-face collaborative session. First, both authors read the list of items, and then removed items that were too specific. For example, "I feel confident in the security measures of our technical infrastructure" is too specific because it's about security in particular. In contrast, some items were too specific in that they applied only to a subset of software professionals, such as developers working in distributed environments. Next, we grouped similar items and selected the most concise and general item from each group, dropping the

others. For example, we chose the item socio2 (see Table 3) over items such as "I find the current infrastructure supports seamless information sharing" and "I find it challenging to share information with the existing tools and processes". This resulted in a pool of 35 items for technical infrastructure and 47 for sociotechnical infrastructure. We provide a reason for every item dropped at this stage in our replication package.

Next, 23 domain experts scored each item's "favorability" toward the target concept. Favorability here does not refer to what the expert believes is the case for them, but rather they are judging "how favorable each item is with respect to the construct of interest" [66]. The domain experts included 14 university professors, 2 postdocs and 7 graduate students, all working in the software engineering domain. Domain experts received instructions, together with examples, on scoring favorability of each item with regard to the construct it attempts to measure, and the instruction to use a rating scale from +5 to -5. The order of the items being scored was randomized for each rater. We also added three attention check items for both scales to the favorability ratings. These attention check items were either negative statements framed as positive or vice versa, for example: "Our sociotechnical infrastructure makes workflow so thorough and rigorous that our deliverables are often late". After a brief trial run, 15 of the experts scored 94.4% of these items correctly, that is, positive statements as favorable and vice versa. The high rate of correct scoring indicates that the experts paid attention to the items when scoring them.

3.4 Selecting items for the preliminary scale

From the favorability ratings, we calculated total-item correlations as advocated by Trochim [66] and performed a Welch's t-test between the lowest and highest quartiles for assessing discriminative power between items. We observed that two raters scored the technical infrastructure items substantially lower than all others, with total scores of 15 and 29 compared to a range of 77 to 154 for the remaining raters. Due to this large discrepancy, which suggested a misunderstanding of the task or scale, we excluded their ratings from the analysis. As a result, selection scores were calculated using ratings from 21 raters for technical infrastructure and 23 for sociotechnical infrastructure.

To select the items for the scale, we first removed all of the items with an item-total correlation below Trochim's recommended cut-off of 0.6 [66]. Then, for each infrastructure dimension, we selected the 10 items with the highest t-test scores. We then reversed some items so that both dimensions would have 5 regularly scored and 5 reverse-scored items, for example, the item "Our technical infrastructure is well-maintained" was changed to "Our technical infrastructure is poorly maintained". Our aim here was to have balanced scales for both dimensions, to reduce response biases, to control for acquiescence (i.e. "yeah saying"), and to disrupt non-substantive responding (straightlining, that is, selecting the same option across items) [70]. Following recommendations in the literature [9], we avoided negations such as "not" and other modifications making items awkward, lengthy, or difficult to interpret.

Table 1: Descriptive Statistics for Scale Items.

Item ID	Mean	SD	Skewness	Kurtosis	Miss %
tech1	3.11	1.12	-0.05	-0.97	2.67
tech2	2.64	1.17	0.52	-0.75	0.89
tech3	2.70	1.20	0.42	-0.86	2.67
tech4	3.66	0.94	-0.58	-0.06	1.78
tech5	3.56	1.00	-0.71	0.06	0.89
tech6	2.47	1.13	0.50	-0.60	1.33
tech7	3.80	0.86	-0.83	0.95	0.44
tech8	3.77	1.00	-0.71	0.06	0.89
tech9	3.91	0.84	-1.01	1.37	0.89
tech10	2.24	1.10	0.80	-0.12	0.00
socio1	2.96	1.09	0.05	-0.97	2.67
socio2	3.70	0.99	-0.71	0.08	0.00
socio3	2.82	1.15	0.31	-0.85	0.44
socio4	2.66	1.10	0.63	-0.51	1.33
socio5	3.56	1.01	-0.56	-0.13	2.22
socio6	3.58	0.95	-0.72	0.32	0.44
socio7	3.60	1.03	-0.67	-0.06	0.44
socio8	2.71	1.05	0.46	-0.64	0.44
socio9	3.51	0.97	-0.68	0.06	4.00
socio10	2.37	0.96	0.80	0.15	0.00

3.5 Assessing the validity of the scale

While Trochim's [66] original guidelines suggest that the scale can now be used in practice, contemporary psychometric standards emphasize the importance of empirically evaluating the scale's factor structure and validity through data collection and analysis [14, 75].

To administer the newly developed scales, we incorporated them into a broader questionnaire used in a larger survey study [43], in which software professionals' attitudes toward their work infrastructure serves as an exogenous variable. The questionnaire also included related constructs, such as job satisfaction (measured using the *5-item Short Index of Job Satisfaction* [37]) and job quality (captured through the autonomy and feedback from the job itself dimensions of the revised *Job Diagnostic Survey* [35]). As shown later in Section 4, including these related constructs allows us to rigorously evaluate the proposed scales in terms of both criterion-related and discriminant validity.

The target population for the questionnaire survey where the scale was administered is software professionals, e.g., programmers, testers, technical artists, and their managers, who are primarily employed by a single organization from which they receive the majority of their income, and who work within the software industry. Freelancers and self-employed individuals were excluded from the study. Only respondents who reported being between the ages of 18 and 65 were included, as these were considered the conventional working-age population.

The questionnaire survey was translated from English into Portuguese, Chinese, Thai, Finnish, Swedish, Arabic, Farsi, Urdu, German, Spanish, Japanese, Russian and Turkish, and implemented using Opinio.² Each translation was localized by a researcher with prior experience in conducting human factors research in software engineering. After receiving permission from the Dalhousie

²<https://www.objectplanet.com/opinio/>

research ethics board, the study was advertised in these locales using social media and company contacts (e.g. LinkedIn, internal company communication channels). The responses were received between December 12, 2024 and June 7, 2025. Incentives such as raffles and gift cards were used depending on the locale, however, caution was used so that gift cards were used only when advertising in closed spaces such as company channels. Our recruiting strategy thus corresponds to convenience and breadth sampling [5].

To improve data quality, all demographic questions were optional. We also included an "I don't know" option for each scale item [see 15, 53] and an attention-check item for the Job Diagnostic Survey [41]. Items scored as "I don't know" were coded as missing, while surveys with a failed attention-check item were discarded. The order of the items was randomized to reduce order effects [48]. The item-level descriptives of the responses can be seen in Table 1. Item means ranged from 2.24 to 3.80, standard deviations from 0.84 to 1.20, all of which are within acceptable bounds. Skewness and kurtosis values also indicated no substantial deviations from normality (all skewness values $<|1.1|$ and all kurtosis values $<|1.4|$). Item-level missingness was mostly below 2.7%. We used listwise deletion in EFA and full information maximum likelihood [18] in CFA.

The sample size needed for both EFA and CFA is debated, but the often-cited work by Worthington and Whittaker [75], states "smaller sample sizes may be adequate... with at least 4 : 1 items per factor and factor loadings greater than $|.6|$ ", and recommends a minimum sample size of 100. Our work satisfies these conditions, with an item to response ratio of 5.65 and a sample size of 122 for the EFA. We split our sample randomly using the base R function sample, resulting in 122 and 123 responses used in EFA and CFA, respectively. For later analyses, we used the whole sample.

To explore and improve the underlying structure of the scales and reduce the number of items, we conducted an exploratory factor analysis (EFA) using the maximum likelihood extraction and oblimin rotation with the R package psych [58] (version 2.5.3). Rotation is applied after extraction to improve interpretability by simplifying the factor structure. We chose oblimin rotation (an oblique method) because the underlying constructs we are attempting to measure should be related, and oblimin rotation allows factors to correlate with one another [19]. Improving factor structure involves iterating among (1) investigating the factor structure using EFA; (2) dropping items with low communality; (3) dropping items with substantial cross-loadings. The analysis continues until all remaining items load highly on the factor they are supposed to load on, and no others. This process is a powerful tool for assessing and enhancing validity because it simultaneously assesses convergent and discriminant validity. However, it can lead to overfitting, which we mitigate by randomly splitting our dataset in two, using half the data for this exploratory analysis, and the other half for confirmatory analysis.

The exploratory analysis was performed on a subset of 122 valid responses randomly selected from the entire data set. The initial EFA suggested a three-factor solution based on eigenvalues > 1 and parallel analysis, with a cumulative variance explained of 56%. However, several items showed substantial cross-loadings or weak factor loadings ($< .40$), and some exhibited high complexity or low communalities ($< .30$), indicating poor fit. Based on these criteria, we iteratively removed items to improve the clarity of the factor

structure. After the final EFA was performed, we assessed the suitability of the data for factor analysis by performing Bartlett’s Test of Sphericity and calculating the Kaiser-Meyer-Olkin (KMO) [73] measure using the R package psych [58] (version 2.5.3). This factor structure from the EFA was retained for the rest of the analyses.

To confirm the factor structure, we performed confirmatory factor analysis with the R package lavaan [60] (version 0.6-18). We performed two confirmatory factor analyses: (1) a model containing only technical and sociotechnical infrastructure factors identified in the EFA and (2) a model including additional factors of job satisfaction [37], as well as autonomy and feedback from the revised job characteristics model [35]. The first model validates the factor structure obtained from the EFA, whereas the second model provides us with factor scores for external constructs to the scale being developed, which were used in later analyses for convergent, criterion-related, and discriminant validity.

We used Job Satisfaction, Autonomy and Feedback from the job itself in later analyses for convergent, criterion-related, and discriminant validity. Job satisfaction is an overall evaluation of a job, and thus we presume it would positively correlate with positive attitudes towards technical and sociotechnical infrastructure. In the job characteristic model introduced by Oldham and Hackman [27], autonomy refers to the degree to which the job provides freedom and discretion to the individual in scheduling their work and in determining the procedures to be used in carrying it out, while feedback refers to the degree to which carrying out work activities results in the employee obtaining direct and clear information about the effectiveness of their performance in the job. Thus, we presume that jobs that are perceived as autonomous and providing feedback in themselves are also on average jobs where technical and sociotechnical infrastructure allow for these through, for example, tooling and job procedures. Hence, we assume that these two constructs are positively associated with higher evaluations of technical and sociotechnical infrastructure.

To assess the criterion-related validity of the newly developed scale, we correlated its factor scores with established constructs of job satisfaction, autonomy and feedback from the job itself using the whole sample. Correlations were calculated using the R package psych [58] version 2.5.3 and regression analysis was performed with the function lm from base R version 4.3.3.

Lastly, to assess the discriminant validity of our scale, we calculated two commonly reported measures, the Heterotrait-Monotrait Ratio of Correlations (HTMT) [33] and the Fornell-Larcker criterion [22]. HTMT was calculated using the function htmt from the R package semTools version 0.5.7. Latent variable correlations used for the Fornell-Larcker discriminant validity analysis were obtained from the CFA model-implied correlation matrix (i.e., function inspect(fit, "cor.lv")). These account for measurement error and reflect the structural relationships between constructs. However, due to recent critique of these measures by Rönkkö & Cho [59], we also assessed discriminant validity by fitting a baseline CFA model with both dimensions of our infrastructure scale, and fitting an alternative model where the correlation between the constructs was set to 1.0. We compared the model fits with the chi-square test, as well as with the Bayesian information criterion (BIC) and the Akaike Information Criterion (AIC) [68].

4 Results

In this section, we first report the demographics of the respondents to our questionnaire study, next we explore the factor structure of the administered scale, and then we confirm it with a split sample approach. Next, we discuss the face and content validity of the scale by analyzing the items remaining after factor analysis. In the last three subsections, we provide evidence for convergent, criterion-related, and discriminant validity of the developed scale.

4.1 Respondent demographics

Table 2: Self-reported demographic information of the respondents.

	Demographic Item	N	%
<i>Gender</i>	Man	157	70%
	Woman	59	26%
	NA / Prefer not to say	9	4%
<i>Country of Residence</i>	Brazil	60	27%
	People’s Republic of China	31	14%
	Finland	25	11%
	Thailand	21	9%
	Saudi Arabia	13	6%
<i>Survey Language</i>	Portuguese	61	27%
	Chinese	48	21%
	Thai	27	12%
	Finnish	23	10%
	English	21	9%
	Other	45	20%
<i>Age (Range: 19–61)</i>	Mean: 35.4, Median: 34		
	18–24	13	6%
	25–30	66	29%
	31–40	87	39%
	41–50	39	17%
	55 or more	20	9%
<i>Years of Experience (Range: 0–40)</i>	Mean: 10.4, Median: 8		
	0–4 years	57	25%
	5–9 years	63	28%
	10–19 years	58	26%
	20+ years	33	15%
<i>Highest Education</i>	Doctorate	15	7%
	Master’s	80	36%
	Bachelor’s	93	41%
	Other or NA	37	16%
<i>Company Size</i>	0–9	11	5%
	10–99	46	20%
	100–999	42	19%
	1000–9999	69	31%
	10,000–99,999	23	10%
	100,000 or more	14	6%
	NA	20	9%

All demographic questions but age were optional to answer.

We received complete responses from 239 software professionals. After removing answers with failed attention check items, we were left with a total of 225 answers. We report demographics in Table 2. Note that none of the demographic questions other than age were mandatory to answer, meaning the number of responses from question to question changes. Respondents reported a wide range of software development experience (0–40 years, mean = 10.4), with the majority holding a bachelor's (41%) or master's degree (36%). The participants were employed in organizations of varying sizes, with nearly one-third working in companies with 1,000–9,999 employees. Respondents were from diverse geographic and linguistic backgrounds, with Brazil, China, and Finland being the most common countries of residence, and Portuguese, Chinese, and Thai among the most common languages the survey was taken in.

4.2 Exploratory factor analysis

Following item reduction, at the end of the iterative EFA process, we fitted an EFA model using maximum likelihood extraction with oblimin rotation on a refined set of 11 items (see Table 3). The results supported a two-factor structure consistent with our theoretical distinction between technical (6 items) and sociotechnical infrastructure (5 items). Items showed strong primary loadings (ranging from .66 to .88), minimal cross-loadings, and high communalities ($h^2 \geq .55$). The two factors were moderately correlated ($r = .63$), supporting the use of an oblimin rotation.

The chi-square test of model fit was statistically significant, $\chi^2(34) = 56.75$, $p = .009$, which is not uncommon given the sample size ($N = 122$). However, alternative fit indices indicated excellent model fit (RMSEA = .077, 90% CI [.039, .112]; RMSR = .03; TLI = .953), supporting the adequacy of the two-factor solution. The two factors accounted for 65% of the total variance, and item complexity was low across the board (mean = 1.0). This 11-item structure was retained for confirmatory factor analysis using a split sample.

We performed post hoc tests for sampling adequacy for the data used in the EFA, using the Kaiser-Meyer-Olkin (KMO) measure [73]. The overall measure of sampling adequacy (MSA) provided by the test was 0.90, which corresponds to marvelous adequacy for factor analysis in the original Kaiser's evaluation levels [38]. All individual item MSAs were above 0.85, ranging from 0.86 to 0.93, supporting the inclusion of all retained items. We also performed Bartlett's Test of Sphericity, which was significant ($\chi^2(45) = 825.59$, $p < .001$), indicating that the correlations among items were sufficiently large for factor analysis.

4.3 Confirmatory factor analysis

Construct validity refers to "the degree to which a test or instrument is capable of measuring a concept, trait, or other theoretical entity." [1]. We conducted a confirmatory factor analysis (CFA) to evaluate the factorial validity of a two-factor model with a separate sample. This was done with the factor structure identified by the EFA described in the previous subsection. Model fit indices are shown in Table 5 and the loadings in Table 4. The model included 11 items loading onto two correlated latent constructs: technical infrastructure and sociotechnical infrastructure. All items loaded significantly ($p < .001$) on their respective factors, with standardized

loadings ranging from 0.588 to 0.853 for technical infrastructure and 0.685 to 0.866 on sociotechnical infrastructure. One item in the technical infrastructure construct exhibited a slightly lower loading (0.588), which we decided to retain due to negatively worded items being helpful for lowering response bias.

The model demonstrated excellent fit to the data ($N = 123$): $\chi^2(43) = 48.062$, $p = .275$, CFI = 0.993, TLI = 0.991, RMSEA = 0.032, 90% CI [0.000, 0.074], with $p = .715$ for RMSEA $\leq .05$, SRMR = 0.041. The two latent constructs were moderately correlated (standardized estimate = 0.601, $p < .001$), supporting their conceptual relatedness while retaining discriminant validity.

4.4 Face and content validity

Content validity refers to "the extent to which a specific set of items reflects a content domain" [14]. As discussed in the background section, attitudes are overall evaluations of objects that comprise three types of information: cognitive, affective, and behavioral [28]. These correspond to beliefs, feelings, and past experiences, respectively. Examining the results of the EFA (Table 3), we find that the items selected through our scale development for both the technical and sociotechnical infrastructure reflect all three components.

For the technical infrastructure items, tech4 and tech7 primarily convey cognitive information, expressing beliefs about the infrastructure being well made and functioning effectively. Item tech5 reflects affective content by referring to satisfaction with the infrastructure. Items tech8, tech9, and tech10 incorporate behavioral information, as they refer to personal experiences using the infrastructure to perform tasks efficiently and with appropriate tools.

The sociotechnical items similarly represent the tripartite structure. Items socio2 and socio5 contain cognitive content, describing beliefs about ease of information sharing and user-friendliness. Items socio6 and socio9 reflect behavioral experiences related to enhanced workflows and responsiveness to diverse team needs. Finally, item socio7 includes an affective element, describing the infrastructure as encouraging continuous learning and improvement.

4.5 Convergent validity

"Convergent validity is evidence of similarity between measures of theoretically related constructs" [14]. To examine convergent validity of the new scale, we computed Pearson correlations between the factor scores of two dimensions, Tech and Sociotech, and factor scores of relevant theoretically related constructs: perceived autonomy, feedback from the job itself, and job satisfaction (see Table 6). As expected, Tech and Sociotech dimensions were strongly correlated with each other ($r = .68$, $p < .001$), supporting the idea that they capture related but distinct aspects of attitudes toward work infrastructure.

Both dimensions also showed positive correlations with job-related attitudes. Specifically, technical infrastructure was moderately associated with job satisfaction ($r = .47$, $p < .001$), feedback from the work itself ($r = .41$, $p < .001$), and autonomy ($r = .34$, $p < .001$). Sociotechnical infrastructure was likewise significantly correlated with job satisfaction ($r = .39$, $p < .001$), feedback ($r = .48$, $p < .001$), and autonomy ($r = .27$, $p < .01$). These results indicate that more favorable attitudes toward technical and sociotechnical infrastructure are associated with greater perceived autonomy, better

Table 3: Items after the iterative exploratory factor analysis process.

ID	Item	F1	F2	Comm.	Uniq.	Comp.	MSA
tech4	The technical infrastructure I use at work is well made	0.88	0.00	0.78	0.22	1.0	0.91
tech5	I am satisfied with the technical infrastructure provided at my workplace	0.88	-0.02	0.75	0.25	1.0	0.91
tech7	Our technical infrastructure works well	0.82	0.09	0.77	0.23	1.0	0.93
tech8	The technical infrastructure I use helps me work more efficiently	0.71	0.13	0.64	0.36	1.0	0.87
tech9	Our infrastructure includes tools appropriate for my work tasks	0.73	0.13	0.64	0.36	1.1	0.93
tech10	The technical infrastructure at my workplace is unreliable	0.79	-0.07	0.56	0.44	1.0	0.91
socio2	Sociotechnical infrastructure makes it easy to share information	-0.05	0.83	0.64	0.36	1.0	0.89
socio5	Sociotechnical infrastructure at work is user-friendly	0.15	0.66	0.58	0.42	1.1	0.89
socio6	Our workflow is enhanced by the sociotechnical systems	0.06	0.77	0.65	0.35	1.0	0.89
socio7	Sociotechnical infrastructure encourages learning and improvement	-0.09	0.79	0.55	0.45	1.0	0.86
socio9	Sociotechnical infrastructure supports different team needs	0.04	0.85	0.78	0.22	1.0	0.86

Model fit: RMSEA = 0.077 (90% CI: 0.039–0.112), RMSR = 0.03, TLI = 0.953

Table 4: Standardized Factor Loadings from CFA Model

Factor	Item	Std. Loading
Technical Infrastructure	tech4	0.831
	tech5	0.845
	tech7	0.853
	tech8	0.764
	tech9	0.754
	tech10	0.588
Sociotechnical Infrastructure	sociotech2	0.701
	sociotech5	0.804
	sociotech6	0.866
	sociotech7	0.685
	sociotech9	0.853

Table 5: Fit indices for CFA models

Fit Index	Two-Factor (CFA Split)	Five-Factor (Full Sample)
χ^2 (df), p	48.1 (43), $p = .275$	297.05 (199), $p < .001$
CFI	0.993	0.962
TLI	0.991	0.956
RMSEA [90% CI]	0.032 [0.000, 0.074]	0.047 [0.035, 0.058]
SRMR	0.041	0.052

Note. Two-factor = technical + sociotechnical infrastructure. Five-factor adds autonomy, feedback, and job satisfaction.

perceived feedback from the work itself, and higher job satisfaction, supporting the convergent validity of the scale.

4.6 Criterion-related validity

For a scale to have criterion (or criterion-related) validity, it needs to have empirical associations with some criterion, whether or not that association is understood theoretically [14]. In this study, we assessed concurrent validity, a subtype of criterion-related validity, by testing whether attitudes toward technical and sociotechnical infrastructure were associated with job characteristics and satisfaction measured at the same time. Specifically, we used a series of linear regression models to examine whether these attitudes

Table 6: Correlations between latent factor scores.

	Tech.	Socio.	Auto.	Feed.	Sat.
Tech.	1.00	.68***	.34***	.41***	.47***
Sociotech		1.00	.27***	.48***	.39***
Autonomy			1.00	.83***	.57***
Feedback				1.00	.64***
Satisfaction					1.00

Note. Tech = Attitude towards technical infrastructure, Socio = Attitude towards sociotechnical infrastructure, Auto = Autonomy, Feed = Feedback, Sat = Job Satisfaction. *** $p < .0001$.

predicted autonomy, feedback from the job itself, and job satisfaction. The results can be seen in Table 7. Technical infrastructure significantly predicted autonomy ($\beta = .31, p = .001$), while sociotechnical infrastructure did not. Both types of infrastructure predicted feedback from the job itself, with sociotechnical infrastructure showing a notably stronger effect ($\beta = .49, p < .001$). In predicting job satisfaction, technical infrastructure was a robust predictor both directly ($\beta = .57, p < .001$) and after adding autonomy and feedback to the model ($\beta = .41, p < .001$). Feedback was the strongest predictor of job satisfaction in the full model ($\beta = .60, p < .001$), while the effect of sociotechnical infrastructure dropped to non-significance, suggesting that its influence is mediated through feedback mechanisms.

4.7 Discriminant validity

Discriminant validity has been defined as “Two measures intended to measure distinct constructs have discriminant validity if the absolute value of the correlation between the measures after correcting for measurement error is low enough for the measures to be regarded as measuring distinct constructs” [59]. To assess discriminant validity of our scale, we first computed the Heterotrait-Monotrait (HTMT) ratios of correlations [33] between infrastructure scale dimensions and constructs used in criterion-related validity in the last subsection. The HTMT value between technical and sociotechnical infrastructure was 0.596. HTMT values between technical infrastructure and autonomy, feedback, and job satisfaction varied from 0.242 to 0.482. The same values for sociotechnical infrastructure were between 0.152 and 0.389. Thus, all HTMT values

Table 7: Four linear regression models predicting autonomy, feedback, and job satisfaction with attitudes toward technical and sociotechnical infrastructure. Estimate followed by the p -value in parenthesis in the line below

Predictor	Autonomy	Feedback	Job Satisfaction	
			Direct	Mediated
Tech	0.310	0.192	0.573	0.416
p	(<0.0001)	(0.045)	(<0.0001)	(0.0001)
Sociotech	0.099	0.489	0.208	-0.095
p	(0.352)	(<0.0001)	(0.129)	(0.452)
Autonomy	—	—	—	0.138
p				(0.293)
Feedback	—	—	—	0.593
p				(<0.0001)
R^2	0.110	0.246	0.221	0.460
Residual SE	0.804	0.815	1.047	0.875
F-statistic	14.83***	35.17***	31.68***	46.88***

were well below the threshold of 0.85 advocated by Henseler [33], supporting discriminant validity among the constructs.

We compared the square roots of the average variance extracted (AVE) for technical and sociotechnical infrastructure ($\sqrt{AVE}_{tech} = 0.793$; $\sqrt{AVE}_{sociotech} = 0.787$) with the correlations from the model-implied latent correlation matrix of the CFA model, where the relationships between latent constructs are adjusted for measurement error. This is also called the Fornell–Larcker criterion [22]. The implied correlation from the CFA model between technical infrastructure and sociotechnical infrastructure was 0.631, while the correlations between other constructs were from 0.293 to 0.431. These results support discriminant validity of our scale, with both constructs sharing more variance with their own indicators than with any other construct.

Lastly, we followed the procedure advocated by Rönkkö and Cho’s [59] by comparing a CFA model in which the correlation between technical and sociotechnical infrastructure was freely estimated to a constrained model in which their correlation was fixed to 1.0. The chi-square difference test was significant, $\chi^2(1) = 170.1$, $p < .001$, and both AIC and BIC were lower for the unconstrained model (AIC: 2394.1 vs. 2562.1; BIC: 2478.6 vs. 2643.9). Taken together, all these results strongly indicate that the constructs of technical and sociotechnical infrastructure are empirically distinct and thus support discriminant validity.

5 Discussion

The presented Software Infrastructure Attitude Scale (SIAS) demonstrates strong psychometric validity, including construct, convergent, criterion-related, and discriminant validity. These properties were supported through expert review, item rating, and validation in a diverse sample of software professionals. Researchers have been interested in sociotechnical aspects of SE for over a decade [50], with recent works focusing on these aspects further underlining their importance in SE research [34, 45]. We believe our work provides a valuable tool for this line of research.

Both infrastructure dimensions performed as expected in relation to established constructs of autonomy, feedback from the job

itself, and job satisfaction, with positive bivariate correlations with all. Attitudes toward technical infrastructure were associated with perceived autonomy, feedback from the job itself, and job satisfaction in multivariate regression analyses. These results are in line with self-determination theory (SDT) [13], which emphasizes that employees who feel supported in their autonomy are more satisfied with their jobs. While autonomy did not predict job satisfaction in our model, this is likely due to shared variance with feedback. From the perspective of SDT, it would be logical that those SE professionals who have more positive attitudes toward technical infrastructure, have, on average, better technical infrastructure, which in turn facilitates higher autonomy, leading to higher job satisfaction.

Attitudes toward sociotechnical infrastructure only predicted feedback from the job itself, the definition of which is “the degree to which carrying out the work activities required by the job results in the individual obtaining direct and clear information about the effectiveness of his or her performance” [27]. It can be speculated that pieces of sociotechnical infrastructure such as coding standards and communication tools facilitate this feedback.

The practical contribution of this work is a scale with psychometric properties that researchers and practitioners can use to measure software professionals’ attitudes toward technical and sociotechnical infrastructure. In empirical software engineering research, the scale can be included in survey studies, either as an independent or dependent variable, or as part of a model that combines it with a set of related constructs. This can mean that the scale can in part explain constructs such as job satisfaction, individual or team performance, well-being, or team dynamics. For example, we have already used the scale as a predictor of job satisfaction in a model explaining voluntary turnover [43]. For industry, the scale can support efforts to evaluate and improve infrastructure-related aspects of the developer experience in organizational settings, for example, by comparing results across remote and on-site workers or by location.

The methodological contribution of this work is an exemplar paper in scale construction for a software engineering specific context. While prior scales for software engineering have been developed, they usually stop at an expert review or by just simply adapt scales from other contexts. We hope our work together with the guidelines by Graziotin et al. [23] inspires further work in software engineering scale development and improvement in measurement.

5.1 Limitations

The developed scale measures attitudes of software professionals toward technical and sociotechnical infrastructure, that is how it is perceived, not the quality of this infrastructure. Further empirical studies are needed to determine the relationship between perception and actual quality. Furthermore, we do not know much what professionals’ attitude towards infrastructure affects, as one has to know how to measure the construct before it can be linked strongly to outcomes.

Overall attitude towards software infrastructure can be inferred by using the scale formatively [57], but attitude toward individual pieces of infrastructure cannot be inferred using this scale. Defining

latent constructs is difficult, and our division into technical and sociotechnical infrastructure, while well justified, is not self-evident to practitioners. Thus, the definitions have to be shown to respondents when using the scale. Furthermore, the boundaries between these two constructs can be partially overlapping, as for example, version control systems can be used in technical and sociotechnical ways. Theorizing this distinction in software engineering would be valuable for future research.

One limitation of the current scale is the exclusion of all but one reverse-worded item. These items were originally included to control for response bias, but most of them exhibited low communalities and factor loadings. Thus, the resulting instrument includes mostly positively worded items, which may limit its robustness to inattentive responding. However, strong model fit and internal consistency of the refined scale mitigate this concern, and we further recommend using attention check items with the scale.

The scale was primarily constructed by academics, with additional items generated by a large language model (LLM). While the impact of using generative AI is unknown, all LLM-generated items underwent the same validation as human-generated ones. We found existing LLM guidelines [69] too exhaustive for our purposes. However, the usage is documented in the replication package. Of the 11 final items, 6 originated from the LLM (2 technical, 4 sociotechnical). Similarly, exploratory factor analysis involves subjective decisions, different cutoffs or their order could yield different results. We report our criteria: first removing cross-loading items, then those with low loadings. As expert evaluation is also subjective, we provide quantitative validation and document our expert review process in the replication package.

The affective content in the selected items, that is, information related to emotions, moods, and feelings, as discussed in Section 4.4, is somewhat limited. Although some items do reflect affective components, they are fewer in numbers than items representing cognitive beliefs or past experiences. This may underrepresent the emotional dimension of attitudes in the resulting scale. This result emerged from following the scale development process, where items were retained based on psychometric properties rather than on theoretical balance. However, the tripartite theory of attitudes does not assume equal contributions from all components, and thus we view our scale as theoretically sound.

Measurement invariance was not assessed in the current study. As our primary goal was to develop and validate the scale structure, future research should examine whether the factor structure and item loadings remain consistent across subgroups (e.g., by gender, region, or job role). Similarly, we did not assess test-retest reliability. It is financially difficult to assess measurement invariance before the scale is in use. Lastly, we have discussed some of the limitations in data gathering in a previous work [42].

6 Conclusions

Software professionals are often interested in the technologies they use and the practices embedded in their teams' ways of working, as evidenced by the popular Stack Overflow Developer Survey³. However, despite infrastructure being widely acknowledged as an important aspect of software development [65], software engineering

research has lacked validated tools to rigorously assess professionals' attitudes toward it. Following the repeated calls for developing software engineering specific measures for behavioral and human factors research [23, 26], we have developed two scales with psychometric properties to measure software professionals attitudes towards technical and sociotechnical infrastructure. Through rigorous item reduction and EFA, we identified a clear two-factor structure representing technical and sociotechnical dimensions, with strong loadings, low item complexity, and good model fit. CFA further supported factorial validity of the scale in a randomized subsample, showing excellent fit.

The content and face validity of the items were supported by their alignment with the cognitive, affective, and behavioral components of attitudes, and by expert judgment during item generation and refinement stages. Convergent validity was demonstrated through moderate to strong correlations with theoretically related constructs of autonomy, feedback from the job diagnostic survey and job characteristics model [27, 35], and job satisfaction. Criterion-related validity analyses showed that attitudes toward technical infrastructure significantly predicted job satisfaction, autonomy, and feedback. Finally, discriminant validity was established with HTMT ratios of correlations, the Fornell–Larcker criterion, and model comparison.

These results provide strong empirical support for the reliability of the scale and validity of its structure. The scale offers a theoretically grounded and psychometrically supported tool for future research examining the role of workplace infrastructure in shaping employee attitudes and outcomes. Our work thus paves the way into theory-testing questionnaire research in the software engineering domain.

Future research can use this instrument to examine how attitudes toward workplace infrastructure shape a range of behaviors and outcomes. For instance, developers with positive attitudes toward technical and sociotechnical infrastructure may exhibit higher productivity, stronger organizational commitment, and more effective team coordination. They may also be more likely to engage in knowledge sharing, or produce higher-quality code. Our scale enables such hypotheses to be empirically tested for the first time in software engineering research.

Data Availability

A replication package comprising (1) a dataset documenting the creation of the scale with item reduction through panel review and favourability ratings; (2) a dataset including data from participants who consented to their data being shared when administering the scale; (3) analysis scripts, is available at Zenodo [44].

Acknowledgments

This project was supported by Natural Sciences and Engineering Research Council of Canada Discovery Grant RGPIN-2020-05001, Discovery Accelerator Supplement RGPAS-2020-00081, and the Kilam Foundation. We would like to thank Prof. Marcos Kalinowski, Prof. Nicole Novielli and Prof. Lucas Gren for their valuable contributions to this research.

³<https://survey.stackoverflow.co/2024/>

References

- [1] 2025. *APA Dictionary of Psychology* (2nd ed.). American Psychological Association. <https://dictionary.apa.org/construct-validity> Entry: [construct validity].
- [2] Pierre-Richard Agénor. 2010. A theory of infrastructure-led development. *Journal of Economic Dynamics and Control* 34, 5 (2010), 932–950. doi:10.1016/j.jedc.2010.01.009
- [3] Icek Ajzen. 1985. *From Intentions to Actions: A Theory of Planned Behavior*. Springer Berlin Heidelberg, Berlin, Heidelberg, 11–39. doi:10.1007/978-3-642-69746-3_2
- [4] John W Atkinson. 1957. Motivational determinants of risk-taking behavior. *Psychological review* 64, 6p1 (1957), 359. doi:10.1037/h0043445
- [5] Sebastian Baltes and Paul Ralph. 2022. Sampling in software engineering research: A critical review and guidelines. *Empirical Software Engineering* 27, 4 (2022), 94. doi:10.1007/s10664-021-10072-8
- [6] Albert Bandura. 1986. *Social foundations of thought and action*. Prentice Hall.
- [7] Herbert Blumer. 1986. *Symbolic interactionism: Perspective and method*. Univ of California Press.
- [8] Walter Buhr. 2003. *What is infrastructure?* Technical Report. Volkswirtschaftliche Diskussionsbeiträge.
- [9] Bernard CK Choi and Anita WP Pak. 2004. A catalog of biases in questionnaires. *Preventing chronic disease* 2, 1 (2004), A13.
- [10] Lee J Cronbach and Paul E Meehl. 1955. Construct validity in psychological tests. *Psychological bulletin* 52, 4 (1955), 281.
- [11] Sarah D'Angelo, Jessica Lin, Jill Dicker, Carolyn D. Egelman, Maggie Hodges, Collin Green, and Ciera Jaspan. 2024. Measuring Developer Experience With a Longitudinal Survey. *IEEE Softw.* 41, 4 (2024), 19–24. doi:10.1109/MS.2024.3386027
- [12] Fred D Davis. 1989. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS quarterly* (1989), 319–340. doi:10.2307/249008
- [13] Edward L Deci, Anja H Olafsen, and Richard M Ryan. 2017. Self-determination theory in work organizations: The state of a science. *Annual review of organizational psychology and organizational behavior* 4 (2017), 19–43. doi:10.1146/annurev-orgpsych-032516-113108
- [14] Robert F DeVellis and Carolyn T Thorpe. 2021. *Scale development: Theory and applications*. Sage publications.
- [15] Sara Dolnicar and Bettina Grün. 2014. Including Don't know answer options in brand image surveys improves data quality. *International Journal of Market Research* 56, 1 (2014), 33–50. doi:10.2501/IJMR-2013-043
- [16] Eliezer Dutra, Patricia Lima, Cristina Cerdeira, Bruna Diirr, and Gleison Santos. 2022. TACT: An Instrument to Assess the organizational Climate of agile teams-A Preliminary Study. *Journal of Software Engineering Research and Development* 10 (2022), 1–1. doi:10.5753/jsred.2021.1973
- [17] Jacquelynne S Eccles. 1983. Expectancies, values, and academic behaviors. In *Achievement and achievement motives*. Freeman, 75–146.
- [18] Craig K Enders and Deborah L Bandalos. 2001. The relative performance of full information maximum likelihood estimation for missing data in structural equation models. *Structural Equation Modeling: A Multidisciplinary Journal* 8, 3 (2001), 430–457. doi:10.1207/S15328007SEM0803_5
- [19] Leandre R Fabrigar, Duane T Wegener, Robert C MacCallum, and Erin J Strahan. 1999. Evaluating the use of exploratory factor analysis in psychological research. *Psychological methods* 4, 3 (1999), 272.
- [20] Leon Festinger. 1957. *A theory of cognitive dissonance*. Stanford university press.
- [21] Martin Fishbein and Icek Ajzen. 1977. *Belief, attitude, intention, and behavior: An introduction to theory and research*. Reading, MA: Addison-Wesley.
- [22] Claes Fornell and David F Larcker. 1981. Evaluating structural equation models with unobservable variables and measurement error. *Journal of marketing research* 18, 1 (1981), 39–50.
- [23] Daniel Graziotin, Per Lenberg, Robert Feldt, and Stefan Wagner. 2021. Psychometrics in behavioral software engineering: A methodological introduction with guidelines. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 31, 1 (2021), 1–36. doi:10.1145/3469888
- [24] Daniel Graziotin, Xiaofeng Wang, and Pekka Abrahamsson. 2014. Happy software developers solve problems better: psychological measurements in empirical software engineering. *PeerJ* 2 (2014), e289. doi:10.7717/peerj.289
- [25] Daniel Graziotin, Xiaofeng Wang, and Pekka Abrahamsson. 2015. Do feelings matter? On the correlation of affects and the self-assessed productivity in software engineering. *Journal of Software: Evolution and Process* 27, 7 (2015), 467–487. doi:10.1002/smr.1673
- [26] Lucas Gren. 2018. Standards of validity and the validity of standards in behavioral software engineering research: the perspective of psychological test theory. In *Proceedings of the 12th ACM/IEEE international symposium on empirical software engineering and measurement*. 1–4.
- [27] J Richard Hackman and Greg R Oldham. 1976. Motivation through the design of work: Test of a theory. *Organizational behavior and human performance* 16, 2 (1976), 250–279. doi:10.1016/0030-5073(76)90016-7
- [28] Geoffrey Haddock, Gregory R Maio, et al. 2004. *Contemporary perspectives on the psychology of attitudes*. Psychology Press Hove.
- [29] Ole Hanseth and Nina Lundberg. 2001. Designing work oriented infrastructures. *Computer Supported Cooperative Work (CSCW)* 10 (2001), 347–372. doi:10.1023/A:1012727708439
- [30] Eddie Harmon-Jones and Judson Mills. 2019. *An introduction to cognitive dissonance theory and an overview of current perspectives on the theory*. American Psychological Association.
- [31] Stephen N Haynes, David Richard, and Edward S Kubany. 1995. Content validity in psychological assessment: A functional approach to concepts and methods. *Psychological assessment* 7, 3 (1995), 238. doi:10.1037/1040-3590.7.3.238
- [32] Fritz Heider. 2013. *The psychology of interpersonal relations*. Psychology Press.
- [33] Jörg Henseler, Christian M Ringle, and Marko Sarstedt. 2015. A new criterion for assessing discriminant validity in variance-based structural equation modeling. *Journal of the academy of marketing science* 43 (2015), 115–135. doi:10.1007/s11747-014-0403-8
- [34] Rashina Hoda. 2021. Socio-technical grounded theory for software engineering. *IEEE Transactions on Software Engineering* 48, 10 (2021), 3808–3832. doi:10.1109/TSE.2021.3106280
- [35] Jacqueline R Idaszak and Fritz Drasgow. 1987. A revision of the Job Diagnostic Survey: Elimination of a measurement artifact. *Journal of applied psychology* 72, 1 (1987), 69.
- [36] Andrew T Jebb, Vincent Ng, and Louis Tay. 2021. A review of key Likert scale development advances: 1995–2019. *Frontiers in psychology* 12 (2021), 637547. doi:10.3389/fpsyg.2021.637547
- [37] Timothy A Judge, Joyce E Bono, and Edwin A Locke. 2000. Personality and job satisfaction: the mediating role of job characteristics. *Journal of applied psychology* 85, 2 (2000), 237. doi:10.1037/0021-9010.85.2.237
- [38] Henry F Kaiser. 1974. An index of factorial simplicity. *Psychometrika* 39, 1 (1974), 31–36. doi:10.1007/BF02291575
- [39] Daniel Katz. 1960. The functional approach to the study of attitudes. *Public opinion quarterly* 24, 2 (1960), 163–204. doi:10.1086/266945
- [40] Philip J Kitchen, Gayle Kerr, Don E Schultz, Rod McColl, and Heather Pals. 2014. The elaboration likelihood model: review, critique and research agenda. *European Journal of marketing* 48, 11/12 (2014), 2033–2050. doi:10.1108/EJM-12-2011-0776
- [41] Franki YH Kung, Navio Kwok, and Douglas J Brown. 2018. Are attention check questions a threat to scale validity? *Applied Psychology* 67, 2 (2018), 264–283. doi:10.1111/apps.12108
- [42] Miikka Kuuttila and Paul Ralph. 2025. Methodological and Practical Challenges in Longitudinal, Large-Scale, Collaborative Questionnaire Survey Research. In *2025 IEEE/ACM International Workshop on Methodological Issues with Empirical Studies in Software Engineering (WSESE)*. IEEE, 64–69. doi:10.1109/WSESE66602.2025.00017
- [43] Miikka Kuuttila, Paul Ralph, Huilian Sophie Qiu, Ronnie de Souza Santos, Morakot Choetkiertikul, Amin Milani Fard, Rana Alkadhi, Xavier Devroey, Gregorio Robles, Hideaki Hata, Sebastian Baltes, Hera Arif, Vladimir Kovalenko, Shalini Chakraborty, Eray Tuzun, and Gianisa Adisaputri. 2026. Staying or Leaving? How Job Satisfaction, Embeddedness and Antecedents Predict Turnover Intentions of Software Professionals. In *Proceedings of the 2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE '26)* (Rio de Janeiro, Brazil). ACM, New York, NY, USA, 1–13. doi:10.1145/3744916.3773124
- [44] Miikka Kuuttila, Paul Ralph, Huilian Sophie Qiu, Ronnie de Souza Santos, Morakot Choetkiertikul, Amin Milani Fard, Rana Alkadhi, Xavier Devroey, Gregorio Robles, Hideaki Hata, Sebastian Baltes, Vladimir Kovalenko, Shalini Chakraborty, Eray Tuzun, Hera Arif, Gianisa Adisaputri, Kelly Garcés, Anielle Severo Lisboa de Andrade, Eyram Amedzor, Bimpe Ayoola, Keisha Gaspard-Chickoree, and Arazoo Hoseyni. 2025. *Replication package for the paper "The Software Infrastructure Attitude Scale (SIAS): A Questionnaire Instrument for Measuring Professionals' Attitudes Toward Technical and Sociotechnical Infrastructure"*. doi:10.5281/zenodo.15962431
- [45] Stefano Lambiasi. 2024. Cultural and socio-technical aspects in software development. In *Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering*. 482–487. doi:10.1145/3661167.3661230
- [46] Per Lenberg, Robert Feldt, and Lars Göran Wallgren. 2015. Behavioral software engineering: A definition and systematic literature review. *Journal of Systems and software* 107 (2015), 15–37. doi:10.1016/j.jss.2015.04.084
- [47] Yao Lyu, Jie Cai, and John M Carroll. 2025. A Systematic Literature Review of Infrastructure Studies in SIGCHI. *arXiv preprint arXiv:2504.09612* (2025).
- [48] Neil Malhotra. 2008. Completion time and response order effects in web surveys. *Public opinion quarterly* 72, 5 (2008), 914–934. doi:10.1093/poq/nfn050
- [49] Claudio Saraiva Mattos, Eliezer Dutra, and Gleison Santos. 2024. An Instrument for Assessing Power Distance in Agile Organizations-Preliminary Results. In *Proceedings of the XXIII Brazilian Symposium on Software Quality*. 136–146. doi:10.1145/3701625.3701653
- [50] Andrew Meneely and Laurie Williams. 2011. Socio-technical developer networks: Should we trust our measurements?. In *Proceedings of the 33rd International Conference on Software Engineering*. 281–290. doi:10.1145/1985793.1985832
- [51] Richard E Nisbett and Timothy D Wilson. 1977. The halo effect: Evidence for unconscious alteration of judgments. *Journal of personality and social psychology* 35, 4 (1977), 250. doi:10.1037/0022-3514.35.4.250

- [52] Nicole Novielli and Alexander Serebrenik. 2019. Sentiment and emotion in software engineering. *IEEE Software* 36, 5 (2019), 6–23. doi:10.1109/MS.2019.2924013
- [53] Abraham Naftali Oppenheim. 2000. *Questionnaire design, interviewing and attitude measurement*. Bloomsbury Publishing.
- [54] Thomas M. Ostrom. 1969. The relationship between the affective, behavioral, and cognitive components of attitude. *Journal of Experimental Social Psychology* 5, 1 (1969), 12–30. doi:10.1016/0022-1031(69)90003-1
- [55] Maarten Ottens, Maarten Franssen, Peter Kroes, and Ibo Van De Poel. 2006. Modelling infrastructures as socio-technical systems. *International journal of critical infrastructures* 2, 2-3 (2006), 133–145. doi:10.1504/IJCIS.2006.009433
- [56] Richard E Petty and John T Cacioppo. 1986. The elaboration likelihood model of persuasion. In *Advances in experimental social psychology*. Vol. 19. Elsevier, 123–205. doi:10.1016/S0065-2601(08)60214-2
- [57] Paul Ralph, Miikka Kuutila, Hera Arif, and Bimpe Ayoola. 2024. Teaching Software Metrology: The Science of Measurement for Software Engineering. In *Handbook on Teaching Empirical Software Engineering*. Springer, 101–154. doi:10.1007/978-3-031-71769-7_5
- [58] William Revelle. 2015. Package ‘psych’. *The comprehensive R archive network* 337, 338 (2015), 161–165.
- [59] Mikko Rönkkö and Eunseong Cho. 2022. An updated guideline for assessing discriminant validity. *Organizational research methods* 25, 1 (2022), 6–14. doi:10.1177/1094428120968614
- [60] Yves Rosseel. 2012. lavaan: An R package for structural equation modeling. *Journal of statistical software* 48 (2012), 1–36.
- [61] James A Russell. 2003. Core affect and the psychological construction of emotion. *Psychological review* 110, 1 (2003), 145.
- [62] Susan Leigh Star. 1999. The ethnography of infrastructure. *American behavioral scientist* 43, 3 (1999), 377–391. doi:10.1177/00027649921955326
- [63] Susan Leigh Star and Karen Ruhleder. 1994. Steps towards an ecology of infrastructure: complex problems in design and access for large-scale collaborative systems. In *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work (CSCW '94)*. 253–256. doi:10.1145/192844.193021
- [64] Paul C Stern, Thomas Dietz, Troy Abel, Gregory A Guagnano, and Linda Kalof. 1999. A value-belief-norm theory of support for social movements: The case of environmentalism. *Human ecology review* (1999), 81–97.
- [65] Margaret-Anne Storey, Neil A Ernst, Courtney Williams, and Eirini Kalliamvakou. 2020. The who, what, how of software engineering research: a socio-technical framework. *Empirical Software Engineering* 25 (2020), 4097–4129. doi:10.1007/s10664-020-09858-z
- [66] William MK Trochim and James P Donnelly. 2001. *Research methods knowledge base*. Vol. 2. Atomic dog publishing Cincinnati, OH.
- [67] Viswanath Venkatesh, Michael G Morris, Gordon B Davis, and Fred D Davis. 2003. User acceptance of information technology: Toward a unified view. *MIS quarterly* (2003), 425–478. doi:10.2307/30036540
- [68] Scott I Vrieze. 2012. Model selection and psychological theory: a discussion of the differences between the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). *Psychological methods* 17, 2 (2012), 228. doi:10.1037/a0027127
- [69] Stefan Wagner, Marvin Muñoz Barón, Davide Falessi, and Sebastian Baltes. 2025. Towards Evaluation Guidelines for Empirical Studies Involving LLMs. In *2025 IEEE/ACM International Workshop on Methodological Issues with Empirical Studies in Software Engineering (WSESE)*. 24–27. doi:10.1109/WSESE66602.2025.00011
- [70] Bert Weijters and Hans Baumgartner. 2012. Misresponse to reversed and negated items in surveys: A review. *Journal of Marketing Research* 49, 5 (2012), 737–747. doi:10.1509/jmr.11.0368
- [71] Jacob Westfall and Tal Yarkoni. 2016. Statistically controlling for confounding constructs is harder than you think. *PLoS one* 11, 3 (2016), e0152719. doi:10.1371/journal.pone.0152719
- [72] Darryl Wilkinson. 2019. Towards an archaeological theory of infrastructure. *Journal of Archaeological Method and Theory* 26, 3 (2019), 1216–1241. doi:10.1007/s10816-018-9410-2
- [73] Brett Williams, Andrys Onsmann, and Ted Brown. 2010. Exploratory factor analysis: A five-step guide for novices. *Australasian journal of paramedicine* 8 (2010), 1–13. doi:10.33151/ajp.8.3.93
- [74] Titus Winters, Tom Manshreck, and Hyrum Wright. 2020. *Software engineering at google: Lessons learned from programming over time*. O'Reilly Media, Inc.
- [75] Roger L Worthington and Tiffany A Whittaker. 2006. Scale development research: A content analysis and recommendations for best practices. *The counseling psychologist* 34, 6 (2006), 806–838. doi:10.1177/0011000006288127