

# Sampling in Software Engineering Research: A Critical Review and Guidelines

SEBASTIAN BALTES, The University of Adelaide, Australia

PAUL RALPH, Dalhousie University, Canada

Representative sampling appears rare in software engineering research. Not all studies need representative samples, but a general lack of representative sampling undermines a scientific field. This study therefore investigates the state of sampling in recent, high-quality software engineering research. The key findings are: (1) random sampling is rare; (2) sophisticated sampling strategies are very rare; (3) sampling, representativeness and randomness do not appear well-understood. To address these problems, the paper synthesizes existing knowledge of sampling into a succinct primer and proposes extensive guidelines for improving the conduct, presentation and evaluation of sampling in software engineering research. It is further recommended that while researchers should strive for more representative samples, disparaging non-probability sampling is generally capricious and particularly misguided for predominately qualitative research.

CCS Concepts: • **General and reference** → **Surveys and overviews**; **Design**; *Reference works*; *Empirical studies*; *Measurement*; *Evaluation*.

Additional Key Words and Phrases: Sampling, Representative sampling, Random sampling, Purposive sampling, Probability sampling, Research methodology

## ACM Reference Format:

Sebastian Baltes and Paul Ralph. 2020. Sampling in Software Engineering Research: A Critical Review and Guidelines. *ACM Trans. Softw. Eng. Methodol.* 0, 0, Article 0 (2020), 21 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

Most research involves selecting some of many possible items to study, i.e. *sampling*. Sampling is crucial for positivist research (e.g. questionnaires, experiments, positivist case studies) because unrepresentative samples bias results. Sampling is just as important for interpretive and constructivist research (e.g. interview studies, grounded theory, interpretivist case studies)—although it is not always called ‘sampling’—because selecting poor research sites hinders data collection and focusing on the wrong topics undermines theory building. Sampling for software engineering (SE) research is particularly interesting and troublesome in at least three ways:

- (1) For many SE phenomena, there is no good list from which to draw a sample.
- (2) Some SE studies adopt poorly understood sampling strategies such as random sampling from a non-representative surrogate population.
- (3) Many SE articles evince deep misunderstandings of *representativeness*—the key criteria for assessing sampling in positivist research (see Section 2.6).

While some previous articles and books discuss sampling in the SE context (e.g., [3, 19, 40, 46]) and sampling specifically for questionnaire studies [11, 30], this paper goes into significantly more depth and discusses more SE-specific sampling problems. More particularly, the purpose of this paper is as follows.

**Purpose:** (1) *investigate the state of sampling in SE research*; (2) *provide a detailed primer—tailored to the SE context—on sampling concerns and techniques*.

---

Authors' addresses: Sebastian Baltes, The University of Adelaide, Adelaide, Australia, [sebastian.baltes@adelaide.edu.au](mailto:sebastian.baltes@adelaide.edu.au); Paul Ralph, Faculty of Computer Science, Dalhousie University, Halifax, Canada, [paulralph@dal.ca](mailto:paulralph@dal.ca).

---

2020. 1049-331X/2020/0-ART0 \$15.00  
<https://doi.org/10.1145/1122445.1122456>

Next, we provide a detailed introduction to sampling concepts (Section 2). Then we describe our empirical methodology (Section 3) and results (Section 4). Section 5 gives specific guidelines for conducting, presenting and evaluating sampling. Section 6 summarizes the paper’s contributions and limitations.

## 2 SAMPLING: A PRIMER

*Sampling* is the process of selecting a smaller group of items to study (a *sample*) from a larger group of items of interest. The group of items of interest is called the *population*, and the (usually imperfect) population list is called the *sampling frame*. Researchers have developed diverse approaches for selecting samples from populations. While there are dozens of sampling strategies, this section focuses on those most applicable to SE research.

Sampling strategies can be divided into two broad categories: *probability* and *non-probability*. This section elaborates some of the approaches in each category, then describes *multi-stage sampling*, and explores the meaning and interpretation of *representativeness*.

The following exposition on sampling is primarily synthesized from several textbooks [8, 27, 54]. To the best of our knowledge, all of these concepts are widely accepted across the research methods literature, except the *cone of sampling* approach, which we propose in Section 2.5. We include this primer before delving into our study because a good grasp of all the different sampling strategies and many of the issues dissected in this section is necessary to understand the study that follows.

### 2.1 Non-probability sampling

Non-probability sampling includes all of the sampling techniques that do not employ randomness. This section discusses four kinds of non-probability sampling: convenience, purposive, referral-chain (or snowball) sampling, and respondent-driven sampling.

While some (positivist) methodologists present non-probability sampling as intrinsically inferior to probability sampling (e.g. [58]), it really depends on the aims and context of the research (see Section 2.6). Each sampling strategy is appropriate in some circumstances and inappropriate in others.

**2.1.1 Convenience sampling.** *Items are selected based on availability or expedience.* When we select people or things to study arbitrarily, or based on them being nearby, available or otherwise easy to study, we adopt convenience sampling. Convenience sampling is controversial because it is very popular, especially in laboratory experiments, despite threatening generalizability (cf. [1, 26]). Being quick and inexpensive, it is most appropriate for pilot studies and studying a universal phenomenon like cognitive biases [38].

**2.1.2 Purposive sampling.** *Items are selected based on their usefulness for achieving the study’s objective.* When we select our objects of study according to some logic or strategy, carefully but not randomly, we employ purposive or purposeful sampling [41].

Guidelines for purposive sampling are often provided in the context of selecting sites or data sources for predominately qualitative, anti-positivist research (e.g. [37]). While non-probability samples can be representative of broader populations, the goal of non-probability sampling is often to find accessible, information rich cases, sites, organizations or contexts from which researchers can learn a lot about their topic of study [41].

Purposive sampling includes a whole range of approaches, for example:

- (1) We study projects hosted on GitHub<sup>1</sup> because it is popular and has good tool support.
- (2) We recruit a panel of experts for a focus group (“expert sampling”).
- (3) We select projects that are as diverse<sup>2</sup> as possible (“heterogeneity sampling”).

<sup>1</sup><https://github.com>

<sup>2</sup>Diversity can be defined along many different axes, gender being one of them [60].

- (4) We extract the sample from a repository using a specific search query, as in a systematic literature review [29] (“search-based sampling”). Search-based sampling is limited by the scope and features of the search engine and the researcher’s ability to construct a good query.

The key advantages of purposive sampling are: (1) the researcher can exercise expert judgment; (2) the researcher can ensure representativeness on specific dimension (see Section 2.6); (3) no sampling frame is needed. The main challenge with purposive sampling is that it is intrinsically subjective and opportunistic.

*2.1.3 Referral-chain (snowball) sampling. Items are selected based on their relationship to previously selected items.* Referral-chain sampling (also called snowball sampling) is useful when there is no good sampling frame for the population of interest [16]. For example, there is no comprehensive list of black-hat hackers or software developers who have experienced sexual harassment. However, members of such “hidden populations” often know each other. Snowball sampling with human participants therefore works by finding a few individuals in the population, studying them, and then asking them to refer other members of the population whom they know.

In SE, snowball sampling is commonly used in systematic literature reviews to supplement search-based sampling. When we begin with an article A, searching the papers A cites is sometimes called *backward snowballing* while searching the papers that cite A is sometimes called *forward snowballing*. We can study software libraries, methods and services (in service-oriented architectures), in much the same way.

The advantage of snowball sampling is that it helps us to identify items that are not in our sampling frame. However, snowball sampling has two major limitations: 1) it biases results toward more connected people (or things); 2) it can lead to sampling a small, highly-interconnected subset of a larger population.

*2.1.4 Respondent-driven sampling. An advanced form of referral-chain sampling designed to mitigate sampling bias.* To mitigate the problems with snowball sampling (of people, not things), *respondent-driven sampling*, a comprehensive approach for mitigating bias in referral chain sampling. For example, researchers might:

- (1) Begin with diverse initial participants (seeds) who (i) have large social networks, (ii) represent different sub-populations, (iii) do not know each other, and (iv) can influence peers to participate.
- (2) Have participants recruit, rather than identify, peers. This reduces selection bias by the researcher.
- (3) Limit recruitment such that each participant can only recruit a small number of peers (typically 3). This prevents highly-connected participants from flooding the sample.
- (4) Require many (e.g. 20) recruitment waves. This generates longer referral chains, decreasing the risk of oversampling from a highly connected subset of the population.
- (5) Prevent individuals from participating more than once.
- (6) Continue recruitment until the sample reaches *equilibrium*, the point where the distribution of variables of interest is stable.
- (7) Apply a mathematical model to account for sampling bias [25, 28].

While the details of the mathematical model used are beyond the scope of this paper, more information and tools are available.<sup>3</sup> Respondent-driven sampling has, however, been criticized for producing optimistic confidence intervals, and generally being “substantially less accurate than generally acknowledged” [22].

## 2.2 Probability sampling

Probability sampling includes all of the sampling techniques that employ randomness. In everyday language, *random* is often used to mean arbitrary or without logic. In sampling and statistics more generally, however, *random* means that each item in the population has an equal probability of selection [10]. Standing on a street corner interviewing “random” pedestrians is not random in the statistical sense. Recruiting participants using email or advertising on social networks is not random. Assigning participants to experimental conditions in the

<sup>3</sup><http://respondentdrivensampling.org/>

order in which they arrive at a laboratory is not random. Practically speaking, any selection without using a random number generator,<sup>4</sup> probably is not random.

Research on and guidelines for probability sampling are often written from the perspective of positivist questionnaire studies. These studies are typically descriptive (e.g. political polling) or explanatory (i.e. testing a theory based on perceptions of respondents). Since examining the entire population is usually impractical, the researcher selects a subset of the population (a sample) and attempts to estimate a property of the population by statistically analyzing the sample. Probability sampling ostensibly facilitates such statistical generalization (cf. [39]).

The overwhelming challenge for applying any kind of probability sampling in SE is the absence of comprehensive sampling frames for common units of analysis (see Section 5). This section describes some probability sampling approaches that are relevant to SE research.

*2.2.1 Whole frame. All items in the sampling frame are selected.* Suppose a researcher wants to assess morale of developers at a specific software development company. The company provides a complete list of developers and their contact information. The researcher creates a survey with questions about job satisfaction, views of the company, employees' future plans, etc. They send the questionnaire to all of the developers—the entire sampling frame. Whether this is technically “sampling” is debatable, but it is an important option to consider, especially when data collection and analysis are largely automated.

*2.2.2 Simple Random Sampling. Items are selected entirely by chance, such that each item has equal chance of inclusion.* Now suppose the results of the above morale survey are less than spectacular. The researcher decides to follow up with some in-depth interviews. However, interviewing all 10,000 developers is clearly impractical, so the researcher assigns each developer a number between 1 and 10,000, uses a random number generator to select 20 numbers in the same range, and interviews those 20 developers. This is simple random sampling because the researcher simply chooses  $n$  random elements from the population.

*2.2.3 Systematic Random Sampling. Given an interval,  $x$ , every  $x$ th item is selected, from a starting point chosen entirely by chance.* To complement the interviews, the researcher decides to review developers' posts on the company's messaging system (e.g. Slack). Suppose there is no easy way to jump to a random message and there are too many messages to read them all. So the researcher generates a random number between 1 and 100 (say, 47) and then reads message 47, 147, 247, etc. until reaching the end of the messages. This is systematic random sampling. Each post still has an equal probability of inclusion; however, the consistent interval could bias the sample if there is a recurring pattern that coincides with the interval (e.g. taking annual weather data in the middle of summer vs. the middle of winter).

*2.2.4 Panel sampling. The same sample is studied two or more times.* Now suppose the researcher implements a program for improving morale, and a year later, re-interviews the same 20 developers to see if their attitudes have changed. This is called *panel sampling* because the same *panel* of developers is sampled multiple times. Panel sampling is probability sampling as long as the panel is selected randomly.

*2.2.5 A repository mining example of probability sampling.* All four of these probability sampling strategies could also be applied in, for example, repository mining. We could (in principle) study every public project on GitHub (whole frame), or we can randomly select 50 projects (simple random sampling) or sort projects by starting date and study every 100th project (systematic random sampling) or take repeated measurements from the same 100 projects over time (panel sampling).

<sup>4</sup>True random number generation is available from numerous sources, including <https://www.random.org/>

### 2.3 Multistage sampling

Methodologists often present multistage sampling as a special case where two or more sampling strategies are intentionally combined (e.g. [57]). Two common approaches are stratified and cluster sampling.

*2.3.1 Stratified/Quota sampling. The sampling frame is divided into sub-frames with proportional representation.* Suppose that the developer morale survey discussed above reveals significant differences between developers who identify as white and those who do not. However, further suppose that 90% of the developers are white. To get more insight into these differences, the researcher might divide developers into two *strata*—white and non-white—and select 10 developers from each strata. If the developers are selected randomly, this is called *stratified random sampling*. If the developers are selected purposively, it is called *quota sampling*. This sampling strategy is interesting because it is *intentionally* non-representative [55].

We conceptualize these strategies as multistage because the researcher purposively chooses the strata (stage 1) before selecting the people or things to study (stage 2). The systematic review reported in this paper uses stratified random sampling.

*2.3.2 Cluster sampling. The sampling frame is divided into groups and items are drawn from a subset of groups.* Suppose that the company from our morale survey example has 20 offices spread around the world. If the researcher wants to conduct face-to-face interviews, traveling to all 20 offices could be prohibitively expensive. Instead, the researcher selects three offices (stage 1) and then selects 7 participants in each of these offices (stage 2). This is called *cluster sampling*. If and only if both selections are random, it is *cluster random sampling*. Cluster sampling works best when the groups (clusters) are similar to each other but internally diverse on the dimensions of interest.

Suppose that the researcher finds that the seven developers at one office seem much happier than developers in the rest of the company. If the researcher decides to conduct extra interviews at that office, in hopes of unraveling the sources of improved morale, this is called *adaptive cluster sampling* [52, 56].

### 2.4 Sampling in qualitative research

Qualitative researchers have to select both sites (e.g. teams, organizations, projects) and data sources (e.g. who to interview, which documents to read, which events to observe). Different qualitative research traditions (e.g. case study, grounded theory, phenomenology, ethnography) talk about this “selection” in significantly different ways [20]. Some qualitative researchers use the term “sampling” (e.g. [21]). Others argue that qualitative researchers should avoid the term “sampling” because it implies statistical generalization to a population [59, 61]. Others argue that there are many kinds of generalization, and qualitative researchers generalize from data to descriptions to concepts to theories, rather than from samples to populations (e.g. [35]).

This paper tries to clarify that sampling is distinct from statistical generalization. Predominately qualitative approaches including case studies, interview studies, grounded theory and action research typically use *non-probability sampling* to support *non-statistical generalization* from data to theory (see [7, 44, 49]). Predominately quantitative studies, especially questionnaire surveys, *sometimes* use *probability sampling* to support *statistical generalization* from samples to populations. As we shall see below, however, many quantitative studies also adopt non-probability sampling.

Selecting sites and data sources is a kind of sampling. When a researcher selects a site because it seems like there is something interesting there, that is purposive sampling. When a researcher interviews whoever will speak on a subject of interest, that is convenience sampling.

As fledgling concepts or theories begin to emerge, however, the researcher may use them to decide what to focus on next. In the grounded theory literature, *theoretical sampling* refers to selecting items to study based on an emerging theory [13]. For example, suppose the researcher from our running example begins generating

a theory of developer morale, which includes a preliminary category, “interpersonal conflict.” The researcher might purposively sample peer code reviews with many back-and-forths, because these reviews might contain evidence of interpersonal conflict.

## 2.5 The Cone of Sampling

Moving to a different example, suppose that we are interested in non-code documents in software projects (e.g. specifications, lists of contributors, diagrams, budgets). But we are especially interested in documents for open source systems. Furthermore, we have had good experiences mining GitHub, so we will limit ourselves to open source projects on GitHub. We only speak English, so we exclude all non-English documents. Now we randomly select 50 English-language open source projects on GitHub and then use simple random sampling to choose up to ten documents from each selected project. We end up with 500 documents. Now suppose we also contact the owners of each selected project to ask if they object to the research, and suppose two of them do, so we delete the corresponding 20 documents.

It is not clear what *sample*, *sampling frame* and *population* refer to in the example above. Is the sample the 500 documents we collected or the 480 we retained? What if we asked permission before collecting the data and never collected the 20 stories from the two objecting projects? Is the sampling frame all GitHub projects or just all GitHub projects that have non-code documents, that are open source, that are in English or some combination thereof? Is the population all documents, English documents, documents in open source systems, documents on GitHub, or some combination thereof? Is there a “study population” (e.g. English documents) that is narrower than a “theoretical population” (e.g. all documents). What do we present as part of the population definition and what do we present as selection criteria? Which part is sampling and which is applying inclusion/exclusion criteria?

Furthermore, is this probability sampling or not? We have already eliminated the vast majority of the world’s software projects before we employ randomness. Does claiming that a sample is good because we used probability sampling make any sense if we have previously excluded 99% of the objects of interest?

We suggest bypassing all this confusion by thinking of most studies as having a multistage *cone of sampling*. For example, Figure 1 (in Section 3) illustrates the sampling strategy used in our systematic literature review. We do not need to label each step in the sampling strategy or to claim that the sample is random or not. What matters is that the sampling procedures are clear enough to replicate and to understand how the sample might be biased.

## 2.6 Representativeness

Kruskal and Mosteller [31] argue, with extensive examples, that the term “representative” has been (mis)used in at least five ways:

- (1) as a “seal of approval bestowed by the writer”
- (2) as the “absence of selective forces in the sampling”
- (3) as a “miniature or small replica of the population”
- (4) as a claim that its members are “typical of the population” or “the ideal case”
- (5) as a claim to heterogeneity or that all subpopulations or classes are included (not necessarily proportionately) in the sample

For our purposes, *representativeness* is the degree to which a sample’s properties (of interest) resemble those of a target population. This section discusses common misunderstandings of representativeness and arguments for representativeness.

First, **representativeness is rooted in positivist epistemology**. Postmodernists, interpretivists, and constructivists reject the entire notion of statistical generalization on numerous grounds, including:

- Broadly applicable (“universal”) theories of social phenomenon simply do not exist [14].

- Each context is unique; therefore, findings from one do not translate wholesale into others [24].
- Statistical generalization precludes deep understanding of a particular social, political economic, cultural, technological context [51].

Representativeness of a sample (or site) is therefore not a valid evaluation criterion under these epistemologies. Contrastingly, in positivism, falsificationism and Bayesian epistemology, the primary purpose of sampling is to support statistically generalizing findings from a sample to a population. Representativeness is therefore the overwhelming quality criterion for sampling: good samples are representative; bad samples are biased.

Second, **representativeness is widely conflated with randomness**. Suppose that two researchers, Kaladin and Shalan, have a sampling frame of 10,000 software projects, with an average size of 750,000 lines of code: 70% open source and 30% closed source. Kaladin randomly selects 10 projects. Suppose 4 of them are open source and 6 are closed source, with an average size of 747,823 lines of code. Meanwhile, Shalan inspects a histogram of project size and discovers a bi-modal distribution, with clear clusters of large and small projects. Shalan purposively selects 7 large open source projects, 7 small open source project, 3 large closed-source projects and 3 small closed source projects.

This example illustrates several key points about representativeness:

- (1) Representativeness is dimension-specific. A sample can be representative on one parameter but not another.
- (2) Probability sampling does not guarantee representativeness on the dimensions of interest.
- (3) Non-random samples can be representative on the dimensions of interest.
- (4) Non-probability sampling can lead to more representative samples than probability sampling.
- (5) Random sampling depends on sample size.

If an unbiased, random sample is large enough, the law of large numbers dictates that its parameters will converge on the parameters of the sampling frame. The parameters of small samples, however, may differ greatly from the sampling frame. What constitutes a “large” sample depends on many factors including the strength of the effects under investigation, the type of statistical tests being used and the number of variables involved.

However, no sample size can overcome bias in the sampling frame. Suppose that Kaladin and Shalan want to draw inferences about all of the software projects conducted in Brazil. However, suppose the sampling frame is a list of all *public sector* projects in Brazil. Further suppose that public sector projects are generally larger and more likely to be open source than private sector projects, so Kaladin’s sample is biased not only toward open source projects but also toward larger projects, and Shalan’s sample is less representative than it first appeared.

Clearly then, *random* is not equal to *representative*. **Rather than defining representativeness, randomness is one of several arguments that sample *should be* representative** (Table 1). None of these approaches guarantee a representative sample, but each has some merit.

As discussed above, we can argue that the sample is representative because individuals were selected randomly. However, random selection will only produce a representative sample *most of the time* if the sample size is large and the sampling frame is unbiased.

Alternatively, suppose we are surveying developers about their perceptions of the relationship between agile practices and morale. We can argue that the larger and broader our sample, the less likely it is to have missed an important subpopulation. The breath argument supports generalization of correlations (e.g. between agility and morale). Breath is the argument of heterogeneity sampling, and can apply to convenience and snowball sampling where oversampling some subpopulations is a key threat.

However, the breath argument does not support point estimates. Suppose only 1% of our sample reports abandoning agile practices because of morale. While the point estimate of 1% is not reliable, abandoning agile practices over morale issues is probably rare. It seems highly unlikely that a survey of 10,000 developers from 100 countries, including thousands of companies and dozens of industries, would miss a large subpopulation of low-morale agile-abandoners.

Table 1. Reasonable Arguments for Representativeness

Name	GPRS* Argument	Threats
Random Selection	No	Each item in sampling frame had an equal probability of inclusion.
Size	No	Sample is so large that missing a subpopulation is unlikely.
Breadth	No	Sample captures a large variance on important parameters.
Parameter Matching	No	Sample and population parameters have similar distributions.
Universality	No	All possible samples are representative because the phenomenon affects the entire population equally
Postmodern Critique	No	The entire logic of statistically generalizing from a sample to a population is flawed.
Practical Critique	No	Generalizing to a population is not the purpose of this kind of study (e.g. case study, experiment).

\*Guaranteed to Produce a Representative Sample?

Another reasonable argument for representativeness is that sample parameters mirror known distributions of population parameters. If we know the approximate distributions for a population of projects' size, age, number of contributors, etc., we can compare the sample parameters (e.g. using the chi-square goodness-of-fit test) to see if they differ significantly. If the sample parameters are close to known population parameters, the sample is representative on those dimensions. If the sample and population match on known dimensions, it seems more likely (but not guaranteed) that they will also match on unknown dimensions. This is the argument of quota sampling.

A quite different argument is the appeal to universality. Suppose we have good reasons to believe that the phenomena of interest affects the entire population equally. For example, Fitt's Law predicts the time required to point at a target based on the target's size and distance [18]. Insofar as Fitt's law is universal, researchers can argue that sampling is irrelevant—all samples are representative. The appeal to universality could apply to many laboratory studies and is related to debates about generalizing from student participants to professionals (see [47]).

Contrastingly, some studies address sampling concerns by simply dismissing them on philosophical grounds (as described above). Others argue that, practically speaking, statistical generalization is not the purpose of the present study (see [48]).

Many studies, especially laboratory studies, simply ignore sampling concerns.<sup>5</sup> Sometimes a single sentence acknowledges the limitation that the results may not generalize (see also Section 4.2). Other studies give dubious arguments for representativeness [32].

An ongoing discussion is centered around students vs. professionals as study participants [17]. Suppose researchers conduct an experiment based on a convenience sample of six American white, male, professional

<sup>5</sup>Here, we do not cite specific examples because we do not wish to cast aspersions on individual researchers. While it is perhaps less credible this way, our credibility is not worth our colleagues' animosity. The truth of this should be obvious to any well-read software engineering researcher.

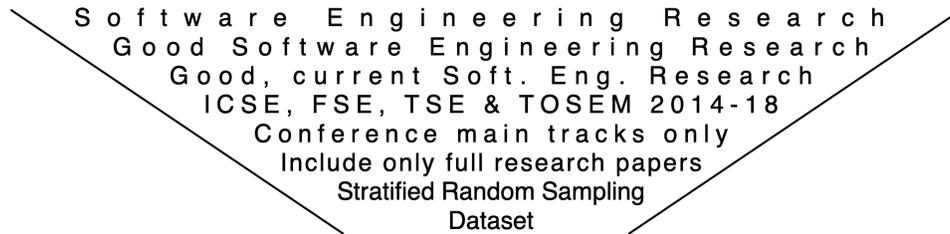


Fig. 1. Cone of sampling for the literature review we conducted (see Section 3).

developers, who have bachelor's degrees in software engineering and are between the ages of 30 and 40. This sample is patently not representative of professional developers in general just because the participants are not students. Small convenience samples do not support statistical generalizability.

### 3 METHOD

To investigate the state of sampling in software engineering, we manually retrieved and analyzed a collection of software engineering papers. This section describes the study's research questions, data collection and data analysis. The study design is based on common guidelines for systematic literature reviews and mapping studies [29, 43]. It is basically a positivist study.

#### 3.1 Objective and research questions

The objective of this study is to investigate the sampling techniques used in software engineering research, and their relationship to research methods and units of analysis. This objective motivates the following research questions. In recent, high-quality software engineering research...

- (RQ1:) ...what *sampling approaches* are most common?
- (RQ2:) ...how do authors *justify* their sampling approaches?
- (RQ3:) ...what *empirical research methodologies* are most common?
- (RQ4:) ...what *units of analysis* are most common?

#### 3.2 Sampling strategy and inclusion/exclusion criteria

Figure 1 summarizes our sampling strategy. Because of the time-intensive nature of the analysis (it took 15-30 minutes to code each paper with longer discussions of difficult cases), we aimed for a sample of 100 articles. The question was, how best to select these 100 articles?

Systematic reviews typically use search-based sampling, followed by various techniques for addressing sampling bias (e.g. reference snowballing) [29]. The idea is to retrieve *all* relevant papers. Since most SE articles involve sampling, search-based sampling will basically just retrieve all empirical SE studies, which is not very helpful. We want to know where the field is headed. This suggests focusing on recent papers in the most influential outlets. Consequently, we limit our sampling frame to articles published between 2014 and 2018 inclusive, in one of four outlets:

- (1) The International Conference on Software Engineering (ICSE)
- (2) Foundations of Software Engineering (FSE), which was held jointly with European Software Engineering Conference in 2015, 2017 and 2018.
- (3) IEEE Transactions on Software Engineering (TSE)

(4) ACM Transactions on Software Engineering and Methodology (TOSEM)

Moreover, we applied several inclusion and exclusion criteria:

- (1) Include: Only full papers.
- (2) Exclude: Front matter, letter from editor, etc.
- (3) For FSE and ICSE: Include only papers in the main technical track (for symmetry).
- (4) For TSE and TOSEM: Exclude journal extensions of papers already in the sample.

We did not evaluate the quality of the articles because they were all accepted at top venues and we were interested in their sampling technique, not their results. Applying these criteria produced a sampling frame of 1,565 full papers. The tool that we implemented to retrieve our sampling frame from DBLP is available on GitHub.<sup>6</sup>

Next, we applied stratified random sampling; that is, we randomly selected five papers from from each outlet-year (e.g. five papers published by TOSEM in 2016). This means that each outlet and each year have equal representation in our sample. We used a true-random number generator<sup>7</sup> to randomly select papers from each outlet-year, and extracted them from the sampling frame we had previously built. We provide the R script implementing our sampling approach as part of our supplementary material [2].

Like many SE studies, we adopt a poorly understood sampling approach. First, we purposively selected “full research papers published in four good outlets over five years” and then randomly selected items to study from this more manageable list. We can make the randomness argument to representativeness; however, the four venues in the sampling frame are obviously not representative of all SE research because they are among the most competitive and are all in English. Other researchers might have chosen a different sampling frame. There is no objective basis on which to select outlets or to study five years of four outlets vs. four years of five outlets vs. one year of 20 outlets. To proceed, we must simply make reasonable choices and explain their implications (more on this below).

### 3.3 Data extraction and analysis

Papers were randomly ordered and assigned an unique identifier from 1 to 100. Below, we use *PXX* to denote the *XX*th primary study (e.g. P35 is the 35th study in our sample). The complete list of primary studies is included in the supplementary material [2].

The first author reviewed each paper and recorded the following data points in a spreadsheet: venue, year, title, authors, length (pages), relevant quotations (paper summary, sampling description, quality attributes, sampling limitations, population), number of studies reported in the paper (usually one), empirical method, number of samples (usually one), sampling stages, sample origin, units of observation, properties analyzed, study population, sampling frame, sampling frame size, sample size, number of items studied. The complete dataset is available as supplementary material [2].

Some papers clearly stated the sampling technique, for example:

- (1) “We invited 739 developers, via e-mail using convenience sampling” (P35)
- (2) “From the top 500 projects, we sampled 30 projects uniformly at random” (P43)
- (3) “We used stratified sampling to identify potential interviewees” (P56).

However, many papers did not clearly explain their sampling technique (specifics below). We therefore had to infer the sampling technique from the text. For example, we inferred purposive sampling from the statement: “We prepared ... buggy code ... and asked workers to describe how to fix it. Only those who passed this qualifying test could proceed to our debugging tasks” (P92).

<sup>6</sup><https://github.com/sbaltes/dblp-retriever>

<sup>7</sup><https://www.random.org/>

Table 2. Frequency of sampling techniques found in SE research papers (n=179 sampling stages)

Type	Strategy	Frequency
Non-probability	Purposive	125 (69.8%)
Non-probability	Convenience	20 (11.2%)
Probability	Simple random	12 (6.7%)
Other	Whole sampling frame	9 (5.0%)
No sampling	No empirical study or generated data	9 (5.0%)
Probability	Stratified random	3 (1.7%)
Non-probability	Snowballing	1 (0.6%)

All ambiguous cases were reviewed by both authors and classified by consensus. For each ambiguous case, we developed a decision rule to guide future classifications. The most important decision rule was to code studies where authors applied certain filtering criteria such as popularity or experience as *purposive samples*. Another important rule was to use the code *sub-sample* to indicate when authors derived a new sample based on a previously introduced one (e.g. selected developers active in projects that had been sampled before). A third rule was to classify studies based on dominant methodology; for instance, a study that was predominately quantitative with a small qualitative component was simply coded as quantitative. For articles reporting multiple studies where some methods were primarily quantitative and others primarily qualitative, we captured both using a separate row for each study.

#### 4 RESULTS AND DISCUSSION

We analyzed 100 articles, of which 96 contained an empirical study. Some articles reported multiple studies; some studies had multiple stages of sampling (Table 3). We examined the results by *article* (100) *study* (127) and *sampling stage* (179). Note that multiple studies can use the same sample and multiple samples can be used by the same study.

The sample sizes ranged from 1 to 700,000 with a median sample size of 20. In two cases, it was not possible to derive the sample size from the descriptions in the paper. It was possible to derive the size of the corresponding sampling frames for only 54 sampling stages. The sizes of the reported sampling frames ranged from 3 to 2,000,000 with a median of 447.5. This section addresses each research question and then comments on observable trends.

##### 4.1 RQ1: Sampling techniques used

Table 2 shows the different sampling techniques used in the primary studies. The frequencies in Table 2 total more than 100 because some papers report multiple studies and some studies use multiple techniques.

The most common strategies were purposive (125) and convenience sampling (20). Only 15 stages utilized probability sampling—12 used simple random sampling and 3 used stratified random sampling. Nine stages analyzed their entire sampling frame. Nine stages did not analyze empirical data.

Half of the sampling stages that involved simple random sampling were sub-samples of previously derived non-random samples. The same applies for stratified random sampling (3 of 3 stages). This is important because random sampling from a nonrandom frame undermines the argument that the sample should be representative because it is random. See the coding of papers P16, P27, P35, P43, P47, P48 and P82 in the supplementary material for more details.

Samples are derived from a variety of sources (Table 4). In 49 cases (27%), the sample was based on an existing dataset, typically from related work. In 21 sampling stages, a subset of a larger sample presented in the

Table 3. Properties of studied papers (n=100)

Property	Count					
	0	1	2	3	4	≥ 5
Studies per paper	4	73	17	5	0	1
Samples per study	0	91	24	5	2	1
Sampling stages per sample	0	112	11	0	0	0

Table 4. Origins of samples in SE research (n=179 sampling stages)

Strategy	Definition	Frequency
Unclear	The strategy was not explained	51
Existing-sample(s)	Study used previously collected data, e.g., reported in related work	49
Sub-sample	Study used a subset of another sample presented in the same paper	21
Online-resources(s)	Sample retrieved from the internet (e.g., websites, mailing lists, LinkedIn)	20
Personal-network	Sample comprises artifacts or people that researchers had access to (e.g., students, industry contacts)	15
Other	For example public or corporate datasets, snowballing.	23
Gen-data	Study used generated data	5
No-study	Paper did not report an empirical study	4

corresponding paper was used. Twenty stages involved sampling from online resources, which was usually done in a purposive manner (17 out of 20 stages). Sources for such samples included the Android developer JavaDoc guide (P16), online Scala tutorials (P45), LinkedIn groups (P49), app stores (P80, P89, P98), or the online CVE database (P91). Fifteen stages involved sampling from the researchers’ personal networks (e.g. their students or colleagues) and five stages used generated data. In 51 cases (29%), the origin of the samples was unclear.

#### 4.2 RQ2: Authors’ justifications

Of the 100 papers in our sample, 74 provided some justification as to whether their sample exhibits certain quality criteria, often despite a questionable—or unexplained—sampling approach. The justifications were often mentioned in “Threats to Validity” or “Limitations” sections. It was common to mention the fact that studied artifacts were “real” (as in “real-world”), mentioned by 22 papers. P17, for instance, described the sample as containing “representative real-world ORM applications”, but did not go into details about the actual strategy followed to select those applications. Further popular adjectives were “representative” (13), “large” (13), and “diverse” (7).

#### 4.3 RQ3: Research methodologies

Predominately quantitative studies outnumber predominately qualitative studies 100 to 13. Of those 100 quantitative studies, 77 report experimental evaluations of software tools, models, or machine learning approaches. Eight studies involve mining software repositories; eight studies report on different kinds of user studies, and five are questionnaire-based surveys. Beyond that, the diversity of approaches defies organization into common methodological categories. For instance, one study involves comparing software metrics; another builds a taxonomy.

Table 5. Units of analysis in samples reported in SE research papers (n=179 sampling stages)

Unit	Examples	Frequency
Code artifacts	Projects, source code, defects, fixes, commits, code smells	115
Non-code artifacts	Bug reports, discussions, pull requests, effort estimates, feature models	27
People	Developers, maintainers, students, interview transcripts	26
Articles	Papers published in SE journals and conferences	2
Generated data	Generated automata, simulated concurrent accesses, generated Java-projects	5
No study	Formal descriptions, proofs	4

#### 4.4 RQ4: Units of analysis

We organized the primary studies' units of analysis into the categories shown in Table 5. Most of the studies investigate *code artifacts* including GitHub projects (e.g. P4, P17, P43), code commits (e.g. P16, P85), and packages (e.g. P56). Examples for *other artifacts* include bug reports (P16), faulty rules in model transformations (P29), and test logs (P99). Besides students (e.g. P69, P94, P96), the category *people* also includes GitHub users (P77), Microsoft developers (P87), and Amazon Mechanical Turk workers (P94).

#### 4.5 Discussion

Perhaps the most salient finding of this study is that **purposive and convenience sampling were the most commonly employed strategies** in both qualitative and quantitative studies. Only ten papers employed probability sampling and 9 out of the 15 corresponding sampling stages were sub-samples of non-probability samples. While this kind of study cannot determine why probability sampling is so rare, at least three factors contribute to this phenomenon:

- (1) Probability sampling is easier for some methodologies (e.g. questionnaires, repository mining) than others (e.g. user studies). However, non-probability sampling is popular even for questionnaires and quantitative data analysis studies.
- (2) Some SE research adopts interpretivism or other philosophical positions incommensurate with statistical generalization. Although we did not analyze philosophical positions (not least because most articles do not state one), a few studies may fall into this group.
- (3) There are no good sampling frames for most SE phenomena.

The significance of this third factor cannot be overstated. Without unbiased sampling frames, the randomness argument for representativeness falls apart. We can only claim that the sample represents the sampling frame, but we do not know how the sampling frame differs from the population.

Sampling frames are usually incomplete. Telephone sampling, for example, is incomplete because not everyone has a phone. However, *most* households have one or more phones and there are techniques to account for unlisted numbers [33]. However, there is nothing like a phone book of all the software developers, projects, products, companies, test suites, embedded systems, design diagrams, user stories, personas, code faults, or code comments in the world or even in a specific country, language or application domain. Instead, we study samples of GitHub projects (e.g. P10), Microsoft developers (e.g. P87) or Huawei test logs (e.g. P99). If we randomly select enough Microsoft developers, we might get a sample representative of all Microsoft developers—but this is obviously not representative of all developers in the world because even if there were such a thing as an average company, Microsoft would not be it.

The closest we can get to publicly available sampling frames of certain sub-populations of software developers is probably the list of registered Stack Overflow users provided by their official data dump or SOTorrent [4] and

the lists of GitHub projects and users provided by the GHTorrent dataset [23]. Those datasets, however, both come with their own challenges and limitations [3].

This raises two questions: how do we get better samples for our research and how should we evaluate sampling when reviewing? We will return to these questions in Section 5.

Meanwhile, there seems to be widespread confusion regarding sampling techniques and terminology. The frequency of articles not explaining where their samples came from is concerning. We cannot evaluate sampling bias without knowing the source of the sample. Beyond that, we see the following archetypal problems in the rhetoric around sampling.<sup>8</sup>

- Incorrectly using the term “random” as a synonym for arbitrary.
- Arguing that a convenience sample of software projects is representative because they are “real-world” projects.
- Assuming that a small random sample is representative because it is random.
- Assuming that a large random sample is representative despite being selected from a previously filtered (and thus biased) sampling frame.
- Implying that results should generalize to large populations without any claim to having a representative sample.
- Dismissing detailed case studies because they only investigate one organization.
- Implying that qualitative research is inferior to quantitative research because of the prevalence of non-probability sampling, as if all quantitative research used representative samples.

## 5 RECOMMENDATIONS

To address the issues outlined above, we present guidelines for researchers and reviewers, grounded in our own understanding of the methodological literature on sampling and the results of our literature review. Note that the purpose of the empirical study we discussed in Section 4 was to motivate the need for sampling guidelines, but the guidelines themselves cannot be traced back to individual observations from the study. We do, however, use the guidelines to assess our own approach (below).

### 5.1 Guidelines for researchers

To improve the conduct and reporting of sampling for any empirical study, we recommend:

- (1) **Clarify your philosophical position.** A treatise on twenty-first century epistemology is not necessary, but one sentence on the study’s perspective—positivist, falsificationist, interpretivist, constructivist, critical realism, etc.—would help. If the reader (or reviewer!) has to guess, they might guess wrong, and mis-evaluate your study.
- (2) **Explain the purpose of sampling.** Clearly state whether you are aiming for a representative sample, or have a different goal (e.g. finding interesting examples).
- (3) **Explain how your sample was selected.** For qualitative studies, the reader should be able to *recover* your reasoning about what to study [7]. For quantitative studies, the reader should be able to replicate your sampling approach. The size of the sample should be evident.
- (4) **Make sure your sampling strategy matches your goal, epistemology, and type of study.** For example, a positivist questionnaire might use respondent-driven sampling; a pilot laboratory experiment might use a convenience sample, and an interpretivist case study might employ purposive sampling.

<sup>8</sup>Here, again, we avoid direct quotations because we do not wish to cast aspersions on individual researchers. While it is perhaps less credible this way, our credibility is not worth our colleagues’ animosity. Anyone who reviews much software engineering research will likely have seen at least some of these attitudes.

- (5) **Avoid defensiveness.** Very few software engineering studies have a strong claim to representative sampling. Overselling the representativeness of your sample is unnecessary and unscientific. Do not misrepresent ad hoc sampling as random; do not pretend small samples are automatically representative because they are random or because they were purposefully selected, and do not ignore the potential differences between sampling frames and populations. Do not pretend a sample is representative because it is “real” (e.g. professionals instead of students, real projects instead of toy examples). Do not admit to sampling bias in your limitations section only to pretend the results are near-universal in your conclusion.

Moreover, if representativeness is the goal of the study, we further recommend:

- (1) **State the theoretical population;** that is, in principle, who or what you would like to generalize to (e.g. professional software developers in Brazil, code faults in cyberphysical systems).
- (2) **Present your cone of sampling** (see Section 2.5). Don’t worry about the population vs. the sampling frame vs. the inclusion/exclusion criteria. Just give a replicable, concise, algorithmic account of how another person can generate the same sample. If the sampling strategy has many phases, consider a diagram like Figure 1.
- (3) Give an **explicit argument for representativeness** (cf. Table 1). Admit the generalizability threats implied by this argument.
- (4) **Clearly explain how the sample could be biased.** For complicated sampling strategies, discuss bias for each step in your cone of sampling. This could be presented in the sampling section or with limitations.
- (5) **Publish your sample** as part of a replication package **if and only if** it does not contain sensitive or protected information. Be very careful of the potential for de-identified data to be re-identified in the future.

*5.1.1 Did we follow our own guidelines?* We clearly stated our philosophical position (positivism, at least for the literature review; the primer is more pragmatic), our purpose (representativeness with some caveats), and our sampling strategy including a cone of sampling (see Section 3). Between our description and supplementary materials, an independent research could replicate exactly what we did. We explained why stratified random sampling was appropriate under these circumstances, and described at length how our approach could bias the sample. Clearly enumerate the study’s limitations (in Sections Section 3 and 6). We clearly explained that our theoretical population is sample of leading, recent SE research, and that our argument to representativeness is randomness. A complete list of the papers in our sample is available as part of the supplementary material [2].

We studied four venues, which is similar to having four sampling frames. We did not use bootstrapping (because our sample is not large enough), techniques for hidden populations (because published research is not a hidden population), or sample coverage (because we were not using heterogeneity sampling).

## 5.2 Mitigating sampling bias in different kinds of studies

Several other techniques can help improving sampling in certain situations. For large samples, we can use bootstrapping to assess stability. For instance, if we have a convenience sample of 10,000 Java classes, we can randomly exclude 1,000 classes and check whether it perturbs the results. If a sampling frame is biased, consider replicating the study using different sampling frames. For example, if we find a pattern in a sample of GitHub projects, replicate the study on a sample of Bitbucket projects. The more diverse the repositories, the more likely the results generalize. This could be a full replication or a limited sanity check with a small sample from a different domain. Developers (or developers with certain experiences) can be treated as hidden populations. If respondent-driven sampling (Section 2.1.4) can reach injecting-drug users and sex workers [36], surely it can help reach software developers. Finally, for studies of software projects, consider using sample coverage; that is “the percentage of projects in a population that are similar to the given sample” [40], to support heterogeneity sampling.

Moreover, many practices can reduce sampling bias and response bias in **questionnaire surveys** [12]. These include starting with important questions that resonate with participants (not demographics), avoiding matrix questions, avoiding mandatory questions, and sending reminders. Offering incentives (e.g. cash, prizes) is also effective. Some research suggests that offering charitable donations does not increase response rates [53]; however, none of this research was done with software developers, and donating to an open source project might be more effective than small cash incentives for motivating open source contributors to complete a questionnaire. There are also myriad techniques for *assessing* response bias [45]. Many questionnaire studies in SE use none of these techniques.<sup>9</sup>

Similarly, sampling bias and publication bias in **systematic reviews** can be addressed by (i) forward and backward snowballing on references, (ii) searching multiple databases, (iii) searching pre-print servers and dissertations, and (iv) requesting unpublished work in an area through relevant mailing lists, and (v) checking websites of prolific researchers in the area.

Evaluating sampling bias in **software repository mining** is fraught. Each repository is likely biased in unpredictable, non-obvious ways. Therefore, we should not assume that random samples from one repository are representative of other repositories or software in general. Purposive or heterogeneity sampling may outperform random sampling in repository mining. Moreover, comparing samples from multiple repositories may help improve representativeness. To assess representativeness, we need research comparing projects stored in public and private repositories, but this is intrinsically difficult. The private code that companies are willing to share might systematically differ from the code they will not share [42].

In the long term, SE research needs better sampling frames. One way to achieve this is to develop curated corpora like the Qualitas corpus, “a large curated collection of open source Java systems” [50]. Similar corpora could be developed for many kinds of code and non-code artifacts used in software projects, including design specifications, requirements specifications, diverse models and diagrams, user stories, scenarios, personas, test cases, closed-source Java systems, systems in all the other common languages, unit tests, and end-user documentation. Creating any one of these corpora is a major undertaking and should be recognized as a significant research contribution in itself. Even without good demographic information, the representativeness of a curated corpus can be improved in numerous ways:

- (1) Including artifacts from diverse domains (e.g. aerospace, finance, personal computing, robotics).
- (2) Including artifacts from diverse software (e.g. embedded systems, enterprise systems, console video games).
- (3) Making the corpus large enough to support heterogeneity sampling and bootstrapping.
- (4) Attempting to match the parameters we can discern; for example, we could attempt to include artifacts from different countries according to the size of each country’s software industry.

Corpora improve reproducibility because one corpus can support many studies. Furthermore, building corpora helps to separate the difficult task of creating and validating a good sampling frame from any particular study of the items in the corpora. This makes research more manageable.

### 5.3 Guidelines for reviewers

From our experience, many reviewers struggle to evaluate sampling. Our advice is to **evaluate sampling in the context of a study’s philosophy, methodology, goals and practical realities**.

For anti-positivist studies, it is sufficient for researchers to justify site selection and explain their data collection. Complaining about low external validity in a case study is typically unreasonable because that is not what a case study is for [48].

When reviewing a positivist study that does not aim for generalization (e.g. a laboratory experiment with human participants) only worry about high-level external validity threats (e.g. using student participants instead

<sup>9</sup>Again, we do not cite specific examples here for fear that it would do more harm than good.

of professionals [15, 17]). Complaining about low external validity in a laboratory experiment is typically unreasonable because that is not what a lab study is for [48].

However, when reviewing a study that does aim for generalization (e.g. a questionnaire study) insist on reasonable attempts to mitigate sampling bias. The whole point of a large questionnaire survey is to sacrifice internal validity for external validity. If external validity is the main priority of a study, it should have a defensible claim that its sample is representative.

For example, suppose we are evaluating a questionnaire study of 3D animators at AAA game companies. The authors recruited animators by posting ads on Facebook, which is basically convenience sampling—end of sampling discussion. This should be rejected not because it uses convenience sampling but because appropriate, practical steps for mitigating sampling bias were not taken. Authors should have used respondent-driven sampling, or found a list of AAA game companies and used stratified random sampling, or advertised on multiple social networks and compared them. They should have reported response rates or bounce rates, compared early responders to late responders and so on.

In contrast, suppose we are evaluating a constructivist grounded theory study of agile practices at a Norwegian software company. We could say “this study has low external validity because we cannot generalize from an  $n$  of 1.” This is simultaneously true and inappropriate. External validity is not an appropriate quality criterion for this kind of study [6] and statistical generalizing is not its aim. Instead, we should be asking why this site was selected, how the researchers went about theoretical sampling, and to what extent the resulting theory seems transferable to other contexts.

Assessing sampling in software repository mining is difficult because it comes from SE, so we are creating the norms. What we can say with confidence is that reviewers should question the assumption that a sample is representative because it was randomly selected from a repository, unless evidence is provided that the software in the repository does not differ from software in general on the dimensions of interest.

The key is to evaluate studies against the norms for that particular kind of study. Pilot and proof of concept studies investigate something under ideal—not representative—conditions. For experiments with human participants, representative sampling is often prohibitively expensive. Most predominately, qualitative research does not seek to generalize to other contexts, so representative sampling is irrelevant, and disparaging “a sample of 1” is merely prejudice against qualitative research. For studies that do not aim for representativeness, reviewers should instead focus on over-generalization. Lab studies and pilot studies under ideal conditions do not show that something works in real life; qualitative field studies do not establish universality.

Reviewers should check whether the sampling strategy is commensurate with the study’s implications. Non-representative sampling should be followed by acknowledging that external validity is limited. Such acknowledgments should not be followed by a sneaky implication that the results are universal. Misusing the term “random” should not be tolerated.

Finally, reviewers should consider whether the sample is large enough. For studies of causal relationships, researchers should be using power analysis to describe desired sample sizes [9]. The natural size of a case study paper is one, in-depth case [34]. Like reporting multiple experiments in one article, multiple case studies should be considered exceptional. Reviewers should also consider local norms, for example average sample sizes [5].

The discussion of representativeness above foreshadows the difficulty of assessing a sampling strategy. The representativeness of a sample is often subjective, and representativeness is not always the goal of the sampling strategy. We suggest the following questions for guiding assessment of a sampling strategy:

- (1) Has the paper specified a philosophical stance?
- (2) Has the paper specified the goal of the sampling strategy (e.g. representativeness, convenience)?
- (3) Has the paper described the sample and sampling strategy sufficiently?
- (4) Is the sampling strategy consistent with the stated goal and philosophical position?

- (5) If the representativeness is the goal, what argument to representativeness is made, and is it reasonable given the type of study and practical constraints?
- (6) Is the sampling strategy reasonable given the context, constraints and maturity of the research?
- (7) Are the limitations of the sample acknowledged?
- (8) Does the sampling strategy match the paper's knowledge claims?

It is very important for a sampling strategy to support the specific knowledge claims of the paper. When an article makes claims about a population, based on sample, or makes generic claims of the form X causes Y, we can and should question the article's argument for representativeness. Much qualitative research, in contrast, seeks to understand one specific context with no attempt to generalize knowledge to other contexts. In such research, challenges to representativeness are far less important.

This is not to say that a paper should be rejected out of hand because some details or missing or the conclusion overreaches. Reviewers can often simply request clarifications or rewording. Even when representativeness is the goal and reasonable attempts to mitigate sampling bias have not been made, such attempts may be possible in a multi-phase review process.

The key phrases here are "reasonable" and "practical constraints." Any criticism of a paper's sampling approach should include suggesting specific, practical techniques to mitigate sampling bias. Complaining that a systematic review should have addressed sampling bias through reference snowballing is reasonable. In contrast, complaining that a study of unit tests should have used probability sampling when no reasonable sampling frame exists is naïve and unreasonable.

## 6 CONCLUSION

This paper makes five contributions:

- (1) An introduction to sampling with examples from SE research. This exposition is more grounded in SE research than previous discussions in reference disciplines, and more comprehensive than previous discussions within SE.
- (2) An analysis of the state of sampling methods in a stratified random sample of recent, SE research in leading venues. It shows that probability sampling is rare, and most probability samples are drawn from unknown or non-representative sampling frames.
- (3) A novel exploration of the arguments for representativeness.
- (4) A novel technique for presenting a sampling strategy as an inverted cone.
- (5) Guidelines for conducting, reporting and reviewing sampling.

A sample is representative of a population when the parameters we care about correspond within a reasonable error tolerance. A random sample can be assumed to be representative if it is sufficiently large and is drawn from an unbiased sampling frame. Few SE studies use random sampling. Of those that do, some samples are too small to assume representativeness and others are drawn from biased sampling frames. Researchers make various arguments for why their samples should be representative: the sample is large, includes diverse items, matches known population parameters, etc.

This creates a paradox: the lack of representative sampling is undermining SE research but rejecting a study over its non-representative sample is capricious because virtually none of the other studies have representative samples. We can escape this paradox by working towards more representative sampling in studies where generalizability is desired. For questionnaires especially, researchers should apply known techniques for mitigating and estimating sampling bias. We also need to develop more curated corpora of SE artifacts and better sampling frames for SE professionals.

Furthermore, we need reciprocal willingness of researchers to present their research more honestly and reviewers to stop capriciously rejecting work over unavoidable sampling bias. This means no more mislabeling

ad hoc sampling as representative, no more pretending small samples are automatically representative because they are random, and no more ignoring the potential differences between sampling frames and populations. It also means no more accepting convenience sampling for experiments while criticizing convenience sampling for case studies and interviews. No more encouraging snowball sampling for literature reviews [29] while rejecting it in questionnaires.

The contributions above should be considered in light of several limitations. We operationalized “recent high-quality software engineering research” as articles published in four top venues during the past five years. Our sample is therefore unlikely to represent the broader field. Studies that were published twice (e.g. a paper in FSE followed by an extended version in TOSEM) have a greater chance of being selected. Moreover, the analysis was hindered by widespread confusion regarding sampling techniques and research methodologies.

Additionally, some of the guidelines suggested in this paper are not directly supported by empirical evidence. The guidelines are meta-science, and like most meta-science, are somewhat polemical. It simply is not practical to conduct experiments to determine whether aligning a study’s sampling strategy with its goals, epistemology and methodology improves scientific outcomes. Rather, meta-science typically relies on the expert judgment of peer reviewers to evaluate face validity and credibility (i.e., the extent to which guidelines align with the wider body of scholarship around the meta-scientific issue).

In conclusion, we hope that this article’s sampling primer, empirical results and recommendations raise awareness of and provide at least some basis for improving sampling in SE research.

## REFERENCES

- [1] Jeffrey J Arnett. 2008. The neglected 95%: why American psychology needs to become less American. *American Psychologist* 63, 7 (2008), 602.
- [2] Baltes and Ralph. 2020. Sampling in Software Engineering Research – Supplementary Material. <https://doi.org/10.5281/zenodo.3666826>. (Feb. 2020). DOI: <http://dx.doi.org/10.5281/zenodo.3666826>
- [3] Sebastian Baltes and Stephan Diehl. 2016. Worse Than Spam: Issues In Sampling Software Developers. In *10th International Symposium on Empirical Software Engineering and Measurement (ESEM 2016)*, Marcela Genero, Andreas Jedlitschka, and Magne Jorgensen (Eds.). ACM, Ciudad Real, Spain, 52:1–52:6. DOI: <http://dx.doi.org/10.1145/2961111.2962628>
- [4] Sebastian Baltes, Lorik Dumani, Christoph Treude, and Stephan Diehl. 2018. SOTorrent: Reconstructing and Analyzing the Evolution Stack Overflow Posts. In *15th International Conference on Mining Software Repositories (MSR 2018)*, Andy Zaidman, Emily Hill, and Yasutaka Kamei (Eds.). ACM, Gothenburg, Sweden, 319–330.
- [5] Kelly Caine. 2016. Local Standards for Sample Size at CHI. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 981–992. DOI: <http://dx.doi.org/10.1145/2858036.2858498>
- [6] Kathy Charmaz. 2014. *Constructing grounded theory*. Sage, London.
- [7] Peter Checkland and Sue Holwell. 1998. Action Research: Its Nature and Validity. *Systemic Practice and Action Research* 11, 1 (1998), 9–21. DOI: <http://dx.doi.org/10.1023/A:1022908820784>
- [8] William G Cochran. 2007. *Sampling techniques*. John Wiley & Sons.
- [9] Jacob Cohen. 1988. *Statistical power analysis for the behavioral sciences*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.
- [10] Johnnie Daniel. 2011. *Sampling essentials: Practical guidelines for making sampling choices*. Sage Publications.
- [11] Rafael Maiani de Mello and Guilherme Horta Travassos. 2016. Surveys in software engineering: Identifying representative samples. In *Proceedings of the 10th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*. ACM, Ciudad Real, Spain, 55.
- [12] Don A Dillman, Jolene D Smyth, and Leah Melani Christian. 2014. *Internet, phone, mail, and mixed-mode surveys: the tailored design method* (4th ed.). John Wiley & Sons, Hoboken, NJ, USA.
- [13] Claire B Draucker, Donna S Martsof, Ratchneewan Ross, and Thomas B Rusk. 2007. Theoretical sampling and category development in grounded theory. *Qualitative health research* 17, 8 (2007), 1137–1148.
- [14] Brian Duignan. 2019. Postmodernism. In *Encyclopedia Britannica*. Encyclopedia Britannica, inc. <https://www.britannica.com/topic/postmodernism-philosophy>
- [15] Davide Falessi, Natalia Juristo, Claes Wohlin, Burak Turhan, Jürgen Münch, Andreas Jedlitschka, and Markku Oivo. 2018. Empirical software engineering experts on the use of students and professionals in experiments. *Empirical Software Engineering* 23, 1 (2018), 452–489.
- [16] Jean Faugier and Mary Sargeant. 1997. Sampling hard to reach populations. *Journal of advanced nursing* 26, 4 (1997), 790–797.

- [17] Robert Feldt, Thomas Zimmermann, Gunnar R Bergersen, Davide Falessi, Andreas Jedlitschka, Natalia Juristo, Jürgen Münch, Markku Oivo, Per Runeson, Martin Shepperd, and others. 2018. Four commentaries on the use of students and professionals in empirical software engineering experiments. *Empirical Software Engineering* 23, 6 (2018), 3801–3820.
- [18] P M Fitts. 1954. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology* 47, 6 (1954), 381–391.
- [19] E Foster. 2014. *Software Engineering: A Methodical Approach*. Apress, New York, USA.
- [20] Stephen J Gentles, Cathy Charles, Jenny Ploeg, K Ann McKibbin, and others. 2015. Sampling in qualitative research: Insights from an overview of the methods literature. *The qualitative report* 20, 11 (2015), 1772–1789.
- [21] Barney G Glaser and Anselm L Strauss. 2017. *Discovery of grounded theory: Strategies for qualitative research*. Routledge.
- [22] Sharad Goel and Matthew J. Salganik. 2010. Assessing respondent-driven sampling. *Proceedings of the National Academy of Sciences* 107, 15 (April 2010), 6743–6747. DOI : <http://dx.doi.org/10.1073/pnas.1000261107>
- [23] Georgios Gousios. 2013. The GHTorrent dataset and tool suite. In *10th International Working Conference on Mining Software Repositories (MSR 2013)*, Thomas Zimmermann, Massimiliano Di Penta, and Sunghun Kim (Eds.). IEEE, San Francisco, CA, USA, 233–236.
- [24] Egon G Guba and Yvonna S Lincoln. 1982. Epistemological and methodological bases of naturalistic inquiry. *ECTJ* 30, 4 (1982), 233–252.
- [25] Douglas D. Heckathorn. 1997. Respondent-Driven Sampling: A New Approach to the Study of Hidden Populations. *Social Problems* 44, 2 (1997), 174–199.
- [26] Joseph Henrich, Steven J Heine, and Ara Norenzayan. 2010. The weirdest people in the world? *Behavioral and brain sciences* 33, 2-3 (2010), 61–83.
- [27] Gary T Henry. 1990. *Practical sampling*. Vol. 21. Sage.
- [28] Lisa G Johnston and Keith Sabin. 2010. Sampling hard-to-reach populations with respondent driven sampling. *Methodological Innovations Online* 5, 2 (2010), 38–48. DOI : <http://dx.doi.org/10.4256/mio.2010.0017>
- [29] Barbara Kitchenham and Stuart Charters. 2007. *Guidelines for performing systematic literature reviews in software engineering*. Technical Report. Keele University and University of Durham.
- [30] Barbara Kitchenham and Shari Lawrence Pfleeger. 2002. Principles of survey research: part 5: populations and samples. *ACM SIGSOFT Software Engineering Notes* 27, 5 (2002), 17–20.
- [31] William Kruskal and Frederick Mosteller. 1979. Representative sampling, I: non-scientific literature. *International Statistical Review/Revue Internationale de Statistique* 47, 1 (Apr 1979), 13–24.
- [32] William Kruskal and Frederick Mosteller. 1979. Representative sampling, III: The current statistical literature. *International Statistical Review/Revue Internationale de Statistique* 47, 3 (Dec 1979), 245–265. DOI : <http://dx.doi.org/10.2307/1402647>
- [33] E Laird Landon Jr. and Sharon K Banks. 1977. Relative Efficiency and Bias of Plus-One Telephone Sampling. *Journal of Marketing Research* 14, 3 (Aug. 1977), 294. DOI : <http://dx.doi.org/10.2307/3150766>
- [34] Allen S Lee. 1989. A scientific methodology for MIS case studies. *MIS quarterly* 13, 1 (1989), 33–50.
- [35] Allen S Lee and Richard L Baskerville. 2003. Generalizing generalizability in information systems research. *Information systems research* 14, 3 (2003), 221–243.
- [36] Mohsen Malekinejad, Lisa Grazina Johnston, Carl Kendall, Ligia Regina Franco Sansigolo Kerr, Marina Raven Rifkin, and George W Rutherford. 2008. Using Respondent-Driven Sampling Methodology for HIV Biological and Behavioral Surveillance in International Settings: A Systematic Review. *AIDS and Behavior* 12, 1 (June 2008), 105–130. DOI : <http://dx.doi.org/10.1007/s10461-008-9421-1>
- [37] Matthew B Miles, A Michael Huberman, and Johnny Saldaña. 2014. *Qualitative data analysis: A methods sourcebook* (4th ed.). Sage, Thousand Oaks, California, USA.
- [38] Rahul Mohanani, Iftaah Salman, Burak Turhan, Pilar Rodriguez, and Paul Ralph. in press. Cognitive biases in software engineering: a systematic mapping study. *IEEE Transactions on Software Engineering* (in press), 24. DOI : <http://dx.doi.org/10.1109/TSE.2018.2877759>
- [39] Kevin J Mullinix, Thomas J Leeper, James N Druckman, and Jeremy Freese. 2015. The generalizability of survey experiments. *Journal of Experimental Political Science* 2, 2 (2015), 109–138.
- [40] Meiyappan Nagappan, Thomas Zimmermann, and Christian Bird. 2013. Diversity in software engineering research. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering*. ACM, 466–476.
- [41] Michael Quinn Patton. 2014. *Qualitative research & evaluation methods: Integrating theory and practice*. Sage publications.
- [42] J W Paulson, G Succi, and A Eberlein. 2004. An empirical study of open-source and closed-source software products. *IEEE Transaction on Software Engineering* 30, 4 (April 2004), 246–256. DOI : <http://dx.doi.org/10.1109/TSE.2004.1274044>
- [43] Kai Petersen, Robert Feldt, Shahid Mujtaba, and Michael Mattsson. 2008. Systematic Mapping Studies in Software Engineering. In *Proceedings of the Evaluation and Assessment in Software Engineering Conference (EASE)*. ACM, Bari, Italy, 68–77.
- [44] Paul Ralph. 2019. Toward methodological guidelines for process theories and taxonomies in software engineering. *IEEE Transactions on Software Engineering* 45, 7 (2019), 712–735.
- [45] Linda J Sax, Shannon K Gilmartin, and Alyssa N Bryant. 2003. Assessing response rates and nonresponse bias in web and paper surveys. *Research in higher education* 44, 4 (2003), 409–432.
- [46] F Shull, J Singer, and Dag Sjøberg. 2007. *Guide to Advanced Empirical Software Engineering*. Springer, London.

- [47] Dag Sjøberg, B Anda, E. Arisholm, T. Dyba, Magne Jørgensen, A Karahasanovic, E F Koren, and M Vokac. 2002. Conducting realistic experiments in software engineering. In *2002 International Symposium on Empirical Software Engineering*. IEEE, Nara, Japan, 17–26. DOI : <http://dx.doi.org/10.1109/ISESE.2002.1166921>
- [48] Klaas-Jan Stol and Brian Fitzgerald. 2018. The ABC of software engineering research. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 27, 3 (2018), 11.
- [49] Klaas-Jan Stol, Paul Ralph, and Brian Fitzgerald. 2016. Grounded Theory in Software Engineering Research: A Critical Review and Guidelines. In *Proceedings of the International Conference on Software Engineering*. IEEE, Austin, TX, USA, 120–131.
- [50] Ewan Tempero, Craig Anslow, Jens Dietrich, Ted Han, Jing Li, Markus Lumpe, Hayden Melton, and James Noble. 2010. The Qualitas Corpus: A Curated Collection of Java Code for Empirical Studies. In *Proceedings of the 17th Asia Pacific Software Engineering Conference*. IEEE, Sydney, Australia, 336–345. DOI : <http://dx.doi.org/10.1109/APSEC.2010.46>
- [51] Gary Thomas and Kevin Myers. 2015. *The anatomy of the case study*. Sage.
- [52] Steven K Thompson. 1990. Adaptive cluster sampling. *J. Amer. Statist. Assoc.* 85, 412 (1990), 1050–1059.
- [53] Vera Toepoel. 2012. Effects of incentives in surveys. In *Handbook of survey methodology for the social sciences*. Springer, 209–223.
- [54] William MK Trochim and James P Donnelly. 2001. *Research methods knowledge base*. Vol. 2. Atomic Dog Publishing, Cincinnati, OH, USA.
- [55] Jan E Trost. 1986. Statistically nonrepresentative stratified sampling: A sampling technique for qualitative studies. *Qualitative sociology* 9, 1 (1986), 54–57.
- [56] Philip Turk and John J Borkowski. 2005. A review of adaptive cluster sampling: 1990–2003. *Environmental and Ecological Statistics* 12, 1 (2005), 55–94.
- [57] Richard Valliant, Jill A Dever, and Frauke Kreuter. 2018. Designing Multistage Samples. In *Practical Tools for Designing and Weighting Survey Samples*. Springer, 209–264.
- [58] Loan R van Hoeven, Mart P Janssen, Kit CB Roes, and Hendrik Koffijberg. 2015. Aiming for a representative sample: Simulating random versus purposive strategies for hospital selection. *BMC medical research methodology* 15, 1 (2015), 90.
- [59] Max Van Manen. 2016. *Phenomenology of practice: Meaning-giving methods in phenomenological research and writing*. Routledge.
- [60] Bogdan Vasilescu, Daryl Posnett, Baishakhi Ray, Mark G.J. van den Brand, Alexander Serebrenik, Premkumar Devanbu, and Vladimir Filkov. 2015. Gender and Tenure Diversity in GitHub Teams. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*. ACM Press, Seoul, Republic of Korea, 3789–3798. DOI : <http://dx.doi.org/10.1145/2702123.2702549>
- [61] Robert K Yin. 2018. *Case study research: Design and methods* (6th ed.). Sage, Thousand Oaks, California, USA.