Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Industry Needs for Quantum Software Engineers

César França

franssa@ufrpe.br

University of Bari Aldo Moro

Maria Teresa Baldassarre

University of Bari Aldo Moro

Gabriel Sousa Santos

Universidade Federal Rural de Pernambuco

Ronnie de Souza Santos

University of Calgary

Research Article

Keywords:

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11543339/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Industry Needs for Quantum Software Engineers

César França[1,2*], Maria Teresa Baldassarre[1†], Gabriel R. Sousa Santos[2†], Ronnie de Souza Santos[3†]

[1*]Dipartimento di Informatica, Università degli Studi di Bari Aldo Moro, Piazza Umberto I , Bari, 70121, Bari, Italy.
[2]Departamento de Computação, Universidade Federal Rural de Pernambuco, R. Dom Manuel de Medeiros, Recife, 52171-900, Pernambuco, Brazil.
[3]Department of Electrical and Software Engineering, University of Calgary, 2500 University Drive NW, Calgary, T2N 1N4, Alberta, Canada.

*Corresponding author(s). E-mail(s): franssa@ufrpe.br;
Contributing authors: mariateresa.baldassarre@uniba.it; gabriel.rsousa@ufrpe.br; ronnie.desouzasantos@ucalgary.ca;
†These authors contributed equally to this work.

Abstract

Purpose: Quantum software engineering has become a professional activity in its own right, but what the job market expects from the professionals who perform it has not yet been examined specifically. This study characterizes the functional roles, classical hard skills and quantum-specific skills demanded from quantum software engineers.

Methods: We retrieved 13,882 job ads from twelve job portals and 78 company career pages between August 2025 and April 2026 and selected 412 ads from 21 countries. Roles and skills were coded manually through thematic synthesis into inductively built taxonomies, the informativeness of each ad was assessed with a quality scale, and the results were crossed with education, seniority, company profile and country.

Results: The market defines six functional roles that job titles do not reveal. For general and infrastructure engineering, the second most frequent role, quantum computing is the context rather than the object of the work. Half of the ads that mention quantum skills request control systems, error correction or compilation, and the classical profile is led by Python, C++ and Rust rather than web technologies. Higher degrees are associated with more quantum and fewer

classical requirements, early-career positions are scarce, and startups and large enterprises demand different layers of the stack.
Conclusion: Quantum software engineering is a family of profiles rather than a single occupation. Education should differentiate quantum literacy for classical engineers from specialized training on the intermediate layers of the stack, and research should give more attention to software for quantum computers.

Keywords: Quantum software engineering, Job ads, Skills, Workforce, Software engineering education, Empirical study

1 Introduction

The emergence of quantum technology as a commercially viable domain has generated a growing demand for a specialized workforce capable of bridging the gap between theoretical advances and industrial application. As quantum processors become accessible through cloud platforms and companies move from laboratory prototypes to products, the development of software for and with quantum computers has become a professional activity in its own right, which led to the proposal of *quantum software engineering* (QSE) as a discipline concerned with the systematic application of software engineering principles, methods, and processes to quantum software systems (Piattini et al. 2020; Akbar et al. 2024).

The growth of this field, however, depends on the availability of professionals able to work in it, and the pool of specialists capable of developing quantum software remains small and disproportionate to the projected growth of the quantum industry (Peterssen 2020; Piattini et al. 2021; Aoun et al. 2021). Under these conditions, knowing what the industry actually expects from quantum software professionals is a prerequisite for educators designing programs, for companies building teams, and for students and engineers planning a transition into the field.

Prior studies have characterized the quantum workforce through interviews and surveys with companies and experts (Fox et al. 2020; Hughes et al. 2022; Greinert et al. 2023; Hasanovic et al. 2022) and through a systematic mapping of the literature (Juárez-Ramírez et al. 2024), showing that classical skills remain central and that quantum-specific skills are not uniformly required across roles (Section 2.2).

However, a consistent limitation across these investigations is their treatment of the workforce demand at a broad level of granularity, encompassing the quantum industry as a whole or quantum computing in general terms (Fox et al. 2020; Hughes et al. 2022; Greinert et al. 2023; Hasanovic et al. 2022; Juárez-Ramírez et al. 2024). To the best of our knowledge, no prior study has focused specifically and systematically on the skills demanded by the job market for professionals working in QSE, a subdomain with distinctive characteristics that differentiate it from quantum hardware development, theoretical research, and classical software engineering alike. Moreover, the analysis of job ads, a well-established method for characterizing skill demands in software engineering (Matturro 2013; França and Mellet 2016; Daneva et al. 2017; Giray 2019;

Meesters et al. 2022; da Silva Lopes et al. 2025), has not yet been applied systematically to this domain.

In this article, we address this gap by characterizing the functional roles, classical hard skills and quantum-specific skills that job ads demand from quantum software engineers. We also examine how the skill demands vary across functional roles and across the educational level, seniority, organizational profile, and country associated with each ad.

We report on the results of a systematic analysis of 412 job vacancies collected from online platforms and filtered specifically for relevance to quantum software engineering roles. Between August 2025 and April 2026, we retrieved 13,882 ads from twelve job portals and from the career pages of 78 companies active in quantum computing, and selected those whose responsibilities and qualifications were consistent with software engineering practice in a quantum computing context. The selected ads, published by employers from 21 countries, were coded manually through an inductive and iterative procedure based on thematic synthesis (Cruzes and Dyba 2011), and their informativeness was assessed with a job posting quality scale developed for this study. Analyzing job ads as primary empirical evidence, rather than relying on surveys directed at companies or academic experts (Fox et al. 2020; Hughes et al. 2022; Hasanovic et al. 2022; Greinert et al. 2023) or on the secondary literature (Juárez-Ramírez et al. 2024), offers a complementary and arguably more direct window into what the market currently demands from quantum software professionals, since job postings represent the operational translation of organizational needs into explicit hiring criteria.

This article extends our earlier study of quantum software testing roles, presented at Q-SE 2026 (de Souza Santos et al. 2026), by broadening its scope from testing to quantum software engineering as a whole, with a corpus almost four times larger, inductively built taxonomies, control variables and a quality assessment of the ads (Section 5.2).

Our results show that *Quantum Software Engineer* is not a single occupation but a family of six functionally distinct roles that job titles do not reveal. For a large share of positions, quantum computing is the context rather than the object of the work. What distinguishes the quantum software engineer lies instead in the intermediate layers of the quantum computing stack (control systems, error correction and compilation) and in a classical profile led by Python, C++ and Rust. Formal qualification is associated with quantum depth and seniority with classical breadth, entry-level positions are scarce, and the type of employer shapes the demand.

The main contributions of this article are the following:

- An empirical characterization of the industrial demand for quantum software engineers, based on 412 job ads from 21 countries.
- Three taxonomies built inductively from the data, covering functional roles, classical hard skills, and quantum-specific skills organized by layers of the quantum computing stack, each accompanied by definitions, examples, and decision rules that allow their reuse in future studies.
- A job posting quality scale that distinguishes differences in what employers require from differences in how completely they report those requirements, which addresses a limitation shared by previous job ad studies.

- A discussion of the implications of these findings for educators, employers, candidates, and researchers.

2 Background

2.1 Analysis of job demands in software development

The analysis of job ads as a method for characterizing the skills demanded by the software industry has constituted a well-established line of empirical research over the past decade (Matturro 2013; França and Mellet 2016; Hiranrat and Harncharnchai 2018; Giray 2019; Jebreen and Nabot 2021; Shadrin et al. 2022; Chévez-Coronel et al. 2023). In this literature, the term *hard skills* refers to teachable, domain-specific, and generally measurable technical competencies, such as knowledge of programming languages, frameworks, databases, and development tools (Hiranrat and Harncharnchai 2018).

Studies in this line differ mainly in scale and in extraction technique. Hiranrat and Harncharnchai (2018) applied keyword-based text mining to 2,229 Thai ads, distinguishing explicit technical skills (tools and technologies) from implicit ones (conceptual and process-level competencies); Giray (2019) combined text mining with Latent Dirichlet Allocation (LDA) topic modeling to cluster 1,597 Turkish ads; Jebreen and Nabot (2021) manually ranked the skills requested in 178 Jordanian ads, finding work experience to be as decisive a criterion as technical competencies; Shadrin et al. (2022) analyzed 23,465 Russian software developer vacancies to derive 22 specialization profiles as a basis for curriculum design; and Chévez-Coronel et al. (2023) compared 9,048 Ecuadorian job offers across work modalities and job profiles. A limitation acknowledged by most of these authors is the reliance on a single country, few platforms and short collection windows, which restricts the generalizability of each snapshot.

Taken together, these studies reveal a convergence that appears to be robust across geographies and methodological approaches: JavaScript and Java consistently appear among the most demanded programming technologies in Thailand (Hiranrat and Harncharnchai 2018), Turkey (Giray 2019), and Ecuador (Chévez-Coronel et al. 2023), with SQL also prominent in Thailand and Turkey, a pattern consistent with global industry surveys (see Table S8 of the Supplementary Material).

Beyond skills, some of these studies report educational requirements (Jebreen and Nabot 2021) and, less often, seniority (Jebreen and Nabot 2021; Leça and Santos 2025; Niva and Markkula 2022), but they rarely disaggregate skill frequencies by career stage, and none of them analyzes salary data.

2.2 Workforce and job market in quantum computing

Fox et al. (2020) interviewed representatives of 21 U.S. companies engaged in quantum-related activities and identified five recurring technical job categories: engineer, experimental scientist, theorist, technician and application researcher. Classical skills such as coding (90% of companies), statistical data analysis (90%), laboratory experience (81%) and electronics knowledge (76%) were valued as much as, if not more than, quantum-specific knowledge (e.g., quantum algorithms, 62%), and industry favored

one- or two-semester quantum awareness courses added to existing engineering programs (33% of companies) over the full master's programs that universities were then planning. Hughes et al. (2022) extended this picture through a quantitative survey of 57 companies affiliated with the Quantum Economic Development Consortium (QED-C). They found that quantum-specific skills are not required for all positions, that most jobs combine quantum awareness with traditional technical competencies, and that skills form three largely independent clusters (hardware, software and business). They also found that, except for highly specialized roles, companies accept a range of degree levels, and that the consensus about which skills are important is substantially stronger than the consensus about which degrees are important.

Other studies approached the workforce from the perspective of curricula and competence frameworks. Greinert et al. (2023) conducted a Delphi-inspired survey with 188 European experts to inform the European Competence Framework for Quantum Technologies (ECFQT), organizing competences into a theoretical background, a practical background and application domains, with a pronounced concentration of demand in quantum computing. Hasanovic et al. (2022) surveyed 24 professionals to define the competences of quantum technicians at the community college level, showing that the workforce pipeline extends beyond advanced degree holders. Finally, Juárez-Ramírez et al. (2024) mapped 51 primary studies on the skills required for quantum computing and concluded that there is no consensus on the topics of quantum computing curricula, that workforce development efforts focus on graduate programs despite the demand for bachelor's-level positions, and that industry job descriptions frequently fail to specify quantum-specific requirements.

Taken together, these studies show that demand is growing and distributed across a heterogeneous set of roles, that classical technical skills remain foundational across most positions, and that quantum-specific knowledge is concentrated in a subset of highly specialized roles. They focus, however, on the quantum workforce in general terms or on quantum computing and hardware roles, with limited attention to quantum *software engineering* as a distinct professional domain. Our earlier study of quantum software testing roles (de Souza Santos et al. 2026) was a first step in that direction.

2.3 Quantum software engineering and its job market

QSE must contend with properties such as superposition, entanglement and the probabilistic nature of quantum computation, which make conventional software engineering techniques insufficient when applied directly to quantum programs and motivate new methods, tools and quality metrics across the whole life cycle (Piattini et al. 2020, 2021; Akbar et al. 2024).

The academic movement toward a formally constituted QSE discipline may be traced to a foundational moment in 2020, when a broad community of researchers and practitioners produced what became known as the Talavera Manifesto (Piattini et al. 2020). The manifesto established that QSE should be agnostic with respect to specific quantum programming languages and technologies, should embrace the coexistence of classical and quantum computing through hybrid system engineering,

and should address quality assurance, reuse, security, and governance as first-class concerns throughout the quantum software lifecycle (Piattini et al. 2020).

Despite four decades of quantum computing progress, the total global pool of specialists capable of developing quantum software remains surprisingly small and disproportionate to the projected growth of the quantum industry (Peterssen 2020). The gap is not merely quantitative but qualitative in nature: the quantum industry increasingly demands specialists with knowledge profiles that do not yet have consolidated educational pathways, combining expertise in quantum physics, mathematics, software engineering, and application-domain knowledge simultaneously. Piattini et al. (2021) analyzed international computing curricula and evidenced that quantum computing topics are either entirely absent or treated only marginally, with quantum cryptography being the sole exception in some curricula. In the same direction, Aoun et al. (2021) analyzed Stack Exchange posts and GitHub issue reports related to QSE, and found that the most prevalent and persistent challenges faced by quantum software developers include bridging the knowledge gap between quantum and classical computing paradigms, interpreting quantum program execution results, managing development environments, and finding adequate learning resources. These findings indicate that, at the practitioner level, the demand for learning support infrastructure is at least as pressing as the demand for technical tooling.

3 Methods

3.1 Overview

To characterize the demand for quantum software engineers, we analyzed publicly available job ads, taking the individual ad as the **unit of analysis**: each posting is treated as an independent observation of industry demand, regardless of the company or platform from which it was retrieved.

Unlike large-scale studies that rely on automated collection and natural language processing over thousands or millions of ads (Cerioli et al. 2020; Shadrin et al. 2022), we followed the small-scale tradition (Grandjean 2017; Wang et al. 2020; Jebreen and Nabot 2021; Meesters et al. 2022; Niva and Markkula 2022; da Silva Lopes et al. 2025), whose manual search, selection and qualitative extraction allow a more precise identification of skills in domains whose vocabulary is not yet standardized. Epistemologically, the study adopts a pragmatist stance: it combines the interpretive coding of qualitative data with descriptive and inferential statistics, choosing each technique for its usefulness in characterizing the demand rather than for its adherence to a single paradigm.

The research procedure was organized into three stages, described in the following subsections and illustrated in Fig. 1. Previous studies also offered us methodological recommendations, which we used to inform our design: selection based on explicit inclusion and exclusion criteria (Papoutsoglou et al. 2019); the vocabulary used to operationalize skills and the scope of what is extracted must be justified and applied consistently (Papoutsoglou et al. 2017; Vassilev et al. 2021). Some limitations of job ad data are discussed in Section 3.5. The following subsections presents the details of each step.

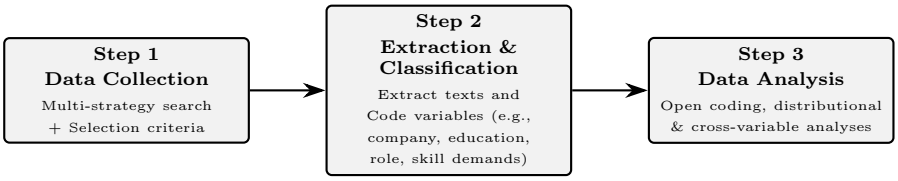


Fig. 1 Overview of the general research procedure encompassing data collection, extraction, and analysis.

3.2 Step 1: Data collection

Job ads were collected through two complementary search strategies. The first strategy consisted of keyword-based searches across general and specialized job portals. Twelve online portals were formally registered for this purpose: Indeed, Glassdoor, ZipRecruiter, Jooble, Flexjobs, JobLeads, Quantum Jobs US, Workday, QubitSok, QuantumJobs List, HK Jobs DB and CV-Library. Portal searches were conducted primarily targeting the English language, but automated translation was also used to access postings written in other languages.

Twelve keywords were defined to guide the job portal searches. These comprised one general term — *Quantum* — and eleven framework- and tool-specific terms drawn from prominent quantum software development platforms and languages: Qiskit, Cirq, Q#, QDK, QuTiP, PennyLane, OceanSDK, D-Wave, Quil, PyQuil, and OpenQASM. Framework-specific terms were included to capture postings that name quantum tools without using the word "quantum"; several of them, however, yielded zero results in most portals, and *Quantum* consistently produced the largest volume of results.

The second strategy consisted of direct searches on the career pages of companies active in the quantum computing industry. The list of companies was itself informed by the first strategy, as employers identified in the portal searches were progressively registered for direct searches, which extended the search beyond the most visible industry players. A total of 78 company websites were searched, ranging from large technology corporations with quantum computing programs (e.g., Google, IBM and Microsoft) and dedicated quantum hardware and software companies (e.g., IonQ, Quantinuum and Xanadu) to defense and financial organizations with quantum initiatives (e.g., Lockheed Martin and JPMorgan Chase & Co.). For company website searches, no additional keyword filter was applied when the company's career page allowed full browsing of all open positions; the quantum-related keywords were used only when the portal required a search term to return results.

Portal searches began in late August 2025 and company website searches ran from early September to December 2025, both conducted by two of the authors. Following a suggestion received when an earlier version of this work was presented (de Souza Santos et al. 2026), an additional round of searches targeting postings written in Chinese was carried out in April 2026. All 123 searches were recorded in a data collection log[1]

Retrieved job ads were subject to a selection process to retain only those relevant to the scope of this study, namely postings that target **software engineering**

[1]The data collection log, the coded dataset, the codebook and the analysis scripts will be made available on Zenodo upon acceptance of this article, to ensure traceability

related activitives in the context of quantum computing. Inclusion and exclusion criteria were applied manually to the retrieved texts. To determine relevance, a SWEBOK-based verification heuristic was systematically adopted (Bourque and Fairley 2004). The SWEBOK Knowledge Areas (software requirements, design, construction, testing, maintenance, configuration management, engineering management, engineering process, models and methods, quality, professional practice, economics, computing foundations, mathematical foundations, and engineering foundations) were used as a reference framework to assess whether the responsibilities and qualifications described in a posting were consistent with software engineering practice in a quantum computing context. Postings clearly unrelated to software engineering activities (e.g., experimental physics positions with no software component, hardware fabrication roles, data engineering and analytics positions whose object of work is data rather than software, or purely administrative positions) were excluded at this stage. Borderline cases, such as research positions in which code is a substantial but secondary output, were kept in the sample to avoid an overly restrictive interpretation of the field. Duplicate ads arise naturally from the aggregation nature of job portals and from the temporal overlap of search campaigns, but these were detected and removed on-the-fly during the selection process.

By the closure of the search process, a total of 13,882 job ads had been retrieved, 12,501 (90.05%) of them from job portals and 1,381 (9.95%) from direct company website searches. A total of 412 were selected for further analysis: 326 (2.61% success rate) from the portal searches and 86 (6.23% success rate) from the company website searches. The substantially higher selection rate for company website searches reflects the targeted nature of that strategy, which focused exclusively on organizations known to be active in quantum computing, thereby reducing irrelevant results from the outset.

The goal of this sampling was not statistical representativeness, since no sampling frame of all quantum software engineering vacancies exists, but the broadest possible coverage of the demand publicly advertised in the selected sources during the collection period. Every retrieved ad was screened, and every ad that met the inclusion criteria was retained, which characterizes an exhaustive purposive sample within the search results.

3.3 Step 2: Extraction and classification

The full text of each selected ad was copied into a Google Spreadsheet, which served as the instrument for extraction and classification. Since postings are usually removed once the position is filled, the analysis relied exclusively on the text captured at the time of extraction.

Because ads vary considerably in how much they describe, and a skill absent from a sparse ad may reflect a lack of detail rather than a lack of demand, we developed a job posting quality scale to distinguish what employers require from how completely they report it. The job posting quality scale comprises six dimensions in two weight tiers: four core dimensions (company description, role description, essential requirements, and preferred requirements), each weighted 2 points, and two supporting dimensions (compensation and benefits, and work arrangement, location and special requirements), each weighted 1 point. Each dimension was rated 0 (absent), 1 (superficial),

or 2 (detailed). The weighted points were summed to a raw score with a maximum of 20, normalized to a 0–100 scale using (raw score ÷ 20) × 100 and classified as low (0–40), medium (41–70), or high (71–100) quality.

Within the delimited content, two types of variables were extracted. The first type comprised **primary variables**, which could in most cases be identified directly from the text: company name, country of origin of the posting, required educational level (coded as *Bachelor's degree, Master's degree and PhD*), seniority level, hard skills and quantum skills. Seniority was operationalized using a classification scheme derived from the title and experiential language of each posting, distinguishing among the following levels: Intern or Student, Entry Level or Junior, Mid-Level, Senior, Staff or Principal, and Management. When any of these attributes was not explicitly stated in the ad, it was recorded as "Not Specified" rather than inferred.

The other type comprised **secondary variables** requiring some level of qualitative analysis to be constructed, as they could not be read directly from the text but instead emerged from a process of interpretation and categorization: role-type, clusters and categories of hard and quantum skills. Although raw information on these variables is present in the job description, it had to be submitted to a coding procedure, and subsequent grouping of terms and expressions identified in the postings, following the analytical procedures described in Section 3.4.

To avoid explicitly naming the organizations in the paper, results are reported by an organizational profile classification. Each employer was assigned to one of four categories based on the researchers' assessment of company size, age, and funding stage: Research / Government Institutions for universities, national laboratories, and public research bodies; Big Tech / Large Enterprises for large, multinational corporations in which quantum computing is one of many activities; Established Companies for specialized quantum firms beyond the early growth phase, as well as mid-sized or large organizations outside the technology sector; and Early-Stage Startups for recently founded firms with small teams.

3.4 Step 3: Data analysis

Primary variables were subjected to simple distributional analyses and cross-variable analyses (see Section 4.1). Secondary variables, in turn, were subjected to an open coding and categorization process, following the thematic synthesis framework proposed by Cruzes and Dyba (2011), which conceptualizes the synthesis of qualitative data as a five-step scientific inquiry: extracting data, coding data, translating codes into themes, creating a model of higher-order themes, and assessing the trustworthiness of the synthesis. In the context of this study, this framework was adapted to the task of building coherent, mutually exclusive, and exhaustive category systems.

The same integrated procedure (Cruzes and Dyba 2011) was applied to the three secondary variables, using the responsibilities and requirements sections delimited in Section 3.3: tentative labels were assigned while reading the first ads, tested against each incoming ad, and created or revised whenever a posting could not be accommodated without distorting an existing category. Each category was defined by its primary differentiating characteristics, so that categories were neither interchangeable nor redundant, and the process continued until saturation, that is, until no new ad

required a new category or a fundamental revision of an existing one. In all cases, whenever the information required to code a given variable could not be found in the raw textual content of the ad, the value was recorded as "Not Specified".

Coding was performed jointly by the two researchers who conducted the searches. They read and coded each ad together, discussing every classification, including borderline cases, until consensus was reached; there was therefore no independent double coding, a limitation discussed in Section 3.5. For the purposes of this study, a **functional role** is the primary organizational and technical function of the advertised position, deliberately separated from the job title, since titles vary considerably even for positions with substantively identical responsibilities (Daneva et al. 2017; Kassab et al. 2021). **Hard skills** are learnable technical competencies related to general software engineering and computing practice that are explicitly required or preferred in the ad, excluding those specific to quantum computing. **Quantum-specific skills** are competencies whose domain of application is particular to quantum computing, quantum information science or quantum hardware engineering, and which would not be expected in a conventional software engineering position, such as quantum SDKs (e.g., Qiskit, CUDA-Q), algorithms (e.g., variational quantum eigensolvers), hardware architectures (e.g., superconducting qubits, trapped ions) and enabling competencies such as pulse-level programming, circuit compilation and noise characterization. Given the volume and diversity of raw skill mentions, both skill variables were organized in a two-level scheme, consistent with the model of higher-order themes of Cruzes and Dyba (2011): *clusters* group terms that serve the same technical purpose, and *categories* group clusters into broader domains.

The coded variables were then analyzed through their distributions and cross-tabulations. *Content variables*, which characterize the demand itself (functional roles, hard skills and quantum-specific skills, with their clusters and categories), were crossed with *control variables*, which provide context (quality tier, company profile, country, education and seniority), and with one another, for example to examine co-occurrences between quantum-specific and hard skill categories. Associations were tested with chi-square, Kruskal–Wallis, Fisher's exact and Spearman tests, and all frequencies are reported with raw counts and percentages.

3.5 Threats to validity

Although Papoutsoglou et al. (2017), Papoutsoglou et al. (2019), and Vassilev et al. (2021) suggest that research of this kind is both methodologically feasible and practically valuable and we meticulously followed the guidelines established by prior small-scale empirical studies, its inherent limitations must be foregrounded.

External validity. The sample of 412 job ads, while adequate for qualitative thematic analysis, does not support statistical generalization. The observed frequencies should be interpreted as indicators of what the quantum software engineering job market chooses to signal through its public hiring documentation during the collection period (August 2025 to April 2026), in a market that changes quickly, rather than as a census of the skills practitioners are expected to possess. Job ad data also carry structural limitations. Not all vacancies are advertised online, since positions

filled through referrals, headhunting or informal networks are absent from any corpus (Vassilev et al. 2021); the searches concentrated on English-language portals, and internships and graduate programs are often advertised through university channels that we did not search. The same vacancy may also be posted on several platforms or reposted over time, inflating skill frequencies, a risk that our manual selection reduces but cannot eliminate (Vassilev et al. 2021).

Construct validity. Postings are written under editorial conventions that may not reflect the day-to-day requirements of the role, and they mix requirements with company descriptions and benefits (Papoutsoglou et al. 2019; Vassilev et al. 2021; Hiranrat and Harncharnchai 2018), which we mitigated by restricting the extraction to the responsibilities and requirements sections. The job posting quality scale was developed for this study and has not been independently validated, and non-English postings were assessed through automated translation, which may lose nuances of the original text.

Reliability and researcher bias. The manual extraction and coding are inherently subjective, since what appears to be a distinct technical skill in one context may indicate a broader competency in another. Because the two researchers coded every ad jointly rather than independently, inter-rater reliability could not be computed, and shared interpretations may have gone unchallenged. Both coders approached the data as software engineering researchers, which may have favored a software-centric reading of borderline roles, for instance research positions in which code is a secondary output, and the vocabulary of software engineering in the naming of categories. We attempted to bound these risks with the SWEBOK-based selection heuristic, explicit decision rules for borderline cases, which are reported in full in the Supplementary Material, and the use of a language model as a second opinion. Generative AI may itself introduce systematic biases in translation, extraction and classification, which is why all of its outputs were checked manually. Finally, the comparison of our results with interview and survey studies (Fox et al. 2020; Hughes et al. 2022) and with a mapping of the literature (Juárez-Ramírez et al. 2024) in Section 5 provides a partial triangulation across data sources: the convergence of these sources on the central role of classical skills and on the hardware/software divide increases our confidence in those findings.

Statistical conclusion validity. The associations between variables were tested with a large number of chi-square, Kruskal–Wallis and Fisher's exact tests without correction for multiple comparisons, and some groups are small (e.g., R6:QATV, low-quality ads and research and government institutions), so that the chi-square approximation is less reliable for them. These tests should therefore be read as exploratory support for the descriptive patterns, and isolated significant results, especially in small groups, should be interpreted with caution.

Ethical considerations. The study relied exclusively on publicly accessible job ads, which were collected manually through the regular interfaces of the portals and career pages, without automated scraping, and used only for non-commercial academic research. The analysis did not involve personal data: the coded variables describe positions and organizations rather than people, results are reported by organizational

profile, no ad is attributed to a named employer, and any personal information occasionally present in the postings, such as recruiters' names or contacts, will be removed from the shared materials. Because the copyright of job ads belongs to the employers and the terms of use of several portals restrict their republication, the full text of the ads will not be redistributed; only the data derived from them will be shared, along with their original URL addresses. This position is consistent with previous job ad studies, which relied on publicly available postings (Kassab et al. 2021; de Souza Santos et al. 2026) and, as Kapitsaki et al. (2024) did for the same reason, made their analyzed data, but not the raw text of the ads, publicly available.

Generative AI was used at several stages, always under the supervision of the researchers. Gemini Flash (version 3 or higher), used directly within Google Sheets, translated non-English postings, served as a double check of the researchers' relevance judgments during selection, supported the cleaning and extraction of the texts, and assisted in naming categories, classifying data and discussing ambiguous cases. The Python scripts that produced the data visualizations were generated with Anthropic Claude Opus 5.5 from specifications written by the researchers, and their outputs were verified against the underlying data. All prompts were shared among the team in research notes, and all outputs were manually checked.

4 Results

4.1 Overview of the search and selection results

Table 1 summarizes the corpus along the control variables defined in Section 3.4. Most ads were classified as *High Quality* (308; 74.8%), which indicates that the selection criteria retained mostly informative and well-structured postings. Established firms and large enterprises account for 334 ads (81.1%), so the demand captured by the corpus is predominantly that of mature organizations. The ads span 21 countries, but the distribution is strongly concentrated: the *U.S.* alone accounts for 43.9% of the corpus, and the five English-speaking countries (U.S., U.K., Canada, Australia and Ireland) for 280 ads (68.0%). Despite deliberately expanding the search to include job ads from China, no postings meeting the study's inclusion criteria were identified there. The Bachelor's degree is the most frequent educational requirement (48.1%), and the demand is concentrated on experienced professionals: early-career positions (*Intern/Student* and *Entry Level/Junior*) account for only 38 ads (9.2%).

The geographical distribution differs across company profiles. Large enterprises are concentrated in the U.S. (74 of 117; 63.2%), whereas early-stage startups are the most dispersed profile and the only one led by the U.K. (19 of 59; 32.2%, against 20.3% in the U.S.), and established companies dominate Canada (24 of 29; 82.8%) and Continental Europe & Ireland (43 of 61; 70.5%).

Company profile is also associated with educational requirements. Startups are the profile least explicit about education, with 49.2% of their ads (29 of 59) stating no requirement, compared with 17.1% in large enterprises and 16.6% in established companies. When startups do specify a degree, however, they are the most demanding: among their 30 ads with a stated requirement, 43.3% ask for a Master's and 16.7% for a PhD. Large enterprises, on the other hand, rely mostly on the Bachelor's degree (69 of

Table 1 Distribution of the control variables in the corpus ($n = 412$ ads).

Variable / value	n	%	Variable / value	n	%
Quality tier			*Country*		
High Quality	308	74.8	U.S.	181	43.9
Medium Quality	69	16.7	U.K.	50	12.1
Low Quality	35	8.5	Canada	29	7.0
Company profile			Germany	17	4.1
Established Companies	217	52.7	France	14	3.4
Big Tech / Large Enterprises	117	28.4	Australia	13	3.2
Early-Stage Startups	59	14.3	Other (15 countries)	54	13.1
Research / Government Inst.	12	2.9	Not determined	54	13.1
Not classified	7	1.7			
Education			*Seniority*		
Bachelor's degree	198	48.1	Intern / Student	19	4.6
Master's degree	83	20.1	Entry Level / Junior	19	4.6
PhD	41	10.0	Mid-Level	171	41.5
Not specified	90	21.8	Senior	123	29.9
			Staff / Principal	51	12.4
			Management	26	6.3
			Not identified	3	0.7

117; 59.0%) and request a PhD in only 7 ads (6.0%), while established companies ask for a PhD in 27 ads (12.4%). No significant association was found between education and seniority, indicating that higher degrees are not systematically required for more senior positions.

4.2 Analysis of roles

Each of the 412 ads was assigned to exactly one role category, corresponding to the dominant function described in the ad. As explained in Section 3.4, these categories were built from the cleaned responsibilities and requirements sections of the ads rather than from their job titles. Job titles proved to be a poor indicator of the actual function: the 182 ads whose title contains the expression *"Software Engineer"* were distributed across all six final categories, and even the 39 titles containing the more specific expression *"Quantum Software Engineer"* were spread across five categories.

For instance, Ad0272 (*Quantum Software Engineer – Qrisp*) asks the candidate to design high-level quantum programming abstractions and to implement algorithms such as QAOA, VQE and Grover in an open-source framework; Ad0350 (*Quantum Software Engineer – Control System*) is about building the control, monitoring and analysis software that operates the company's experimental hardware; Ad0386 (*Quantum Software Engineer – QEC*) asks the candidate to pursue research that advances QEC theory and to develop new QEC protocols for neutral-atom devices; and Ad0174 (*Quantum Software Engineer*) describes generic feature development, automated testing and CI/CD integration. Although all four share essentially the same title, they were assigned to four different categories.

The categories were therefore defined by the *primary object of work*, that is, what the professional is expected to produce or be accountable for: quantum algorithms and programming tools, classical software infrastructure, the software layer that drives

the physical device, new scientific knowledge, people and programs, or quality evidence. Three decision rules guided borderline cases: formal responsibility for people, budgets or programs prevails over technical work (e.g., Ad0352, *Quantum Error Correction Manager – Software*); ads whose main deliverable is new knowledge rather than production software are classified as research, even when the title refers to software engineering (e.g., Ad0085); and ads that target the verification rather than the construction of a quantum software artifact are classified as quality assurance (e.g., Ad0328, *Quantum Compiler Test Engineer*). This process resulted in the six categories summarized in Table 2; illustrative examples from the corpus are given in Table S1 of the Supplementary Material.

Table 2 Role categories derived from the job ads, with their definitions, typical activities and number of ads (n=412; % of the corpus).

Role	n (%)	Definition	Typical activities
Quantum Software & Algorithm Developer (R1:QSAD)	137 (33.3)	Professionals who design, implement and optimize quantum algorithms, quantum applications, and the programming tools used to express and compile them (SDKs, compilers, libraries, and frameworks), working at the logical, hardware-agnostic or hardware-aware level of the quantum software stack and delivering production-oriented code.	Developing quantum algorithms for application domains such as optimization, chemistry, or machine learning; implementing and benchmarking them on simulators and real hardware; building compilers, intermediate representations, and circuit optimization passes; and developing and maintaining quantum SDKs and open-source libraries.
Software Engineering – General & Infrastructure (R2:SEGI)	111 (26.9)	Professionals who build and operate the classical software that surrounds quantum computers and quantum products, such as cloud platforms, back-end and front-end services, DevOps and reliability engineering, HPC integration and security. In these roles, knowledge of quantum computing is contextual rather than central to the work.	Designing microservices and APIs for quantum cloud access; maintaining CI/CD pipelines, observability and monitoring; integrating quantum resources with HPC schedulers and third-party clouds; and developing internal tools and user interfaces.
Quantum Control & Hardware Interface (R3:QCHI)	87 (21.1)	Professionals who develop the software closest to the physical quantum device, namely control systems, calibration and characterization, pulse-level control, firmware and embedded software, and instrumentation. These roles bridge experimental physics and software engineering.	Implementing calibration and closed-loop optimization routines; developing control-stack and pulse-level software; writing firmware and drivers for control electronics; supporting bring-up, measurement and troubleshooting in the laboratory; and implementing error-suppression techniques.
Quantum Research & Applied Science (R4:QRAS)	42 (10.2)	Professionals whose main deliverable is new scientific or technical knowledge, such as new algorithms, error-correction schemes, simulation methods or architectural insights, typically validated through publications, patents or prototypes, with programming used as a research instrument.	Theoretical and numerical research on quantum algorithms and quantum error correction; resource estimation; modeling and simulation of quantum devices; running experiments on hardware; and publishing and presenting results.
Management & Leadership – Technical & Program (R5:MLTP)	26 (6.3)	Professionals primarily accountable for coordinating people, projects, programs or entire engineering organizations, with the technical domain acting as the context of management rather than as its object.	Leading and growing teams; defining technical roadmaps and architecture strategies; planning schedules, budgets and risks; and coordinating cross-functional delivery among hardware, software and research groups.
Quality Assurance / Testing & Validation (R6:QATV)	9 (2.2)	Professionals responsible for verifying that quantum software, compilers or systems behave correctly and perform as expected, by designing test strategies, frameworks and validation procedures.	Designing test plans and automated test suites; validating compiled circuits and system integration; analyzing measurement data to detect misbehavior; and reporting and triaging defects.

The six categories combine two dimensions. Three of them correspond to a layer of the technical stack: R1:QSAD covers the logical layer of quantum programs and algorithms, R3:QCHI the layer closest to the physical device, and R2:SEGI the classical

software that surrounds both. The other three correspond to a type of contribution that can occur at any layer: new knowledge (R4:QRAS), coordination of people and programs (R5:MLTP) and verification (R6:QATV). Overlaps between categories are therefore expected, because a researcher, a manager or a tester can work on compilers, control systems or cloud services alike.

R1:QSAD and R4:QRAS, for example, both deal with quantum algorithms, and job titles do not separate them: of the 24 titles containing the word *Scientist*, 14 were classified as R1:QSAD and 7 as R4:QRAS. However, in R1:QSAD the output is an algorithm, application or tool that will be delivered to a customer or integrated into a product (as in Ad0087 and Ad0088), while in R4:QRAS the output is new knowledge addressed to the research community or to the company's long-term R&D (as in Ad0002 and Ad0085). Likewise, the word *Lead* appears in 27 titles, but only 10 of them carry the formal responsibility for people, budgets or programs that defines R5:MLTP. Finally, some subjects cut across categories rather than defining one. Quantum error correction is the clearest example: of the 20 titles that mention it, 12 were classified as R4:QRAS, 3 as R1:QSAD, 3 as R3:QCHI, 1 as R2:SEGI and 1 as R5:MLTP, depending on whether the position develops new QEC protocols, implements decoders in a software product, runs decoding in real time on the control hardware, provides infrastructure for QEC experiments or manages a QEC team.

In quantitative terms (Table 2), R1:QSAD is the most frequent role (33.3%), followed by R2:SEGI (26.9%) and R3:QCHI (21.1%), while R6:QATV is marginal (2.2%); the three quantum-specific technical roles (R1:QSAD, R3:QCHI and R4:QRAS) together account for 266 ads (64.6%).

Seniority and education are both associated with the role (Table 3). Part of the seniority association is structural, since 23 of the 26 management-level ads belong to R5:MLTP. Among the technical roles, R1:QSAD and R4:QRAS concentrate on mid-level positions, while R3:QCHI and R2:SEGI lean towards senior and staff profiles. Early-career opportunities are scarce in every role: R1:QSAD and R2:SEGI offer the largest absolute numbers (15 and 13 of the 38 early-career ads), R5:MLTP offers none, and R4:QRAS offers internships but no entry-level positions. Education forms a gradient of formal qualification. R4:QRAS is the most demanding role, with a PhD required in 31.0% of its ads, which alone account for 31.7% of all PhD requirements in the corpus; R1:QSAD occupies an intermediate position and accounts for 45.8% of all Master's requirements; and R2:SEGI is the least demanding, with a Bachelor's degree in 64.0% of its ads and postgraduate degrees in only 12.6%. Overall, the more an ad is oriented towards producing new quantum knowledge, the higher the formal qualification required; the more it is oriented towards classical infrastructure, the lower.

Roles are weakly associated with company profiles. R1:QSAD is slightly underrepresented in established companies (62 of its 137 ads; 45.3%), with 44 (32.1%) in Big Tech / Large Enterprises, and it accounts for the largest share of both startup ads (22 of 59; 37.3%) and research and government ads (7 of 12; 58.3%).

R2:SEGI and R5:MLTP, in contrast, are concentrated in established companies (69 of 111, 62.2%; and 17 of 26, 65.4%, respectively). R2:SEGI is rare in early-stage startups (7; 6.3%) and absent from research and government institutions. In startups,

Table 3 Seniority and educational requirements by functional role (% of the ads of each role). Both distributions differ significantly across roles (chi-square test, $p < .001$). Three ads with unspecified seniority are not shown.

		Seniority				Education			
Role	**n**	**Early**	**Mid**	**Sen./Staff**	**Mgmt.**	**Bach.**	**Mast.**	**PhD**	**N.S.**
R1:QSAD	137	10.9	53.3	35.8	0.0	35.0	27.7	13.9	23.4
R2:SEGI	111	11.7	33.3	52.3	1.8	64.0	9.0	3.6	23.4
R3:QCHI	87	5.7	37.9	52.9	1.1	52.9	19.5	3.4	24.1
R4:QRAS	42	9.5	54.8	35.7	0.0	35.7	26.2	31.0	7.1
R5:MLTP	26	0.0	0.0	11.5	88.5	42.3	23.1	7.7	26.9
R6:QATV	9	11.1	55.6	33.3	0.0	77.8	11.1	0.0	11.1
All ads	412	9.2	41.5	42.2	6.3	48.1	20.1	10.0	21.8

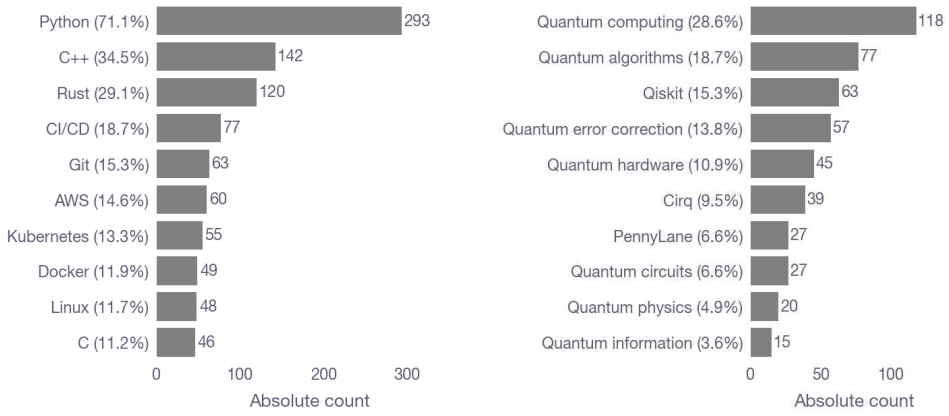
the most frequent role after R1:QSAD is R3:QCHI (18; 30.5%), and the three quantum-specific technical roles together account for 47 of the 59 startup ads (79.7%), which suggests that young companies hire mostly for the quantum-specific core of their technology and less for classical infrastructure or management. R3:QCHI, in turn, is less frequent in Big Tech / Large Enterprises (17 of its 87 ads; 19.5%).

No significant association was found between roles and geographical region. The U.S. is the leading location for every main role, ranging from 35.7% of R4:QRAS ads (15 of 42) to 65.4% of R5:MLTP ads (17 of 26).

4.3 Analysis of hard skills

Hard skills were extracted from the cleaned responsibilities and requirements sections of each ad, as delimited in Section 3.3, and recorded as the literal term used by the employer (e.g., *Kubernetes*, *linear algebra*, *Zynq UltraScale+*). Quantum-specific technologies, such as quantum SDKs, algorithms and hardware platforms, were deliberately excluded from this variable and are analyzed separately in Section 4.4. Of the 412 ads, 376 (91.3%) mention at least one hard skill, while 36 (8.7%) do not name any specific technical competence beyond quantum-specific knowledge and were recorded as *Not Specified*. The 376 ads yielded 2,667 hard skill mentions, an average of 7.1 per ad (median 6; maximum 38), corresponding to 644 distinct raw terms.

Fig. 2a shows the top 10 raw hard skills mentioned in the ads. This raw material was strongly heterogeneous in three respects, which shaped the construction of the category system. First, the same skill frequently appears under different spellings or labels (e.g., *Go* and *Golang*; *Kubernetes* and *K8s*; *CI/CD*, *CI pipelines* and *Continuous integration*). Second, the terms vary widely in their level of abstraction: some designate a body of knowledge (e.g., *distributed systems*, *linear algebra*), others a product or vendor platform (e.g., *AWS*, *Xilinx*), and others a specific version or device (e.g., *C++20*, *AMD ZCU216*). Third, the vocabulary extends well beyond what is usually reported in studies of the general software labor market (Hiranrat and Harncharnchai 2018; Giray 2019; Chévez-Coronel et al. 2023), reaching into domains such as programmable logic, high-performance computing, radio-frequency instrumentation and electromagnetic simulation.



(a) Hard skills (b) Quantum-specific skills

Fig. 2 Top 10 hard skills (a) and quantum-specific skills (b) (n=412)

For this reason, the terms were organized in the two-level hierarchy described in Section 3.4, after spelling and labeling variants of the same skill were merged into a single canonical term. Four decision rules resolved borderline cases: language- and vendor-agnostic knowledge was separated from the named tools that implement it (*concept versus implementation*, e.g., *compiler design* versus *LLVM*); the environments where software is hosted and executed were separated from the tools that build, deliver and monitor it (*platform versus delivery*, e.g., *Kubernetes* versus *Docker* and *Terraform*); languages were separated from the libraries built on them (*language versus library*, e.g., *Python* versus *NumPy*); and terms related to circuitry, programmable logic and firmware were kept in a dedicated category to preserve the visibility of a domain uncommon in the general software labor market (*hardware specificity*). The full rules are given in the Supplementary Material. The final scheme comprises 67 clusters organized into eight categories. Table 4 defines each category; illustrative examples and the clusters that compose each category are given in Tables S2 and S3 of the Supplementary Material.

Two features of this taxonomy deserve attention. First, the categories are not mutually exclusive at the level of the ad: although each term belongs to a single cluster and category, a typical ad spans several categories, reflecting the composite technical profile expected of these professionals. Second, the taxonomy itself already indicates the breadth of the domain. Five of the eight categories (*Programming Languages*, *Frameworks & Libraries*, *Infrastructure & Cloud*, *Dev Tools & DevOps* and *Data & Storage*) closely resemble those found in studies of the general software labor market (Hiranrat and Harncharnchai 2018; Giray 2019; Shadrin et al. 2022; Chévez-Coronel et al. 2023), whereas *Systems & Hardware*, together with clusters such as *Hardware Description*, *High Performance & Grid* and *Scientific & Mathematical*, capture competences that are rarely, if ever, reported in that literature and that reflect the specific position of quantum software between classical computing infrastructure and experimental physical devices.

Table 4 Hard skill categories derived from the job ads, with their definitions and number of ads that mention each category (n=412; % of the corpus).

Category	n (%)	Definition
Concepts & Fundamentals (H1:C&F)	163 (39.6)	The theoretical, mathematical and architectural foundations of computer science and engineering, that is, language-agnostic principles that remain valid regardless of the specific tool used.
Programming Languages (H2:PRG)	326 (79.1)	The formal languages used to give instructions to a computer, to describe hardware or to query data, including high-level software languages, low-level systems languages, hardware description languages and query languages.
Frameworks & Libraries (H3:F&L)	116 (28.2)	Pre-written code and specialized toolkits that run on top of a programming language and provide functionality for specific tasks, such as web development, machine learning, numerical computing, optimization or testing.
Infrastructure & Cloud (H4:I&C)	117 (28.4)	The environments and resources where software is hosted and executed, including cloud providers, virtualization, container orchestration, HPC and grid schedulers, and the security infrastructure of these environments.
Dev Tools & DevOps (H5:DEV)	215 (52.2)	The ecosystem of tools that supports the software development lifecycle, including version control, continuous integration and delivery, Infrastructure as Code, build and compilation systems, testing and debugging tools, and observability.
Data & Storage (H6:D&S)	49 (11.9)	Technologies for the management, persistence, serialization and movement of information, including relational and non-relational databases, data warehouses, streaming platforms, data pipelines and business intelligence tools.
Systems & Hardware (H7:S&H)	103 (25.0)	Low-level engineering and the physical components of computing, bridging software and circuitry, including programmable logic, chip design, embedded systems and firmware, operating systems, interconnects, instrument interfaces and specialized hardware.
Management & Process (H8:M&P)	19 (4.6)	The methodologies, standards and organizational tools used to manage people, time and projects across the software development lifecycle.

As expected, ads combine multiple hard skill categories, on average 3 (median 3). Only 71 (17.2%) are restricted to a single category and 58 (14.1%) span five or more, but none covers all eight categories at the same time. As shown in Table 4, H2:PRG is by far the most frequent category, present in 326 ads (79.1%). It is also the most frequent single profile (47 ads; 12.5%). Notice that H7:S&H, the category with the least correspondence in the general software labor market, appears in a quarter of the ads (25.0%), at a frequency close to that of cloud infrastructure and software frameworks.

The share of ads in each category for each value of the control variables is given in Table S4 of the Supplementary Material. *Quality tier* is, unsurprisingly, the variable most strongly associated with whether an ad mentions any hard skill at all: 25 of the 35 low-quality ads (71.4%) name none, against 10 of the 308 high-quality ads (3.2%). Among the ads that do mention hard skills, quality affects mainly the *amount* of detail (median of 6, 4 and 2 hard skills in high-, medium- and low-quality ads) rather than the *type* of competence requested, which supports analyzing the remaining crossings on the 376 ads that mention hard skills. *Company profile* shows a distinctive pattern for large enterprises, which most often request H1:C&F (60.2%, against 37.2% in established companies and 36.4% in startups) and least often name specific H2:PRG (75.7%, against 91.5% and 87.3%); established companies, in contrast, list concrete technology stacks, with the highest shares of H5:DEV (62.8%), H4:I&C (35.7%) and H6:D&S (17.1%). *Education* shows an inverse gradient: the median number of hard skills falls from 7 in ads requiring a Bachelor's degree to 6 for a Master's and 4 for a PhD, with H1:C&F (51.5%, 33.3% and 24.2%) and H4:I&C (37.6%, 20.5% and 9.1%) as the most sensitive categories. *Seniority*, in contrast, is not associated with the breadth of the requirements, but at the top of the technical career *Staff/Principal*

and *Management* positions are the least likely to name specific H2:PRG (73.5% and 61.5%, against more than 88% at the other levels) and the most likely to mention H4:I&C (38.8% and 50.0%), which suggests a shift from specific tools to architectural and platform-level responsibilities.

Functional role is the variable most strongly associated with the hard skill profile (Table 5). Six of the eight categories vary significantly across roles, with the strongest associations found for H4:I&C, H7:S&H and H2:PRG, and the number of hard skills per ad also differs markedly. The exceptions are H3:F&L and H5:DEV; the latter is requested by 41.7% to 66.7% of the ads in every role and can therefore be regarded as a transversal competence of the field. The role is also weakly related to the absence of hard skills: the *Not Specified* cases are concentrated in R1:QSAD (18 of 137 ads; 13.1%) and R4:QRAS (6 of 42; 14.3%), whereas every R5:MLTP and R6:QATV ad mentions at least one hard skill.

Table 5 Share of ads (%) mentioning each hard skill category, by functional role ($n = 376$ ads that mention at least one hard skill). n: number of ads; Med.: median number of hard skills per ad; Cat.: average number of hard skill categories per ad. Category abbreviations as in Table 4. $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$ in the column header (chi-square test for each category and Kruskal–Wallis test for the number of hard skills).

Role	n	Med.***	Cat.	H1: C&F**	H2: PRG***	H3: F&L	H4: I&C***	H5: DEV	H6: D&S***	H7: S&H***	H8: M&P***
All ads	376	6	2.9	43.4	86.7	30.9	31.1	57.2	13.0	27.4	5.1
R1:QSAD	119	4	2.5	49.6	89.9	26.9	13.4	52.9	4.2	16.8	0.8
R2:SEGI	103	9	3.7	54.4	84.5	38.8	68.0	64.1	25.2	23.3	11.7
R3:QCHI	83	6	3.1	25.3	97.6	31.3	16.9	62.7	14.5	57.8	2.4
R4:QRAS	36	3.5	2.2	41.7	77.8	27.8	8.3	41.7	11.1	8.3	0.0
R5:MLTP	26	4	2.5	34.6	57.7	15.4	50.0	50.0	7.7	23.1	15.4
R6:QATV	9	4	2.7	33.3	88.9	44.4	11.1	66.7	0.0	22.2	0.0

The roles differ mainly in the breadth and orientation of their classical profiles. R2:SEGI has the broadest profile (median of 9 hard skills): it alone accounts for 59.8% of all H4:I&C ads, and its most frequent terms after Python (*AWS*, *Kubernetes*, *Docker*, *GCP*, *Terraform*) are typical of cloud and platform engineering in any industry, which indicates that, for this role, the quantum domain is the context rather than the object of the work. R3:QCHI defines the hardware-proximate profile, with H2:PRG in 97.6% of its ads, mostly Python, C++, Rust and C, and H7:S&H in 57.8%, more than twice the rate of any other role. R1:QSAD and R4:QRAS have the narrowest classical profiles (medians of 4 and 3.5 hard skills), centered on programming and on mathematical and compiler-related foundations, with compiler toolchains (*LLVM*, *MLIR*) and machine learning frameworks (*PyTorch*, *JAX*) as their most characteristic technologies; for both, a substantial part of the technical demand is expressed through quantum-specific skills (Section 4.4). R5:MLTP is the role that least often names specific programming languages (57.7%), while half of its ads mention infrastructure and build platforms (*Bazel*, *Gerrit*); R6:QATV is too small to support firm conclusions.

4.4 Analysis of quantum-specific skills

Quantum-specific skills were recorded, like hard skills, as the literal terms used by the employer, but under the definition given in Section 3.4 (see Fig. 2b). Of the 412 ads,

356 (86.4%) mention at least one quantum-specific skill, and 56 (13.6%) were recorded as *Not Specified*. The 356 ads yielded 1,629 mentions, an average of 4.6 per ad (median 4). The richest profile, with 21 terms, does not come from a research or algorithm position but from a management one: Ad0384 (*VP of Quantum Software Engineering*) expects familiarity with everything from *qubit shuttling* and *hybrid quantum processors* to *lattice surgery*, *magic state distillation* and *QEC decoding*.

The vocabulary is markedly more fragmented than that of hard skills. The 1,629 mentions correspond to 743 distinct raw terms, more than the 644 terms obtained from the 2,667 hard skill mentions, which means that each quantum-specific term is repeated, on average, about half as often. This fragmentation has three main sources. First, as with hard skills, the same item appears under several spellings (*Qiskit* and *QisKit*; *CUDA-Q* and *CUDA Quantum*; *TKET*, *tket* and *pytket*). Second, the range of abstraction is even wider: the most frequent term, *quantum computing*, appears in 118 ads (28.6%), and in nine of them it is the only quantum-specific requirement, whereas at the other end employers name highly specialized techniques such as *union-find* decoders, *Rydberg-based entangling gates* or the *Quantum Schur Transform*. Between these extremes, the most frequent terms after *quantum computing* combine broad subjects and a few named frameworks: *quantum algorithms* (77 ads), *Qiskit* (63), *quantum error correction* (57), *quantum hardware* (45), *Cirq* (39) and *PennyLane* (27). The three SDKs tend to be requested together rather than as alternatives: of the 68 ads that mention at least one of them, 42 (61.8%) mention two or more and 19 mention all three, and every ad that mentions Cirq also mentions Qiskit, which is consistent with the recommendation of the Talavera Manifesto that QSE be agnostic with respect to specific programming languages and technologies (Piattini et al. 2020). Third, and more specific to this variable, the same word often refers to different parts of the technology. *Simulation* is the clearest case: it denotes a class of problems in Ad0126 (*quantum simulation* problems solved by algorithms for neutral-atom computers), a family of software tools in Ad0261 (*simulation platforms* used to run quantum circuits), and a hardware design activity in Ad0222 (*quantum hardware simulation* to accelerate the design cycles of fault-tolerant computers).

For these reasons, the categories were not organized around domains of practice, as was done for hard skills, but around the *layers of the quantum computing stack*, from the physical device to the application and the infrastructure through which it is accessed, with an additional layer for the theoretical foundations that underlie all of them. Broad terms were kept in dedicated *general* clusters within the corresponding layer (e.g., *quantum hardware* in *General Systems & Testbeds*), so that even vague requirements still indicate which part of the stack the employer has in mind. Four decision rules resolved terms at the boundary between two layers: named tools were classified by the function they perform rather than by their vendor or product family (*function over product name*, e.g., *Qiskit* as an SDK but *Qiskit Pulse* as pulse-level control); theoretical frameworks were separated from the problems they are used to solve, except when a formalism is requested as the language of error correction (*theory versus application*); the qubit technology itself was separated from the electronics, firmware and routines that drive, calibrate and read it out (*device versus operation*);

and terms were assigned according to whether they concern writing a program, transforming it for a target device, or accessing the resources on which it runs (*position in the execution pipeline*). The full rules, with examples, are given in the Supplementary Material. The final scheme comprises 62 clusters organized into eight categories. Table 6 defines each category, from the bottom to the top of the stack; illustrative examples and the clusters that compose each category are given in Tables S5 and S6 of the Supplementary Material.

Table 6 Quantum-specific skill categories derived from the job ads, ordered from the bottom to the top of the quantum computing stack, with their definitions and number of ads that mention each category (n=412; % of the corpus).

Category	n (%)	Definition
Fundamentals & Quantum Theory (Q1:FQT)	231 (56.1)	The theoretical, physical and mathematical foundations of quantum computing, including quantum mechanics and quantum information, atomic, molecular and optical (AMO) and condensed matter physics, the quantum circuit model (circuits, gates and qubit operations), the numerical and mathematical frameworks used to model quantum systems, and broad field-level terms that signal domain knowledge without pointing to a specific layer.
Physical Hardware & QPU Architectures (Q2:PHW)	170 (41.3)	The physical realization of quantum processing units, covering the qubit modalities (superconducting, trapped-ion, neutral-atom, photonic, spin-based, topological, as well as annealing and analog systems), the design and fabrication of quantum devices, including chip-design (EDA) tools, and system-level architectures such as gate-based, measurement-based and fusion-based computation.
Control Systems & Electronics (Q3:CSE)	79 (19.2)	The classical hardware, firmware and low-level software that bridge the classical and quantum layers, used to drive, calibrate and read out qubits, including cryogenic engineering, RF and microwave electronics, real-time control, pulse-level programming, calibration and characterization routines, control stacks and quantum operating systems, and laboratory automation.
Error Correction & Noise Mitigation (Q4:ECN)	88 (21.4)	The strategies used to identify, characterize, correct and mitigate errors in quantum computers, spanning both the fault-tolerant and the near-term (NISQ) paradigms: error-correcting codes and fault-tolerant schemes, decoding and syndrome extraction, logical-qubit operations, error mitigation and suppression, noise modeling, benchmarking and characterization protocols, and the software tools that support them.
Compilation & Circuit Optimization (Q5:CCO)	70 (17.0)	The pipeline that translates a high-level quantum program into gates or pulses executable on a specific device, and the optimization required for efficient execution, including compiler engineering, transpilation and synthesis, gate decomposition, qubit mapping and routing, intermediate representations, scheduling and resource estimation.
Quantum Software & SDKs (Q6:SDK)	136 (33.0)	The high-level programming environments, languages and libraries used to write, simulate and manage quantum programs, including mainstream SDKs, quantum programming languages, simulators and emulators, domain-specific libraries, vendor-specific tools and GPU-accelerated platforms.
Algorithms & Applications (Q7:ALG)	130 (31.6)	The algorithms designed to run on quantum computers and their application to concrete problems, that is, the user layer of the stack, including algorithmic paradigms, quantum simulation and chemistry, quantum machine learning, optimization and finance, cryptography and quantum security, quantum communication and networking, and sensing and metrology.
Cloud & Hybrid Infrastructure (Q8:CLD)	43 (10.4)	The delivery, access and orchestration of quantum computing resources, including cloud platforms for remote access to hardware, hybrid HPC–quantum systems, quantum data centers, distributed quantum computing, workflow orchestration and run times, and vendor-specific access platforms.

Two aspects of this taxonomy are worth noting. First, as with hard skills, the categories describe terms rather than ads: a single ad typically spans several layers of the stack. Ad0020 is a representative case, combining the physics of superconducting circuits (*Fundamentals*), simulation packages (*Quantum Software & SDKs*) and the *characterization and calibration of quantum processors* (*Control Systems & Electronics*). Second, the taxonomy refines the hardware/software division reported in previous workforce studies (Hughes et al. 2022) and the theory- and hardware-oriented topics identified in the curricular literature (Juárez-Ramírez et al. 2024). Between the device and the algorithm, three intermediate layers (*Control Systems & Electronics*, *Error Correction & Noise Mitigation* and *Compilation & Circuit Optimization*) concentrate

competences that are neither pure physics nor conventional application development, and it is in these layers that the work of the quantum software engineer appears most distinctive.

The 356 ads that mention at least one quantum-specific skill span, on average, 2.7 categories (median 2, and up to seven): 65 (18.3%) are restricted to a single category, and only 24 (6.7%) span five or more. As shown in Table 6, the two lowest layers are the most frequent, Q1:FQT (56.1% of the corpus) and Q2:PHW (41.3%), the three intermediate layers appear in 17.0% to 21.4% of the ads, and Q8:CLD is the least frequent (10.4%).

The prominence of the two lowest layers should, however, be read in light of the general clusters described above. In 101 of the 231 Q1:FQT ads (43.7%), the only foundational term is a broad one, such as *quantum computing* or *quantum physics*, and in 51 of the 170 Q2:PHW ads (30.0%) the only hardware term is a generic reference such as *quantum hardware*. Q1:FQT alone is also the most frequent single-category profile (29 ads), and 32 ads (9.0% of those that mention quantum skills) express their quantum demand exclusively through general terms. At the same time, half of the ads that mention quantum skills (179; 50.3%) request at least one of the three intermediate layers, which confirms that the competences that most clearly distinguish the quantum software engineer are a substantial, rather than marginal, part of the demand.

The co-occurrence of categories reveals two broad profiles. On the one hand, Q2:PHW and Q3:CSE tend to appear together (Q2:PHW + Q3:CSE is the fourth most frequent profile, with 15 ads). On the other hand, Q5:CCO, Q6:SDK and Q7:ALG form a software-oriented block (Q6:SDK–Q7:ALG; Q5:CCO–Q6:SDK), to which Q4:ECN is linked mainly through compilation. The two profiles are mutually exclusive to some extent: Q3:CSE is negatively associated with Q7:ALG and with Q6:SDK, reproducing within the quantum layer the hardware/software divide reported in previous workforce studies (Hughes et al. 2022).

The crossings with the control variables, computed on the 356 ads that mention at least one quantum-specific skill, are given in Table S7 of the Supplementary Material. *Quality tier* affects the depth rather than the presence of quantum requirements: only 7 of the 35 low-quality ads (20.0%) omit quantum-specific skills, but only 14.3% of those that mention them reach an intermediate layer, against 52.9% of the high-quality ads, so the intermediate layers are the categories that only well-described ads reveal. *Company profile* is associated only with Q6:SDK and Q7:ALG, which large enterprises request more often (47.7% and 45.9%) than startups (25.5% and 23.5%); startups, in turn, most often request Q4:ECN (35.3%) and Q3:CSE (29.4%), and 68.6% of their ads mention an intermediate layer. *Education* is the control variable most clearly associated with the quantum-specific profile, in the opposite direction to hard skills: the median number of quantum skills rises from 4 (Bachelor's) and 5 (Master's) to 7 (PhD), with Q7:ALG (29.3%, 45.5% and 63.2%) and Q6:SDK (35.9%, 41.6% and 57.9%) as the most sensitive categories, although part of this association reflects the role composition of each level, since 31 of the 38 PhD ads belong to R4:QRAS or R1:QSAD. *Seniority*, finally, affects mainly the amount of quantum requirements, which peaks at the mid level, and management positions show the highest share of

Q1:FQT (77.3%), which suggests that managers are expected to understand the field broadly rather than master any particular layer.

The quantum-specific categories were then crossed with the hard skill categories, in order to examine whether the layer of the quantum stack in which an ad is positioned is associated with a particular classical technical profile. Of the 412 ads, 323 (78.4%) mention both types of skills, 33 (8.0%) mention only quantum-specific skills, 53 (12.9%) only hard skills and 3 (0.7%) neither. Table 7 reports, for the 323 ads that mention both, the share of ads in each hard skill category among those that request each quantum-specific category.

Table 7 Share of ads (%) mentioning each hard skill category, among the ads that mention each quantum-specific skill category ($n = 323$ ads that mention both types of skills). n: number of ads; Med.: median number of hard skills per ad; Cat.: average number of hard skill categories per ad. Hard skill abbreviations as in Table 4; quantum-specific abbreviations as in Table 6. $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$ (two-sided Fisher's exact test comparing the ads that mention the quantum-specific category with those that do not).

Quantum category	n	Med.	Cat.	H1: C&F	H2: PRG	H3: F&L	H4: I&C	H5: DEV	H6: D&S	H7: S&H	H8: M&P
All ads	323	6	2.9	43.3	87.9	30.0	28.5	57.3	11.8	28.5	4.3
Q1:FQT	203	6	2.9	44.8	86.2	32.0	28.6	56.2	12.8	23.6*	5.4
Q2:PHW	144	5	2.9	43.1	89.6	29.9	22.2*	52.8	12.5	31.9	4.2
Q3:CSE	77	5	2.9	31.2*	89.6	28.6	20.8	53.2	13.0	46.8***	2.6
Q4:ECN	81	4	2.4	38.3	85.2	19.8*	12.3***	45.7*	9.9	27.2	0.0*
Q5:CCO	67	5	2.8	46.3	91.0	31.3	26.9	58.2	10.4	19.4	0.0
Q6:SDK	120	5	2.8	47.5	93.3*	36.7	23.3	50.0*	6.7*	22.5	2.5
Q7:ALG	108	5	2.8	52.8*	85.2	31.5	25.0	54.6	8.3	20.4*	2.8
Q8:CLD	40	9.5	3.9	47.5	90.0	52.5**	80.0***	70.0	20.0	30.0	2.5

The first result concerns the amount of requirements. The breadth of the two profiles is inversely related: the more quantum-specific skills an ad requests, the fewer hard skills it requests (Spearman's $\rho = -0.15$, $p = .006$; the same holds for the number of categories, $\rho = -0.15$, $p = .007$). The median number of hard skills falls from 8 in ads restricted to a single quantum-specific category to 6 in those with two and 5 in those with three or more. The 53 ads that mention hard skills but no quantum-specific skill, most of them R2:SEGI positions, show the broadest classical profiles (median of 8 hard skills), with higher shares of H4:I&C (47.2%) and H6:D&S (20.8%) than the ads that mention both types of skills (28.5% and 11.8%). Employers therefore seem to describe a position either mainly through its classical stack or mainly through its quantum content, and rarely in detail through both.

The second result concerns the type of requirements. Most associations between the two sets of categories are weak, and H2:PRG remains above 85% for every quantum-specific category, which confirms its role as the common core of the field. Two clear associations, however, connect the ends of the quantum stack with the classical categories that correspond to them. At the top, Q8:CLD is strongly associated with H4:I&C: 32 of the 40 ads in this category (80.0%) also request classical cloud technologies, against 28.5% overall, and they show the broadest hard skill profile of the table (median of 9.5 hard skills and 3.9 categories), with H3:F&L also over-represented (52.5%). At the bottom, Q3:CSE is strongly associated with H7:S&H: 36 of the 77 ads in this category (46.8%) also request classical systems and hardware competences, against 28.5% overall ($p < .001$), while H1:C&F is under-represented (31.2%). In both

cases, the quantum layer and the classical category describe the same technical object from two sides: the platforms that give access to quantum resources and the cloud infrastructure on which they run, and the control electronics and the embedded and programmable hardware on which they are implemented.

Finally, functional role is, again, the variable most strongly associated with the profile (Table 8). All categories except Q1:FQT vary significantly across roles, with the strongest associations found for Q6:SDK, Q3:CSE and Q5:CCO, and the number of quantum skills per ad also differs markedly. Q1:FQT, requested by 50.0% to 76.3% of the ads in every role, can therefore be regarded as the transversal quantum competence of the field, in the same way as H5:DEV for hard skills. The role is also strongly related to the absence of quantum skills: 33 of the 111 R2:SEGI ads (29.7%) do not mention any, and this single role accounts for 33 of the 56 *Not Specified* cases (58.9%), while the rate remains below 10% in R1:QSAD (9 of 137), R3:QCHI (5 of 87) and R4:QRAS (4 of 42).

Table 8 Share of ads (%) mentioning each quantum-specific skill category, by functional role ($n = 356$ ads that mention at least one quantum-specific skill). n: number of ads; Med.: median number of quantum-specific skills per ad; Cat.: average number of quantum-specific categories per ad. Category abbreviations as in Table 6. $^{*}p < .05$, $^{**}p < .01$, $^{***}p < .001$ in the column header (chi-square test for each category and Kruskal–Wallis test for the number of skills and categories).

Role	n	Med.***	Cat.***	Q1: FQT	Q2: PHW*	Q3: CSE***	Q4: ECN**	Q5: CCO***	Q6: SDK***	Q7: ALG***	Q8: CLD***
All ads	356	4	2.7	64.9	47.8	22.2	24.7	19.7	38.2	36.5	12.1
R1:QSAD	128	5	3.2	68.8	43.0	10.9	28.1	38.3	62.5	55.5	10.2
R2:SEGI	78	2	2.2	64.1	37.2	14.1	9.0	10.3	24.4	25.6	30.8
R3:QCHI	82	4	2.3	53.7	63.4	48.8	25.6	3.7	15.9	18.3	3.7
R4:QRAS	38	5.5	2.8	76.3	57.9	15.8	44.7	10.5	31.6	39.5	2.6
R5:MLTP	22	3	2.4	72.7	40.9	22.7	22.7	13.6	31.8	27.3	9.1
R6:QATV	8	5.5	2.9	50.0	37.5	37.5	25.0	37.5	62.5	37.5	0.0

R1:QSAD has the broadest quantum-specific profile (median of 5 quantum skills and 3.2 categories) and defines the software-oriented block identified above: it alone accounts for 70.0% of all compilation ads, 58.8% of the SDK ads and 54.6% of the algorithm ads, which complements its comparatively narrow classical profile. R3:QCHI defines the hardware-proximate block: it most often requests Q2:PHW (63.4%) and Q3:CSE (48.8%), accounts for half of all control-system ads (40 of 79), rarely mentions the upper layers of the stack, and is the role in which the intermediate layers are most present (64.6%), with *quantum error correction* as its second most frequent term, after *quantum computing*. R4:QRAS, which has the narrowest classical profile, shows one of the richest quantum-specific ones (median of 5.5 skills), with the highest shares of Q1:FQT (76.3%) and Q4:ECN (44.7%); none of its ads expresses its quantum demand only through general terms, which suggests that research positions are defined mainly by their quantum-specific scientific depth. R2:SEGI, in contrast, has the narrowest quantum-specific profile (median of 2 skills), with 37.2% of its ads restricted to a single category and only 21.8% mentioning any intermediate layer; its exception is Q8:CLD (30.8% of its ads, 55.8% of all ads in this category), which indicates that, when general software engineers are asked for quantum-specific knowledge, it is mostly familiarity with the field and with the platforms through which quantum resources

are accessed. R5:MLTP combines a high share of Q1:FQT (72.7%) with a moderate median of 3 quantum skills, and R6:QATV relies mainly on Q6:SDK and on the testing of compilers and control stacks (e.g., Ad0328), but its count is too small to support firm conclusions.

5 Discussion

5.1 Main findings

Read together, the analyses of roles, hard skills and quantum-specific skills point to five main findings:

F1. Quantum software engineering is a family of functionally distinct roles, and job titles do not reveal them.

As seen through the job market, *Quantum Software Engineer* is not a single occupation but a family of six roles with different deliverables, qualifications and skill profiles, and the hard and quantum-specific profiles vary more across these roles than across any control variable (Sections 4.3 and 4.4). Any aggregate statement about "what quantum software engineers need" therefore averages over profiles as different as a site reliability engineer for a quantum cloud and a researcher developing new error-correcting codes. Title–function mismatches have been reported for requirements engineers, software architects and AI engineers (Daneva et al. 2017; da Silva Lopes et al. 2025; Meesters et al. 2022), but they appear here in a more extreme form, since the same title covers SDK development, control software, error correction research and generic feature development (Section 4.2). Within software, the broad "engineer" and "software" categories of previous quantum workforce studies (Fox et al. 2020; Hughes et al. 2022) conceal at least three stack positions (R1:QSAD, R3:QCHI and R2:SEGI) that differ in qualification, seniority and skills. The two-dimensional structure of the classification (layer of the stack versus type of contribution) suggests that analyses of the QSE workforce, often framed around the life cycle of quantum programs (Piattini et al. 2020, 2021; Akbar et al. 2024), should also consider roles that work on other layers of the stack, and it offers a common reference for future studies, which currently use incompatible role lists.

F2. For a large share of positions, quantum computing is the context rather than the object of the work.

General and infrastructure software engineering (R2:SEGI) is the second most frequent role, its hard skill profile is typical of cloud engineering in any industry, and almost a third of its ads state no quantum-specific skill (Sections 4.3 and 4.4). This confirms, from what employers write, what earlier studies found through interviews and surveys (Fox et al. 2020; Hughes et al. 2022), and adds precision: the classical demand is concentrated in R2:SEGI, where quantum knowledge, when requested, consists mostly of familiarity with the field and with cloud and hybrid platforms, and some ads explicitly frame it as something the candidate will "learn along the way" (Ad0137, Ad0362). Our results only partially agree with the observation that industry job descriptions fail

to specify quantum-specific requirements (Juárez-Ramírez et al. 2024): the underspecification is real, but R2:SEGI accounts for more than half of the ads without quantum skills, so it is better understood as a property of positions in which quantum computing is the context of the work. The inverse relationship between the breadth of classical and quantum requirements points in the same direction, suggesting a division of labor in which few positions require depth in both, and in which the absence of a skill from an ad should not be read as the absence of a need. Theoretically, these results suggest that *quantum depth*, the extent to which quantum knowledge is the object rather than the context of the work, is a useful dimension for describing this workforce, which cuts across the hardware, software and business clusters proposed by Hughes et al. (2022).

F3. What distinguishes the quantum software engineer lies in the middle and lower layers of the stack.

Half of the ads that mention quantum skills request control systems, error correction or compilation, and more than a quarter of those that mention hard skills request systems and hardware competences such as FPGA programming, firmware and instrument-control interfaces (Sections 4.3 and 4.4). The contrast with the general software labor market is clear: whereas previous studies consistently report JavaScript, SQL, Java and web frameworks among the most demanded skills (Table S8 of the Supplementary Material), the classical profile of our corpus is led by Python (293 ads; 71.1%), C++ (142; 34.5%) and Rust (120; 29.1%), and resembles embedded, scientific and high-performance software engineering. The combination of software and experimental work that our earlier study found in quantum software testing (de Souza Santos et al. 2026) defines an entire role here (R3:QCHI), which echoes the difficulty of hiring engineers with analog electronics experience reported by Fox et al. (2020). Our data reproduce within the quantum layer the hardware/software divide reported by Hughes et al. (2022), but they also show that the boundary itself is the object of a role, for which specializing in either side is not an option. These results have implications for the QSE research agenda. Much of the existing work focuses on the design, testing and maintenance of quantum programs (Piattini et al. 2020; Dwivedi et al. 2024; Akbar et al. 2024; Stefano et al. 2022), whereas the market shows that a large share of the work is software *for* quantum computers: control stacks, calibration routines, compilers and decoders that must meet real-time and hardware constraints, closer to embedded and research software engineering than to application development. The quantum-specific taxonomy also refines the curricular categories identified by Juárez-Ramírez et al. (2024) by making explicit three intermediate layers in which software engineering is central. We speculate that, as hardware platforms approach fault tolerance, error correction will move from research to engineering, a shift that the presence of QEC terms in R3:QCHI and R1:QSAD ads suggests has already begun.

F4. Quantum depth follows formal qualification, not seniority, and the market offers few entry points.

As the required degree increases, the number of hard skills falls and the number of quantum skills rises, whereas seniority is not associated with education and mainly shifts requirements towards infrastructure and platform-level responsibilities

(Sections 4.3 and 4.4). A degree and experience thus seem to work as different currencies, signalling quantum depth and classical breadth, respectively. This is consistent with Hughes et al. (2022), who found that companies accept a range of degree levels except for highly specialized jobs, and with Fox et al. (2020), who reported that quantum information theorists are among the hardest profiles to hire; it also supports the argument that workforce development focuses on graduate programs despite the demand for Bachelor's-level professionals (Juárez-Ramírez et al. 2024), since more than twice as many ads ask for a Bachelor's degree as for a Master's. The most worrying result is the scarcity of early-career positions, fewer than one in ten ads. The field appears to depend on experienced professionals while being too young to have produced many of them, and plausibly fills its positions with experienced classical engineers who learn the quantum context on the job and with PhD holders who bring quantum depth but limited software engineering experience, a hypothesis consistent with the hiring difficulties reported by Fox et al. (2020). This scarcity should nevertheless be interpreted with caution, because internships are often advertised through university channels that our search did not cover.

F5. Who hires shapes what is asked.

Startups hire mostly for the quantum-specific core of their technology and most often request error correction, control systems and other intermediate layers, whereas large enterprises describe their needs through foundations rather than specific stacks and concentrate their quantum demand on SDKs and applications (Sections 4.2 and 4.4). One possible explanation lies in the economics of training: companies without the resources to train new staff expect candidates to be productive from the start (Fox et al. 2020), and startups are indeed the most demanding profile when they state a degree, while large enterprises can recruit on foundations, develop specific knowledge internally, and own the main SDKs and cloud platforms. Previous quantum workforce studies did not analyze organizational profiles (Fox et al. 2020; Hughes et al. 2022; Greinert et al. 2023), and this result shows that the type of employer is a relevant dimension of the demand. By contrast, the distribution of roles does not vary across regions, which suggests that the role structure reflects the technology itself rather than local labor markets. The absence of Chinese ads, despite a dedicated search, more likely reflects hiring channels that our search did not reach than an absence of demand.

5.2 Revisiting our earlier study

Our earlier study analyzed 110 postings with testing-related activities and concluded that quantum software testing combines classical quality assurance with experimental validation, that the techniques dominating academic research on quantum software testing were absent from the postings, and that the hard skills combined Python, C++, Rust and automation with control and calibration (de Souza Santos et al. 2026). Several of these findings hold in the present corpus, which is almost four times larger. Testing remains a responsibility embedded in other roles (dedicated positions represented 4.5% of the earlier sample and R6:QATV represents 2.2% of the present one), with verification tasks appearing in compiler and control-stack positions such

as Ad0328 and Ad0168; although Piattini et al. (2021) expected quantum computing to have a heavy influence on software testing and quality, testing has not yet become a career track of its own, as happened with quality assurance in classical software (Kassab et al. 2021). The research–practice gap also persists: no ad mentions metamorphic, mutation-based or property-based testing or test oracles, formal methods appear in only two, and industrial verification is expressed mainly through benchmarking (64 ads) and system-level validation. The broader corpus also qualifies the earlier conclusions. Selecting postings through testing activities favored hardware-proximate positions, whereas R3:QCHI is only the third most frequent role in the full corpus, and the share of U.S. ads fell from 75.5% to 43.9%, which shows how sensitive the profile of job-ad studies is to the search strategy. We speculate that quantum software engineering, today strongly anchored at the boundary with the physical device, will move up the stack as the hardware stabilizes, with a growing share of compiler, platform and application work and of the verification techniques proposed by research, a hypothesis that periodic replications with the same instruments could test.

6 Conclusion

This article characterized the industrial demand for quantum software engineers through the analysis of 412 job ads from 21 countries. The market does not define a single occupation but a family of six functional roles, organized along two dimensions, the layer of the stack (algorithms, classical infrastructure and device control) and the type of contribution (research, management and verification), which job titles systematically fail to reveal. The classical profile is led by Python, C++ and Rust, shares a common core of programming and DevOps tooling across all roles, and includes systems and hardware competences that are rarely reported in the general software labor market; for general and infrastructure engineering positions, which form the second largest role, quantum computing is the context rather than the object of the work. Quantum-specific demand is organized along the layers of the quantum computing stack, and half of the ads that mention quantum skills request at least one of its intermediate layers, namely control systems, error correction and compilation, which is where the work of the quantum software engineer appears most distinctive. Across these analyses, functional role is the variable most strongly associated with the skill profile, formal qualification is associated with quantum depth rather than with seniority, early-career positions are scarce, and startups and large enterprises demand different layers of the stack.

Beyond these findings, the study offers instruments that can be reused in future research. The three taxonomies were built inductively from the data and come with definitions, examples and decision rules for borderline cases; the quantum-specific one, organized by layers of the stack rather than by domains of practice, makes it possible to locate even vague requirements, such as *quantum hardware*, in a specific part of the technology, and may support the mapping of curricula and competence frameworks, such as the European Competence Framework for Quantum Technologies (Greinert et al. 2023), against the demand of the market. The job posting quality scale addresses a limitation shared by previous job ad studies: quality was strongly associated with

whether and how many skills an ad mentions, but much less with which ones, and it behaved differently for classical and quantum-specific skills. Without such a measure, differences in informativeness across employers, platforms or countries could be mistaken for differences in demand. We therefore recommend that future studies assess and report the quality of their samples, or at least restrict frequency comparisons to ads that mention the variable being analyzed.

The heterogeneity of roles (F1) and the weight of classical engineering (F2) suggest that there is no single curriculum for the quantum software engineer. For positions in which quantum computing is the context of the work, the findings support the lighter educational response recommended by Fox et al. (2020) and Hughes et al. (2022), one or two broad quantum courses rather than a full degree, covering the basic concepts of the field and the hybrid execution model, including cloud access, job pipelines and the integration of quantum processors with HPC resources. The intermediate layers of the stack (F3), in contrast, require a combination of compiler construction, embedded systems, signal processing and quantum information that no single traditional program provides, and they are where the gap between the quantum and classical paradigms reported by practitioners is widest (Aoun et al. 2021). Computing programs could offer electives on quantum compilation and on control software for experimental hardware, while physics and engineering programs could add software engineering practices to their laboratory training. Since the main bottleneck for newcomers seems to be the lack of experience rather than of advanced degrees (F4), universities could also combine undergraduate quantum courses with internships, capstone projects and contributions to open-source quantum software.

Employers may use the role structure (F1) to write more informative titles, so that a "quantum software engineer" position signals whether it expects compiler design, FPGA programming or a publication record, and programs that fund internships, apprenticeships or residencies could help to widen a pipeline that the market itself seems to keep narrow (F4). Candidates, in turn, need to read the responsibilities rather than the title of a position. The findings point to an entry path for classical software engineers, who can learn the quantum context on the job (F2), and to a choice between two routes: a postgraduate degree for positions in which quantum depth is the core of the work, and professional experience in classical engineering for positions in which quantum computing is the context or in which hardware-proximate skills are decisive, for which physics or electrical engineering graduates with solid software engineering training are natural candidates (F3, F4). Candidates interested in the layers closest to the device are more likely to find them in startups, and those interested in SDKs, applications or platform engineering in large enterprises (F5).

The results reflect what employers choose to state in public ads during a limited period, and they should be complemented by other sources of evidence. The mismatch between titles and functions (F1) also suggests that automated studies relying on titles or keyword searches would likely misclassify a substantial part of this market. Future work could repeat the analysis periodically with the same instruments to test whether quantum software engineering moves up the stack as hardware matures, interview hiring managers to understand how companies compensate for the scarcity of experienced professionals, and investigate the engineering of software *for* quantum

computers, such as how to test calibration routines, how to manage the co-evolution of firmware and quantum hardware, and how to engineer real-time decoders, which the market already demands but QSE research has only begun to address.

Acknowledgements

This work has been partially supported by the following projects: "QUASAR: QUAntum software engineering for Secure, Affordable, and Reliable systems", grant 2022T2E39C, under the PRIN 2022 MUR program funded by the EU - NGEU; SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union – NextGenerationEU; Patto territoriale "Sistema universitario pugliese" – CUP F61B23000370006

References

Akbar, M.A., Khan, A.A., Mahmood, S., Rafi, S.: Quantum software engineering: A new genre of computing. In: QSE-NE 2024: Proceedings of the 1st ACM International Workshop on Quantum Software Engineering: The Next Evolution, pp. 1–6. Association for Computing Machinery (ACM), ??? (2024). https://doi.org/10.1145/3663531.3664750

Aoun, M.R.E., Li, H., Khomh, F., Openja, M.: Understanding quantum software engineering challenges an empirical study on stack exchange forums and github issues. In: Proceedings - 2021 IEEE International Conference on Software Maintenance and Evolution, ICSME 2021, pp. 343–354. Institute of Electrical and Electronics Engineers Inc., ??? (2021). https://doi.org/10.1109/ICSME52107.2021.00037

Bourque, P., Fairley, R.: Swebok. Nd: IEEE Computer society, 2 (2004)

Chévez-Coronel, K., Saraguro-Bravo, R., Mawyin-Koriakova, M., Rugel-Diaz, P., Jaramillo-Saltos, J.: A comparative analysis of hard and soft skills required by software development companies in ecuador. In: Proceedings - 2023 11th International Conference in Software Engineering Research and Innovation, CONISOFT 2023, pp. 75–81. Institute of Electrical and Electronics Engineers Inc., ??? (2023). https://doi.org/10.1109/CONISOFT58849.2023.00019

Cruzes, D.S., Dyba, T.: Recommended steps for thematic synthesis in software engineering. In: 2011 International Symposium on Empirical Software Engineering and Measurement, pp. 275–284. IEEE, ??? (2011). https://doi.org/10.1109/ESEM.2011.36 . http://ieeexplore.ieee.org/document/6092576/

Cerioli, M., Leotta, M., Ricca, F.: What 5 million job advertisements tell us about testing: A preliminary empirical investigation. In: Proceedings of the ACM Symposium on Applied Computing, pp. 1586–1594. Association for Computing Machinery, ??? (2020). https://doi.org/10.1145/3341105.3373961

Dwivedi, K., Haghparast, M., Mikkonen, T.: Quantum software engineering and quantum software development lifecycle: a survey. Cluster Computing **27**, 7127–7145 (2024) https://doi.org/10.1007/s10586-024-04362-1

Silva Lopes, L.C.R., Carvalho, K.C., Kassab, M., Kalinowski, M., OliveiraJr, E., Neto, V.V.G.: Investigating the professional profile of software architects in brazil using job advertisements. In: Anais do XXXIX Simpósio Brasileiro de Engenharia de Software (SBES 2025), pp. 643–649. Sociedade Brasileira de Computação, ??? (2025). https://doi.org/10.5753/sbes.2025.10448 . https://sol.sbc.org.br/index.php/sbes/article/view/37039

Souza Santos, R., Baldassarre, M.T., França, C.: Industry expectations and skill demands in quantum software testing. In: Seventh International Workshop on Quantum Software Engineering (Q-SE 2026) (2026). https://doi.org/10.1145/3786150.3788614 . http://arxiv.org/abs/2512.14861

Daneva, M., Wang, C., Hoener, P.: What the job market wants from requirements engineers? an empirical analysis of online job ads from the netherlands. In: International Symposium on Empirical Software Engineering and Measurement, vol. 2017-November, pp. 448–453. IEEE Computer Society, ??? (2017). https://doi.org/10.1109/ESEM.2017.60

França, C., Mellet, D.: Soft skills required! uma análise da demanda por competências não-técnicas de profissionais para a indústria de software e serviços. In: IX Fórum de Educação em Engenharia de Software (FEES 2016) do XXX Simpósio Brasileiro de Engenharia de Software (SBES 2016) (2016)

Fox, M.F.J., Zwickl, B.M., Lewandowski, H.J.: Preparing for the quantum revolution: What is the role of higher education? Physical Review Physics Education Research **16** (2020) https://doi.org/10.1103/PhysRevPhysEducRes.16.020131

Giray, G.: An analysis of desired skills for software development jobs in turkey using text mining. In: Turkish National Software Engineering Symposium (2019). https://www.researchgate.net/publication/335234799_An_Analysis_of_Desired_Skills_for_Software_Development_Jobs_in_Turkey_Using_Text_Mining

Greinert, F., Müller, R., Bitzenbauer, P., Ubben, M.S., Weber, K.A.: Future quantum workforce: Competences, requirements, and forecasts. Physical Review Physics Education Research **19** (2023) https://doi.org/10.1103/PhysRevPhysEducRes.19.010137

Grandjean, N.: Soft skills in online job advertisements - an analysis of british online job ads for project managers and system engineers. PhD thesis, Université catholique de Louvain (2017). https://doi.org/https://hdl.handle.net/2078.2/6175 . https://hdl.handle.net/2078.2/6175

Hughes, C., Finke, D., German, D.A., Merzbacher, C., Vora, P.M., Lewandowski, H.J.:

Assessing the needs of the quantum industry. IEEE Transactions on Education **65**, 592–601 (2022) https://doi.org/10.1109/TE.2022.3153841

Hiranrat, C., Harncharnchai, A.: Using text mining to discover skills demanded in software development jobs in thailand. In: ACM International Conference Proceeding Series, pp. 112–116. Association for Computing Machinery, ??? (2018). https://doi.org/10.1145/3206129.3239426

Hasanovic, M., Panayiotou, C., Silberman, D., Stimers, P., Merzbacher, C.: Quantum technician skills and competencies for the emerging quantum 2.0 industry. Optical Engineering **61** (2022) https://doi.org/10.1117/1.oe.61.8.081803

Jebreen, I., Nabot, A.: An analysis of job advertisements for software engineering employability skills and knowledge: Jordan as a case study. Asia-Pacific Journal of Science and Technology **26**(4) (2021) https://doi.org/10.14456/apst.2021.46

Juárez-Ramírez, R., Jiménez, S., Navarro, C.X., Guerra-García, C., Perez-Gonzalez, H.G., Fernández-y-Fernández, C., Ortiz-Hernández, J., Cancino, K.: Skills required for quantum computing: A comprehensive review of recent studies. Programming and Computer Software **50**, 844–874 (2024) https://doi.org/10.1134/S0361768824700804

Kapitsaki, G., Chatzivasili, L., Papoutsoglou, M., Galster, M.: An exploratory study on soft skills present in software positions in cyprus: a quasi-replication study. In: Proceedings of the 18th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM '24) (2024). https://doi.org/10.1145/3674805.3686681

Kassab, M., Laplante, P., Defranco, J., Neto, V.V.G., Destefanis, G.: Exploring the profiles of software testing jobs in the united states. IEEE Access **9**, 68905–68916 (2021) https://doi.org/10.1109/ACCESS.2021.3077755

Leça, M.D.M., Santos, R.D.S.: Curious, critical thinker, empathetic, and ethically responsible: Essential soft skills for data scientists in software engineering. In: 2025 IEEE/ACM 47th International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS), pp. 151–162. IEEE, ??? (2025). https://doi.org/10.1109/ICSE-SEIS66351.2025.00021 . https://ieeexplore.ieee.org/document/11023902/

Matturro, G.: Soft skills in software engineering: A study of its demand by software companies in uruguay. In: 2013 6th International Workshop on Cooperative and Human Aspects of Software Engineering, CHASE 2013 - Proceedings, pp. 133–136 (2013). https://doi.org/10.1109/CHASE.2013.6614749

Meesters, M., Heck, P., Serebrenik, A.: What is an ai engineer? an empirical analysis of job ads in the netherlands. In: Proceedings - 1st International Conference on AI Engineering - Software Engineering for AI, CAIN 2022, pp. 136–144. Institute

of Electrical and Electronics Engineers Inc., ??? (2022). https://doi.org/10.1145/3522664.3528594

Niva, A., Markkula, J.: Communication skills requirements of junior software engineers- analysis of job ads. In: International Conference on Product-Focused Software Process Improvement, pp. 201–216 (2022). Springer

Papoutsoglou, M., Ampatzoglou, A., Mittas, N., Angelis, L.: Extracting knowledge from on-line sources for software engineering labor market: A mapping study. IEEE Access **7**, 157595–157613 (2019) https://doi.org/10.1109/ACCESS.2019.2949905

Peterssen, G.: Quantum technology impact: the necessary workforce for developing quantum software. In: Proceedings of the First International Workshop on QuANtum SoftWare Engineering & pRogramming (QANSWER). CEUR Workshop Proceedings, vol. 2561, pp. 6–22 (2020)

Papoutsoglou, M., Mittas, N., Angelis, L.: Mining people analytics from stackoverflow job advertisements. In: Proceedings - 43rd Euromicro Conference on Software Engineering and Advanced Applications, SEAA 2017, pp. 108–115. Institute of Electrical and Electronics Engineers Inc., ??? (2017). https://doi.org/10.1109/SEAA.2017.50

Piattini, M., Peterssen, G., Pérez-Castillo, R., Hevia, J.L., Serrano, M.A., Hernández, G., Guzmán, I., Paradela, C.A., Polo, M., Murina, E., Jiménez, L., Marqueño, J.C., Gallego, R., Tura, J., Phillipson, F., Murillo, J.M., Niño, A., Rodríguez, M.: The Talavera manifesto for quantum software engineering and programming. In: Proceedings of the First International Workshop on QuANtum SoftWare Engineering & pRogramming (QANSWER) (2020). Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0)

Piattini, M., Serrano, M., Perez-Castillo, R., Petersen, G., Hevia, J.L.: Toward a quantum software engineering. IT Professional **23**(1), 62–66 (2021) https://doi.org/10.1109/MITP.2020.3019522

Stefano, M.D., Pecorelli, F., Nucci, D.D., Palomba, F., Lucia, A.D.: Software engineering for quantum programming: How far are we? Journal of Systems and Software **190** (2022) https://doi.org/10.1016/j.jss.2022.111326

Shadrin, D., Sozykin, A., Kuklin, E., Bersenev, A.: Classification of software developer skills using analysis of russian job ads. In: Proceedings - 2022 Ural Symposium on Biomedical Engineering, Radioelectronics and Information Technology, USBEREIT 2022, pp. 236–239. Institute of Electrical and Electronics Engineers Inc., ??? (2022). https://doi.org/10.1109/USBEREIT56278.2022.9923408

Vassilev, G., Romanko, O., Evans, K.: What's in a job? measuring skills from online job adverts. In: 36th IARIW Virtual General Conference (2021). https://www.researchgate.net/publication/357807535

Wang, C., Tang, Y., Liang, P., Daneva, M., Sinderen, M.V.: What industry wants from requirements engineers in china? an exploratory and comparative study on re job ads. In: International Symposium on Empirical Software Engineering and Measurement. IEEE Computer Society, ??? (2020). https://doi.org/10.1145/3382494.3410672

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Supplementarymaterial.zip