

# A Cross-Continental Analysis of How Regional Cues Shape Top Stack Overflow Contributors

Elijah Zolduoarrati\*

University of Otago, elijah.zolduoarrati@otago.ac.nz

Sherlock A. Licorish

University of Otago, sherlock.licorish@otago.ac.nz

John Grundy

Monash University, john.grundy@monash.edu

## ABSTRACT

Stack Overflow offers valuable knowledge for software developers, but studies suggest digital information tends to cluster geographically, limiting access to necessary knowledge for innovation. This study explores posts of top contributors on Stack Overflow across the United States, Brazil, India, Egypt, the United Kingdom, and Australia. We analyse platform activities, conduct social network analysis, employ topic modelling paired with thematic analysis, before dissecting their knowledge sharing patterns via directed content analysis. Results indicate that cultural factors, entrepreneurial activities, tech ecosystem maturity, as well as workforce diversity in a region were found to shape how top contributors contribute. For instance, individualistic users communicate directly whilst collectivistic users prefer subtle communication and socio-emotional cues. Moreover, top contributors in nascent technology ecosystems were more likely to discuss fundamental concepts, while those in mature ecosystems focus on specialised niches. This study sheds light on how diversity in human aspects may influence the dynamics of CQA settings, where future researchers can explicate the extent of which latent contextual factors affect user contributions and community structure.

**Keywords:** contribution, cross-regional behaviour, repository mining, social network analysis, topic modelling, Stack Overflow, user/contributor interactions

## 1 INTRODUCTION

Stack Overflow has emerged as key knowledge source for many software engineers to navigating a multitude of coding challenges since its inception in 2008. Through a unique model of peer-to-peer support and collective problem-solving, Stack Overflow has transformed the landscape of software development, empowering developers to collaboratively tackle programming obstacles and advance the industry as a whole [1]. Despite the recent emergence of large language models (LLMs), such as GitHub Copilot and ChatGPT, its importance has still remained high within the software development community to date [2]. Owing to its collaborative environment and vast information repository, the platform has continuously attracted a diverse audience of software developers. Individuals from diverse backgrounds share expertise and seek assistance to bolster software development rigours [3]. Stack Overflow thus serves as a microcosm of the software ecosystem, reflecting the diverse ways in which individuals from different countries engage in collaborative problem-solving. Naturally, the website also reflected certain diversity-related trends observed within the larger IT field [4, 5]. For instance, gender segregation is as widespread in the platform as it is in real life [5, 6]. Despite an emerging body of research, pertinent literature on Stack Overflow primarily focusses on intra-community determinants, and those that do study factors related to geographical domains only do so in isolation [7]. Examples include culture and education diversity which were investigated by two different research endeavours [8, 9]. Compounding this issue, the geographical clustering of digital information [10] creates a complex interplay between human diversity, geographic location, and contribution levels that warrants further investigation.

This paper describes a comprehensive exploration of question-answering practices on Stack Overflow, focussing on variations observed across countries situated on different continents. We focussed on the United States, Brazil, India, Egypt, the United Kingdom, and Australia. These countries were chosen

---

\*Corresponding author

given their high user<sup>1</sup> bases within their respective continents: North America, South America, Asia, Africa, Europe, and Oceania. We believe this set allows us to explore potential associations between users' online behaviour and latent regional characteristics related to diversity in human aspects (e.g., education, culture, language, and tech presence), which is the goal of our study. Our investigation focussed on the behaviour of top contributors within each country. This approach was predicated on the rationale that highly active contributors exhibit distinct characteristics compared to those with sporadic participation. For example, gamification factors driving their contributions, such as the Stack Overflow reputation system, may vary from those with occasional contributions. Previous research has also shown those members to be most integral to teamwork [11, 12], especially regarding feature development and functionality extensions [13]. Thus, understanding their dynamics can be noteworthy. We employ the Stack Overflow Database from the Stack Exchange Data Dump, which has been employed effectively in prior research [14].

Firstly, we begin with performing a quantitative exploration of platform usage patterns within each country, subsequently assessing the statistical significance of observed fluctuations. This first research question (RQ1) aims to cast a wide net and highlights any prevailing trends before narrowing our focus to the top contributors in each country. Next, we conducted a social network analysis (SNA) to observe intra-country user interactivity, elucidating prominent users based on sentiment and contribution towards their fellow countrymen. We then used topic modelling to investigate the content of participant discussions and identify prominent topics of interest within the top contributors' exchanges of each country. To achieve this, we harness Latent Dirichlet Allocation (LDA), which has been used to good effect in previous literature of similar design [15]. Topics were then categorised systematically using thematic analysis. After the identification of prevalent discussion topics within each country, a directed content analysis was employed to systematically investigate deeper meaning in top contributors' exchanges. We then examine the interaction and knowledge-sharing behaviours of software developers, identifying key contribution patterns and behaviours across countries to supplement our quantitative findings.

Previous works have explored the link between regions and user contributions on Stack Overflow, yet such studies were primarily focussed on higher-level comparisons [9, 16]. There is a lack of in-depth analyses on intra-country user interactions, predominance of topics, and how knowledge exchange practices differ across regional contexts. This gap requires a deeper investigation that examines the gradations within and between regions. We wanted to answer the following RQs:

**RQ1:** How do users' Stack Overflow contribution levels compare across different countries?

**RQ2:** How do users within specific countries interact with each other via Stack Overflow?

**RQ3:** What are the predominant Stack Overflow topics of discourse among users from different countries?

**RQ4:** How do Stack Overflow users' knowledge exchange practices differ across countries?

Our first RQ aims to highlight any prevailing patterns of users' contributions. A broad examination of cross-country contribution patterns is carried out, revealing notable fluctuations across countries. The second RQ examines the role of social influence, expertise, and collaborative behaviours in shaping the knowledge-sharing landscape among peers. It also presents a nuanced comprehension of user community structures within countries. This enables us to discern which users within each country exhibit higher levels of contribution, revealing patterns of activity and influence. Our third RQ is motivated by the need to discern the thematic landscape of discussions, elucidating key areas of interest within each country's user community. Finally, our fourth RQ aims to capture the subtleties of communication styles within each country and the degree to which they differ from one another.

This study makes several novel contributions. From a theoretical standpoint, our study enriches both qualitative and quantitative discourse on how top contributors from diverse backgrounds engage in collaborative problem-solving within a global digital space. Furthermore, our study contributes to the current literature on the social network dynamics of Stack Overflow, subsequently elucidating diverse topics and users of interest that emerge within localised contexts. To the best of our knowledge, our

---

<sup>1</sup> The terms 'user' and 'contributor' are used interchangeably given that those participating on Stack Overflow are seen to occupy both roles.

study is the first to use both qualitative and quantitative insights to illuminate the dynamics of cross-continental Stack Overflow user interactions. Insights gleaned from our analyses may benefit the scientific community in studying the evolution of Stack Overflow (or other CQA platforms) and how different regions contribute to their success. For industry practitioners, insights from our study present several real-world applications. First, practitioners can tailor communication strategies for optimising effectiveness and encouraging collective problem-solving among online collaborators, particularly for similar CQA websites. Additionally, recognising topics of interest within each country can guide industry practitioners in staying informed about emerging trends, potential challenges, and towards addressing knowledge gaps.

The remainder of this manuscript is organised as follows. Section 2 offers background information pertinent to our research regarding related work, Section 3 outlines the research questions, variables, and measurements. Section 4 delineates our study’s experimental design, while Section 5 details the methodologies employed for data collection, processing, and analysis, Section 6 presents the study’s findings in relation to each RQ. Section 7 discusses and contextualises these findings, reflecting on their implications for theory and practice. Section 8 considers the study’s limitations, and we conclude with final remarks and suggestions for future research in Section 9. Additionally, a replication package is provided for those interested in further examining our research methodology [17]<sup>2</sup>.

## 2 RELATED WORK

Programming-related Q&A online platforms, such as Stack Overflow, have become hubs for knowledge diffusion within the software engineering (SE) field [2, 18]. Initially a regional community in North America and Europe, Stack Overflow has evolved into a global platform attracting students and professionals alike [19]. The emergence of generative language models like ChatGPT and Gemini have been associated with a decline in Stack Overflow participation [20], yet the prevalence of errors in over half of GPT-generated solutions underscores the continued relevance of Stack Overflow as a reliable resource for programming-related inquiries [20]. In fact, human-written answers were found to be more relevant, utilisable, clearer and yet more concise compared to those generated by GPT models [21], suggesting that these technologies, while evolving, still require improvements. Many research studies have thus delved into knowledge exchange practices on Stack Overflow, offering valuable insights into participatory and crowdsourced knowledge creation [22]. Zagalsky et al. [22] proposed a dual framework, distinguishing between participatory collaboration and independent crowdsourced contributions. Zhang et al. [23] explored the longevity of exchanged knowledge, revealing notable trends in obsolescence. Compared to answers pertaining to web development or database systems, those related to mobile app development exhibited a higher likelihood of becoming obsolete, presumably due to such fields being continuously evolving [23]. Research using Stack Overflow has also explored specific topics, as the multifaceted nature of SE encompasses a diverse array of languages, tools, and technologies [24]. Prevalent topics of discourse include web and mobile applications [24], as well as dynamics in continuous integration/continuous delivery (CI/CD), with a focus on error logging and build issues [1]. Uddin et al. [25] introduces an additional perspective derived from discussions related to the Internet of Things (IoT), where predominant topics include software and network systems, coinciding with the increasing popularity of open-source systems.

Cultural influences have been found to impact user participation on Stack Overflow [9, 16, 26]. Notably, the breadth of the SE community’s understanding of cultural influences is shaped by Hofstede’s cultural framework which incorporates a crucial dimension known as individualism-collectivism [27]. This spectrum measures the extent to which individuals prioritise their self-reliance and autonomy (individualistic) compared to prioritising their group membership and obligations (collectivistic). Recent studies have investigated user contributions in relation to these cultural variations. In one instance, it was found that individuals from collectivist societies are less likely to contribute given the misalignment of the platform’s core design against their cultural expectations [16]. For example, certain users desire social engagement, yet the platform’s community guidelines accentuate productivity, discouraging informal discourse and off-topic conversations [16]. Moreover, developers from individualistic cultures tend to accrue higher reputations and greater use of the pronoun “I,” paired with a stronger task-oriented

---

<sup>2</sup> <https://zenodo.org/records/13994284>

focus. In contrast, collectivistic developers demonstrate more frequent use of the pronouns “we” and “you,” and are more inclined towards information exchange behaviours [9].

Linguistic aspects have also been found to exert an influence on Stack Overflow contributions. Native speakers displayed a greater tendency to ask more questions and leave more comments, suggesting a language-mediated effect on user behaviour online [28]. An anthropological theory that classifies cultures based on communication styles is Hall’s context theory, distinguishing between low-context and high-context cultures [29]. Individuals in high-context cultures heavily rely on nonverbal cues, such as body language and shared experiences, to convey meaning. Conversely, low-context cultures emphasise clear and direct verbal communication [30]. While this framework has been extensively applied in various fields, research specifically investigating its influence on Stack Overflow remains scarce, which we aim to address in the current study.

Researchers have begun to develop unifying indices that aim to provide a more nuanced understanding of diversity issues on Stack Overflow. Inglehart and Baker [31] devised the Self-Expression Index. Higher values signify a greater prioritisation of subjective well-being and quality of life over basic survival values, such as economic and physical security (commonly taken for granted in most post-industrial societies) [31]. Within Stack Overflow, contributors from nations with higher indices (i.e., prioritising quality of life and well-being), such as the USA and Canada, exhibit slightly more tendency to provide answers, indicating voluntary participation without anticipated incentives [28]. Levine and Norenzayan [32] propose the Pace of Life metric, encompassing the pace of one’s movements, work intensity, and daily experience density [33]. Countries with colder climates, economically productive nations, and individualistic cultures tend to also have the highest pace of life, whereas economically undeveloped countries would yield the slowest [32]. Such variations largely propagate to Stack Overflow, where contributors from countries with a higher pace of life (e.g., Japan or USA) tend to participate more [16].

Despite its inherent design as a community question-answering (CQA) platform, Stack Overflow appears to function akin to a social network, where knowledge exchange practices are conducted irrespective of users’ geographic location, and researchers have thus employed social network graphs to model the interactions and relationships among users [8, 34, 35, 36]. Odiete et al. [8] employed SNA to map latent connections among experts within various programming language communities, revealing a distinct association between the age and level of abstraction of a language and the profile of its user base. Older, lower-level languages like C, Fortran, and Assembly Language tend to attract a higher concentration of experienced developers, while newer, higher-level general-purpose languages like Java and Python are more popular among developers comfortable with a multitude of programming paradigms. Menshikova [34] employed a similar methodology to identify and characterise overlapping communities within the platform, revealing the existence of interconnected groups of users with shared interests, such as the overlap between the Python and TensorFlow communities. Brooke [36] brings gender-based segregation into light, applying SNA to investigate peer parity and reciprocity. Interestingly, contributions of female users were found to be undervalued compared to male and anonymous counterparts [36].

While prior studies have identified discrepancies in participation to Stack Overflow based on country, language, tech presence, and culture, a holistic examination of these factors is lacking. Our study aims to fill this gap and provide a deeper analysis regarding the collective impact of these factors. We provide an inclusive perspective on how top contributors from different backgrounds interact with each other, elucidating otherwise subtle variations across different user archetypes.

### **3 RESEARCH QUESTIONS**

To address the RQs introduced in Section 1, it is imperative to clearly define what variables the questions are investigating, and how such constructs are eventually measured. Firstly, RQ1 will be answered by quantifying users’ contribution levels, which encompasses a set of comparative quantitative analyses of platform-related variables such as total questions, answers, and account age. RQ2 then focusses on analysing users’ interaction patterns. The study uses SNA, reflecting on pertinent literature [36]. The core construct for RQ3 is users’ discussion topics, harnessing LDA and eventually triangulating keywords using thematic analysis. RQ4 delves into the examination of users’ knowledge exchange,

quantified using content analysis. Table 1 provides a comprehensive overview of all RQs, variables, and approaches employed in the study.

Table 1. Description of RQs

RQ#	Research Question	Variable	Approach/Methodes
RQ1	How do users' Stack Overflow contribution levels compare across different countries?	Contribution levels.	Descriptive statistics.
RQ2	How do users within specific countries interact with each other via Stack Overflow?	Interaction patterns.	Social network analysis.
RQ3	What are the predominant Stack Overflow topics of discourse among users from different countries?	Topics of discourse.	Topic modelling. Thematic analysis.
RQ4	How do Stack Overflow users' knowledge exchange practices differ across countries?	Knowledge exchange practices.	Directed content analysis.

In the next section, we summarise our experiment design regarding data collection and location inference approaches. Finally, we outline each RQ's appropriate methods as outlined in Table 1.

#### 4 EXPERIMENT DESIGN

Following an examination of the RQs in Table 1, we concluded that the required constructs and corresponding measurements can be derived from the Stack Overflow (SO) Database. To ensure data integrity and suitability for analysis, a series of preprocessing steps were applied to prepare the dataset for addressing all four RQs. Firstly, the four RQs' emphasis is directed on users' geographic origins to enable comparative analysis. Thus, our initial step was directed to determining and recording users' locations, which will eventually be extracted and stored for subsequent analysis and querying. Afterwards, for RQ1, the concept of "*user contribution*" is somewhat multifaceted and therefore remains an abstract construct. As proxies, we seek to operationalise platform variables, thereby enabling quantifiable measurement with respect to the previously-inferred locations. To examine users' interaction patterns through SNA for RQ2, the quantity of answers exchanged between users will be analysed. This encompasses answers given (outgoing) and received (incoming). Additionally, to enrich the understanding of user interactions, the degree of polarity expressed in answers will be quantified. This sentiment analysis, in conjunction with the volume of exchanged answers, will provide a more nuanced perspective on user relationships within the social network.

To investigate users' topics of discourse for RQ3, Latent Dirichlet Allocation will be employed to identify the most relevant keywords pertaining to a topic. These keywords will subsequently undergo thematic analysis following the framework outlined by Braun and Clarke [38]. To ensure the accuracy of inferred themes, thematic analysis will be guided with relevant literature. Finally, in examining users' knowledge exchange practices as outlined in RQ4, a directed content analysis will be conducted to scrutinise users' actual posts and identify latent cross-country variations. This analysis will determine whether observed thematic differences in user exchanges are a byproduct of chance, or attributable to broader diversity-related contextual factors. Our experimental design is summarised in Figure 1. The next section explains these approaches in detail.

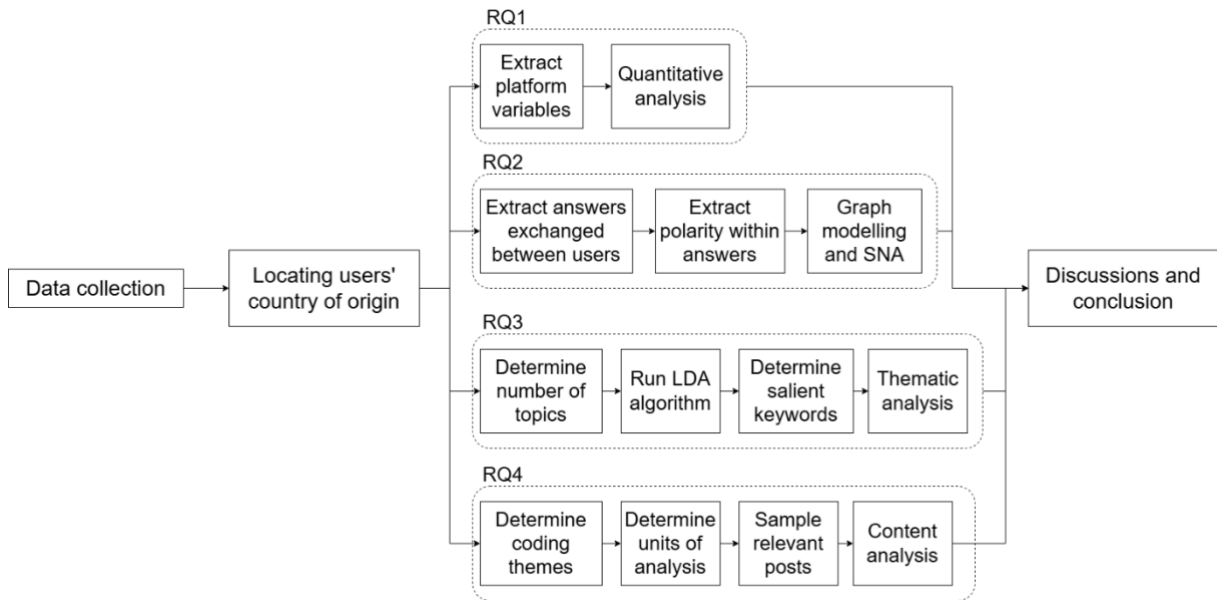


Figure 1. Experiment design

## 5 EXPERIMENT SETUP

### 5.1 Data Collection and Processing

We firstly downloaded the SO Database from the Stack Exchange Data Dump<sup>3</sup> as done in previous research projects [39]. This relational database, stored locally on Microsoft SQL Server, includes diverse activities and raw metrics in tables such as users’ posts, comments, badges, and edits. We used the June 2022 release<sup>4</sup>, being the latest at the time of our analysis. The database included 56,264,788 posts, 85,467,182 comments, and 22,796,157 edits from 17,922,426 unique users. It is worth noting that our study exclusively employed data from this database, refraining from incorporating any external data to answer the RQs. The total rows for each table are presented in Table 2. We used these records to explore the contribution profiles on the Stack Overflow platform. The interactive database schema in the *Figures* folder of our replication package [17] shows the complete database structure<sup>5</sup>.

Table 2. Stack Overflow Database properties

Table Name	Total Rows
Posts	56,264,788
Questions	22,634,239
Answers	33,520,483
Non-QnA*	110,066
Comments	85,467,182
Edits	22,796,157
Users	17,922,426
Badges	48,088,681

\*) Non-QnA refers to posts that are neither questions nor answers (e.g., tag wiki or moderator nominations)

### 5.2 Locating Users’ Country of Origin

To determine users’ country of origin, we excluded 14,174,843 users with null location fields, retaining data for 3,814,901 unique users. Second, inconsistencies in naming conventions were found, such as variations like “SF Bay Area” and “San Fran” both denoting San Francisco, but treated as distinct locations. To address such mismatches and enable aggregation, we explored 15 proprietary geocoding APIs from numerous GIS providers, including ArcGIS, Bing Maps, GeocodeEarth, Geodict, Geograpy, Geonames, Geopy, Geotext, Google Maps, MapQuest, Nominatim, OpenCage, OpenStreetMaps, Photon, and TomTom. The selection considered factors such as suitability for the study, commercial availability, and insights from recent literature. Previous literature has outlined that geocoding APIs employ complex algorithms that leverage machine learning techniques, thereby enhancing their

<sup>3</sup> <https://archive.org/details/stackexchange>

<sup>4</sup> [https://archive.org/details/stackexchange\\_20220606](https://archive.org/details/stackexchange_20220606)

<sup>5</sup> Replication package » Figures » SO Database Schema.html

reliability and precision [40]. As an example, Peterson [41] compared Google Maps API, Bing Maps, MapQuest, and OpenStreetMap, revealing disparities in cartographic functionality. Certain features available in one API were absent in others e.g. the listener object to override zoom level on user click is present in Google Maps API but not in Bing Maps. After a comprehensive assessment of these factors, we deemed the fitness of 15 APIs for our use-case, enabling us to conduct a thorough evaluation to facilitate our final selection.

To yield a 95% confidence level with a 5% error margin, a ‘ground truth’ dataset containing 384 distinct cities and their respective states was curated, where each location was manually labelled by hand. This ‘ground truth’ enables a performance comparison of each API in a standardised manner, ensuring a fair evaluation. Specifically, outputs from the API were cross-referenced with our ‘ground truth’ dataset using NLTK and word2vec packages for automation. This facilitated assessing the alignment between API results and our manually-validated ‘ground truth’ dataset. For example, if 384 observations were checked where Bing’s geocoding API matched 377 of the manual observations, then Bing’s accuracy would be 98.18%. The accuracy of each API is shown in Table 3. For more detailed information on the evaluation of these GeoAPIs, please refer to the *GeoAPI Evaluation* section within our replication package [17].

Table 3. Accuracy for each API

API	Accuracy		
	City	State	Country
ArcGIS	94.55%	90.39%	99.22%
Bing	98.18%	79.22%	99.22%
GeocodeEarth	98.18%	92.21%	99.22%
Geodict	64.68%	43.64%	63.12%
Geograpy	96.10%	48.83%	74.03%
Geonames	97.66%	90.39%	87.01%
Geopy	97.40%	69.87%	97.14%
Geotext	91.69%	43.64%	71.69%
Google	95.06%	90.91%	98.44%
MapQuest	98.18%	92.08%	99.22%
Nominatim	96.10%	84.16%	88.31%
OpenCage	91.95%	89.09%	99.48%
OpenStreetMaps	53.25%	89.35%	96.88%
Photon	50.13%	79.74%	95.06%
TomTom	96.62%	86.75%	96.88%

Table 3 shows us that MapQuest provided us with the most accurate API, which demonstrated satisfactory performance across city (98.18% accuracy), state (92.08%), and country-level queries (99.22%). In contrast, certain APIs, including OpenStreetMaps and Photon, demonstrated subpar accuracy when it came to city-level queries. These APIs exhibited the lowest accuracy among all the tested APIs, with OpenStreetMaps yielding 53.25% accuracy and Photon returning 50.13%. Additionally, Geodict showed the least reliability across all APIs, having an accuracy of 64.68% for cities, 43.64% for states, and 63.12% for countries. Given these findings, a decision was made to standardise all user “Location” entries and detect invalid entities by utilising MapQuest, such as “Tatooine” (Star Wars) for Users.Id<sup>6</sup> 11896718 and “Valhalla” (Norse mythology) for Users.Id 10938505. Moreover, some users provided non-geographical locations like “Mostly at home” for Users.Id 2736499. As 67,318 invalid location entries were then discarded, we simultaneously ensure that city, state, and country entries are entirely accurate. As a demonstration, the string “Manhattan, NYC” would return “New York City” as the city, “New York” as the state, and “United States” as the country.

<sup>6</sup> Prefixes denote the table that a particular column is in. For instance, “Users.Id” refers to the column Id within the Users table of the Stack Overflow Database (refer to our replication package [17] for the full database schema).

To ensure the representativeness of our downstream analyses, we selected the top country with the largest user base per continent on Stack Overflow. Our rationale behind this decision is that we wanted to ensure a representative sample, enhancing the reliability of subsequent findings whilst reducing the risk of drawing erroneous conclusions with otherwise small, unrepresentative populations. Prior processing steps have allowed us to tally the total users per country, eventually picking the ones with the most users per continent, as seen in Table 4.

Table 4. Top countries per continent, in terms of total users

Continent	Country	Total Users	Total Males	Total Females
North America	United States	760,809	630,622	130,187
Asia	India	665,163	552,911	112,252
Europe	United Kingdom	146,152	124,187	21,965
South America	Brazil	95,968	83,763	12,205
Oceania	Australia	61,464	51,572	9,892
Africa	Egypt	33,189	27,582	5,607

This provided us with a total of six countries for analysis, with the United States having the largest user base (760,809 users; 130,187 females) and Egypt having the smallest (33,189 users; 5,607 females) among them. The inclusion of all six countries also represent equal proportions on both low-context individualistic and high-context collectivistic cultures, thereby allowing for a comprehensive understanding of developer communities across regions.

### 5.3 User Profile and Users' Contribution Levels (RQ1)

In order to answer our RQ1, we extracted several platform variables for each user profile, as shown in Table 5. These variables have been extracted in prior work, and have been proven to yield accurate data after manual analysis [26].

Table 5. Extracted Stack Overflow platform variables

Platform Variable	Description
<i>Gender</i>	The user's gender, inferred by Genderize <sup>7</sup> [39].
<i>ProfileLength</i>	Number of words users wrote about themselves in their profile.
<i>YearlyDurationUsage</i>	Number of years between their initial registration to their last access date.
<i>UpVotes</i>	Total upvotes casted by a user, indicating usefulness of posts.
<i>DownVotes</i>	Total downvotes casted by a user, indicating posts of minimal value.
<i>Views</i>	Total number of times a user profile has been viewed.
<i>Reputation</i>	Total number of scores from a user's posts.
<i>Questions</i>	Total questions asked by a user.
<i>Answers</i>	Total answers provided by a user.
<i>Comments</i>	Total comments provided by a user.
<i>Edits</i>	Total edits a user made to existing posts/comments.
<i>Badges</i>	Total badges earned by a user, regardless of bronze, silver, or gold tiers.

Our quantitative analysis entails calculating the arithmetic mean for each platform variable, allowing us to observe pertinent trends across countries. Validation of results entails performing Kruskal-Wallis test to ascertain whether observed fluctuations are indeed statistically significant [42]. Upon rejecting the null hypothesis, our study proceeds with the Conover-Iman post-hoc test to identify specific pairwise differences between groups [43], eventually enabling us to designate which countries differ from the others. Our implementation applies the Bonferroni correction to adjust yielded p-values to limit spurious false positives [44]. Given our limited scope of only six countries, trends stemming from such tests would be more distinctive.

### 5.4 Users' Interaction Patterns (RQ2)

To answer RQ2, we narrow our focus to select the top 100 users with the highest aggregate of posted and received answers, considering the impracticality of visualising all user interactions in a country. Our

<sup>7</sup> <https://genderize.io>



approach draws inspiration from Licorish and MacDonell [45], where they employed a social network analysis technique with respect to influential developers within the IBM Jazz project based on their communication patterns. Selecting this subset of contributors offers a sufficient volume of interactions, reasonably representing each country’s SE landscape, whilst maintaining brevity. The subsequent RQs would therefore also be focussed on these selected top 100 contributors per country, ensuring sample stability and more concentrated analyses. While this strategy strengthens internal validity and facilitates detailed examination within this subset, we acknowledge the inherent trade-off of excluding potentially valuable interactions beyond the chosen participants (see Section 8 for further discussion). It is important to note that the subsequent inferences drawn pertain specifically to the top contributors within each country, and that our findings should not be generalised to profile the entire user base of the countries studied.

Afterwards, on these top contributors, we performed SNA with graph models which were proven to yield good efficacy in prior works of similar design [35, 46, 47]. Graphs are widely employed in SNA due to their effectiveness in representing complex relational structures [48]. Thus, given a graph  $G(V, E)$ , the node set  $V$  represents users whereas the edge set  $E$  represents the interactions between them [49]. This allows us to investigate intra-country user interaction by analysing their responses to each other’s questions, providing insight into the interaction patterns among the top contributors within a specific country. We omit analysis of questions with global participation beyond our selected six countries. For example, a question from a user in Australia might elicit responses from countries like Japan or South Korea, which are not part of our sample. For question threads that accrue answers from both the six target countries and other regions, data were selectively extracted to include only answers originating from the designated six countries. For instance, if a question garnered answers from Egypt, India, and Singapore, the responses from Egypt and India would be retained while the ones from Singapore were to be excluded. Finally, to maintain focus on the core RQ regarding top contributors’ interaction patterns, comments were excluded from this analysis. While comments offer valuable context, suggestions, and clarifications, they primarily serve as supplementary information rather than direct answers<sup>8</sup>. Answers are also considered the core content in CQA sites, offering richer information compared to other content forms [50]. As such, we direct our focus to answers only.

We harness signed directed graphs to illustrate the sentiment in interactions, where positive signs refer to affirmation (e.g., esteem and praise), and negative signs to contradiction (e.g., dislike and blame) [47]. Directed graphs were used based on our primary focus for this RQ, which was to characterise the behaviours of top contributors, rather than analysing the patterns of their reciprocal continuous interactions. Next, edges in the graph are weighted based on the total number of answers a user has posted to another user. For example, if user  $A$  provides more answers to user  $B$ ’s questions than to user  $C$ ’s, the weight assigned to the edge  $(A, B)$  will be greater than that of  $(A, C)$ . It is essential to recognise that an edge  $(B, A)$  does not constitute the same back-and-forth conversation as the edge  $(A, B)$ . Rather, these represent distinct communicative interactions within separate question-answer contexts. Unlike platforms facilitating direct messaging, Stack Overflow’s structure mandates a question-response paradigm, thereby limiting communication to this format. Thus, the edges  $(B, A)$  and  $(A, B)$ , despite involving the same individuals, would typically represent distinct interactions under different question threads that are unrelated to each other. Node sizes would be directly proportional to its *degree centrality*, which is the sum of *outdegree centrality* (interactions posted) and *indegree centrality* (interactions received) [51, 52]. For example, if a certain user  $A$  posted 10 answers and received 34 answers in total, their *degree centrality* would be 44. This implies larger node size also indicates higher user activity. Moreover, using such a measure shows users’ prevalence in community interactions, taking into account both directions of interactions (i.e., giving and receiving) [51].

Signs are visualised through colours in the graph, aligning with other practices [45, 53]. Edges are coloured blue if the corresponding emotion is positive, and red otherwise. However, if the corresponding emotion is neutral, we colour the edge black. Layout of the graph uses the Fruchterman-Reingold force-directed graph layout algorithm, where edges are treated like dampeners that attract nodes, while nodes repel each other [54, 55]. The iterative algorithm gradually converges towards an equilibrium state, characterised by a balance between repulsive and attractive forces, thereby optimising node

---

<sup>8</sup> <https://stackoverflow.com/help/privileges/comment>

arrangement. The number of iterations was set to 100 based on empirical evidence demonstrating its effectiveness in achieving an optimal node distribution [56].

### 5.5 Users’ Topics of Discourse (RQ3)

To answer RQ3, we conduct topic modelling of the top contributors’ contributions using Latent Dirichlet Allocation (LDA). LDA, being one of the most widely-used techniques to model topics, has demonstrated good results in various studies employing similar research designs owing to its three-level hierarchical Bayesian structure [46]. Previous works have thus incorporated topic modelling to observe latent topics in users’ discussions in Stack Overflow [24, 57]. However, studies that perform a cross-regional analysis are scarce. To address this gap, we aim to uncover the dominant topics within users’ discussions in each country, enabling us to gain insights into the prevailing themes of discourse among users within a particular geographical context. To maintain a similar coverage with the previous RQ, we seek to model topics derived from the answers provided by the top 100 users within each respective country.

Hyperparameter optimisation associated with the Dirichlet distribution was used to derive an optimal parameter set to effectively model users’ answers [46]. First, the number of topics  $k$  should be determined optimally, as too large values may lead to overly fine-grained, complex topics that are difficult to analyse. Conversely, too small values of  $k$  may result in overlapping topics that may not be distinctive towards each other [58]. In light of this dilemma, we iterate over a range of 2 to 100  $k$  values and evaluate *Topic Coherence* and *Jaccard similarity index* for each [1, 59]. *Topic Coherence* measures the semantic similarity among its high-scoring words (i.e., whether words tend to co-occur along the corpus) [60], wherein the *Jaccard index* is used to measure similarities between two text segments, enabling calculation of entity overlap [61, 62]. These measures allowed us to pick  $k$  that yields the highest *Topic Coherence* and the lowest *Jaccard index*, as the intuition is to cover a diverse, non-overlapping range of topics where each topic is semantically meaningful to human observers [1]. Secondly, we define the question-topic density parameter  $\alpha$  (higher values imply broader topics) and the topic-word density parameter  $\beta$  (higher values imply more topics are required to describe one document) [46]. We adhere to *de facto* standard heuristics where  $\alpha = 50/k$  and  $\beta = 0.01$ , proven to yield good effect in previous studies [63, 64]. Third, we designated Kullback-Leibler (KL) Divergence score as a distance metric to calculate the likelihood that the observed term  $w$  was generated by the latent topic  $T$ , popularised due to its computational efficiency [46, 59]. Afterwards, we set the maximum number of iterations to 2,000 and number of passes to 50, with continuous perplexity measurement to gauge convergence [65]. Perplexity is defined as the inverse of the geometric mean per-word likelihood, where lower perplexity values indicates a better fit to the data [66]. If convergence occurs before 2,000 iterations, signifying satisfactory model generalisation, we terminate remaining iterations.

Words that occur in all topics provide minimal information about the main topical composition of an answer [67]. Thus we compute topic *distinctiveness* and *saliency* as suggested by Chuang et al. [67] to avoid the use of incoherent or insignificant term groupings. First, we compute topic *distinctiveness* wherein we first calculate the conditional probability  $P(T|w)$  for a given word  $w$  and latent topic  $T$ .  $P(T|w)$  dictates the likelihood that  $w$  was generated by  $T$ . We also compute the marginal probability  $P(T)$  which conveys the probability that a randomly selected word  $w'$  was generated by  $T$ . The *distinctiveness* of a given word  $w$  is then the KL Divergence score between  $P(T|w)$  and  $P(T)$  [67]:

$$\mathbf{distinctiveness}(w) = \sum_T P(T|w) \log \frac{P(T|w)}{P(T)} \quad (1)$$

This formulation quantifies the information gain obtained by observing a specific term  $w$  in a document, compared to the average information gain achieved by observing any random term  $w'$ . In other words, it measures the relative informativeness of term  $w$  in identifying the document’s underlying topic [67]. This approach is particularly useful as it uncovers terms that are both commonly encountered yet specific to a certain thematic domain. Therefore, the *saliency* of a term  $w$  is defined by the product of its *distinctiveness* and probability of occurrence [67]:

$$\mathbf{saliency}(w) = \mathbf{distinctiveness}(w) \times P(w) \quad (2)$$

Determining the term *saliency* enables swift classification and resolution of thematic ambiguities [67]. However, LDA typically lacks an automatic mechanism for generating meaningful topic names, necessitating manual inference and labelling [57]. To this end, we employ thematic analysis to derive pertinent themes, in accordance to prior works of similar design [1, 68]. We adopt the inductive (bottom-up) approach outlined by Braun and Clarke [38]. Two coders employ an inductive coding approach to determine the thematic context of each word cluster, and subsequently, infer the overarching conceptual theme represented by these words. Categories are successively clustered into broader and broader categories, culminating in a hierarchical taxonomy by consulting pertinent literature and SE-related textbooks. For example, the keywords set “*element, event, attribute, jquery, selector, tag, div, use, child, html*” were found to strongly relate to jQuery, corroborated by their prevalence on jQuery-related textbooks [69]. All thematic observations undergo a similar confirmatory process to ascertain our observations are correct. Refer to our replication package [17] for the full taxonomy<sup>9</sup> and the resulting sunburst plots per country<sup>10</sup>. All calculations were done using the Gensim<sup>11</sup> and pyLDAvis<sup>12</sup> Python packages, owing to their robust and scalable implementations [59].

## 5.6 Users’ Knowledge Exchange (RQ4)

To answer RQ4, we conduct directed content analysis (CA) to investigate actual user-generated content. This RQ shifts the focus to the non-technical domain, investigating the knowledge exchange practices among the top contributors. Prior quantitative findings are thus supplemented by adding a qualitative standpoint. Our study builds upon previous work [39], employing a classification protocol based on Licorish’s framework [37], which includes 13 coding themes refined through rigorous iterations [70, 71], seen in Table 6. This coding schema has been consistently proven to yield good insights [72]. While previous work [39] used a similar approach (including studies authored by Licorish and MacDonell [73, 74]), findings in that study are primarily focussed on gender-related behavioural differences and did not incorporate geographical locations, which we aim to explore in the current study.

Table 6. Descriptions and examples of coding themes

Scale	Coding Theme	Description	Example
1	Type I Question	Requests solution or answer due to a knowledge deficit.	Which class contains the implementation and deployment for screen hibernation feature?
2	Type II Question	Initiates discussion and starts a dialogue.	Let us talk about the new python method that discards the unique index and has even weirder side effects in more complicated cases.
3	Answer	Provides relevant answers for information-seeking questions.	You can find hibernation features implemented across these classes (HMC1, CYMH, DHH and 3HC), I would suggest that you follow a similar approach.
4	Information exchange	Shares relevant information.	You do not need to do that since the API team were able to crack down issue number 315 yesterday.
5	Discussion	Provides additional context that expresses ideas or thoughts.	Solving issue #138 helped to solve the error produced in the ( <i>field_automation</i> ) class since it took care of all refactoring problems
6	Comment	Provides relevant statements that does not directly answer any inquiry.	I highly believe that test should be implemented first using a test-driven approach where tests fail at the start.
7	Reflection	Provides an appraisal, self-evaluation, or convey personal experiences.	I have noticed that the MVC framework from last year project can be applied to the current one besides including the useful techniques learnt in that challenging project.

<sup>9</sup> Replication package » Figures » Topics Taxonomy.png

<sup>10</sup> Replication package » Figures » Topics Sunburst Plots

<sup>11</sup> <https://radimrehurek.com/gensim/index.html>

<sup>12</sup> <https://pypi.org/project/pyLDAvis>

Scale	Coding Theme	Description	Example
8	Scaffolding	Proposes advice and implementation details to others.	I think it is a better idea to use clear and elaborative comments when we code to help in the final production of the documentation file.
9	Instruction/Command	Provides directive statements.	Fix or delete your posted answer since your provided pattern is neither a mixture of regex nor like-clause.
10	Gratitude	Provides appreciation, thankfulness, or praise.	Nice, your solution actually worked, thanks for the post.
11	Off task	Provides an unrelated transmission of messages regarding the current task or post.	Admin it has been a while.
12	Apology	Expresses remorse and regret.	I do apologise for posting a question with the wrong...
13	Not Coded	Communications that cannot be assigned to any of the preceding twelve categories.	N/A

We consider this set of coding themes suitable for our RQ as our study examines knowledge-sharing and information dissemination on Stack Overflow, while also aligning with previous literature [39]. Prior to sampling, we firstly define three units of analysis: the sampling unit, recording unit, and context unit [75]. The sampling unit refers to each individual data instance that is chosen for examination. The recording unit is the text component that is subjected to categorisation during the coding phase, whereas the context unit denotes the text necessary to establish additional context for the recording unit [76]. Firstly, the sampling unit was determined to be individual answers. Answers are considered the main information sources in CQA sites, offering a richer knowledge pool than other textual content forms (e.g., comments or users' About Me sections) [50]. Mining knowledge from these primary means for information dissemination allows us to gain insights into the site's main purpose [77]. Our recording unit is thus sentences, following other studies with similar design [39]. Finally, we designated question threads to be the context unit in our study, as they infer additional meaning of the recording unit [76]. This approach provides more content relevance, which may be otherwise impalpable if the context unit was narrower (e.g., only the corresponding paragraph). There may also be several answers referring to other prior answerers (e.g., "*There are some good answers already. I'll focus mainly on what I think they lack ...*"), where we may only understand its full nuance if we inspect the full question thread.

Probability sampling was done considering the extensive nature of the content being too extensive to be analysed as a population [78]. To ensure comparability across results, we employed random sampling based on Cochran's formula and sampled 385 posts for each country. Yielding a total of 2,310 posts, we tokenised sentences using the NLTK<sup>13</sup> Python package, which presented a sum of 6,201 sentences. To assess inter-coder reliability, we harness Cohen's Kappa ( $\kappa$ ) coefficient [79]. Prior to coding, coders undergo two iterations of a pilot reliability test to calibrate inter-coder consistency, which would be done on the initial 10% of all samples [80, 81]. Our first iteration yielded a  $\kappa$  of 0.754, while the second iteration produced a  $\kappa$  of 0.809. Afterwards, both coders divided the remaining samples equally, thereby presenting a final  $\kappa$  value of 0.852 for the entire set of samples. This signifies strong agreement between both coders [79]. Our replication package contains all reliability test results<sup>14</sup>. Coding results were then analysed with chi-squared test of homogeneity, allowing us to observe whether thematic prevalence is generally the same across regions. In the event where the null hypothesis would be rejected, our plan entails calculating Pearson's standardised residuals to find the source of significance [82]. Finally, we synthesise our findings with frequency counts and percentages for each coding theme. These series of measures allow us to effectively identify and quantify the degree to which countries diverge in their adherence to each of the 13 coding themes.

## 6 EXPERIMENT RESULTS

<sup>13</sup> <https://www.nltk.org>

<sup>14</sup> Replication package » Results » RQ4 - Content Analysis » Reliability Calibration

## 6.1 Quantitative Results

### 6.1.1 Users' Contribution Levels (RQ1)

Results for key platform variables are documented in Table 7. Firstly, we observe **differing males-to-females ratios** (last column). A higher ratio signifies a larger male population relative to females. Among the countries analysed, Brazil exhibited the highest males-to-females ratio, with a mean of 6.86, hinting that there are approximately 6.86 males for every female. The United Kingdom followed with an average of 5.65 males per female, while Australia presented a ratio of 5.21. India and Egypt had similar ratios of 4.93 and 4.92, respectively. The United States had the lowest males-to-females ratio, with an average of 4.84, suggesting a relatively more balanced gender distribution.

Table 7. Countries' contribution levels.

Country Name	<i>ProfileLength</i>	<i>YearlyDurationUsage</i>	<i>UpVotes</i>	<i>DownVotes</i>	<i>Views</i>	<i>Reputation</i>	<i>Questions</i>	<i>Answers</i>	<i>Comments</i>	<i>Edits</i>	<i>Badges</i>	<i>Males-to-females</i>
Australia	24.67	2.95	70.22	16.35	97.34	630.14	9.73	20.05	27.31	54.46	10.44	5.21
Brazil	6.50	0.23	50.18	8.11	44.82	187.84	7.38	12.14	10.67	18.92	5.25	6.86
Egypt	17.51	2.46	41.21	18.50	35.95	74.55	8.69	11.65	4.29	33.63	2.88	4.92
India	43.56	0.32	39.72	16.52	73.52	251.37	13.26	20.67	6.40	42.74	5.41	4.93
United Kingdom	18.15	0.49	70.75	18.01	135.01	847.23	9.35	23.22	35.58	64.27	13.53	5.65
United States	23.96	1.64	50.20	5.62	191.27	1420.22	6.07	41.53	18.81	55.04	24.15	4.84

India stands out with the **longthiest profile bios**, averaging 43.56 words within the *ProfileLength* variable, suggesting a preference for detailed self-descriptions. Conversely, users of Brazil were found to favour concise introductions, averaging just 6.5 words. **Engagement duration** with respect to average *YearlyDurationUsage* also reveals regional disparities. Users from Australia have the highest average at nearly 2.95 years, while their counterparts in Brazil exhibit the shortest average duration of just 0.23 years. Results potentially indicate a deeper commitment and sustained interest among the Australian user base. **Voting patterns** show further regional variations. While users from the United Kingdom and Australia lead in upvoting activity at 70.75 and 70.22 *UpVotes* cast, respectively, those residing in India and Egypt demonstrated a more 'neutral' approach – casting the least *UpVotes* at 39.72 and 41.21, respectively. Those from the United States retained the lowest average of 5.62 *DownVotes*. Interestingly, users from Egypt exhibit the highest downvoting activity at 18.5 *DownVotes* on average, which warrants further exploration to understand the potential underlying motives.

Differences across regions were also found in terms **user visibility** and **overall engagement**. Users from the United States lead in profile views, attracting an average of 191.27 *Views* compared to users from other countries like Egypt, with only 35.95 *Views* on average. Results suggest a stronger personal branding and profile visibility among the former. This trend is mirrored in **reputation**, with users from the United States holding the highest average *Reputation* score (1,420.22) compared to other countries such as Brazil at 187.84 and India at 251.37, potentially indicating that users from the United States are generally more established and accrue more renown in the community. Interestingly, while India displays the highest average number of *Questions* asked (13.26), signifying a **more inquisitive** user base, users from the United States provide the most *Answers* on average (41.53), suggesting more active responders. This juxtaposition hints that contextual factors across continents shape **user engagement** not only in terms of information consumption (as evidenced by the higher average *Questions* in India) but also in their inclination to actively and **voluntarily share knowledge** (as seen in the higher *Answers* provided by users of United States).

Analysing **degree of engagement** through *Comments* and *Edits* reveals further regional differences. Users from the United Kingdom were found to leave the most *Comments*, averaging 35.58, closely followed by those from Australia at 27.31. Findings could be indicative of a preference for discussion-based learning among these users. Conversely, users from India and Egypt tend to leave fewer *Comments*, averaging 6.4 and 4.29, respectively. Such figures might suggest a preference for other forms of engagement. Users from the United States, despite falling within the middle of the *Comments* spectrum (18.81), tend to lead in editing activity, averaging 55.04 *Edits* per user. Results may signify a focus on **content refinement** to uphold the site’s standard within the United States user base. Trends largely propagate with respect to *Badges*, where users from the United States earned 24.15 Badges on average. Results are indicative of a **higher level of achievement and expertise** within the United States.

**Finding 1:** Our analysis sheds light on a rich tapestry of regional disparities in contribution. From gender disparity in Egypt, content curation patterns (*UpVotes* in the United Kingdom, *DownVotes* in Egypt) to active content generation (*Questions* in India, *Edits* in the United States), geographical aspects were found to influence how users interact with the platform.

Our Kruskal-Wallis test revealed statistically significant differences ( $p < 0.05$ ) in all platform variables across countries, with the exception of *ProfileLength* ( $p = 0.275$ ). This suggests that, while there are significant variations in platform usage patterns between countries, ***ProfileLength* variable remains largely consistent**. Such a finding suggests a universal tendency among users to maintain similar levels of detail in their profile bios, regardless of their geographical location. As outlined in Section 5.3, Conover-Iman post-hoc tests with Bonferroni corrections were subsequently conducted for platform variables yielding  $p \leq 0.05$ . Outcomes of all pairwise comparisons yielded p-values that were less than 0.05, indicating statistically significant differences between the groups being compared. Further results for our may be seen in our accompanying replication package [17]<sup>15</sup>.

### 6.1.2 Users’ Interaction Patterns (RQ2)

We define users’ interactions as the **aggregate sum of answers given and received** for each user. Subsequently, the top 100 users with the highest total interactions were identified per country, allowing us to render their respective interaction graphs. Summary statistics for each graph is seen in Table 8. Total interactions for all users within each country’s top 100 subset is provided in our replication package [17]<sup>16</sup>.

Table 8. Summary statistics for network graphs.

Graph Property	Country					
	AU	BR	EG	IN	GB	US
Nodes	88	27	16	97	98	99
Edges	195	17	9	288	548	898
Avg. Polarity	0.219	0.156	0.288	0.217	0.241	0.264
Density [83]	0.025	0.024	0.037	0.031	0.057	0.093
Clustering coef. [84]	0.072	0	0	0.128	0.105	0.14
Transitivity [85]	0.045	0	0	0.136	0.060	0.139

An examination of network size reveals that **graphs of the United States and the United Kingdom possess the greatest number of nodes and edges**, suggesting the existence of the most interconnected social network. The former presented 99 nodes connected with 898 edges, whereas the latter yield 98 nodes connected with 548 edges. Compared to the aforementioned two countries, **Australia and India exhibit fewer nodes and edges** – albeit still demonstrating relatively large numbers. Australia presents 88 nodes with 195 edges, wherein India retained 97 nodes and 288 edges. This implies that, while not as extensive as the interactions within the United States and the United Kingdom, the intra-country engagements in these two countries remain noteworthy. Finally, it is evident that **Egypt and Brazil exhibit the least number of nodes and edges**, indicating a lack of substantial intra-country interactions in both nations, even among their top 100 contributors. The former consisted of 16 nodes, but only 9

<sup>15</sup> Replication package » Results » RQ1 - Conover-Iman Test

<sup>16</sup> Replication package » Results » RQ2 - Users’ Total Interactions

edges were present. Moreover, the latter network had 27 nodes, yet only 17 edges. Both exemplifies a relatively sparse connection pattern. Patterns are similarly evident in terms of density, with the United States and United Kingdom exhibiting the highest density, while Egypt and Brazil display the lowest.

We can also discern that **the average polarity values for all of the graphs are positive**, with the smallest by Brazil at 0.156, and the largest by the United States with 0.264. Findings suggest a **general trend towards positive sentiment among the top contributors** within the same country. In terms of clustering coefficients, the United States and India demonstrate the highest values at 0.14 and 0.128, respectively. This suggests a higher likelihood of nodes in these graphs interacting with other nodes that have previously engaged with each other. Conversely, coefficients were zero for both Brazil and Egypt, indicating that no local clustering were apparent within these two countries. This is also true for transitivity wherein Brazil and Egypt both demonstrated zeros. India's transitivity of 0.136 stands out as the highest among the analysed networks, indicating a strong propensity for triadic closure. This implies that when two nodes are connected to a third node, they are also highly likely to be directly connected to each other, forming a closed triad. In the interest of conciseness, this section solely presents social network visualisations for the United States (Figure 2) and India (Figure 3). This choice stems from the observation that social networks in other countries exhibited similarity to these two focal points. Appendix A provides visualisations of the social networks for the other four countries.




Figure 2. Social network for United States top contributors.

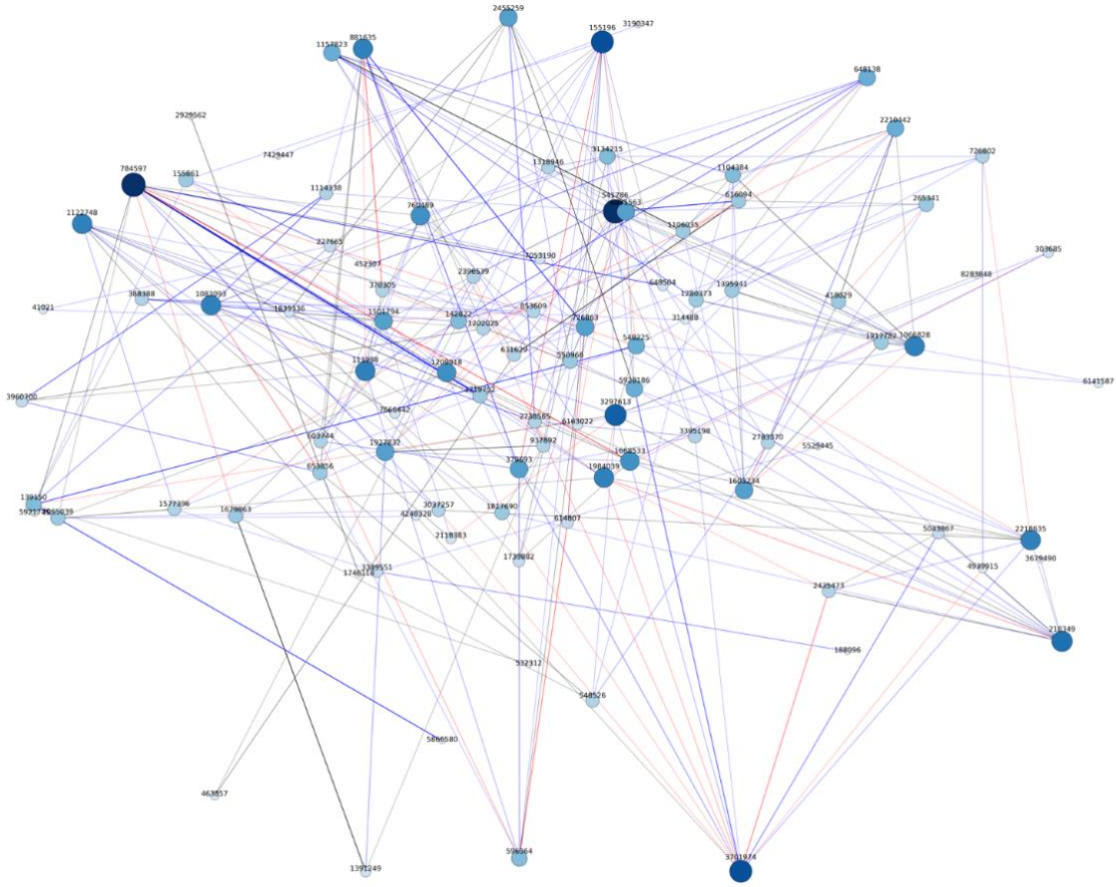


Figure 3. Social network for India top contributors.

As noted in Section 5.4, the presence of reciprocal edges between two nodes does not indicate continuous dialogue between the individuals involved. Instead, these edges signify distinct communications where one does not overlap with another. From Figure 2 depicting the social network within the United States, we observe Users.Id 4653, 117700, 39677, and 34397 are the most popular in the graph, with the largest node sizes and darkest colours. These users likely represent **the most active participants within the network**. Additionally, a dense cluster of users in the central region suggests the presence of tightly-knit communities with frequent interactions between members [86]. Further investigation reveals that two representative members (Users.Id 9034 and 1144035) from this cluster consistently provide answers related to SQL and databases, suggesting a potential community of experts. Notably, **the majority of edges (blue) indicate positive interactions**, while only a small number of negative and neutral interactions (red and black) are evident.

Figure 3, showing India's rendered network, exhibits several different patterns. Users.Id 784597, 541786, and 3701974 stand out due to their large node size and dark colour, indicating popularity. Interestingly, **interactions directed towards Users.Id 3701974 appear predominantly negative** (red edges), suggesting potential animosity from their own countrymen. Conversely, Users.Id 6141587, 41021 and 463857 **exhibit less engagement** (smaller, lighter nodes), hinting at a lower level of participation in both providing and receiving answers. Unlike the network for the United States, **the social network of India lacks distinct clusters**, suggesting minimal community formation.

We observe the **emergence of a centralised network structure for the United Kingdom**. Users.Id 22656, 23354, 157247, and 19068 occupy central positions, acting as hubs for interactions. However, the network exhibits a core-periphery structure that depicts relatively even distribution of activity beyond these central figures. For example, Users.Id 2287576 and 439688 (smaller, lighter nodes) demonstrate considerably lower interaction levels compared to the central hub. Findings suggest a less prominent central group and a more balanced distribution of engagement. The social network of **Australia exhibits a core-periphery pattern similar to that of the United Kingdom**. A central cluster



of highly active users dominates the network, evident by their larger, darker nodes and the abundance of thick blue edges connecting them. Notable examples include Users.Id 16076, 256196, 30674, and 139985. Observations suggest frequent positive interactions within this tightly-knit community. However, there were also smaller, less prominent groups with frequent internal interactions. Interestingly, the presence of red edges, particularly between users on the network’s periphery, suggests instances of negativity. A notable example is the thick red edge from Users.Id 174777 towards Users.Id 5421539, highlighting a potentially negative interaction despite the significant distance between their nodes.

The network for **Brazil exhibits a significantly sparser structure** compared to the other countries. Patterns suggest a weaker interconnectedness even among the top contributors, where social influence is distributed more evenly across the network. The maximum observed interaction for a single user is just three, with Users.Id 131874 holding this distinction. Notably, the abundance of black edges indicates a high prevalence of neutral interactions among the top contributors of Brazil. Finally, the network for **Egypt displays the least density of connections**. Despite the sparsity, edges of the network are predominantly blue, indicating a positive overall sentiment. The limited number of nodes further suggests that these top contributors primarily engage with individuals outside this subset, which may entail users from other countries. Network visualisations are also included in our replication package [17] to allow a more detailed inspection<sup>17</sup>.

**Finding 2:** SNA reveals strong regional differences in user interactivity. The United States and the United Kingdom have the most densely connected networks, while Egypt and Brazil exhibit the least. Network for the United States shows a central community of experts, those of the United Kingdom and Australia shows a core-periphery structure, while that of India suggests isolated negativity towards certain users.

### 6.1.3 Users’ Topics of Discourse (RQ3)

To determine the optimal number of topics ( $k$ ) for each country, we conducted experiments with  $k$  values ranging from 2 to 100. The results of these experiments revealed that the optimal  $k$  value for each country varied. Subsequently, for each country, we selected  $k$  that yielded the highest absolute difference between *Topic Coherence* and *Jaccard similarity index* (see Section 5.4), presented in Table 9. Full results of our experiments may be seen in our replication package [17]<sup>18</sup>.

Table 9. Optimal  $K$  per country.

Country	$K$
Australia	32
Brazil	28
Egypt	26
India	35
United Kingdom	42
United States	41

Topic modelling was then performed for each country using the optimal  $k$  identified in Table 9, yielding a total of 204 distinct topics. For each country, the top 10 topics that accounted for the most tokens are presented, and the remaining 144 topics may be found in the replication package [17]<sup>19</sup>. We present the top 10 most salient keywords (see Eq. 2 in Section 5.5), in accordance to literature [87]. Our subsequent thematic analysis commenced with an exploration of potential areas of interest within the data, where initial observations are documented in Table 10.

<sup>17</sup> Replication package » Figures » Social Networks

<sup>18</sup> Replication package » Results » RQ3 - LDA » Overlap-Coherence Results

<sup>19</sup> Replication package » Results » RQ3 - LDA » Final LDA Results.xlsx

Table 10. LDA results for all countries

Country	Topic No.	% of Tokens	Thematic Observation	Top 10 Most Salient Terms
Australia	1	6.7%	Performance Suggestions	performance, time, code, good, thing, would, lot, much, problem, better
	2	5.1%	Swing Painting (GUI)	swing, paint, painting, component, timer, ui, context, paintcomponent, graphic, detail
	3	4.4%	Form Templates	page, work, following, form, php, code, fine, tried, jquery, working
	4	4.7%	Query Statements	table, query, column, row, sql, database, join, record, transaction, data
	5	4.4%	Graphical User Interface	view, button, model, event, click, user, action, control, dialog, menu
	6	4.1%	Array Data Structures	like, list, result, output, filter, item, something, want, try, loop
	7	4.0%	Object-Oriented Programming	method, call, interface, code, foo, implement, like, class, signature, parameter
	8	3.8%	ASP.NET Web Framework	web, database, url, site, server, asp, website, domain, page, net
	9	9.0%	Concurrent Programming	thread, task, lock, async, await, wait, reactive, observable, safe, call
	10	3.6%	C-family Error Handling	class, exception, static, constructor, throw, catch, field, thrown, instance, declared
Brazil	1	8.4%	Functional Programming	function, language, would, use, way, thing, much, javascript, good, algorithm
	2	5.9%	Error Handling	error, code, message, type, json, following, trying, get, getting, wrong
	3	5.1%	Python Packages	package, window, install, run, installed, python, linux, machine, version, py
	4	5.1%	User Authentication	server, app, user, connection, service, password, docker, application, token, build
	5	4.6%	Query Statements	table, query, database, id, model, sql, user, data, insert, mysql
	6	4.4%	AWK GNU	awk, gnu, sed, r, char, multi, regexp, gawk, string, unix
	7	4.2%	Array Data Structures	array, string, value, number, return, index, integer, key, null, int
	8	4.0%	Android Development	android, app, thread, activity, widget, screen, fragment, crash, phone, layout
	9	4.0%	Styling and Layout	text, answer, left, font, problem, solved, div, width, title, solution
	10	3.9%	Graphical User Interface	event, button, form, click, page, jquery, browser, chrome, firefox, ajax
Egypt	1	6.9%	Error Handling	error, trying, got, tried, following, run, code, problem, using, getting
	2	6.3%	Array Data Structures	array, function, pointer, variable, element, object, address, value, loop, memory
	3	4.7%	Query Statements	query, like, select, join, sql, group, clause, use, count, give
	4	4.7%	Relational Databases	table, key, try, nan, output, database, entity, relation, product, two
	5	4.5%	Image Display	image, parent, child, like, code, screen, look, want, size, following
	6	4.5%	String Formatting	string, char, character, int, integer, byte, null, format, literal, value
	7	4.5%	User Login Page	user, app, facebook, php, password, application, login, post, website, account
	8	4.4%	Microsoft Visual Studio	net, window, application, dll, visual, system, studio, microsoft, framework, assembly
	9	4.4%	PHP Sessions	page, request, browser, session, ajax, url, redirect, load, javascript, post
	10	4.0%	Graphical User Interface	view, form, item, field, menu, controller, list, bind, model, adapter
India	1	5.6%	Performance Suggestions	would, good, answer, question, much, memory, performance, best, better, way
	2	4.8%	jQuery	element, event, attribute, jquery, selector, tag, div, use, child, html
	3	4.5%	Pointer Behaviour	type, pointer, int, undefined, behavior, memory, char, variable, compiler, chapter
	4	4.2%	ASP.NET Web Pages	page, user, form, button, control, net, login, click, asp, show

Country	Topic No.	% of Tokens	Thematic Observation	Top 10 Most Salient Terms
	5	3.8%	Error Handling	error, getting, issue, trying, exception, code, tried, log, throw, missing
	6	3.8%	Styling and Layout	style, color, height, width, screen, cell, view, bar, display, scroll
	7	3.8%	Object-Oriented Programming	class, object, instance, constructor, property, create, static, method, interface, reference
	8	3.6%	Functional Programming	value, date, true, number, false, condition, return, format, operator, expression
	9	3.4%	Regular Expressions	regex, match, demo, group, character, lookahead, use, pattern, capture, negative
	10	3.4%	Query Statements	table, query, column, join, row, record, sql, clause, entity, insert
United Kingdom	1	6.3%	Question-Answering Practices	good, question, answer, people, thing, lot, idea, bad, think, really
	2	5.6%	LINQ Preference	would, could, use, linq, want, really, using, write, way, simpler
	3	3.9%	Directory Traversal	file, project, directory, path, command, package, build, folder, install, studio
	4	3.6%	Error Handling	error, code, problem, delphi, warning, bug, fix, issue, message, debugger
	5	3.3%	Object-Oriented Programming	class, method, instance, constructor, static, interface, base, member, private, subclass
	6	3.2%	Loops and Control Flow	list, loop, key, dictionary, item, collection, map, generator, comprehension, tuple
	7	3.1%	Variables and Scope	object, property, variable, reference, value, prototype, assign, new, assigned, copy
	8	2.9%	jQuery	element, id, attribute, selector, jquery, dom, use, div, select, try
	9	2.9%	Performance Suggestions	performance, time, faster, difference, cost, cache, fast, slow, cpu, overhead
	10	2.8%	Functions	function, call, scope, global, variable, called, var, closure, inside, setTimeout
United States	1	6.9%	Version Control	git, commit, branch, commits, name, merge, hash, repository, master, id
	2	6.2%	Performance Suggestions	much, performance, faster, lot, better, good, time, thing, even, le
	3	4.4%	Question-Answering Practices	want, question, think, would, answer, way, really, might, probably, one
	4	3.9%	Query Statements	row, column, table, join, query, clause, key, subquery, select, unique
	5	3.6%	.NET Framework	project, version, library, net, build, framework, studio, visual, feature, tool
	6	3.4%	JSON Parsing	like, something, look, json, try, code, work, sound, would, could
	7	3.4%	C-family Compilers	compiler, standard, implementation, language, template, header, defined, declaration, definition, compile
	8	3.1%	Error Handling	error, python, exception, module, code, catch, throw, problem, import, fix
	9	3.1%	Variables and Scope	object, variable, property, reference, name, value, scope, global, assign, local
	10	2.9%	Memory Allocation	memory, stack, dll, allocated, assembly, code, garbage, debugger, heap, allocation

Next, we engaged in an iterative process of reviewing the relationships between topics and their associated keywords, traversing the entire set of 204 topics multiple times before tallying the results. To refine the topic categorisation and establish the hierarchical structure of the extracted information, multiple meetings were held for in-depth discussions. Our process captured three top-level categories across all countries: *Applied Computer Science*, *Theoretical Computer Science*, and *Computer Systems*. Each category encompasses a collection of subcategories, which in turn may further branch into sub-subcategories. These, in turn, give rise to specific topics, with some topics even branching into sub-topics. For example, as illustrated in Figure 4, the *Computer Systems* category consisted of several sub-categories, including *Computer Architecture*. The latter is then branched into several sub-subcategories including *Memory Operations*, where one of the main topics of discourse is *I/O Operations*. Under it, *Serialisation* becomes a specific focus. Refer to our replication package [17] for the full taxonomy<sup>20</sup>.

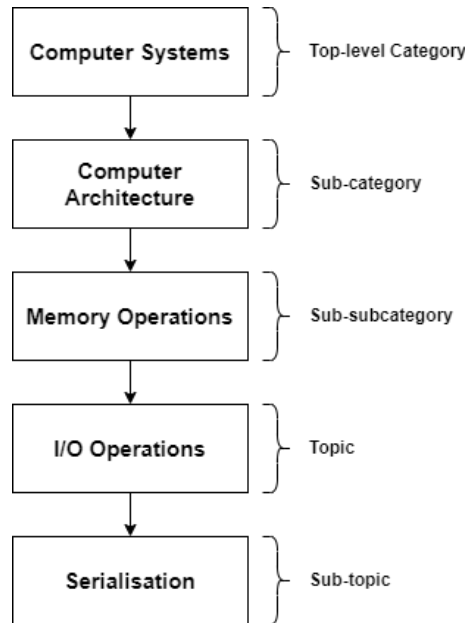


Figure 4. Example topic structure.

The observed hierarchical structure can be attributed to the varying levels of specificity with which different countries explore a particular category. While some countries may focus on broader aspects of the topic, others may delve into more granular details. To extend our example in Figure 4, we further observe that the top contributors from the United States tend to engage in more general discussions surrounding *I/O Operations*, while those from Australia exhibit a tendency to delve specifically into *Serialisation*. Across all categories, our analysis identified a total of 11 unique subcategories, 39 sub-subcategories, 66 topics, 40 subtopics, and 2 sub-subtopics. This hierarchical structure encompasses a broad spectrum of concepts related to computer science. Table 11 depicts the sub-category distribution per country, and Figure 5 serves as a visualisation to provide additional nuance.

Table 11. Sub-category distribution per country

Category	Sub-category	AU	BR	EG	IN	GB	US	Total
Theoretical Computer Science	Data Structures	2	1	1	1	2	2	9
	Algorithms	1	1	0	1	2	1	6
Computer Systems	Database Operations	1	1	2	2	2	3	11
	Cybersecurity	1	1	0	0	1	1	4
	Computer Architecture	4	1	0	0	4	5	14
Applied Computer Science	Web Development	0	1	4	8	5	2	20
	Third-Party Resources	3	2	1	1	0	2	9
	Software Engineering	10	10	9	11	13	14	67
	Programming Paradigms	2	2	1	4	3	2	14
	Human-Computer Interaction	6	5	5	5	6	4	31

<sup>20</sup> Replication package » Figures » Topics Taxonomy.png

Category	Sub-category	AU	BR	EG	IN	GB	US	Total
	Data Science	2	3	3	2	4	5	19




Figure 5. Bar chart of sub-category distribution per country.

From Table 11, we observe that **topics that fall under *Software Engineering* are among the most widely-discussed**, occurring a total of 67 times across all six countries. Discourses pertaining to ***Human-Computer Interaction* occupy the second position**, with 31 topics. However, our manual analysis reveals that topics within this sub-category primarily represent instances of fellow users exchanging solutions. Topics related to ***Web Development* were found 20 times**, whereas the ones related to ***Data Science* follows closely** in 19 instances. On the other hand, our thematic analysis revealed only 14 topics related to ***Computer Architecture*** and another 14 related to ***Programming Paradigms***. Next, we found 11 topics that fall under the ***Database Operations*** sub-category, and 9 topics for both ***Data Structures*** and ***Third-Party Resources***. The **least topics were found pertaining to *Algorithms* and *Cybersecurity***, yielding only 6 topics for the former and 4 for the latter.

**Finding 3:** Thematic analysis results from LDA reveal variations across different specialisations within computer science. Notably, ***Software Engineering*** and ***Human-Computer Interaction*** are among the most discussed subcategories, whilst ***Cybersecurity*** and ***Algorithms*** were found to be among the least.

## 6.2 Qualitative Results

### 6.2.1 Users' Knowledge Exchange (RQ4)

Results from our CA were tallied by each country, allowing us to analyse them using frequency counts and percentages. Such an approach enabled us to identify the most and least prevalent categories of behaviour for each country. We document our quantitative results in Table 12, whereas visualisations per country may be seen in our replication package [17]<sup>21</sup>.

Table 12. CA results per country.

Category	AU	BR	EG	IN	GB	US	Total
Type I Question	49 (1.84%)	11 (0.45%)	8 (0.33%)	6 (0.27%)	42 (1.13%)	46 (1.4%)	162
Type II Question	62 (2.33%)	96 (3.93%)	162 (6.72%)	12 (0.53%)	49 (1.32%)	45 (1.37%)	426
Answer	826 (31.08%)	390 (15.96%)	449 (18.62%)	322 (14.24%)	803 (21.59%)	737 (22.4%)	3,527
Information Exchange	716 (26.94%)	349 (14.28%)	408 (16.92%)	336 (14.86%)	748 (20.11%)	736 (22.37%)	3,293
Discussion	158 (5.94%)	465 (19.03%)	558 (23.13%)	373 (16.5%)	394 (10.59%)	346 (10.52%)	2,294
Comment	210 (7.9%)	527 (21.56%)	292 (12.11%)	444 (19.64%)	334 (8.98%)	287 (8.72%)	2,094
Reflection	58 (2.18%)	208 (8.51%)	111 (4.6%)	275 (12.16%)	158 (4.25%)	130 (3.95%)	940
Scaffolding	274 (10.31%)	180 (7.36%)	148 (6.14%)	185 (8.18%)	628 (16.89%)	273 (8.3%)	1,688
Instruction/Command	272 (10.23%)	189 (7.73%)	216 (8.96%)	280 (12.38%)	544 (14.63%)	659 (20.03%)	2,160
Gratitude	10 (0.38%)	17 (0.7%)	29 (1.2%)	11 (0.49%)	1 (0.03%)	3 (0.09%)	71
Off Task	5 (0.19%)	1 (0.04%)	0 (0%)	1 (0.04%)	4 (0.11%)	3 (0.09%)	14
Apology	2 (0.08%)	1 (0.04%)	8 (0.33%)	0 (0%)	2 (0.05%)	4 (0.12%)	17
Not Coded	16 (0.6%)	10 (0.41%)	23 (0.95%)	16 (0.71%)	12 (0.32%)	21 (0.64%)	98

Across all countries, **Answers was the most-occurring theme**, which had occurred 3,527 times. Initial findings demonstrated that users effectively addressed the key aspects of the given inquiry. This was **closely followed by Information Exchange** which was present in 3,293 times **and Discussion** in 2,294. Notably, users also exhibited a tendency to provide directives within their answers, as evidenced by the 2,160 occurrences of the *Instruction/Command* theme. Furthermore, the coding theme *Comment* was flagged 2,094 times, while *Scaffolding* appeared 1,688 times. This suggests a propensity for users to leave comments and provide suggestions – albeit not as extensive as their tendency to discuss and exchange information. Additionally, *Reflection* occurred 940 times, and *Type II Question* occurred 426 times, hinting users' subtle tendency to convey appraisal and incite discussions within question threads. **The low frequency of Type I Question** (162 times) suggests that users are generally less inclined to pose questions within their answers. However, it is worth noting that sparse instances of questioning behaviour were still observed. *Not Coded* occurred 98 times. **Themes associated with socio-emotional expressions occur only a few times**, such as *Gratitude* (71 occurrences) and *Apology* (17 occurrences). This suggests a tendency among users to engage in less emotionally charged interactions. Finally, *Off Task* occurred 14 times.

With respect to the category *Type I Question*, **users from Australia were the most likely to ask questions within their answer posts** (they asked 49 questions (1.84%)). United States was next in line with 46 occurrences (1.4%), then United Kingdom with 42 (1.13%). **Users from Brazil, Egypt, and India were much less likely to ask questions within their answer posts** (refer to Table 12). Chi-squared test exhibit statistically significant results ( $\chi^2 = 46.98$ ,  $p < 0.001$ ), hinting that these disparities

<sup>21</sup> Replication package » Figures » Content Analysis

in question-asking behaviour is not due to chance. In terms of *Type II Question*, users from Egypt exhibited the highest tendency to incite discussions with other users, doing so 162 times (6.72%). This was followed by users of Brazil at 96 times (3.93%) and Australia at 62 times (2.33%). On the other hand, those from India displayed the least likelihood in this regard (asking indirect questions 12 times at 0.53%). Our chi-squared results confirm that difference noted was statistically significant ( $\chi^2 = 302.41, p < 0.001$ ).

**Users from Australia were also the most likely to directly address the inquiry at hand**, answering a total of 826 times (31.08%) followed by those from the United States with 737 times (22.4%). On the other hand, **those from Brazil gave direct answers only on 390 instances** (15.96%), and users from India answered even less at 322 (14.24%). Chi-squared results for *Answer* indicates that such fluctuations were statistically significant ( $\chi^2 = 264.99, p < 0.001$ ). Patterns also propagate in terms of *Information Exchange*, where the strongest inclination was exhibited by those from Australia, tallying a total of 716 times (26.94%). Users from the United States and United Kingdom followed suit at 736 (22.37%) and 748 (20.11%), respectively. Such tendencies are less pronounced, however, in India, where they were the least likely to share information (doing so 336 times at 14.86%) within their answer posts. Differences in *Information Exchange* were found to be statistically significant ( $\chi^2 = 226.30, p < 0.001$ ). **Users from Egypt were more likely to provide additional context** that expresses ideas or thoughts, doing so 558 times (23.13%). Brazil followed in second, doing so at 558 occurrences (23.13%) and India at 444 (19.64%). The *Discussion* coding theme were less prevalent within users from United States, United Kingdom (refer to Table 12), and even less in Australia, where they exhibiting such a behaviour only 158 times (5.94%). Findings **suggest a more reserved approach to expressing their opinions**. Chi-squared test of the *Discussion* coding theme affirmed that differences are statistically significant ( $\chi^2 = 800.04, p < 0.001$ ).

Users from Brazil (providing *Comments* 527 times at 21.56%) and India (444 times at 19.64%) **are the most likely to give statements that does not directly answer the inquiry in context**. Table 12 also tells us that users of Egypt and Australia exhibit this tendency to a lesser extent. Interestingly, users from the United Kingdom (334 times; 8.98%) and United States (287; 8.72%) were **the least likely to engage in this type of interaction**. This suggests a tendency for users in both countries to prioritise direct relevant responses to the posed question. Chi-squared test on the *Comment* coding category revealed statistically significant differences ( $\chi^2 = 634.44, p < 0.001$ ). In addition, those from India exhibit **a notable inclination towards incorporating appraisal, self-evaluation, or personal experiences** within their answers, doing so on 275 occasions (12.16%), where Brazil followed suit on 208 occurrences (8.51%). Users from Egypt, the United Kingdom, and the United States demonstrate this tendency to a lesser extent, seen in Table 12. However, users of Australia were the least likely to engage in self-reflection, doing so 58 times (2.18%). Chi-squared test on *Reflection* yielded statistically significant results ( $\chi^2 = 397.87, p < 0.001$ ).

**Users from the United Kingdom scaffolded a total of 628 times** (16.89%), followed by those from Australia with 274 times (10.31%). Though not as intense as the previous two countries, the United States, India, and Brazil still exhibit a notable tendency to provide suggestions, as seen in Table 12. **Users from Egypt were the least likely to scaffold others' content** – doing so 148 times which accounts for 6.14% of their coded exchanges. Chi-squared test for the *Scaffolding* coding theme produced statistically significant results ( $\chi^2 = 432.29, p < 0.001$ ). With respect to the *Instruction/Command* coding theme, **users from the United States were the most likely to provide directives** (659 occurrences at 20.03%). This was followed by United Kingdom at 544 occurrences (14.63%). Users from India, Australia, and Egypt exhibited this behaviour to a lesser extent (see Table 12), whereas those from Brazil were the least likely to give directives when answering questions, doing so 189 times (7.73%). Chi-squared test applied to the *Instruction/Command* coding theme yielded statistically significant results ( $\chi^2 = 539.16, p < 0.001$ ).

**Across all countries, users exhibited a remarkably low inclination towards expressing gratitude within their answer posts**. That said, users from Egypt (29 times, 1.2%) and India (11 times; 0.49%) exhibit a relatively higher proclivity for this behaviour. Those from Australia and Brazil did so to a lesser extent, whilst users from the United States (3 times; 0.09%) and United Kingdom (once; 0.03%) exhibit near-zero tendency to express gratitude within their answer posts. Chi-squared test for the

*Gratitude* coding theme revealed statistically significant results ( $\chi^2 = 66.54$ ,  $p < 0.001$ ). **Users from all countries were likely to refrain from providing responses unrelated to the specified inquiry.** However, those of Australia stood out for exhibiting this behaviour with the highest frequency (5 occurrences at 0.19%). The United States, the United Kingdom, India, and Brazil were less likely to provide unrelated responses, as seen in Table 12. Interestingly, users of Egypt did not display such behaviour at all. Chi-squared test for *Off Task* confirmed that difference was not statistically significant however ( $\chi^2 = 4.91$ ,  $p = 0.426$ ).

**The *Apology* coding theme revealed minimal occurrences across all countries**, suggesting a general reluctance among users to apologise within their answer posts. However, users of Egypt were found to make the most apologetic responses (8 times, 0.33%). This was immediately followed by the United States with 4 (0.12%). As outlined in Table 12, those from Australia, United Kingdom, and Brazil were less likely to do so. It is interesting to note that users of India did not convey any. Chi-squared results for *Apology* confirmed that the difference noted was statistically significant ( $\chi^2 = 19.06$ ,  $p < 0.05$ ). Finally, the *Not Coded* coding theme served as a catch-all for any behaviours that fell outside the scope of the other 12 predefined categories, or those not explicitly captured by the established coding framework [39]. Results indicate few codes were assigned to this category for all countries (i.e.,  $< 1\%$ ).

**Finding 4:** Generally, top contributors prioritise direct answers and information exchange, yet there are certain subtle differences across countries. For example, top contributors from Australia ask clarifying questions, those from Egypt use indirect language to spark discussions, whereas those from India tend to express personal experiences.

Calculations for Pearson's absolute adjusted residuals are seen in Table 13, where values higher than 3 were highlighted as red cells to convey pairs that exhibit statistically significant differences [88]. From Table 13, we can see that the standardised residuals for nearly all country pairs surpass the threshold of 3, implying that all countries tend to statistically differ between each other. For the *Type I Question* coding theme, the most significant difference was between Australia and India with 5.25 standardised residual, whereas Egypt and the United Kingdom were observed to be the pair that differs most with respect to *Type II Question*. In terms of *Answer*, Australia and India were observed to yield the most significant difference, presenting an absolute adjusted residual value of 13.92. For *Information Exchange*, the most substantial difference was observed between Australia and Brazil, with a standardised residual of 11.11. Notably, Australia also exhibited significant difference with Egypt in terms of *Discussion*, with a standardised residual of 17.55. In regards to *Comment*, the highest standardised residual was between Brazil and United Kingdom (13.94). For *Reflection*, however, the highest was between Australia and India with 13.88. Our subsequent findings focus on the Scaffolding coding theme, where the highest standardised residual was observed between Egypt and the United Kingdom (12.37). The pair Brazil and the United States presents the highest standardised residual within the *Instruction/Command* coding theme (12.97), whereas the pair Egypt and the United Kingdom leads *Gratitude* (6.44). As differences in the *Off-Task* coding theme was not statistically significant, no pairs exhibit standardised residuals greater than 3. Moreover, despite statistically significant fluctuations observed in the *Apology* coding theme, no country pair exhibited standardised residuals exceeding 3. The highest residual is 2.63 which was presented between Egypt and the United Kingdom. This intriguing pattern suggests the potential presence of Simpson's paradox, which will be further explored in Section 8.



Table 13. Pearson's absolute adjusted residuals for all country pairs.

Pair	Type I Question	Type II Question	Answer	Information Exchange	Discussion	Comment	Reflection	Scaffolding	Instruction/Command	Gratitude	Off Task	Apology	Not Coded
Australia - Brazil	4.612	3.286	12.662	11.113	14.257	13.867	10.158	3.689	3.112	1.571	1.533	0.505	0.966
Australia - Egypt	5.100	7.584	10.217	8.583	17.550	5.004	4.792	5.373	1.543	3.362	1.808	2.055	1.430
Australia - India	5.247	5.175	13.915	10.299	11.884	12.079	13.884	2.558	2.379	0.591	1.441	0.824	0.459
Australia - United Kingdom	2.371	3.056	8.562	6.389	6.511	1.523	4.497	7.432	5.179	3.314	0.845	0.338	1.663
Australia - United States	1.362	2.783	7.557	4.076	6.296	1.140	3.878	2.668	10.339	2.340	1.014	0.560	0.177
Brazil - Egypt	0.661	4.330	2.446	2.528	3.506	8.803	5.499	1.708	1.538	1.822	0.400	2.355	2.308
Brazil - India	1.055	7.775	1.644	0.561	2.269	1.635	4.121	1.045	5.318	0.932	0.055	0.361	1.379
Brazil - United Kingdom	2.825	6.614	5.477	5.856	9.350	13.938	6.925	10.834	8.179	4.759	0.899	0.224	0.557
Brazil - United States	3.579	6.190	6.072	7.735	9.144	13.776	7.249	1.295	12.972	3.839	0.713	1.023	1.170
Egypt - India	0.414	11.161	4.026	1.919	5.677	7.062	9.383	2.717	3.802	2.654	0.448	2.482	0.923
Egypt - United Kingdom	3.393	11.327	2.828	3.131	13.240	3.945	0.659	12.370	6.587	6.443	1.227	2.634	3.202
Egypt - United States	4.109	10.666	3.484	5.086	12.884	4.175	1.205	3.086	11.466	5.548	1.061	1.710	1.343
India - United Kingdom	3.631	2.937	7.056	5.116	6.616	11.875	11.449	9.526	2.444	3.851	0.822	0.581	2.114
India - United States	4.305	3.040	7.606	6.968	6.517	11.811	11.556	0.156	7.470	2.885	0.641	1.280	0.311
United Kingdom - United States	1.009	0.182	0.817	2.309	0.105	0.379	0.625	10.721	5.987	1.125	0.217	0.969	1.926

## 7 DISCUSSION

### 7.1 Users' Contribution Levels (RQ1)

Initial observations suggest that gender disparity is still prevalent in Stack Overflow, aligning with findings from previous works [5, 36, 89]. Specifically, male users were found to be more abundant than their female counterparts, synthesising the work by Vasilescu et al. [5]. Such disparities may be attributed to participation-related confounders such as reluctance to contribute, as well as discrepancies in rewards [89]. Delving into the gender disparities in countries, results indicate that in the user base of Brazil, there exists six to seven times more males for each female. We surmise this significant disparity is due to the underrepresentation of women in STEM education in Brazil, in which women constitute only around 30% of individuals enrolled in STEM-related programs [90]. Similar gender disparities are prevalent in the United Kingdom, Australia, India, and Egypt, reflected in their gender imbalance. Numerous research efforts have delved into the underlying causes of this gender disparity, with a substantial portion of the findings attributing it to factors related to STEM education [91]. Our findings corroborate the notion that the prevailing perception of SE as a male-dominated field [92] extends to programming-related CQAs such as Stack Overflow, influencing gender participation and engagement.

Regarding the analysed platform variables, higher *ProfileLength* is seen within India and Australia. Joglekar's [93] findings regarding the expressive and outgoing nature of individuals from India provide a potential explanation for this observed trend. This includes user willingness to provide details about their professional journey and personal interests, resulting in more extensive self-descriptions. For example, Users.Id 1140579 from India wrote in their profile "*My expertise lies in VB.Net + Excel Automation, VSTO and VBA,*" voluntarily disclosing their technical areas of interest. Similarly, there is a perception that Australians are generally happy-go-lucky [94], where such behaviours may ripple towards their online behaviour. Australians' inclination towards openness and informality may translate into a greater propensity to disclose personal information, anecdotes, and supplementary details within their online profiles. For example, Users.Id 398670 from Australia wrote in their profile "*I've turned into a git fiend and no longer understand why anybody uses anything else. I'm also becoming increasingly obsessive about testing,*" where hyperbolic language and self-deprecating humour suggest a jesting attitude. Interestingly, those from the United States and the United Kingdom tend to write less within their profiles. This difference in behaviour could be explained by cultural variations in self-expression, as Americans have been seen as more reserved [95], and less likely to share personal information online. On the other hand, unspoken norms in the United Kingdom often emphasise self-effacement, which may discourage individuals from engaging in overt self-promotion [96]. Such aversion may thus lead to a preference for brevity in their online profiles, emphasising the substance of their expertise and background rather than an elaborate self-description.

Despite having the largest and second-largest user bases, the United States and India exhibit relatively lower *YearlyDurationUsage* compared to other countries on the platform. Hence, despite having substantial user base, users from both countries exhibit a relatively recent registration pattern where the majority of users only joined the platform within the last two years. In contrast, users of Australia demonstrate a substantial *YearlyDurationUsage*, suggesting a more engaged, longer tenure, and higher levels of temporal commitment within their user base despite their smaller overall number. Moreover, observed disparity between Egypt's *YearlyDurationUsage* and its relatively less developed tech ecosystem suggests a resilience and adaptability among Egyptian users, indicating a strong intrinsic interest in programming-related CQA that transcends broader technological limitations.

With respect to *UpVotes*, our analysis reveals that users from Australia and the United States presented higher numbers. These countries fall under low-context cultures [97]. In low-context cultures, communication tends to be explicit and direct [98]. This cultural emphasis on direct expressions could manifest in user interactions on Stack Overflow, leading to a higher propensity for using *UpVotes* as a clear, unambiguous way to acknowledge helpful contributions. Conversely, higher-context cultures such as Brazil, India, and Egypt may cast fewer *UpVotes* due to a preference for more nuanced and context-dependent communication [99]. The patterns observed for *DownVotes* are less distinct compared to those for *UpVotes*. For one, Egypt presented higher *DownVotes* than United Kingdom, but the latter also ranked above India in terms of *DownVotes*. Contrary to the initial assumption that the United Kingdom is a low-context culture, empirical evidence suggests otherwise. Specifically, United Kingdom exhibits a blend of high-context and low-context cultural characteristics (mid-context), rather than falling neatly

into either category [29]. Our findings thus provide preliminary evidence that high-context cultures utilise *DownVotes* more frequently. In high-context societies, communication often relies on implicit understanding [100]. We posit this may extend towards disagreement expressions via *DownVotes* as a means of signalling.

The observed dominance of users from the United States and the United Kingdom in terms of *Views* and *Reputation* aligns with the presence of prominent technology clusters in these countries. Users hailing from countries with smaller and less mature tech ecosystem, such as Brazil and Egypt, exhibited lower averages of *Views* and *Reputation*. Early access to cutting-edge tech in countries with notable tech ecosystems could attract global Stack Overflow users to see how they harness such technologies, fostering interest, renown, and visibility. For example, the development and global launch of GraphQL by Facebook (now Meta) in the United States [101] may position American users, particularly those involved in GraphQL development, as initiators of such topic discussions on Stack Overflow. Their active participation may lead to increased profile views, as they contribute to resolving challenges faced by others that may still adapt to GraphQL. Consequently, their worldwide assistance may contribute to a higher accumulation of reputation.

Another interesting observation pertains to India having the highest number of *Questions*, yet the United States retained the highest average *Answers*. This may be attributed to several factors. We postulate that in India, despite its burgeoning tech community, a large number of developers may be relatively early in their careers [102]. This higher rate of questioning may reflect a proactive approach to seeking assistance within the Indian user community. In contrast, Americans are generally seen to exhibit ethos towards volunteering, altruism, and care for others [95]. It may thus be surmised that their preference in contributing to the well-being of others extends to Stack Overflow, where such users may find satisfaction in providing assistance or sharing their expertise. With respect to *Comments*, our analysis indicates that users from the United Kingdom presented the highest. This trend could be attributed to their general inclination towards indirect communication, characterised by subtle language nuances and implicit meanings [29]. This aligns with the purpose of comments, which often serve to provide additional context, suggestions, or clarifications without answering the core inquiry. Users from the United Kingdom also exhibit the largest average *Edits*, a phenomenon that may be inherently linked to their tendency to value diplomacy and tact [29]. Editing may be perceived as a subtle, yet consequential means of enriching the community's knowledge base, aligning with the British general emphasis on polished communication [103]. Finally, we observe users from developing countries – such as Egypt and Brazil – have less *Badges* than their developed counterparts like United Kingdom or United States. This disparity could be attributed to a difference in perceived value and motivation towards gamification elements. For example, the former group may see gamification as meaningless and therefore stopped chasing such elements altogether [104].

## 7.2 Users' Interaction Patterns (RQ2)

Looking at the results, the graph layout of top contributors generated for the United States reveals a clustering in the centre. However, these users do not necessarily exhibit the highest levels of engagement, hinted by their lighter hues and relatively smaller node sizes. In fact, more active users (i.e., darker and larger nodes) tend to be located outside of this centre cluster. Such a dichotomy implies central cluster users engaging in more frequent answer exchange within their group, despite their lower activity levels, while users outside this cluster engage in a broader range of interactions. Large node sizes in the latter indicate their tendency to give and receive answers across diverse question threads rather than sustained interactions, unlike those in the central cluster. Given the knowledge-intensive nature of SE [105], it is likely that highly active users may possess broad technical skillset [106], facilitating engagement across manifold domains rather than being constrained to specific topics. A notable example is Users.Id 34397. Upon manual examination of their profile, we discovered that this user exhibits strong expertise in C#, JavaScript, and ASP.NET MVC, along with other domains.

The social network graph for India reveals a less clustered distribution of top contributors compared to the United States, exhibiting decentralised and loosely-connected network structure. This observation implies that the top contributors of India engage in interactions more evenly across diverse topics rather than forming tightly-connected communities centred on specific niches. While certain users may exhibit higher levels of activity (e.g., Users.Id 541786 and 3701974), the overall interaction landscape among

these top contributors remains relatively uniform. The cultural characteristics of India, encompassing both collectivism and individualism [107], may owe to this phenomenon (though India leans more towards the former). Specifically, their individualistic aspect might encourage users to pursue a variety of personal interest [108], thereby contributing to a diverse array of topics. Simultaneously, their collectivistic values of group harmony may discourage the formation of exclusive clusters [109]. Examination of sentiment patterns reveals that Users.Id 3701974 consistently encounters a higher frequency of negative-toned interactions compared to positive ones. A manual review of their profile confirmed that a considerable number of answers directed towards this user convey negative emotional tones. For example, another user (Users.Id 5043867) answered as such within the subject’s question post: “*You have done various mistakes few of them are as mentioned here...*” The willingness of Indian top contributors to provide critical feedback, even within a strong sense of cultural solidarity [110, 111], suggests a balance between collective loyalty and individual responsibility.

The social network graph for the United Kingdom displays a centralised structure akin to that of the United States. However, a key difference is with the positioning of highly active users. In the former graph, these users occupy the central cluster, while in the latter, they are located outside the core cluster. The observed hub-and-spoke structure [112] indicates that highly active users – even among the top 100 contributors – serve as influencers, mediating interactions and shaping the dissemination of information within the network. It is also discernible that social interactions within the United Kingdom exhibits a core-periphery structure where highly active users exhibit a strong reciprocal relationship, engaging in a high volume of answer exchanges among themselves [113]. Conversely, a lower level of connectivity may be observed among less active participants. This may be due to users from the United Kingdom being task-based [29], where highly-active users are often driven by the goal of accomplishing objectives, subsequently seeking to benefit and contribute from the central core where information is exchanged rapidly. Such measures allow tasks to be efficiently completed. Moreover, anecdotal evidence indicate that individuals with a high degree of social embeddedness are more likely to make valuable contributions to the community [114]. This is exemplified by Users.Id 22656, who consistently ranked among the top contributors on Stack Overflow for a decade [115]. In recognition of his exceptional contributions, Stack Overflow commemorated his remarkable achievement of reaching one million reputation points by publishing a dedicated blog post<sup>22</sup>. However, it is not to say that less active participants contribute less knowledge to the platform [114]. Their knowledge may still supplement those that are more active, even if the former’s participation is less visible [116].

The social network graph for Australia exhibits similarities to that of the United Kingdom, displaying a core-periphery structure. This similarity may be attributed to their shared position on the ‘applications-first’ end of the persuading scale [29], where persuasion is primarily driven by practical considerations rather than abstract reasoning or principles. However, unlike those of the United Kingdom, Australian users with high levels of activity tend to be less centralised, with their distribution spanning across distinct segments of the graph. Such interaction patterns could be indicative of Australians’ inclination for diverse and socially-engaging interactions [117], as opposed to more rigid approach of community formation. Despite this dispersion, highly active users remain directed towards the central region, aligning with their pragmatic ‘applications-first’ approach where practical insights are exchanged efficiently [118]. Less active users, likely prioritising observation or adopting a ‘lurker’ engagement style [16], tend to occupy the network’s outskirts. These passive preferences may still lead them to consume information, but they might not be as centrally involved in the dynamic core of problem-solving discussions. Rather, they may gravitate towards frequent interactions within their own segment of the network, as evidenced by the recurring positive interactions between Users.Id 283366 and Users.Id 356282, both of which may be classified as less active users (smaller and lighter nodes; see Appendix A). A closer examination of their user profiles indicates a shared interest in Vue.js and jQuery, suggesting that their frequent interactions may stem from a commonality in their technical interests. Overall, our findings align with the work by Safadi et al. [114] where, within online communities, individuals attain social status not through inherent characteristics (e.g., formal roles or personal profiles) but through their active engagement with others. The nature of these interactions, particularly the selection of conversations and the attraction of peers, is influenced by shared interests [114].

---

<sup>22</sup> <https://stackoverflow.blog/2018/01/15/thanks-million-jon-skeet>

The social network for Brazil shows a markedly sparser graph with a significantly reduced number of nodes and edges compared to other countries. This suggests that even among the top 100 users of Brazil, interactions within this subset are limited, and their primary engagement may occur with individuals outside this subset or with users from other countries. Research on cultural perceptions consistently demonstrates Brazilians as warm, welcoming, joyful, and positive-minded individuals [119]. This cultural predisposition may lead users from Brazil to conduct interactions with a diverse range of international participants rather than solely engaging with fellow Brazilians. For example, Users.Id 131874 recently provided an answer to a question posted by Users.Id 3509 from Israel. Another example pertains to Users.Id 276959 where their latest answer is on a question thread by Users.Id 3622471 from Dallas, TX (see Appendix A). These instances of cross-cultural engagement illustrate the tendency of Brazilian users to interact with individuals from diverse backgrounds. Our findings align with recent research demonstrating that Brazilians exhibit a strong propensity for acculturation, actively assimilating aspects of manifold cultures [120, 121].

Similar sparsity is observed in the social network graph for Egypt, and this has the least nodes and connections among all six observed countries. This observed sparsity, even among its top 100 users, suggests limited internal engagement, possibly attributed to a relatively underdeveloped tech ecosystem in the region [122, 123]. Instead, top contributors of Egypt may opt into knowledge exchange with users from regions with more established tech ecosystems. This pattern is parallel with the notion that individuals in less developed environments may gravitate towards external knowledge sources to compensate for local limitations [124]. For example, the most *UpVoted* answer by Users.Id 722783, represented by the central node in Figure A4, was provided in response to a question posed by Users.Id 815244, a user located in Charlotte, NC. This exemplifies that top contributors from Egypt exhibit a propensity for disseminating information with individuals from other countries. However, this pattern does not necessarily imply an aversion to interacting with fellow Egyptians – but rather highlights a tendency to broaden their knowledge base and engage with global standpoints.

### 7.3 Users' Topics of Discourse (RQ3)

Thematic analysis conducted on the results obtained from our LDA model revealed several insights. Within Australia, it is notable that the most prominent topic is performance suggestions, hinting at a strong inclination among Australian top contributors to provide performance-enhancing recommendations within their answers. One example is noted for Users.Id 16076, who explicitly expressed in their profile “*I'm passionate about writing clear, concise code and performance tuning...*,” hinting at the user's general preference for performant codes. Aside from performance, the identified topics suggest a diverse range of interests among users, including graphical user interfaces (GUI), form templates, web frameworks, and language-specific topics such as ASP.NET and the C programming family. This absence of a single dominant topic among Australian top contributors suggests a diversified user community on Stack Overflow, where their contributions span a wide range of areas, rather than being concentrated in a specific domain – thereby reflecting their breadth of expertise. This diversity may be attributed to Australia's rapid advancement in the IT industry, which necessitates a customer-centric and market-driven approach for successful commercialisation [125]. To create software products that meet the needs of both customers and developers, developers may need to adopt versatile roles [126, 127], requiring them to engage in discussions that encompass a multitude of topics. Our results indicate that this trend may propagate to their Stack Overflow user base. Moreover, the resurgence of startups within Australia may contribute to the prevalence of general *Software Development* topics among Australian top contributors on Stack Overflow [128, 129], as compared to other areas like *Human-Computer Interaction* or *Cybersecurity*.

Topics discussed by the top contributors from Brazil exhibit a narrower range of diversity. Specifically, they demonstrated a stronger inclination towards answering questions related to functional programming and error handling, both of which are fundamental concepts in SE [130]. The inherent ease of implementing functional programming concepts in Python has contributed to its strong association with this paradigm [131], which aligns with the observed prevalence of Python package-related topics among Brazilian top contributors' answers. Despite the existence of other languages being discussed in certain topics (e.g., Java in Android development or JavaScript in GUIs), the pervasiveness of Python-related topics suggests that Python is their dominant language of discourse. Top contributors from Brazil may

be incentivised by the comparatively simpler natural language-like syntax of Python, compared to low-level languages like Java or C++ [132]. Moreover, as Brazil's tech ecosystem is leaving its nascent stage [133], aspiring developers may find Python's beginner-friendliness appealing [134], making it an attractive choice for those seeking to enter the tech industry. This aligns with the observed predominance of topics within the *Software Development* sub-category among Brazilian top contributors, where startups and small tech enterprises may prioritise general software development over specific niches [126, 127].

Top contributors from Egypt displayed a similar pattern with error handling and array-like data structures emerging as prominent topics of discussion. These topics, both foundational in software development [135], underscore the focus on fundamental concepts among answers originating from Egypt. However, top contributors from Egypt did not exhibit a notable focus on any particular programming language. Even the observed discussions related to PHP were grouped alongside JavaScript under the broader topic of *PHP Sessions*. Thus, we surmise Egyptian top contributors' contributions may encompass a wider range of programming languages (or perhaps language-agnostic), rather than being concentrated in a specific language domain. Such phenomenon may be attributed to the embryonic state of Egypt's tech ecosystem [122, 123] that limits the scope of technical discussions and restricts implementations primarily to less complicated problems. Consistent with the previous two countries, Egyptian top contributors exhibit a stronger inclination towards *Software Development*-related topics compared to other sub-categories. However, they do not engage in discussions related to *Cybersecurity*, *Algorithms*, or *Computer Architecture* – outlining that such users are primarily focussed in practical foundational knowledge in software development, rather than specific implementations.

Interestingly, top contributors from India demonstrate a shared inclination with those of Australia, discussing performance-enhancing improvements in their answer posts. This suggests an advocacy for performant coding practices within the user communities of both countries. Our findings are particularly noteworthy given India's established position as a prominent tech hub [136], fostering a thriving ecosystem of startups and IT enterprises. In addition, the blooming 'startup culture' in India may further contribute to the performance-oriented coding practices among their users on Stack Overflow [137], as the dynamic and competitive nature of startups necessitates efficiency and optimal performance as deciding success factors [138]. Indian top contributors, immersed in this culture, are likely to be driven towards adopting (and therefore advocating) for performant coding practices, ensuring its alignment with the demands of a rapidly evolving and innovation-driven tech sphere. The observed variety of topics and programming languages answered by those from India, including jQuery (JavaScript), pointer behaviour (C-family), and ASP.NET web pages, further corroborates the influence of startup culture. In other words, users from India may be more inclined to wear multiple hats due to exposure and demands of their startup ecosystem [126, 127], subsequently conducting knowledge exchange spanning a wider range of technical areas. Notably, Indian top contributors exhibit limited engagement in discussions related to *Cybersecurity* and *Computer Architecture*. This suggests a narrower focus on *Software Development* compared to other sub-categories.

In contrast to the prevalence of technical discussions observed in other regions, the most prominent topic among top contributors in the United Kingdom centres on question-answering practices. Analysis of keywords in Table 10 suggests that top contributors from the United Kingdom engage in evaluating the given question, which primarily focusses on the asker's implementation strategies. For example, Posts.Id 195615 authored by Users.Id 22656 said “*It's not a good idea to assume that an offset is a number of hours or half-hours...*”, and Posts.Id 13254652 by Users.Id 12960 said “*I really can't imagine this is a good idea. Each thread takes a reasonable amount of resource...*” Observations suggest that top contributors from the United Kingdom tend to express their negative feedback in a subtle manner, employing phrases like “*I really can't imagine this is a good idea*” rather than directly dismissing the proposed implementation. This finding aligns with the discoveries by Meyer [29], which categorise the United Kingdom as exhibiting a preference for indirect negative feedback. In such ecosystems, negative feedback is typically conveyed tactfully, sometimes even delivered privately to maintain the recipient's rapport [29]. We postulate that this specific characteristic for users from United Kingdom extends to their online interactions on Stack Overflow. Interestingly, other prevalent topics range from programming basics such as loops, control flow, functions, as well as variables – to specific niches such as LINQ, jQuery, or directory traversal algorithms. One potential explanation for such a high degree of

variety in topics could be attributed to the proliferation of bootcamps. Popular ones such as Makers<sup>23</sup> and Le Wagon<sup>24</sup> could contribute to a broader, yet foundational level of technical proficiency among developers, thereby resulting in a wider array of topics being addressed in their online discussions. In terms of sub-categories, top contributors from the United Kingdom exhibit a preference for discussions on *Software Development*-related topics, with no apparent engagement in topics related to *Third-Party Resources*.

For top contributors from the United States, the most prevalent topic in their answers is centred on version control, as evidenced by the high frequency of keywords such as *git*, *commit*, and *branch*. A manual inspection of answers under this topic revealed that they primarily focus on providing explanations for Git's general functionalities and common pitfalls. For example, Posts.Id 68750921 by Users.Id 1256452 which said "*You cannot change the commit you made. You can make a new commit...*" which highlights the limitation of Git, and its potential workaround. Another example within Posts.Id 7520329 by Users.Id 148870 which wrote "*First, let's just make something clear: there is no single "correct" workflow for Git...*" clearing potential misconceptions in Git usage. The nature of answers falling under this topic led us to infer that they are primarily aimed at clarifying users' general understanding of Git usage, its capabilities, and limitations. When clarifying technical concepts, the inherent assertiveness often attributed to Americans could extend to their online interactions [139]. Such a direct approach may prove useful in minimising ambiguity and ensure that all participants are on the same page. Moreover, they may be more likely to take the initiative to clarify users' understanding because they are more comfortable with taking on leadership roles [140]. Interestingly, top contributors from the United States demonstrate a similar inclination to offer performance suggestions (observed among the top contributors of Australia and India), as well as a tendency to evaluate users' approaches (observed among those from the United Kingdom). This convergence of behavioural traits could be attributed to the nation's longstanding history of cultural diversity [141], fostered by a continuous influx of immigrants from diverse backgrounds across the globe [142]. We postulate this cultural diversity also affects the fact that users from the United States have engaged in discourses across all sub-categories, with *Software Engineering* being the most frequently discussed (14 instances) and *Cybersecurity* being the least discussed (1 instance).

#### 7.4 Users' Knowledge Exchange (RQ4)

Our findings reveal that top contributors from Australia, United Kingdom, and the United States were more likely to ask direct questions (*Type I Question*) to resolve a knowledge deficit. This preference could be attributed to the individualistic cultural norms prevalent in these societies [143, 144]. Individualistic cultures tend to emphasise directness, which may manifest in online interactions through a preference for straightforward communication [145, 146]. Conversely, those from collectivistic societies such as Egypt, India, and Brazil exhibited a lower propensity for direct question-asking. Instead, such users (except those of India), demonstrated a tendency towards initiating discussions and posing indirect questions – coded by the theme *Type II Question* in Table 12. We postulate this is due to the collectivists' nature in indirectness and maintaining social harmony [147]. Finally, the observed deviation by users from India may be owed to the cultural emphasis on silence and understated communication [148]. Thus, such preferences may lead users to refrain from asking direct and indirect questions altogether.

Patterns regarding the individualism-collectivism spectrum also propagate towards other coding themes. We observe countries leaning towards individualistic societies (namely Australia, United Kingdom, and the United States) exhibit a greater inclination to directly answer questions and conduct information exchange compared to their collectivistic counterparts. Again, cultural emphasis may play a part, in which their preference for explicit communication would contribute such tendencies in directly addressing questions and providing relevant information [149]. In one example, Users.Id 527702 from Australia stated "*Here's the simplest way to get the info you want using jQuery and YQL...*" which emphasises clarity and brevity. However, this does not imply that the top contributors from collectivistic cultures (e.g. India, Egypt, and Brazil) lack knowledge. Instead, their contributions often extend beyond direct answers, expressing their ideas and providing comments that oftentimes do not directly answer

---

<sup>23</sup> <https://makers.tech>

<sup>24</sup> <https://www.lewagon.com>

the given inquiry. This aligns with the notion that collectivistic cultures tend to foster relational connections and cultivate a sense of community spirit, focussing on collective well-being and the common good [150]. For example, Users.Id 458204 from Egypt wrote “*It is very good and easy to use,*” providing a subjective opinion rather than directly addressing the core inquiry.

Drawing upon these observed characteristics, we postulate that inherent cultural attributes of collectivistic cultures may also contribute to their dominant presence of self-evaluation. Top contributors from India, Egypt, and Brazil provide more reflective statements than those from Australia, United Kingdom, and the United States. For example, Users.Id 839211 from Brazil wrote “*I believe that this is important because you will always have a "rollback" option if you discover a bug in your system later on...*” which tells of the user’s recommendation if they were in the asker’s shoes. Another example is by Users.Id 999885 from Egypt, having written “*I didn’t notice it was that easy...*” admitting that the user have missed certain implementation shortcuts. The observed inclination towards self-reflection among collectivistic individuals is largely warranted, serving as an essential mechanism for aligning one’s behaviour with the implicit social norms that govern these societies [151]. Our findings also corroborate cross-cultural studies that examine these countries in isolation. For example, Bashir and Khan [152] suggested that Egyptians tend to be more bureaucratic, resulting in a slower pace of consensus-building. Similarly, Gupta and Sukamto [148] indicated that Indians prioritise maintaining social harmony, sometimes at the expense of directly addressing issues or exchanging information with complete accuracy.

We also explored how countries differ in terms of giving directives and scaffold content. Top contributors from individualistic cultures were more likely to provide directives and offer advice. For example, Users.Id 13198 from Australia wrote “*Get the image in DoWork handler and assign the image to PictureBox in the RunWorkerCompleted handler...*” which hints a more instruction-like tone. Another example is from Users.Id 7964527 from the United States, having written “*You can using pandas describe...*” which offers one way to address the asker’s question – while hinting that there may be other alternatives. However, the distinction in scaffolding behaviour is less pronounced, with Indian top contributors proposing more implementation alternatives compared to those from Australia. While India adheres to a collectivistic cultural framework [153], its thriving tech industry may foster an environment where diverse approaches to problem-solving are encouraged [153, 154]. This inclination towards exploring multiple alternatives to address a single issue may thus be reflected in their tendency to provide manifold implementation suggestions. As an example, Users.Id 548225 from India noted “*You can modify your awk to take care of it...*,” suggesting that the user proposed a solution beyond the scope of the asker’s initial consideration.

Cultural dimensions seem to resonate towards how top contributors convey gratitude and apology. For one, those from collectivistic societies were more likely to express apologies and gratitude in their answers, compared to those from individualistic ones. For example, a certain Users.Id 2977164 from Brazil wrote “*Thanks for the good tips @HubertL, problem solved!*” which expresses thankfulness for another’s advice. Another example is the post by Users.Id 3240583 from Egypt, having written “*my bad I updated the ingress file without run...*” showing both self-evaluation and remorse. These trends may be attributed to the crucial role of expressing appreciation and offering apologies in sustaining group cohesiveness, particularly as prioritised in collectivistic societies [9]. However, it is not to say that individualistic societies never apologise or express their gratitude – they still do, albeit less frequently (and not as explicit) than their collectivist counterparts. For example, Users.Id 57159 from the United Kingdom wrote “*Sorry, if you think about the “net present value” of any software you write today, it has no effect what the software does in 2038...*” which conveys a brief remorse before promptly delving into the substantive response. Findings also align with previous work [9] in which users from collectivistic cultures tend to convey more gratitude and apology in online written communications. Finally, we observe the *Off Task* and *Not Coded* coding themes. Chi-squared test on the former indicates that the differences are not statistically significant, further hinting that fluctuations across countries may be due to chance, while the *Not Coded* category was minimal, supporting the relevance of the coding scheme. Limitations for both the *Off Task* and *Not Coded* coding themes are outlined within Section 8.

Stack Overflow top contributors were found to generally exhibit consistent behavioural patterns. This observation is supported by the magnitude of the *Answer* and *Information Exchange* coding themes



across all countries, indicating a general tendency among users to address posed inquiries and engage in knowledge sharing as primary modes of knowledge exchange. This convergence in their behaviour could be attributed to the shared identity of software developers as adept problem-solvers [155, 156]. Hence, as the underlying objective of providing answers and exchanging novel knowledge are largely similar across all users, our results suggest that cultural backgrounds influence the manner in which this objective is achieved, evident in fluctuations between *Instruction/Command*, *Comment*, or *Scaffolding* coding themes. Specifically, top contributors from collectivistic cultures were more group-oriented and often employed indirect communication strategies. In contrast, those from individualistic cultures were found to use more assertive, direct, and task-oriented communication, often avoiding unnecessary circumlocution. Our results corroborate cross-cultural studies regarding both dimensions [145, 146, 149]. Despite being highly influential, cultural dimensions are not the sole determinant of variations in users' knowledge exchange. Rather, specific contextual factors within a country, such as the experience of the user and prevalence of technology, also play a significant role. Finally, patterns conceptualise the dynamics of the global knowledge ecosystem on Stack Overflow, revealing that individuals from diverse backgrounds contribute in their own unique ways.

## 7.5 Implications

The high gender disparity is apparent in Stack Overflow, mirroring the broader dilemma faced by the broader tech industry. Encouragingly, initiatives like PyLadies<sup>25</sup> and Women Who Code<sup>26</sup> have emerged as pioneers for inspiring more women in technology. However, these programs are primarily concentrated within the United States, which could be a contributing factor to the nation having a lower male-to-female user ratio compared to other countries. Building on this observation, we advocate expanding such programs or adapting them to fit the specific needs of countries with higher male-to-female ratios (e.g., Brazil). Broader implementation of such groups has the potential to significantly improve female representation within both Stack Overflow and the general tech industry. Next, users from developed nations generally receive more profile views and amass more reputation on the platform, attributable to their higher answer count. Additionally, a noteworthy pattern emerges within countries experiencing tech industry growth, characterised by an abundance of questions, indicating a strong inclination towards acquiring technical knowledge. This insight may guide teams in strategically utilising platforms like Stack Overflow to cultivate latent talent within their ranks, fostering a culture of expertise development. For example, senior team members from the United States can facilitate knowledge spillover to their less experienced counterparts, as they may have gained earlier exposure to specific technologies. Finally, Egypt's ongoing initiatives to bridge the digital divide have demonstrably increased user retention (higher *YearlyDurationUsage*), evidenced by factors like the growth of high-tech startups and advancements in infrastructure [157]. We thus advocate for other developing countries to adopt such programs, which we believe will promote equitable access to technology.

To enhance Stack Overflow's appeal to newcomers, we recommend a multifaceted approach that combines user-friendly onboarding materials and clear community expectations. Specifically, the provision of a simplified, user-friendly guidelines, video tutorials, and workshops tailored for new users (instead of the general guidelines that currently exists) would be of great benefit. This onboarding process should introduce best practices and expectations for effective community engagement, where this simplified guideline would explicitly warn of potential consequences for non-compliance, such as downvotes or unanswered questions. Additionally, clear and concise community guidelines that emphasise respectful interactions should be developed – even for established users to maintain a culture of civility. Figure 6 highlights such a concern regarding the treatment of new contributors. The figure depicts a scenario where a user, clearly identified as new, receives downvotes without any accompanying explanation. Given that the user may be unaware of the specific issue with their contribution, such issues fail to address the root cause of the negative feedback and could discourage them from future participation.

---

<sup>25</sup> <https://pyladies.com>

<sup>26</sup> <https://womenwhocode.com>

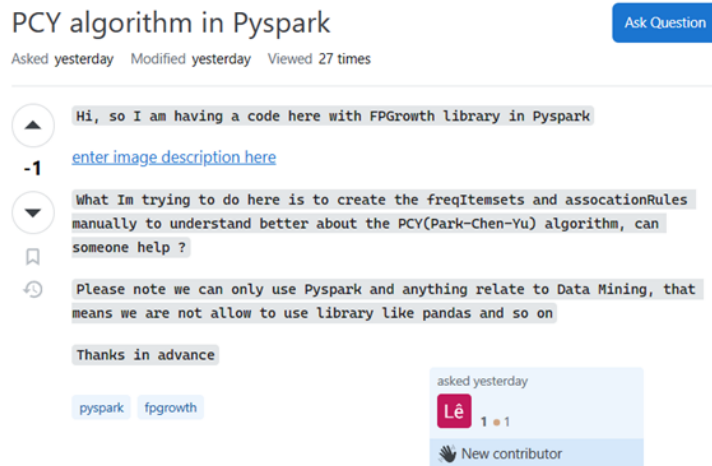


Figure 6. Example downvoted post in Stack Overflow

Based on this issue, we advocate to supplement the downvoting mechanism that discourages anonymous downvotes, instead prompting users to provide constructive feedback directly to the original poster (e.g., “off-topic” or “unclear question”). We believe such rework would enable newcomers to identify areas for improvement in their future contributions, as they represent the potential for growth, innovation, and expertise. However, if such newcomers are ostracised (e.g., through hateful comments or unexplained downvotes), their discouragement could lead to a shrinking pool of future contributors. Finally, we also advocate for the development of online semi-formal communities, akin to Stack Overflow Collectives<sup>27</sup>, which holds potential for fostering a more inclusive environment. These safe spaces would cater specifically to underrepresented groups, such as women, the LGBTQ+ community, or non-English speaking users by offering them opportunities to network, socialise, and foster a stronger sense of belonging within the broader CQA community.

We found that the top 100 contributors within individualistic societies (United States, United Kingdom, and Australia) exhibit a core-periphery structure. In contrast, those from collectivist societies (India, Egypt, and Brazil) demonstrate a more decentralised structure. This insight is relevant for Stack Overflow – as well as similar CQA site designers – to develop effective acculturation mechanisms to enhance productivity while respecting members’ cultural norms. As an example, integrating egalitarian-based approaches developed in collectivist cultures with the hierarchical, task-oriented focus of individualistic societies could foster a more collaborative yet goal-directed community atmosphere. Such an approach may simultaneously mitigate each other’s weaknesses and create a synergistic environment. Furthermore, our analysis uncovers latent topics of interest within each country’s top contributors, which could enable industry practitioners to utilise these insights to stay abreast of emerging trends and cultivate developers’ learning preferences. We postulate such measures would augment their overall adaptability within the team. Bootcamp companies may harness these trends to identify knowledge gaps and determine the specific niches required to strengthen the local tech cluster within their respective regions.

In line with their cultural norms, top contributors from gregarious societies exhibit similar behaviours in their online interactions, characterised by a focus on social connections and informal communication. Conversely, those from task-oriented cultures tend to adopt a more direct and goal-oriented approach in their online interactions. Efficient communication channels should be devised so that both ends of the spectrum are facilitated, in turn bridging culture-related discrepancies [29]. Finally, research has demonstrated the ineffectiveness of one-size-fits-all solutions that solely focus on either collectivist or individualist norms [158]. Hence, our results largely accentuate the need to consider socio-cultural variations in online communication channels, so that members’ societal norms are facilitated, thereby nurturing innovation and team cohesiveness [9]. We expect doing so would enhance the happiness of software developers, in turn begetting higher problem-solving aptitude and creativity [156]. Overall, our findings confirmed the existence of a complex interplay between diversity in human aspects, geographic

<sup>27</sup> <https://stackoverflow.com/collectives>

location, and contribution levels within Stack Overflow. These results underscore the necessity for future research to qualitatively gauge how these complexities manifest beyond the platform. For instance, user experiments may be conducted to evaluate the practical utility of our findings (see Section 9).

## 8 THREATS TO VALIDITY

### 8.1 Internal Validity

From a global user base of 17,922,426, we excluded 14,242,161 users due to missing or invalid location data, retaining a smaller subset of 3,680,265 users. From this reduced subset, we further selected 1,762,745 users to represent each continent’s user base, effectively discarding 1,917,520 users. Moreover, several users who did not explicitly state their location might have originated from either Australia, United Kingdom, United States, India, Brazil, or Egypt. Conversely, certain users may have provided inaccurate location information. Our data cleaning procedures, while having undergone extensive tests, may have inadvertently excluded relevant data (or included ones that are irrelevant). However, a manual examination of user profiles following our directed CA ensured that this potential limitation has been mitigated as much as possible. In fact, similar biases were held to be evident in a previous study [26], and thus, we employed mitigation strategies suggested by these authors to address them effectively.

With regards to our statistical tests in RQ1, we employed Bonferroni corrections to control the experiment-wise error rate. However, this approach is known to be overly conservative, potentially leading to a higher likelihood of type II errors (i.e., failing to detect true effects) compared to alternative methods like Holm or Benjamini-Hochberg corrections [159]. Our statistical outcomes should thus be interpreted with certainty given the conservative approach. For RQ2, our SNA yielded insights regarding highly influential users. While we recognise the potential for deeper qualitative inquiry into the factors contributing to their popularity, the exploratory nature of our study mainly seeks to establish any preliminary causations between diversity and user contributions. Thus, in-depth qualitative methods apart from the employed directed CA have been intentionally excluded from this research. Additionally, our decision to analyse only the top 100 contributors per nation for RQ2 onwards (see Section 5.4) might have introduced sampling bias, potentially including outliers. However, given the general alignment of our results with existing literature, we believe that such biases – if any – exert minimal impact on our findings. Next, the topic modelling (RQ3) and directed CA (RQ4) phases employed conventional methodologies (when considering content generators available today). We recognise the potential of more advanced LLM-based techniques to generate results that are unaffected by researcher bias (e.g., misinterpreting themes on LDA based on extracted keywords). For example, topic modelling could employ transformer-based models such as BERTopic [160], and deductive coding can be supported by GPT-3.5 [161]. Using such tools can be particularly helpful, given their demonstrated proficiency in code documentation and conceptual understanding on Stack Overflow data [162]. We should note that these methods are not perfect, and are in fact still evolving. For instance, LDA was found to produce more coherent topics and yield less outliers than BERTopic [163, 164]. Also, due to the study’s exploratory design, more conventional methods were adopted with hopes that it will shed light into any patterns worth further investigating. For CA, we leverage the readily accessible Stack Overflow Database, which does not require obtaining explicit consent from users. Nevertheless, user data have been omitted during the sampling process to safeguard the anonymity of individuals and their private information.

Regarding gender classifications provided by Genderize, our findings are limited to binary cis-male and cis-female demographics. Thus, results may not adequately represent the LGBTQ+ community within Stack Overflow, despite an emerging body of research regarding biases toward non-binary genders [36]. There were challenges in standardising user-provided location data. At first, we tried conventional tools such as fuzzy string matching and packages like GeoText<sup>28</sup>, but they presented several limitations [39]. For example, GeoText struggled to recognise locations written in non-Latin characters, such as “Москва” (Moscow) or “北京” (Beijing), and was incapable of identifying locations without an explicit city or state name (e.g., “14 Downing Street, SW1A 2AH”), rendering it unsuitable for our study. As standardising users’ location data manually would require a significant amount of effort and time, we

---

<sup>28</sup> <https://pypi.org/project/geotext>

turned to proprietary geocoding APIs which often leverage complex machine learning algorithms, offering enhanced reliability and accuracy [165]. Empirical evaluation revealed that MapQuest outperformed the other APIs tested. For example, while OpenStreetMap excelled in handling English characters within Western countries such as Australia and USA, it struggled with transliterating non-Latin characters back to their Latin equivalents. In one instance, a query for the Indian district “गौतम बुद्ध नगर” correctly pointed to “Gautam Buddh Nagar,” but the same query with the added location context “Gautam Buddh Nagar India” erroneously yielded “Bridgenorth” in Canada. Additionally, ArcGIS’ daily limit of 2,000 addresses was insufficient for our dataset, which exceeded 3 million rows and thus would require thousands of days to finish. We thus picked MapQuest as it presents satisfactory results. However, notwithstanding our extensive testing regime, MapQuest itself is not flawless. Despite boasting a commendable 99.22% accuracy rate, the remaining 0.78% margin for error may have a small impact on our findings.

## 8.2 External Validity

Our sample of 1,762,745 users only accounts for 9.84% of the global Stack Overflow user base. Consequently, our findings may not be directly generalisable to the broader Stack Overflow community or to other CQA websites like Cross Validated<sup>29</sup> or Quora<sup>30</sup>. Despite our sample encompassing countries from all continents to capture a range of cultural, linguistic, and contextual standpoints, the applicability of our findings to other countries within those continents may be limited. For example, while India represents Asia in our sample, our results may not be directly relevant to other sub-regions such as Malaysia (Southeast Asia) or Japan (East Asia). Nonetheless, by rigorously aligning our discoveries with established literature, we assert that these limitations are likely to exert minimal impact.

Stack Overflow is a dynamic platform where users’ contribution patterns can evolve over time due to changes in technology trends, platform features, and user demographics. For example, OverflowAI<sup>31</sup> was announced several months following the conclusion of our analyses, which might have influenced how users contribute. Moreover, the emergence of generative chat Artificial Intelligence (AI) tools – such as Google’s Gemini and OpenAI’s ChatGPT – has demonstrably helped SE-related rigours. Since the time of this study, such tools have accrued popularity, where this may influence the broader SE landscape and extend to which information dissemination practices were conducted on Stack Overflow. As our study does not incorporate longitudinal analysis to capture how observed patterns evolve over time [166], our results thus may not fully reflect latent changes that may have evolved since the time of writing. Our intra-country social network analysis to address RQ2 is confined to the top 100 users having the highest aggregate sum of answers given and received. Due to this magnitude, the behaviours of these top users may exert a disproportionate yet significant influence on the countries’ overall social network dynamics. We acknowledge our findings may not be generalisable to users beyond this subset, where the majority of users may exhibit different engagement levels from those of the top users. For instance, the extent of which gamification mechanisms (e.g., badges and reputation) drive or deter contributions may vary from person to person, especially those beyond the top 100 subset that we sampled from. Similarly, the utilisation of conventional methods for studying RQ3 and RQ4 limited the scope to only the top 100 contributors per country. Consequently, the generalisability of our findings might be limited compared to if one were to utilise techniques like BERTopic for topic modelling and GPT-3.5 for CA involving a larger subset of contributors’ data [160, 161]. Employing such methods may easily allow for the inclusion of users beyond the top 100, thereby presenting more comprehensive findings. However, this study is designed as an initial exploration into the interactions between diversity, geographical regions, and contributions. Consequently, the selection of the top 100 users from six countries was considered sufficient for deriving preliminary insights. Finally, it is important to acknowledge that our findings are primarily exploratory in nature, focussing on identifying trends that have emerged within the data. Our findings thus should not in any way be used to promote regional superiority or profiling users based on location. In fact, given our limited data, generalisability to broader regions is also limited, as areas with diverse sub-cultural compositions (e.g., Australia and the United States) may not fit into a single, one-size-fits-all construct. However, at the time of writing, we have

---

<sup>29</sup> <https://stats.stackexchange.com>

<sup>30</sup> <https://www.quora.com>

<sup>31</sup> <https://stackoverflow.blog/2023/07/27/announcing-overflowai>

endeavoured our best to avoid any interpretations and discussions that could construe such threats, and therefore we have minimised this to the greatest extent applicable.

### 8.3 Construct Validity

Due to the exploratory nature of this analysis and the understanding that a single region cannot fully encompass a monolithic culture, we refrain from using countries as a proxy to explicitly capture the construct of *culture*. In fact, they may not be accurate proxies to paint the construct of contribution, given the absence of a measurable, comprehensive composite measure for assessing users' participatory behaviour in CQA platforms. However, it is worth noting that the selection of all six countries (each representing one continent), were chosen carefully based on having the highest user representation per continent, a selection intended to mitigate the potential impact of this limitation to the greatest extent possible. Moreover, all six countries exhibited heterogeneity in cultural composition, subcultures, languages, and technological presence. In this manner, patterns presented in our study are not likely to derive from culture alone. Instead, languages spoken, educational attainment, and access to technology are likely to be some of many contributing factors.

When conducting thematic analysis on RQ3, we have exerted significant effort to adhere to established literature [38], yet our inferences derived from salient keywords extracted through LDA may contain potential inaccuracies. However, we have ensured their validity by consulting relevant literature and thoroughly examining individual posts. This meticulous process has enabled us to confidently assert that each keyword group corresponds to the inferred theme. For CA, our chi-squared tests reveal potential instances of Simpson's paradox, where pairs exhibited significance within specific groups but not when examined in detail [167]. We attribute this limitation to the lack of *a priori* identification of confounding variables and causal relationships, owing to the scarcity of literature [168]. Lastly, while inter-coder reliability ( $\kappa = 0.852$ ) indicates strong agreement, it exceeds the 64–81% range reported in the literature for reliable coding [79].

## 9 SUMMARY AND FUTURE WORK

This study explores the diverse patterns of top contributors' contributions on Stack Overflow across a representative sample of countries spanning multiple continents: the United States, Brazil, India, Egypt, the United Kingdom, and Australia. We commence with a quantitative analysis to identify the predominant platform activities among each user base that does not only encompass the top contributors, such as answering, asking, commenting, or voting. Next, we narrow our scope to the top 100 contributors and performed social network analysis to delve into intra-country user interactions and collaborative behaviours, taking into account polarity and intensity. Following the identification of user interaction patterns, we apply topic modelling to uncover the primary topics of discussion among the top contributors in each region, where we followed through with thematic analysis to pinpoint their key areas of interest. Finally, we conclude with a directed content analysis to complement the quantitative findings with a qualitative perspective. While topic modelling and thematic analysis focus on extracting the technical aspects of user interactions, our content analysis delves into the non-technical aspects, specifically knowledge exchange practices.

Our results provide insights into the user archetypes prevalent within each country. However, these archetypes are not entirely dichotomous; rather, they represent a multifaceted spectrum where complementary interactions between regions foster a holistic ecosystem of information dissemination. Firstly, we observe that contextual characteristics of countries exert great degrees of influence towards their user behaviour on the platform. Gregarious and communicative cultures tend to craft lengthier profile bios and demonstrate higher cross-cultural engagement compared to intra-cultural interaction. More collectivist societies tend to exhibit a more decentralised social network where there are no dominating users. Their answer posts display heightened socio-emotional cues, including apologies, self-evaluation, and expressions of gratitude. Additionally, they tend to employ indirect directives and favour more subtle contributions, such as comments that does not directly address the specific inquiry. Conversely, more individualistic societies display a core-periphery structure in their social networks, suggesting a tendency to concentrate on specific areas of interest and favour a task-oriented approach – rather than generalising their expertise compared to their collectivistic counterparts. Individualistic societies prioritise straightforwardness, hence they tend to engage in more reciprocal questioning within their answer posts, opting for more direct and directive responses over subtle communication through

comments. However, these archetypes are not always true, since we also observed indications that these ecosystems tend to evaluate the askers' proposed approaches rather than directly providing what the askers are seeking. Additionally, they exhibit a tendency to minimise socio-emotional cues such as apologising or giving thanks. We found some country-specific characteristics attributable to less explicit factors. For example, top contributors from Australia tend to provide performance-enhancing recommendations within their answers, often delving into language-specific niches. This inclination likely stems from their pragmatic 'applications-first' nature, where persuasion is predominantly driven by practical considerations rather than abstract principles. In contrast, top contributors from countries with nascent tech ecosystems, like Brazil and Egypt, focus on programming fundamentals and language-agnostic topics, rather than delving into specific language domains. Overall, while cultural ethos was found to influence user interactions, other factors like tech ecosystem maturity and startup emergence also play a similar, if not more significant, role.

Building upon our findings, future studies can delve deeper into user-generated content and how they differ across countries, or perform an inductive content analysis to supplement our top-down approach. The utilisation of transformer-based models (e.g., BERTopic) and LLMs (e.g., GPT-3.5) to aid in topic modelling and content analysis also remains as a worthwhile research endeavour, as doing so would move beyond conventional methods as we have employed in our study. In fact, given that our study has hinted certain nuanced relationships between diversity, geographical locations, and contribution levels, scholars might pursue this line of inquiry by including all users from all countries to more comprehensively examine how these variations manifest on a broader form. Next, as we have demonstrated user preferences towards specific languages and technologies, scholars can also examine the quality of which users implement their solutions. We propose that this research would deepen the understanding of the platform's coding ecosystem, potentially influencing general software development rigours. Finally, future works may conduct large-scale user experiments to qualitatively assess the extent of which unveiled findings may be applicable to the broader SE community. Research may also build on the insights gathered from our study to investigate user characteristics further. For instance, a qualitative exploration regarding what motivates highly popular users holds significant investigative merit.

## REFERENCES

- [1] Zahedi, M., Rajapakse, R. N., & Babar, M. A. (2020). Mining Questions Asked About Continuous Software Engineering: A Case Study of Stack Overflow. *Proceedings of the 24th International Conference on Evaluation and Assessment in Software Engineering*. 41–50. doi:10.1145/3383219.3383224
- [2] Ahmad, A., Feng, C., Ge, S., & Yousif, A. (2018). A Survey on Mining Stack Overflow: Question and Answering (Q&a) Community. *Data Technologies and Applications*, 52(2), 190-247. doi:10.1108/DTA-07-2017-0054
- [3] Fischer, F., Böttinger, K., Xiao, H., Stransky, C., Acar, Y., Backes, M., & Fahl, S. (2017, 22-26 May 2017). Stack Overflow Considered Harmful? The Impact of Copy&Paste on Android Application Security. *2017 IEEE Symposium on Security and Privacy (SP)*. 121-136. doi:10.1109/SP.2017.31
- [4] Morrison, P., & Murphy-Hill, E. (2013, 18-19 May 2013). Is Programming Knowledge Related to Age? An Exploration of Stack Overflow. *2013 10th Working Conference on Mining Software Repositories (MSR)*. 69-72. doi:10.1109/MSR.2013.6624008
- [5] Vasilescu, B., Capiluppi, A., & Serebrenik, A. (2012, 14-16 Dec. 2012). Gender, Representation and Online Participation: A Quantitative Study of Stackoverflow. *2012 International Conference on Social Informatics*. 332-338. doi:10.1109/SocialInformatics.2012.81
- [6] Wang, Y., & Redmiles, D. (2019, 25-31 May 2019). Implicit Gender Biases in Professional Software Development: An Empirical Study. *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*. 1-10. doi:10.1109/ICSE-SEIS.2019.00009
- [7] Zolduoarrati, E., Licorish, S. A., & Stanger, N. (2023). Secondary Studies on Human Aspects in Software Engineering: A Tertiary Study. *Journal of Systems and Software*, 200, 111654. doi:10.1016/j.jss.2023.111654
- [8] Odiete, O., Jain, T., Adaji, I., Vassileva, J., & Deters, R. (2017). Recommending Programming Languages by Identifying Skill Gaps Using Analysis of Experts. A Study of Stack Overflow. *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization*. 159–164. doi:10.1145/3099023.3099040
- [9] Zolduoarrati, E., Licorish, S. A., & Stanger, N. (2022). Impact of Individualism and Collectivism Cultural Profiles on the Behaviour of Software Developers: A Study of Stack Overflow. *Journal of Systems and Software*, 192, 111427. doi:10.1016/j.jss.2022.111427

- [10] Graham, M., De Sabbata, S., & Zook, M. A. (2015). Towards a Study of Information Geographies: (Im)Mutable Augmentations and a Mapping of the Geographies of Information. *Geo: Geography and Environment*, 2(1), 88-105. doi:10.1002/geo.2.8
- [11] Licorish, S. A., & MacDonell, S. G. (2015). Communication and Personality Profiles of Global Software Developers. *Information and Software Technology*, 64(C), 113–131. doi:10.1016/j.infsof.2015.02.004
- [12] Licorish, S. A., & MacDonell, S. G. (2014). Personality Profiles of Global Software Developers. *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*. Article 45. doi:10.1145/2601248.2601265
- [13] Licorish, S. A., & MacDonell, S. G. (2017). Exploring Software Developers' Work Practices: Task Differences, Participation, Engagement, and Speed of Task Resolution. *Information & Management*, 54(3), 364-382. doi:10.1016/j.im.2016.09.005
- [14] Bacchelli, A., Ponzanelli, L., & Lanza, M. (2012). Harnessing Stack Overflow for the Ide. *2012 Third International Workshop on Recommendation Systems for Software Engineering (RSSE)*. 26-30. doi:10.1109/RSSE.2012.6233404
- [15] Shao, B., & Yan, J. (2017). Recommending Answerers for Stack Overflow with Lda Model. *Proceedings of the 12th Chinese Conference on Computer Supported Cooperative Work and Social Computing*. 80–86. doi:10.1145/3127404.3127426
- [16] Oliveira, N., Muller, M., Andrade, N., & Reinecke, K. (2018). The Exchange in Stackexchange: Divergences between Stack Overflow and Its Culturally Diverse Participants. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW), Article 130. doi:10.1145/3274399
- [17] Zolduoarrati, E., Licorish, S. A., & Grundy, J. (2024). *A Cross-Continental Analysis of How Regional Cues Shape Developers' Stack Overflow Contributions – Replication Package* [Data set]. Zenodo. doi:10.5281/zenodo.13994284
- [18] Shi, Y., Chen, S., & Kang, L. (2021). Which Questions Are Valuable in Online Q&a Communities? A Question's Position in a Knowledge Network Matters. *Scientometrics*, 126(10), 8239-8258. doi:10.1007/s11192-021-04135-2
- [19] Nivala, M., Seredko, A., Osborne, T., & Hillman, T. (2020, 27-30 April 2020). Stack Overflow – Informal Learning and the Global Expansion of Professional Development and Opportunities in Programming? *2020 IEEE Global Engineering Education Conference (EDUCON)*. 402-408. doi:10.1109/EDUCON45650.2020.9125165
- [20] Kabir, S., Udo-Imeh, D. N., Kou, B., & Zhang, T. (2024). Is Stack Overflow Obsolete? An Empirical Study of the Characteristics of Chatgpt Answers to Stack Overflow Questions. *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Article 935. doi:10.1145/3613904.3642596
- [21] Xu, B., Nguyen, T. D., Le-Cong, T., Hoang, T., Liu, J., Kim, K., . . . Lo, D. (2023, 11-15 Sept. 2023). Are We Ready to Embrace Generative Ai for Software Q&A? *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 1713-1717. doi:10.1109/ASE56229.2023.00023
- [22] Zagalsky, A., Teshima, C. G., German, D. M., Storey, M.-A., & Poo-Caamaño, G. (2016). How the R Community Creates and Curates Knowledge: A Comparative Study of Stack Overflow and Mailing Lists. *Proceedings of the 13th International Conference on Mining Software Repositories*. 441–451. doi:10.1145/2901739.2901772
- [23] Zhang, H., Wang, S., Chen, T. H., Zou, Y., & Hassan, A. E. (2021). An Empirical Study of Obsolete Answers on Stack Overflow. *IEEE Transactions on Software Engineering*, 47(4), 850-862. doi:10.1109/TSE.2019.2906315
- [24] Barua, A., Thomas, S. W., & Hassan, A. E. (2014). What Are Developers Talking About? An Analysis of Topics and Trends in Stack Overflow. *Empirical Software Engineering*, 19(3), 619-654. doi:10.1007/s10664-012-9231-y
- [25] Uddin, G., Sabir, F., Guéhéneuc, Y.-G., Alam, O., & Khomh, F. (2021). An Empirical Study of Iot Topics in Iot Developer Discussions on Stack Overflow. *Empirical Software Engineering*, 26(6), 121. doi:10.1007/s10664-021-10021-5
- [26] Zolduoarrati, E., Licorish, S. A., & Stanger, N. (2024). Harmonising Contributions: Exploring Diversity in Software Engineering through Cqa Mining on Stack Overflow. *ACM Transactions on Software Engineering and Methodology*. doi:10.1145/3672453
- [27] Hofstede, G. (1980). Culture and Organizations. *International Studies of Management & Organization*, 10(4), 15-41. doi:10.1080/00208825.1980.11656300
- [28] Oliveira, N., Andrade, N., & Reinecke, K. (2016). Participation Differences in Q&a Sites across Countries: Opportunities for Cultural Adaptation. *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. Article 6. doi:10.1145/2971485.2971520
- [29] Meyer, E. (2014). *The Culture Map*. New York City, NY: PublicAffairs.
- [30] Ting-Toomey, S., & Tenzin, D. (2019). *Communicating across Cultures* (2nd ed.). New York City, NY: Guilford Press.
- [31] Inglehart, R., & Baker, W. E. (2000). Modernization, Cultural Change, and the Persistence of Traditional Values. *American Sociological Review*, 65(1), 19-51. doi:10.2307/2657288

- [32] Levine, R. V., & Norenzayan, A. (1999). The Pace of Life in 31 Countries. *Journal of Cross-Cultural Psychology*, 30(2), 178-205. doi:10.1177/0022022199030002003
- [33] Levine, R. (1997). *A Geography of Time: The Temporal Misadventures of a Social Psychologist, or How Every Culture Keeps Time Just a Little Bit Differently*. New York City, NY: Basic Books.
- [34] Menshikova, A. (2018). Evaluation of Expertise in a Virtual Community of Practice: The Case of Stack Overflow. In D. A. Alexandrov, et al. (Eds.), *Digital Transformation and Global Society. International Conference on Digital Transformation and Global Society*. 483-491. doi:10.1007/978-3-030-02843-5\_40
- [35] Choetkiertikul, M., Avery, D., Dam, H. K., Tran, T., & Ghose, A. (2015, 28 Sept.-1 Oct. 2015). Who Will Answer My Question on Stack Overflow? *2015 24th Australasian Software Engineering Conference*. 155-164. doi:10.1109/ASWEC.2015.28
- [36] Brooke, S. J. (2021). Trouble in Programmer's Paradise: Gender-Biases in Sharing and Recognising Technical Knowledge on Stack Overflow. *Information, Communication & Society*, 24(14), 2091-2112. doi:10.1080/1369118X.2021.1962943
- [37] Licorish, S. A. (2013). *Collaboration Patterns of Successful Globally Distributed Agile Software Teams: The Role of Core Developers*. (Doctoral Theses, Auckland University of Technology, Auckland, New Zealand). Retrieved 18 October, 2023, from <https://hdl.handle.net/10292/5973>
- [38] Braun, V., & Clarke, V. (2006). Using Thematic Analysis in Psychology. *Qualitative Research in Psychology*, 3(2), 77-101. doi:10.1191/1478088706qp063oa
- [39] Zolduoarrati, E., & Licorish, S. A. (2021). On the Value of Encouraging Gender Tolerance and Inclusiveness in Software Engineering Communities. *Information and Software Technology*, 139, 106667. doi:10.1016/j.infsof.2021.106667
- [40] Panasyuk, A., Yu, E. S. L., & Mehrotra, K. G. (2019, 30 Jan.-1 Feb. 2019). Improving Geocoding for City-Level Locations. *2019 IEEE 13th International Conference on Semantic Computing (ICSC)*. 416-421. doi:10.1109/ICOSC.2019.8665524
- [41] Peterson, M. P. (2014). Evaluating Mapping Apis. In *Modern Trends in Cartography: Selected Papers of Cartocon 2014* (pp. 183-197): Springer. doi:10.1007/978-3-319-07926-4\_15
- [42] Ostertagova, E., Ostertag, O., & Kováč, J. (2014). Methodology and Application of the Kruskal-Wallis Test. *Applied mechanics and materials*. 115-120. doi:10.4028/[www.scientific.net/AMM.611.115](http://www.scientific.net/AMM.611.115)
- [43] Conover, W. J. (1999). *Practical Nonparametric Statistics* (3 ed.). Hoboken, NJ: John Wiley & Sons.
- [44] Osman, M., Parnell, A. C., & Haley, C. (2017). "Suicide Shall Cease to Be a Crime": Suicide and Undetermined Death Trends 1970–2000 before and after the Decriminalization of Suicide in Ireland 1993. *Irish Journal of Medical Science*, 186(1), 201-205. doi:10.1007/s11845-016-1468-9
- [45] Licorish, S. A., & MacDonell, S. G. (2014). Understanding the Attitudes, Knowledge Sharing Behaviors and Task Performance of Core Developers: A Longitudinal Study. *Information and Software Technology*, 56(12), 1578-1596. doi:10.1016/j.infsof.2014.02.004
- [46] Blanco, G., Pérez-López, R., Fdez-Riverola, F., & Lourenço, A. M. G. (2020). Understanding the Social Evolution of the Java Community in Stack Overflow: A 10-Year Study of Developer Interactions. *Future Generation Computer Systems*, 105, 446-454. doi:10.1016/j.future.2019.12.021
- [47] Krzysztof, S., & Mikołaj, M. (2014). Signed Graphs. In R. Alhajj & J. Rokne (Eds.), *Encyclopedia of Social Network Analysis and Mining* (pp. 1726-1734). New York, NY: Springer New York. doi:10.1007/978-1-4614-6170-8\_251
- [48] Truong, Q. D., Truong, Q. B., & Dkaki, T. (2016). Graph Methods for Social Network Analysis. In P. C. Vinh & L. Barolli (Eds.), *Nature of Computation and Communication*. 276-286. doi:10.1007/978-3-319-46909-6\_25
- [49] Chowdhury, R. R., Gupta, S., & Chede, S. (2021). World War Iii Analysis Using Signed Social Networks. *Social Network Analysis and Mining*, 11(1), 107. doi:10.1007/s13278-021-00822-3
- [50] Liu, Z., Xia, Y., Liu, Q., He, Q., Zhang, C., & Zimmermann, R. (2018). Toward Personalized Activity Level Prediction in Community Question Answering Websites. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 14(2s), Article 41. doi:10.1145/3187011
- [51] Ting, S. L., & Tsang, A. H. C. (2014). Using Social Network Analysis to Combat Counterfeiting. *International Journal of Production Research*, 52(15), 4456-4468. doi:10.1080/00207543.2013.861947
- [52] Krupa, M., Cenek, M., Powell, J., & Trammell, E. J. (2018). Mapping the Stakeholders: Using Social Network Analysis to Increase the Legitimacy and Transparency of Participatory Scenario Planning. *Society & Natural Resources*, 31(1), 136-141. doi:10.1080/08941920.2017.1376140
- [53] Mohammad, S., Dunne, C., & Dorr, B. (2009). Generating High-Coverage Semantic Orientation Lexicons from Overtly Marked Words and a Thesaurus. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 2 - Volume 2*. 599–608. Retrieved 4 April, 2024, from <https://aclanthology.org/D09-1063>
- [54] Fruchterman, T. M. J., & Reingold, E. M. (1991). Graph Drawing by Force-Directed Placement. *Software: Practice and Experience*, 21(11), 1129-1164. doi:10.1002/spe.4380211102
- [55] Jonker, D., Langevin, S., Giesbrecht, D., Crouch, M., & Kronenfeld, N. (2016). Graph Mapping: Multi-Scale Community Visualization of Massive Graph Data. *Information Visualization*, 16(3), 190-204. doi:10.1177/1473871616661195



- [56] Jänicke, S., Heine, C., Hellmuth, M., Stadler, P. F., & Scheuermann, G. (2010). Visualization of Graph Products. *IEEE Transactions on Visualization and Computer Graphics*, 16(6), 1082-1089. doi:10.1109/TVCG.2010.217
- [57] Ahmed, S., & Bagherzadeh, M. (2018). What Do Concurrency Developers Ask About? A Large-Scale Study Using Stack Overflow. *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*. Article 30. doi:10.1145/3239235.3239524
- [58] Gan, J., & Qi, Y. (2021). Selection of the Optimal Number of Topics for Lda Topic Model—Taking Patent Policy Analysis as an Example. *Entropy*, 23(10). doi:10.3390/e23101301
- [59] Sokolovsky, A., Gross, T., & Bacardit, J. (2021). Is It Feasible to Detect Floss Version Release Events from Textual Messages? A Case Study on Stack Overflow. *PLOS ONE*, 16(2), e0246464. doi:10.1371/journal.pone.0246464
- [60] Stevens, K., Kegelmeyer, P., Andrzejewski, D., & Buttler, D. (2012). Exploring Topic Coherence over Many Models and Many Topics. *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. 952–961. doi:10.5555/2390948.2391052
- [61] Ahasanuzzaman, M., Asaduzzaman, M., Roy, C. K., & Schneider, K. A. (2016). Mining Duplicate Questions in Stack Overflow. *Proceedings of the 13th International Conference on Mining Software Repositories*. 402–412. doi:10.1145/2901739.2901770
- [62] Wu, Y., Zhang, Q., & Huang, X.-J. (2011). Efficient near-Duplicate Detection for Q&a Forum. *Proceedings of 5th International Joint Conference on Natural Language Processing*. 1001-1009. Retrieved 25 October, 2023, from <https://aclanthology.org/I11-1112.pdf>
- [63] Wei, X., & Croft, W. B. (2006). Lda-Based Document Models for Ad-Hoc Retrieval. *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. 178–185. doi:10.1145/1148170.1148204
- [64] Lukins, S. K., Kraft, N. A., & Etzkorn, L. H. (2010). Bug Localization Using Latent Dirichlet Allocation. *Information and Software Technology*, 52(9), 972-990. doi:10.1016/j.infsof.2010.04.002
- [65] Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. *The Journal of Machine Learning Research*, 3, 993–1022. Retrieved 24 October, 2023, from <https://ai.stanford.edu/~ang/papers/jair03-lda.pdf>
- [66] Naim, S. M., Boedihardjo, A. P., & Hossain, M. S. (2017, 11-14 Dec. 2017). A Scalable Model for Tracking Topical Evolution in Large Document Collections. *2017 IEEE International Conference on Big Data (Big Data)*. 726-735. doi:10.1109/BigData.2017.8257988
- [67] Chuang, J., Manning, C. D., & Heer, J. (2012). Termite: Visualization Techniques for Assessing Textual Topic Models. *Proceedings of the International Working Conference on Advanced Visual Interfaces*. 74–77. doi:10.1145/2254556.2254572
- [68] Tahaei, M., Vaniea, K., & Saphra, N. (2020). Understanding Privacy-Related Questions on Stack Overflow. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14. doi:10.1145/3313831.3376768
- [69] Chaffer, J. (2007). *Learning JQuery: Better Interaction Design and Web Development with Simple Javascript Techniques* (1st ed.). Birmingham, United Kingdom: Packt Publishing.
- [70] Zhu, E. (1996). Meaning Negotiation, Knowledge Construction, and Mentoring in a Distance Learning Course. Retrieved 18 October, 2023, from <https://eric.ed.gov/?id=ED397849>
- [71] Henri, F. (1992). Computer Conferencing and Content Analysis. In A. R. Kaye (Ed.) *Collaborative Learning Through Computer Conferencing*. 117-136. doi:10.1007/978-3-642-77684-7\_8
- [72] Licorish, S. A., & MacDonell, S. G. (2021). Self-Organising Roles in Agile Globally Distributed Teams. *arXiv preprint arXiv:2106.06154*. doi:10.48550/arXiv.2106.06154
- [73] Licorish, S. A., & MacDonell, S. G. (2013, 25-25 May 2013). Differences in Jazz Project Leaders' Competencies and Behaviors: A Preliminary Empirical Investigation. *2013 6th International Workshop on Cooperative and Human Aspects of Software Engineering (CHASE)*. 1-8. doi:10.1109/CHASE.2013.6614725
- [74] Licorish, S. A., & MacDonell, S. G. (2013). Adopting Softer Approaches in the Study of Repository Data: A Comparative Analysis. *Proceedings of the 17th International Conference on Evaluation and Assessment in Software Engineering*. 240–245. doi:10.1145/2460999.2461035
- [75] Vourvachis, P. (2007). On the Use of Content Analysis (Ca) in Corporate Social Reporting (Csr): Revisiting the Debate on the Units of Analysis and the Ways to Define Them. *British Accounting Association Annual Conference 2007*. Retrieved 22 September, 2023, from <https://eprints.kingston.ac.uk/id/eprint/4129>
- [76] Anandarajan, M., Hill, C., & Nolan, T. (2019). The Fundamentals of Content Analysis. In *Practical Text Analytics: Maximizing the Value of Text Data* (pp. 15-25). Cham: Springer International Publishing. doi:10.1007/978-3-319-95663-3\_2
- [77] Ghasemi, N., Fatourehchi, R., & Momtazi, S. (2021). User Embedding for Expert Finding in Community Question Answering. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(4), Article 70. doi:10.1145/3441302

- [78] United States Government Accountability Office (1996). Content Analysis: A Methodology for Structuring and Analyzing Written Material. Retrieved 21 September, 2023, from <https://www.gao.gov/assets/pemd-10.3.1.pdf>
- [79] McHugh, M. L. (2012). Interrater Reliability: The Kappa Statistic. *Biochemia Medica*, 22(3), 276-282. Retrieved 27 September, 2023, from <https://hrcak.srce.hr/89395>
- [80] Neuendorf, K. A. (2017). *The Content Analysis Guidebook* (2nd ed.). Thousand Oaks, CA: SAGE.
- [81] Danielson, W. A., & Mullen, J. J. (1965). A Basic Space Unit for Newspaper Content Analysis. *Journalism Quarterly*, 42(1), 108-110. doi:10.1177/107769906504200114
- [82] Agresti, A. (2007). *An Introduction to Categorical Data Analysis*. Hoboken, NJ: John Wiley & Sons, Inc. doi:10.1002/0470114754
- [83] Blythe, J., McGrath, C., & Krackhardt, D. (1996). The Effect of Graph Layout on Inference from Social Network Data. In F. J. Brandenburg (Ed.) *Graph Drawing. International Symposium on Graph Drawing*. 40-51. doi:10.1007/BFb0021789
- [84] Luce, R. D., & Perry, A. D. (1949). A Method of Matrix Analysis of Group Structure. *Psychometrika*, 14(2), 95-116. doi:10.1007/BF02289146
- [85] Vasques Filho, D., & O'Neale, D. R. J. (2020). Transitivity and Degree Assortativity Explained: The Bipartite Structure of Social Networks. *Physical Review E*, 101(5), 052305. doi:10.1103/PhysRevE.101.052305
- [86] Hoffman, M., Steinley, D., Gates, K. M., Prinstein, M. J., & Brusco, M. J. (2018). Detecting Clusters/Communities in Social Networks. *Multivariate Behavioral Research*, 53(1), 57-73. doi:10.1080/00273171.2017.1391682
- [87] Wu, Z., He, Q., Chen, Q., Xue, H., & Li, S. (2021). A Topical Network Based Analysis and Visualization of Global Research Trends on Green Building from 1990 to 2020. *Journal of Cleaner Production*, 320, 128818. doi:10.1016/j.jclepro.2021.128818
- [88] Sharpe, D. (2015). Chi-Square Test Is Statistically Significant: Now What? *Practical Assessment, Research, and Evaluation*, 20(1), 8. Retrieved 4 April, 2024, from <https://scholarworks.umass.edu/pare/vol20/iss1/8/>
- [89] May, A., Wachs, J., & Hannák, A. (2019). Gender Differences in Participation and Reward on Stack Overflow. *Empirical Software Engineering*, 24(4), 1997-2019. doi:10.1007/s10664-019-09685-x
- [90] Vieira, C. C., & Vasconcelos, M. (2021). Using Facebook Ads Data to Assess Gender Balance in Stem: Evidence from Brazil. *Companion Proceedings of the Web Conference 2021*. 145-153. doi:10.1145/3442442.3453456
- [91] Amirtham S, N., & Kumar, A. (2021). Gender Parity in Stem Higher Education in India: A Trend Analysis. *International Journal of Science Education*, 43(12), 1950-1964. doi:10.1080/09500693.2021.1946867
- [92] Trinkenreich, B., Britto, R., Gerosa, M. A., & Steinmacher, I. (2022). An Empirical Investigation on the Challenges Faced by Women in the Software Industry: A Case Study. *Proceedings of the 2022 ACM/IEEE 44th International Conference on Software Engineering: Software Engineering in Society*. 24-35. doi:10.1145/3510458.3513018
- [93] Joglekar, P. J. (1984). Metropolitan Communications. *IETE Journal of Education*, 25(1), 12-16. doi:10.1080/09747338.1984.11436001
- [94] Haslam, S. A., Oakes, P. J., McGarty, C., Turner, J. C., Reynolds, K. J., & Eggins, R. A. (1996). Stereotyping and Social Influence: The Mediation of Stereotype Applicability and Sharedness by the Views of in-Group and out-Group Members. *British Journal of Social Psychology*, 35(3), 369-397. doi:10.1111/j.2044-8309.1996.tb01103.x
- [95] Shufflebarger Snell, A. M. (2020). A Dialogic Approach to Exploring Culture in Community-Based Adult Esl Classrooms. *TESOL Journal*, 11(1), e00450. doi:10.1002/tesj.450
- [96] Lewandowska-Tomaszczyk, B., & Wilson, P. A. (2021). Expressive and Reserved Cultural Linguistic Schemas: British and American Pride Clusters. In M. Sadeghpour & F. Sharifian (Eds.), *Cultural Linguistics and World Englishes* (pp. 261-293). Singapore: Springer Singapore. doi:10.1007/978-981-15-4696-9\_13
- [97] Harris, J. (2022). Adult English Learners with Limited or Interrupted Formal Education in Diverse Learning Settings. In L. J. Pentón Herrera (Ed.), *English and Students with Limited or Interrupted Formal Education: Global Perspectives on Teacher Preparation and Classroom Practices* (pp. 43-59). Cham: Springer International Publishing. doi:10.1007/978-3-030-86963-2\_4
- [98] McKay-Semmler, K. L. High- and Low-Context Cultures. In *The International Encyclopedia of Intercultural Communication* (pp. 1-5). doi:10.1002/9781118783665.ieicc0106
- [99] Adair, W. L., Buchan, N. R., Chen, X.-P., & Liu, D. (2016). A Model of Communication Context and Measure of Context Dependence. *Academy of Management Discoveries*, 2(2), 198-217. doi:10.5465/amd.2014.0018
- [100] Korac-Kakabadse, N., Kouzmin, A., Korac-Kakabadse, A., & Savery, L. (2001). Low- and High-Context Communication Patterns: Towards Mapping Cross-Cultural Encounters. *Cross Cultural Management: An International Journal*, 8(2), 3-24. doi:10.1108/13527600110797218

- [101] Gözneli, B. (2020). *Identification and Evaluation of a Process for Transitioning from Rest Apis to GraphQL Apis in the Context of Microservices Architecture*. (Master's thesis, Technische Universität München, München). Retrieved 16 July, 2023, from [https://www.matthes.in.tum.de/file/dxn59kqrdvi3/Sebis-Public-Website/-/Master-s-Thesis-Berke-Goezneli/MA\\_thesis\\_Berke\\_Goezneli.pdf](https://www.matthes.in.tum.de/file/dxn59kqrdvi3/Sebis-Public-Website/-/Master-s-Thesis-Berke-Goezneli/MA_thesis_Berke_Goezneli.pdf)
- [102] Cusick, J. J., Prasad, A., & Tepfenhart, W. M. (2008). Global Software Development: Origins, Practices, and Directions. In *Advances in Computers* (Vol. 74, pp. 201-269): Elsevier. doi:10.1016/S0065-2458(08)00606-2
- [103] Groschl, S., & Doherty, L. (2006). The Complexity of Culture: Using the Appraisal Process to Compare French and British Managers in a Uk-Based International Hotel Organisation. *International Journal of Hospitality Management*, 25(2), 313-334. doi:10.1016/j.ijhm.2005.04.002
- [104] Mogavi, R. H., Haq, E.-U., Gujar, S., Hui, P., & Ma, X. (2022). More Gamification Is Not Always Better: A Case Study of Promotional Gamification in a Question Answering Website. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2), Article 452. doi:10.1145/3555553
- [105] Ye, D., Xing, Z., & Kapre, N. (2017). The Structure and Dynamics of Knowledge Network in Domain-Specific Q&a Sites: A Case Study of Stack Overflow. *Empirical Software Engineering*, 22(1), 375-406. doi:10.1007/s10664-016-9430-z
- [106] Vadlamani, S. L., & Baysal, O. (2020, 28 Sept.-2 Oct. 2020). Studying Software Developer Expertise and Contributions in Stack Overflow and Github. *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 312-323. doi:10.1109/ICSME46990.2020.00038
- [107] Kapoor, S., Hughes, P. C., Baldwin, J. R., & Blue, J. (2003). The Relationship of Individualism–Collectivism and Self-Construals to Communication Styles in India and the United States. *International Journal of Intercultural Relations*, 27(6), 683-700. doi:10.1016/j.ijintrel.2003.08.002
- [108] Baker, D. S., & Delpechitre, D. (2013). Collectivistic and Individualistic Performance Expectancy in the Utilization of Sales Automation Technology in an International Field Sales Setting. *Journal of Personal Selling & Sales Management*, 33(3), 277-288. doi:10.2753/PSS0885-3134330303
- [109] Quek, K. M.-T., Knudson-Martin, C., Rue, D., & Alabiso, C. (2010). Relational Harmony: A New Model of Collectivism and Gender Equality among Chinese American Couples. *Journal of Family Issues*, 31(3), 358-380. doi:10.1177/0192513x09351162
- [110] Lambert, N. M., Stillman, T. F., Hicks, J. A., Kamble, S., Baumeister, R. F., & Fincham, F. D. (2013). To Belong Is to Matter: Sense of Belonging Enhances Meaning in Life. *Personality and Social Psychology Bulletin*, 39(11), 1418-1427. doi:10.1177/0146167213499186
- [111] Pande, A. (2017). India and Its Diaspora: Charting New Avenues of Engagement. *International Studies*, 54(1-4), 180-195. doi:10.1177/0020881718777592
- [112] An, Y., Zhang, Y., & Zeng, B. (2015). The Reliable Hub-and-Spoke Design Problem: Models and Algorithms. *Transportation Research Part B: Methodological*, 77, 103-122. doi:10.1016/j.trb.2015.02.006
- [113] Held, F., Hawe, P., Roberts, N., Conte, K., & Riley, T. (2021). Core and Peripheral Organisations in Prevention: Insights from Social Network Analysis. *Health Promotion Journal of Australia*, 32(3), 492-502. doi:10.1002/hpja.374
- [114] Safadi, H., Johnson, S. L., & Faraj, S. (2021). Who Contributes Knowledge? Core-Periphery Tension in Online Innovation Communities. *Organization Science*, 32(3), 752-775. doi:10.1287/orsc.2020.1364
- [115] Calefato, F., Lanubile, F., & Novielli, N. (2018). How to Ask for Technical Help? Evidence-Based Guidelines for Writing Questions on Stack Overflow. *Information and Software Technology*, 94, 186-207. doi:10.1016/j.infsof.2017.10.009
- [116] Lifshitz-Assaf, H. (2018). Dismantling Knowledge Boundaries at Nasa: The Critical Role of Professional Identity in Open Innovation. *Administrative Science Quarterly*, 63(4), 746-782. doi:10.1177/0001839217747876
- [117] Giambra, L. M., & Stone, B. S. (1982). Australian-American Differences in Daydreaming, Attentional Processes, and Curiosity: First Findings Based on Retrospective Reports. *Imagination, Cognition and Personality*, 2(1), 23-35. doi:10.2190/b205-xeqd-tbjr-nqyw
- [118] Hyman, A. (2016). *Global Project Management: The Availability and Applicability of International Project Frameworks When Traversing Geography and Culture*. Retrieved 1 December, 2023, from <http://hdl.handle.net/20.500.12648/6933>
- [119] Sardinha, J. (2011). Highlighting the Contrasts, Downplaying the Divergences: Insertion and Visibility Tactics of Brazilians in Portugal. *Ethnic and Racial Studies*, 34(6), 986-1005. doi:10.1080/01419870.2010.526237
- [120] Canfield, M., Worrell, M., & Gilvarry, C. (2017). Determinants of Substance Use Amongst Brazilians Residing in the UK: The Role of Acculturation. *Drug and Alcohol Review*, 36(6), 751-760. doi:10.1111/dar.12530
- [121] Scottham, K. M., & Dias, R. H. (2010). Acculturative Strategies and the Psychological Adaptation of Brazilian Migrants to Japan. *Identity*, 10(4), 284-303. doi:10.1080/15283488.2010.523587

- [122] El Elj, M., & Abassi, B. (2014). The Determinants of Innovation: An Empirical Analysis in Egypt, Jordan, Syria and Turkey. *Canadian Journal of Development Studies / Revue canadienne d'études du développement*, 35(4), 560-578. doi:10.1080/02255189.2014.934787
- [123] Becheikh, N. (2013). The Impact of Knowledge Acquisition and Absorptive Capacity on Technological Innovations in Developing Countries: Evidence from Egyptian Small and Medium-Sized Enterprises. *Journal of African Business*, 14(3), 127-140. doi:10.1080/15228916.2013.843997
- [124] Lawan Ngoma, A., & Wana Ismail, N. (2013). The Determinants of Brain Drain in Developing Countries. *International Journal of Social Economics*, 40(8), 744-754. doi:10.1108/IJSE-05-2013-0109
- [125] Reboud, S., Mazzarol, T., & Soutar, G. (2014). Low-Tech Vs High-Tech Entrepreneurship: A Study in France and Australia. *Journal of Innovation Economics & Management*, 14(2), 121-141. doi:10.3917/jie.014.0121
- [126] Sink, E. (2006). Small Isvs: You Need Developers, Not Programmers. In *Eric Sink on the Business of Software* (pp. 77-82). Berkeley, CA: Apress. doi:10.1007/978-1-4302-0143-4\_8
- [127] Marion, T. J., Friar, J. H., & Simpson, T. W. (2012). New Product Development Practices and Early-Stage Firms: Two in-Depth Case Studies. *Journal of Product Innovation Management*, 29(4), 639-654. doi:10.1111/j.1540-5885.2012.00930.x
- [128] Cetindamar, D., Lammers, T., & Zhang, Y. (2020). Exploring the Knowledge Spillovers of a Technology in an Entrepreneurial Ecosystem—the Case of Artificial Intelligence in Sydney. *Thunderbird International Business Review*, 62(5), 457-474. doi:10.1002/tie.22158
- [129] Nefae, M. A., Muthaly, S., & Khan, S. (2023). Digital Transformation Strategies for the Sustainable Growth of Startups in Australia. *CARMA 2023 - 5th International Conference on Advanced Research Methods and Analytics*. 27. Retrieved 2 December, 2023, from <http://ocs.editorial.upv.es/index.php/CARMA/CARMA2023/paper/view/16495>
- [130] Krafft, M., Fraser, G., & Walkinshaw, N. (2020). Motivating Adult Learners by Introducing Programming Concepts with Scratch. *Proceedings of the 4th European Conference on Software Engineering Education*. 22–26. doi:10.1145/3396802.3396818
- [131] Mueller, J. (2019). *Functional Programming for Dummies*. Hoboken, NJ: John Wiley & Sons. Retrieved 2 December, 2023, from <https://cir.nii.ac.jp/crid/1130282273306333952>
- [132] Agrawal, J. P., & Farook, O. (2013, 2013/06/23). A Case for Python Scripting in Undergraduate Engineering Technology. doi:10.18260/1-2--19036
- [133] Minardi, A. M. A. F., Kanitz, R., & Bassani, R. H. (2020). An Overview of the Private Equity and Venture Capital Industry in Brazil. In D. Klonowski (Ed.), *Entrepreneurial Finance in Emerging Markets: Exploring Tools, Techniques, and Innovative Technologies* (pp. 145-160). Cham: Springer International Publishing. doi:10.1007/978-3-030-46220-8\_10
- [134] Lokkila, E., Christopoulos, A., & Laakso, M.-J. (2023). A Data-Driven Approach to Compare the Syntactic Difficulty of Programming Languages. *Journal of Information Systems Education*, 34(1), 84-93. Retrieved 2 December, 2023, from <https://aisel.aisnet.org/jise/vol34/iss1/7/>
- [135] Vakaliuk, T. A., Chyzhmotria, O. V., Chyzhmotria, O. H., Didkivska, S. O., & Kontsedailo, V. V. (2023). The Use of Massive Open Online Courses in Teaching the Fundamentals of Programming to Software Engineers. *Educational Technology Quarterly*, 2023(1), 106-120. doi:10.55056/etq.37
- [136] Subramanya, M. H. B. (2017). How Did Bangalore Emerge as a Global Hub of Tech Start-Ups in India? Entrepreneurial Ecosystem - Evolution, Structure and Role., *Journal of Developmental Entrepreneurship*, 22(01), 1750006. doi:10.1142/s1084946717500066
- [137] Sachdev, N., & Singh, K. (2021). Fintech Environment and Funding Activity in India. *Vidyabharati International Interdisciplinary Research Journal*, 13(1), 11-20. Retrieved 2 December, 2023, from <https://www.viirj.org/vol13issue1/3.pdf>
- [138] Cantamessa, M., Gatteschi, V., Perboli, G., & Rosano, M. (2018). Startups' Roads to Failure. *Sustainability*, 10(7), 2346. doi:10.3390/su10072346
- [139] McCrae, R. R., Terracciano, A., De Fruyt, F., De Bolle, M., Gelfand, M. J., Costa Jr., P. T., & Project, C. o. t. A. P. P. o. C. (2010). The Validity and Structure of Culture-Level Personality Scores: Data from Ratings of Young Adolescents. *Journal of Personality*, 78(3), 815-838. doi:10.1111/j.1467-6494.2010.00634.x
- [140] Goddard, I., & Horowitz, J. M. (2023). More Than Half of Americans Say They Took on Leadership Roles When Growing Up. <https://pewrsr.ch/405u4On>
- [141] Lowe, J., & Archibald, C. (2009). Cultural Diversity: The Intention of Nursing. *Nursing Forum*, 44(1), 11-18. doi:10.1111/j.1744-6198.2009.00122.x
- [142] Abramitzky, R., Boustan, L., Jacome, E., & Perez, S. (2021). Intergenerational Mobility of Immigrants in the United States over Two Centuries. *American Economic Review*, 111(2), 580-608. doi:10.1257/aer.20191586
- [143] Gupta-Dame, N., Macdonald, D., Ross-White, A., & Snelgrove-Clarke, E. (2023). Postnatal Experiences of South Asian Immigrant Women in Australia, Canada, the United Kingdom, and the United States: A Qualitative Systematic Review Protocol. *JBI Evidence Synthesis*, 21(6), 1310-1317. doi:10.11124/jbies-22-00402

- [144] Tsai, J. L. (2021). Why Does Passion Matter More in Individualistic Cultures? *Proceedings of the National Academy of Sciences*, 118(14), e2102055118. doi:10.1073/pnas.2102055118
- [145] Yuan, Y. C., Liao, W., & Bazarova, N. N. (2019). Judging Expertise through Communication Styles in Intercultural Collaboration. *Management Communication Quarterly*, 33(2), 238-271. doi:10.1177/0893318918824674
- [146] Barkema, H. G., Chen, X.-P., George, G., Luo, Y., & Tsui, A. S. (2015). West Meets East: New Concepts and Theories. *Academy of Management Journal*, 58(2), 460-479. doi:10.5465/amj.2015.4021
- [147] Gudykunst, W. B., Matsumoto, Y., Ting-Toomey, S., Nishida, T., Kim, K., & Heyman, S. (2006). The Influence of Cultural Individualism-Collectivism, Self Construals, and Individual Values on Communication Styles across Cultures. *Human Communication Research*, 22(4), 510-543. doi:10.1111/j.1468-2958.1996.tb00377.x
- [148] Gupta, M., & Sukanto, K. (2020). Cultural Communicative Styles: The Case of India and Indonesia. *International Journal of Society, Culture & Language*, 8(2), 105-120. Retrieved 3 December, 2023, from [https://www.ijscel.net/article\\_39000.html](https://www.ijscel.net/article_39000.html)
- [149] Hall, M., de Jong, M., & Steehouder, M. (2004). Cultural Differences and Usability Evaluation: Individualistic and Collectivistic Participants Compared. *Technical Communication*, 51(4), 489-503. Retrieved 3 December, 2023, from <https://www.ingentaconnect.com/content/stc/tc/2004/00000051/00000004/art00003>
- [150] MacDonald, H. A., Sulsky, L. M., Spence, J. R., & Brown, D. J. (2013). Cultural Differences in the Motivation to Seek Performance Feedback: A Comparative Policy-Capturing Study. *Human Performance*, 26(3), 211-235. doi:10.1080/08959285.2013.795572
- [151] Aruta, J. J. B. R., Antazo, B. G., & Pacey, J. L. (2021). Self-Stigma Is Associated with Depression and Anxiety in a Collectivistic Context: The Adaptive Cultural Function of Self-Criticism. *The Journal of Psychology*, 155(2), 238-256. doi:10.1080/00223980.2021.1876620
- [152] Bashir, G. M., & Khan, H. U. (2016, 8-10 Sept. 2016). Factors Affecting Learning Capacity of Information Technology Concepts in a Classroom Environment of Adult Learner. *2016 15th International Conference on Information Technology Based Higher Education and Training (ITHET)*. 1-6. doi:10.1109/ITHET.2016.7760729
- [153] Bhattacharya, I., & Sharma, K. (2007). India in the Knowledge Economy – an Electronic Paradigm. *International Journal of Educational Management*, 21(6), 543-568. doi:10.1108/09513540710780055
- [154] Güss, C. D., & Wiley, B. (2007). Metacognition of Problem-Solving Strategies in Brazil, India, and the United States. *Journal of Cognition and Culture*, 7(1-2), 1-25. doi:10.1163/156853707X171793
- [155] Reifer, D. J. (2005). Educating Software Engineers: An Industry Viewpoint. *SIGSOFT Softw. Eng. Notes*, 30(3), 8-9. doi:10.1145/1061874.1061876
- [156] Graziotin, D., & Fagerholm, F. (2019). Happiness and the Productivity of Software Engineers. In C. Sadowski & T. Zimmermann (Eds.), *Rethinking Productivity in Software Engineering* (pp. 109-124). Berkeley, CA: Apress. doi:10.1007/978-1-4842-4221-6\_10
- [157] Kamel, S. H., & Rizk, N. (2019). The Role of Innovative and Digital Technologies in Transforming Egypt into a Knowledge-Based Economy. In M. Habib (Ed.), *Handbook of Research on the Evolution of It and the Rise of E-Society* (pp. 386-400). Hershey, PA, USA: IGI Global. doi:10.4018/978-1-5225-7214-5.ch017
- [158] MacGregor, E., Hsieh, Y., & Kruchten, P. (2005). Cultural Patterns in Software Process Mishaps: Incidents in Global Projects. *Proceedings of the 2005 Workshop on Human and Social Factors of Software Engineering*. 1-5. doi:10.1145/1083106.1083116
- [159] Nakagawa, S. (2004). A Farewell to Bonferroni: The Problems of Low Statistical Power and Publication Bias. *Behavioral Ecology*, 15(6), 1044-1045. doi:10.1093/beheco/arh107
- [160] Grootendorst, M. (2022). Bertopic: Neural Topic Modeling with a Class-Based Tf-Idf Procedure. *arXiv Computation and Language*. doi:10.48550/arXiv.2203.05794
- [161] Tai, R. H., Bentley, L. R., Xia, X., Sitt, J. M., Fankhauser, S. C., Chicas-Mosier, A. M., & Monteith, B. G. (2024). An Examination of the Use of Large Language Models to Aid Analysis of Textual Data. *International Journal of Qualitative Methods*, 23, 16094069241231168. doi:10.1177/16094069241231168
- [162] Oishwee, S. J., Stakhanova, N., & Codabux, Z. (2024). Large Language Model Vs. Stack Overflow in Addressing Android Permission Related Challenges. *Proceedings of the 21st International Conference on Mining Software Repositories*. 373-383. doi:10.1145/3643991.3644933
- [163] Axelborn, H., & Berggren, J. (2023). *Topic Modeling for Customer Insights: A Comparative Analysis of Lda and Bertopic in Categorizing Customer Calls*. (Student thesis). Retrieved 13 August, 2024, from <http://urn.kb.se/resolve?urn=urn:nbn:se:umu:diva-209225>
- [164] Egger, R., & Yu, J. (2022). A Topic Modeling Comparison between Lda, Nmf, Top2vec, and Bertopic to Demystify Twitter Posts. *Frontiers in Sociology*, 7, 886498. doi:10.3389/fsoc.2022.886498
- [165] Lee, K., Claridades, A. R. C., & Lee, J. (2020). Improving a Street-Based Geocoding Algorithm Using Machine Learning Techniques. *Applied Sciences*, 10(16), 5628. doi:10.3390/app10165628

- [166] Bornfeld, B., & Rafaeli, S. (2019). When Interaction Is Valuable: Feedback, Churn and Survival on Community Question and Answer Sites: The Case of Stack Exchange. *Social and Psychological Perspectives in Collaboration Research*. Retrieved 8 August, 2023, from [https://aisel.aisnet.org/hicss-52/cl/social\\_and\\_psychological\\_perspectives/8/](https://aisel.aisnet.org/hicss-52/cl/social_and_psychological_perspectives/8/)
- [167] Pearl, J. (2022). Comment: Understanding Simpson's Paradox. In *Probabilistic and Causal Inference: The Works of Judea Pearl* (Vol. 36, pp. 399–412): Association for Computing Machinery. doi:10.1145/3501714.3501738
- [168] Ameringer, S., Serlin, R. C., & Ward, S. (2009). Simpson's Paradox and Experimental Research. *Nursing Research*, 58(2), 123-127. doi:10.1097/NNR.0b013e318199b517

# APPENDIX A




Figure A1. Social network for United Kingdom top contributors

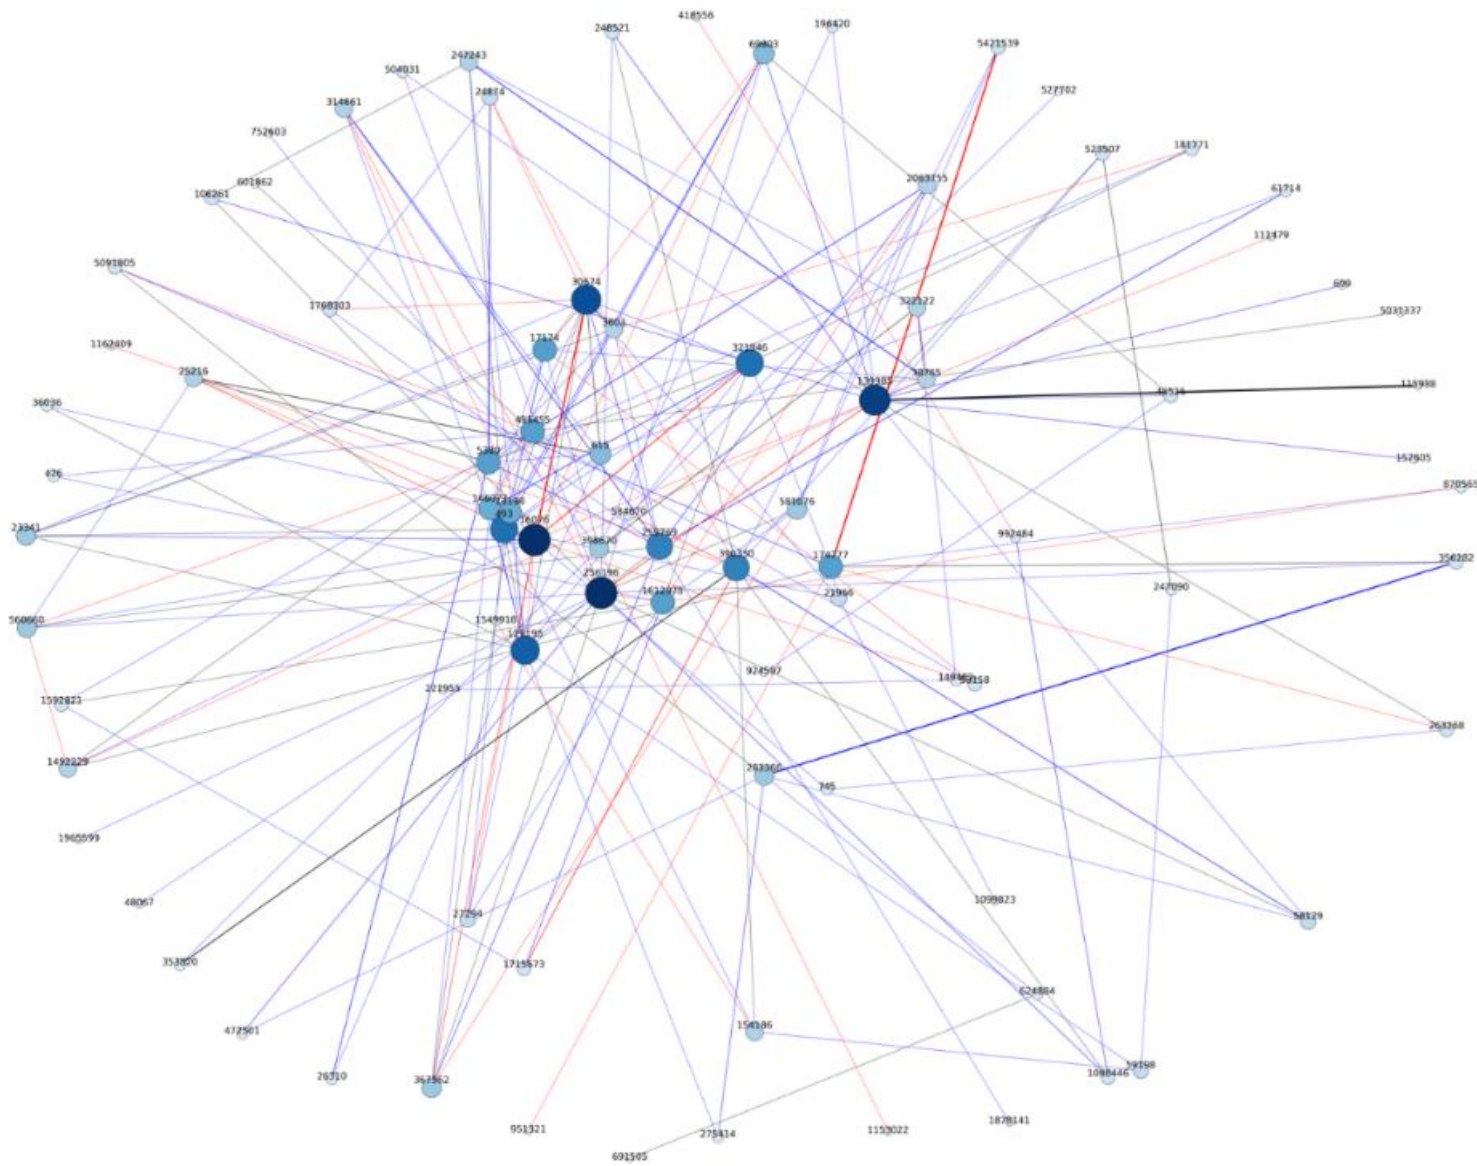


Figure A2. Social network for Australia top contributors



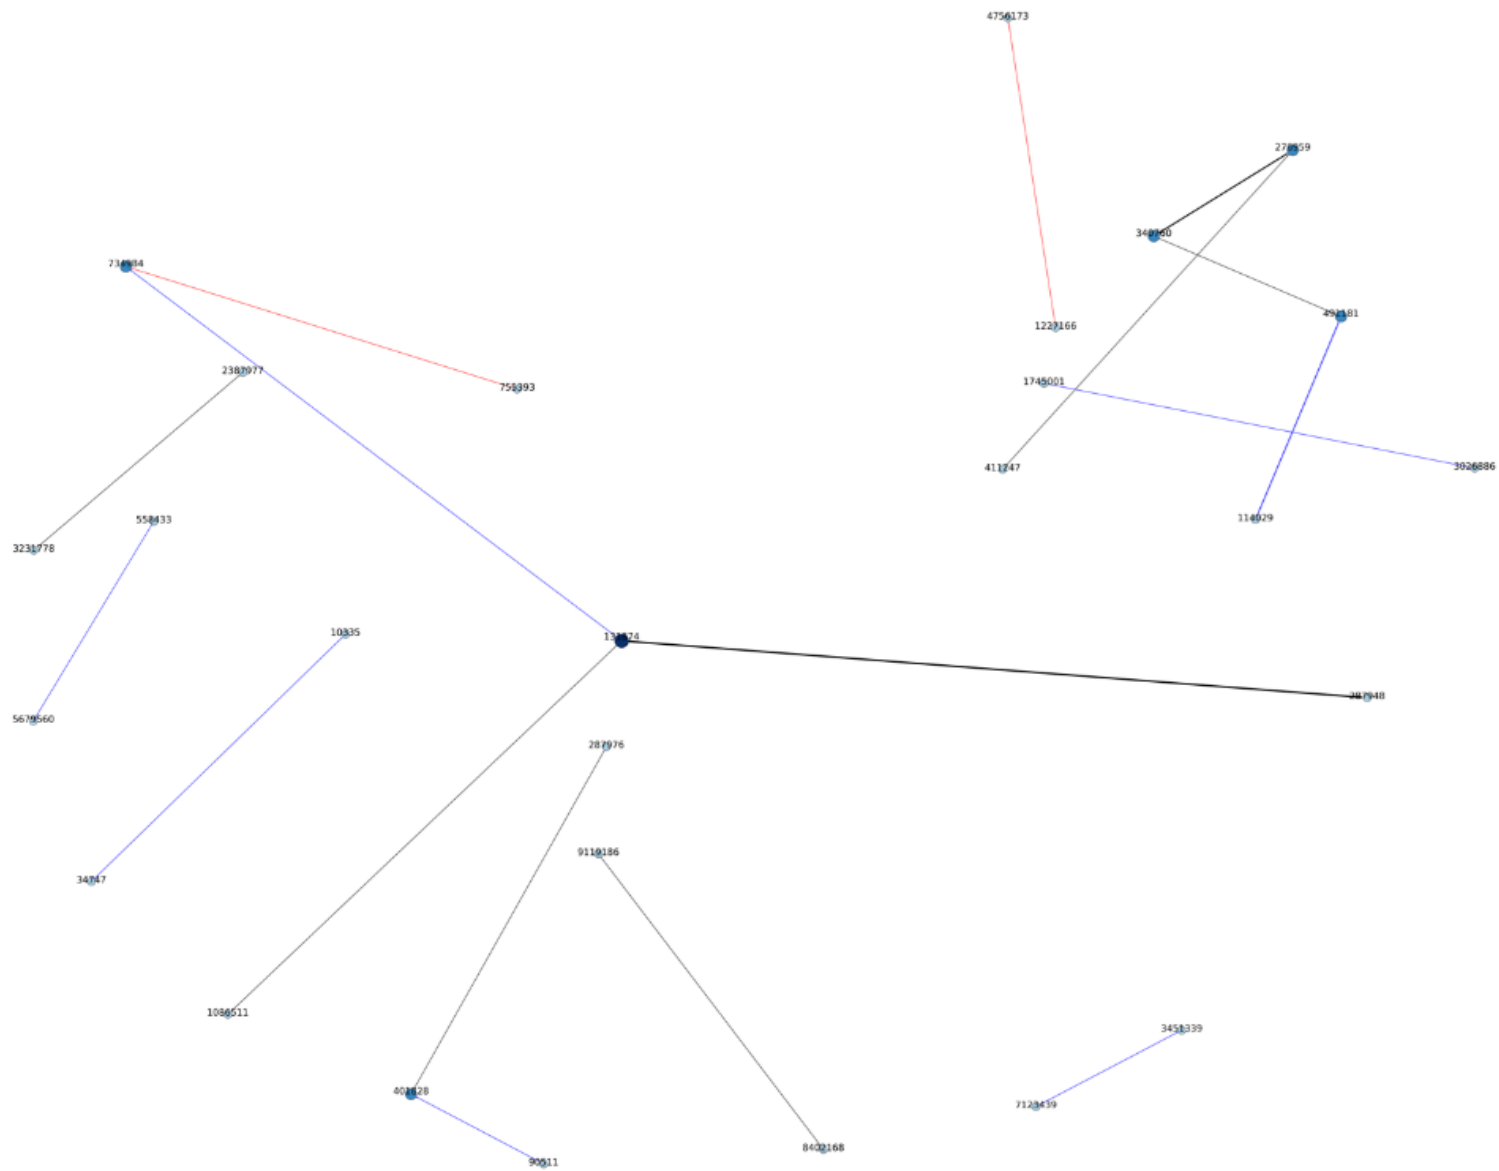


Figure A3. Social network for Brazil top contributors




Figure A4. Social network for Egypt top contributors