Social Network Analysis and Mining (2025) 15:90
https://doi.org/10.1007/s13278-025-01516-w

ORIGINAL ARTICLE

Analysing social network structures and thematic engagement on X audio spaces

Roba Darwish[1,2] · **Alia I. Abdelmoty**[1] · **Liam D. Turner**[1]

Received: 17 April 2025 / Revised: 9 July 2025 / Accepted: 31 July 2025

Abstract
The rise of audio-based platforms such as X Spaces introduces synchronous, ephemeral modes of interaction that differ from traditional text-based social media. This study investigates user connectivity and thematic engagement by analysing a substantial dataset of X Spaces events. We model participation as a bipartite user-Space network, projected into user-user connections to assess community structure and interaction patterns. Our network analysis confirms broad-scale network characteristics and provides novel insights into local structures and user influence. Beyond structural analysis, we examine how users engage thematically across modalities by comparing creator-assigned topics, conversation-derived topics extracted from transcribed audio summaries, and textual posts. Our results reveal that user interests are only weakly aligned across modalities, indicating distinct communicative roles for audio and text. We also introduce a hybrid method combining BERT embeddings, spaCy similarity, and expert validation to assess the alignment between creator-assigned topics and actual conversation content. While most Spaces exhibit high topic coverage, 44% introduce additional themes, suggesting that live audio conversations often diverge from their predefined scope. These findings contribute to understanding interest expression, structural cohesion, and topical drift in emerging audio-based social platforms.

Keywords Social network analysis · Audio-based platforms · Topic classification · X Spaces

1 Introduction

The digital social landscape is rapidly evolving, with platforms such as X Spaces (formerly Twitter Spaces) introducing synchronous, audio-based communication that contrasts with the asynchronous nature of traditional text-based social networks (Li and Penaranda Valdivia 2022; Bajpai et al. 2022; Jung et al. 2022). This shift toward more immersive, real-time interaction enables spontaneous dialogue and immediate feedback, shaping how users engage, present themselves, and form connections online.

Despite growing interest in live audio, most empirical studies in social media research continue to focus on text-based platforms. In contrast, real-time voice communication has been extensively studied in domains such as online gaming, where it is shown to foster interpersonal trust and facilitate efficient collaboration (Jung et al. 2022; Williams et al. 2007). These findings motivate the need for dedicated investigation of how audio-based interaction influences the structure, content, and dynamics of online social networks.

This paper addresses that need through a twofold integrated analysis of X Spaces:

1. Network Analysis: We model the X Spaces ecosystem as a bipartite user-Space network and project this into a user-user graph to explore clustering, influence, and structural variation.
2. Thematic Engagement Analysis: We compare user interests derived from three sources-platform-assigned top-

Alia I. Abdelmoty and Liam D. Turner have contributed equally to this work.

Roba Darwish
DarwishR@cardiff.ac.uk

Alia I. Abdelmoty
AbdelmotyAI@cardiff.ac.uk

Liam D. Turner
TurnerL9@cardiff.ac.uk

[1] School of Computer Science and Informatics, Cardiff University, Senghennydd Rd., CF24 4AG Cardiff, UK

[2] Computer Science and Engineering Department, Yanbu Industrial College, King Faisal Rd., 46435 Yanbu, Saudi Arabia

Published online: 25 August 2025

ics, conversation-derived topics extracted from Otter.ai summaries using TweetNLP, and user-generated textual posts. To assess alignment between assigned and actual conversation content, we propose a hybrid method combining BERT embeddings, spaCy similarity, and expert validation.

A distinctive feature of X Spaces is that creators assign topical labels to their events from a fixed set of platform-defined categories. However, whether these assigned topics accurately reflect the content of the conversation is unknown. We apply a topic classification pipeline to Otter.ai summaries of the conversations and to text-post histories of participating users, allowing for intra- and cross-modal comparison of user interests. We also develop a semantic alignment framework that integrates BERT-based embeddings, spaCy similarity scoring, and expert annotation to measure the match between intended and emergent topics. Compared to our previously published conference version (Darwish et al. 2025), this study significantly expands both the scale and depth of analysis. Our dataset includes 3435 users across 619 Spaces, compared to 834 users previously. We add a third layer of interest comparison using transcribed audio summaries, and offer a more nuanced assessment of topical alignment. These enhancements allow us to explore misalignment across modalities, topical drift during conversations, and variation in local versus global user influence in the X Spaces network.

To guide our investigation, we pose the following questions:

- **RQ1.** Do the broad-scale structural properties of the X Spaces network persist at larger scales, and what do they reveal about patterns of community formation and user participation?
- **RQ2.** How do local network structures vary across users in X Spaces, and what do these differences suggest about user engagement and network position?
- **RQ3.** To what extent do topic similarity patterns among users' direct connections in X Spaces correlate across audio and text modalities?
- **RQ4.** How well do the topics assigned by creators to their X Spaces align with the actual topics discussed during live conversations?

This work makes the following contributions:

1. A large-scale structural analysis of the X Spaces network, confirming broad-scale network characteristics and identifying patterns of uneven participation and clustered interest.
2. A local network analysis using two-hop ego networks, revealing heterogeneity in user roles and influence within the network.
3. A tri-modal interest comparison using platform-assigned, audio-derived, and text-derived topics, showing weak but significant cross-modal alignment.
4. A hybrid semantic mapping framework for evaluating alignment between intended and actual conversation topics, demonstrating that over 44% of Spaces introduce unanticipated themes during discussion.

These findings offer methodological insights for multimodal social network analysis and contribute to a deeper understanding of how communication modality shapes thematic expression and social connectivity.

The remainder of the paper is organised as follows: Sect. 2 surveys related work on audio-based platforms, network structures, and topic classification. Section 3 outlines the dataset and topic extraction. Section 4 outlines the social network analysis methodology, including network construction and the assessment of structural patterns. Section 5 presents our approach to thematic engagement analysis, including methods for comparing topic similarity and interest alignment. Section 6 offers a critical discussion of findings, limitations, and future directions. Section 7 concludes with implications for platform design and multimodal social analytics.

2 Related work

Social media platforms have undergone a profound evolution from text-dominant channels to multimodal spaces supporting visual, video, and more recently, synchronous audio-based interaction. Platforms such as Clubhouse, Discord Stage Channels, and X Spaces (formerly Twitter Spaces) represent a shift toward real-time, ephemeral communication, which challenges traditional models of social engagement and interest modelling (Li and Penaranda Valdivia 2022; Bajpai et al. 2022).

Audio-Based Social Platforms Audio-based platforms allow participants to engage in live, voice-only discussions, often with minimal textual support. These affordances can enable more spontaneous, emotionally expressive, and socially immersive interactions (Jung et al. 2022). Studies from online gaming and livestreaming have shown that voice-based media increases interpersonal trust and coordination (Williams et al. 2007; Jung et al. 2022), yet such findings remain underexplored in the context of social networks like X Spaces.

Two core features shape user behaviour on these platforms: synchronicity, which affords immediate feedback and social presence (Bajpai et al. 2022; Kung-Ming and

Khoon-Seng 2005; Knott 2024; Clark 1993) and ephemerality, which can reduce self-presentation pressure and encourage informal or exploratory conversation (Breznitz 2011; Schlesinger et al. 2017; Bayer et al. 2016).

Despite this, few studies have empirically examined how network structures emerge in these spaces or how thematic expression varies across modalities. Our work contributes to this gap by analysing both the network topology and thematic engagement in X Spaces, while highlighting how the ephemeral and synchronous nature of audio alters interactional dynamics.

Social Network Structure in Online Platforms Despite the distinctive features of audio-based platforms discussed in the preceding section, the fundamental user interactions and relationships established within these settings can still be effectively analysed through the prism of Social Network Analysis (SNA). The principles and methodologies of SNA, which have been successfully employed to comprehend the structure and dynamics of diverse online platforms, including the text-based interactions on X, furnish a valuable framework for investigating X Spaces.

Real-world networks, notwithstanding their considerable scale, frequently exhibit 'small-world' properties, characterised by high interconnectedness amongst neighbours (strong clustering) and relatively short pathways between any two nodes (Villazon-Terrazas et al. 2015; Albert and Barabási 2002; Boccaletti et al. 2006). Prior research has indicated that various real-world network structures often adhere to power-law distributions, a pattern observed in the network of the World Wide Web constructed from web pages (Villazon-Terrazas et al. 2015). Analysis of the structural attributes of the X network (Villazon-Terrazas et al. 2015) has similarly revealed intricate network features and a scale-free nature, displaying significant connectivity and clustering that aligns with a power-law distribution. Comprehending these network properties is crucial for understanding the operational dynamics of X as a social media platform. Furthermore, (Malik 2022) explored a network formed by users and their text posts, discovering through various network metric analyses that its structure also displays scale-free characteristics. Our work builds upon these established SNA approaches to examine the novel social structures emerging within the audio-centric environment of X Spaces.

However, most prior studies focus on text-based or follow-based graphs. Audio-based social spaces-particularly those constructed from co-attendance ties in ephemeral events-have not been studied extensively. Our use of a bipartite user-Space network and ego-centric projections offers insight into how temporary, conversational settings shape persistent network structure. We also investigate local variation in user positioning, reflecting different engagement patterns.

Topic Classification and Cross-Modal Thematic Analysis Understanding how users express and align their interests across modalities is critical for studying emergent social platforms. While prior research in social media content analysis has often focused on topic modelling techniques, such as Latent Dirichlet Allocation (LDA) (Ramage et al. 2010; Mei et al. 2007), for unsupervised discovery of themes in text posts, our work leverages approaches more suited for text classification.

Recent studies have explored differences in thematic content across media types, showing that user expressions often diverge in how topics are framed or discussed in text versus audio or visual formats. For instance, a related investigation on the X platform examined variations in topics discussed in audio messages compared to text and visual posts, finding considerable differences in thematic content based on the media format employed (Li and Penaranda Valdivia 2022). However, such work rarely examines platform-assigned metadata (e.g., X Spaces topic labels), nor does it attempt to semantically align these intended themes with the emergent conversation content.

To achieve this, we utilise robust topic classification techniques. Specifically, we employ the TweetNLP framework (Antypas et al. 2022; Camacho-Collados et al. 2022), which provides state-of-the-art neural models designed to effectively classify topics in informal social media language, capturing nuances like hashtags and abbreviations. Our study applies TweetNLP to classify topics in both users' recent text posts and transcribed X Space conversations, enabling cross-modal thematic comparison.

Our study contributes to this area by developing a hybrid semantic alignment framework-combining BERT embeddings, spaCy similarity, and expert annotation to compare creator-assigned topic labels with conversation-derived topics. This approach accounts for the fact that the two topic sets originate from distinct ontological schemes, and that ephemeral audio conversation may naturally drift away from formal expectations.

Moreover, we explore cross-modal topic similarity by comparing themes derived from users' text posts with those extracted from their audio participation, enabling a triangulated view of user interests. These analyses address ongoing challenges in modelling user heterogeneity and modality-specific behaviours-highlighting how users may adapt their participation to the format and audience of the medium.

3 Dataset

This section details the comprehensive methodology employed in our study, from data acquisition to the analytical frameworks used. A high-level overview of the entire research pipeline is presented in Figure 1.

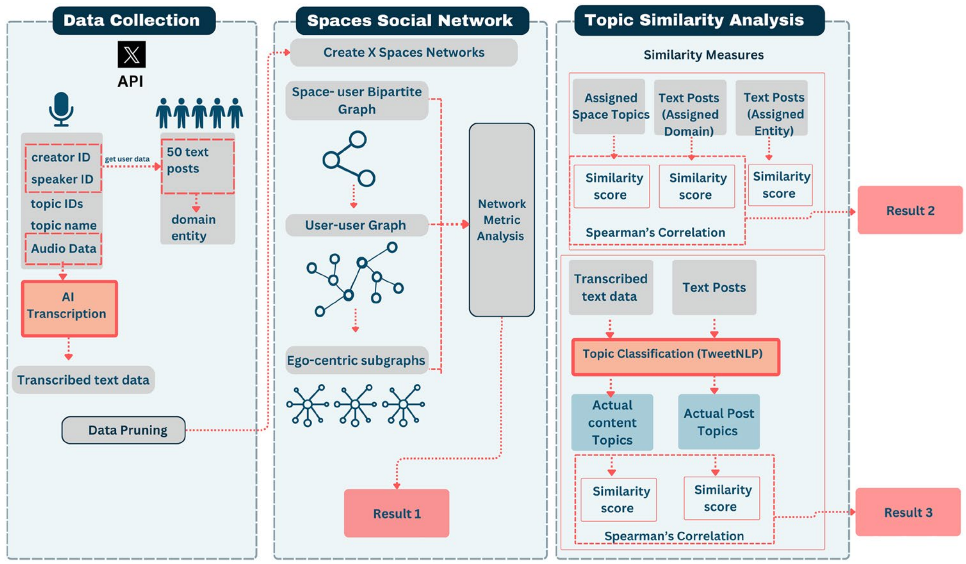

Fig. 1 Overview of the research methodology, illustrating the data collection process, the construction of the X Spaces social network, and the subsequent topic similarity analysis. Adapted from (Darwish et al. 2025)

This study leverages data obtained from the X platform, encompassing both text and audio interactions, collected via the X API over a six-month period from January to June 2024. Our initial dataset included 619 Spaces that had assigned topics across 16 different categories defined by X, including arts, sports, entertainment, technology, business, and news. This broad representation of public interest topics forms the basis of our structural analysis, which includes 3435 unique users who participated as creators, speakers, or hosts within these Spaces. Due to API restrictions, information on the accounts of listeners was not available. Therefore, the scope of this study is viewed through the lens of users that contribute content both via hosting/speaking in Spaces and in the text that they post.

For the subsequent topic similarity analysis, we initially considered a subset of 101 Spaces from the 619. This initial narrowing was based on the practical requirement of focusing on a manageable number of Spaces for in-depth topic analysis. We then applied a duration criterion, excluding four Spaces (including four users not active in other spaces) because their duration was less than 10 min. This resulted in a pool of 97 Spaces that met our basic criteria (assigned topic and sufficient duration).

For each of these 97 Spaces, we proceeded to collect detailed data, including Space IDs, topic IDs, topic names, user IDs, and their roles. These user IDs were then used to retrieve up to 50 of each user's recent text posts within the data collection period. Additionally, we obtained context annotations for these text posts using the X API *(X Developer Platform, June 2022)*, which provided information about the domains such as 'technology' and 'finance,' and entities like 'apple' and 'bitcoin', serving as a representative indicator of users' interests in their text posts.

To ensure the quality and analysability of the text data for this topic analysis, we applied further exclusion criteria to the users within these 97 Spaces: we excluded users who had no text posts associated with their accounts (1 user) and users whose text posts, upon retrieval, contained no extractable contextual information relevant to our analysis (5 users, encompassing 65 text posts). Following these exclusions, the final dataset for topic analysis comprised 97 Spaces and 834 unique users, along with a total of 46,081 text posts, 79 unique domains identified through context annotations, and 3843 unique entities.

3.1 Audio content acquisition

To gain a more comprehensive understanding of user engagement and topic discussion within X Spaces, we incorporated audio content from publicly available recordings. Of the 97 spaces, 86 were transcribed using Otter.ai, an Automatic Speech Recognition (ASR) service. Recordings for the remaining 10 spaces were unavailable.

Otter.ai was selected for several key reasons. First, a comparative evaluation of multiple ASR tools, detailed in (Gaber et al. 2020), identified Otter.ai as having superior accuracy. This accuracy was crucial for our research, enabling reliable transcription of the nuanced and complex discussions within X Spaces, particularly given the multi-speaker format. Second, Otter.ai's commitment to robust data privacy and security practices was a paramount consideration (Otter 2025). Furthermore, Otter.ai provides an automatic summarisation feature that generates concise summaries of the transcribed audio. We utilised these Otter.ai-generated summaries as the input for our topic classification model. These summaries capture the key themes and main points discussed within each Space, providing a condensed representation of the conversation.

The total duration of the transcribed X Spaces recordings was approximately 150 h, with individual sessions ranging from 15 min to 4 h.

Our goal in incorporating audio content was to extract topics based on the actual conversations within the spaces, rather than relying solely on the pre-selected topics chosen by the creators.

3.1.1 Topic analysis of space conversations

To extract meaningful themes from the transcribed X Spaces conversations, we applied the pre-trained topic classification model provided by the TweetNLP library (Antypas et al. 2022). TweetNLP is a Python library that provides state-of-the-art natural language processing (NLP) models specifically trained on social media text. The TweetNLP model's training on X text posts makes it particularly relevant for analysing X Spaces conversations (Antypas et al. 2022). Moreover, its utilisation of social media-specific categories (19 distinct categories, see TweetNLP documentation) enhances our understanding of various discussion themes, facilitating direct comparisons between topics from text posts and those from X Spaces conversations. Specifically, we accessed the model via `tweetnlp.load_model("topic-classification")` (Camacho-Collados et al. 2022). Before applying this model, we processed the concise summaries generated by Otter.ai to ensure they were compatible with the RoBERTa model's input length. The TweetNLP topic classification model is built on the RoBERTa-base architecture, which has a maximum input sequence length of 512 tokens. To ensure compatibility, we used the RoBERTa tokenizer associated with this model to count the number of tokens in each summary and confirmed that none exceeded this limit. While acknowledging that summaries inherently involve some loss of conversational nuance, this step was necessary due to the model's architectural constraint. By keeping the summaries within this limit, we were able to classify the topics in a structured manner, providing a foundation for understanding the thematic structures and topics discussed during the conversations. We then applied the TweetNLP multi-label topic classification model to these summarised transcripts. For each Space, the model assigned one or more relevant topic labels from the 19 available categories based on whether the probability for each topic exceeded the model's internal threshold of 0.5. These assigned sets of topic labels were then used in our subsequent analysis to calculate the similarity of topics across Spaces and between these Space topic sets and user text posts.

4 Social network analysis

Understanding the structural properties of the X Spaces network is crucial for analysing user interactions. In this section, we construct and analyse multiple network representations, including bipartite, user-user, and ego-centric graphs. We apply key network metrics such as clustering coefficients, degree distributions, and centrality measures to assess connectivity, clustering, and overall network structure, providing insights into user engagement patterns. We also examine the small-world characteristics of the user-user network and fit degree distributions to evaluate whether the X Spaces network mirrors other types of online social networks. This section is divided into two parts: methods, which detail the network construction and analysis techniques, and results, which present the findings of our analysis.

4.1 Methods

4.1.1 Network representations

Bipartite Space-User Graph To analyse the social structure of X Spaces, we constructed a bipartite network where nodes belong to two disjoint sets, U and V, with edges only connecting nodes from different sets. Bipartite graphs are particularly useful for representing relationships between two distinct sets of entities, making them well-suited for modelling user participation in X Spaces. According to Anandha et al. (Anandha Jothi and Sankar 2023), bipartite graphs effectively depict such connections, and Barabási (Barabási 2016) identifies them as a fundamental concept in network science. Formally, we define our bipartite graph as $G = (U, V, E)$, where U represents the set of users, each corresponding to a unique account that participated in at least one X Space during our data collection period. V denotes the set of X Spaces, representing individual audio sessions hosted on the platform. Finally, E consists of undirected, unweighted edges representing user participation; an edge (u, v) exists if user u participated in space v, as defined by the X API (see Sect. 3).

Projection into a User-User Graph Although bipartite networks are highly effective for capturing relationships between two distinct sets, such as users and X Spaces, they are often projected into unimodal networks to analyse interactions within a single node type (Hansen et al. 2011). Building on the bipartite graph that represents user participation in X Spaces, we now focus on examining the social structure among users themselves. To achieve this, we project the bipartite graph into a user-user graph. In this projection, nodes represent users, and edges signify co-participation in at least one X Space. Edge weights quantify connection strength, reflecting the number of shared spaces. We construct this graph using all X Space participants, regardless of their roles (e.g., creator, speaker, host), focusing on co-participation to examine interaction frequency and intensity.

Figure 2 provides a simplified illustration of constructing a user-user network from a bipartite Space-User graph. In the bipartite graph (top), nodes are partitioned into two sets: Users and Spaces, where edges indicate user participation in specific X Spaces. This bipartite structure is then projected into a unimodal user-user graph (bottom) to analyse user interactions. In this projection, edges between users represent shared participation in the same X Space.

Ego-Centric Subgraphs To delve deeper into individual user behaviour within the social structure of X Spaces, we examined 2-degree ego-centric subgraphs. This approach, as highlighted by (Golbeck 2013), is particularly useful for revealing individual behaviours within the broader network. We chose 2-degree ego networks to extend beyond immediate connections, capturing a wider local neighbourhood around each user. In the context of X Spaces, this means we analyse not only the users a given individual directly shares Spaces with, but also the connections those individuals have formed through participation in other X Spaces.

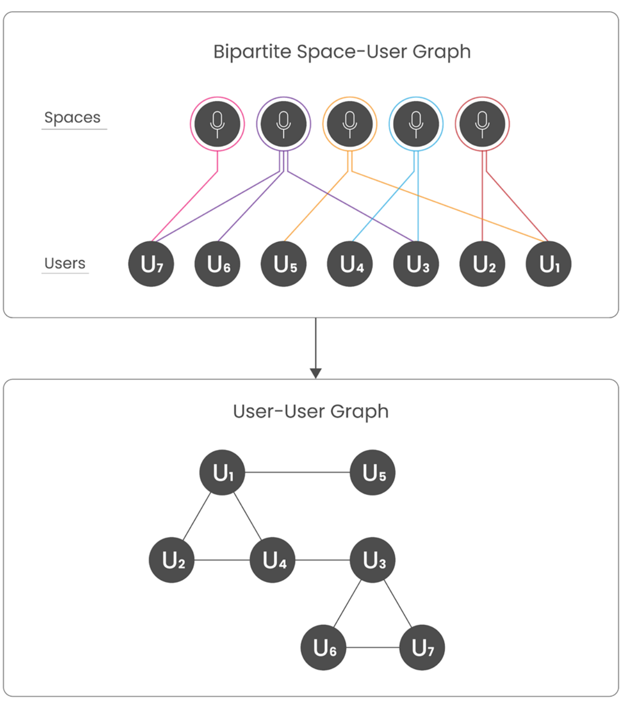

Fig. 2 Illustration of Network Construction. The top panel shows a conceptual bipartite Space-User graph, with connections representing user participation in Spaces. The bottom panel depicts the resulting User-User graph, where edges indicate co-participation in at least one Space

This effectively captures a layer of indirect associations and potential influence, even if there's no direct co-participation in a specific Space between the ego and the extended connections. This approach allows us to look at increasingly larger, but still 'local' neighborhoods around a particular individual in the social network, as noted in prior research (Hansen et al. 2011). This is crucial for understanding the structural aspects of user engagement.

By focusing on these extended connections, we can gain deeper insights into their local network interactions and participation patterns, allowing us to understand the nuances of individual user engagement within the overall network structure.

However, it's important to note that for our subsequent analysis of topic similarity (Sect. 5), we will focus exclusively on the immediate, direct connections between users. This shift in focus allows us to examine the specific topic-related relationships that exist between users who directly share X Spaces. While the 2-degree ego networks provide a valuable perspective on the broader structural context, the analysis of direct connections will enable us to analyse the topical similarity of the user's immediate social circle.

4.1.2 Network analysis

To gain a comprehensive understanding of the structural properties and connectivity patterns of the Spaces Network, we conducted a detailed analysis of several key network metrics. These calculations were performed across the various network representations we constructed: the bipartite user-Space network, the projected user-user network, and the ego-centric subgraphs.

Bipartite Network Metrics For the bipartite user-Space network, we focused on fundamental properties such as the number of nodes (users and Spaces) and the number of edges (user-Space participation connections). These basic counts provide a foundational understanding of the network's size. Additionally, we computed network density, the number of connected components, and degree statistics for both user and Space nodes. To assess local connectivity, we calculated the bipartite clustering coefficient, which measures the likelihood of nodes forming tightly interconnected structures in the absence of regular triangles (Lind et al. 2005). We computed the overall mean clustering coefficient of the bipartite network, as well as the average clustering coefficient for each node type separately to highlight structural differences between users and Spaces.

User-User Network Metrics To analyse the projected user-user network, we calculated a range of network metrics. These included basic network properties (number of nodes and edges), connection strength (average edge weight), network density, fragmentation (number of components), local clustering (transitivity and mean clustering coefficient), and global connectivity. Given that the user-user network was not fully connected, we assessed global connectivity by calculating the mean shortest path length exclusively within its giant component. To evaluate small-world properties, we compared the network's structural characteristics to those of 1000 random Erdős-Rényi (ER) networks with the same number of nodes and edges. The ER model served as a standard random baseline, allowing us to identify deviations from randomness in clustering and path length, as per (Watts and Watts 1998) approach. Specifically, we considered the network to exhibit small-world characteristics if it displayed a significantly higher clustering coefficient and a relatively short mean shortest path length compared to the ER networks.

Ego-Centric Subgraph Metrics For the ego-centric subgraphs, we performed a comprehensive analysis of each user's local network. Specifically, for each ego-centric subgraph, we computed the number of nodes and edges, as these metrics provide a fundamental understanding of the subgraph's size and complexity. We also calculated centrality measures, including degree centrality, betweenness centrality, and closeness centrality, to characterise the unique local network structure surrounding each user. Furthermore, we computed mean node strength, mean degree, clustering coefficient, and density to provide a detailed profile of the structural and cohesive properties inherent within each ego-centric subgraph.

Unweighted measures were used to capture the fundamental structural patterns of the network. This approach focuses on overall connectivity rather than the intensity of interactions. (Abdallah 2011) suggests that unweighted measures can be effective when interactions are 'less focused' or when edge weights are relatively uniform. Furthermore, unweighted measures are preferred for their simplicity, ease of interpretation, and computational efficiency.

Degree Distribution Analysis and Fitting To examine the topological structure and connectivity patterns of the user-user network, we modelled its degree distribution using distributions commonly observed in complex systems: power law, truncated power law, and exponential distributions (Villazon-Terrazas et al. 2015; Malik 2022). This approach aimed to identify the most accurate statistical representation of network connectivity (Turner et al. 2019). As demonstrated in studies of online social networks, including Twitter, analysing degree distributions is a crucial step in understanding the underlying structure of these complex systems (Villazon-Terrazas et al. 2015). For power law and truncated power law fits, we fitted each model to the empirical data with $X_{\min} = 1$, including all degrees from the lowest positive degree. This sets the lower bound for the fit, indicating that all observed degrees were considered in our analysis of these distributions. Subsequently, goodness-of-fit tests were employed to determine how well each model explained the observed degree distribution (Alstott et al. 2014). By fitting these distributions, we aimed to identify whether the Spaces Network exhibits characteristics indicative of scale-free networks (power law), networks with an upper limit on node degree (truncated power law), or networks with a more uniform degree distribution (exponential).

Furthermore, we employed likelihood ratio tests to compare the relative performance of the different models, enabling us to identify the model that provided the most accurate and statistically significant representation of the Spaces Network's degree distribution (Clauset et al. 2009). This comprehensive approach allowed us to gain a deeper understanding of the network's underlying structure and connectivity patterns.

4.2 Results

4.2.1 Network characteristics

Bipartite User-Space Network Analysis To understand user participation patterns on the X Spaces platform, we analysed the bipartite user-Space network. This network comprised 3534 users and 619 Spaces, totalling 4153 nodes. It exhibited a sparse density of 0.003, common in large social networks (Golbeck 2013), indicating limited global connectivity. The network's fragmentation into 101 distinct components suggested a lack of uniform connectivity across the platform. Table 1 summarises these key structural properties. The network displayed a mean clustering coefficient of 0.5181, indicating moderate local clustering. Notably, users exhibited a significantly higher clustering coefficient (0.560) than Spaces (0.345), suggesting that users who participate in the same Spaces tend to form tightly knit communities based on shared interests. This highlights the formation of

Table 1 Bipartite network properties

Property	Value
Number of Nodes	4153 (3534 Users, 619 Spaces)
Number of Edges	6665
Density	0.003
Number of Components	101
Mean Clustering Coefficient (Overall)	0.518[1]
Mean Clustering Coefficient (Users)	0.560
Mean Clustering Coefficient (Spaces)	0.345
Mean Degree (Users)	1.886
Mean Degree (Spaces)	10.767

[1]Computed using NetworkX. The overall clustering coefficient considers all nodes in the bipartite network, while the separate values correspond to the clustering within the user and space node sets

localised user clusters, while Spaces function as hubs that aggregate diverse user groups.

We examined user and Space degree distributions to further explore participation patterns. Users participated in an average of 1.886 Spaces, indicating selective engagement, while Spaces hosted an average of 10.767 users, reflecting multi-user discussions. Figure 3 visually supports these observations, revealing a steep decline in user degree distribution, indicating that most users engage with only a few Spaces, while a small subset of users are highly active. This pattern underscores the variable nature of user engagement on the platform. The relatively uniform average edge weight ($\mu = 1.21$, $\sigma = 1.31$) supports the use of unweighted measures, aligning with our methodological approach, to capture the network's core connectivity.

User-User Network Analysis Building on the insights gained from the bipartite analysis, we proceeded to examine the user-user network to explore direct relational patterns. This network revealed a sparse density of 0.006, consistent with large social networks where not all potential connections are realised (Golbeck 2013), indicating a similar pattern of limited global connectivity. The network's mean degree of 24.19 (Mdn = 17, $\sigma = 23.89$) and high clustering coefficient of 0.85 suggested substantial local clustering and significant variability in user connections, implying the presence of influential hubs. The mean shortest path length of 3.94 (Mdn = 4, $\sigma = 1.15$) indicated that users were typically only a few relational steps apart, despite the network's fragmentation into 123 components, which suggests the existence of distinct communities with limited interaction between them.

The complete network analysis validated our prior findings from the subset of 834 users (Darwish et al. 2025), particularly regarding clustering coefficient, mean degree, and average edge weight. However, the subset exhibited a considerably higher transitivity (0.858) compared to the complete network (0.579) and a substantially higher density (0.05 compared to 0.006). This difference likely arose from random sampling selecting a more cohesive portion of the network. Despite these variations, the overall network topology remained consistent. A comprehensive summary of the network metrics for both the full dataset and the subset of 834 users is provided in Table 2.

The user-user network exhibited characteristics consistent with small-world phenomena. It demonstrated a high clustering coefficient of 0.85, significantly exceeding the average of 0.007 found in 1000 generated Erdős-ényi (ER) random networks. Although the mean shortest path length in the user-user network (3.94) was slightly longer than the ER benchmark (2.86), it remained relatively short considering the network's size. This combination of high clustering and a relatively short mean shortest path length suggests that the user-user network displays small-world characteristics.

Ego-centric Subgraphs Analysis The analysis of two-degree ego-centric subgraphs, constructed for each of the 3534 users within the X Spaces platform, revealed significant variations in local network structure and user behaviour. The structural properties of ego-centric networks in X Spaces exhibited significant variability across users, reflecting a range of network configurations. Figure 4 presents a sample two-degree ego-centric network, illustrating one such configuration. This network, centred on a specific user

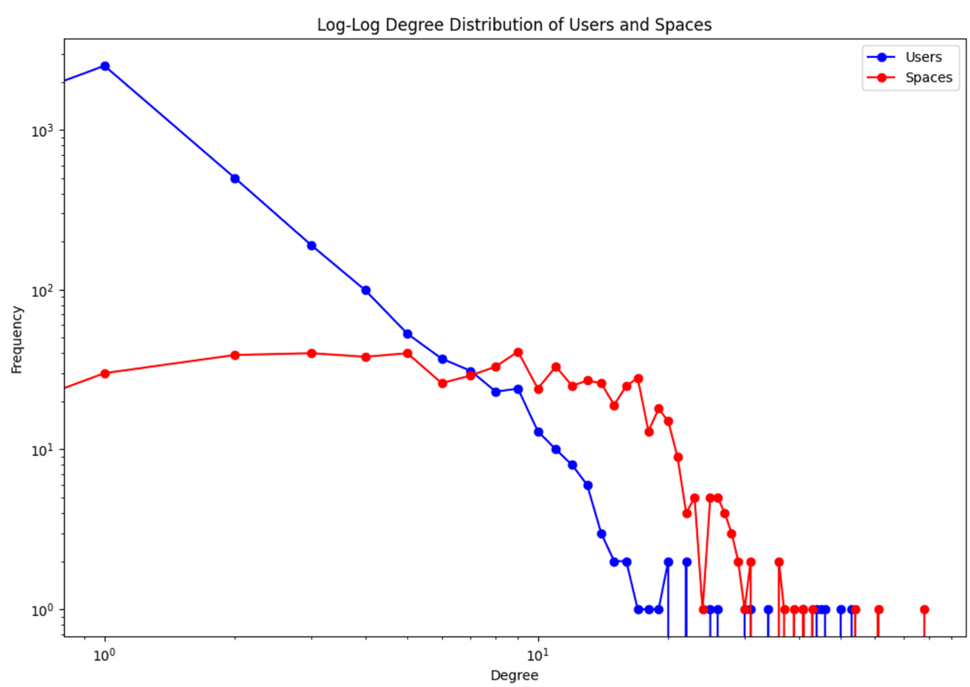

Fig. 3 Log-log degree distribution of users and Spaces in the bipartite network

Table 2 Summary statistics for the user-user network. The left section displays metrics from the full dataset, while the right section shows results from the initial sample. n denotes the number of users; μ represents the mean, Mdn the median, and σ the standard deviation of the key characteristics

Network Properties	Full Dataset (n = 3534)	Initial Sample (n = 834)
Number of Nodes	3534	834
Number of Edges	42,738	18,964
Average Edge Weight	1.20	1.10
Density	0.006	0.05
Transitivity	0.579	0.858
Number of Components	123	27

	Full Dataset (n = 3534)			Initial Sample (n = 834)		
Statistic	μ	Mdn	σ	μ	Mdn	σ
Mean Degree	24.19	17	23.89	25.17	15	25.92
Mean Clustering Coefficient	0.85	1.00	0.25	0.92	1.00	0.17
Mean Shortest Path	3.94	4	1.15	2.97	3.0	1.34

Table 3 Structural properties of ego-centric networks

Statistic	μ	Mdn	σ
Number of Nodes	215	177	205
Number of Edges	3272	2532	3521
Density	0.32	0.15	0.32
Mean Clustering Coefficient	0.86	0.85	0.11
Mean Degree Centrality	0.32	0.15	0.32
Mean Betweenness Centrality	0.008	0.004	0.01
Mean Node Strength	32.12	30.00	19.05

(ego node, dark blue), displays the local connections formed within two degrees of separation and exemplifies the type of network structure observed. Table 3 summarises the key network metrics calculated for these subgraphs, presenting the mean (μ), median (Mdn), and standard deviation (σ) for each statistic.

The average ego network in X Spaces consisted of 215 nodes ($\mu = 215$, Mdn = 177, $\sigma = 205$) and 3272 edges ($\mu = 3272$, Mdn=2532, $\sigma = 3521$). The high standard deviations indicate considerable variability in the size and complexity of users' social circles within X Spaces. Notably, the mean clustering coefficient was high ($\mu = 0.86$, Mdn = 0.85, $\sigma = 0.11$), suggesting strong local interconnectedness within users' immediate social circles and indicating that users sharing connections within X Spaces tend to form closely-knit groups.

In terms of centrality, the mean degree centrality was 0.32 ($\mu = 0.32$, Mdn = 0.15, σ=0.32), and the mean betweenness centrality was 0.008 ($\mu = 0.008$, Mdn = 0.004, σ=0.01). These values, along with the high standard deviations, highlight the variability in user connectivity and bridging roles within their ego networks. The mean node strength, reflecting the average interaction intensity, was 32.12 ($\mu = 32.12$, Mdn = 30.00, $\sigma = 19.05$), suggesting a high level of user engagement.

The density of the ego networks, representing the proportion of possible connections, was 0.32 ($\mu = 0.32$, Mdn = 0.15, $\sigma = 0.32$). This indicates a moderate level of connectedness within the ego networks.

Overall, the results show that while users in X Spaces tend to form highly clustered local communities, there is substantial heterogeneity in the size, complexity, and connectivity of their ego networks.

Degree Distribution and Model Fitting In examining the degree distribution, we compared three models: power law, truncated power law, and exponential. The power law model yielded an α value of 1.28, indicating that node degrees are varied, with some users holding a significantly greater number of connections. This is reflected in the high standard deviation of the mean degree. The truncated power law model, with an α value of 1.00000001, provided

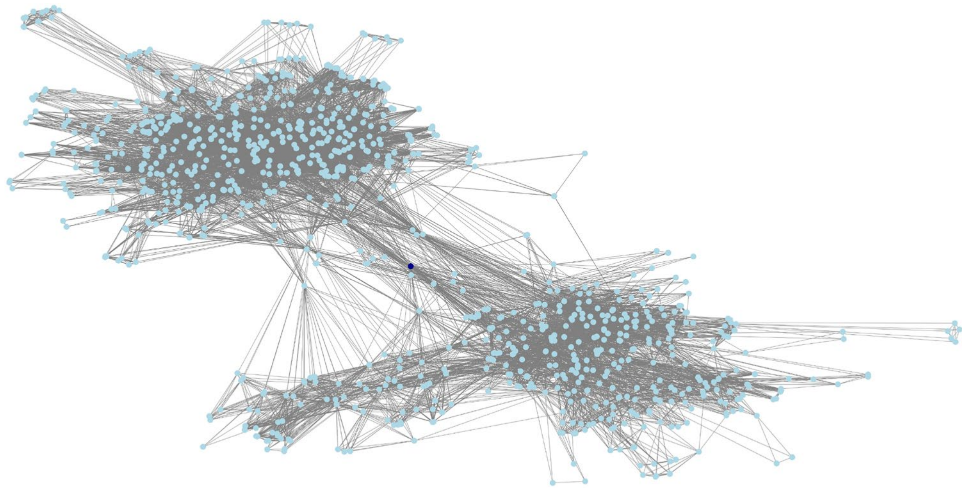

Fig. 4 A sample two-degree ego-centric network from X Spaces, illustrating the interconnectedness of users. This visualisation centres on a specific user (ego node, dark blue) and displays their direct (one-hop) and indirect (two-hop) connections. The network reveals a structure characterised by two distinct clusters of users (light blue nodes), connected by the ego node. This pattern suggests the presence of strong community formations within X Spaces, where users who share common connections tend to form tightly-knit groups

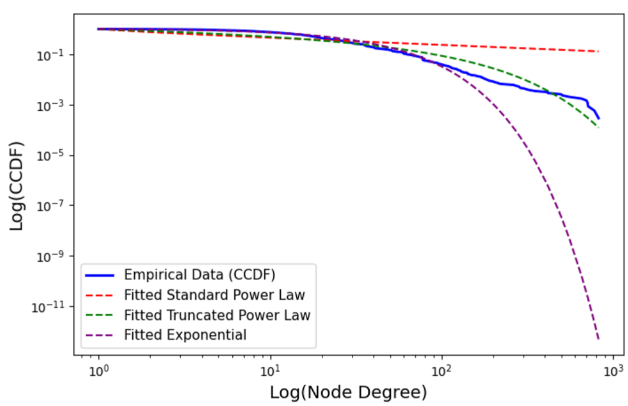

Fig. 5 Degree distribution fitting for the X Spaces network, showing the comparison between the power law, truncated power law, and exponential models. The truncated power law model (α = 1.00000001) provides the best fit, with a natural cutoff in the distribution, as confirmed by likelihood ratio tests

a better fit, suggesting a practical limit to the number of connections users can sustain. On the other hand, the exponential model, with a λ of 0.035, was less effective in capturing the network's structure. Likelihood ratio tests confirmed the superiority of the truncated power law over both the power law and exponential models. Specifically, the truncated power law was significantly favoured over the power law ($LLR = -103.19$, p-value < 0.001), pointing to a natural cutoff in the distribution. Likewise, the comparison with the exponential model ($LLR = -19.53$, p -value < 0.001) further confirmed the truncated power law as the most appropriate fit (See Figure 5).

These results align with the earlier analysis of a smaller data subset, where the truncated power law also showed the best fit ($\alpha = 1.00000004$, $\lambda = 0.040$). The consistency of this fit across both the subset and the complete dataset suggests that the network's distributional properties remain robust despite the increase in data size.

4.2.2 Summary of findings

The following key findings summarise the main observations from our social network analysis.

- *Sparse bipartite network with localised clustering.* The user-Space network exhibits sparse connectivity (density = 0.003) and fragmentation (101 components), suggesting a diverse but not highly interconnected interaction landscape (Table 1). However, a moderate overall clustering coefficient (0.518), with higher clustering among users (0.560) compared to Spaces (0.345), indicates the formation of tightly-knit user groups within specific Spaces, driven by shared interests and fostering localised communities.
- *Uneven user and Space engagement.* The degree distribution reveals selective user engagement, with most participating in few Spaces while a small fraction are highly active (Figure 3). This long-tailed distribution highlights an uneven contribution to platform interactions, potentially reflecting user specialisation or platform design encouraging focus on specific topics.
- *Heterogeneous ego-centric network structures.* Substantial variability in local network configurations (high standard deviations in node and edge counts) is evident from the analysis of ego-centric subgraphs (Table 3). This suggests diverse user roles within their immediate networks, ranging from central to peripheral.
- *Varied user centrality and active engagement.* Centrality measures reveal significant differences in users' degree and betweenness centrality, highlighting some users as highly connected or crucial bridges between communities. The relatively high mean node strength

(32.12) observed in ego-centric networks (Table 3) underscores active user engagement and contribution to platform interactions.

The preceding network analysis has revealed the structural characteristics and engagement patterns within X Spaces, providing a framework for understanding user connections. However, to delve deeper into the underlying drivers of these connections, particularly the role of shared interests, Sect. 5 will examine the similarity of topics associated with user activity within these networks. Specifically, we will investigate the similarity of topics among connected users, both in terms of their Space participation and their textual expressions, to understand the nature of the shared interests driving these network formations.

5 Thematic engagement analysis

5.1 Methods

To determine whether users share similar interests with the people they attend Spaces with, we measured the similarity of the topics associated with the Spaces each user attended. By examining the overlap in Space topic participation, we can assess the extent to which users choose to engage in discussions with others who have attended similar Spaces. Additionally, we calculated the similarity of their interests derived from their text posts. This analysis aims to understand whether users are more likely to interact with like-minded individuals who share common interests in specific Space topics and whether they exhibit common behaviours in their text interactions.

This analysis investigates whether users in X Spaces interact with like-minded individuals by examining the similarity of their Space participation and text post topics. We aim to quantify this similarity within users' ego-centric networks, which represent their direct connections. Specifically, we analyse first-degree connections for a subset of users (n=834) with associated text data. This focus on first-degree connections ensures the topic similarity calculations accurately reflect the ego's direct relationships, minimising the potential influence of unrelated topics introduced by second-degree connections. We analysed 'topics' from Space participation and text posts, using creator-assigned, conversation-extracted (TweetNLP), and X-provided categories. Detailed methodologies for representing these topics are described in the following sections.

For every pair of users, i_1 and i_2, within an ego's network, the Jaccard similarity is calculated as follows:

$$J(T_{i_1}, T_{i_2}) = \frac{|T_{i_1} \cap T_{i_2}|}{|T_{i_1} \cup T_{i_2}|}$$

where T_{i_1} and T_{i_2} represent the sets of topics associated with users i_1 and i_2, respectively. We chose Jaccard similarity for its effectiveness in measuring set overlap, which aligns with our topic representation.

To summarise each ego's network alignment, we computed a mean similarity score:

$$\text{Mean Similarity}_i = \frac{1}{n} \sum_{j=1}^{n} \text{Similarity}(T_i, T_j)$$

where T_i is the ego's topic set, T_j is each alter's topic set, and n is the number of alters.

The similarity calculations described in the Methods section are consistently applied across all analyses-whether examining creator-assigned topics, conversation-derived topics, or text post topics. As such, the following sections apply the Jaccard similarity approach to different topic sets.

5.1.1 Similarity in audio spaces

In this section, we explore the thematic connections within the network based on topics associated with X Spaces. We conduct two distinct analyses to measure topic similarity among users within their ego-centric networks.

Creator-Assigned Topics To evaluate the similarity of interests among users, we assessed the overlap between the topics associated with each ego's attended Spaces-represented by topic IDs-and those of their direct neighbours. The topic set T_i for a user i includes the topic IDs of all Spaces they attended.

Using the Jaccard similarity approach outlined earlier, we quantified the degree of thematic overlap. For example, if User i_1 attended Spaces with topics $\{t_1, t_2, t_3\}$ (representing "Gaming," "News," "Business and Investment") and User i_2 attended Spaces with $\{t_1, t_2, t_4, t_5\}$ (representing "Gaming," "News," "Technology," "Art and Culture"), their Jaccard similarity is calculated as follows:

$$J(T_{i_1}, T_{i_2}) = \frac{|\{t_1, t_2\}|}{|\{t_1, t_2, t_3, t_4, t_5\}|} = \frac{2}{5} = 0.4$$

This similarity score reflects the extent to which users share topical interests based on their attended Spaces. Higher similarity scores suggest stronger thematic alignment, while lower scores indicate a more diverse range of interests.

To better understand the significance of the results, we calculated the mean, median, and standard deviation of similarity scores across all ego networks. These statistics provide insight into the distribution and variability of thematic alignment in the network. A high mean similarity score suggests strong thematic alignment, while a larger standard deviation indicates more diverse thematic engagement among users.

Conversation-Extracted Topics Following the analysis of similarities based on creator-assigned topics, we extended our investigation to include topics derived directly from the content of the conversations held in X Spaces. This approach enables us to move beyond the limitations of pre-assigned topics and gain insight into the actual themes discussed during real-time interactions.

The primary motivation for incorporating audio content into our study was to extract and analyse topics directly from the conversations themselves, thereby going beyond the constraints of pre-selected topics provided by creators and gaining a deeper understanding of the authentic interests and discussions within the Spaces. This approach allows us to examine how accurately the pre-assigned topics reflect the actual content of the discussions, offering a comparison between the intended and emergent themes. Our hypothesis is that the conversations would provide a more realistic view of the topics discussed in each Space, which may or may not overlap with the assigned topics.

Building on the topic extraction from individual Space conversations described earlier (see Sect. 3.1.1), we used the sets of topic labels extracted from each Space using the TweetNLP classification model to assess thematic alignment between users. For each user, we constructed a topic set by aggregating the TweetNLP-assigned labels from all the Spaces they attended. These aggregated sets capture the themes users were exposed to during live conversations. Similar to the analysis using creator-assigned topics, we employed the Jaccard similarity index to quantify the overlap between topic sets. However, in this case, the sets T_{i_1} and T_{i_2} represent the topic categories derived from the conversations of users i_1 and i_2, respectively, using the TweetNLP model. We then calculated the Jaccard similarity between each ego and their direct neighbours within the network, following the same approach as in the previous subsection. Additionally, to offer a complementary perspective on thematic overlap less sensitive to differing topic set sizes, we also computed the Overlap Coefficient for these conversation-derived topic sets.

5.1.2 Similarity in text posts

In addition to analysing topic similarity within X Spaces, we explored whether users who participate in Spaces with similar themes also express aligned interests in their public text-based posts. This parallel analysis enables a broader investigation into the consistency of user interests across distinct modes of engagement on the same platform. To achieve this, we implemented two complementary methods for extracting and comparing the topical content of user text posts. The first draws on X's native semantic annotations, which assign domain and entity labels to user-generated content (Tweet Annotation 2024). This approach, previously introduced in (Darwish et al. 2025), allows for the comparison of user interests in text posts based on semantic features extracted directly from the X platform. The second applies the TweetNLP classification model to extract overarching themes from user text posts, providing a basis for comparison with the topics identified in X Spaces conversations.

X-Provided Domains and Entities The identification of user interests benefits significantly from examining both domains and entities. Entities found in users' text posts are particularly insightful for discovering what they are interested in, as these directly indicate the specific subjects of their engagement (Shen et al. 2013). As demonstrated by researchers (Michelson and Macskassy 2010), the detection and linking of named entities in user text plays a crucial role in identifying interest topics and developing a user's topic profile. While entities offer specificity, domains provide a broader classification, grouping related entities under common themes. By examining both domains and entities, we aimed to capture a comprehensive view of user interests, ranging from general themes to particular subjects. By representing each user's interests as a collection of their domains and entities, we measured the overlap of these elements to assess the similarity of interest between a user and their immediate connections. This mirrors the analytical strategy used for Space topics.

TweetNLP-Extracted Topics To further complement the insights from X-provided domains and entities, we also employed an alternative approach by applying the TweetNLP topic classification model to the same set of user text posts. The text posts were cleaned and preprocessed using the standard procedures implemented within the TweetNLP model, ensuring consistency in the input data for topic extraction. While X's domains offer a thematic categorisation of user interests, exploring TweetNLP's classification allows us to capture user interests through a distinct, predefined framework of social media-centric topics. This provides a different lens on user interests, potentially highlighting overarching themes that might be organised differently or captured with varying emphasis compared to X's internally derived domains. Following the same aggregation and similarity computation method used for X-provided data, we constructed a set of TweetNLP-assigned topic labels for each user based on all the labels in their posts and compared these sets to those of their network neighbours. For these comparisons, the Overlap Coefficient was also calculated as a complementary metric to provide a robust measure of thematic overlap, particularly where topic set cardinalities may vary.

This dual approach to modelling text-based interests-one leveraging X's semantic annotations (domains/entities) and the other employing an external classification model with its own thematic structure (TweetNLP topics)-enables a comparative assessment of how different frameworks reflect user interests and how these relate to audio-based participation in Spaces. A comparison of the similarity scores derived from both approaches is presented in the results section.

5.1.3 Cross-platform similarity and correlation

We assessed the correlation between Space discussion topics and user text post topics using Spearman's rank correlation coefficient (ρ) (Silva et al. 2021), a non-parametric measure suitable for our non-normally distributed Space topic similarity data, as confirmed by Kolmogorov-Smirnov (KS) tests.

We examined the correlation between pre-assigned Space topics and user text post domains, as in our previous work (Darwish et al. 2025). This analysis was supported by the non-normal distribution of pre-assigned Space topic similarities,, evidenced by a KS statistic of $D = 0.2098$ ($p < 0.00001$, Figure 6a). Spearman's ρ was used to assess the strength and direction of the association between the ranked mean Jaccard similarities of these two variables. Furthermore, the correlation between conversation-derived Space topics and text post topics (TweetNLP-extracted) was examined using Spearman's ρ. This choice was justified by the non-normal distribution of the conversation-derived Space topic similarity data, supported by a KS statistic of $D = 0.2013$ ($p < 0.00001$, Figure 6b). Spearman's ρ was computed on the ranked mean Jaccard similarities of these two variables to determine the strength and direction of their association.

For both correlation analyses, a positive ρ indicates a tendency for the ranked similarities of both topic types to increase together, while a negative ρ suggests an inverse relationship. The statistical significance of all correlations was evaluated at a significance level of $p < 0.05$, with values below this threshold considered unlikely to have occurred by chance.

5.1.4 Alignment of intended and actual topics

This section presents the methodology developed to quantify the alignment between the intended topics of X Spaces, as assigned by their creators, and the actual themes that emerged during live audio conversations. This analysis addresses a key methodological challenge: comparing two sets of topics derived from fundamentally different classification schemes-one defined by platform taxonomy and the other extracted from conversational content.

The aim of this alignment analysis is to characterise the relationship between pre-defined conversational intents and the spontaneous thematic developments that occur in real-time discussions. By systematically measuring the overlap between intended and emergent topics, we assess the extent to which Spaces remain on-topic or diverge into novel, unanticipated directions.

Creator-Assigned Topics: These topics represent the intended focus of each X Space, designated by the creator at the time of session setup. During the creation process, users select one or more topics from a predefined set offered by the Twitter API. These labels are used for content discovery and serve as indicators of the conversation's advertised focus.

Conversation-Derived Topics: To capture the actual thematic content, we applied the TweetNLP topic classification model to the transcribed summaries of Space conversations. This model assigns categories based on the linguistic characteristics of the spoken content, producing emergent topic labels that reflect the discussed themes.

Comparability Challenge and Hybrid Alignment Framework

A core methodological challenge in this analysis arises from the incompatible classification schemes underlying the two topic sets. Creator-assigned labels are drawn from a curated platform-specific taxonomy (the Twitter API topic list), whereas conversation-derived categories are generated independently by the TweetNLP classifier and may exhibit different granularity, terminology, or thematic boundaries. To enable a meaningful comparison across

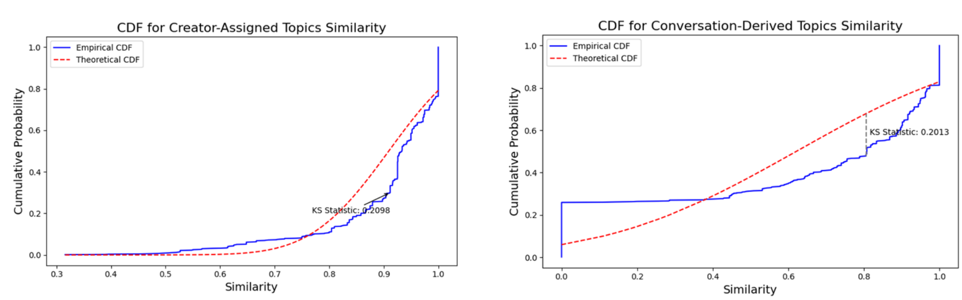

(a) CDF for creator-assigned topics (KS = 0.2098).

(b) CDF for conversation-derived topics (KS = 0.2013).

Fig. 6 CDF plots illustrating the deviations between the empirical distribution of Space topic similarity and a theoretical normal distribution for creator-assigned topics (a, adapted from (Darwish et al. 2025)) and conversation-derived (b), confirming non-normality these heterogeneous sources, we developed a three-stage hybrid semantic alignment framework that integrates:

1. *Contextual Semantic Similarity (BERT-Based)* We first applied the Sentence-BERT model `all-MiniLM-L6-v2`, which has demonstrated state-of-the-art performance in semantic similarity tasks (Reimers and Gurevych 2019). Both creator-assigned and conversation-derived topic labels were encoded into vector representations, and cosine similarity was computed across all label pairs. Pairs with a similarity score of ≥ 0.50 were automatically accepted as semantically aligned. This threshold was selected based on prior benchmarking and empirical validation to ensure both inclusiveness and precision.
2. *Lexical Semantic Similarity (spaCy-Based)* For topic pairs not meeting the initial BERT similarity threshold, we employed spaCy's `en_core_web_lg` model as a complementary semantic comparison method. This model generates vector representations based on lexical co-occurrence patterns and syntactic features. We computed topic similarity by averaging word vectors for each phrase and measuring cosine similarity between them (Explosion 2025). Pairs scoring ≥ 0.65 were accepted as matches. While Sentence-BERT embeddings excel at capturing deep contextual semantic relationships, transformer-based models can occasionally overlook meaningful lexical overlaps, particularly when embeddings fall just below an acceptance threshold (e.g., BERT < 0.50). Given that many topic labels in both the Twitter API and TweetNLP schemes are relatively short, noun-heavy phrases (e.g., "music festivals," "designer fashion"), spaCy's lexical vector comparisons served as a crucial complementary check. This allowed us to recover relevant matches that BERT may have missed. All accepted matches from these automated stages (BERT or spaCy) were subsequently reviewed and verified through manual annotation to ensure initial consistency.
3. *Expert Human Annotation:* Remaining unmatched pairs were evaluated manually by three expert annotators (the study authors). This step addressed ambiguous or novel topics (e.g., "NFTs," "supernatural," "MSFT") not reliably handled by either automated method. Topic pairs were accepted if at least two of the three annotators independently agreed on the match, ensuring inclusion through majority consensus.

Alignment Quantification Once semantic mappings were established, we calculated two alignment metrics for each Space:

- Jaccard Similarity: Measuring the proportion of shared elements between the mapped creator-assigned and conversation-derived topic sets.
- Overlap Coefficient: Capturing the extent to which the smaller of the two sets is entirely covered by the other, providing a sensitivity measure for partial matches.

Together, these metrics provide a robust quantification of how closely the actual discussions aligned with the advertised topics.

Identifying Emergent Themes To assess the novelty of conversational content, we flagged any conversation-derived topic not covered by the mapped creator-assigned categories as an emergent theme. A Space was classified as containing emergent themes if it included at least one such unmatched label. The proportion of such Spaces was then computed to determine how frequently conversations diverged from their originally declared topical scope.

5.2 Results

The following subsections detail our findings regarding topic similarity within X Spaces (based on both creator-assigned and conversation-derived topics), topic similarity in user text posts (analysed using X-provided domains and entities, and TweetNLP), and the correlation of user interests expressed across these two engagement modes.

5.2.1 Topic similarity in audio spaces

We found a high overlap in user interests within X Spaces when analysing topic similarity based on creator-assigned topics. This strong interest alignment is evident in Figure 7 (a) and detailed in Table 4, which shows a mean Space topic similarity score of 0.90 (Mdn = 0.93, σ= 0.11).

To gain a more comprehensive understanding, we also analysed topics derived directly from Space conversations. Figure 7 (b) and Table 4 further compares the similarity scores obtained using both the creator-assigned and conversation-derived methods. As shown in the table, we observed a high mean similarity score (0.85) for conversation-derived topics, which was only slightly lower compared to the creator-assigned topics (0.90).

Table 4 Statistics for Space Topic Similarity

Statistic	μ	Mdn	σ
Mean space similarity (by assigned topics)	0.90	0.93	0.11
Mean space similarity (by actual content)	0.85	0.90	0.10

This small 0.05 difference in mean similarity scores indicates that users who exhibit high thematic alignment based on creator-assigned topics also tend to show high alignment when their interests are derived directly from conversation content. While creator-assigned topics might reflect users' idealised or initial interests for joining a Space, the consistently high similarity for conversation-derived topics implies that users attending similar Spaces do indeed share strong thematic commonalities in the actual discussions they engage in. This suggests a robust internal thematic focus within these conversation-derived interest profiles.

The statistics further highlight this consistency. The high mean (0.90) and median (0.93) similarity scores for creator-assigned topics, coupled with a small standard deviation (0.11), indicate a strong thematic alignment within users' networks based on these initial labels. Similarly, the high mean (0.85) and small standard deviation (0.10) for conversation-derived topics reflect the consistent thematic focus and lower variability that emerges in real-time discussions.

5.2.2 Topic similarity in text posts

Our analysis of topic similarity in users' public text posts revealed variations depending on the method of categorisation. When considering assigned domains, which represent broader interest categories provided by X, users exhibited a moderate level of similarity within their networks ($\mu = 0.43$, Mdn $= 0.43$, $\sigma = 0.09$), as illustrated in Figure 9 (a) and detailed in Table 5.

In contrast, similarity was significantly lower when examining assigned entities, which represent more specific interests provided by X ($\mu = 0.16$, Mdn $= 0.15$, $\sigma = 0.06$). The distribution of these entity-based similarities is shown in Figure 9 (b) and further detailed in Table 5. This suggests that while users may share broader interest areas with their network neighbours, their specific interests, as reflected in named entities, have less overlap.

Interestingly, analysis using the TweetNLP model for topic extraction revealed a moderate level of similarity ($\mu = 0.58$, Mdn $= 0.58$, $\sigma = 0.09$), as shown in Table 5 under the label "actual post" and depicted in Figure 9 (c). This level of similarity is closely aligned with that observed at the domain level. This suggests that TweetNLP's topic classification, while providing a different framework, captures a level of shared interest in text posts that is more akin to broader thematic domains than to specific entities.

Table 5 Statistics for Text Post Similarity

Statistic	μ	Mdn	σ
Mean text post similarity (assigned domains)	0.43	0.43	0.09
Mean text post similarity (assigned entities)	0.16	0.15	0.06
Mean text post similarity (actual post)	0.58	0.58	0.09

5.2.3 Cross-platforms interest correlation

The relationship between creator-assigned Space topic similarity and user text post domain similarity is characterised by a Spearman's ρ of 0.18 ($p < 0.00001$). This indicates a weak positive correlation, suggesting only a slight tendency for users interested in similar topics to share related text post interests.

Similarly, the Spearman's ρ between the ranked average Jaccard similarities of conversation-derived Space topics and user text post topics (TweetNLP-extracted) was found to be 0.26 ($p < 0.00001$). This also indicates a weak positive correlation, suggesting a modest tendency for users engaging in similar real-time Space discussions to also share related interests in their text posts. This weak association is visually represented in Figure 8, a scatter plot comparing the mean topic similarity in ego networks for conversation-derived Space topics and TweetNLP-extracted text post topics.

Thus, both analyses revealed statistically significant but weak positive correlations, highlighting a limited overlap between user interests expressed in Spaces (both creator-assigned and conversation-derived) and their text posts. This suggests that while some alignment exists, users tend

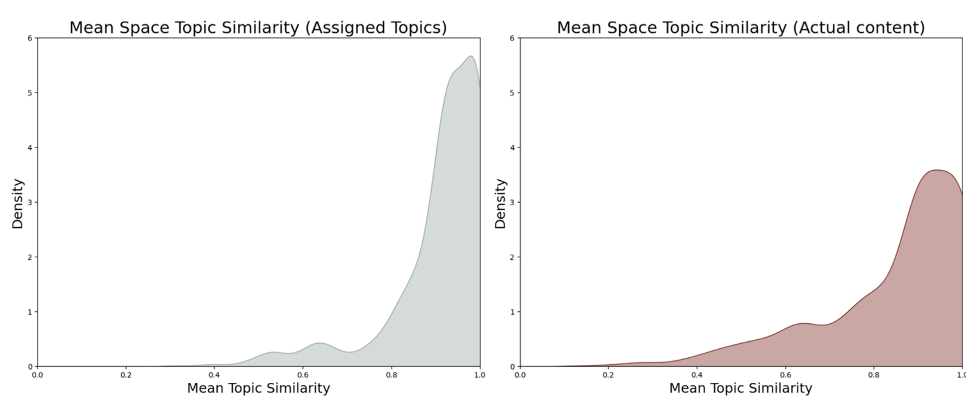

Fig. 7 Density plots of mean topic similarity within ego networks for Spaces. (a) Mean topic similarity using creator-assigned topics. (b) Mean topic similarity using conversation-derived "actual content"

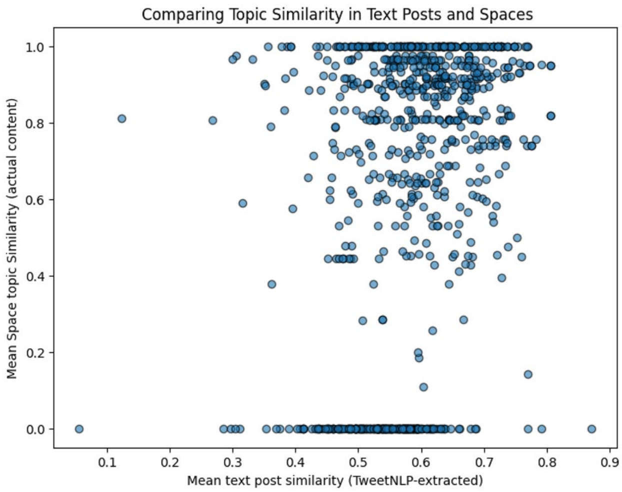

Fig. 8 Scatter plot illustrating the relationship between the mean Jaccard similarity of TweetNLP-extracted topics within a user's ego network (for text posts) and the mean Jaccard similarity of topics (actual content) within the same ego network (X Spaces)

to exhibit a broader spectrum of interests in their text post content compared to the more focused interactions within Spaces.

5.2.4 Alignment of intended and actual topics

We applied our hybrid alignment framework to 86 X Spaces, successfully mapping 42 unique creator-assigned topics to 15 conversation-derived categories generated by the TweetNLP classifier.

Of the 42 creator-assigned topics, 61.9% (n = 26) were automatically aligned at the BERT-based semantic similarity stage, reflecting the effectiveness of contextual embeddings in identifying conceptual equivalences. An additional 7.1% (n = 3) were resolved using the spaCy lexical model, which successfully recovered alignments missed by the BERT threshold. The remaining 31.0% (n = 13) required manual annotation, underscoring the necessity of human judgement for interpreting nuanced or emerging topics-such as NFTs, MSFT, paranormal, and supernatural-that fell below automated similarity thresholds but were contextually relevant.

Figure 10 visualises the multi-stage alignment process. The left-to-right flows represent topic mappings from creator-assigned to conversation-derived categories, with wider flows at the BERT stage indicating its dominant contribution. The comparatively narrower paths through spaCy and manual review highlight their more limited, but essential, roles. Flow widths are scaled by similarity scores, conveying each method's relative impact on the final mappings.

For a summarised view, Figure 11 presents a two-column Sankey diagram that links each creator-assigned topic directly to its final mapped category, omitting the alignment stages for visual clarity.

Across all 86 analysed Spaces, the mean Jaccard similarity between creator-assigned and conversation-derived topic sets was 0.35 ($\sigma = 0.31$), indicating only partial thematic overlap. In contrast, the mean overlap coefficient was 0.71 ($\sigma = 0.39$), suggesting that while creators' labels typically covered a substantial portion of the discussed content, additional or more specific themes often emerged during live interactions. Notably, 44.19% of Spaces (38 out of 86) introduced novel themes not present in the mapped creator assignments. This finding underscores the dynamic and exploratory nature of live audio discussions, which frequently evolve beyond their initially advertised topical scope.

5.2.5 Summary of findings

The following key findings summarise the main observations from our topic similarity analysis.

- *Space Topic Similarity.* The analysis revealed high mean similarities for both creator-assigned and conversation-derived topics within X Spaces (Table 4, Figure 7). Creator-assigned topics exhibited a mean similarity of 0.90, and conversation-derived topics also showed a high mean similarity of 0.85 (with a standard deviation of 0.10). This indicates a strong level of thematic consistency among users engaging in similar Spaces. While creator-assigned topics demonstrate slightly higher uniformity, the high similarity observed in conversation-derived topics suggests that actual discussions within these Spaces generally remain well-aligned and focused.

Fig. 9 Density plots of mean topic similarity within ego networks for text posts. (a) Mean text post similarity using assigned domain labels. (b) Mean text post similarity using assigned entity labels. (c) Mean text post similarity using TweetNLP-derived topics

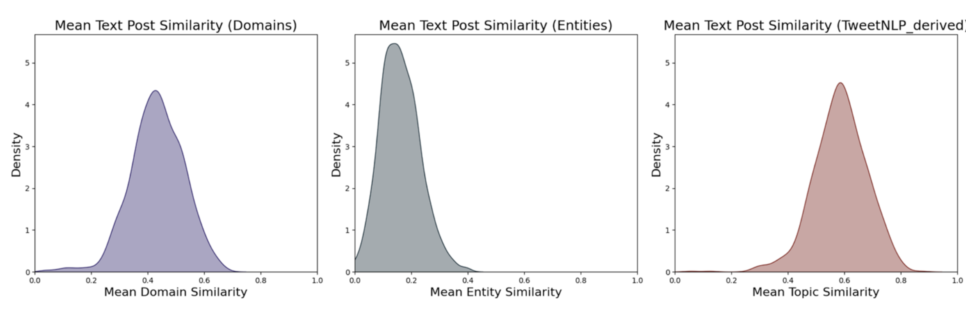

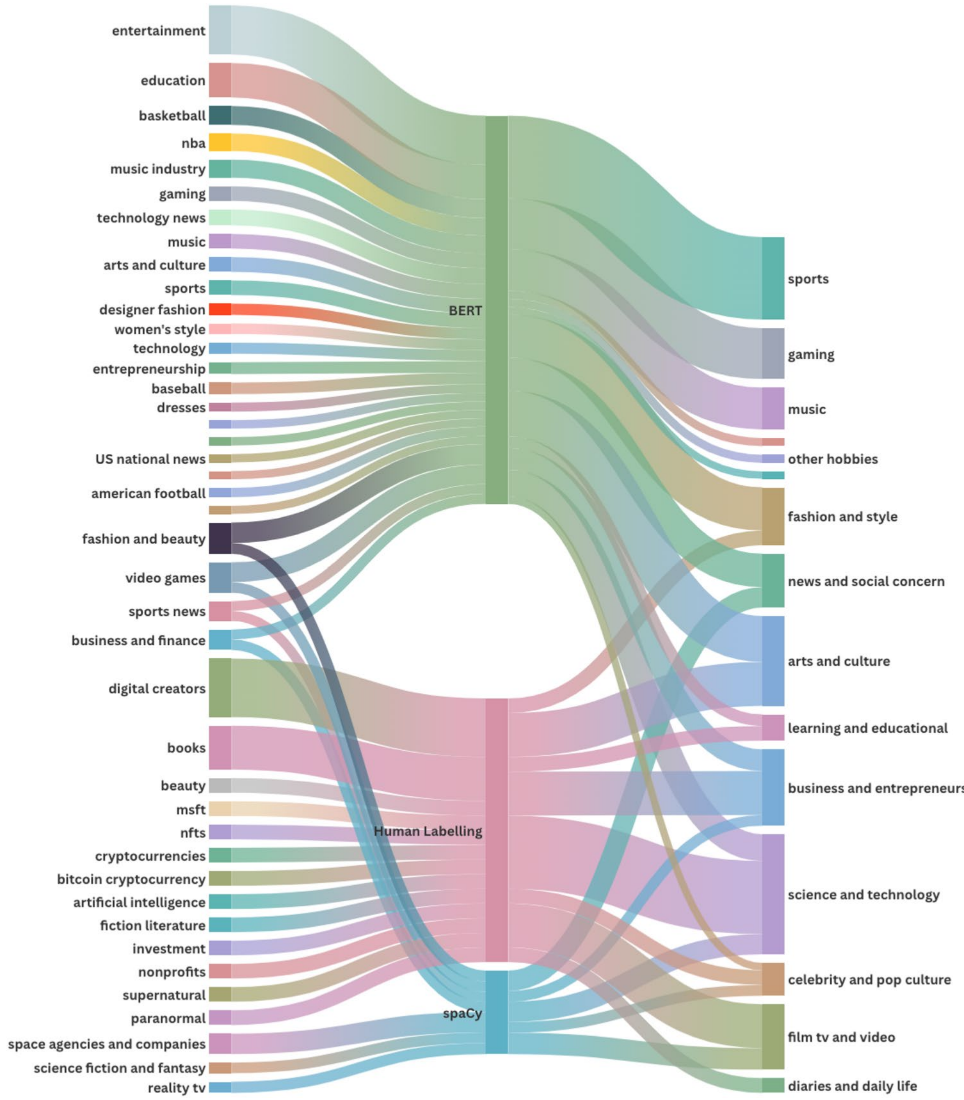

Fig. 10 Flow diagram showing how initial Twitter API topics (left) align with conversation-derived topics (right) through a three step process: BERT, spaCy and manual review, with stream widths proportional to their similarity scores. For any topic pairs matched via expert human annotation, we assign a default similarity score of 1, reflecting majority agreement among the annotators

- *Variation in text post topic similarity across categorisation methods.* The similarity of topics in user text posts differed depending on the approach used for categorisation. Domain-based analysis, reflecting broader interests, yielded a moderate similarity of ($\mu = 0.43$). Entity-based analysis, focusing on specific interests, resulted in a low similarity of ($\mu = 0.16$). TweetNLP-extracted topics (actual post), providing an alternative perspective, showed a moderate similarity of ($\mu = 0.58$), interestingly close to the moderate similarity observed at the domain level. (Table 5, Figure 9).
- *Weak correlation between Space topics and text post topics suggests platform-specific interest expression.* The weak positive correlations between Space topics (both creator-assigned, $\rho = 0.18$, and conversation-derived, $\rho = 0.26$) and user text post content (Sect. 5.2.3) indicate that users tend to express their interests differently on X Spaces compared to their text posts, suggesting distinct communicative roles for these features.
- *Partial coverage and emergent themes.* While the mean overlap coefficient of 0.71 suggests that creator-assigned topics generally encompass the majority of emergent themes, a significant proportion of Spaces (approximately 44%) introduced novel topics not initially anticipated. This highlights the dynamic nature of live conversations, where discussions frequently extend beyond their pre-defined scope.

6 Discussion

The preceding sections have presented distinct yet interconnected analyses of user behaviour on X Spaces. Sect. 4 examined the structural characteristics of user interactions through network analysis, while Sect. 5 explored thematic

Fig. 11 Sankey diagram mapping each creator-assigned Twitter API topic (left) directly to its corresponding conversation-derived TweetNLP category (right)

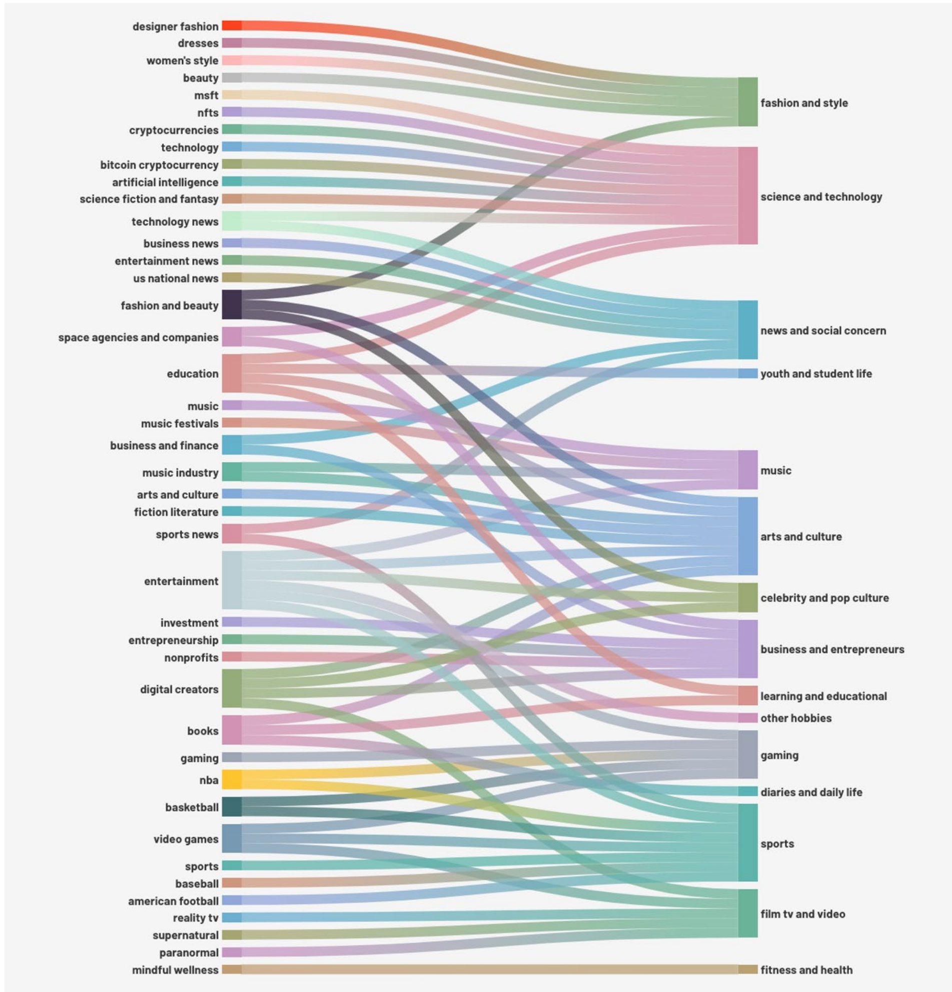

similarities in user interests across audio and textual modalities. This section integrates these structural and thematic insights to provide a cohesive interpretation of user dynamics within this evolving social audio environment.

6.1 Network structure and connectivity

The structural analyses offer key insights into how users engage with X Spaces and form social ties within the platform. By examining the bipartite user-Space network, the projected user-user network, ego-centric subgraphs, and degree distributions, we identified both broad patterns and localised variations in user connectivity.

The bipartite user-Space network exhibited a low overall density (0.003) and significant fragmentation into 101 components, reflecting sparse global connectivity-an expected feature of large-scale digital networks (Golbeck 2013). Despite this sparseness, the network displayed a moderate overall clustering coefficient (0.518), with clustering notably higher among users (0.560) than Spaces (0.345). This suggests that users tend to participate in similar sets of Spaces, forming tightly connected clusters of shared interest, even as the global structure remains fragmented. These findings align with prior studies that highlight the formation of micro-communities within fragmented social platforms (Golbeck 2013).

The degree distribution further underscores this local clustering dynamic. While the majority of users participate in a small number of Spaces, a minority exhibit disproportionately high levels of engagement. This long-tailed pattern is characteristic of selective participation and echoes findings from other social platforms, where a small subset of users accounts for much of the interaction volume. These behavioural asymmetries may be driven by platform affordances, topical relevance, or self-selection into interest-driven communities.

The projection of the bipartite structure into a user-user network revealed further evidence of local cohesion despite global sparseness. Although the user-user network wasn't fully connected, its giant component exhibited a relatively short mean shortest path length and a high clustering coefficient, indicative of a small-world structure. However, it's important to note that our co-attendance-based network construction assumed mutual connections among all users within a Space. Therefore, the observed high clustering coefficient and the small-world properties are best understood as reflecting the opportunity for interaction and the formation of communities based on shared situational or thematic contexts among speakers, rather than necessarily representing established, direct social ties. This configuration still enables efficient information diffusion and supports both tightly knit interaction clusters and indirect reachability across the network.

The analysis of ego-centric networks provided a more granular view of user positioning. By constructing two-hop ego networks based on Space co-attendance, we were able to assess both direct and indirect social proximity for each user. These ego subgraphs revealed substantial variation in local network structure: while some users were embedded in dense, tightly clustered communities, others spanned broader, less cohesive neighbourhoods. This variation suggests differentiated user roles-from peripheral participants to structurally central connectors-highlighting distinct patterns of engagement and potential influence within the platform.

Centrality metrics reinforced this heterogeneity. Variations in degree and betweenness centrality suggest that some users occupy pivotal structural positions-either through direct interactions with many others or by bridging otherwise disconnected subgroups. The relatively high mean node strength (32.12) observed within ego networks further reflects active participation by certain users, indicating their importance in shaping interaction flows and sustaining engagement across the network.

6.2 User engagement and topic interactions

Our analysis demonstrates that users adapt their engagement strategies according to the modality offered by the platform, with distinct patterns emerging between audio-based and text-based interactions. Specifically, we found that topic similarity is generally higher within X Spaces than in users' public text posts. This divergence invites closer examination of how thematic alignment varies not only across modalities but also within the structure of audio Spaces themselves. This adaptation suggests that users strategically leverage each modality for distinct communicative purposes, leading to varied expressions of interest and engagement.

Conversation-derived topics, extracted using automated classification of transcribed audio content from X Spaces, provide a grounded account of real-time discussions. These themes often reflect spontaneous shifts, tangential explorations, or sub-community interests. Indeed, while most Spaces exhibit high topic coverage of their primary themes, our analysis found that 44% introduce additional themes, strongly suggesting that live audio conversations often diverge from their predefined scope. The dynamic nature of discourse during live audio events-responsive to the flow of interaction rather than a rigid initial framing-is likely encouraged by the ephemeral, synchronous, and socially intimate nature of audio Spaces, which have been shown to foster more authentic or focused interactions than asynchronous formats (Li and Penaranda Valdivia 2022; Bajpai et al. 2022). Unlike the more curated nature of text posts, the live, unscripted environment of audio conversations allows for rapid topic evolution driven by immediate participant engagement and a less constrained discourse.

A similar degree of modality-specific interest expression is evident in the analysis of user text posts. Here, the degree of topic similarity varied according to the categorisation framework: domain-level annotations (e.g., "Politics," "Technology") captured broader thematic interests and yielded moderate similarity scores, while entity-level labels (e.g., specific people, organisations, or products) resulted in lower overlap among users. Interestingly, topic similarity from TweetNLP-extracted categories fell closer to the domain-level metrics, suggesting that this model captures generalised themes that align with higher-level interest signals.

The low correlations observed between topic similarities in Spaces and in text posts (both domain- and TweetNLP-based) further underscore this cross-modality divergence. These weak associations suggest that users do not express a uniform set of interests across all features of the platform. Rather, user behaviour appears to be context-sensitive, shaped by the structural and temporal affordances of each modality. Text posts, with their asynchronous and persistent nature, are well-suited for general commentary and broader self-expression. In contrast, live Spaces support immediate, dialogic interactions, often focused on more specific or time-sensitive topics. This behavioural differentiation aligns with prior studies noting how platform design influences content production and thematic consistency across modalities (e.g., Li and Penaranda Valdivia 2022; Bajpai et al. 2022,). This highlights how individuals strategically adapt their communication styles and thematic focus to best suit the unique affordances and social expectations of each format.

Given Jaccard similarity's sensitivity to varying topic set sizes, we also computed the Overlap Coefficient for a complementary perspective. For X Spaces, the Mean Overlap Coefficient ($\mu = 0.99$, Mdn $= 0.86$, $\sigma = 0.07$) was notably higher than Jaccard ($\mu = 0.85$, Mdn$= 0.90$, $\sigma = 0.10$) (Table 4). Similarly, for text posts, Overlap ($\mu = 0.83$, Mdn $= 0.84$, $\sigma = 0.07$) consistently exceeded Jaccard ($\mu = 0.58$, Mdn$= 0.58$, $\sigma = 0.09$) (Table 5). This pattern suggests that Overlap, by focusing on shared elements relative to the smaller set, offers a less penalising and more robust measure of thematic similarity when set sizes vary. While lower Jaccard values for text posts point to broader individual interests, the higher Overlap Coefficient reveals a consistent core of shared themes. Together, these results show that Spaces foster tighter topical convergence, whereas text posts support common interests within a more diverse topical range.

These findings carry several implications for platform design, recommendation systems, and user modelling. Understanding user interests in X Spaces-or any multi-modal platform-requires an appreciation of how users distribute their attention and content across different engagement channels. Personalisation algorithms that rely exclusively on Space participation may fail to capture the full breadth of a user's interests as expressed in other content types, such as text posts, and vice versa. Our results suggest that multimodal modelling of user interests is essential to developing more accurate, context-aware systems that reflect the complexity of user behaviour. Recognizing this heterogeneity in user engagement across modalities is crucial for building more comprehensive user profiles that reflect their varied online personas and interests.

In summary, our investigation into both the network structure and thematic alignment within X Spaces presents a nuanced picture of user interaction. Network analysis revealed localised community structures and uneven engagement patterns, suggesting that users participate in concentrated clusters shaped by shared interest. Topic similarity analysis supported this, showing that users within a given Space often engage in thematically coherent discussions-but only to the extent permitted by the modality and structure of that feature. Meanwhile, the divergence in interest expression between synchronous audio and asynchronous text interactions highlights the adaptive strategies users employ when navigating platform affordances. These insights underscore the importance of designing systems that account for the complex interplay between structure, modality, and user intent in social media environments.

6.3 Limitations and future work

This study has offered valuable insights into user behaviour on X Spaces by examining both social network structures and thematic engagement across modalities. However, several methodological limitations should be acknowledged, each of which opens avenues for future research.

Network Construction Assumptions. Our social network analysis was based on user co-attendance in Spaces to construct both user-user and ego-centric networks. While this approach effectively captures shared topical contexts and potential community structures, it inherently assumes that all users participating in the same Space are mutually connected. This assumption does not account for the strength or directionality of interactions, potentially inflating measures such as the clustering coefficient and overestimating the density of triadic closures. As a result, the derived networks may reflect opportunity for connection rather than actual engagement or influence.

However, this methodological choice was informed by the unique nature of our dataset, which exclusively comprised users participating as speakers in Spaces. Given that these platforms are explicitly designed for verbal engagement and the speaking floor is limited (typically a maximum of 13 speakers at once), there is a reasonable, albeit indirect, assumption that speakers within the same Space share a common focus or engage in some form of interaction. Future work should consider alternative network construction techniques that incorporate interaction frequency, verbal exchanges, temporal co-presence, or content alignment, which would allow for a more fine-grained and realistic representation of social ties within audio-based platforms.

Constraints in Topic Representation. Our analysis of conversation-derived topics relied on summarised transcriptions due to the 512-token input limit of the RoBERTa-based TweetNLP topic classifier. Although these summaries were designed to capture core themes, they may omit contextual nuances and temporal developments that occur over the full course of a Space. This approach, while effective for identifying major themes, risks a bias towards more generalised topics, potentially underrepresenting thematic shifts or less prominent narratives within the conversation. Future work could address this limitation by segmenting full transcripts-e.g., by speaker turns or conversational phases-and applying topic models to these smaller units. This would support temporal or multi-level topic modelling, allowing researchers to track how discussions evolve over time and how different participants contribute to the thematic flow.

Opportunities for Enhancing Topic Alignment and Dynamics. Several additional directions arise from the current topic analysis framework. First, a quantitative mapping between NLP-derived conversation topics and the pre-assigned creator topics-beyond the hybrid alignment framework used here-could facilitate automated evaluation of alignment accuracy and semantic drift. Second, investigating topic transitions within Spaces may uncover recurring patterns in how discussions shift between themes or respond to participant inputs. This could lead to the development of real-time topic tagging systems that reflect actual conversational content more faithfully than static pre-defined labels.

Impact on User Engagement. Finally, future research could examine how the alignment or misalignment between pre-assigned and emergent topics affects user behaviour, such as participation duration, speaker activity, or audience retention. Understanding this relationship could inform the design of more responsive and context-aware recommendation or moderation systems, ultimately enhancing the user experience on social audio platforms.

Overall, while the methodological choices in this study were appropriate for the available data and research goals, we recognise that alternative designs and more granular data representations could yield deeper insights into user dynamics. Future work that builds on these foundations-particularly by integrating richer interaction signals and more flexible topic modelling techniques-will be well-positioned to advance understanding of engagement and discourse in real-time social platforms.

7 Conclusion

This study has provided a comprehensive investigation into user behaviour within the evolving social audio environment of X Spaces, integrating structural network analysis with thematic modelling of user interests across synchronous (audio) and asynchronous (text) modalities. Leveraging an expanded dataset of 3435 users across 619 Spaces, we demonstrated that the X Spaces social network exhibits sparse global connectivity, yet maintains key broad-scale structural properties typical of social systems-namely, heterogeneous degree distributions, high clustering, and short average path lengths. These features persist at scale and confirm a platform where localised, tightly-knit user communities emerge within individual Spaces, supported by a small subset of highly active users embedded within a more diffuse, peripheral user base.

Ego-centric network analysis further uncovered substantial variation in how users are embedded within their local network environments. Some users are situated within dense, cohesive clusters-reflecting focused community participation-while others span more dispersed structures, indicative of broader topical engagement and potentially more influential or bridging roles. These distinct local configurations reflect the diverse modes of user participation and suggest multiple pathways through which users contribute to, and benefit from, the platform's social dynamics.

Beyond structural properties, our thematic analysis highlighted a key divergence in user interests expressed through audio-based versus text-based engagement. Topic similarity was consistently higher within users' Space participation networks than in their public text posts, and correlations across these modalities were weak. This suggests that users tailor their communication to the affordances of each modality. Particularly notable was the substantial discrepancy between creator-assigned Space topics and the themes that actually emerged in live conversations. This finding underscores the limitations of static topic labels and the value of analysing real-time content to better understand user intent and discourse.

Together, these findings point to the importance of multi-modal and multi-level approaches in studying social media platforms. From a design perspective, recognising the divergence between intended and actual user interests-and between networked and discursive forms of engagement-can inform the development of personalisation algorithms, recommendation systems, and community support tools that are better aligned with user behaviours and needs.

Future work can build upon these contributions in several directions. First, incorporating richer interaction data (e.g., verbal exchanges, frequency, or time-stamped engagement) would yield more accurate models of social connectivity. Second, adopting temporal analysis of topic shifts within Spaces could capture conversational dynamics beyond summary-level content. Third, leveraging full transcripts and employing segmented or streaming NLP techniques would allow deeper examination of how themes unfold. Finally, comparative analysis across multiple audio-based platforms would provide a broader understanding of how platform affordances shape user experience and social formation.

In summary, this study contributes empirical insights into the interplay between network structure and thematic content in shaping user experience on synchronous audio platforms. It highlights the value of integrated approaches for understanding emergent online behaviours and offers a methodological framework for future investigations into the evolving landscape of real-time social media.

Author Contributions Conceptualization: R.D. , A.A. and L.T. ; Methodology: R.D. , A.A. and L.T. ; Formal analysis and investigation: R.D. ; Writing - original draft preparation: R.D.; Writing - review and editing: R.D. , A.A. and L.T.; Resources: R.D. ; Supervision: A.A. and L.T.

Funding No funding was received for conducting this study.

Data Availability The data analysed in this study are publicly available through the X API. Due to the potential for individual user identification from specific collections of text posts and Space content, the raw datasets analysed in this study will not be directly shared. However, researchers can access and collect similar publicly available data by adhering to X's developer terms and conditions and by following the detailed data collection procedures outlined in the Methods section of this paper. The code supporting the analyses in this study is available at https://github.com/RobaDarwish/xspaces-analysis-code.

Declarations

Conflict of interest The authors declare no Conflict of interest.

Ethics approval and Consent to participate This study utilised publicly available data collected through the X API, including both text posts and audio replays of Spaces. The research protocol was reviewed and approved by the School of Computer Science and Informatics Research Ethics Committee at Cardiff University (COMSC/Ethics/2023/114). As the data analysed was publicly available and did not involve direct interaction with human subjects or the collection of private identifiable information beyond what was publicly disclosed, individual informed consent was not directly sought from participants.

References

Albert R, Barabási A-L (2002) Statistical mechanics of complex networks. Rev Mod Phys 74(1):47. https://doi.org/10.1103/RevModPhys.74.47

Abdallah S (2011) Generalizing unweighted network measures to capture the focus in interactions. Soc Netw Anal Min 1(4):255–269. https://doi.org/10.1007/s13278-011-0018-8

Alstott J, Bullmore E, Plenz D (2014) powerlaw: a python package for analysis of heavy-tailed distributions. PLoS ONE 9(1):85777. https://doi.org/10.1371/journal.pone.0095816

Anandha Jothi M, Sankar K (2023) On the metric dimension of bipartite graphs. AKCE Int J Graphs Comb 20(3):287–290. https://doi.org/10.1080/09728600.2023.2223248

Antypas D, Ushio A, Camacho-Collados J, Neves L, Silva V, Barbieri F (2022) Twitter topic classification. In: Proceedings - International Conference on Computational Linguistics COLING 29:3386–3400. https://doi.org/10.48550/arXiv.2209.09824

Bajpai T, Asher D, Goswami A, Chandrasekharan E (2022) Harmonizing the cacophony with mic: An affordance-aware framework for platform moderation. In: Proceedings of the ACM on Human-Computer Interaction 6(CSCW2):1–22. https://doi.org/10.1145/3555110

Barabási A-L (2016) Graph theory. In: Network science, Chap 2. Retrieved from https://networksciencebook.com/

Bayer JB, Ellison NB, Schoenebeck SY, Falk EB (2016) Sharing the small moments: ephemeral social interaction on snapchat. Inf Commun Soc 19(7):956–977. https://doi.org/10.1080/1369118X.2015.1084349

Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang D-U (2006) Complex networks: structure and dynamics. Phys Rep 424(4–5):175–308. https://doi.org/10.1016/j.physrep.2005.10.009

Breznitz D (2011) Delete: the virtue of forgetting in the digital age. Rev Pol Res 28(3):307–309. https://doi.org/10.1111/j.1541-1338.2011.00497_2.x

Camacho-Collados J, Rezaee K, Riahi T, Ushio A, Loureiro D, Antypas D, Boisson J, Espinosa-Anke L, Liu F, Martínez-Cámara E, *et al.* (2022) TweetNLP: Cutting-Edge Natural Language Processing for Social Media. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. Association for Computational Linguistics, Abu Dhabi, U.A.E. .https://doi.org/10.48550/arXiv.2206.14774

Clark M (1993) Reactions to and strategic self-presentation of happiness, sadness, and anger. In: an Invited Symposium at the Meeting of the American Psychological Society, Chicago, IL

Clauset A, Shalizi CR, Newman ME (2009) Power-law distributions in empirical data. SIAM Rev 51(4):661–703. https://doi.org/10.1137/070710111

Darwish R, Abdelmoty A, Turner L.D (2025) Investigating social network structures and interests on x audio spaces. In: Cherifi, H., Donduran, M., Rocha, L.M., Cherifi, C., Varol, O. (eds.) Complex Networks & Their Applications XIII, pp. 397–408. Springer, Cham . https://doi.org/10.1007/978-3-031-82431-9_33

Silva LOJ, Maldonado G, Brigham T, Mullan AF, Utengen A, Cabrera D et al (2021) Evaluating scholars' impact and influence: Cross-sectional study of the correlation between a novel social media-based score and an author-level citation metric. J Med Internet Res 23(5):28859. https://doi.org/10.2196/28859

Explosion AI: Linguistic Features: Word vectors and semantic similarity. https://spacy.io/usage/linguistic-features#vectors-similarity. spaCy Documentation (Accessed: June, 2025)

Golbeck J (2013) Analyzing the Social Web, 1st edn. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA . https://doi.org/10.1016/C2012-0-00171-8

Gaber M, Pastor G.C, Omer A (2020) Speech-to-text technology as a documentation tool for interpreters A new approach to compiling an ad hoc corpus and extracting terminology from video-recorded speeches. TRANS revista de traductología (24), 263–281 https://doi.org/10.24310/TRANS.2020.v0i24.7876

Hansen D, Shneiderman B, Smith M, Himelboim I (2011) Chapter 3-social network analysis: measuring, mapping, and modeling collections of connections In: Analyzing social media networks with node XL. Boston: Morgan Kaufmann . https://doi.org/10.1016/B978-0-12-817756-3.00003-0

Jung K, Park Y, Kim H, Lee J (2022) Let's talk@ clubhouse: Exploring voice-centered social media platform and its opportunities, challenges, and design guidelines. In: CHI Conference on Human Factors in Computing Systems Extended Abstracts, pp. 1–6 . https://doi.org/10.1145/3491101.3519774

Knott R (2024) Synchronous vs. asynchronous communication: how to use both to dominate remote work. TechSmith. https://www.techsmith.com/blog/synchronous-vs-asynchronous-communication/

Kung-Ming T, Khoon-Seng S (2005) Asynchronous vs. synchronous interaction. In: Encyclopedia of distance learning, pp 104–113. IGI Global, Hershey, PA, USA. https://doi.org/10.4018/978-1-59140-555-9.ch016

Lind PG, Gonzalez MC, Herrmann HJ (2005) Cycles and clustering in bipartite networks. Phys Rev E-Stat Nonlinear Soft Matter Phys 72(5):056127. https://doi.org/10.1103/PhysRevE.72.056127

Li J, Penaranda Valdivia K (2022) Does media format matter? investigating the toxicity, sentiment and topic of audio versus text social media messages. In: Proceedings of the 10th International Conference on Human-Agent Interaction. HAI '22, pp. 13–22. Association for Computing Machinery, New York, NY, USA . https://doi.org/10.1145/3527188.3561927

Malik HAM (2022) Complex network formation and analysis of online social media systems. Cmes-Comr Model Engg Sci 130(3):1737–1750. https://doi.org/10.32604/cmes.2022.018015

Mei Q, Ling X, Wondra M, Su H, Zhai C (2007) Topic sentiment mixture: modeling facets and opinions in weblogs. In: Proceedings of the 16th International Conference on World Wide Web, pp. 171–180 . https://doi.org/10.1145/1242572.1242596

Michelson M, Macskassy S.A (2010) Discovering users' topics of interest on twitter: a first look. In: Proceedings of the Fourth Workshop on Analytics for Noisy Unstructured Text Data, pp. 73–80. https://doi.org/10.1145/1871840.1871852

Otter.AI: Privacy and Security. Accessed: February, 2025. https://otter.ai/privacy-security

Ramage D, Dumais S, Liebling D (2010) Characterizing microblogs with topic models. In: Proceedings of the International AAAI Conference on Web and Social Media, vol. 4, pp. 130–137 . https://doi.org/10.1609/icwsm.v4i1.14026

Reimers N, Gurevych I (2019) Sentence-bert: Sentence embeddings using siamese bert-networks. https://doi.org/10.48550/arXiv.1908.10084

Schlesinger A, Chandrasekharan E, Masden C.A, Bruckman A.S, Edwards W.K, Grinter R.E (2017) Situated anonymity: Impacts of anonymity, ephemerality, and hyper-locality on social media. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, pp. 6912–6924 . https://doi.org/10.1145/3025453.3025682

Shen W, Wang J, Luo P, Wang M (2013) Linking named entities in tweets with knowledge base via user interest modeling. In: Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 68–76 . https://doi.org/10.1145/2487575.2487686

Turner LD, Whitaker RM, Allen SM, Linden DE, Tu K, Li J, Towsley D (2019) Evidence to support common application switching behaviour on smartphones. R Soc Open Sci 6(3):190018. https://doi.org/10.1098/rsos.190018

Villazon-Terrazas J, Aparicio S, Alvarez G (2015) Study on twitter as a complex network. In: The Third International Conference on Digital Enterprise and Information Systems (DEIS2015), vol. 54, p. 19

Williams D, Caplan S, Xiong L (2007) Can you hear me now? the impact of voice in an online gaming community. Hum Commun Res 33(4):427–449. https://doi.org/10.1111/j.1468-2958.2007.00306.x

Watts DJ, Watts SH (1998) Collective dynamics of 'small-world' networks. nature 393:440–442. https://doi.org/10.1038/30918

X: Tweet Annotation. Accessed: July. 8, 2024. https://developer.x.com/en/docs/x-api/annotations/overview