

Elephant in the Room: Dissecting and Reflecting on the Evolution of Online Social Network Research

Luca Pajola,^{1*} Saskia Laura Schröer,² Pier Paolo Tricomi,^{1*} Mauro Conti,¹ Giovanni Apruzzese²

¹ University of Padua, Padua, Italy

² University of Liechtenstein, Vaduz, Liechtenstein

{luca.pajola, mauro.conti}@unipd.it, {saskia.schroeer, giovanni.apruzzese}@uni.li, pierpaolo.tricomi@spritzmatter.com

Abstract

Billions of individuals engage with Online Social Networks (OSN) daily. The owners of OSN try to meet the demands of their end-users while complying with business necessities. Such necessities may, however, lead to the adoption of restrictive data access policies that hinder research activities from “external” scientists—who may, in turn, resort to other means (e.g., rely on static datasets) for their studies. Given the abundance of literature on OSN, we—as academics—should take a step back and reflect on what we have done so far, after having written thousands of papers on OSN.

This is the first paper that provides a holistic outlook to the entire body of research that focused on OSN—since the seminal work by Acquisti and Gross (2006). First, we search through over 1 million peer-reviewed publications, and derive 13,842 papers that focus on OSN: we organize the metadata of these works in the Minerva-OSN dataset, the first of its kind—which we publicly release. Next, by analyzing Minerva-OSN, we provide factual evidence elucidating trends and aspects that deserve to be brought to light—such as the predominant focus on Twitter or the difficulty in obtaining OSN data. Finally, as a constructive step to guide future research, we carry out an expert survey (n=50) with established scientists in this field, and coalesce suggestions to improve the status quo—such as an increased involvement of OSN owners. Our findings should inspire a reflection to “rescue” research on OSN. Doing so would improve the overall OSN ecosystem, benefiting both their owners and end-users—and, hence, our society.

Introduction

Online Social Networks (OSN) are platforms where **end-users** create virtual connections, share their experiences, communicate, and build social relations. Since the rollout of major OSN in the late 2000s – such as Facebook and Twitter – OSN have established a non-negligible space in the routines of billions of human beings (Kim and Hastak 2018). Indeed, we now use OSN *also* for purposes revolving around aspects of our professional life. Examples include following and publishing news, finding work opportunities, purchasing or selling items, or even for full-time and OSN-specific jobs like “social media manager” and “social influencer” (Jin, Muqaddam, and Ryu 2019).

*These authors are affiliated with Spritz Matter srl.
Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The growth of OSN was simply impossible to overlook, and the **research community** did not stand idly by. Starting from the seminal work by Acquisti and Gross (2006) on Facebook, over the years thousands of scientific papers have been published, focusing on a plethora of themes revolving around various OSN and having plenty of societal implications. For instance, Chu et al. (2010) started to design algorithms to detect “bots” in OSN, and years later, Bessi and Ferrara (2016) found that such bots played a crucial role in the outcome of the US elections. Since then, abundant papers have focused on detecting fake news (Shu et al. 2017).

Simply put, OSN are a (perpetually increasing) data source with infinite potential—and therefore are a valuable business for OSN **owners**. To sustain any OSN there is the effort of thousands of employees (at various levels), who work to provide a *fulfilling experience to end-users* while *sustainably generating revenue* for their own companies, and *ensuring compliance* with modern norms (e.g., regulations and moral customs). Although some OSN owners may appreciate collaborating with researchers, some decisions may also put “barriers” that are difficult to overcome by academics—such as allowing access to OSN data only through thousands-of-USD-worth of fees (Calma 2023).

The interplay among these three stakeholders (OSN owners, researchers, and end-users) delineates a complex ecosystem. Ideally, the aggregated contributions of every player should improve the ecosystem as a whole—which would ultimately lead to an overall enhancement of our quality of life. **As researchers, we want to facilitate such an improvement.** To this purpose, we ask ourselves a broad question: “*what has happened since the first research paper on OSN?*” The OSN ecosystem has made enormous advances also thanks to researchers: in this context, are there some elements that have received a greater (potentially overshadowing others) degree of attention—and, if so, why?

In this paper, we perform a reflective exercise and seek to scrutinize **the entire body of OSN research**. Our goal is to examine whether prior literature, as a whole, conceals certain “trends” which, if not known, could be harmful to future research—and understand potential root causes. To this purpose, we first systematically collect the metadata of over 1 million publications, and identify 13,842 works related to OSN. We organize these papers in the Minerva-OSN dataset, the first of its kind for OSN-research, which we publicly




Figure 1: **OSN prevalence over time.** We highlight the discrepancy between OSN investigated in research and their real-world popularity (website ranking based on Alexa [25]).

release. Next, through a meta-analysis, we investigate our Minerva-OSN dataset across four dimensions: *OSN prevalence*, wherein we show that some OSN are significantly more studied than others (see Figure 1); *topic popularity*, wherein we reveal the topmost research topics considered in prior work; *author analysis*, wherein we elucidate some aspects of the authors of the papers in Minerva-OSN; *regulation and policies*, wherein we attempt to explain some of our findings in light of major changes, unrelated to research but which deeply affect OSN-related research. Finally, to further validate some of the “issues” we uncovered, and to improve the status quo, we reach out to some of the authors of the papers in Minerva-OSN: through an (anonymous) expert survey, we collect the opinions of 50 renowned OSN researchers, highlighting the (often hidden) difficulties of carrying out impacting research in the OSN domain.

Contributions. We advance the state of OSN research by:

- Carrying out the largest meta-analysis of OSN-related literature, highlighting trends and aspects of prior research, and tying them with the OSN ecosystem.
- Releasing the dataset used for our meta-analysis, Minerva-OSN, comprising of 13,842 papers (published between 2006 and 2023) included in 135 peer-reviewed venues and discussing 91 OSN.
- Conducting an expert survey with renowned OSN researchers (n=50), inquiring for their perspective on the present state of research and ways to enhance it.

All resources are publicly released in our repository.¹ (Note: in the Appendix we provide a list of all non-bibliographic sources and their URLs—e.g., [1] refers to our repository.)

POSITIVE LIGHT. This paper is a reflective exercise in which we constructively scrutinize years of endeavors on OSN. Throughout our paper, we elucidate issues underpinning systematic problems—but which have not a single culprit. Hence, we are not pointing the finger at anyone.

¹<https://github.com/pajola/Minerva-OSN>

Related Work

We are not aware of any paper whose contributions overlap with ours—despite thousands of works on OSN.

Most papers on OSN, such as Tang et al. (2022), focus on a given phenomenon on one OSN; others, such as Fiesler, Beard, and Keegan (2020), focus on a *specific problem* across many OSN. Neither of these types of works attempts to provide a holistic view of OSN research.

Papers that analyze prior literature also tend to focus on *one aspect*: for instance, the review by Andreassen (2015) considers works on OSN addiction, whereas the SoK by Singhal et al. (2023) focuses on content moderation. Similarly, various reviews consider a single OSN: this is the case of Antonakaki, Fragopoulou, and Ioannidis (2021) for Twitter; and of Wilson, Gosling, and Graham (2012) for Facebook. In contrast, our analyses are agnostic to the focus of any analyzed paper (as long as it pertains to OSN) and we do not restrict our vision to any specific OSN.

Furthermore, many papers collect the opinion of experts to guide future research—either through questionnaires, such as Dwivedi et al. (2023); or through interviews, such as Bieringer et al. (2022). However, we are unaware of any work that carried out an expert study in the OSN context.

Finally, to derive our conclusions, we analyze the metadata (e.g., titles and abstract) of prior work. Such an approach is widely adopted in some research domains – e.g., management and business (Donald 2022; Gaur and Kumar 2018) – but it has never been used before in the OSN domain. This is why our research also encompasses original methods. For instance, we combine automated “AI-driven” techniques with manual “human-driven” analyses. As a practical example of our technical innovation, we: (i) use a topic model to determine which topic best describes any given paper, (ii) augment the model’s accuracy via our domain expertise, and (iii) systematically validate the model’s quality via a tri-party reviewer system—we are not aware of works adopting a similar quality assurance procedure.

The Minerva-OSN Dataset

Among our major contributions, we introduce the Minerva-OSN dataset. (Minerva was the goddess of wisdom in Roman mythology.) We use the dataset as a basis to derive our findings. Here, we first elucidate the methodology employed in the creation of Minerva-OSN dataset—which is rooted in the guidelines of systematic literature reviews (Booth et al. 2021). Then, we explain how we carried out the analysis on topic modeling—representing a pivotal part of our research, as well as a substantial technical contribution in this domain.

Data Collection (creation of Minerva-OSN)

We design Minerva-OSN with the intention of collecting “all” *OSN-related* and *peer-reviewed* papers (in computer science) that have ever been published. Hence, the first step entails identifying research-focused venues (i.e., journals and conference proceedings) that contain such publications.

Venues and candidate papers. We begin our search with Google Scholar (among the most well-known repositories of scientific articles), and consider the subcategories within

the “Engineering and Science” category [31]. For each sub-category, Google Scholar lists the 20 most popular venues. We integrate these 20 top-tier venues with all those listed in the Springer’s Lecture Notes in Computer Science and the International Conference Proceedings Series of the Association of Computing Machinery. We manually review all identified venues, remove duplicates and merge venues that changed publisher or name (e.g., WWW changed to TheWebConf). In total, we identify 135 relevant venues (reported in the Appendix). Then, for each venue, we query the Scopus database for *all* papers that appeared therein—starting from 2006 (the year in which Facebook was released to the general public), as also done by Antonakaki, Fragopoulou, and Ioannidis (2021). For every paper, we collect its metadata, i.e. title, authors (and affiliations), abstract, publication date, DOI. In total, we extract over 1 million “candidate” papers.

Identifying papers on OSN. There is no straightforward way to determine if any given paper is related to (any) OSN. Hence, we design a sensible (and original) *heuristic*: for each paper (among our 1+ million), we automatically analyze its abstract, looking for *mentions of the name of at least one OSN* (intuitively, a paper on OSN “should” mention an OSN in its abstract). Applying our heuristic requires us to determine a list of OSN that we can use to check the abstracts. To this end, we are inspired by the procedure followed by Fiesler, Beard, and Keegan (2020), and rely on the Wikipedia page on OSN [26]; moreover, to provide a more comprehensive coverage, we expand such a list by including alt-tech OSN [27] as well as YouTube—which, surprisingly, was not considered by Fiesler, Beard, and Keegan (2020) despite its popularity. Overall, we identify 296 unique OSN that we use to apply our heuristic,² leading to 15,770 papers.

Filtering and Validation. Our heuristic may yield some false positives. To provide more accurate results, after removing duplicates, we first exclude all papers that were published before the rollout of the “matching” OSN (e.g., “Discord” was released in 2015, and clearly any paper mentioning “discord” in the abstract before 2015 could not be on this OSN). Then, we scrutinize papers having “homonyms,” i.e., words that can be either used to denote an OSN, or a completely different concept (e.g., “band”): in these cases, we only include the paper if the abstract also mentions other common terms on OSN (e.g., “social network”). Lastly, we perform a manual inspection: for each OSN (out of 296), we inspect up to 50 matched papers, discarding those that, despite meeting the abovementioned criteria, do not focus on OSN. Such a procedure is useful to remove papers on less-known OSN (e.g., in one abstract, the OSN “Ning” was associated with a person named Ning): out of 74 excluded papers, 73 pertained to OSN with less than 50 papers.

SUMMARY. The Minerva-OSN after filtering contains 13,842 OSN-related papers published in 2006–2023 across 135 venues. Our Minerva-OSN is the largest curated dataset of this kind (and adheres to FAIR principles [30]).

Limitations. We cannot claim that Minerva-OSN contains *all* publications on OSN. Some papers may not be in-

²We also compiled a list of OSN synonyms; for instance, “Sina Weibo” is often referred to just as “Weibo” (Yang et al. 2022).

dexed in Scopus. Moreover, our selection of venues, despite its comprehensiveness, may have omitted some journals or conferences. Finally, our heuristic may have not captured papers that did not mention any specific OSN in the abstract.

Topics Identification (technical analysis)

To derive some of our findings, we leverage simple scripts (e.g., cross referencing among various papers) applied to Minerva-OSN. However, one of our goals cannot be feasibly achieved via human effort alone: to identify the topmost “topics” of the 13,842 papers in Minerva-OSN, we rely on artificial-intelligence techniques—and, particularly, on *topic modeling* (Dieng, Ruiz, and Blei 2020). Yet, contrarily to most prior work – also in the OSN context, e.g., Sartori et al. (2023) – we do not “blindly” apply such AI-based techniques: we also (i) guide the topic model with our own expertise in the OSN domain, and (ii) conduct a thorough manual analysis focused on assessing the results of our topic model. Let us explain our procedure, which is systematic in nature.

Topic modeling. The main idea of topic modeling is to use a model to identify “topics” (i.e., thematic patterns) in a collection of documents—which, in our case, are the abstracts of the papers in Minerva-OSN. Our model of choice is the well-known BERTopic (Grootendorst 2022), which we apply through the following steps: (i) we create a sentence embedding for each abstract by using “bge-small-en-v1.5” (one of the best open-source models [2]); (ii) apply dimensionality reduction to such sentence embeddings via UMAP (McInnes, Healy, and Melville 2018); (iii) cluster the reduced sentence embeddings via HDBSCAN (Campello, Moulavi, and Sander 2013); (iv) create topic representations based on the resulting clusters, i.e., we identify the *most relevant words describing the topics clustered*, which we do by using Count Vectorizer, removing English stopwords, and applying c-TF-IDF; (v) finally, we assign a label on the resulting topics by using LLAMA-2 (Touvron et al. 2023).³ Indeed, as recommended by Grootendorst (2022), we use Large Language Models (LLM) to facilitate the human-interpretability of our findings. Specifically, we use LLM in two ways. First, as we mentioned, we use LLAMA-2 for *topic labeling* (i.e., “given a set of keywords that denote a topic, what is a name that can denote them?”); we do this by issuing the prompt in Listing 2 (in the Appendix). Second, we use GPT-4 to *describe the content* of each topic identified by BERTopic (i.e., “given the keywords and the label of a topic, what is a human-readable description of such a topic?”); we do this by issuing the prompt in Listing 1.

Describe in two sentences the research area named “[TOPIC NAME]” and defined by the following keywords: [KEYWORDS].

Listing 1: Topic description prompt to GPT-4.

³**Trial-and-Error and Alternatives.** The procedure described above was derived after extensive testing of various alternatives. E.g., for step (i) we considered also “all-mpnet-base-v2” and “bge-small-en”; for (ii) we considered also PCA; for (iii) we considered also K-Means; and for (v) we considered KeyBERT and MMR.

Human guidance and Results. We found that the “blind” application of the abovementioned workflow led to misleading results.⁴ Hence, we use our domain expertise to guide the topic modeling in providing informative results. Specifically, to prevent the topic model from being biased towards any OSN (typically, popular ones) and to cluster abstracts only based on their actual contents, we replace the specific OSN name of each abstract with a neutral tag (e.g., we replace “Twitter” with “OSN”). After applying such a preprocessing step and re-doing the operations discussed in the previous paragraph, our topic model assigned a topic to 83.9% papers in Minerva-OSN. Specifically, 11,606 papers were associated to one among 17 topics (automatically identified by our model). Such an outcome is in line with observations of prior work (Karabacak and Margetis 2023), since the presence of “outliers” is expected in such a sparse collection of documents.⁵ We report in Figure 2 distribution of the papers in Minerva-OSN across the 17 topics identified by our topic model; in the Appendix, we report the description of each topic (Table 2) and a visualization of our clustering (Figure 10). However, before we present these findings, we must ascertain the quality of our topic model: this is necessary to ensure that our following analysis is well founded.

Tri-party validation. There is no ground truth available that we can use to “automatically” measure the accuracy of our topic model—e.g., we do not know (a priori) if a paper put in the “eLearning” cluster is indeed about eLearning. Hence, following the recommendations by Hagen (2018), we must inspect the results ourselves. However, we found no prior work that systematically assessed the combination of BERTopic and LLAMA2 for the specific purpose of topic modeling for abstracts of OSN-related papers. Actually, most related work (e.g., Koloski et al. (2024)) merely mention some manual review process but without describing the exact methodology. Therefore, to provide a foundation for future work, we designed an original quality-assurance pipeline that encompasses a tri-party (human-driven) reviewing system. Specifically, we proceed as follows. First, for each of the 17 topics identified by our model, the abstracts of 10 different papers are randomly selected. Next, these 170 abstracts are provided to three experts with an academic background (in OSN and AI): these experts *do not know* which topic was assigned to each abstract by the topic model, but they know the list of possible topics (in Table 2) and their description. The experts must then independently select two topics that they think better capture the content of each abstract, specifying their first and second choices. Afterward, we assess (a) how much the experts agreed among each other, and (b) how much the experts agreed with the topic model. The results are positive. Specifically, for (a), we found a remarkable 73.7% agreement among the first choice of each reviewer—denoting that the reviewers share similar views (O’Connor and Joffe 2020).

⁴For instance, the topic model was assigning hundreds of papers to a single cluster labeled “Twitter topics”, which was not useful.

⁵We will still consider these 2,236 outliers in the remainder of our paper—albeit not for topic-related analyses (given that our focus is on the “topmost” topics studied by prior OSN literature).




Figure 2: **Distribution of topics in Minerva-OSN.** After processing the 13,842 papers in dataset, our topic model identified 17 topics which described the abstracts of 83.9% of the papers in Minerva-OSN (topic description in Table 2).

For (b), we found that the topic assigned by BERTopic was placed first by the reviewers in 72.5% of the cases, and first or second in the 80.2% of the cases. We further examined the responses, scrutinizing major occurrences of agreement/disagreement. We found that for more specific topics like “Covid-19 Social Media Analysis” and “Hate Speech Detection” the consensus was higher (first choice in >90% of the cases). Conversely, broader topics such as “Multimedia Retrieval and Tagging” or “Online Video Quality and User Experience” lead to greater disagreement (accuracy slightly >50%). However, even in these challenging cases, we appreciate that the agreement increases by considering also the second choice (accuracy >70%). We can hence conclude that the topics assigned by our topic model are, in general, appropriate—thereby validating our future analyses.

FINDINGS (TOPIC PREVALENCE). From Figure 2, we see that the papers in Minerva-OSN span across a variety of topics. The three most prevalent topics are “Multimedia Retrieval and Tagging” (1,426 papers), “Knowledge Mining” (1,305), “Security and Privacy Risks” (1,110). This result is somewhat expected (e.g., also Antonakaki, Fragopoulou, and Ioannidis (2021) mention many of the themes captured in Figure 2) but we are the first to quantify such prevalence in the whole body of OSN literature.

We stress that our repository [1] includes the entire source code we developed for our analyses, as well as the outcome of manual validation of the topic assignment.

Analyses (“what have we done so far?”)

We now focus on our primary contribution and analyze the Minerva-OSN dataset to derive novel findings on OSN research. We begin by investigating *generic research trends* (RQ1), and then focus on the *authors* of prior works (RQ2). Next, via deeper analyses, we elucidate the diversity of topics studied in specific OSN (RQ3). Then, inspired by the advent of country-specific regulation such as




Figure 3: Yearly distribution of the most prevalent OSN.

the GDPR, we will examine potential differences from a geographical perspective (RQ4). Finally, we explore the relationship between an OSN real-world popularity and the degree of research activity devoted to such an OSN (RQ5).

RQ1: Major OSN research trends

What has OSN literature mostly focused on in 2006–2023? We extend our preliminary analysis (focused on the topmost topics) on Minerva-OSN to gain a better understanding of major research trends over time. Such a temporal analysis represents the cornerstone of our paper. To answer this RQ, we define the two following sub-RQs:

- RQ1. a: *Temporal OSN prevalence*, where we analyze the most researched OSN over time (2006–2023).
- RQ1. b: *Topics over Time*, where we investigate how the main topics evolved throughout the years 2006–2023.

Methodology. For RQ1. a, we analyze the abstracts of each paper in our dataset and identify the OSN considered. We aggregate the results on a yearly basis and create a stacked histogram plot. For RQ1. b we investigate how the 17 topics found by our topic model have evolved over time by analyzing the yearly frequency of their occurrences.

Results. [RQ1. a] Overall, of the 296 OSN we considered when creating Minerva-OSN, only 91 (30.6%) have been investigated by academic work. Among these, 77 OSN have been researched by less than 100 papers, and 52 OSN have even been studied by less than ten papers (we report these details in the Appendix). Figure 3 shows the evolution of the seven most prevalent OSN. In general, the peak of OSN research activity was in 2018 and has been slightly decreasing since.⁶ Twitter is the most prevalent OSN in research, followed by Facebook, Wikipedia, and YouTube. We also observe that Twitter has been predominant in OSN research since 2012 and that Flickr’s share has decreased since 2018.⁷ This clearly highlights that *OSN research is not equally distributed across different OSN but concentrates*

⁶The observed decline in the number of research publications in the year 2023 is due to Scopus not being updated with the latest publications of 2023 (we created Minerva-OSN in Dec. 2023).

⁷This might be related to the fact that Flickr announced the end of its 1 TB free storage plan in November 2018 [3].

on a few, with Twitter being predominant. Even though it was known that Twitter was well-studied in research (Antonakaki, Fragopoulou, and Ioannidis 2021), we are the first to provide quantitative results on the whole OSN panorama.

[RQ1. b] Figure 4 reports the evolution of topics across OSN over time. Different trends can be observed for each specific topic. For instance, consider “Covid-19 Social Media Analysis”: this topic emerged in 2020, peaked in 2021, and has declined since (this expected result confirms the validity of our prior analyses). There are also emerging topics that have found continuous interest from the community. This is the case of “Hate Speech Detection” and “Fake News Detection” for instance. Conversely, topics like “Sentiment Analysis” and “Security and Privacy Risks” are topics that have been always studied by the community. These findings suggest a community that adapts its interest frequently with the evolution of different trends of OSN mass utilization.

RQ1: What has OSN literature mostly focused on in 2006–2023? There is a tendency to focus on a few OSN, Twitter being the most popular. Researchers study OSN for different reasons, ranging from sentiment analysis and recommendation systems to sensitive topics such as hate speech detection. Interest in these topics follows temporal trends that vary—likely tied to real-world attention.




Figure 4: **Topics over Time** (for 8 topics). The y-axis reports the number of papers on each topic. Thick lines show the trend of that specific topic, thin lines show all the others.




Figure 5: Authors’ affiliation country (yearly distribution).

RQ2: Analysis of the Authors of OSN papers

What are the predominant patterns among authors researching OSN? With this research question, we investigate the traits of scientists contributing to the field of OSN. In this way, we provide another perspective on the state of research. We disentangle this RQ with the following analyses:

- RQ2. a: How has the geographical distribution of authors evolved over time?
- RQ2. b: Do OSN researchers tend to interact directly with the platforms that are studied in their papers?
- RQ2. c: How diverse are the topics studied by OSN researchers?

Through this RQ, we seek to examine if OSN research is becoming more inclusive or exclusive to specific communities.

Methodology. We answer RQ2 by analyzing the author’s information contained in our dataset. We recall that Scopus provides some author information for each indexed paper. To answer RQ2. a, we retrieve general statistics such as the total number of authors per year and the country of their affiliation. For RQ2. b, we use regular expressions to identify whether any co-author of a paper is working for a specific OSN.⁸ We report the percentage of papers focusing on an OSN with at least one author affiliated to such an OSN. For RQ2. c, we conduct a cross-analysis of the authors w.r.t. the 17 topics found by the topic model. We examine how many unique topics are (on average) studied by each author.

Results. [RQ2. a] We analyze the contributions of each country to OSN research from 2006 to 2023 by examining the affiliations of the authors for each paper. The findings are presented in Figure 5, showing the 6 most represented countries (we aggregate countries outside the top-6 in an “other” bin; the detailed ranking is in the Appendix). We observe that the US is the leading country in terms of quantitative contributions since the dawn of OSN research. In contrast, for China and India, the number of OSN papers has increased over time (especially recently). Other countries like the UK and Germany provide a stable contribution.

⁸A limitation of this approach is that an author may be affiliated to a given OSN, but the affiliation reported on the paper is a different one. Hence, our numbers may underestimate the reality.

Additionally, we record the number of unique authors who have published at least one OSN paper per year according to Minerva-OSN. The results show a growing interest in the research community: while in 2006 only 48 researchers focused on OSN, the number increased to 1,135 in 2010; 2,707 in 2014; and reached the maximum of 4,111 in 2022. After 2016, the number of unique authors remains stable around 3,900 per year. This finding suggests that OSN research, after over 10 years of growth, reached a plateau after 2016.

[RQ2. b] With respect to authors affiliated with OSN companies, we observe the following: 5.2% of papers focusing on Facebook or Instagram have at least one author affiliated with (i.e., employed by) Facebook or Instagram (in total 2,860 papers focus on Facebook and Instagram). Twitter affiliates are listed as authors in 2.0% of the 5,248 papers on Twitter. Wikimedia (Wikipedia) affiliates are authors in 1.5% of 2,139 papers. We also identify an interesting trend on LinkedIn:⁹ *almost 16.6% of 253 papers on LinkedIn have at least one co-author affiliated with LinkedIn.* Nevertheless, we also found that the number of publications co-authored by an Instagram or Facebook employee is twice as high as those co-authored by a Twitter employee: this is intriguing given that the number of research papers on Twitter is twice as large as that of Facebook or Instagram (see Figure 1). Such a phenomenon is probably driven by Twitter’s API being (until 2023 [13]) freely available to use, thereby enabling “anyone” to obtain Twitter data usable for research on Twitter. Overall, these trends suggest that research on a specific OSN, such as LinkedIn, is only accessible with the direct involvement of the OSN—highlighting the relevance of “data-access policies” enforced by OSN owners (we will elaborate on this aspect in the next section).

[RQ2. c] Nearly 32,000 distinct researchers have co-authored papers included in Minerva-OSN (notably, 2.6% wrote more than five papers on OSN). We analyze how many topics are studied by these authors (on average). We now report statistics for authors with varying impact, i.e., for authors that contributed to at least i publications. For $i = 1$, the average is 1.2 topics; for $i = 5$, it increases to 3.6; for $i = 10$ and $i = 15$, the average number of topics is 5.3 and 6.5, respectively. This finding indicates that authors predominantly focus their efforts on specific research domains, and that more experienced authors tend to expand their coverage of topics eventually—which is a positive result.

RQ2. What are the predominant patterns among authors researching OSN? The US is the leading country for OSN research. Some OSN (e.g., LinkedIn) have a significantly higher percentage of papers co-authored by employees of the corresponding OSN. In general, researchers focus on a few select topics, but authors are likely to expand their coverage as they publish more papers.

RQ3: Topic Distribution across OSN

How do research topics differ across OSN? Motivated by our previous findings showing that most academic work focuses

⁹We did not include Microsoft affiliations since LinkedIn was only acquired by Microsoft in 2016.

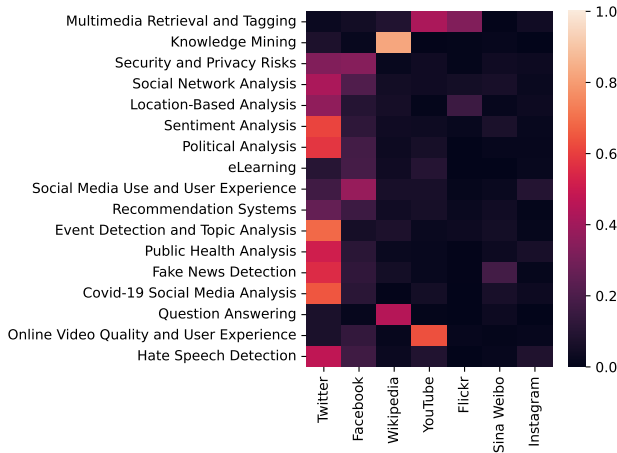


Figure 6: Topic distribution across 7 most prevalent OSN.

on a few OSN, we want to further understand the relationship between research topics and OSN. We examine whether certain topics are primarily addressed using a single OSN (indicating strong dependency) or investigated across multiple OSN (weak dependency). RQ3 has no sub-RQs.

Methodology. We conduct this analysis considering the 17 topics identified in RQ1. b. Recall that our topic model will yield one topic for each paper in our Minerva-OSN dataset; at the same time, each paper is associated to one or more OSN. We then create a function that maps every topic to all the OSN for which there is at least one paper that focuses on that topic. Next, we count the number of unique OSN associated to each topic. Finally, each topic distribution is normalized between 0 and 1.

Results. Figure 6 illustrates the distribution of topics across the seven OSN with the highest number of papers in Minerva-OSN (see Table 1). Intriguingly, Twitter, the most prevalent OSN, monopolizes research topics, reaffirming its foundational role in OSN research. Facebook appears to be relevant on topics such as “Social Media Use and User Experience,” and “Security and Privacy Risks.” Similarly, Wikipedia appears to be strongly related to topics such as “Knowledge Mining,” and “Question Answering.” Also, we note that there are some topics that are not investigated by any of these seven OSN. For instance, “eLearning” is strongly connected to Moodle. These results show that some research topics are strictly dependent on specific OSN.

RQ3. How do research topics differ across OSN?

Most research topics are highly dependent on a specific OSN. Twitter is the OSN with the highest variety of topics; Wikipedia and YouTube dominate on two specific topics, while Flickr dominates on one topic—in common with YouTube (“Multimedia Retrieval and Tagging”).

RQ4: OSN research across the World

How does the OSN-research output differ across geographical areas? This RQ has a twofold purpose. On the one hand,




Figure 7: Left: papers from EEA/non-EEA researchers. Right: the distribution of unique OSN over the years.

we want to investigate if there are some notable disparities in the OSN research pursued in different areas of the World. On the other hand, we want to use our analysis as a proxy to investigate the potential impact that the GDPR had on OSN research.¹⁰ This aspect was touched by Kotsios et al. in 2019, i.e., more than five years ago. Hence, we seek to investigate whether the enactment of the GDPR (in 2018) led to some “changes” in the trends pursued by researchers in different areas of the World—and, specifically, in four major geographical areas: the US, China, India, and the European Economic Area (EEA). For RQ4 we consider two sub-RQ:

RQ4. a: *Regional publication trends*, where we explore whether there has been any shift in publications among researchers from the four areas of choice.

RQ4. b: *Regional OSN diversity*, where we analyze potential changes in the OSN studied by researchers from in the four areas of choice.

Methodology. For RQ4. a, we had extracted the country of the authors’ affiliations (in RQ2. a): we re-use this information as-is for authors from the US, China and India; for the EEA we simply aggregate data of authors from EEA countries. For RQ4. b, we compute the total occurrences of *different* OSN by year for each group of authors from the four considered regions. Additionally, we compute the number of publications among the chosen groups for four popular OSN in 2024: Twitter, Facebook, Reddit and TikTok.

Results. [RQ4. a] Figure 7 (left) visualizes the distribution of research papers over years for authors from the four considered geographical areas. Overall, we can observe a stable rise of publications for each group—until 2017. Indeed, after 2017, we see two contrasting trends: researchers from the EEA and US have a stark drop in the number of OSN-related publications; in contrast, the OSN-research output of researchers from India or China continues to rise.

[RQ4. b] We plot the distribution of the *different* OSN studied every year by researchers from the four considered

¹⁰**Background.** The GDPR (General Data Protection Regulation) is a European regulation released in April 2016, and implemented in May 2018 to protect natural persons in regard to the processing of their personal data. The GDPR inevitably also influenced OSN owners, requiring them to increase user protection and asking for user consent (Ahmed et al. 2020). Despite the GDPR being meant for European citizens, OSN services not located in the EU must still abide by the GDPR if they want to operate in the EU.




Figure 8: **Papers (related to Twitter, Facebook, Reddit, TikTok) authored by EEA and non-EEA researchers over time.** The y-axis of each subfigure has a different scale because we compare trends between EEA and other countries for each OSN (we do not aim at carrying out cross-platform analyses). TikTok was released in 2018 so there is no data before that year.

geographical areas in Figure 7 (right). Researchers from India and China study a smaller number of OSN w.r.t. researchers from the EEA or the US; we observe a small decline of OSN diversity for EEA researchers in 2018–2019. Figure 8 visualizes the yearly changes for Twitter, Facebook, Reddit, and TikTok. We observe that research on TikTok is scarce (and is mostly led by China) and that there has been a decline in Facebook and Twitter after 2018. For Reddit, instead, the number of publications is increasing.

[Discussion.] We conjecture that the findings of RQ3 may be due to the GDPR (effective since 2018), which led to more restrictive data-access policies from OSN-owners—which made certain platforms “harder to study” by many researchers from the EEA or the US. Notably, researchers from the EEA and the US were also those with the highest research output: we posit the reason why OSN-related publications from China or India did not decrease is because of a different “ecosystem” (as also evidenced by, e.g., a smaller diversity of studied OSN) that was not significantly affected by the GDPR. However, these are just hypotheses: there can be other reasons behind the decrease in publications from the EEA/US (e.g., maybe researchers simply turned their attention to other domains). Nonetheless, the fact that the GDPR influenced research (not necessarily on OSN) is recognized (Greene et al. 2019; Siegert et al. 2020).

RQ4. How does the OSN-research output differ across geographical areas? Starting from 2018, researchers from the US and the EEA published less papers on OSN. China and India study less OSN w.r.t. the US and EEA. Researchers from the EEA and the US have similar trends for Twitter and Facebook; however, few researchers from the EEA (and none from India) study TikTok.

RQ5. OSN Popularity: Users vs. Researchers

How does the popularity of OSN (among end-users) relate to their prevalence in research? The goal of RQ5 (which has no sub-questions) is understanding whether some OSN are over- or under-studied w.r.t. their popularity among users—thereby reflecting their importance in the real world.

Methodology. To carry out a meaningful and feasible analysis, we restrict our investigation to the 20 most prevalent OSN in our Minerva-OSN. We then proceed as follows: (1) We take a snapshot of the current OSN popularity based

OSN	Rollout Date	Total papers	Papers since '16	Domain Rank	MAU
Twitter	2006	5,248	3,672	42	611 [18]
Facebook	2004	2,545	1,598	6	3,065 [18]
Wikipedia	2001	2,139	1,006	71	0.13 [19]
YouTube	2005	1,438	960	14	2,000 [32]
Flickr	2004	985	350	>1000	60 [20]
Weibo	2009	573	438	>1000	582 [33]
Instagram	2010	482	445	16	2,000 [18]
Reddit	2005	358	337	Top1000	1,212 [23]
Yelp	2004	319	268	>1000	178 [21]
Moodle	2002	269	175	>1000	-
Foursquare	2009	255	158	>1000	-
LinkedIn	2003	253	159	67	310 [4]
StackOverflow	2008	159	144	>1000	100 [35]
Second Life	2003	119	11	>1000	-
Google+	2011	89	40	>1000	-
Last.fm	2002	85	32	>1000	-
TikTok	2016	82	82	8	1,582 [18]
Telegram	2013	80	79	Top1000	800 [34]
Quora	2010	78	66	Top1000	400 [22]

Table 1: **OSN prevalence in research and real-world popularity.** We report the number of papers of the 20 most prevalent OSN across the entire Minerva-OSN, and the number of papers published since the release of the youngest OSN (TikTok, in 2016); the rank of Cloudflare’s global website ranking for each OSN [29]; the number of monthly active users (in millions); a “-” denotes that we could not find any reliable/exhaustive source. We highlight (in boldface) ranks below the top 100 and active users if above 1 billion.

on its usage (as of April 2024). We refer to two sources: *website ranking* [29], and *monthly active users*. The website ranking reflects the number of unique users accessing a domain over a specific period of time, and allows conclusions about the overall popularity of a certain OSN compared to other domains. Monthly active users is a platform-specific metric about the registered users actively using the OSN on a monthly basis. (2) We capture a historical snapshot of OSN popularity from 2009–2022 based on website ranking to identify how OSN popularity has changed over time by using Alexa’s top website list [25]. For comparison, we compute the number of published papers per OSN for 2006–2023. We capture the historical data for 6 OSN. We exclude TikTok since no historical data is available.

Results. We report the snapshot of OSN popularity in Table 1. Facebook is still very popular in 2024, and TikTok is also among the most popular OSN. Regarding monthly active users (MAU), four OSN have more than a billion active

users: Facebook, Instagram, Reddit and TikTok. Considering that only 82 papers have been published about TikTok, it seems understudied compared to its current impact.

We report the historical analysis of OSN popularity in Figure 1; the x-axis ends in 2022 because we did not find 2023 data for some OSN. Interestingly, Twitter has been the most studied OSN in research papers, while its popularity based on website ranking is constantly decreasing. Conversely, while Facebook’s popularity has been generally stable (i.e., it has been one of the top-ranked domains globally over many years), we observe a declining trend in research.

Altogether, these results reveal an alarming trend: *OSN research is not representative of the real world anymore*. Consider the following: in 2023, among people aged 18–29, the most popular OSN were BeReal, Instagram, Snapchat, and TikTok [28]. These OSN are not well covered in the literature. Therefore, research on OSN can present *nonrepresentative bias*, e.g., Twitter and Facebook are not the most suitable OSN to conduct studies in younger generations, such as on cyberbullying (Whittaker and Kowalski 2015).

RQ5. How does the popularity of OSN (among end-users) relate to their prevalence in research? The existing body of OSN research does not accurately reflect the dynamics of the real world. Studies utilizing Twitter as a primary data source are disproportionately high, notwithstanding its observed decrease in popularity in recent years. Consequently, research topics pertaining to younger demographics may be inadvertently overlooked.

Data-access Policies

We attempt to explain some of the trends that we brought to light by analysing the “data-access policies” that each OSN enacts. In particular, we focus on policies related to the access to data pertaining to the OSN by end-users.

Motivation. One of the main goals of research is advancing our knowledge by analysing factual evidence—i.e., data. In the context of OSN, end-users typically have access only to a small subset¹¹ of the whole data gravitating around the OSN. Hence, from the perspective of a researcher who wants to derive new “general” insights from a given OSN, or who wants to test the general effectiveness of a new method, it is impossible to carry out similar analyses if the data they have access to only pertains to a single (or a restricted subset of) user. It then follows that OSN that provide easy access to their data (e.g., via custom APIs, or public and up-to-date datasets) are more appealing to researchers—thereby explaining why more papers focus on such OSN (e.g., Antonakaki, Fragopoulou, and Ioannidis 2021).

Method. We follow an approach similar to the one by Fiesler, Beard, and Keegan (2020). Specifically, in December 2023, we manually review the policies of eight major OSN: Facebook, Reddit, Snapchat, StackExchange, Steam, TikTok, Twitter, YouTube. We scrutinize whether these OSN (i) *allow end-users to retrieve data* from the OSN

¹¹E.g., a given Facebook user can see the posts shared by their friends and which appear on their news feed; but can hardly observe the activities performed by complete strangers.

in a “legitimate” way, i.e., by using APIs provided by the OSN owners—and not, e.g., via scraping or creating fake accounts; we also review (ii) *what data* can be retrieved. Finally, we scrutinize (iii) *milestones that changed* the policies of a given OSN. To provide long-term evidence of our search, we have saved the web pages documenting our statements in permanent links (reported in the appendix).

Results. Let us report our findings for each OSN:

- *Facebook* provides an API [5], which is free (albeit it has a limit of 200 queries per hour) but it only allows to retrieve information pertaining to the account owner.
- *Reddit* provides an API [6] which has a limit of 1,000 queries per minute for the free version. An important change on Reddit occurred in July 2023 [7]: before this date, there was no limit to the queries that could be performed through this API. To provide some context, a popular app would require over 1.7M USD per month to sustain its traffic with the current pricing [8].
- *Snapchat* provides an API to create posts [9], but it cannot be used to retrieve data from the OSN for research.
- *StackExchange* provides an API [10] which allows free access to its data but with a limit of 10k requests per day.
- *Steam* provides a free API [11] which allows to retrieve of information about news/stats on a given video game; however, it does not allow to retrieve of data pertaining to social interactions among users (e.g., posts on boards).
- *TikTok* provides a free API [12] usable to customize a user’s own posts (similar to SnapChat’s).
- *Twitter* provides an API to read the content posted on its platform [13], but it is not free: the basic version costs 100 USD per month, and allows to read at most 10k posts; whereas the pro version costs 5,000 USD per month and allows 1M posts per month. A significant change in Twitter’s API occurred on March 30th, 2023: before this date, there was a free API that allowed to perform read requests (around 100 every 15 minutes).
- *YouTube* has a free API [14] (with limited quotas) to read data from specific channels/videos or for search queries.

Considerations. It is apparent that these popular OSN hardly enable “independent” researchers to analyze the data posted on their platform—at least today. Indeed, some OSN, such as Twitter and Reddit, were more open to research activities—but this changed in 2023. There are, however, some OSN that do provide APIs for exclusive use to researchers. Yet, obtaining access to such APIs is not straightforward. We provide some examples of such “difficulties”:

- for Tiktok, one must submit an application form, for which a response is provided “within 3–4 weeks” [15];
- for Facebook, it is also required to “sign up to learn more Information about Facebook Open Research and Transparency including access and eligibility” [16];
- following Reddit’s policy change, it was stated that unlimited usage of its API would be given to researchers: yet, as of March 2024, it is still unclear how to do so [17].

Put simply, obtaining access to the data of OSN for research purposes is tough for researchers, given the current

policies—a finding which echoes the one by Fiesler, Beard, and Keegan (2020). We acknowledge, however, that such policies can be well founded. OSN have the data of millions of individuals, and it is reasonable that OSN owners want to ensure that only “trustworthy” parties have access to their data; moreover, accessing such data has a cost, and excessive queries may put a lot of stress on the OSN owners’ servers. Finally, it is also within the OSN owners’ rights not to give away “for free” high-value data that can be monetized [36].

Researchers’ Opinion

What do researchers think about OSN research? Using Minerva-OSN as a starting point, we contact other scientists to grasp their opinions on the state of OSN research.

Methodology

Participant Selection. Due to privacy restrictions in Scopus, authors’ contact information are not readily available (which is why it is not included in our Minerva-OSN). Hence, we manually retrieve contact details. To make our analysis feasible and meaningful, we first identify researchers that co-authored *at least 5 papers* in our Minerva-OSN: this criterion was met by 1,138 authors. We then analyze the most recent publications of these authors to pinpoint the most plausible contact information—which we find after manually reviewing 788 papers. Additionally, recognizing the value of diverse perspectives, we include emails of co-authors of these 788 papers, given that contact authors often differ from the primary author of interest. This approach facilitates broader participation and enriches the survey with insights from younger researchers. In total, we contacted 2,367 authors via an anonymous online questionnaire.

Questionnaire. Our questionnaire is *anonymous*, but we ask for some basic (and optional) information about each participant, such as age, country, and employment status; we also inquire about the participant’s research background, such as which OSN they work on, years of experience in OSN research, and if they are still actively researching this field. We report a copy of the survey in our repository. The questionnaire has the following structure:

- **Data Gathering: Ideal vs Real World.** We ask authors to rate (1=less relevant to 5=more relevant) what option they would prefer when gathering data for OSN research. We specifically mention using open-source datasets, collecting new datasets through official APIs, collecting datasets through methods not involving APIs, and entering contact with OSN organizations. Importantly, for this part, we first ask participants to provide their answer by assuming an “ideal world” (i.e., with no constraints or unlimited resources), and then ask the same question but by considering the implications of the real world.
- **Partnerships with OSN.** We ask closed questions about the authors’ perceptions of the difficulty in collaborating with OSN providers and whether a more straightforward process to (i) partner up with OSN providers and (ii) obtain OSN data would facilitate future research.




Figure 9: **Researchers preference in data gathering: ideal vs. real world.** We indicate preference (1 is low, 5 is high) to rely on open source datasets (“Open”), creating a new dataset through official APIs (“API”) or by scraping (“scrape”), or via a direct channel with the OSN (“OSN”).

- **Reproducibility and Outlook** We ask closed questions about the authors’ perception of the degree of reproducibility of prior research on OSN. The questionnaire ends with an open question in which we ask for comments and feedback on the current state of OSN research.

Sample Description and Limitations. We sent emails to 2,367 authors in May 2024 (we only sent one email to avoid annoying fellow researchers) and collected responses for one week. In total, we received 50 answers (a participation rate of 2%). We acknowledge our sample is small (we do not claim generality) but their opinion is still useful in driving future research. Our participants are mainly from Europe (64%), North America (18%), and Asia (12%); such a distribution may create bias (e.g., Europeans are more affected by GDPR) in their responses. Most respondents are associate professors or above (74%), stressing that the answers we received come from experts in the field. Indeed, 86% of participants have been doing research for more than 5 years, while only 4% is relatively new to the field, with less than 2 years of experience (but their opinion is still scientifically valid). Most have worked on Twitter (86%) and Facebook (72%), followed by Wikipedia, Instagram, and Reddit (around 35% each), which partially reflects the popularity of such OSN in general research (and may, hence, cause some bias). 74% of participants are still doing research on OSN.

Results

Data Gathering (Ideal vs. Real-world). We compare “what researchers would like to use” (in an ideal world) against “what researchers are forced to use” (by accounting for the constraints of the real world). Figure 9 presents the comparative results for both scenarios, with a rating scale from 1 (lowest preference) to 5 (highest). Notably, alternative methods (e.g., scraping) are not among our participants’ priorities. Using APIs was the preferred choice in an ideal scenario, but reusing open-source datasets is the first option in the real world. More participants consider contacting OSN owners in an ideal world w.r.t. the real world.

Partnerships with OSN. When asked if they had engaged in direct research collaborations with any OSN, 78% of par-

participants responded negatively, with 98% of them confirming that such collaborations are challenging to set up (30% consider it “hard,” while 50% “very hard”). These findings may explain why partnerships with OSN have relatively low ratings in the ideal-vs-real-world questions. As a matter of fact, 92% of participants agreed that there should be a more straightforward process to partner up with OSN, given that it could improve the state of OSN research (96% said so).

Reproducibility and Outlook. In terms of reproducibility, while 22% of participants admitted to never having conducted a reproducibility study, the majority found the process challenging. Specifically, 34% described it as “difficult,” while 26% found it “very difficult.” Another 12% remained neutral on the matter. Only 2 participants considered it “easy.” With regards to the last (open and optional question), we report below five excerpts of the statements we received, which we find particularly inspiring:

- “Twitter closed their academic collaboration program, we lost access to any meaningful data for research. Around the same time, Reddit also created restrictions, leaving us with no data to use and almost impossible to do any research unless there is lots of money to pay.”
- “OSN will work with you if you are a paid contractor.”
- “I hope the government will force OSN companies to share their data with the researchers.”
- “OSN should make their data available for academic purposes—research and teaching. OSN are popular media nowadays, especially among young people, and valuable insights can be extracted from them.”
- “IRB and Office of Sponsored Projects tend to have a lot of legal issues and hurdles for university researchers.”

TAKEAWAY. OSN researchers perceive obtaining data (and reproducing prior works) as challenging tasks, and that making them easier would facilitate future research.

Discussion

Our investigation highlights that researchers encounter substantial challenges in studying OSN—a problem that is exacerbated by the numerous restrictions that OSN owners enacted over time. To fix these issues, our findings should become a subject of discussion not only among stakeholders within the OSN ecosystem (i.e., researchers and OSN owners¹²), but also among policymakers. Given the billion-sized userbase of OSN, and the potential of OSN to influence our society, *a smooth development of such platforms should be within the interest of political institutions.* According to our study, researchers have been advocating for such a change—but their “voices” have not been adequately heard, it appears.

Throughout the course of our study, we identified numerous indicators suggesting that the priorities of researchers may not be fully aligned with the real world. The most evident sign of this phenomenon can be appreciated in Figure 1, showing that *Twitter has been the most represented*

¹²We argue that research on OSN should transcend simplistic (but fair and legitimate) visions such as “career trajectories of researchers” or “a yet another source of profit for OSN owners.”

OSN in research since 2011, despite being outside the top-3 OSN (popularity-wise) since 2015. Such an asymmetry is particularly detrimental to younger demographics (who tend to use OSN that are less prevalent in research, such as Instagram and TikTok), as potential risks remain unexplored due to researchers’ propensity to focus on OSN that are “easier to study” – despite the fact that such OSN may not fully capture the entire spectrum of our society.

Let us elaborate on our previous claim. Consider the heatmap in Figure 6: Twitter is a key resource in 10 out of the 17 most commonly studied research topics. Such a finding suggests that Twitter is—or rather, *was*—a cornerstone for OSN research, owing to its open API and policies that permit straightforward access to its data, thereby fostering the collection and distribution of new datasets to the community. Such a characteristic, however, is not valid anymore: the easy accessibility to Twitter’s data was curtailed in April 2023 with the introduction of (expensive!) paid tiers in its API. In contrast, other OSN, such as Instagram and TikTok, have always posed challenges to researchers interested in analyzing their data. We stress, however, that *there is nothing wrong in carrying out research on an “overly-studied” OSN.* However, it is factual that some OSN have a different user base, and lack of research on such platforms is detrimental to their communities (and, hence, to our whole society).

Our findings have also highlighted a decreased “interest” from researchers towards OSN—despite the skyrocketing growth of these platforms’ userbase. Indeed, *starting from 2018, less researchers has been attracted in OSN research, and less peer-reviewed papers on OSN have appeared* (see RQ1. a, RQ2. a, RQ4). This can be linked to the enactment of country-wide policies such as the GDPR, as well as by – potentially – a reduced attractiveness of “overused” OSN to researchers (which may have also negatively affected the desirability of certain submissions to some venues). Hence, the question arises: how will the recent “restrictive” policy changes on well-researched OSN (such as Reddit and Twitter) affect future research on OSN?¹³

Recommendations and Future Work

We have provided a holistic overlook of extant research on OSN by analyzing the metadata of 13,842 publications—which we organized in the Minerva-OSN dataset, which we publicly release. Our results highlight some fundamental issues in the OSN ecosystem. To “rescue” OSN-related research, we make the following recommendations:

- *OSN owners should “lend a hand” to OSN researchers.* Restrictions to OSN data impair research, which is detrimental to all OSN stakeholders.
- *OSN researchers should explore “other” OSN.* Each OSN caters to a specific segment of the world’s population, and investigating such platform may reveal aspects of our society that deserve to be brought to light.
- *Policymakers (at all levels) should contribute, too.* OSN are a pivotal component of our society, and certain policies can be waived for research purposes.

¹³An intriguing initiative that seeks to mitigate the poor access to OSN data is Perspective [24], but it only pertains to toxic content.

Our contributions are meant to guide future work on OSN, inspiring to, e.g., explore new topics (or on new OSN); or to raise awareness of the issues faced by researchers to other relevant stakeholders. Finally, our Minerva-OSN dataset, codebase, and analytical methods can be used by downstream research to investigate other aspects related to OSN.

References

- Acquisti, A.; and Gross, R. 2006. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *PETS*.
- Ahmed, J.; Yildirim, S.; Nowostaki, M.; Ramachandra, R.; Elezaj, O.; and Abomohara, M. 2020. GDPR compliant consent driven data protection in online social networks: A blockchain-based approach. In *ICICT*.
- Andreassen, C. S. 2015. Online social network site addiction: A comprehensive review. *Current addiction reports*.
- Antonakaki, D.; Fragopoulou, P.; and Ioannidis, S. 2021. A survey of Twitter research: Data model, graph structure, sentiment analysis and attacks. *Exp. Syst. Appl.*
- Bessi, A.; and Ferrara, E. 2016. Social bots distort the 2016 US Presidential election online discussion. *First monday*.
- Bieringer, L.; Grosse, K.; Backes, M.; Biggio, B.; and Krombholz, K. 2022. Industrial practitioners' mental models of adversarial machine learning. In *SOUPS*.
- Booth, A.; Sutton, A.; Clowes, M.; and Martyn-St James, M. 2021. *Systematic approaches to a successful literature review*. Sage.
- Calma, J. 2023. Twitter just closed the book on academic research. *The Verge*.
- Campello, R. J.; Moulavi, D.; and Sander, J. 2013. Density-based clustering based on hierarchical density estimates. In *PAKDD*.
- Chu, Z.; Gianvecchio, S.; Wang, H.; and Jajodia, S. 2010. Who is tweeting on Twitter: human, bot, or cyborg? In *AC-SAC*.
- Dieng, A. B.; Ruiz, F. J.; and Blei, D. M. 2020. Topic modeling in embedding spaces. *TACL*.
- Donald, W. E. 2022. Content analysis of metadata, titles, and abstracts (CAMTA): application of the method to business and management research. *Management Research Review*.
- Dwivedi, Y. K.; et al. 2023. "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *IJIM*.
- Fiesler, C.; Beard, N.; and Keegan, B. C. 2020. No robots, spiders, or scrapers: Legal and ethical regulation of data collection methods in social media terms of service. In *ICWSM*.
- Gaur, A.; and Kumar, M. 2018. A systematic approach to conducting review studies: An assessment of content analysis in 25 years of IB research. *Journal of World Business*.
- Greene, T.; Shmueli, G.; Ray, S.; and Fell, J. 2019. Adjusting to the GDPR: The impact on data scientists and behavioral researchers. *Big data*.
- Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv:2203.05794*.
- Hagen, L. 2018. Content analysis of e-petitions with topic modeling: How to train and evaluate LDA models? *Information Processing & Management*.
- Jin, S. V.; Muqaddam, A.; and Ryu, E. 2019. Instafamous and social media influencer marketing. *MIP*.
- Karabacak, M.; and Margetis, K. 2023. Natural language processing reveals research trends and topics over two decades: a topic modeling study. *The Spine Journal*.
- Kim, J.; and Hastak, M. 2018. Social network analysis: Characteristics of online social networks after a disaster. *IJIM*.
- Koloski, B.; Lavrač, N.; Cestnik, B.; Pollak, S.; Škrlić, B.; and Kastrin, A. 2024. AHAM: Adapt, Help, Ask, Model Harvesting LLMs for Literature Mining. In *IDA*.
- Kotsios, A.; Magnani, M.; Vega, D.; Rossi, L.; and Shklovski, I. 2019. An analysis of the consequences of the general data protection regulation on social network research. *TSC*.
- McInnes, L.; Healy, J.; and Melville, J. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv:1802.03426*.
- O'Connor, C.; and Joffe, H. 2020. Intercoder reliability in qualitative research: debates and practical guidelines. *International journal of qualitative methods*.
- Sartori, E.; Pajola, L.; Da San Martino, G.; and Conti, M. 2023. The impact of Covid-19 on online discussions: the case study of the sanctioned suicide forum. In *WWW*.
- Shu, K.; Sliva, A.; Wang, S.; Tang, J.; and Liu, H. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter*.
- Siebert, I.; Varod, V. S.; Carmi, N.; and Kamocki, P. 2020. Personal data protection and academia: Gdpr issues and multi-modal data-collections. *OJAKM*.
- Singhal, M.; Ling, C.; Paudel, P.; Thota, P.; Kumarswamy, N.; Stringhini, G.; and Nilizadeh, S. 2023. SoK: Content moderation in social media, from guidelines to enforcement, and research to practice. In *EuroS&P*.
- Tang, S.; Mi, X.; Li, Y.; Wang, X.; and Chen, K. 2022. Clues in tweets: Twitter-guided discovery and analysis of SMS spam. In *ACM CCS*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv:2307.09288*.
- Whittaker, E.; and Kowalski, R. M. 2015. Cyberbullying via social media. *Journal of school violence*.
- Wilson, R. E.; Gosling, S. D.; and Graham, L. T. 2012. A review of Facebook research in the social sciences. *Perspectives on psychological science*.
- Yang, W.; Wu, Z.; Mok, N. Y.; and Ma, X. 2022. How to Save Lives with Microblogs? Lessons From the Usage of Weibo for Requests for Medical Assistance During COVID-19. In *ACM CHI*.

Acknowledgment

We thank the reviewers for their comments: addressing your feedback improved our paper immensely. We want to extend our gratitude to all the researchers who participated in our study. Part of this research has been funded by the Hilti Foundation.

Paper Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
- (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
- (e) Did you describe the limitations of your work? **Yes**
- (f) Did you discuss any potential negative societal impacts of your work? **Yes, in the Ethics and Limitations.**
- (g) Did you discuss any potential misuse of your work? **Yes, in the Ethics and Limitations.**
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes**
- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**

2. Additionally, if your study involves hypotheses testing...

- (a) Did you clearly state the assumptions underlying all theoretical results? **NA**
- (b) Have you provided justifications for all theoretical results? **NA**
- (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
- (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA**
- (e) Did you address potential biases or limitations in your theoretical framework? **NA**
- (f) Have you related your theoretical results to the existing literature in social science? **NA**
- (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **NA**

3. Additionally, if you are including theoretical proofs...

- (a) Did you state the full set of assumptions of all theoretical results? **NA**

- (b) Did you include complete proofs of all theoretical results? **NA**

4. Additionally, if you ran machine learning experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes, we provided a URL with source code and data.**
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes**
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **NA, as we only utilize unsupervised machine learning.**
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **Yes, specific details are reported in our repository.**
- (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes, the analyses we conducted are confirmed by researchers' opinion in our survey.**
- (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? **NA, as we only utilize unsupervised machine learning.**

5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity...**

- (a) If your work uses existing assets, did you cite the creators? **NA, as we only utilized the dataset we collected.**
- (b) Did you mention the license of the assets? **NA**
- (c) Did you include any new assets in the supplemental material or as a URL? **NA**
- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **Yes, this information is available in the survey description.**
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **NA**
- (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? **Yes. The discussion can be found in the dataset section.**
- (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **Yes**

6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity...**

- (a) Did you include the full text of instructions given to participants and screenshots? **Yes, in the online repository we share a copy of the survey.**
- (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **Yes, we explained in the ethical considerations section that IRB was not required, participants were not exposed to risks, and their privacy was preserved.**

- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA
- (d) Did you discuss how data is stored, shared, and de-identified? Yes

Ethical Considerations and Limitations

Our institutions are aware of the research carried out in this paper, and deemed that a formal IRB approval was not required. Our Minerva-OSN was composed by collecting publicly available information, hence making it public does not violate any copyright. The participants of our expert survey were informed of the purpose of our (anonymous) questionnaire, and they have the option to withdraw their data if they so desire. We discussed the limitations of our contributions in our main paper.

Additional Information

We report in this section additional technical information.

Lists

In this section, we report the lists utilized in our analyses.

Google Scholar Subcategories. Artificial intelligence, Computational Linguistics, Computer Security and Cryptography, Computer Vision and Pattern Recognition, Data Mining and Analyses, Databases and Information Systems, Human Computer Interaction, Multimedia.

Venues. ACM Asia Conference on Computer and Communications Security; ACM Conference on Recommender Systems; ACM International Conference on Information and Knowledge Management; ACM International Conference on Multimedia; ACM International Conference on Multimedia Retrieval; ACM International Conference on Web Search and Data Mining; ACM International Conference Proceeding Series; ACM Multimedia Systems Conference; ACM SIGIR Conference on Research and Development in Information Retrieval ACM SIGKDD International Conference on Knowledge Discovery & Data Mining ACM SIGMOD International Conference on Management of Data; ACM Symposium on Computer and Communications Security; ACM Symposium on User Interface Software and Technology; ACM Transactions on Computer-Human Interaction; ACM Transactions on Intelligent Systems and Technology; ACM Transactions on Knowledge Discovery from Data; ACM Transactions on Multimedia Computing, Communications, and Applications; ACM/IEEE International Conference on Human Robot Interaction; Annual Meeting of the Special Interest Group on Discourse and Dialogue; Applied Intelligence; Applied Soft Computing; Artificial Intelligence Review; Behaviour & Information Technology; Big Data Mining and Analytics; BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP; British Machine Vision Conference; Computational Linguistics; Computational Visual Media; Computer Human Interaction; Computer Speech & Language; Computer Vision and Image Understanding; Computers & Security Conference of the European Chapter of the Association for

Computational Linguistics; Conference of the North American Chapter of the Association for Computational Linguistics; Human Language Technologies; Conference on Computational Natural Language Learning; Conference on Empirical Methods in Natural Language Processing; Conference on Robot Learning; Data Mining and Knowledge Discovery; Engineering Applications of Artificial Intelligence; EURASIP Journal on Image and Video Processing; Expert Systems with Applications; IACR Transactions on Cryptographic Hardware and Embedded Systems; IEEE Conference on Multimedia Information Processing and Retrieval; IEEE European Symposium on Security and Privacy; IEEE International Conference on Advanced Video and Signal-Based Surveillance; IEEE International Conference on Automatic Face & Gesture Recognition; IEEE International Conference on Big Data; IEEE International Conference on Data Mining; IEEE International Conference on Image Processing; IEEE International Conference on Multimedia and Expo; IEEE International Conference on Trust, Security and Privacy in Computing and Communications; IEEE Security & Privacy; IEEE Spoken Language Technology Workshop; IEEE Symposium on Security and Privacy; IEEE Transactions on Affective Computing; IEEE Transactions on Big Data; IEEE Transactions on Circuits and Systems for Video Technology; IEEE Transactions on Dependable and Secure Computing; IEEE Transactions on Fuzzy Systems; IEEE Transactions on Human-Machine Systems; IEEE Transactions on Image Processing; IEEE Transactions on Information Forensics and Security; IEEE Transactions on Knowledge and Data Engineering; IEEE Transactions on Multimedia; IEEE Transactions on Neural Networks and Learning Systems; IEEE Transactions on Pattern Analysis and Machine Intelligence; IEEE Transactions On Systems, Man And Cybernetics Part B, Cybernetics; IEEE/CVF Conference on Computer Vision and Pattern Recognition; IEEE/CVF International Conference on Computer Vision; IEEE/CVF International Conference on Computer Vision Workshops; IEEE/CVF Winter Conference on Applications of Computer Vision; IET Image Processing; Information Processing & Management; Information Systems; International Conference on 3D Vision; International Conference on Advances in Social Networks Analysis and Mining; International Conference on Artificial Intelligence and Statistics; International Conference on Computational Linguistics; International Conference on Data Engineering; International Conference on Intelligent User Interfaces; International Conference on Language Resources and Evaluation; International Conference on Learning Representations; International Conference on Machine Learning; International Conference on Natural Language Generation; International Conference on Pattern Recognition; International Conference on Web and Social Media; International Joint Conference on Artificial Intelligence; International Journal of Child-Computer Interaction; International Journal of Computer Vision; International Journal of Data Science and Analytics; International Journal of Human Computer Studies; International Journal of Human-Computer Interaction; International Journal of Interactive Mobile Technologies; International Society for Music Information Retrieval Con-

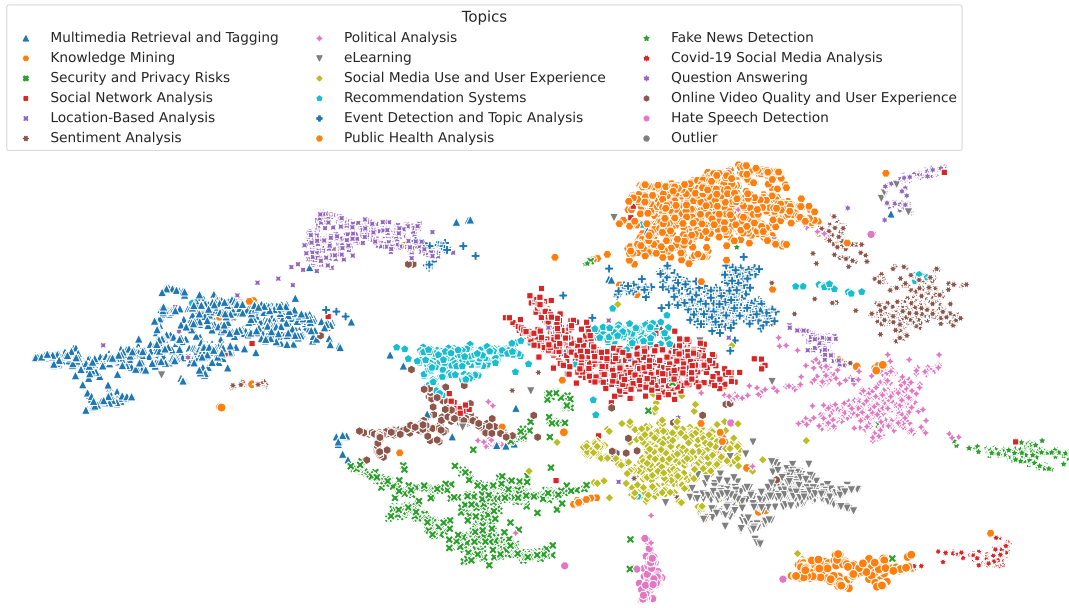


Figure 10: BERTopic visualization.

Topic Name	Description	Size
Multimedia Retrieval and Tagging	It focuses on developing methods to efficiently search and categorize visual content, such as images and videos, often leveraging techniques like face recognition and object detection. This research area is crucial for enhancing user experiences on OSN by enabling more accurate and relevant content discovery.	1,426
Knowledge Mining	It is a research area focused on identifying and associating textual mentions of entities with their corresponding entries in a knowledge base, ensuring accurate integration of information from diverse sources. This field leverages semantic analysis and data integration techniques to enhance the understanding and organization of information across the web and various datasets.	1,305
Security and Privacy Risks	The research area focuses on identifying and mitigating the privacy and security threats that users face on OSN. This includes protecting user information from unauthorized access, preventing data breaches, and addressing issues like spam and fake accounts.	1,110
Social Network Analysis	It is a research area that uses networks and graph theory to study social structures by mapping and measuring relationships and flows among people, groups, or organizations. It helps identify influential users, information dissemination patterns, and the overall connectivity within OSN.	1,082
Location-Based Analysis	It integrates social networking with location-based services, allowing users to share and access geo-related information such as locations, trajectories, and points of interest (POIs) in real-time. This research area leverages vast amounts of spatiotemporal data to analyze user behavior and interactions, offering applications in marketing, recommendations, and urban planning.	876
Sentiment Analysis	It is a research area focused on identifying and classifying emotions expressed in text, particularly on OSN like Twitter. It involves using language-based models to analyze and categorize sentiments in tweets and other text data.	756
Political Analysis	This research area examines how OSN influence political behavior and public opinion, particularly during elections. This research area involves studying user interactions, media consumption, and the dissemination of political information to understand and analyze the impact on public discourse and electoral outcomes.	746
eLearning	This research area explores how students interact with and participate in virtual educational settings, focusing on factors like course design, teacher involvement, and the use of OSN. This research aims to enhance learning outcomes by understanding and improving the ways students engage with online courses and educational tools.	731
Social Media Use and User Experience	The research area focuses on studying how people interact with OSN and the impact of these interactions on their personalities and behaviors. It examines various aspects of social media use, including user behavior, online engagement, and the psychological effects on individuals.	698
Recommendation Systems	Recommendation Systems are algorithms designed to suggest relevant items to users based on their preferences and behaviors, often leveraging OSN and user interactions. These systems analyze data from users and items to provide personalized recommendations, enhancing user experience across various platforms.	635
Event Detection and Topic Analysis	This research area focuses on identifying and analyzing significant occurrences (events) from social media platforms like Twitter by examining user-generated content such as tweets, hashtags, and other data. This research area leverages techniques to detect real-time events, track topics, and analyze trends within OSN to provide insights into social dynamics and emergent phenomena.	614
Public Health Analysis	This research area involves examining data from OSN like tweets to understand and address public health issues such as mental health, drug use, and depression. This research area leverages social media data to identify trends, monitor health-related behaviors, and inform public health interventions.	333
Fake News Detection	This research area focuses on identifying and mitigating the spread of false information, rumors, and misinformation across social media and news platforms. It involves analyzing the credibility of information and sources to ensure the accuracy and reliability of the content shared with the public.	318
Covid-19 Social Media Analysis	This research area examines how social media platforms were used to disseminate information, track public sentiment, and spread misinformation during the Covid-19 pandemic. This research area focuses on analyzing tweets and other social media content to understand public health communication, vaccine perceptions, and the impact of misinformation on public behavior.	282
Question Answering	It is a field of research focused on developing systems that can automatically respond to questions posed in natural language by retrieving relevant information from a knowledge base or generating answers based on learned models. This area encompasses various techniques and models to understand, interpret, and generate human-like dialogue, aiming to provide accurate and contextually appropriate answers.	251
Online Video Quality and User Experience	The research area focuses on evaluating and enhancing the quality of video streaming services to ensure optimal user satisfaction. It involves analyzing various factors such as network conditions, content delivery, and user feedback to improve the Quality of Experience (QoE) for live and on-demand video content.	235
Hate Speech Detection	This research area focuses on developing algorithms and tools to identify and reduce harmful language, including hate speech and cyberbullying, on OSN. This involves analyzing text for offensive and hateful content to create safer and more inclusive digital environments.	208

Table 2: Description of the topic identified in Minerva-OSN.

ference; International Workshop on Quality of Multimedia Experience; International Workshop on Semantic Evaluation; International World Wide Web Conferences; Journal of Big Data; Journal of Information Security and Applications; Journal of Visual Communication and Image Representation; Knowledge and Information Systems; Knowledge-Based Systems; Language Resources and Evaluation; Lecture Notes in Computer Science; Medical Image Analysis; Meeting of the Association for Computational Linguistics; Multimedia Systems; Multimedia Tools and Applications; Multimodal Technologies and Interaction; Neural Computing and Applications; Neural Information Processing Systems; Neural Networks; Neurocomputing; Online Social Networks and Media; Pattern Recognition; Pattern Recognition Letters; Proceedings of the ACM on Human-Computer Interaction; Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies; Security and Communication Networks; Semantic Web; SIAM International Conference on Data Mining; Signal Processing: Image Communication; Social Network Analysis and Mining; Symposium On Usable Privacy and Security; The Journal of Machine Learning Research; Transactions of the Association for Computational Linguistics; Transactions on Asian and Low-Resource Language Information Processing; Universal Access in the Information Society; USENIX Security Symposium; Virtual Reality; Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery; Workshop on Innovative Use of NLP for Building Educational Applications; Workshop on Machine Translation; Workshop on Representation Learning for NLP; World Wide Web.

OSN. We now report information about the OSN we analyzed accompanied by their number of mentions in OSN research abstracts.

- *Widely studied OSN (at least 100 mentions):* Twitter (5248), Facebook (2545), Wikipedia (2139), YouTube (1438), Flickr (985), Sina Weibo (573), Instagram (482), Reddit (358), Yelp (319), Moodle (269), Foursquare City Guide (255), LinkedIn (253), Stack Overflow (159), Second Life (119).
- *Frequently studied OSN (between 20 and 100):* Google+ (89), Last.fm (85), TikTok (82), Telegram (80), Quora (78), douban (74), Pinterest (64), Myspace (63), Delicious (55), Tumblr (47), Meetup (46), WordPress (45), Snapchat (39), LiveJournal (33), ResearchGate (29), Flixster (23), LibraryThing (23).
- *Rarely studied OSN (less than 20 mentions):* Goodreads (19), Duolingo (17), Discord (16), VK (16), Gab (14), Renren (12), 4chan (12), XING (12), ICQ (7), Stumble-

Upon (7), CaringBridge (6), Periscope (6), Yammer (6), Mastodon (6), CouchSurfing (6), Diaspora* (6), Parler (6), PatientsLikeMe (5), Steemit (5), Blogger (5), Plurk (5), Grindr (5), mixi (5), hi5 (4), Tuenti (4), Skyrock (4), aNobii (4), SoundCloud (4), Steam (3), Habbo (3), Ask.fm (3), DailyStrength (3), Cyworld (3), Koo (2), Academia.edu (2), Fotolog (2), BitChute (2), Chictopia (2), Ning (2), Buzznet (2), FilmAffinity (2), Friendica (2), Qzone (1), Odnoklassniki (1), Voat (1), eToro (1), Identi.ca (1), Blind (1), StudiVZ (1), DLive (1), Inspire (1), Xanga (1), Foursquare Swarm (1), Epik (1), Bebo (1), 8kun (1), Fetlife (1), About.me (1), Tagged (1), Fictocracy (1).

- *Never studied OSN:* 23snaps, 8tracks.com, Amikumu, Anphabe.com, AsianAve.com, Athlinks, Band, BeReal, Biip.no, BitClout, BlackPlanet, Brainly, Busuu, CafeMom, Care2, Cellufun, Chess.com, Classmates.com, Cloob, CloutHub, CozyCot, CrossFit, Crunchyroll, Cucumbertown, DTube, DXY.cn, Dayviews, Dead Runners Society, DeviantArt, Dol2day, Doximity, Draugiem.lv, Dreamwidth, Dulwich On-View, Elixio, Ello, English, baby!, Experts Exchange, Faces, Fark, FictionCity, Fieldoo, Fillos de Galicia, Filmow, Fishbrain, Flickchart, Focus.com, Fotki, Frank, Fyuse, Gaia Online, GameFAQs, GameTZ.com, Gapo, Gapyear.com, Gays.com, Gaysir, Geni.com, Gentlemint, Gettr, GirlsAskGuys, GiveSendGo, Glee.com, GovLoop, HR.com, Hello, Hospitality Club, Houseparty, Hub Culture, I Had Cancer, IMVU, IRC-Galleria, Ibibo, Indaba Music, Influenster, Infogalactic, ItsMy/GameCloud, JamiiForums, JustPaste.it, Kaixin001, KakaoStory, Kiwibox, KizlarSoruyor, Kobo, Koofers, Kroogi, Labroots, LambdaMOO, Letterboxd, LifeKnot, Likee, LimeWire, LinguaLeo, LinkExpats, Listography, Lunchclub, MEETin, Marco Polo, MeWe, Metapedia, Miaopai, Micro.blog, Minds, MocoSpace, MouthShut.com, Mubi, My World@Mail.Ru, MyHeritage, MyLife, NK.pl, Nasza-klasa.pl, Nearby, Newgrounds, Nexopia, Nextdoor, Odyssee, Open Diary, PEERtrainer, Partyflock, Patriots.win, Peach, PewTube, Pink Petro, Pixnet, Plaxo, Play.fm, Playlist.com, Poa.st or Poast, Portfolium, Postcrossing, Quechup, RallyPoint, Raptr, RateItAll, Ravelry, Readgeek, Reverbnation, Rooster Teeth, Rumble, Ryze, Sarahah, Sgrouples, Sharesome, Skoob, Snow, Solaborate, Something Awful, Soup.io, Spacehey, Spaces, Spot.IM, Spoutible, Stage 32, SubscribeStar, TV Time, TV Tropes, TakingITGlobal, Tal Canal, Talenhouse, TalkBizNow, Taringa!, TermWiki, The Meta Network, The Sphere, The Student Room, The WELL, Thinkspot, Threads, Total Recut, TravBuddy.com, Travellerspoint, Triller, Trombi, Truth Social, Twoo.com (netlog), Untappd, Vampirefreaks.com, Vero, Viadeo, Vingle, Virb, WAYN, WT Social, Warm Showers, Wappad, We Heart It, Werkenntwen, Whisper, WikiWikiWeb, Woxie, Wozworld, WriteAPrisoner.com, Wykop.pl, Xt3, YTMND, Yo, aSmallWorld, alimero, beBee, iNaturalist, italki, zoo.gr, İnci Sözlük. .

```
<s>[INST] <<SYS>> You are a helpful, respectful and honest assistant for labeling topics. <</SYS>> I have a topic that contains the following documents: [DOCUMENTS] The topic is described by the following keywords: '[KEYWORDS]'. Based on the information about the topic above, please create a short label of this topic. Make sure you to only return the label and nothing more. [/INST]
```

Listing 2: Topic labeling prompt.

Top50 authors' countries (in the entire Minerva-OSN).

United States (6220), China (4151), India (1240), Germany (1079), United Kingdom (1036), Japan (892), Italy (841), Australia (789), Spain (714), France (616), Canada (604), Singapore (571), Brazil (487), Taiwan (469), South Korea (449), Netherlands (391), Greece (368), Hong Kong (358), Portugal (277), Switzerland (249), Austria (240), Ireland (196), Saudi Arabia (187), Turkey (178), Russian Federation (167), Iran (159), Finland (156), Israel (148), Poland (145), Malaysia (143), Norway (138), Indonesia (137), Tunisia (122), Mexico (120), Pakistan (120), Viet Nam (116), Qatar (114), New Zealand (101), Belgium (99), Denmark (97), Sweden (96), Thailand (87), United Arab Emirates (75), Egypt (71), Chile (61), Czech Republic (59), Argentina (53), Philippines (53), Romania (52), Bangladesh (47).

List of External Links

We provide the full list of links referenced in the paper:

- [1] <https://anonymous.4open.science/r/Minerva-OSN-3841>
- [2] <https://huggingface.co/spaces/mteb/leaderboard>
- [3] <https://blog.flickr.net/en/2018/11/01/changing-flickr-free-accounts-1000-photos/>
- [4] <https://www.linkedin.com/pulse/important-linkedin-statistics-data-trends-oleksii-bondar-pqliie/>
- [5] <https://web.archive.org/web/20231130194404/https://developers.facebook.com/docs/graph-api/>
- [6] <https://web.archive.org/web/20231123193950/https://www.reddit.com/dev/api/>
- [7] <https://web.archive.org/web/20231210011040/https://www.redditinc.com/blog/apifacts>
- [8] https://web.archive.org/web/20231210090528/https://www.reddit.com/r/apolloapp/comments/13ws4w3/had_a_call_with_reddit_to_discuss_pricing_bad/
- [9] <https://web.archive.org/web/20231129021701/https://developers.snap.com/creative-kit>
- [10] <https://web.archive.org/web/20231212231535/https://api.stackexchange.com/docs>
- [11] <https://web.archive.org/web/20231204034413/https://steamcommunity.com/dev>
- [12] <https://web.archive.org/web/20231210002644/https://developers.tiktok.com/>
- [13] <https://web.archive.org/web/20231130070701/https://developer.twitter.com/en/docs/twitter-api>
- [14] <https://web.archive.org/web/20240303190109/https://developers.google.com/youtube/v3/docs/search/list>
- [15] <https://web.archive.org/web/20231128205404/https://developers.tiktok.com/products/research-api/>
- [16] <https://web.archive.org/web/20231210080026/https://fort.fb.com/researcher-apis>
- [17] https://web.archive.org/web/20240512003912/https://www.reddit.com/r/AskAcademia/comments/1b32i9q/accessing_reddit_data_for_academic_purposes/?rdt=37952
- [18] <https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/>

- [19] <https://en.wikipedia.org/wiki/Wikipedia:Wikipedians>
- [20] <https://photutorial.com/flickr-statistics/>
- [21] <https://thesmallbusinessblog.net/how-many-people-use-yelp/>
- [22] <https://business.quora.com/resources/reach-over-400-million-monthly-unique-visitors-on-quora/>
- [23] <https://www.statista.com/forecasts/1309791/reddit-mau-worldwide>
- [24] <https://perspectiveapi.com/research/>
- [25] https://web.archive.org/web/*/https://www.alexa.com/topsites
- [26] https://en.wikipedia.org/wiki/List_of_social_networking_services
- [27] <https://en.wikipedia.org/wiki/Alt-tech>
- [28] <https://www.statista.com/statistics/1337525/us-distribution-leading-social-media-platforms-by-age-group/>
- [29] <https://radar.cloudflare.com/domains>
- [30] <https://www.go-fair.org/fair-principles/>
- [31] https://scholar.google.com/citations?view_op=top_venues&hl=en&vq=eng
- [32] <https://www.statista.com/statistics/1314183/youtube-shorts-performance-worldwide/>
- [33] <https://www.statista.com/statistics/795303/china-mau-of-sina-weibo/>
- [34] <https://www.statista.com/statistics/234038/telegram-messenger-mau-users/>
- [35] <https://usesignhouse.com/blog/stack-overflow-stats/#:~:text=How%20many%20registered%20users%20of%20Stack%20Overflow%20are%20there%3F,-Want%20a%20link&text=As%20of%20November%202022%2C%20there%20are%20more%20than%20100%20million,and%2023%20million%20registered%20users.>
- [36] <https://web.archive.org/web/2/https://www.nytimes.com/2023/04/18/technology/reddit-ai-openai-google.html>