

Join the Chat: How Curiosity Sparks Participation in Telegram Groups

Original

Join the Chat: How Curiosity Sparks Participation in Telegram Groups / Paoletti, Giordano; Almeida, Jussara M.; Vassio, Luca; Gonçalves, Marcos André; Mellia, Marco. - 19:(2025), pp. 1493-1509. (International AAAI Conference on Web and Social Media Copenhagen (DK) June 23–26, 2025) [10.1609/icwsm.v19i1.35884].

Availability:

This version is available at: 11583/3000951 since: 2025-06-16T08:09:22Z

Publisher:

AAAI

Published

DOI:10.1609/icwsm.v19i1.35884

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Join the Chat: How Curiosity Sparks Participation in Telegram Groups

Giordano Paoletti¹, Jussara M. Almeida², Luca Vassio¹, Marcos André Gonçalves², Marco Mellia¹

¹ Politecnico di Torino, Turin, Italy

² Universidade Federal de Minas Gerais, Belo Horizonte, Brazil

{giordano.paoletti, luca.vassio, marco.mellia}@polito.it, {jussara, mgoncalv}@dcc.ufmg.br

Abstract

This study delves into the mechanisms that spark user curiosity driving *active* engagement within public Telegram groups. By analyzing approximately 6 million messages from 29,196 users across 409 groups, we identify and quantify the key factors that stimulate users to actively participate (i.e., send messages) in group discussions. These factors include social influence, novelty, complexity, uncertainty, and conflict, all measured through metrics derived from message sequences and user participation over time. After clustering the messages, we apply explainability techniques to assign meaningful labels to the clusters. This approach uncovers macro categories representing distinct curiosity stimulation profiles, each characterized by a unique combination of various stimuli. Social influence from peers and influencers drives engagement for some users, while for others, rare media types or a diverse range of senders and media spark curiosity. Analyzing patterns, we found that user curiosity stimuli are mostly stable, but, as the time between the initial message increases, curiosity occasionally shifts. A graph-based analysis of influence networks reveals that users motivated by direct social influence tend to occupy more peripheral positions, while those who are not stimulated by any specific factors are often more central, potentially acting as initiators and conversation catalysts. These findings contribute to understanding information dissemination and spread processes on social media networks, potentially contributing to more effective communication strategies.

Introduction

Telegram is an instant messaging platform that has gained significant popularity since its launch in 2013. As of March 2025, Telegram has more than 1 billion monthly active users¹. One of Telegram distinctive features is its support for large groups that can accommodate up to 200,000 members making it a powerful tool for discussion within a vast audience.

In this complex scenario, studying the dynamics that drive user engagement - i.e., the active posting of content in groups - is key to understanding the processes that guide information dissemination. Prior studies have analyzed *user's engagement* in social media platforms (Aldous, An, and Jansen 2019; Vassio et al. 2022) using various metrics (e.g., number of likes,

comments, shares) to capture how users *consume content and interact with others*. Specifically on Telegram, Hashemi and Zare Chahooki (2019) proposed indicators to distinguish high-quality from low-quality groups, marked by irregular activity, frequent spam, and unclear topics. Hoseini et al. (2024) studied the dynamics of information dissemination in groups, examining message forwarding and lifespan.

The large range of features makes Telegram an interesting platform for studying online social movements, extremism, and misinformation spread (Baumgartner et al. 2020; Urman and Katz 2022; Venâncio et al. 2024; Kloo, Cruickshank, and Carley 2024). Evidence suggests that some groups even use Telegram to coordinate illicit activities, such as pump-and-dump schemes in the cryptocurrency market (Xu and Livshits 2019), as well as form pods where members interact with each other's content to artificially boost the popularity of their Instagram accounts (Weerasinghe et al. 2020).

We here look into user engagement on Telegram from a novel and complementary point of view: we study what drives user participation in group discussions by posting content, rather than as a passive observer. We call this *curiosity*. It is a fundamental human cognition element that, when adequately stimulated, becomes a driving force behind a person's choice to consume and, in the present case, *to produce information* (Sousa, Almeida, and Figueiredo 2022). As such, by focusing on factors that stimulate user curiosity to join discussions, we offer a novel perspective of the processes contributing to information spread on online media platforms such as Telegram.

Human curiosity modeling finds its roots in the Psychology domain, with theories such as the Information Gap Theory (Loewenstein 1994; Kidd and Hayden 2015), which tries to explain how one's curiosity responds to external stimuli, as one of its foundational pillars. Also, according to the literature (Berlyne 1960), one's curiosity stimulation process has many facets, manifested as stimulus components, referred to as *collative variables*. In Computer Science, a number of recent studies have proposed metrics to operationalize (some of) those variables in various scenarios, often using them as components of information services, notably recommendation systems (e.g., Fu and Niu (2023); Xu et al. (2019)).

In this work, we aim at studying *curiosity stimulation driving user participating in discussions* in public Telegram groups. Our investigation builds upon the work of Sousa,

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹"Du Rove's Channel". <https://t.me/durov/404>. Accessed April 26, 2025.

Curiosity Stimulation

Background

Previous works on Psychology have provided theoretical background to better understand human curiosity. According to (Loewenstein 1994), our organism receives various stimuli from the environment and selectively responds to those that generate more pleasure and raise our curiosity. Berlyne (1960) identified four collative variables pertaining to the process of curiosity stimulation: *novelty*, *complexity*, *uncertainty*, and *conflict*. These variables have been extensively validated in Psychology as main drivers of human curiosity and engagement, forming a cornerstone of modern cognitive and behavioral research field. The author argued that these variables can be quantified using information theory metrics, an idea later reinforced by Silvia (2006), who demonstrated the mathematical feasibility of such quantifications. This formalization bridges theoretical constructs and quantifiable metrics, ensuring a rigorous grounding in Psychology while enabling computational applications.

Inspired by these concepts, recent Computer Science studies proposed metrics to capture one or more of these variables, using them as components of information services (e.g., recommendation functions) (Zhao and Lee 2016; Xu et al. 2019; Fu and Niu 2023; Yan and Chen 2023; Deng et al. 2023; Wang and Chen 2023; Tang and Liao 2022). Sousa, Almeida, and Figueiredo (2019) presented one of the first efforts to operationalize all four of Berlyne’s collative variables to study curiosity stimulation in online music consumption. Similarly, Fu and Niu (2023) proposed metrics based on information theoretical measures, following the principles argued by Berlyne (1960) and Silvia (2006). By leveraging these well-established principles, these prior studies demonstrated the validity and reliability of Information Theory metrics in operationalizing curiosity stimulation across diverse contexts. These metrics are quantified in bits, facilitating direct comparisons across systems (Timme and Lapish 2018).

Subsequently, Sousa, Almeida, and Figueiredo (2022) introduced metrics to quantify curiosity stimulation that motivates users to share content in public, politically oriented WhatsApp groups. They emphasized the importance of incorporating *social influence* as a fifth factor in curiosity stimulation, complementing the four collative variables previously discussed. The inclusion of social influence aligns with the broader psychological framework defined by Kashdan et al. (2018), which emphasizes the exploratory behaviors and interpersonal dynamics driving curiosity in group settings. Following the findings of Ver Steeg and Galstyan (2013), which demonstrated the utility of information-theoretical metrics for measuring social influence, they proposed metrics to account for all these factors. Nevertheless, their analysis focused exclusively on social influence metrics to identify *profiles* of curiosity stimulation. Our work adopts these metrics to explore curiosity stimulation within the domain of Telegram group discussions, building on the psychological and computational insights from these prior studies.

Our goal is to study curiosity stimulation as a key factor which drives *users to actively engage in Telegram group discussions* by posting messages, i.e., from the perspective

Almeida, and Figueiredo (2022), which proposed metrics to measure curiosity stimulation in Brazilian politically oriented WhatsApp groups. Our study extends and diverges significantly from this prior work in four key aspects:

- (i) The prior work explored only metrics related to “social influence”. Instead, here we take a comprehensive approach by considering multiple complementary facets of curiosity stimuli, which allows us to fully characterize user’s curiosity and its role in driving their participation in group discussion.
- (ii) Unlike prior work, which was limited to a single topic (politics) within a specific cultural and linguistic context (Brazilian election), our study spans diverse topics, across multiple groups, and a heterogeneous population. This broad scope mitigates biases from local effects, enhancing the generalization of our findings and enabling the exploration of how curiosity profiles vary across different discussion topics.
- (iii) We analyze the connection between curiosity and the user’s engagement (i.e., the number of messages they post), a critical factor for understanding the role of curiosity in information spread that was not considered in the prior work. This analysis sheds light on how curiosity influences user activity participation in group discussions.
- (iv) Telegram groups offer a unique and novel environment for this analysis due to their much larger membership², which creates a richer and more varied set of potential sources of curiosity stimulation driving the discussions.

Our study is guided by the following research questions:

1. **RQ1:** Which curiosity *stimuli* more often drive user participation? Do the stimuli experienced by users change over time as they join the group discussions? Does the discussion topic affect the prevalence of certain stimuli?
2. **RQ2:** What types of *stimuli* activate the most engaged users? What types relate to those who are most influential in initiating and perpetuating conversations?

Regarding RQ1, our analysis identifies six major *curiosity stimulation profiles*. Users generally maintain stable curiosity profiles, though, over time, they may shift between profiles in response to varying stimuli. The distribution of curiosity stimulation profiles varies across topics. For instance, users in political groups are mostly motivated by diverse and balanced discussions. To answer RQ2, the impact of specific profiles on participation and influence is notable: users drawn to direct peer interactions often engage less, while those motivated by uncertainty are more active. Users without specific curiosity-driven stimuli tend to hold central roles in discussions, often initiating and steering conversations.

In the following, we introduce the background on curiosity stimulation and the metrics to capture its main components (Sec. *Curiosity Stimulation*). We then describe the data (Sec. *Telegram Groups Dataset*) we use in our study and discuss the main results (Sec. *Analysis of Curiosity Stimulation*). Finally, we offer conclusions and directions for future work (Sec. *Conclusions and Future Work*). Appendix A delivers additional results.

²Whatsapp groups were limited to 250 users in the prior work (the limit has grown to 1000 users).

of *content production*. This contrasts with a body of studies of user engagement from the perspective of content consumption (e.g., Shivaram et al. (2024); Aldous, An, and Jansen (2019); Vassio et al. (2021)). Our contribution lies in conducting a comprehensive analysis of all factors influencing curiosity stimulation, investigating a broader range of topics across a heterogeneous user population, and exploring how curiosity stimulation operates in the larger user base of Telegram groups, which differs from the smaller user bases in previous studies. In addition, for the first time, we study the relationship between curiosity and user engagement.

Curiosity Stimulation Metrics

One’s curiosity is constantly stimulated by signals from the environment (Loewenstein 1994). In this work, we study the user curiosity stimulation within a given Telegram group. Our key assumption is that the messages recently posted in the group act as a stimulus for driving a user to post a message, i.e., past messages stimulate the user’s curiosity.³ As such, we model each Telegram group as the sequence of messages posted over time and assume that stimulation may evolve as users post messages. Thus, for a given group, we measure curiosity stimulation by computing metrics associated with the collative variables at each *message* (also referred to as *posting event*, hereafter). For this, we define the *window of interaction* as a time interval ΔT preceding the current posting event. Based on all messages posted during $[t - \Delta T, t]$ in the group, we compute metrics of curiosity stimuli driving the user to post a message at time t .

To select ΔT , we build on prior studies of collective attention in WhatsApp groups, which adopted a window of attention of 30 minutes (Caetano et al. 2021). In the absence of a corresponding analysis for Telegram, we use explicit replies⁴ as a proxy for attention, observing that 93.7% of users have a median reply time below 30 minutes. The choice $\Delta T = 30$ minutes ensures that 85.5% of messages have at least 10 preceding messages within the same time window, meeting the minimum volume requirement for reliable metric computation (see next Section). Thus, setting ΔT to 30 minutes strikes a good balance between filtering noise from the user’s attention period and maintaining sufficient data coverage (see Fig. 10 in Appendix A).

Towards computing the metrics, we represent each message by the (anonymized) sender ID, the timestamp, and the set of media types (i.e., text, URL, image, video, audio, poll, or other⁵) in the message. We refer to the latter as the media *categories* associated with the message. Note that we do not exploit the message content itself which may depend on the context (e.g., topic, language, slang) associated with a particular group. By only using macroscopic category properties, we aim to compare curiosity metrics across groups, language, topics, etc., irrespective of such contextual

³If the same user is active in multiple groups, their curiosity stimulation is analyzed separately and independently for each group.

⁴On Telegram, a user can explicitly reply to a message by mentioning it.

⁵*Other* aggregates all other media that can be posted on Telegram but were found to be rarely used (stickers, gifs, documents, etc.).

properties. Note that we can apply this approach even when the actual content is not available (e.g., due to privacy issues, or because the message itself contains no text).

To quantify the collative variables of (i) novelty, (ii) uncertainty, (iii) conflict, (iv) complexity and (v) social influence, we compute 9 metrics in total, extending the derivations proposed in Sousa, Almeida, and Figueiredo (2022). Given a message posting event, we compute each metric taking either all the users who posted messages in ΔT (*user* metrics) or the categories of media used in those messages (*category* metrics). Table 1 provides an overview of the 9 metrics. Below, we summarize their operational formulation and intuition.

Operational Definitions

The curiosity driving the posting of a message by a given user u at a given time t depends on the past messages posted in the group (including messages posted by u himself) in the interaction window $[t; t - \Delta T]$.

We define the sets of users, groups, and types of media as U, G, M respectively. We do not assume access to message content and use only the different categories of media used to compose the message to represent it. Each message under analysis is then a tuple (u, C_m, g, t) indicating that it was posted by user $u \in U$ in group $g \in G$ at time t with media belonging to set $C_m \subseteq M$. Let U_g be the set of users in group g . Given these definitions, we define the different collative metrics as follows:

Novelty: The *userNovelty* metric captures the novelty of the posting experience for user u . It reflects that users who have posted frequently in recent times find the discussed topic less novel since they are already engaged in the debate, which in turn affects their curiosity. The metric is detailed in Tab. 1, where $n_{v|t,g}$ is the number of messages sent in the window of interaction $[t; t - \Delta T]$ by any user v (including u) in group g . If no previous message is sent by u , we set the novelty to its maximum value possible, which corresponds to the uniform distribution of users, i.e. $-\log_2(1/|U_{t,g}|)$ where $U_{t,g}$ is the set of distinct users who posted content in group g during the current window of interaction.

Similarly, we define *catNovelty*, a novelty metric related to the message categories, based on the surprisal (i.e., Shannon information content) associated with that variable. It captures how surprising it is that the message posted at time t contains the media categories in C_m . As shown in Tab. 1, it depends on $P_{t,g}(C_m)$, the average probability of the media categories posted in the message. This measure is estimated as

$$P_{t,g}(C_m) = \frac{1}{|C_m|} \sum_{c \in C_m} \frac{n_{c|t,g}}{\sum_{c' \in C_{t,g}} n_{c'|t,g}}$$

where $C_{t,g}$ is the set of media categories of all messages posted during the window ΔT . Both novelty metrics vary between 0 and $+\infty$.

Uncertainty: Both novelty-related metrics focus on the message posted at time t . The aggregation of either metric for all messages posted during the window of interaction, using entropy, captures the uncertainty driving user curiosity. The idea behind these metrics is that the curiosity of the user may

| Collative Variable | Metric | Description | Definition |
|--------------------|------------------------|---|---|
| Novelty | <i>userNovelty</i> | Novelty in user u participation wrt. others who posted during ΔT | $\begin{cases} -\log_2(\frac{n_{v t,g}}{\sum_v n_{v t,g}}), & \text{if } n_{u t,g} > 0 \\ -\log_2(1/ U_{t,g}), & \text{otherwise} \end{cases}$ |
| | <i>catNovelty</i> | Novelty in the category of m wrt. message categories posted during ΔT | $\begin{cases} -\log_2(P_{t,g}(C_m)), & \text{if } P_{t,g}(C_m) > 0 \\ -\log_2(1/ C_{t,g}), & \text{otherwise} \end{cases}$ |
| Uncertainty | <i>userUncertainty</i> | Variety (entropy) in user participation during ΔT | $-\sum_{v \in U_g} \frac{n_{v t,g}}{\sum_{v'} n_{v' t,g}} \log_2(\frac{n_{v t,g}}{\sum_{v'} n_{v' t,g}})$ |
| | <i>catUncertainty</i> | Variety (entropy) in media category usage in messages posted during ΔT | $-\sum_{c \in C_{t,g}} \frac{n_{c t,g}}{\sum_{c'} n_{c' t,g}} \log_2(\frac{n_{c t,g}}{\sum_{c'} n_{c' t,g}})$ |
| Conflict | <i>userConflict</i> | Diversity of distinct users posting messages during ΔT | $-\log_2(\frac{1}{ U_{t,g} })$ |
| | <i>catConflict</i> | Diversity of distinct media categories used during ΔT | $-\log_2(\frac{1}{ C_{t,g} })$ |
| Complexity | <i>catComplex</i> | Complexity wrt. media categories used during ΔT | $-\log_2(\frac{ C_{t,g} }{ M })$ |
| Social Influence | <i>maxDirInf</i> | Maximum direct influence to u from users who posted during ΔT | $\max_{o \in O_{d,t,g}} PMI_{t,g}^+(d, o)$ |
| | <i>maxIndInf</i> | Maximum indirect influence from users who posted during ΔT | $\max_{o' \in O_{d,t,g}} MI_{t,g}^+(D, o')$ |

Table 1: Metrics of curiosity stimulation computed at each posting event of a message m sent by u , considering a window of interaction ΔT preceding the message. Refer to section *Operational Definitions* for details on the terms in the equations.

be more/less stimulated by the greater/less variety in the users posting messages (and in the categories of these messages) during the window of interaction. Thus, we define metrics of uncertainty related to users and categories as detailed in Tab. 1. Both uncertainty metrics vary between 0 and 1.

Conflict: Conflict measures the diversity in the stimulus. The basic idea is that the distinct elements (categories/users) that appear in the window of interaction represent the different responses stimulating the curiosity of the user, and the strength of each response is captured by the probability of occurrence of each element (category/user). Thus, we define metrics of conflict related to users and categories to quantify the difference in terms of number of distinct users or distinct categories present in the window of interaction, as detailed in Table 1. Both conflicts vary between 0 and $+\infty$: more users (categories) lead to greater conflict, while a single user (category) leads to no conflict (i.e., conflict equal to 0).

Complexity: An alternative form of capturing the diversity of message categories in the current window of interaction is by exploiting the unique occurrences of the media types. This metric, also based on surprisal, captures how many distinct categories of messages are posted during the current window of interaction, out of the total number of categories $|M|$. Following the definition in Table 1, *catComplex* varies between 0 (when all categories are present in the time window) to $\log_2(|M|)$ (when only a single category is present). We note that complexity and conflict are functions of each other: $catComplex(t, g) = \log_2(\frac{1}{|C_{t,g}|}) + \log_2(|M|) = -catConflict(t, g) + \log_2(|M|)$.

Social Influence : Social Influence encompasses two types of stimuli - direct and indirect. *Direct influence* measures the impact from peers, while *indirect influence* refers to the impact from influencers or activists, i.e., users who often

post content and sustain discussions by encouraging broader group participation rather than influencing any single user. Thus, we use metrics for both types of stimuli. To distinguish between the user whose curiosity we are analyzing, i.e., user u who posted a message at time t in group g and the other users who may stimulate u 's curiosity through social influence, we refer to the former as the destination d and the latter as origins o of social influence.

We define:

- (i) Probability of observing an origin o on group g : $P_{t,g}(O = o)$;
- (ii) Probability of observing a destination d on group g : $P_{t,g}(D = d)$;
- (iii) Probability of observing a destination d given an origin o : $P_{t,g}(D = d | O = o)$;
- (iv) Joint probability of observing a destination d and an origin o : $P_{t,g}(D = d, O = o)$.

All these probabilities are conditioned on the group g and timestamp t of the current message and can be estimated given the message sequence history up to t .

The *direct* social influence metric, defined in Tab. 1, is based on the concept of Pointwise Mutual Information (*PMI*), which is computed for a particular destination d and origin o as:

$$PMI_{t,g}(d, o) = \log_2 \left(\frac{P_{t,g}(d|o)}{P_{t,g}(d)} \right) \text{ if } P_{t,g}(d) > 0$$

or 0 otherwise. The *PMI* of d and o measures the reduction in the uncertainty of destination d posting a message due to the knowledge that o also posted a message within the window of interaction. It means that, based on historical patterns, the behavior of user d in the group is influenced by the recent behavior of user o . Note that *PMI* is negative if d and o occur less frequently than would be expected under the assumption of independent behavior, then we clip negative values to 0, i.e. $PMI_{t,g}^+(d, o) = \max(PMI_{t,g}(d, o), 0)$

Since the current interaction window may have multiple origins stimulating curiosity destination d , we compute the maximum *direct* social influence for all origins o on d as detailed in Tab. 1, where $O_{d,t,g}$ is the set of origins o on d at time t on group g for d . In addition to direct influence, we argue that some users may have a natural ability to influence others. Users who often post content and, in doing so, keep the discussion going by driving others to also post content. We consider that, even in the absence of prior experiences such influencers may still stimulate the curiosity of a user (e.g., a newcomer) in the group. We identify these *indirect influencers* by searching for origins o that tend to have a strong influence towards any destination $d' \in D$. Given this rationale, we define a metric based on the mutual information of all the destinations conditioned on a particular origin, that is:

$$MI_{t,g}(D, o) = \sum_{d' \in U_g} P_{t,g}(d', o) PMI_{t,g}(d', o)$$

Again, we clip negative values of mutual information at 0 (i.e., MI^+) as a reflection of no social influence from origin o on any destination. Smaller values of MI^+ suggest weaker and no clear influence of o over users. We also aggregate the metric for all origins in the window ΔT by taking the maximum, as detailed in Tab. 1.


Both social influence metrics vary between 0 and $+\infty$. Note that Sousa, Almeida, and Figueiredo (2022) proposed two variants of each social influence metric, based on the maximum and the average across all origins. Yet, they found both variants to be strongly correlated with each other, keeping only the maximum in their study. As such, we also choose to use only the maximum.

Telegram Groups Dataset


Our study relies on a dataset collected by Perlo et al. (2025). The authors used *TGStat*, an open catalogue of worldwide Telegram public groups⁶, to discover popular groups. Using the Telethon Python package (Exo 2017), the authors collected data from groups across 10 topics: *Education*, *Bookmaking*, *Cryptocurrencies*, *Technologies*, *Darknet*, *Software and apps*, *Video and Films*, *Politics*, *Erotic* and *Linguistics*. Group topics were labelled by TGStat and manually validated by the authors. In Appendix A, Table 3 offers a qualitative description of the topics. We use a subset of the dataset that covers the whole month of March 2024, with around 22 million messages posted by more than 900 thousand users. The authors kindly shared the data with us, which consists of the ordered sequence of messages for each group with anonymized user identifiers.

Although we cannot tell when each user became a member of the group, we can assess when they first appeared in our data. We found that the average rate of first-time user appearances reaches 10% on the first day before declining and stabilizing at roughly 2.5% (see Figure 9 in Appendix A). Considering that our analysis focuses on active users (see below), we find that the average time between the first and last messages of these active users in our dataset


⁶<https://tgstat.com>.



(a) Distribution of the number of users (900k) who sent at least one message in the period under consideration among the 409 analysed groups .



(b) Distribution of the number of users per group for which we compute stimulus metrics (29k users).



(c) Distribution of the number of messages (6 million events) per user for which we compute stimulus metrics.

Figure 1: Overview of analyzed Telegram dataset.

is more than 15 days. These observations suggest that, although the context surrounding the analyzed groups changes dynamically over time, the behavior of involved users does not change abruptly. Our focus is on users who engage consistently in group discussions over time, rather than those who exhibit a burst of activity and then disengage.

As shown in Figure 1a, there is great diversity across the groups in terms of number of active users (i.e., users who posted at least one message). This diversity is also clear across the various topics. For example, groups in *Linguistics* tend to have fewer active users, whereas groups in *Education* and *Cryptocurrencies* tend to have larger numbers of users actively posting content. In particular, note the presence of many groups (35%) with at least 1,000 active members, which is much larger than the limit of 250 simultaneous members (not necessarily actively posting content) allowed




Figure 2: Correlation among curiosity metrics for all messages.

by WhatsApp when the dataset analyzed by Sousa, Almeida, and Figueiredo (2022) was gathered. In fact, we note that the median number of active users in that dataset⁷, computed during the month of peak activity (October 2018), was only 34 users (as indicated by the dashed black vertical line in Fig. 1a), whereas here, groups have a median of 493 active users (red vertical line). Such a large difference underscores the importance of studying curiosity on Telegram, given the novelty of its larger and more diverse social spaces.

To mitigate data sparsity issues, we follow Sousa, Almeida, and Figueiredo (2022) and compute the metrics only for users who posted at least 30 messages over the entire period and only for posting events with at least 10 messages within the *window of interaction*. Messages from service bots are considered in the computation of metrics, as they are inherently part of the conversation. Yet, the final analysis focused only on metrics computed for users.

After the filtering, we end up with roughly 6 million messages. 29,196 users were active in 409 distinct groups. As shown in Figs. 1b and 1c, there is great heterogeneity in terms of the number of users per group and the number of messages per user for which we can compute the metrics. Such diversity reflects the wide range of activity levels among users. However, the coverage across topics is fairly homogeneous as the number of groups per topic ranges from 25 for Erotic to 60 for Education.

Analysis of Curiosity Stimulation

Metric Selection

Recall that we have 9 metrics to capture components (collative variables) of curiosity stimulation. As initial step, we assess whether these metrics are correlated one to each other. To that end, we first compute all 9 metrics, standardize them according to a Z-score normalization, then compute the correlation matrix among the features (see Fig. 2). Our findings indicate that direct and indirect social influence metrics capture distinct aspects of social-related curiosity, and both metrics have weak correlations with the other metrics. This

⁷Publicly available at <https://doi.org/10.5281/zenodo.5790153>

shows that these metrics provide complementary dimensions of curiosity stimulation compared to the traditional collative variables. Among the other metrics, we find that: (i) *userConflict* and *userUncertainty* are strongly correlated (Pearson’s correlation ρ above 0.8), and (ii) *catConflict* and *catComplexity* has a perfect inverse correlation ($\rho = -1$), as expected from their formulation (see previous section). We then keep only one of these strongly correlated pairs of metrics. We end up with 7 complementary metrics, namely: *userNovelty*, *catNovelty*, *userUncertainty*, *catUncertainty*, *catComplex*, *maxIndInf* and *maxDirInf*.

These observations align with previous findings for WhatsApp groups (Sousa, Almeida, and Figueiredo 2022). However, our study differs in a key methodological aspect: we consider all seven metrics to infer curiosity profiles in Telegram groups, while prior work used only two social influence metrics to infer curiosity profiles in WhatsApp groups. Our broader approach allows for a more comprehensive analysis, revealing richer and more diverse profiles, as discussed next.

Message-Level Curiosity Stimulation

We start by investigating the curiosity stimuli that more often drive users to post a new message. To that end, we infer profiles by clustering messages (independently from users and groups), according to the 7 identified metrics. Given the large dataset size (6 million messages), we employ an iterative stochastic clustering approach. We use MiniBatch K-means (Sculley 2010), running 10 iterations for each tested number of clusters (k). We define the best number of clusters by identifying the elbow point in the average inertia curve, which measures the sum of squared distances of samples to their nearest cluster centroid. Following this method, we identify 7 clusters, each reflecting a different profile of curiosity stimulation.

We proceed to characterize the identified clusters with respect to any distinctive behavior in terms of metric values. To that end, we adopt a combined approach by comparing boxplots of metric distributions within each cluster and analyzing the SHAP (SHapley Additive exPlanations) values derived from a decision tree classifier (Lundberg et al. 2020). We omit the plots here due to space constraints (deferring them to Figs. 12 and 13 in Appendix A). In turn, Tab. 2 offers a high-level description of each cluster through the median values of all metrics, highlighting those that are higher with statistical significance. We label each cluster (*Description* column of the table) according to those distinctive patterns.

Clusters 0 and 1 refer to messages whose postings were mostly stimulated by (indirect and direct) social influence. They account for roughly 28% of all messages in our dataset. Clusters 3 and 5 consist of messages where the traditional collative variables, notably novelty and uncertainty, play a stronger role in curiosity stimulation. Clusters 2 and 6 have mixed stimuli. Given their similarities in terms of metric distributions (and having the closest centroids), we choose to merge them into a single cluster (Cluster 2, mixed stimuli) to enhance interpretability. Finally, cluster 4 consists of messages for which no clear evidence of curiosity stimulation was found (i.e., curiosity independent).

Here we compare our findings with those of Sousa,

| ID | Description | <i>userNov</i> | <i>catNov</i> | <i>userUncert</i> | <i>catUncert</i> | <i>catComplex</i> | <i>maxDirInf</i> | <i>maxIndInf</i> | % Messages |
|----|---------------------|----------------|---------------|-------------------|------------------|-------------------|------------------|------------------|------------|
| 0 | Indirect Influence | -0.82 | -0.56 | -0.91 | -0.92 | 1.30 | -0.64 | 0.14* | 14.8% |
| 1 | Direct Influence | -0.23 | -0.47 | -0.06 | -0.48 | -0.08 | 1.47* | -0.08 | 13.2% |
| 2 | Mixed stimuli | 0.61 | -0.40 | 0.43 | -0.18 | -0.08 | -0.13 | -0.34 | 19.8% |
| 3 | Categorical Novelty | -0.15 | 2.93* | -0.29 | 1.28 | -0.08 | -0.16 | -0.25 | 6.8% |
| 4 | Independent | -0.83 | -0.35 | -0.58 | -0.04 | -0.08 | -0.54 | -0.25 | 21.7% |
| 5 | Uncertainty | 0.82* | 0.98 | 1.21* | 1.57* | -1.15 | 0.21 | -0.45 | 11.3% |
| 6 | Mixed stimuli | 0.45 | -0.56 | 0.55 | -0.92 | 1.30* | -0.12 | -0.36 | 12.5% |

Table 2: Median of z-score normalized curiosity stimulus metrics for message clusters. An asterisk (*) indicates when a metric distribution is significantly higher than those in each other cluster (Mann-Whitney test, $p < 0.01$).




Figure 3: Distributions of Euclidean distances in the curiosity stimulus space between the messages sent by a user and (i) the barycenter of all messages by the same user (blue), (ii) the centroid of each message’s profile (orange) and (iii) the barycenter of messages sent by other users in the last 15 minutes (green).

Almeida, and Figueiredo (2022) that analyzed political groups in Whatsapp. For this, we restrict our analysis to Politics groups on Telegram. Results reveal striking differences: We find six significant clusters for Telegram while only three emerged in Whatsapp. On Telegram, Direct and indirect social influences dominate only 2.6% and 8% of the messages, versus 13% and 14.4% on WhatsApp. Messages with no clear curiosity stimulation make up only 21% in our data, while those where 73% in WhatsApp. These differences stem from the focus of the prior work restricted to social influence. By incorporating all curiosity metrics, we reveal other profiles that are also important for Politics groups.

Takeaway: (i) Messages are grouped into 6 clusters in the curiosity stimulus space, identifying more diverse and complex curiosity profiles, compared to prior analysis of WhatsApp. (ii) Clusters 0 and 1, characterized primarily by (indirect and direct) social influence, and Cluster 4, with no clear evidence of curiosity stimulation, resemble findings from the WhatsApp study, but with a much lower presence, notably in Politics groups. (iii) Roughly 50% of the messages in all analyzed groups are dominated by other curiosity stimulation factors, highlighting the broader and more nuanced role of curiosity stimuli in influencing user engagement.

User-Level Curiosity Stimulation

After characterizing curiosity stimulation at the message level, we shift our focus on studying it at the user level. To




Figure 4: Distribution of the probability p_{trans} that the next message from the same user in a group belongs to a different curiosity stimulus cluster, conditioned on whether the time interval δ_t between them exceeds a certain threshold.

that end, we first analyze the distribution of messages sent by the same user within the 7-dimension curiosity stimulus space. In Fig.3 we compare the distributions of the Euclidean distances of messages within this space. We observe that the distances between a user’s messages and the user’s barycenter (i.e. the average of all messages sent by the same user) in this space are statistically smaller (Mann-Whitney test, $p < 0.01$) than the distances between those messages and the centroids of their respective clusters (i.e., profiles identified in the previous section). This observation suggests that the messages sent by the same user tend to be concentrated in a localized area of the stimulus space, implying that their curiosity is consistently triggered by similar types of stimuli.

We also observe that the nature of curiosity stimuli seems to be specific to each user since messages posted by different users within the same group and in close temporal proximity can exhibit markedly different stimulation profiles. This is illustrated by the green distribution in Fig. 3, which consists of distances between a user’s message and the barycenter of messages sent by other users within the preceding 15 minutes. It is important to note that the choice of this time period is not arbitrary, but is directly tied to the selection of the window of interaction, ensuring that at least 50% of the windows overlap. This means that users are influenced by a common context for at least half of their window of interaction. Thus, the messages offering stimuli to them may be quite similar. Yet, the green distribution is biased towards larger distances, i.e., those users are far apart in the curiosity stimulation space, compared to the current user under analysis.




Figure 5: Average fraction of messages sent in each stimulus curiosity cluster from users belonging to a group of the topic (blue solid line). The dotted grey line is the average distribution over the entire population. Dashed green/red lines highlight that the distribution in the topic population is statistically higher/lower than the general one.

Despite such a concentration of messages from the same user in the curiosity stimulation space, we also observe some changes with time. Specifically, we measure the probability that two consecutive messages posted by the same user belong to different curiosity stimulation profiles, conditioned on whether the time elapsed between their postings exceeds (or not) a given threshold. To ensure robust results, we only estimate this probability for samples (users) with at least 10 messages. Fig. 4 illustrates, for two such thresholds, that the probability of transitioning into a different curiosity profile increases with the time interval between them. These results suggest that, while curiosity-driven behaviour tends to be mostly stable, it can also exhibit occasional shifts in response to changing stimuli over time.

In general, we observe that some changes in curiosity profile are more recurrent than others. For example, when in clusters 0 and 1, for which curiosity is mostly stimulated by indirect and direct influence, respectively, users more often change to cluster 2, with mixed stimuli (roughly 40% and 52% of the time, respectively). For users in cluster 2, in turn, the most common transition is back to cluster 1 (35%). Yet, often such transitions are only temporary fluctuations, as 23% of transitions revert to the previous cluster in the next message, and 11% revert after two messages. That is, despite some occasional changes, each user can be fairly well characterized by one curiosity stimulation profile that dominates his behaviour across most of the messages posted during the period of analysis⁸. Indeed, when analyzing the messages sent by a user collectively, for 76.6% of the users, the majority of their messages (more than 50%) fall within a single curiosity profile.

Next, we analyze deviations of user curiosity stimulation across different topics, considering the 10 group topics available in our dataset. To that end, we take only user messages posted in groups of the same topic and compute the average (across all users) fractions of messages in each curiosity profile. We do so for groups in each topic and compare the results against those of the overall population, so as to infer patterns that are specific to each topic. Fig. 5 shows radar plots with the average fractions of messages (blue) in each curiosity profile (axes) for three selected

topics, namely *Politics*, *Cryptocurrencies* and *Linguistics*. Results for the overall population are shown in grey. Plots for the other topics are presented in Appendix A.

We observe that users in groups discussing *Politics* (Fig. 5a) show a preference for messages driven mostly by uncertainty (cluster 5), revealing a tendency for peer discussions with a broad and balanced audience and using different media types (categories). This contrasts with prior work on WhatsApp, where this aspect was overlooked due to a focus solely on social influence. The same was also observed for groups in *Technologies*. In contrast, *Cryptocurrencies* (Fig. 5b) (as well as *Bookmaking*), which likely involve trust-building with individuals who have previously provided successful advice, exhibit a significantly higher proportion of messages mostly influenced by direct social influence (cluster 1). In the *Linguistics* groups (Fig. 5c), in turn, there is a marked increase in messages in cluster 0, which are strongly stimulated by indirect social influence. These results align with our manual observation that discussions in such groups are often led by language teachers who encourage conversation by posing questions and introducing subjects for debate.

Finally, we consider whether user curiosity varies based on group membership and topic. Our dataset provides limited opportunities to explore this phenomenon, as only 2% (596) of the users are active in multiple groups, limiting the generalizability of these findings. Nevertheless, focusing on those users active in multiple groups, we observe that a user’s dominant curiosity profile differs 52% of the time between groups within the same topic. This percentage increases to 62% when comparing groups in different topics. These results hint that a user’s curiosity depends on the specific group they participate in, with topic diversity playing a role in shaping how their curiosity is stimulated.

Takeaway: (i) Users’ messages focus tightly around their barycenter, showing consistent curiosity-driven behavior. (ii) Different users in the same group and timeframe may exhibit distinct curiosity profiles, emphasizing the individual nature of curiosity, even within shared environments. (iii) While users typically follow one curiosity profile, transitions occur over time, though many are temporary. (iv) We are the first to identify that curiosity stimulation patterns vary by topic — e.g., while uncertainty is a dominant factor driving user curiosity in *Politics* groups, direct social influence is a much

⁸Such observation may change for longer periods of analysis.




Figure 6: OLS coefficients modeling user engagement (message count) based on average curiosity metrics, number of active users and average number of messages by others in the group, for groups in various topics. Grey cells mark non-significant coefficients ($p > 0.01$); black cells mark metrics excluded due to multicollinearity.

stronger factor in *Cryptocurrencies*.

Active Engagement and Influence

To address our second research question, we study how the curiosity profile of a user relates to user engagement in group discussions. To the best of our knowledge, this analysis is the first of its kind in the context of instant messaging social networks, offering new insights into the role of curiosity in information spread processes.

We start by representing the curiosity profile through the barycenter of the user’s messages in the curiosity stimulus space, and we capture the user engagement by the number of messages sent by the user. Then, we fit an OLS (Ordinary Least Squares) regression for each topic, with the number of messages sent as the response variable. As predictor variables, we use the barycenter of the user’s messages in terms of curiosity metrics along with two control variables: the number of active users in the group and the average number of messages sent by other users in the group. These control variables account for the heterogeneity in group activity levels. For each topic, we log-transform the control features and the response variable as their distributions span multiple orders of magnitude and standardize all values according to a Z-score normalization. Variables with a Variance Inflation Factor (VIF) greater than 10 are removed to avoid multicollinearity, as recommended by Neter, Wasserman, and Kutner (1983). We opt to exclude the *Software and Applications* topic due to the poor fitting (adjusted $R^2 < 0.2$), focusing our analysis of the significant OLS coefficients on the remaining 9 topics.

The heatmap in Fig 6 illustrates the OLS coefficients between the predictors (curiosity stimulation metrics and control variables) and the number of messages sent by users across different topics. Positive values (red) suggest that higher values of certain features are associated with more frequent active engagement by the user, while negative values (blue) suggest the opposite relationship. Metrics for which the

relationship is not statistically significant (grey) or that were removed due to multicollinearity (black) are also highlighted.

We observe that both *userNovelty* and *maxDirInf* consistently show a negative correlation with user activity across all topics. This suggests a general pattern where users who are typically motivated by the novelty of their contributions, or by interactions with peers they frequently engage with, tend to be less actively engaged in the conversation. In contrast, *userUncertainty* exhibits a consistent positive correlation across topics, particularly in *Politics* and *Darknet*, indicating that users who prefer conversations with varied and balanced participation may be involved more frequently. The other metrics show mixed effects across topics indicating the variability of certain relations by the surrounding context.

We also note the strong effect of group activity levels on user’s propensity to participate. The average number of messages sent by other users in the group shows a strong positive correlation with user engagement, suggesting that higher overall activity may encourage users to contribute more often, possibly due to the stimulating effect of a more dynamic environment. Similarly, the number of active users in the group also tends to have some positive influence on individual participation, across topics: as more users actively engage in the group, individual users may perceive their contributions as part of a broader and more meaningful discussion.

As a complementary analysis of the relationship between a user’s curiosity stimulation profile and the activity she engages in the groups, we study their influence on other users in the chat, in an orthogonal approach to the proposed metrics, by using graphs. More specifically, to represent the influence dynamics in group conversations, we build a *Influenced-by* network for each group as follows. We keep a 30-minute time window ΔT consistently with the window of interaction. For each message sent by user i at time t , we draw direct edges from i to all other users j (different from i) who posted messages in the period $[t - \Delta T, t]$, that is, users who potentially influenced i . The weight of each edge (i, j) is given by the fraction of messages sent by user j out of the total messages sent during the time window. To avoid overestimating the influence of a single user who has sent multiple consecutive messages with no interleaving messages from others (*chain of messages*), we count all messages in a *chain* as a single one. Thus, for each message sent by i , the total influence from other users adds up to 1. Note that we do not add any edge if no other user sent a message within $[t - \Delta T, t]$. We illustrate the creation of an *Influence-by* network with a toy example in Fig. 7. Finally, we merge all edges with the same source i and target j into one, with a weight equal to the sum of their weights. Note that while an edge from i to j indicates that i is influenced by j , it also implies that users who are more central in the network tend to exert more influence over others in the group. We build one network for each group but focus our analyses on groups with at least 100 users in their *Influence-by* networks (346 groups), to ensure robust computation of graph metrics. For each such network, we compute the PageRank centrality scores of its nodes (i.e., users).

Recall that, aside from temporary fluctuations, the majority of a user’s messages fall within the same cluster. This highlights the predominance of a single curiosity

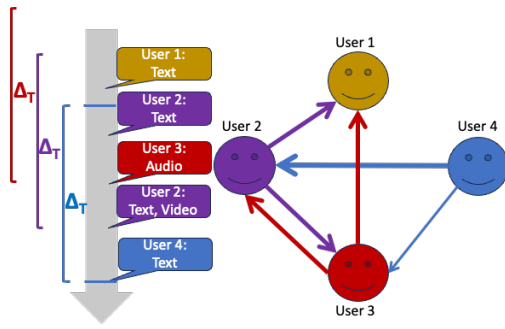


Figure 7: A toy example of an *Influenced-by* network. *User 4* (blue) posts a message at time t . Within the preceding time window $[t - \Delta T, t]$, *User 2* (purple) sent two messages, and *User 3* (red), one. For the message posted by *User 4*, we create an edge from *User 4* to *User 2* with weight of $\frac{2}{3}$, and an edge from *User 4* to *User 3* with weight $\frac{1}{3}$. Similarly, for the second message sent by *User 2*, we add two edges with weights $\frac{1}{2}$ to *User 1* and *User 3* (no self-loops).

stimulation profile over all others (see Fig. 11 in Appendix A), which accounts for the 63% of the explained variance in the distribution of the user’s messages across the different curiosity stimulus profiles. Accordingly, we take this profile as representative of the user’s curiosity stimulation and analyze the distribution of PageRank scores across users within each profile. These distributions are shown in Fig. 8.

As observed, users who mainly post messages belonging to the profiles strongly driven by *direct social influence* (cluster 1) and *uncertainty* (cluster 5) tend to occupy more peripheral positions within the *Influenced-by* networks. This could be attributed to the fact that users with high *direct social influence* tend to engage less actively, generating fewer messages that drive ongoing discussions, thus reducing their centrality within the network. This observation is indeed consistent with the results on the OLS coefficients, shown in Fig. 6. On the other hand, users characterized by high *uncertainty* may see their influence diluted within the broader conversation, making them less likely to be pivotal in shaping the direction of the discourse. In contrast, users often driven by indirect influence (cluster 0), tend to occupy more central positions in the network, suggesting that these users participate more often, also influencing others. Indeed, as shown in Fig. 6, this relationship can be observed, to some degree, for groups in some (but not all) topics, notably Darknet, Erotic and Linguistics. Finally, it is noteworthy that users whose majority of messages are not influenced by any curiosity stimuli (i.e., independent, cluster 4) are often the most central in the *Influenced-by* networks. It is plausible to infer that such users may serve as initiators of conversations.

Takeaway: *This study is the first to reveal key relationships between curiosity and active engagement in group discussions, notably: (i) Users driven by novelty or recurrent peer interaction (direct social influence) tend to be less participative. (ii) Those preferring balanced participation or indirect influence are more active and central in group interactions. (iii) High overall group activity consistently encourages individual participation. (iv) Independent users often occupy*




Figure 8: Distribution of node PageRank in the *Influenced-by* network, categorizing users based on the message curiosity-stimulus cluster where they hold the highest share. Percentages show how many users predominantly sent messages in each cluster.

central roles, potentially acting as conversation initiators.

Conclusions and Future Work

This study investigates how the intrinsic human trait of curiosity influences user active engagement, specifically content posting, in Telegram public groups. To our knowledge, this is the first research to focus on curiosity-driven engagement within Telegram.

By employing different metrics which quantify aspects of curiosity stimulation, we identified 6 distinct stimulus profiles. Two profiles are primarily characterized by direct or indirect social influence, while three others are dominated by factors such as novelty, uncertainty, or a combination of multiple stimuli. By contrast, the last profile (accounting for nearly 22% of the messages) is characterized by the absence of any of the other identified stimuli, suggesting that these messages may be driven by other factors not captured by our metrics, such as the specific content of the messages, or that they may not be influenced by curiosity at all. This underscores the complex social dynamics at play in online communities and highlights the need for further investigation into the forces driving user participation.

Interestingly, we found that social influence—both direct and indirect—is a dominant component of curiosity stimulation in 28% of all messages, particularly in groups related to topics such as *Bookmaking*, *Cryptocurrencies*, *Linguistics*, *Darknet*, and *Software & Application*. Yet, our results suggest that other forms of curiosity stimulation—such as novelty—also play a significant role in user participation. For example, interest in novel media types drives around 7% of messages, while mixed stimuli account for 32%, disputing previous studies that have only focused on social influence.

We also observed that as the time between messages increases, users are more likely to switch to a different curiosity stimulation profile. This supports previous findings (Sousa, Almeida, and Figueiredo 2019, 2022) that human curiosity is dynamic and influenced by changes in both the external environment and effective stimuli. However, despite these fluctuations, the overall curiosity stimulus profile that motivates a user to post content remains relatively stable

over time, considering the limited analyzed time frame. This allowed us to examine the dominant curiosity stimulation profiles at the user level across different group topics.

The analysis revealed a strong relationship between specific curiosity stimulus profiles and user engagement patterns. For example, users driven by uncertainty were notably more active, particularly in groups discussing diverse and dynamic topics such as *Politics*, where the variety of discussions encourages more frequent participation. In contrast, users influenced by direct social influence tended to post less often and typically occupied more peripheral roles within the group. Additionally, users who were not driven by any specific curiosity stimulus frequently emerged as conversation starters and central group figures. This finding emphasizes the importance of understanding not only individual curiosity profiles but also the broader social dynamics that influence participation and leadership within online communities.

Furthermore, the diverse discussion topics present in the dataset allowed us to identify distinct patterns in curiosity stimulation based on the subject of the discussion. As noted, social influence plays a significant role in some topics, while uncertainty and other factors drive engagement in areas such as *Politics* and *Technology*. Interestingly, in groups related to *Video & Films* and *Software & Applications*, the independent curiosity profile dominates, suggesting that messages in these groups may not be primarily driven by curiosity.

In conclusion, this research offers valuable insights into the dynamics of online communities by focusing on curiosity-driven engagement. It may enhance our understanding of the role of curiosity — a fundamental human behavioural trait linked to cognitive processes — in the context of information dissemination and spreading, aspects largely overlooked by prior works.

Future research will further explore how the content of specific messages triggers different curiosity types and how these triggers influence user engagement across various platforms. In this paper, we prioritized a breadth-oriented approach, analyzing a diverse range of topics, groups, and languages to mitigate potential biases associated with short-term observations. We believe that examining the long-term impact of different curiosity stimuli on sustained participation and community growth will help to better understand the broader implications of these dynamics on user involvement and the evolution of digital communities, enhancing the generalizability of the findings.

Limitations

- This study relies on data from public Telegram groups identified by TGStat. While we cannot guarantee that TGStat is bias-free, we note that it has been used in previous research (Tikhomirova and Makarov 2021; Urman and Katz 2022; Morgia et al. 2023), which supports its reliability as a data source. Our analysis is confined to the data collected from the selected Telegram groups, so we do not claim broad generalizability, such as to private or less active groups. Yet, the strength of our study lies in its large dataset, which includes over 6 million messages across 409 groups and spans 10 diverse topics—far exceeding the scope of prior work.

- Metrics used in our study operationalize cognitive concepts. Thus, validation is extremely hard, as it might involve a multi-disciplinary approach. Nevertheless, metric derivation followed the procedures of Sousa, Almeida, and Figueiredo (2022), which, in turn, followed the seminal works by Berlyne (1960); Silvia (2006), which proposed a methodology to quantify collative variables related to curiosity stimulation based on information theoretical metrics. Thus, despite the lack of direct validation, the used metrics are rooted in the foundational theories of curiosity stimulation.
- Our analysis of curiosity stimulation is focused on a subset of users and messages (i.e., users who shared at least 30 messages in the period, and messages with at least 10 other messages in the window of interaction). Whereas this filtering decision was made for the sake of robust metric computation, we argue that removed users can be considered too inactive to be studied with respect to curiosity stimulation. Similarly, the curiosity driven by the posting of messages after periods of low or no activity might be less dependent on these (few or none) previous messages.
- We deliberately choose not to analyze the message content (other than the type of media). We do recognize that content is likely to play a direct role in shaping user curiosity. Yet, our choice allows for a more universal comparison of curiosity metrics across different groups and even social network platforms, free from the influence of specific contexts such as the discussion topic or language. This strategy also ensures broader applicability, particularly when content is inaccessible due to privacy concerns, making the method both flexible and ethically sound.

Acknowledgments

This work has been partially supported by the Spoke 1 “FutureHPC & BigData” of ICSC — Centro Nazionale di Ricerca in High-Performance-Computing, Big Data and Quantum Computing, funded by the European Union — NextGenerationEU. It has also been partially supported by *Conselho Nacional de Desenvolvimento Científico e Tecnológico* (CNPq) and *Fundação de Amparo à Pesquisa do Estado de Minas Gerais* (FAPEMIG), both in Brazil.

References

- Aldous, K. K.; An, J.; and Jansen, B. J. 2019. View, like, comment, post: Analyzing user engagement by topic at 4 levels across 5 social media platforms for 53 news organizations. In *ICWSM*, volume 13, 47–57.
- Baumgartner, J.; Zannettou, S.; Squire, M.; and Blackburn, J. 2020. The Pushshift Telegram Dataset. *ICWSM*, 14: 840–847.
- Berlyne, D. E. 1960. *Conflict, arousal, and curiosity*. McGraw-Hill Book Company.
- Caetano, J. A.; Almeida, J.; Gonçalves, M.; Meira Jr, W.; Marques-Neto, H. T.; and Almeida, V. 2021. Analyzing topic attention in online small groups. In *ASONAM*, 64–68.

- Deng, Y.; Ding, Q.; Wu, X.; and Cai, Y. 2023. Curiosity Enhanced Bayesian Personalized Ranking for Recommender Systems. In *NeurIPS*, 312–324. Springer.
- Exo, L. 2017. Telethon: Async Python 3 Telegram client API. <https://docs.telethon.dev/>. Version 1.28, Accessed: 2024-09-15.
- Fu, Z.; and Niu, X. 2023. Modeling Users’ Curiosity in Recommender Systems. *ACM TKDD*, 18(1): 1–23.
- Hashemi, A.; and Zare Chahooki, M. A. 2019. Telegram group quality measurement by user behavior analysis. *Social Network Analysis and Mining*, 9(1): 33.
- Hoseini, M.; de Freitas Melo, P.; Benevenuto, F.; Feldmann, A.; and Zannettou, S. 2024. Characterizing Information Propagation in Fringe Communities on Telegram. In *ICWSM*, volume 18, 583–595.
- Kashdan, T. B.; Stiksma, M. C.; Disabato, D. J.; McKnight, P. E.; Bekier, J.; Kaji, J.; and Lazarus, R. 2018. The five-dimensional curiosity scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people. *Journal of Research in Personality*, 73: 130–149.
- Kidd, C.; and Hayden, B. Y. 2015. The psychology and neuroscience of curiosity. *Neuron*, 88(3): 449–460.
- Kloo, I.; Cruickshank, I. J.; and Carley, K. M. 2024. A Cross-Platform Topic Analysis of the Nazi Narrative on Twitter and Telegram During the 2022 Russian Invasion of Ukraine. In *ICWSM*, volume 18, 839–850.
- Loewenstein, G. 1994. The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, 116(1): 75.
- Lundberg, S. M.; Erion, G.; Chen, H.; DeGrave, A.; Prutkin, J. M.; Nair, B.; Katz, R.; Himmelfarb, J.; Bansal, N.; and Lee, S.-I. 2020. From local explanations to global understanding with explainable AI for trees. *Nature Mach. Intell.*, 2(1): 2522–5839.
- Morgia, M. L.; Mei, A.; Mongardini, A. M.; and Wu, J. 2023. It’s a Trap! Detection and Analysis of Fake Channels on Telegram. In *IEEE Conf. on Web Services*, 97–104.
- Neter, J.; Wasserman, W.; and Kutner, M. H. 1983. *Applied linear regression models*. Richard D. Irwin.
- Perlo, A.; Paoletti, G.; Jha, N.; Vassio, L.; Almeida, J.; and Mellia, M. 2025. Topic-wise Exploration of the Telegram Group-verse. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion ’25)*, 10. ACM.
- Sculley, D. 2010. Web-scale k-means clustering. In *WWW Conf.*, 1177–1178.
- Shivaram, K.; Bilgic, M.; Shapiro, M.; and Culotta, A. 2024. Forecasting Political News Engagement on Social Media. In *ICWSM*, volume 18, 1451–1462.
- Silvia, P. J. 2006. *Exploring the psychology of interest*. oxford university Press.
- Sousa, A. M.; Almeida, J. M.; and Figueiredo, F. 2019. Analyzing and modeling user curiosity in online content consumption: a LastFM case study. In *ASONAM*, 426–431.
- Sousa, A. M.; Almeida, J. M.; and Figueiredo, F. 2022. Metrics of social curiosity: The WhatsApp case. *Online Social Networks and Media*, 29: 100200.
- Tang, M.-C.; and Liao, I.-H. 2022. Preference diversity and openness to novelty: Scales construction from the perspective of movie recommendation. *JASIST*, 73(9): 1222–1235.
- Tikhomirova, K.; and Makarov, I. 2021. Community detection based on the nodes role in a network: The telegram platform case. In *Analysis of Images, Social Networks and Texts*, 294–302. Springer.
- Timme, N. M.; and Lapish, C. 2018. A tutorial for information theory in neuroscience. *eneuro*, 5(3).
- Urman, A.; and Katz, S. 2022. What they do in the shadows: examining the far-right networks on Telegram. *Information, Communication & Society*, 25(7): 904–923.
- Vassio, L.; Garetto, M.; Chiasserini, C.; and Leonardi, E. 2021. Temporal dynamics of posts and user engagement of influencers on Facebook and Instagram. In *ASONAM*, 129–133.
- Vassio, L.; Garetto, M.; Leonardi, E.; and Chiasserini, C. F. 2022. Mining and modelling temporal dynamics of followers’ engagement on online social networks. *Social Network Analysis and Mining*, 12(1): 96.
- Venâncio, O. R.; G. Ferreira, C. H.; Almeida, J. M.; and C. da Silva, A. P. 2024. Unraveling User Coordination on Telegram: A Comprehensive Analysis of Political Mobilization during the 2022 Brazilian Presidential Election. In *ICWSM*.
- Ver Steeg, G.; and Galstyan, A. 2013. Information-theoretic measures of influence based on content dynamics. In *Proceedings of the sixth ACM international conference on Web search and data mining*, 3–12.
- Wang, N.; and Chen, L. 2023. How do item features and user characteristics affect users’ perceptions of recommendation serendipity? A cross-domain analysis. *User Modeling and User-Adapted Interaction*, 33(3): 727–765.
- Weerasinghe, J.; Flanigan, B.; Stein, A.; McCoy, D.; and Greenstadt, R. 2020. The Pod People: Understanding Manipulation of Social Media Popularity via Reciprocity Abuse. In *WWW Conf.*, 1874–1884.
- Xu, J.; and Livshits, B. 2019. The anatomy of a cryptocurrency pump-and-dump scheme. In *USENIX*, 1609–1625.
- Xu, K.; Mo, J.; Cai, Y.; and Min, H. 2019. Enhancing recommender systems with a stimulus-evoked curiosity mechanism. *IEEE TKDE*, 33(6): 2437–2451.
- Yan, D.; and Chen, L. 2023. The influence of personality traits on user interaction with recommendation interfaces. *ACM THIS*, 13(1): 1–39.
- Zhao, P.; and Lee, D. L. 2016. How much novelty is relevant? it depends on your curiosity. In *SIGIR*, 315–324.

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, see the Analysis of Curiosity Stimulation.**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **NA**
- (e) Did you describe the limitations of your work? **Yes, see the Limitations.**
- (f) Did you discuss any potential negative societal impacts of your work? **NA**
- (g) Did you discuss any potential misuse of your work? **NA**
- (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? **NA**
- (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? **Yes**
- (b) Have you provided justifications for all theoretical results? **Yes**
- (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? **Yes, see the Limitations**
- (f) Have you related your theoretical results to the existing literature in social science? **Yes, partially. Although curiosity has never been studied on Telegram before, we compared our results with those obtained in Whatsapp.**
- (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **Yes**
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**
- (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**
- (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)? **No**
- (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **NA**
- (f) Do you discuss what is “the cost” of misclassification and fault (in)tolerance? **NA**
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? **Yes**
- (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? **No**
- (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **No, the code used to compute the curiosity metrics is public and the dataset of Telegram messages was kindly shared with us by the authors who are fully aware of its use for the purposes of this paper only**
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes, see Ethical Statement.**
- (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? **NA**
- (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **NA**
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? **NA**
- (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **NA**
- (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 deidentified? **NA**

Ethical Statement

The data in this paper was collected and posted by Perlo et al. (2025). The authors state that it was gathered using the Telethon API, which in turn utilizes the official Telegram API. The dataset includes messages from public groups from which the collectors have not been removed and where the admins have not set a time limit for viewing such messages.

Appendix A - Supplementary materials

To complement the main findings presented in this paper, additional materials have been included in this appendix. These supplementary materials provide detailed explanations,




Figure 9: Percentage of users actively engaging for the first time by sending a message during the study period, shown with average values over groups and 95% confidence intervals.

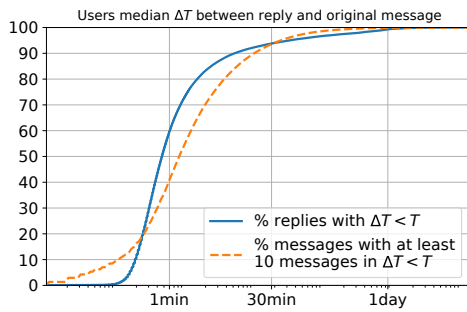


Figure 10: Cumulative distributions of user median reply times (solid blue) and the proportion of messages with at least 10 preceding messages in the same ΔT window (dashed orange). The intersection at 30 min. represents a sweet spot, balancing the trade-off between filtering noise from the user’s attention period and maintaining sufficient data coverage.

extended analyses, and supporting results that further substantiate the claims made in the study. They aim to offer deeper insights and transparency, ensuring the reproducibility of the experiments and facilitating future research in the field.

Table 3 offers a qualitative description of the topics, as provided by Perlo et al. (2024), and we included it here for completeness.

Figure 9 illustrates the percentage of users actively engaging for the first time by sending a message during the study period. Although we lack information regarding the group’s activity duration or the age of its members, observing the timing of users’ first activation reveals an initial peak on the first day, averaging around 10%, which then stabilizes at approximately 2.5%.

To address the choice of the interaction time window, we analyzed user reply times and message density within varying timeframes. Using reply times as a proxy for user attention spans, we observed that over 90% of users reply within 30 minutes. Additionally, we examined how the proportion of messages with at least 10 preceding messages changes with different time windows. The choice of a 30-minute window




Figure 11: Share of messages in the dominant curiosity profile (where the user posted most frequently) compared to the share in the second-most active curiosity profile.

balances the trade-off between including too much noise (with longer windows) and filtering out excessive data (with shorter windows). This intersection of trends is illustrated in Figure 10, highlighting 30 minutes as the ideal window length for our analysis.

To enhance explainability, we present both the distribution of features for each cluster and an evaluation of feature importance. Figure 12 illustrates the feature distributions (normalized Z-score) for the identified message clusters, providing insights into the distinguishing characteristics of each cluster. Additionally, we assess feature importance using SHAP values derived from a shallow decision tree. The SHAP beeswarm plot in Figure 13 ranks features by their mean absolute SHAP values, highlighting those with consistent contributions to the model while deprioritizing rare, high-magnitude effects.

Figure 14 depicts the average fraction of messages sent by users in each curiosity stimulus profile for groups within a specific topic (solid blue line). The dotted grey line represents the average distribution across the entire population, providing a baseline for comparison. Dashed green and red lines indicate categories where the topic-specific distribution is statistically higher or lower, respectively, than the general population average. This visualization highlights distinctive curiosity patterns within topic-focused groups compared to the broader population.

Figure 11 highlights the dominance of the most frequently used curiosity profile over the second-most active profile and, consequently, over all other profiles.

| Topic | Description |
|--------------------------|---|
| Education | Discussion about college and university courses and exams. |
| Bookmaking | Discussion about online betting and similar topics. |
| Cryptocurrencies | Discussion about cryptocurrencies, market stock and similar topics. Some groups offer official support for crypto exchanges. |
| Technologies | Discussions about consumer electronics, mostly smartphones. Some groups are second-hand marketplaces for consumer electronics. |
| Darknet | Trading of mostly illegal content, i.e., credit card numbers, accounts of media platforms, etc. |
| Software and apps | Discussion about usage of software, mostly regarding Android modding and app piracy. Some groups discuss software development for specific languages or technologies. |
| Video and Films | Discussion about movies and video sharing of movies and TV series and (possibly piracy). |
| Politics | Discussion about political news at large in different countries. |
| Erotic | Sharing and suggestion of adult content or services. |
| Linguistics | Community of users practising a particular language for educational purposes. |

Table 3: Topic description taken from Perlo et al. (2025).







Figure 12: Boxplots of the curiosity stimulus metrics (Z-score normalized) for message clusters.




(a) Cluster 0, exhibiting high indirect influence.




(b) Cluster 1, exhibiting high direct influence.




(c) Cluster 2, exhibiting mixed stimuli.




(d) Cluster 3, exhibiting high categorical novelty.



(e) Cluster 4, exhibiting low stimuli values.



(f) Cluster 5, exhibiting high uncertainty metrics.



(g) Cluster 6, exhibiting mixed stimuli.

Figure 13: SHAP beeswarm plot of message clusters. Features are ranked by the mean absolute SHAP values, prioritizing average impact over rare, high-magnitude effects.

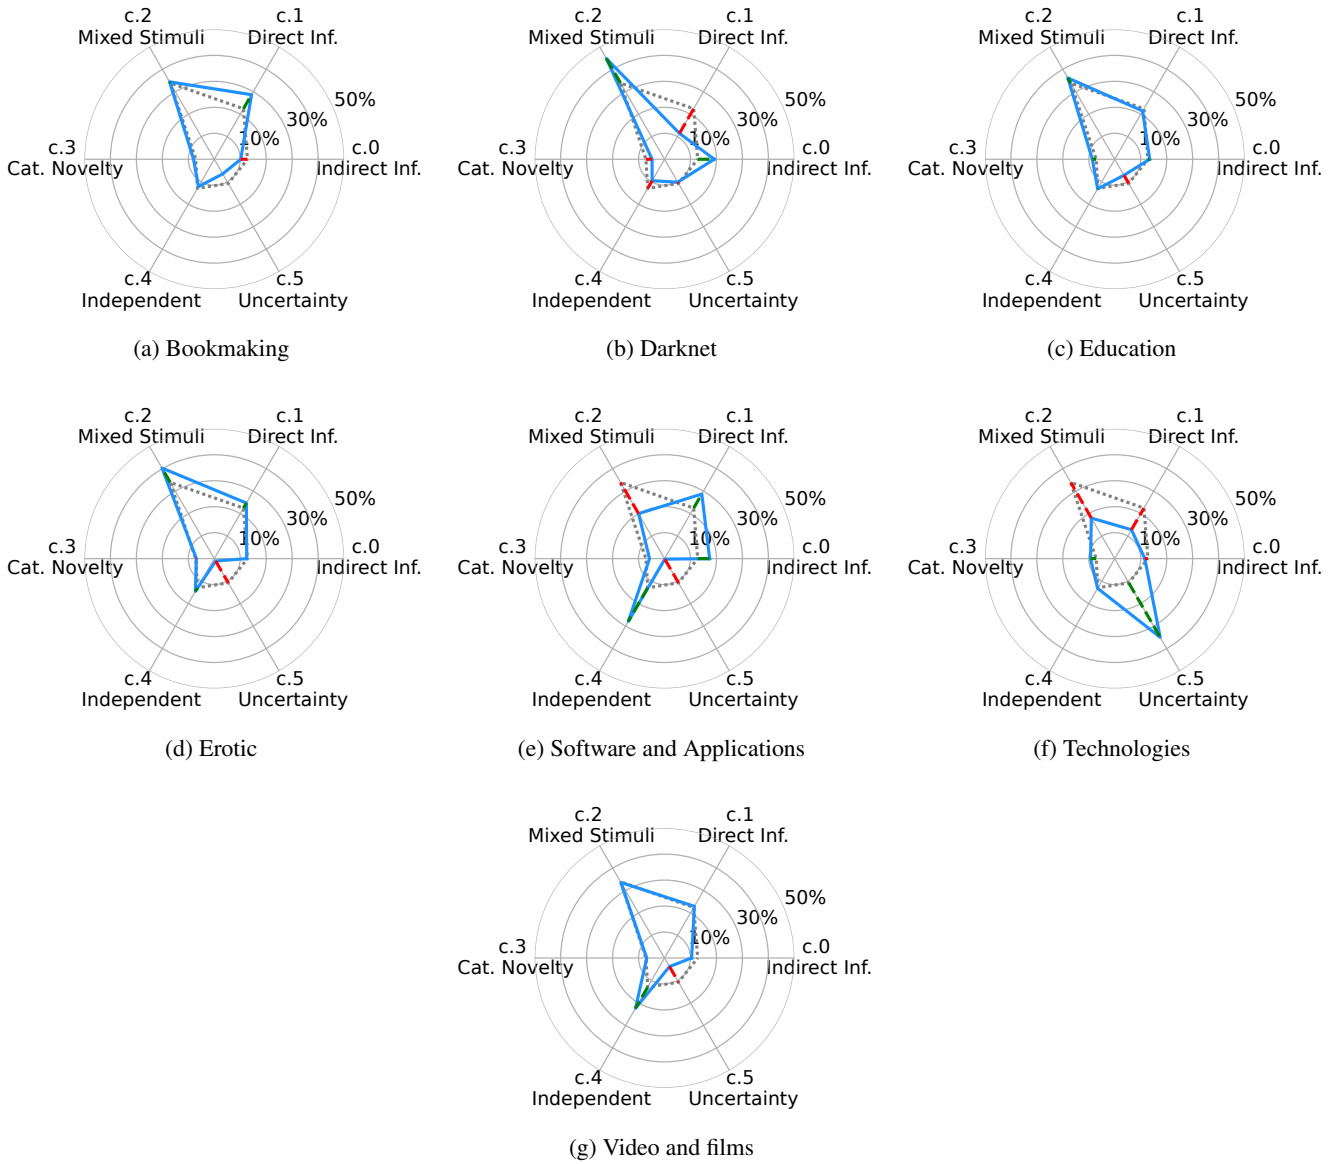


Figure 14: Average fraction of messages sent in each stimulus curiosity profile from users belonging to a group of the topic (blue solid line). The dotted grey line is the average distribution over the entire population. Dashed green/red lines highlight that the distribution in the topic population is statistically greater/lower than the general one.