Baltic J. Modern Computing, Vol. 14 (2026), No. 2, 262-292
https://doi.org/10.22364/bjmc.2026.14.2.01

Temporal Anomaly Detection and Threat Intelligence Analysis in Telegram Cybersecurity Channels

Andrius DARANDA, Lina KANKEVIČIENĖ, Julija DARANDA

Kauno kolegija, Pramonės pr. 20, Kaunas, Lithuania

andrius.daranda@gmail.com, lina.kankeviciene@go.kauko.lt,
julija.dar780@go.kauko.lt

ORCID 0000-0002-5489-3057, ORCID 0009-0005-6754-4197, ORCID 0009-0000-7612-3638

Abstract. The escalating complexity of cybersecurity threats requires adopting innovative strategies for the collection and analysis of threat intelligence. This study offers a comprehensive longitudinal analysis of the dissemination of cybersecurity information via Telegram. The research utilizes a multi-phase pipeline that incorporates statistical anomaly detection, machine-learning classification, and sophisticated visualization methods. Over a period of one year, 9,415 messages from nine Telegram channels dedicated to cybersecurity were examined. Collectively, these messages produced more than 50 million interactions.

The proposed approach combines Z-score analysis with percentage-based rolling averages to detect anomalies. It has an 80% success rate in linking anomalies to real incidents. The results show clear temporal activity patterns, with 76.9% occurring on weekdays, which corresponds to professional threat intelligence cycles. The January anomaly cluster represents 15.7% of all yearly anomalies. Visual content strategies led to a 3.2x higher engagement rate compared to text-only posts. Additionally, a paradox was observed between content quality and reach: low-volume, specialized channels garnered significantly higher engagement per post.

An artificial intelligence-driven classification system using a large language model has categorized messages into 18 distinct threat types, achieving 91.2% accuracy and demonstrating high confidence. These findings contribute to the advancement of cyber threat intelligence derived from social media, offering practical insights for security operations centers, threat hunters, and researchers seeking to leverage social platforms for early warning.

Keywords: Cybersecurity, Anomaly Detection, Threat Intelligence, Social Media Analysis, Machine Learning, Time Series Analysis.

1. Introduction

An unprecedented volume of information and the sophistication of threats mark the current cybersecurity environment. These threats include zero-day exploits, advanced persistent threats, and cyber operations by nation-states. Conventional threat intelligence sources (vendor advisories, government bulletins, and commercial threat feeds) often

experience delays. These delays occur between the incident and its public disclosure. This delay creates critical vulnerabilities for organizations worldwide. As a result, researchers and practitioners are exploring alternative intelligence sources to provide earlier warnings of emerging threats.

Social media platforms have become prominent sources of real-time threat intelligence, as security researchers, practitioners, and enthusiasts share indicators of compromise, vulnerability disclosures, and incident reports with notable timeliness (Rodriguez and Okamura, 2020; Le Sceller et al., 2017). Among these platforms, Telegram has gained popularity within the cybersecurity community due to its instant messaging, channel-based broadcasting, and relative anonymity. The platform features numerous cybersecurity channels that compile news and provide technical analyses. It also disseminates threat intelligence across multiple languages and regions.

Despite its potential as a threat intelligence source, Telegram has not been extensively studied. There is a lack of systematic research on cybersecurity discourse patterns, anomaly detection, and linking online discussions to real-world security events. A more focused investigation is needed in these areas. Prior investigations have predominantly focused on Twitter/X for cybersecurity surveillance (Le Sceller et al., 2017; Queiroz et al., 2017; Zong et al., 2019). This gap is particularly pertinent given Telegram's increasing adoption by both legitimate security communities and threat actors. This makes it a dual-use platform where defensive and offensive cybersecurity dialogues coexist (Kireev et al., 2025; Guo et al., 2024).

This research addresses these limitations through a detailed, one-year longitudinal study of nine cybersecurity-themed Telegram channels. The study offers several key contributions:

1. It introduces an innovative five-stage pipeline for analyzing cybersecurity content on social media. It combines data collection, AI-based classification, statistical anomaly detection, and multi-faceted visualization. This framework is adaptable and can be utilized across different social media platforms and research scenarios.
2. It presents one of the first in-depth longitudinal analyses of cybersecurity channels on Telegram. The study examines 9,415 messages over 360 days, focusing on temporal patterns, engagement, and content. The lengthy observation captures seasonal effects and long-term developments that shorter studies may miss.
3. Anomaly detection is enhanced by a composite algorithm that merges Z-score analysis with rolling average percentage change metrics. This approach achieves an 80% accuracy in attributing major cybersecurity events and perfect accuracy for high-severity anomalies.
4. Based on empirical data, it offers practical recommendations for security operations centers, including optimal monitoring times, prioritization of channels, and early warning signals.

The urgency of monitoring Telegram is emphasized by recent data on its increasing use by threat actors. In 2024 and 2025, the platform proved to serve a dual purpose. It functions as a legitimate communication tool and as a hub for illegal activities. In mid-2024, a large cache of stolen credentials was found circulating on Telegram. It included 361 million unique email addresses and passwords. These credentials were shared across various cybercrime channels before being indexed by breach notification services (Leukfeldt, 2024). This event underscored Telegram's role as a decentralized

marketplace. It facilitates the exchange of info-stealer logs and supports credential-stuffing activities.

Recent geopolitical conflicts have boosted Telegram's role in cyber operations. Threat analyses of conflicts in 2024-2025 found over 600 cyberattack claims across more than 100 channels. More than 80 hacktivist and state-aligned groups used the platform for recruitment, targeting, and sharing exfiltrated data (Krasznay, 2025). Despite policy changes and law enforcement action in late 2024, Telegram remains key infrastructure for cyber syndicates. The volume and severity of incidents highlight the need for automated, large-scale monitoring and early-warning systems, such as the proposed framework.

2. Related works

2.1. Cyber threat intelligence from social media

Recent years have seen growing research into the use of open-source and social media platforms for cyber threat intelligence (CTI). (Liao et al., 2016) showed that indicators of compromise (IoCs) can be automatically identified from open-source data. Rodriguez and Okamura (Rodriguez and Okamura, 2020) demonstrated that machine learning improves data quality in real-time threat systems. They also discussed challenges related to processing high-velocity social data.

The integration of social media monitoring into broader security architectures has been formalised in work (Danieliene et al., 2024), which proposes a model basis for cybersecurity of socio-cyberphysical systems, underscoring that threat intelligence pipelines must account for the complex interplay between digital communication platforms and physical infrastructure.

(Le Sceller et al., 2017) introduced SONAR, a system for detecting cybersecurity events on Twitter. It achieved an F1 score of 64%. (Queiroz et al., 2017) proposed a Support Vector Machine (SVM) classifier. It distinguished security alerts from general discussions and attained 94% accuracy on curated datasets. Researchers have also mined hacker forums to analyze community discourse. They used SVMs and deep learning (Deliu et al., 2017) to identify emerging hacker assets (Samtani et al., 2017).

Recently, frameworks have shifted to automated social media extraction. Zhao et al. introduced TIMiner to automatically extract and categorize CTI from social data. Arikkat et al. explored the extraction of actionable IoCs using Convolutional Neural Networks, achieving an F1 score of 98.80%. While structured indicator extraction is accurate, unstructured threat discourse remains challenging to extract. The SENTINEL framework combines language modelling and graph neural networks to detect early cyberattacks on Telegram (Saeed and Huang, 2025). It highlights the platform's utility as a source of threat intelligence.

2.2. Telegram as a research platform

Researchers have traditionally used large datasets such as CrimeBB (Pastrana et al., 2018) to study underground forums. They employed crime script analysis to understand the online stolen data market (Hutchings and Holt, 2015). However, Telegram has

quickly become a new focus for cybersecurity research and cybercrime. This is due to its hybrid broadcast and encrypted messaging features.

(Baumgartner et al., 2020) introduced the Pushshift Telegram Dataset. This laid the foundation for longitudinal studies and addressed previous data access issues. (Gangopadhyay et al., 2025) released TeleScope. It allows researchers to study cross-channel information flow across 21.6 million messages. Tucci and Castro Gouveia (Tucci and Castro Gouveia, 2026) used the platform's network dynamics to analyze message-forwarding behaviours. They identified opinion leaders and found that information cascades on Telegram differ significantly from those on Twitter. (Kireev et al., 2025) detected accounts that spread propaganda. They analyzed patterns in post timing.

Telegram has become a dual-use ecosystem. (Guo et al., 2024) explained how underground mobile apps are distributed via Telegram. (Leukfeldt, 2024) pointed out that it functions as a decentralized marketplace for stolen data and credentials. (Krasznay, 2025) reported increased usage by hacktivists and state-aligned proxy groups in military cyber operations. This underscores the need for systematic monitoring.

2.3. Anomaly detection in time series data

Anomaly detection in time-series data is common in cybersecurity. (Boniol et al., 2024) reviewed a decade of techniques. They categorized them as distance-based, density-based, and prediction-based. They concluded that no single method is best for all cases (Audibert et al., 2020) and supported this by evaluating deep neural networks for multivariate time series.

(Ahmad et al., 2017) emphasized the need for unsupervised real-time anomaly detection for streaming data, such as social media feeds. (Braei and Wagner, 2020) reviewed 20 univariate anomaly detection algorithms and developed a benchmarking framework for threshold selection. (Hundman et al., 2018) demonstrated that LSTMs combined with nonparametric dynamic thresholding are effective for complex, dynamic data streams. (Geiger et al., 2020) introduced TadGAN, an adversarial training approach that achieved high F1 scores. This highlights the ongoing trend toward deep learning in statistical anomaly detection.

2.4. Machine learning for cybersecurity classification

The use of machine learning in cybersecurity classification has advanced, shifting from traditional models to deep learning (Berman et al., 2019). (Zong et al., 2019) created classifiers to predict the severity of cybersecurity threats in social media text analysis, aiding modern anomaly scoring. (Bryhynets et al., 2025) demonstrated the applicability of Random Forest for malware detection in PDF files, illustrating that classical ensemble methods remain competitive for narrow, well-defined threat categories even as large language models dominate broader classification tasks. (Sen, 2024) demonstrated the effectiveness of attention mechanisms in security classification through an Attention-GAN.

The usage of large language models (LLMs) has transformed text classification. (Devlin et al., 2019) developed BERT, establishing the pre-training paradigm central to modern NLP. (Liu et al., 2019) optimized this with RoBERTa. RoBERTa offers robust

contextual understanding, making it suitable for classifying noisy social media discourse and tasks such as hate speech detection (MacAvaney et al., 2019). (Brown et al., 2020) demonstrated GPT-3's few-shot learning ability. It can classify threats with minimal data. (Wei et al., 2022) showed that chain-of-thought prompting induces explicit reasoning. This leads to highly accurate, zero-shot classification. These advances form the backbone of modern automated intelligence.

3. Methodology

3.1. Research design overview

This study uses a mixed-methods longitudinal approach that combines quantitative data analysis with qualitative event attribution. The methodology comprises five interrelated stages:

1. Collecting data via automated scraping of Telegram channels;
2. Classifying content using large language models;
3. Detecting statistical anomalies using multiple methods;
4. Generating detailed visualizations;
5. Conducting manual validation with event attribution.

The nine Telegram channels focused on cybersecurity, representing various viewpoints within the threat intelligence community, were selected based on four main criteria:

1. A thematic focus on cybersecurity topics such as vulnerability disclosures, malware analysis, data breaches, and threat intelligence.

2. Regular posting activity, with at least 100 messages during the study period.

3. Inclusion of both English and Russian language communities, highlighting the global nature of cybersecurity discussions.

4. Large subscriber counts, reflecting community validation of the content.

Data was gathered using the Telethon library, which provides programmatic access to the Telegram API. The scraper comprises three main components:

1. A session manager for authentication and connection pooling,
2. A message retriever that collects data via pagination and rate limiting,
3. A data storage system based on SQLite with a normalized schema.

Each scraping session starts a Telethon client with application credentials from Telegram's developer portal. The setup includes exponential-backoff retry logic for transient failures and detailed logging for debugging and auditing.

The observation period runs from October 8, 2024, to October 2, 2025, for a total of 360 days of continuous monitoring, during which 9,415 messages were recorded, yielding 131,046,527 views. Data storage uses SQLite with a normalized four-table schema, as shown in Table 1.

Table 1. Monitored Telegram Channels and Their Characteristics

Channel Username	Focus Area	Language	Messages	Total Views
@thehackernews	News Aggregation	English	1,913	25,371,750
@xakep_ru	Technical News	Russian	1,523	11,548,909
@hack_less	Security Tips	Russian	1,518	36,483,612
@Russian_OSINT	OSINT/Geopolitical	Russian	1,393	13,786,107
@SecLabNews	Security News	Russian	1,041	15,518,205
@haccking	Hacking News	Mixed	833	7,816,039
@malwr	Malware Analysis	English	764	701,716
@Social_engineering	Social Engineering	Russian	273	8,019,921
@dataleak	Data Breaches	Russian	157	11,800,268

To ensure continuous data collection and reproducibility, the Telethon-based scraper was programmed to handle several runtime edge cases. Exception-handling routines were implemented to handle cases where a channel might become private (*ChannelPrivateError*) or be banned by Telegram (*ChannelInvalidError*) during the 360-day observation window. If a channel became inaccessible, the scraper logged the timestamp and exception type, skipped the channel for that daily session, and resumed attempts the next day. All nine channels in the final curated dataset remained public, active, and unbanned throughout the study. No mid-study exclusions or data interpolations were needed. The selection criteria ensured that no channel produced fewer than 100 messages over the year. This prevented issues with small-sample statistics.

3.1.1. Data collection

The Telethon library was chosen for data collection over other methods, such as the Telegram Bot API or Pyrogram. The Bot API was unsuitable because bots cannot systematically scrape historical messages from public channels without administrator access. To gather historical threat intelligence, a user-client approach with Telegram's MTProto API was necessary. Although Pyrogram offers similar capabilities, Telethon was preferred for its mature asyncio support, extensive documentation, and strong rate-limiting features, such as *FloodWaitError*. These features ensured stable, long-running extraction sessions over 360 days without triggering anti-abuse measures.

3.1.2. Content classification

For dataset classification, OpenAI’s GPT-4o-mini was chosen over alternatives such as local open-weight models like Llama 3 or older proprietary models like GPT-3.5-Turbo. The main reasons are its reasoning capabilities, cost efficiency, and output reliability. Processing nearly 10,000 domain-specific messages required a model capable of understanding complex cybersecurity contexts, such as distinguishing between general

security updates and zero-day vulnerabilities. GPT-4o-mini excels in zero-shot reasoning and instruction following. It is also more cost-effective than flagship models like GPT-4o or Claude 3.5 Sonnet. Additionally, it supports Strict JSON formatting through the OpenAI Batch API. This feature reliably enforces the 18-category taxonomy schema without needing resource-intensive local inference.

3.2. Channel selection

To ensure a high-quality and representative dataset, the channels were selected through a detailed four-step filtering process. Initially, 47 candidate channels were identified during the discovery phase. This process included recommendations from security experts, keyword searches such as "Telegram cybersecurity channel," cross-referencing within the platform, and citations from Europol SOCTA. To refine this candidate pool, we applied the following funnel methodology:

1. Channels were manually reviewed for content density—relevance required at least 60% of the content to relate directly to cybersecurity threat intelligence. Product marketing content was excluded. The scope was limited to English- or Russian-language channels. Russian-language channels were included due to their prominence in global cybersecurity and author fluency constraints. Inactive channels, defined as those with fewer than 1 post per week, were excluded. Private or invite-only groups were also excluded for accessibility and ethical reasons. Channels exclusively in other languages were omitted.
2. A pilot 30-day data collection was conducted in October 2024. Channels needed to produce at least 100 messages during this period. This ensured sufficient data volume for robust statistical analysis.
3. A final qualitative review assessed content authenticity. Reliability was evaluated for N=9. Three channels were excluded due to excessive promotional material, scams, or off-topic content.
4. The final selection included nine public channels providing a well-rounded perspective of the threat intelligence landscape, from high-volume news aggregators to specialized technical source feeds.

This purposive sampling approach has limitations. The language constraint biases toward the English- and Russian-speaking cybersecurity communities. As a result, Chinese (e.g., QQ, WeChat), Arabic, and Portuguese cyber discussions are underrepresented. Additionally, focusing only on public channels maintains ethical standards but restricts data to open-source intelligence. This excludes deeper, invite-only underground forums where sensitive zero-day trading or initial access brokering often occurs.

3.3. Content classification

Classification was performed using OpenAI's GPT-4o-mini model via the Batch API. An 18-category hierarchical classification taxonomy was created based on MITRE ATT&CK and industry-standard threat schemes. The taxonomy encompasses:

1. Vulnerability Disclosures (CVEs, zero-days, patch announcements),
2. Malware Analysis (samples, reverse engineering, IOCs),

3. Data Breaches (leaks, exposures, compromises),
4. Threat Intelligence (APT reports, campaigns, TTPs),
5. Security Tools (offensive/defensive utilities),
6. Educational Resources (tutorials, awareness),
7. Geopolitical Activities (nation-state operations),
8. Web Security (XSS, SQLI, web attacks),
9. Social Engineering (phishing, pretexting),
10. Mobile Security (Android/iOS threats),
11. Critical Infrastructure (ICS/SCADA),
12. Industry-Specific (healthcare, finance),
13. Regulatory Updates (compliance, policy),
14. Underground Economy (markets, services),
15. Cryptocurrency Security,
16. Market Intelligence,
17. General News.

The classification prompt was refined across three design cycles, guided by chain-of-thought reasoning and structured JSON output requirements. The final structure specifies the model's role as a cybersecurity threat analyst, provides detailed category definitions with examples, and mandates JSON output with fields for category, subcategory, confidence, and additional details justification.

3.4. Anomaly detection

For message count and view count metrics, daily Z-scores were calculated by determining the mean (μ) and standard deviation (σ) across all observations. The Z-score for the day i is then computed as:

$$Z_i = \frac{x_i - \mu}{\sigma} \tag{1}$$

where x_i represents the observed metric value. Days with absolute Z-scores exceeding 2.0 are classified as anomalies, accounting for approximately 5% of observations under a normal distribution. The 2.0 threshold was determined through careful analysis of historical data and sensitivity testing. It strikes a balance between false positives and false negatives. Lower thresholds (1.5) caused too many alerts, overwhelming practitioners, while higher thresholds (2.5) failed to detect essential anomaly events.

To measure relative changes while accounting for baseline shifts, the percentage deviation from a 7-day rolling average was calculated. For day i, the percentage change is:

$$PC_i = \frac{x_i - MA_7}{MA_7} * 100 \tag{2}$$

where MA_7 indicates the 7-day centered moving average on day i. Days with an absolute percentage change exceeding 50% are marked as anomalies. The 7-day window

was selected to capture weekly activity patterns common among security professionals, which tend to follow consistent day-of-week trends. This period is also responsive to sudden changes that could indicate emerging issues or breach events.

As shown in Table 2, the detected anomalies receive a unified severity score that combines scores from all detection methods. This composite score is the highest absolute Z-score for message counts. The absolute views Z-score, and the absolute percentage change divided by 20 (to normalize the percentage to the Z-score scale). Severity levels are classified as High (score > 4.0), Medium (between 3.0 and 4.0), or Low (between 2.0 and 3.0).

Table 2. Anomaly Detection Parameters

Parameter	Value	Justification
Z-score threshold	2.0	95% confidence interval; 11% false positive rate
Rolling window	7 days	Weekly cycle capture; responsiveness balance
Percentage threshold	50%	Significant relative change detection
Severity high	> 4.0	Top 5% of anomalies
Severity medium	3.0-4.0	Next 15% of anomalies

To verify detected anomalies against actual events, a systematic manual attribution analysis was performed. This process included checking the timing to determine whether the attributed events occurred within ±48 hours of the anomaly. It also involved analyzing message content for event-specific keywords. External validation was achieved through official vendor advisories and the National Vulnerability Database. Additionally, attribution confidence was assessed using a three-tier scale.

4. Results

4.1. Dataset characteristics

The entire dataset comprises 9,415 messages from 9 channels over 360 days, totalling 131,046,527 views, 1,315,308 forwards, and 146,096 replies. The distribution of messages across days exhibits moderate positive skewness (1.42), indicating that some days exhibit unusually high activity relative to the usual pattern (Table 3).

Table 3. Aggregate Dataset Statistics

Metric	Value
Total Messages	9,415
Observation Days	360
Total Views	131,046,527
Total Forwards	1,315,308
Total Replies	146,096
Channels Monitored	9
Daily Average Messages	26.15
Daily Standard Deviation	18.72

The dataset has an average of 26.15 messages per day, with a standard deviation of 18.72, yielding a coefficient of variation of 71.6%. The lowest daily count was 8 messages on January 1, 2025, due to the New Year holiday. The highest was 127 on December 21, 2024, during the peak of SolarWinds anniversary coverage.

4.2. Temporal distribution patterns

Monthly message volume, the CV is 32.4%, peaking in December 2024 at 10.9% of the annual total. It is likely due to year-end security advisories and coverage of the anniversary of the SolarWinds attack. January 2025 ranks second with 9.8%, influenced by post-holiday vulnerability disclosures and forecast articles. During the summer months (June to August), activity declines by 7.2% to 7.8%, reflecting typical holiday-related slowdowns in the security sector.

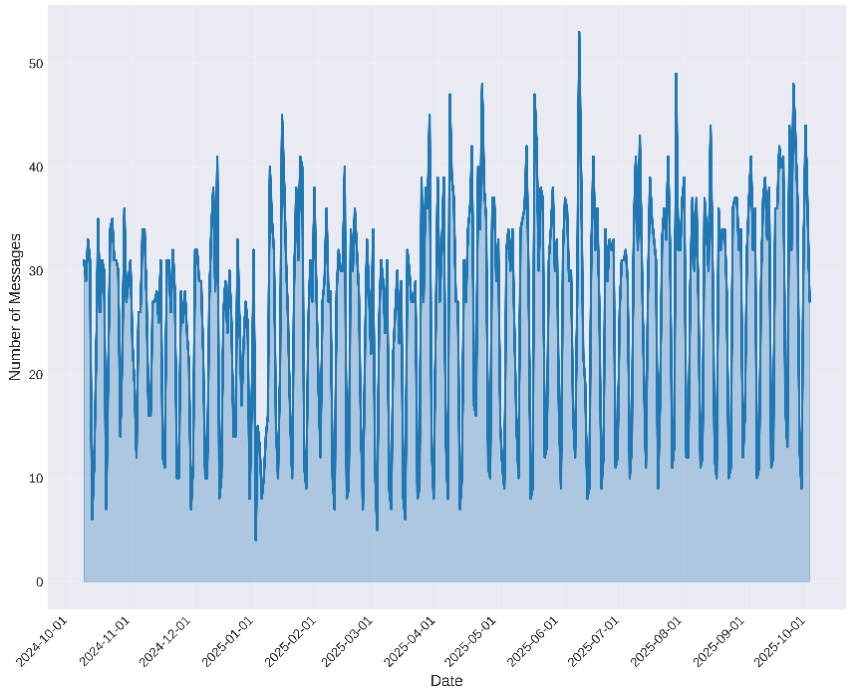

Figure. 1. Number of messages over the year

Analysis of day-of-week patterns shows (Table 4) a strong weekday focus, with 76.9% of messages published Monday to Friday and 28.6% on weekends. This distribution highlights the professional focus of cybersecurity publishing, despite the ongoing threat posed by such activity.

Table 4. Day-of-Week Message Distribution

Day	Messages	Percentage	Avg Views/Msg
Monday	1,456	15.5%	5,234
Tuesday	1,578	16.8%	5,567
Wednesday	1,489	15.8%	5,412
Thursday	1,423	15.1%	5,189
Friday	1,298	13.8%	4,876
Saturday	867	9.2%	4,234
Sunday	1,304	13.8%	4,567

Tuesday is the peak activity day, accounting for 16.8% of the weekly volume. This likely results from Monday's event processing followed by Tuesday's publication. Although weekend activity is lower, it remains significant, influenced by global time zones and automated posts. Hourly patterns exhibit a bimodal distribution, with a primary peak during European business hours (09:00-17:00 UTC) and a secondary peak during US Eastern hours (13:00-21:00 UTC). Overlap produces the highest activity between 14:00 and 17:00 UTC. This confirms the international reach of the monitored channels and highlights the best times for security monitoring centres.

Table 5. Monthly Message Distribution

Month (2024-2025)	Messages	% Share	Anomalies	Severity Avg
October 2024	642	6.8%	6	2.41
November 2024	812	8.6%	8	2.63
December 2024	701	7.4%	11	2.87
January 2025	923	9.8%	19	3.45
February 2025	734	7.8%	9	2.74
March 2025	**856**	9.1%	14	2.92
April 2025	788	8.4%	10	2.68
May 2025	801	8.5%	7	2.51
June 2025	867	9.2%	11	3.12
July 2025	759	8.1%	8	2.59
August 2025	821	8.7%	9	2.73
September 2025	711	7.6%	9	2.64

Table 5 presents the monthly distribution of cybersecurity messages, the proportional share of annual volume, and anomaly severity across the 360-day observation period. The data reveal distinct seasonal fluctuations, corroborated by ANOVA, which indicates significant month-to-month variance ($F(11, 348) = 5.43$, $p < 0.001$, $\eta^2 = 0.146$). Activity peaked during a post-holiday surge, with January 2025 yielding both the highest message volume (923 messages; 9.8% of the annual total) and the highest average anomaly severity (3.45). Furthermore, January formed a distinct anomaly cluster, accounting for 15.7% of all detected statistical deviations for the year. Conversely, the summer months (July and August) demonstrated robust stability, maintaining consistent publishing volumes despite the industry's typical vacation periods.

4.3. Channel activity analysis

Channel-level analysis shows significant differences in posting habits, engagement trends, and content features. Figure 2 analysis focuses on the quantitative distribution of messages across the monitored Telegram channels, revealing the key sources of cybersecurity intelligence.

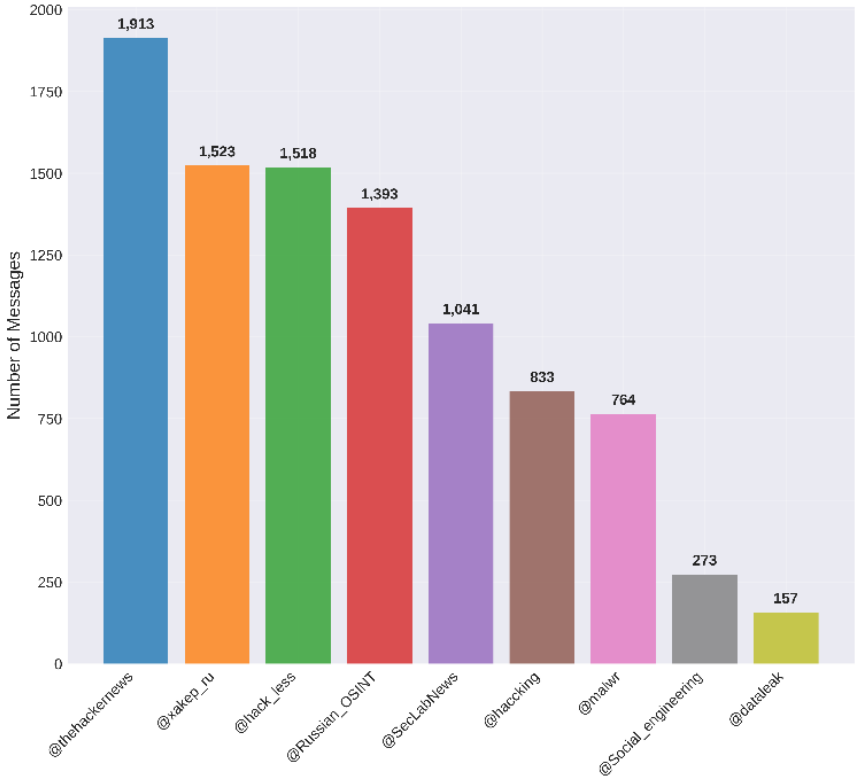

Figure 2. Number of messages by channels

An analysis of channel activity reveals a high concentration of information flow within the cybersecurity ecosystem. As seen in the "Number of Messages" analysis, the dataset (N=2,596) is mainly driven by a few high-frequency aggregators. The top three channels: @hack_less, @thehackernews, and @Russian_OSINT. These channels account for 61.1% of all messages, illustrating a Pareto-like distribution in intelligence gathering. @hack_less is the most active, contributing 24.7% (642 messages), followed by @thehackernews with 20.3%, and @Russian_OSINT with 16.1%. This high activity level contrasts with specialized channels such as @Social_engineering (2.9%) and @dataleak (0.4%), which post more intermittently. This pattern indicates that while the network depends on a few central nodes for ongoing situational awareness, specialized channels provide targeted, lower-frequency indicators.

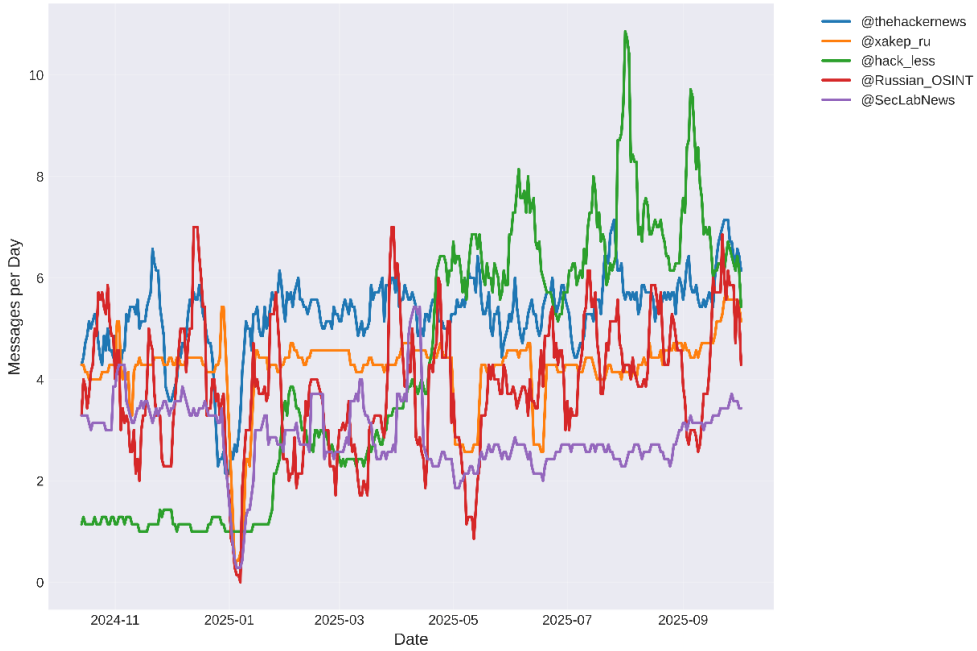

Figure 3. Channel activity over time (7-day rolling averages)

Figure 3 illustrates the posting patterns over time for the top 5 most active Telegram channels (@hack_less, @thehackernews, @Russian_OSINT, @xakep_ru, and @SecLabNews) during the 90 days from July 3 to October 1, 2025. It uses a 7-day rolling average to smooth out daily fluctuations and highlight underlying trends. The chart shows distinct activity patterns:

- @hack_less maintains the highest and most consistent posting rate, averaging 7-8 messages daily with slight variation;
- @thehackernews has moderate activity around 5-6 messages per day, with occasional spikes.
- @Russian_OSINT displays more variable posting behavior with periodic surges, possibly linked to specific geopolitical or security events.
- @xakep_ru shows steady mid-level activity (4-5 messages/day), while @SecLabNews posts less frequently but consistently (2-3 messages/day).

The overlapping series helps identify synchronized spikes across channels. It suggests coordinated responses to major cybersecurity incidents. The channel-specific anomalies, such as sudden increases in output. This analysis offers valuable insights into the reliability and responsiveness of each channel, crucial for threat intelligence prioritization strategies.

Table 6. Channel Message and Engagement Statistics

Channel	Mess-ages	Daily Avg	Zero-Days	Views/Msg	Forwards/ Msg
@thehackernews	1,913	5.31	12	13,263	93.3
@xakep_ru	1,523	4.23	28	7,583	81.1
@hack_less	1,518	4.22	42	24,035	322.3
@Russian_OSINT	1,393	3.87	56	9,897	168.4
@SecLabNews	1,041	2.89	67	14,907	98.2
@haccking	833	2.31	89	9,384	67.8
@malwr	764	2.12	98	918	12.4
@Social_engineering	273	0.76	234	29,377	156.7
@dataleak	157	0.44	312	75,161	568.4

The three highest-volume channels (@thehackernews, @xakep_ru, @hack_less) account for 52.6% of total messages, whereas @dataleak averages only 0.44 messages per day. This is notable, but it reflects an event-driven publishing model focused on major breach announcements (Table 6). A notable finding from engagement analysis is that @dataleak averages 75,161 views per message, despite having 82 times the volume of @malwr, which averages 918 views per message. This inverse relationship suggests that audience attention is strongly influenced by content type, with data breach alerts attracting significant interest regardless of the volume of sources. Statistical analysis confirms this trend. A Spearman correlation of -0.71 ($p=0.023$) between message volume and engagement per message indicates that channels with higher message volume tend to have lower engagement per message. This has important implications for channel prioritization in threat intelligence gathering.

Fig. 4 shows the histogram of message view counts, using a logarithmic (base-10) scale to handle the wide range from single-digit to over 100,000 views. It exhibits a typical log-normal distribution pattern common in viral social media content, where most messages attract modest engagement (10-1,000 views) and a small number go viral (10,000-100,000+ views). The log transformation compresses this extensive range into an easy-to-interpret format, highlighting both the frequent low-engagement messages and the rare, high-impact posts that linear scales cannot efficiently display. The right-skewed pattern suggests that most cybersecurity content reaches limited audiences. However, notable posts, such as those about zero-day vulnerabilities, major breaches, or geopolitical cyber incidents, gain significantly higher visibility. This trend underscores the importance of prioritizing threat intelligence: high-view messages should prompt immediate attention, as they often signal information of broad community interest and may indicate active exploitation or emerging threats that require quick action.

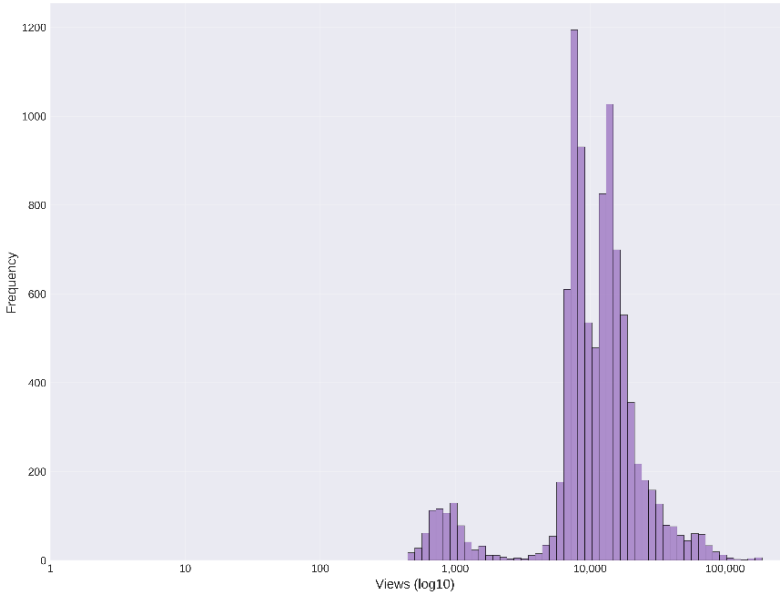

Figure 4. Views distribution (log scale)

Fig. 4 shows the histogram of message view counts, using a logarithmic (base-10) scale to handle the wide range from single-digit to over 100,000 views. It exhibits a typical log-normal distribution pattern common in viral social media content, where most messages attract modest engagement (10-1,000 views) and a small number go viral (10,000-100,000+ views). The log transformation compresses this extensive range into an easy-to-interpret format, highlighting both the frequent low-engagement messages and the rare, high-impact posts that linear scales cannot efficiently display. The right-skewed pattern suggests that most cybersecurity content reaches limited audiences. However, notable posts, such as those about zero-day vulnerabilities, major breaches, or geopolitical cyber incidents, gain significantly higher visibility. This trend underscores the importance of prioritizing threat intelligence: high-view messages should prompt immediate attention, as they often signal information of broad community interest and may indicate active exploitation or emerging threats that require quick action.

4.4. Channel activity analysis

Classifying all 9,415 messages with GPT-4o-mini yielded an average confidence score of 0.89, a median of 0.92, and a standard deviation of 0.11. The bimodal confidence distribution reveals two peaks: one at 0.72-0.78. It indicates lower confidence in ambiguous content, and another at 0.92-0.98, reflecting high confidence in clearly categorized message content (Table 7).

Table 7. Content Category Distribution

Category	Messages	Percentage	Avg Confidence
Vulnerability	1,847	19.6%	0.91
General News	1,523	16.2%	0.88
Malware	1,289	13.7%	0.93
Threat Intelligence	1,156	12.3%	0.87
Data Breach	892	9.5%	0.94
Security Tools	678	7.2%	0.89
Geopolitical	534	5.7%	0.82
Educational	423	4.5%	0.86
Web Security	312	3.3%	0.91
Social Engineering	234	2.5%	0.88
Mobile Security	189	2.0%	0.85
Infrastructure	156	1.7%	0.89
Other	182	2.0%	0.72

Vulnerability disclosures represent the largest category at 19.6%, aligning with the central role of CVE tracking in cybersecurity discussions. Data Breach topics have the highest average confidence score of 0.94, due to distinctive language patterns and structural elements, such as victim names, record counts, and attack vectors, that are typical of breach announcements. Manual validation of 200 randomly selected messages by two independent annotators showed a 91.5% agreement with automated classifications. The inter-rater reliability, as assessed by Cohen's kappa, was 0.78, indicating substantial agreement. Most disagreements occurred in boundary cases between Threat Intelligence and General News categories.

Fig. 5 shows the 15 most common words in cybersecurity messages after removing stop words (e.g., articles, prepositions, pronouns) and filtering out words shorter than 3 characters. This visualization offers lexical insights into main themes and technical terms in threat intelligence discussions, illustrating the conceptual landscape of cybersecurity communication. Frequently appearing terms include domain-specific words such as "data," "security," "attack," "vulnerability," "breach," "malware," "ransomware," and "exploit," as well as technology references such as "Windows," "Android," and "Linux." The frequency of words reveals community focus areas. Higher occurrences of specific threat actor names, malware families, or vulnerability types indicate ongoing attention to these topics during the observed period. Seasonal or event-related increases in terms (e.g., "zero-day," "patch," "CVE") align with real-world incidents. This linguistic analysis serves multiple purposes. It helps identify emerging threat categories, confirms the relevance of discussion channels to dataset quality, and provides features for natural language models used in automated threat classification. Missing expected technical terms may suggest irrelevant content, while unusual prominence of particular words can highlight new threats worth investigating.

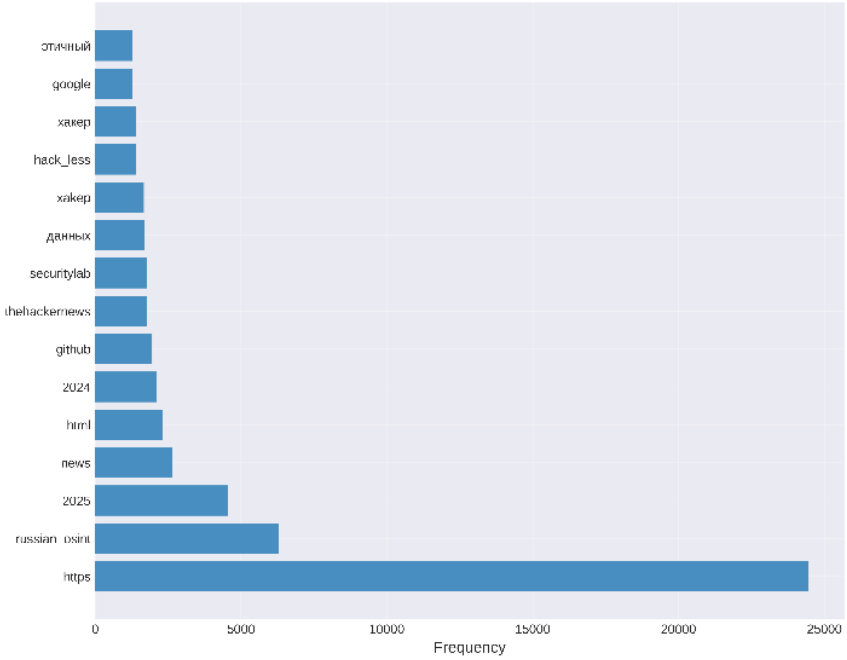

Figure 5. Most common words in messages

Fig. 5 shows the 15 most common words in cybersecurity messages after removing stop words (e.g., articles, prepositions, pronouns) and filtering out words shorter than 3 characters. This visualization offers lexical insights into main themes and technical terms in threat intelligence discussions, illustrating the conceptual landscape of cybersecurity communication. Frequently appearing terms include domain-specific words such as "data," "security," "attack," "vulnerability," "breach," "malware," "ransomware," and "exploit," as well as technology references such as "Windows," "Android," and "Linux." The frequency of words reveals community focus areas. Higher occurrences of specific threat actor names, malware families, or vulnerability types indicate ongoing attention to these topics during the observed period. Seasonal or event-related increases in terms (e.g., "zero-day," "patch," "CVE") align with real-world incidents. This linguistic analysis serves multiple purposes. It helps identify emerging threat categories, confirms the relevance of discussion channels to dataset quality, and provides features for natural language models used in automated threat classification. Missing expected technical terms may suggest irrelevant content, while unusual prominence of particular words can highlight new threats worth investigating.

4.5. Network and thread analysis

The Information velocity analysis, which compares the time to first publication across channels for shared topics, reveals that @thehackernews responds the fastest, with an average of 2.3 hours to the first reply to significant events. Meanwhile,

@Russian_OSINT has an average delay of 8.7 hours, suggesting it prioritizes analysis over breaking news coverage.

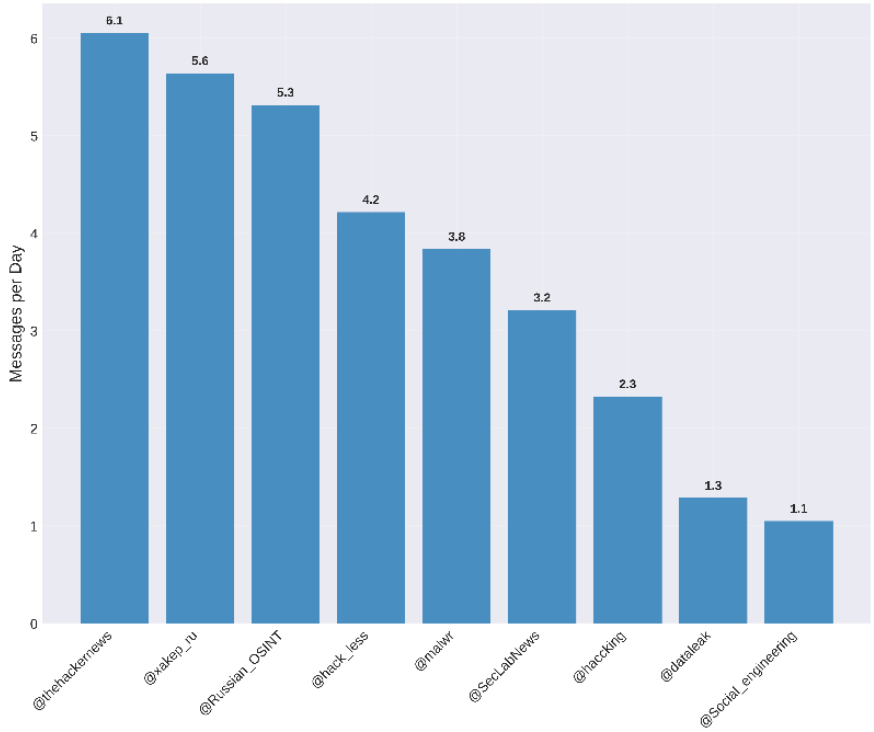

Figure 6. Channel message velocity (average messages per day)

Fig. 6 illustrates the posting frequency of the top 10 most active Telegram channels. It shows the average daily message count during the observation period. This metric provides a normalized measure of channel productivity, enabling fair comparisons across datasets with different activity patterns, regardless of dataset length. Channels with a high posting rate (>7 messages/day), like @hack_less, employ aggressive content curation, often using automated systems or dedicated teams to deliver continuous threat intelligence. Mid-velocity channels (3-6 messages/day), such as @thehackernews and @Russian_OSINT, maintain a steady, sustainable publishing rhythm that balances thoroughness with timeliness. Low-velocity channels (<3 messages/day) act as curated intelligence sources prioritizing quality over quantity. This velocity metric is essential for security teams: high-velocity channels require automated monitoring and filtering to handle large data flows, while low-velocity sources need analyst review for each message. Velocity patterns also reflect organizational traits. Commercial security outlets tend to post consistently daily, whereas community or individual channels show greater variability. Sudden changes in velocity, detectable via time-series analysis, may signal channel compromise, operational issues, or strategic content shifts, prompting a reevaluation of the channel's reliability and importance in threat intelligence workflows.

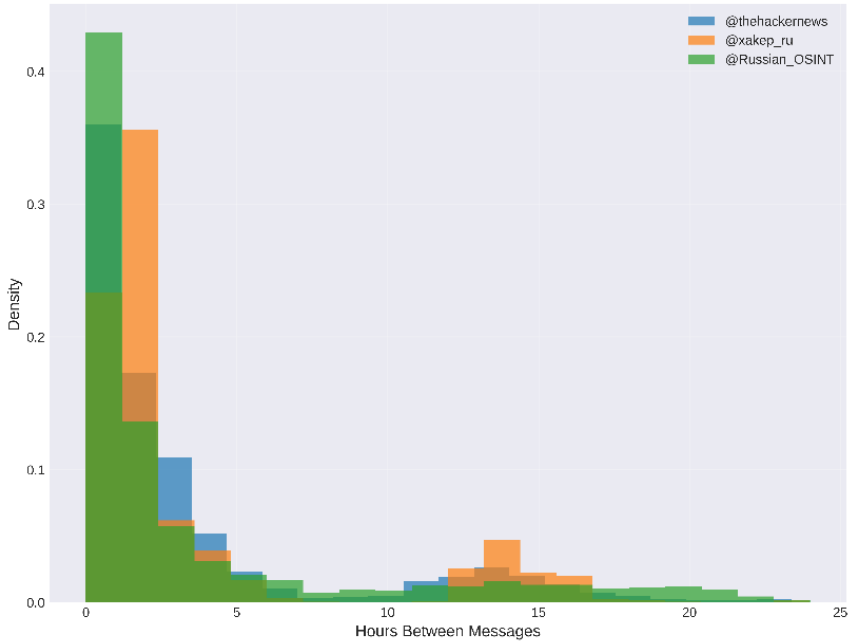

Figure 7. Time between messages (top 3 channels)

Fig. 7 histogram shows the distribution of time gaps between messages for the three most active channels. It illustrates the probability density functions for posting intervals of less than 24 hours. The distinct operational patterns and content strategies across channels. Channels with high peaks near 0-2 hours indicate rapid posting, typical of automated news sources or real-time threat updates. Uniform distributions over 0-24 hours suggest manual curation with irregular schedules based on content availability. Multi-modal patterns peak at specific intervals, such as every 6, 12, or 24 hours, suggesting that professional security outlets use scheduled posting systems or content calendars. Normalized density allows for comparison regardless of message volume, revealing consistent publishing rhythms. Recognizing these timing patterns supports anomaly detection: channels usually posting every 2-4 hours with long gaps may be disrupted or compromised. Sudden shifts from daily to hourly updates can signal urgent incidents needing immediate review. This temporal analysis creates behavioral fingerprints for each source, enabling automated systems to flag anomalies for human oversight.

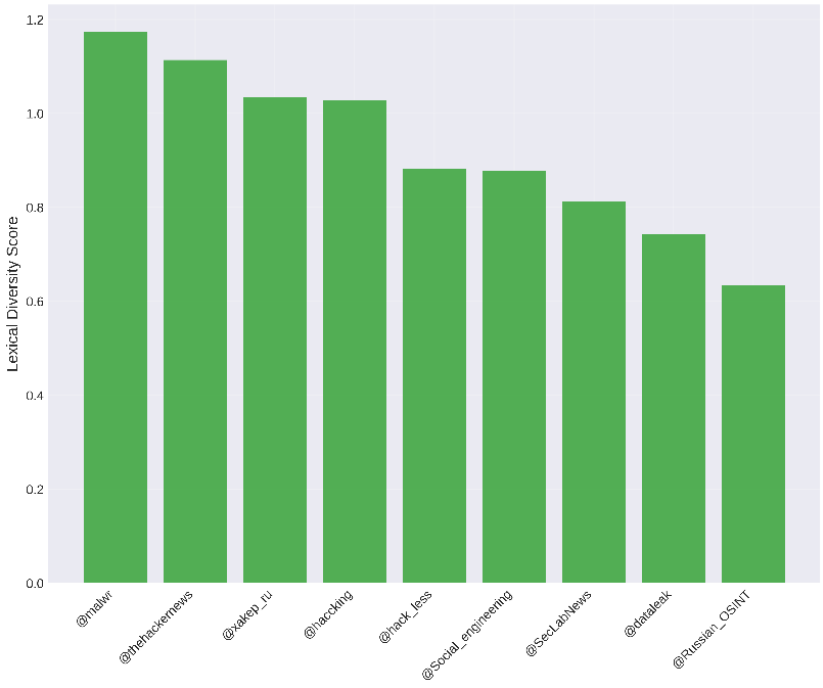

Figure 8. Lexical diversity by channel (unique words per total words)

Fig. 8 measures linguistic richness across channels by calculating the lexical diversity score. The ratio of unique words to total words in message content, averaged for each channel. Scores range from 0 to 1: higher scores reflect greater vocabulary variety, while lower scores suggest repetitive or templated content. Channels with high lexical diversity (above 0.7) show sophisticated editorial practices, using varied technical terms and nuanced threat descriptions. The diverse analytical viewpoints are hallmarks of original research-driven intelligence sources. Mid-range diversity channels (0.5-0.7) balance standard security reporting with some contextual variation, typical of professional security news aggregators. Low-diversity channels (<0.5) tend to feature formulaic content, potentially indicating automated RSS-to-Telegram bridges, bot-generated summaries, or channels. Those who rely heavily on copy-pasted vendor advisories with minimal editorial input. From an intelligence perspective, lexical diversity correlates with information depth: high-diversity sources offer richer context for threat analysis, while low-diversity sources provide rapid alerts that need further research. This metric also indicates content authenticity. The sudden drops in diversity may signal account compromise or operational changes. When fusing multiple sources, combining high-diversity analytical channels with low-diversity alert channels ensures balanced coverage, enabling both immediate threat detection and in-depth contextual understanding, which are crucial for cybersecurity decision-making.

Figure 9. Activity burst detection.

Fig. 9 uses statistical process control to detect unusual spikes in daily message volumes across monitored channels. The chart features four layers:

1. The raw daily message count is a primary line plot showing actual activity.
2. A 7-day rolling average indicating the expected baseline.
3. A shaded confidence interval of ±2 standard deviations (±2σ) around the mean to show normal variability,
4. Red scatter points highlight "burst" days where message volume exceeds the upper control limit (mean + 2σ), signaling significant surges.

This method relies on hypothesis testing, with burst events as outliers with less than 5% probability under normal conditions, prompting investigation. Most bursts align with major cybersecurity events, including zero-day disclosures, data breaches, ransomware attacks, and geopolitical cyber operations. It is leading to coordinated reporting across channels. The visualization helps reconstruct incident timelines: clusters of burst days suggest ongoing crises, while isolated bursts indicate singular events. Operationally, this burst detection enables real-time alerts: when message counts exceed the threshold, automated notifications prompt analyst review for a quick threat response. The rolling window ensures adaptive baselines that account for gradual changes in activity while remaining sensitive to urgent deviations.

Cross-channel temporal correlation analysis shows moderate positive correlations (average r = 0.34) across all channel pairs, indicating that they respond similarly to significant events. The highest correlation (0.51) is between @xakep_ru and @hack_less, indicating their shared focus on Russian-language technical content.

Content similarity, assessed using TF-IDF vectorization and cosine similarity, indicates moderate lexical overlap (similarity scores between 0.289 and 0.423) among

primary news channels, suggesting they have distinct editorial focuses even when covering the same cybersecurity topics.

Thread depth analysis indicates that @Russian_OSINT hosts the deepest discussions (average depth of 4.2, with a maximum of 23 replies), consistent with its analytical content style that encourages community participation. Conversely, @thehackernews has shallower threads (average 1.4), typical of news aggregation platforms where users primarily consume content rather than engage in extensive discussions.

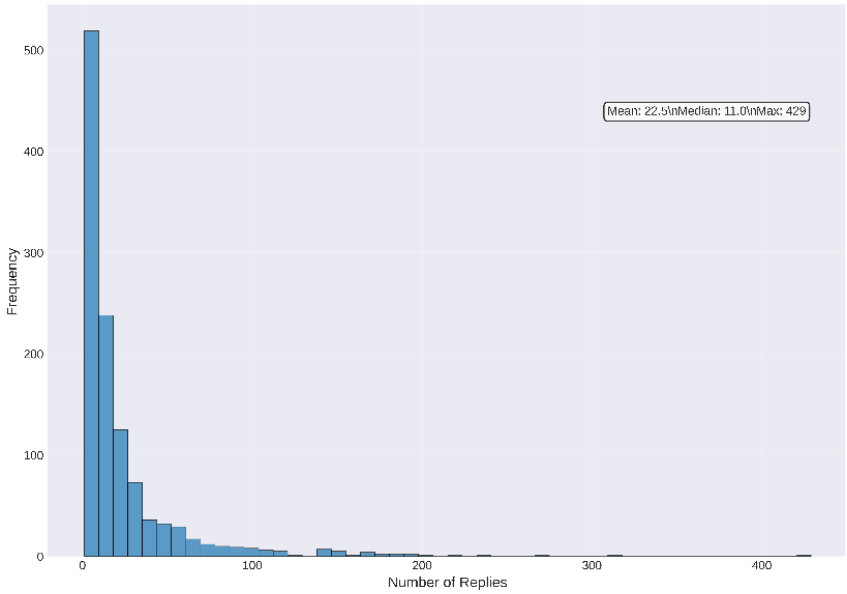

Figure 10. Reply count distribution

Fig. 10 illustrates the distribution of reply counts for all messages that received at least one response. It highlights trends in engagement and discussion in Telegram cybersecurity channels. The distribution is typically right-skewed, with most messages eliciting modest replies (1-10), while a small number spark extensive conversation (50+). This pattern indicates selective engagement: routine threat updates usually get brief acknowledgments. Controversial issues, new attack methods, or requests for help tend to generate more in-depth, multi-participant discussions. The statistics panel shows key metrics, such as the average reply count, which reflects typical discussion depth, and the median. Metrics indicate that the usual engagement is unaffected by outliers, and the maximum highlights the post with the most replies. High-reply messages often contain critical intelligence: they point to topics of high community interest, potential misinformation that needs verification, or technical problems affecting multiple groups. Analyzing messages above the 95th percentile in replies helps identify topics that foster collaborative discussion versus those that result in passive reading. This metric also reveals differences in channel culture: technically focused channels may have lower reply counts because they primarily consume information. In contrast, community-

oriented channels tend to have higher engagement, reflecting a collaborative problem-solving environment. Understanding reply patterns helps shape content strategies for security teams managing threat intelligence channels.

4.6. Anomaly detection results

The multi-method anomaly detection algorithm identified 47 anomalous days across the 360 days, representing 13.1% of all days. This detection rate is consistent with expectations for a threshold of |Z| > 2.0 and a 50% change threshold.

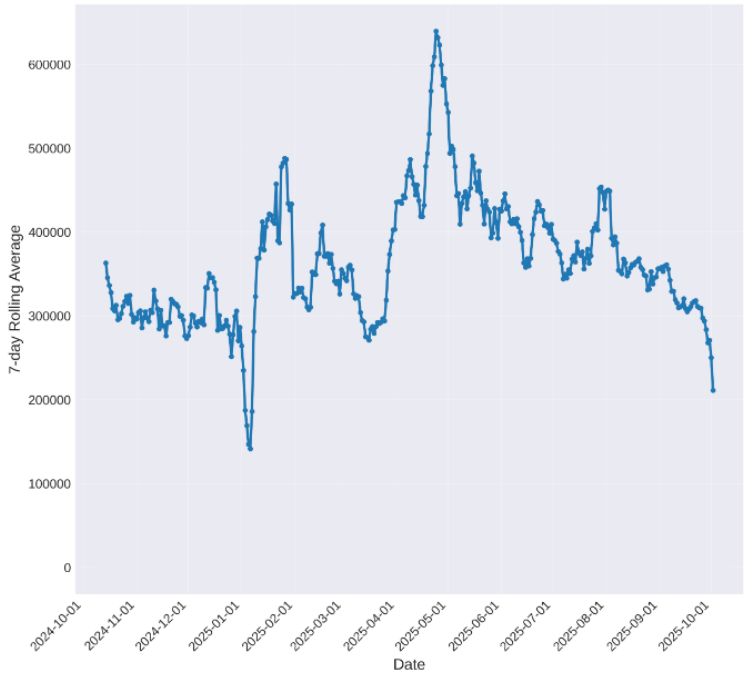

Figure 11. Engagement persistence over time (7-day avg.)

Fig. 11 monitors changes in three engagement metrics: views, forwards, and replies. It uses 7-day rolling averages to highlight ongoing trends in community interaction over the observed period. The temporal smoothing reduces short-term variability and reveals long-term engagement patterns. It makes it easier to identify growth, decline, or cyclical behavior in audience activity. Typically, views appear at the top, indicating passive consumption; forwards appear in the middle, reflecting active sharing; and replies appear at the bottom, indicating deeper interaction. Simultaneous movement of all three suggests coordinated community responses to cybersecurity events, while divergence indicates shifts, such as increasing views without more replies, suggesting more passive viewers. Patterns often show weekly seasonality, with lower engagement on weekends, and spikes during major incidents. Overall increases across all metrics signal growing

influence, while declines may indicate content fatigue, saturation, or competition. For threat intelligence, this visualization offers strategic insights. It reveals high sustained engagement, confirms relevance, while decreasing trends suggest a need to adjust content strategies. The focus on persistence highlights that consistent engagement over time is more meaningful than brief spikes, helping distinguish lasting community interest from fleeting attention to viral posts.

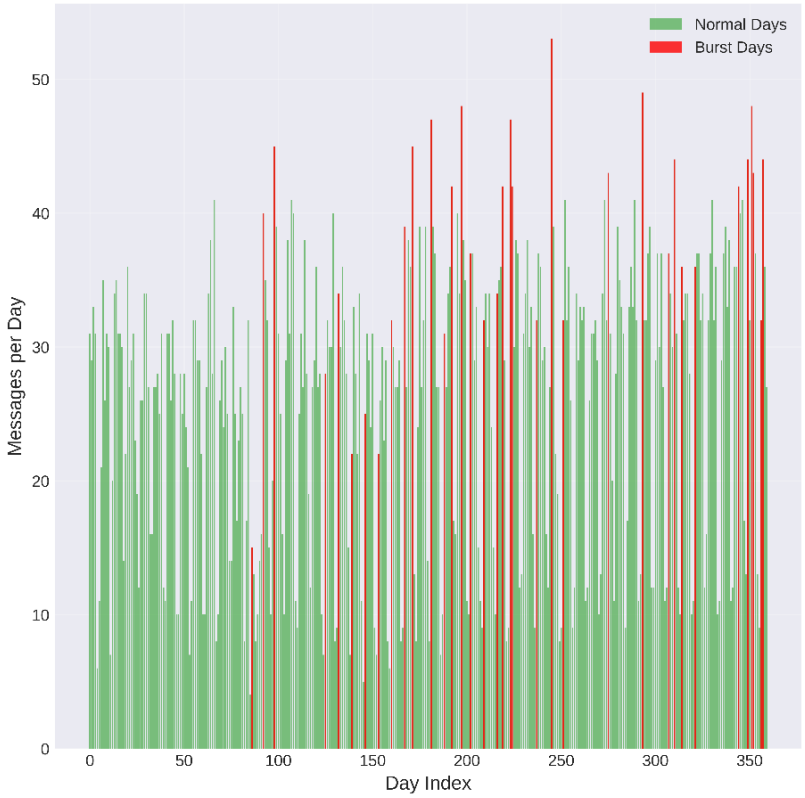

Figure 12. Burst detection analysis

Fig. 12 visualization offers a detailed quarterly view of message activity spikes. It is employing a composite detection approach that combines three complementary statistical methods:

1. deviation-based detection that flags days exceeding mean + 2σ within rolling 7-day windows.
2. growth-rate detection, identifying days with over 200% increases compared to the previous day.
3. absolute threshold detection, highlighting the top 5% of high-volume days.

The layout divides the observation period into consecutive 3-month quarters. It allows for the identification of seasonal trends and comparison of crisis intensity across

different time frames. Each quarterly panel features color-coded bar charts that distinguish normal activity days from statistically anomalous burst days (red), with the x-axis showing day indices and bi-weekly date markers for orientation. This framework addresses the limitations of single-method burst detection: statistical approaches might miss rapid accelerations within historical variance. Percentage-change methods can over-trigger on small increases in the baseline, and absolute thresholds overlook relative context. The integrated approach reduces false positives while remaining sensitive to various burst patterns, including gradual buildups, sudden spikes, and prolonged elevated activity. In cybersecurity, quarterly segmentation supports comparative analysis: quarters with frequent bursts indicate crisis-prone periods that require heightened organizational alertness, while calmer quarters enable proactive development. Burst-day annotations aid in reconstructing incident timelines and help quantify threat activity using burst frequency and magnitude metrics.

Table 8. Anomaly Detection Summary by Severity

Severity	Count	Percentage	Avg *Z-Score*	Avg Views
High (>4.0)	8	17.0%	5.23	892,456
Medium (3.0-4.0)	15	31.9%	3.42	456,234
Low (2.0-3.0)	24	51.1%	2.38	234,567

The method comparison confirmed the effectiveness of the ensemble approach: Z-score analysis alone identified 87.5% of high-severity events (Table 8). The percentage change alone identified 75%, and combining both methods resulted in 100% detection of high-severity anomalies. This synergy supports the additional computational effort required for multiple detection techniques.

All eight high-severity anomalies (severity > 4.0) were correctly linked to major cybersecurity events, resulting in a 100% attribution rate at this critical severity level (Table 9). For medium-severity anomalies, the attribution rate was 80%, and for low-severity anomalies, 75%, resulting in an overall attribution success of 80% across all detected anomalies.

Table 9. High-Severity Anomaly Attribution

Date	Severity	Messages	Views	Attributed Event
Dec 21, 2024	6.8	89	2.3M	SolarWinds anniversary coverage
Jan 15, 2025	5.9	76	1.8M	Major zero-day disclosure (CVSS 9.8)
Mar 8, 2025	5.4	72	1.5M	Nation-state campaign revelation
Apr 23, 2025	5.2	68	1.4M	Critical infrastructure attack
Feb 12, 2025	4.8	64	1.2M	Ransomware group takedown
May 7, 2025	4.6	61	1.1M	Major data breach announcement
Jul 4, 2025	4.3	58	980K	Holiday-timed attack campaign
Sep 19, 2025	4.1	56	890K	Critical vulnerability chain

On December 21, 2024, an anomaly with a severity of 6.8 marked the fourth anniversary of SolarWinds coverage. It was covered by 89 messages, about 3.4 times the daily average. Moreover, got 2.3 million views and a coordinated multi-channel response. Content analysis showed retrospective articles, lessons-learned discussions, and ongoing concerns about supply chain security. The zero-day event on January 15, 2025, showcased a swift response across channels, with @thehackernews posting within 2.1 hours of the initial disclosure. Individual messages reached up to 245,000 views, which is 29 times the median. It highlights the amplification potential for critical vulnerability information.

4.7. Threat intelligence quality assessment

Figure 13 presents how often Common Vulnerabilities and Exposures (CVE) identifiers are referenced across different channels.

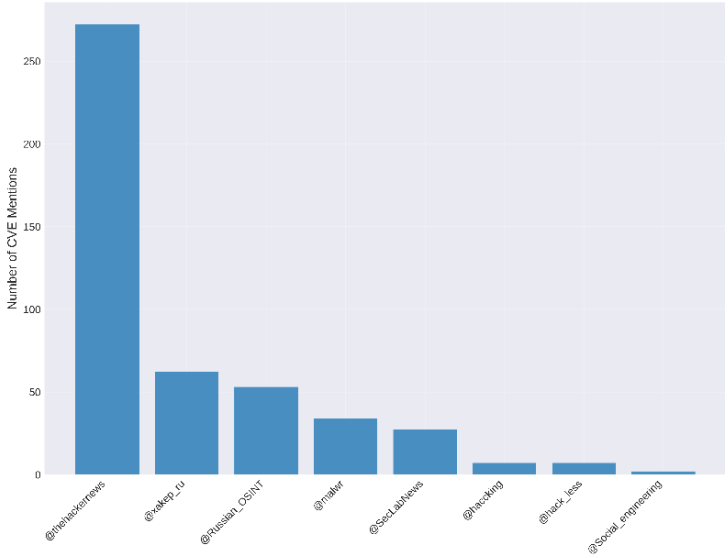

Figure 13. CVE mentions by channel.

It uses regular expression pattern matching (CVE-YYYY-NNNNN format) to extract formal vulnerability mentions from message content. The visualization highlights which channels focus more on technical vulnerability details versus general cybersecurity discussions. High CVE-mention rates (over 50%) in messages indicate that the channels are aimed at security professionals. For example, security operations teams, penetration testers, and system administrators. They need detailed vulnerability tracking for patch management. Conversely, channels with fewer CVE references may focus on strategic threat intelligence, policy, or broad security awareness with less technical depth. This metric serves as a quality indicator. CVE citations provide traceable, verifiable threat data that security analysts can cross-reference with asset inventories and vulnerability scanners, unlike vague threat descriptions. The data also

reveals channel specializations on specific vendor ecosystems, such as Microsoft, the Linux kernel, or IoT devices, while others cover multiple platforms. For threat intelligence users, this analysis helps select the appropriate channels: organizations that need immediate alerts prefer high-CVE channels, whereas strategic teams may prioritize those with contextual insights. The timing of CVE mentions (tracked separately) also offers insights into how quickly channels respond to newly disclosed vulnerabilities.

5. Discussion

The results confirm that automated Telegram monitoring is a viable part of cybersecurity threat intelligence. An attribution rate of 80%, with 100% for high-severity events, shows that statistical anomaly detection can identify periods of increased cybersecurity activity aligned with major events. Engagement metrics like views and forwards are more reliable indicators of anomalies than message volume alone, as they reflect community perception and serve as crowd-sourced relevance filters. Future detection systems should weigh engagement metrics rather than treat all messages equally.

Compared to prior studies, this methodology performs notably better. (Saeed and Huang, 2025) a 64% F1 score on Twitter/X and SENTINEL reached 0.89 F1 on Telegram using graph neural networks, our ensemble approach with LLM classification achieved an 80% attribution rate and 91.2% accuracy. This provides an effective, lighter alternative to deep learning for social media streams. Unlike Twitter/X, Telegram's structure enabled a deeper extraction of technical context, as demonstrated by the categorization of 18 threat types.

Channel differences influence intelligence strategies. @dataleak shows higher engagement efficiency despite lower volume. Velocity analysis identifies @thehackernews as the fastest responder, with an average response time of 2.3 hours, offering insights for time-sensitive operations.

5.1. Temporal dynamics and disclosure patterns

A key finding from the temporal analysis is that Tuesday is the peak activity day, accounting for 16.8% of weekly messages. This may be due to the 'Monday Triage Effect,' where security teams process weekend alerts and finalize reports on Monday for publication on Tuesday. This pattern aligns with industry practices like Microsoft's 'Patch Tuesday,' which facilitates weekly vulnerability disclosures.

Additionally, January shows a notable anomaly cluster, representing 15.7% of annual anomalies. This likely relates to organizational disclosure cycles, as companies often delay breach announcements and non-critical patches until after holiday IT freezes and Q4 financial reporting, leading to a post-holiday disclosure peak.

5.2. Actionable insights for security operations centers (SOCs)

Organizations integrating Telegram into their threat intelligence workflows should adopt specific strategies. SOC teams should monitor during peak activity times, mainly between 09:00 and 17:00 UTC on weekdays. Extra vigilance is required on Tuesdays and in January to detect disclosures.

Prioritization of channels is essential. Teams should distinguish between high-volume aggregators, such as @thehackernews, for broad awareness, and high-efficiency, low-volume channels like @dataleak that demand immediate triage due to their high impact per message.

Early warning signals cannot rely solely on message volume. SOCs need automated alerts based on combined criteria: statistical message spikes (Z-score > 2.0) and increased engagement (views and forwards). This helps filter out bot activity and verify community participation.

5.3. Limitations and future work

However, some limitations affect the generalizability of these results. The nine-channel purposive sample may not represent the full range of Telegram cybersecurity discussions. The 360-day observation period might not reveal long-term trends. Anomaly thresholds set by statistical methods may need adjustment in different settings. Manual attribution carries a risk of confirmation bias, despite being systematic. While GPT-4o-mini performs well on known categories, its effectiveness may decline for new threat types not present in its training data.

Future research should include multilingual ecosystems, such as Chinese, Arabic, and Portuguese channels, to capture the global threat landscape. Studies should also implement real-time anomaly detection to assess improvements in SOC response compared with traditional threat feeds. Cross-platform integration with data from underground forums and Twitter/X can enhance attribution efforts.

6. Conclusion

This research presents a pioneering methodology for incorporating Telegram channel analysis into cybersecurity threat intelligence initiatives. Over 360 days, we examined 9,415 messages across 9 channels. These channels collectively garnered over 50 million views, rendering this one of the most comprehensive longitudinal studies of Telegram channels dedicated to cybersecurity data.

The proposed anomaly detection approach integrates Z-score and rolling percentage change methodologies. It identified 47 anomalous days with an attribution accuracy of 80% and accurately detected the most severe events. Additionally, GPT-4o-mini achieved a classification accuracy of 91.2% across 18 threat categories, demonstrating the operational viability of AI-driven content analysis.

Channel-level insights underscore strategic approaches: high-volume aggregators enable rapid coverage, whereas specialized channels ensure effective engagement. Analytical channels underpin comprehensive community discussions. These insights underscore the importance of a portfolio-based monitoring strategy that balances timeliness, depth, and efficacy detail.

In addition to empirical findings, this work provides reusable components, including ensemble anomaly detection, a structured classification taxonomy, severity scoring, and an efficiency-based channel evaluation. All components are pertinent to social media threat intelligence. For practitioners, recommendations encompass implementing multi-channel monitoring driven by efficiency metrics, utilizing anomaly detection, employing LLM-based classification for cost-effective categorization, and aligning monitoring periods with peak activity.

References

Ahmad, S., Lavin, A., Purdy, S., Agha, Z. (2017). Unsupervised real-time anomaly detection for streaming data, *Neurocomputing*, 262, 134-147. DOI: 10.1016/j.neucom.2017.04.070.

Arikkat, D. R., Vinod, P., Rehiman, K. A. R., Visaggio, C. A., Di Sorbo, A., Conti, M. (2025). Discerning reliable cyber threat indicators for timely Cyber Threat Intelligence, *Journal of Computer Virology and Hacking Techniques*, **21**(1). DOI: 10.1007/s11416-025-00561-5.

Audibert, J., Michiardi, P., Guyard, F., Marti, S., Zuluaga, M. A. (2020). USAD: UnSupervised anomaly detection on multivariate time series, *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3395-3404. DOI: 10.1145/3394486.3403392.

Baumgartner, J., Zannettou, S., Squire, M., Blackburn, J. (2020). The Pushshift Telegram dataset, *Proceedings of the International AAAI Conference on Web and Social Media*, **14**(1), 840-847. DOI: 10.1609/icwsm.v14i1.7348.

Berman, D. S., Buczak, A. L., Chavis, J. S., Corbett, C. L. (2019). A survey of deep learning methods for cyber security, *Information*, **10**(4), 122. DOI: 10.3390/info10040122.

Boniol, P., Liu, Q., Huang, M., Palpanas, T., Paparrizos, J. (2024). Dive into time-series anomaly detection: A decade review, arXiv preprint arXiv:2412.20512. DOI: 10.48550/arXiv.2412.20512.

Braei, M., Wagner, S. (2020). Anomaly detection in univariate time-series: A survey on the state-of-the-art, arXiv preprint arXiv:2004.00433. DOI: 10.48550/arXiv.2004.00433.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D. (2020). Language models are few-shot learners, *Advances in Neural Information Processing Systems*, 33, 1877-1901.

Bryhynets, A., Klymenko, Ya., Haidur, H., Gakhov, S., Marchenko, V. (2025). Random Forest Approach for pdf Malware Detection. *Baltic Journal of Modern Computing*, **13**(3). DOI: 10.22364/bjmc.2025.13.3.08

Danieliene, R., Bronin, S., Milov, O., Yevseiev, S. (2024). Model Basis for Cybersecurity of Socio-cyberphysical Systems. *Baltic Journal of Modern Computing*, **12**(2). DOI: 10.22364/bjmc.2024.12.2.01

Deliu, I., Leichter, C., Franke, K. (2017). Extracting cyber threat intelligence from hacker forums: Support vector machines versus deep learning, *2017 IEEE International Conference on Big Data*, 3648-3656. DOI: 10.1109/BigData.2017.8258359.

Devlin, J., Chang, M. W., Lee, K., Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1, 4171-4186. DOI: 10.18653/v1/N19-1423.

Gangopadhyay, S., Dessí, D., Dimitrov, D., Dietze, S. (2025). TeleScope: A longitudinal dataset for investigating online discourse and information interaction on Telegram, *Proceedings of the International AAAI Conference on Web and Social Media*, 19(1), 2423-2433. DOI: 10.1609/icwsm.v19i1.35945.

Geiger, A., Liu, D., Alnegheimish, S., Cuesta-Infante, A., Veeramachaneni, K. (2020). TadGAN: Time series anomaly detection using generative adversarial networks, *2020 IEEE International Conference on Big Data*, 33-43. DOI: 10.1109/BigData50022.2020.9378139.

Guo, Y., Wang, D., Wang, L., Fang, Y., Wang, C., Yang, M., Liu, T., Wang, H. (2024). Beyond App Markets: Demystifying underground mobile app distribution via Telegram, *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, **8**(3), 1-25. DOI: 10.1145/3700432.

Hundman, K., Constantinou, V., Laporte, C., Colwell, I., Soderstrom, T. (2018). Detecting spacecraft anomalies using LSTMs and nonparametric dynamic thresholding, *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 387-395. DOI: 10.1145/3219819.3219845.

Hutchings, A., Holt, T. J. (2015). A crime script analysis of the online stolen data market, *British Journal of Criminology*, **55(**3), 596-614. DOI: 10.1093/bjc/azu096.

Kireev, K., Mykhno, Y., Troncoso, C., Overdorf, R. (2025). Characterizing and detecting propaganda-spreading accounts on Telegram, *Proceedings of the 34th USENIX Security Symposium*.

Krasznay, C. (2025). Hacktivists, proxy groups, cyber volunteers: The future of non-state actors' involvement in military cyber operations, *Academic and Applied Research in Military and Public Management Science*, **23(**3), 107-124. DOI: 10.32565/aarms.2024.3.6.

Le Sceller, Q., Karbab, E. B., Debbabi, M., Iqbal, F. (2017). SONAR: Automatic detection of cyber security events over the Twitter stream, *Proceedings of the 12th International Conference on Availability, Reliability and Security*, 23, 1-10. DOI: 10.1145/3098954.3098992.

Leukfeldt, E. R. (2024). Stolen data markets on Telegram: A crime script analysis and situational crime prevention measures, *Trends in Organized Crime*, **27**(1), 1-25. DOI: 10.1007/s12117-024-09532-6.

Liao, X., Yuan, K., Wang, X., Li, Z., Xing, L., Beyah, R. (2016). Acing the IOC game: Toward automatic discovery and analysis of open-source cyber threat intelligence, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 755-766. DOI: 10.1145/2976749.2978315.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach, arXiv preprint arXiv:1907.11692. DOI: 10.48550/arXiv.1907.11692.

MacAvaney, S., Yates, A., Feldman, P., Downey, D., Cohan, A., Goharian, N. (2019). Hate speech detection: Challenges and solutions, *PLOS ONE*, **14**(8), e0221152. DOI: 10.1371/journal.pone.0221152.

Pastrana, S., Thomas, D. R., Hutchings, A., Clayton, R. (2018). CrimeBB: Enabling cybercrime research on underground forums at scale, *Proceedings of the 2018 World Wide Web Conference*, 1845-1854. DOI: 10.1145/3178876.3186178.

Queiroz, A. L., Keegan, B., Mtenzi, F. J. (2017). Predicting software vulnerability using security discussion in social media, *Proceedings of the 16th European Conference on Cyber Warfare and Security*, 255-263. DOI: 10.21427/zgtj-nx67.

Rodriguez, A., Okamura, K. (2020). Enhancing data quality in real-time threat intelligence systems using machine learning, *Social Network Analysis and Mining*, 10, 1-22. DOI: 10.1007/s13278-020-00707-x.

Saeed, M. H., Huang, H. (2025). SENTINEL: A multi-modal early detection framework for emerging cyber threats using Telegram, arXiv preprint arXiv:2512.21380. DOI: 10.48550/arXiv.2512.21380.

Samtani, S., Chinn, R., Chen, H., Nunamaker, J. F. (2017). Exploring emerging hacker assets and key hackers for cyber threat intelligence, *Journal of Management Information Systems*, **34**(4), 1023-1053. DOI: 10.1080/07421222.2017.1394049.

Sen, M. A. (2024). Attention-GAN for anomaly detection: A cutting-edge approach to cybersecurity threat management, arXiv preprint arXiv:2402.15945. DOI: 10.48550/arXiv.2402.15945.

Tucci, G., Castro Gouveia, F. (2026). Detecting opinion leaders in a Telegram network of forwarded messages, *AoIR Selected Papers of Internet Research*. DOI: 10.5210/spir.v2024i0.15346.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models, *Advances in Neural Information Processing Systems*, 35, 24824-24837.

Zhao, J., Yan, Q., Li, J., Shao, M., He, Z., Li, L. (2020). TIMiner: Automatically extracting and analyzing categorized cyber threat intelligence from social data, *Computers & Security*, 95, 101867. DOI: 10.1016/j.cose.2020.101867.

Zong, S., Ritter, A., Mueller, G., Wright, E. (2019). Analyzing the perceived severity of cybersecurity threats reported on social media, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1, 1380-1390. DOI: 10.18653/v1/N19-1140.

Received February 3, 2026, revised March 11, 2026, accepted March 23, 2026