Kinetik: Game Technology, Information System, Computer Network, Computing, Electronics, and Control
Journal homepage: http://kinetik.umm.ac.id
ISSN: 2503-2267
Vol. 11, No. 3, August, Pp. 619-628

Predicting social media post engagement and virality using graph neural network approaches and content-based features

Fathimah Az Zahrah*[1], Riska Dhenabayu[1], Muhammad Fajar Wahyudi Rahman[1], Renny Sari Dewi[1], Zamabhungane Hadebe Aminah[2]
State University of Surabaya, Indonesia[1]
University of Johannesburg, South Africa[2]

Article Info

Keywords:
Engagement Regression, Graph Attention Network, GraphSAGE, Post Similarity Graph, Virality Screening

Article history:
Received: January 14, 2026
Accepted: May 16, 2026
Published: August 01, 2026

Cite:
F. Az Zahrah, Riska Dhenabayu, Muhammad Fajar Wahyudi Rahman, Renny Sari Dewi, and Zamabhungane Hadebe Aminah, "Predicting Social Media Post Engagement and Virality Using Graph Neural Network Approaches and Content-Based Features", *KINETIK*, vol. 11, no. 3, Aug. 2026. https://doi.org/10.22219/kinetik.v11i3.2686

*Corresponding author.
Fathimah Az Zahrah
E-mail address:
fathimah.22119@mhs.unesa.ac.id

Abstract

Social media teams increasingly rely on early signals to prioritize content, yet forecasting engagement and identifying viral posts remain difficult under temporal drift and heavy-tailed interaction counts. This study evaluated Graph Neural Network (GNN) approaches for predicting post engagement and virality from pre-posting content-based and contextual features. The Social Media Engagement Report dataset, which contained 100,000 posts across Twitter, LinkedIn, Facebook, and Instagram spanning March 2021–March 2024, was used. Post-release variables (impressions, reach, engagement rate) were excluded to prevent leakage. A homogeneous post–post graph was constructed using k-nearest-neighbor similarity in an embedding space and exact-match links on low-cardinality context. Ridge/Logistic Regression, Random Forest, and XGBoost as the baselines were compared against GraphSAGE and GAT under a chronological train, validation, and test split. Regression used MAE, RMSE, and R^2, while virality classification used ROC-AUC, PR-AUC, and Precision at the top 1% ranked posts. GraphSAGE yielded the strongest virality screening, achieving ROC-AUC = 0.66, PR-AUC = 0.54–0.56, and Precision@1% up to 0.75, substantially above non-graph baselines. For regression, GAT produced the lowest errors despite a negative R^2, indicating limited explained variance. Overall, similarity-graph GNNs are most effective for early virality identification, whereas exact count prediction remains challenging in a strictly pre-posting, time-aware setting. These findings contribute empirical evidence that, under leakage-safe pre-posting conditions, graph-based relational modeling is more useful for ranking and screening potentially viral posts than for precise engagement-count forecasting.

1. Introduction

Social media is a major channel for communication, commerce, and information dissemination, with more than five billion users worldwide and substantial daily usage [1]. Across digital media contexts, creators, public institutions, activists, and newsrooms often treat viral reach as a strategic route to attention and influence [2][3][4]. In practice, content that fails to achieve outsized reach is often deprioritized operationally, reinforcing a practical heuristic of "no viral, no attention" in digital strategy and resource allocation [4][5][6][7]. Engagement distributions are typically heavy-tailed, where a small share of posts accounts for most interactions [2]. Engagement is shaped by multiple factors, including content type, posting time, platform, audience demographics, and network effects. Recent studies of viral events on social media have emphasized that simple metrics, such as audience size or posting frequency, are poor predictors of outlier spread [2][8][9]. Furthermore, most viral posts do not significantly increase long-term engagement and audience growth, and their effects tend to decay quickly [2]. This complexity creates a practical prediction problem. Before post-release signals are observed, practitioners must estimate which posts are likely to achieve high engagement, the likely magnitude of engagement counts, and how limited advertising budget and cross-posting capacity should be prioritized [10]. A model that relies on post-release indicators would have limited value for planning because those indicators are only observed after publication and promotion decisions have already been made. Accordingly, the central research problem addressed in this study is how to predict engagement counts and high-engagement outcomes using only pre-posting features, while excluding post-release indicators and preserving realistic deployment through time-aware evaluation.

Prior studies have addressed social media popularity using supervised regression or classification on tabular features, often using Random Forests or XGBoost, which process each post individually and do not explicitly model relational structure among posts, users, or campaigns [11][12][13][14]. In reality, social media data are relational, where posts connect through topical similarity, shared audiences, and shared contextual patterns [15][16]. Consequently, the unresolved methodological problem is whether implicit post-level relations can be constructed from standard pre-posting tabular features and used to improve prediction beyond traditional machine learning models that treat each post as an isolated observation. Recent applications have shown that Graph Neural Networks, especially GraphSAGE and Graph Attention Networks (GAT), are effective for social network analysis, click-through rate prediction, and recommendation systems [17][18][19][20]. However, evidence from those domains does not yet establish whether similar benefits transfer to post-level social media engagement prediction when only pre-posting features are available and post-release feedback signals are excluded.

Despite these advances, three research gaps remain in the specific context of post-level social media engagement prediction. First, although GNNs are widely studied in recommender systems and broader social network analysis, fewer studies directly test whether a post–post graph constructed from content and posting metadata can improve pre-posting prediction of post-level engagement relative to strong tabular baselines [15][21]. Second, prior evaluations are often narrow, either focusing on a single task or a limited metric view, so there is still limited evidence on model behavior across both engagement-count prediction and high-engagement classification, including ranking-oriented metrics that are useful for practical post screening [18][22]. Third, even when temporal aspects are acknowledged, direct comparisons between graph-based models and strong tabular baselines under leakage-safe, time-aware evaluation remain limited, even though this setting is closer to operational deployment because models are trained on past posts and evaluated on future posts only. Consequently, it remains unclear whether GraphSAGE and GAT can provide predictive advantages over strong tabular baselines when applied to a post–post similarity graph derived only from pre-posting content and contextual metadata under leakage-safe, time-aware evaluation. Addressing these gaps offers empirical insight into whether graph construction adds practical value for engagement prediction, particularly when the task involves both high-engagement screening and exact engagement-count prediction.

This research addresses these gaps by clarifying the practical and empirical value of post–post similarity graphs for social media engagement prediction, specifically by examining whether graph-based relational modeling improves early virality screening and engagement-count prediction under leakage-safe, time-aware conditions. Specifically, this study constructs a homogeneous post–post similarity graph from pre-posting content and contextual features, then evaluates GraphSAGE and GAT against strong feature-based baselines under chronological train, validation, and test splitting. Accordingly, this study asks three research questions: (RQ1) to what extent can pre-posting content-based and contextual features support the prediction of engagement counts and high-engagement outcomes; (RQ2) do graph-based models built on a post–post similarity graph outperform strong feature-based baselines under leakage-safe, time-aware evaluation; and (RQ3) under this pre-posting setting, which prediction formulation is more practically informative for decision-making: high-engagement classification or engagement regression? The novelty of this study does not lie in proposing a new graph architecture. Rather, it lies in providing a unified leakage-safe and time-aware empirical evaluation of GraphSAGE and GAT on a post–post similarity graph for both engagement regression and high-engagement classification, alongside strong tabular baselines. This evaluation provides a diagnostic contribution by identifying which prediction formulation benefits most from relational modeling in realistic pre-posting conditions. This study contributes: (1) empirical evidence that post–post similarity graphs can improve early high-engagement screening compared with strong tabular baselines under leakage-safe temporal evaluation; (2) a task-level clarification that graph-based relational modeling is more beneficial for virality ranking than for exact engagement-count prediction; and (3) practical evidence that ranking-oriented classification metrics, particularly PR-AUC and Precision@k, provide more actionable guidance for social media decision-making than regression accuracy alone. These contributions aim to bridge methodological advances in graph neural networks with the practical needs of engagement-driven social media strategy.

This research is limited to the use of a publicly available dataset that aggregates post-level data from multiple platforms and countries. The study is constrained to pre-posting features such as content attributes, audience demographics, and posting metadata, excluding any post-release indicators such as impressions, reach, or engagement rate to prevent information leakage. The analysis focuses solely on predicting Likes, Comments, and Shares without incorporating sentiment text analysis, network propagation dynamics, or external event factors.

This paper is organized as follows. Section 2 describes the dataset, preprocessing, feature engineering, graph construction, experimental setup, and evaluation protocol. Section 3 reports the results for regression and virality classification and discusses the findings in relation to prior studies and practical implications. Section 4 concludes the study and outlines limitations and future research directions.

2. Research Method

2.1 Dataset Collection and Preprocessing

The Social Media Engagement Report dataset was used in this study [23]. The dataset contained 100,000 unique posts distributed across Twitter/X, LinkedIn, Facebook, and Instagram, covering 243 locations across seven continents, and spanning March 2021 to March 2024. It provided 24 features covering content attributes, engagement metrics, audience demographics, temporal information, campaign identifiers, and influencer identifiers. The dataset was partitioned with a chronological time-aware split into train, validation, and test by time to approximate deployment and avoid look-ahead [18], [24]. The dataset was sorted by 'Date' and then divided into three sets by time: training comprising the earliest data (2021-2022 or 66%); validation (2023 or 17%); and test (2024 or 17%). All data-dependent transforms, including label quantiles and scalers, were fit on the train only and then applied to validation/test.

Two prediction tasks were defined: (1) Engagement regression predicted the engagement counts as Likes, Comments, and Shares. Engagement counts are typically skewed and heavy-tailed; therefore, regression targets were modelled as log1p(y) during training and were transformed back or exponentiated with expm1 to compute evaluation metrics on the original scale during evaluation [2][25][26]; (2) Virality classification predicted whether a post is high engagement for each target (High_Likes, High_Comments, High_Shares) using a binary label based on the p60 quantile, meaning that the top 40% of posts were labelled as viral/high-engagement [27]. This threshold maintained sufficient positive samples while operationalizing virality as a high-engagement screening label rather than as an extreme diffusion event, consistent with percentile-based definitions of unusually high engagement within a given observation window [28].

Cyclic encodings were used for time-of-day and day-of-week to preserve cyclical proximity (e.g., 23:00 is close to 00:00). Hour-of-day and weekday indices were mapped into sine–cosine coordinates, enabling models to capture periodic temporal patterns without discontinuities at cycle boundaries [29]. Categorical low-cardinality features, such as Platform, Post Type, Weekday Type, Time Periods, Age Group, Audience Gender, Audience Continent, and Sentiment, were encoded with One-Hot Encoder (OHE) [26]. High-cardinality features, such as Post Content, Audience Interests, and Audience Location, were encoded using Feature Hashing (HashingVectorizer) followed by dimensionality reduction (TruncatedSVD) to form compact dense embeddings [16], [30]. Numerical features were standardized using the Standard Scaler Z-score normalization for calculating the mean and standard deviation of the training data [26]. Post–posting metrics, including Impressions, Reach, and Engagement Rate, were not used as input features in any model to avoid target leakage. Similarly, Likes, Comments, and Shares were used only as regression targets and as the basis of virality labels and were not included in the feature matrix. Identifier columns, such as Post ID, were excluded from model inputs to prevent memorization effects. Columns with extreme missingness, including Campaign ID and Influencer ID, were not used as model inputs. No additional imputation was applied beyond dataset-provided cleaning. All processed features were concatenated into a single feature matrix used by baselines and Graph Neural Networks.

2.2 Graph Construction

A homogeneous post–post graph was constructed from pre-posting representations, where each post was instantiated as a node. Edges were established based on two mechanisms: first, content-based similarity between posts through k-nearest neighbors in an embedding space, and second, exact-match links based on identical values on low-cardinality contextual features, e.g., same 'Platform', 'Post Type', or 'Weekday Type' [15][31][32]. All edges were merged, duplicates were removed, and the graph was converted to an undirected representation for message passing. The final graph was a single PyTorch Geometric (PyG) Data object with 100,000 nodes and 4,867,132 edges. Although several edge-construction rules were used, the graph was treated as homogeneous because all nodes represented posts. Figure 1 illustrates a local ego graph centered on one post node to show how a post is connected to its neighboring posts through the constructed similarity links. The visualization represents only a local neighborhood, not the entire graph. This graph structure enabled models to learn from both a post's own features and the features of its connected neighbors.

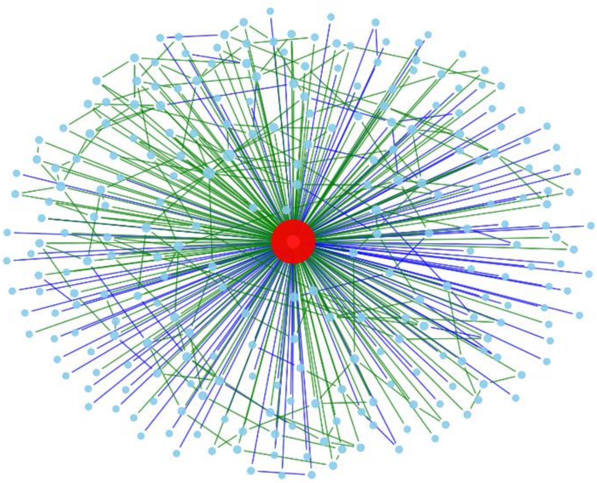

Figure 1. Ego Graph Visualization of the Post-post Similarity Graph Centered on Node 52499

2.3 Models

Baselines and GNNs were trained on the same feature vectors as illustrated in Figure 2. For the baselines, these models used tabular features and did not explicitly use graph structure. Regression used Ridge Regression, Random Forest Regressor (RF) with approximately 200 trees, and Extreme Gradient Boosting Regressor (XGBoost) with around 400 estimators, maximum depth of 8, learning rate of 0.05, subsample of 0.8, and column subsampling of 0.8. Additionally, Logistic Regression (LR), Random Forest Classifier (RF) with approximately 300 trees, and Extreme Gradient Boosting Classifier (XGBoost) with around 500 trees, depth 8, learning rate 0.05, and subsample 0.8 were employed for classification.

GraphSAGE (Graph Sample and Aggregate) is an inductive GNN that learns node embeddings by sampling and aggregating features from a node's local neighborhood. In this study, two SAGEConv layers with tuned hidden dimensionality were used. After each aggregation, ReLU activations and dropout were applied to improve stability and reduce overfitting. Formally, for a node i at a layer $k+1$,

$$h_i^{(k+1)} = \sigma\left(W^{(k)} \cdot AGG\left(\left\{h_i^{(k)}\right\} \cup \left\{h_j^{(k)}: j \in N(i)\right\}\right)\right) \tag{1}$$

Where $h_i^{(0)} = x_i$ is the input feature vector and $N(i)$ denotes the neighbours of the node i, $AGG(\cdot)$ is a mean aggregator, and σ is the ReLU activation, which is shown in Equation 1. The final node embedding was passed into two task-specific models: a binary classifier for virality logits (High_Likes, High_Comments, High_Shares) and a regressor for log-transformed engagement counts.

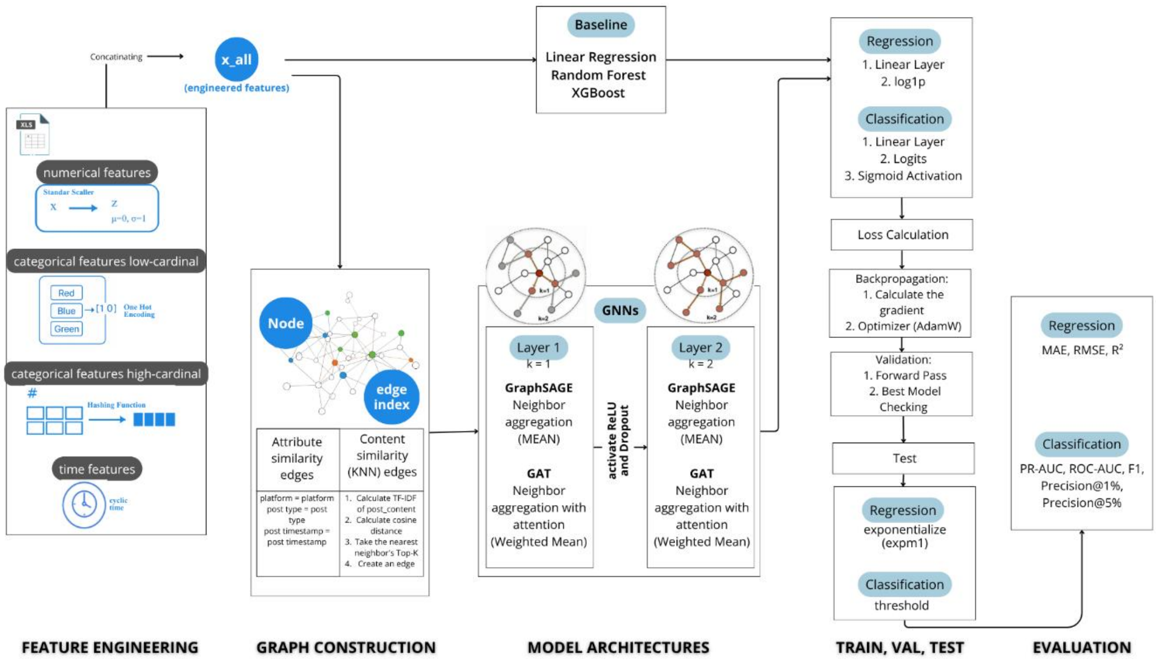

Figure 2. Framework of the Overall Proposed Model

GAT (Graph Attention Network) uses masked self-attention to weight neighbor contributions during aggregation. In this study, GATConv layers with 4 attention heads in the first layer, ELU activations, and dropout were used. For each node i and neighbour $j \in N(i)$, the attention coefficients were computed and were normalized over the neighborhood as shown in Equation 2. Aggregation step is given in Equation 3.

$$a_{ij} = softmax_j(LeakyReLU(a^T[Wh_i \| Wh_j] \tag{2}$$

$$h_j^{\prime} = \sigma(\sum_{j \in N(i)} a_{ij} W h_j) \tag{3}$$

The final node embeddings were passed to two linear heads: one for binary classification logits and one for regression on log-transformed engagement counts. All deep models used AdamW with gradient clipping (1.0) as defined in Equation 4. Neighbor sampling was applied to enable scalable mini-batch training by processing subsets of nodes and their neighborhoods.

$$AdamW = \theta_{t+1} = \theta_t - \alpha \cdot \frac{m_t}{\sqrt{v_t + \epsilon}} - \lambda \cdot \theta_t \tag{4}$$

Where m_t and v_t are the bias-corrected first and second moment estimates of the gradients, ϵ is a small constant for numerical stability, α is the learning rate, and λ is the weight decay.

For classification, Binary Cross-Entropy with Logits was used with a positive-class weight computed from the training distribution to address label imbalance. For regression, L1 loss (MAE) was applied on log-transformed targets for robustness to heavy-tailed outliers [33]. Early stopping was used based on the best validation performance, with the checkpoints selected using PR-AUC for classification and MAE for regression.

2.4 Evaluation Metrics

Model performance was assessed using standard metrics for each task. Regression performance was measured using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-Squared (R^2), computed on the original scale after inverse-transform with expm1. The regression metrics are defined in Equation 5 to 7 [11][12].

$$MAE = \frac{1}{N} \sum_{i=1}^{n} |y_i - \hat{y}_i| \tag{5}$$

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n} (\hat{y}_i - y_i)^2}{n}} \tag{6}$$

$$R^2 = 1 - \frac{Sum\ of\ Squares\ for\ Regression\ (SSR)}{Total\ Sum\ of\ Squares\ (SST)} \tag{7}$$

This definition of Equation 7 allows negative values when the model performs worse than the mean-based reference predictor. Classification performance was measured using Area Under the Precision-Recall Curve (PR-AUC), ROC-AUC, and Precision@k at the top 1% of predictions. PR-AUC and Precision@k were emphasized because practical decision-making often focuses on prioritizing a small subset of high-potential posts. Precision@k is defined in Equation 8, where Top_k denotes the set of posts ranked in the top k% by predicted confidence and $Z_i = 1$ if post i is truly high-engagement [18].

$$Precision@k = \frac{1}{[kN]} \sum_{i \in Top-k\%} z_i \tag{8}$$

3. Results and Discussion

3.1 Engagement Regression Performance

The regression task estimated the absolute counts of Likes, Comments, or Shares for each post. Table 1 reports the performance of each model. Across all three targets, GAT achieved the lowest MAE and RMSE, while GraphSAGE was competitive on Shares but did not consistently outperform Ridge on Comments and Likes. Compared with the strongest non-graph baseline, GAT reduced MAE by 4.02–6.50% and RMSE by 5.44–8.94% across the three engagement targets. However, the negative R^2 values across all models indicate that none of the regression models explained the out-of-sample variance in engagement counts better than a mean-based reference. This pattern is substantively important. It shows that precise engagement-count prediction remained weak even when GAT reduced MAE and RMSE. One plausible cause is that the leakage-safe feature set deliberately excluded post-release exposure indicators, namely Impressions, Reach, and Engagement Rate, while the remaining pre-posting features mainly represented content attributes, posting context, audience profile, and sentiment. These variables can support relative error reduction, but they may not fully capture unobserved exposure dynamics, platform ranking effects, or temporal shifts in audience response. In addition, the constructed post–post graph represents similarity among posts rather than an explicit diffusion, follower, or interaction network. Therefore, graph aggregation can smooth local signals and reduce average errors, but it cannot fully explain the variance of engagement counts. For this reason, the regression findings should be interpreted as relative error improvements rather than reliable variance explanation.

Table 1. Engagement Regression Test Result

Target	Model	MAE	RMSE	R^2
Comments	GAT	125.93	145.73	-0.016
Comments	GraphSAGE	134.05	159.02	-0.209
Comments	Ridge	133.55	158.36	-0.208
Comments	XGBoost	135.60	161.55	-0.257
Comments	RF	138.61	166.27	-0.332
Likes	GAT	257.33	301.02	-0.084
Likes	GraphSAGE	275.18	329.48	-0.299
Likes	Ridge	268.12	318.33	-0.209
Likes	XGBoost	271.04	323.21	-0.246
Likes	RF	278.58	334.05	-0.331
Shares	GAT	50.17	58.00	-0.003
Shares	GraphSAGE	52.35	61.64	-0.132
Shares	Ridge	53.66	63.69	-0.203
Shares	XGBoost	54.46	64.88	-0.248
Shares	RF	55.83	66.90	-0.327

Figure 3 provides a normalized comparison across targets and metrics, showing that GAT occupied the most favorable relative position in the regression setting across Comments, Likes, and Shares. This pattern indicates a consistent relative advantage in error reduction of attention-based neighborhood aggregation in the current post–post similarity graph setting. Overall, the regression results indicate that while graph-based models improved robustness in error terms, precise numerical engagement prediction remained challenging under a multi-platform dataset, a strictly pre-posting feature setting, and a chronological evaluation [2][8].

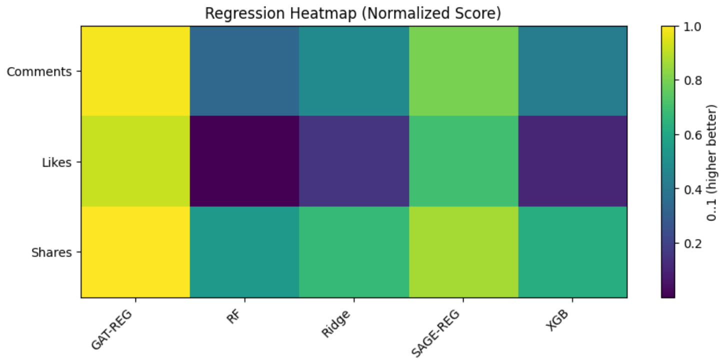

Figure 3. Regression Heatmap Based on Normalized Model Performance

3.2 Virality Classification Performance

The classification tasks predicted whether a post belonged to the high-engagement class for each target. Table 2 reports metrics, including ROC-AUC, PR-AUC, and precision at the top 1% of posts ranked by model confidence. Across all tasks, GraphSAGE consistently achieved the best classification performance for High_Likes, High_Comments, and High_Shares, with ROC-AUC around 0.66 and PR-AUC around 0.54–0.56. In addition, top-1% precision reached 0.66–0.75, indicating that the small subset of prioritized posts contained a substantially higher proportion of true high-engagement cases than the baseline-ranked subsets. Compared with the best non-graph baseline, GraphSAGE improved PR-AUC by 0.132 for High_Comments, 0.146 for High_Likes, and 0.155 for High_Shares. The gain was larger at the top-ranked decision point, where Precision@1% improved by 0.230, 0.317, and 0.319 for the same targets. This top-ranked improvement is practically important because promotion resources are usually allocated to a limited number of posts. In contrast, the baseline classifiers remained close to random ranking on ROC-AUC around 0.49–0.50 and achieved lower PR-AUC around 0.40. GAT performed only slightly above or close to the strongest baseline models, which suggests that attention-based aggregation did not extract additional discriminative signal from the constructed similarity graph. A likely explanation is that the graph connected posts based on similarity in content and contextual metadata, rather than on explicit social propagation or interaction pathways. Under this type of homogeneous post–post similarity graph, GraphSAGE may benefit from stable neighborhood smoothing, whereas GAT may be less effective because learned attention weights have limited relational variation to exploit. This difference helps explain why GraphSAGE was more suitable for high-engagement classification, while GAT was more useful for reducing regression errors.

Table 2. Virality Classification Test Result

Target	Model	ROC-AUC	PR-AUC	Prec@1%
High_Comments	GraphSAGE	0.658	0.537	0.656
High_Comments	GAT	0.509	0.408	0.442
High_Comments	LogReg	0.504	0.405	0.426
High_Comments	RF	0.503	0.399	0.367
High_Comments	XGBoost	0.495	0.395	0.349
High_Likes	GraphSAGE	0.662	0.547	0.713
High_Likes	GAT	0.508	0.407	0.425
High_Likes	LogReg	0.493	0.400	0.349
High_Likes	RF	0.495	0.401	0.385
High_Likes	XGBoost	0.495	0.401	0.396
High_Shares	GraphSAGE	0.665	0.557	0.751
High_Shares	GAT	0.509	0.410	0.423
High_Shares	RF	0.494	0.402	0.402
High_Shares	LogReg	0.492	0.400	0.432
High_Shares	XGBoost	0.492	0.400	0.408

Figure 4 summarizes the normalized classification comparison across targets and models, making the cross-target pattern easier to compare. The heatmap supports the conclusion that the benefit of graph-based learning is task-dependent: GraphSAGE provides consistent gains for high-engagement classification, whereas GAT and feature-only baselines remain clustered at lower relative performance levels. Thus, the classification visualization supports the practical conclusion that graph-based neighborhood aggregation is more actionable for screening potentially high-engagement posts than for exact count forecasting.

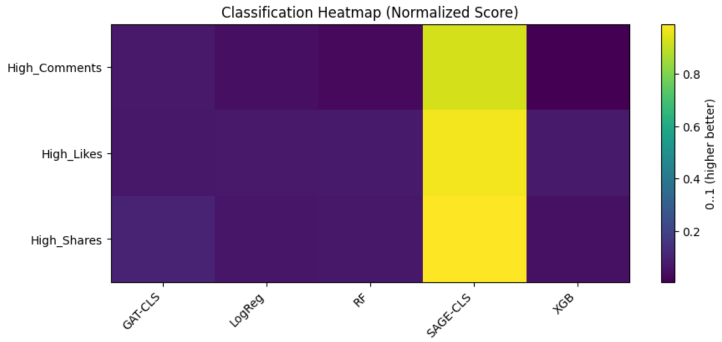

Figure 4. Classification Heatmap Based on Normalized Model Performance

3.3 Interpretation and Practical Implications

The empirical findings can be interpreted with the central question of whether graph neural networks, using relations derived from content similarity and low-cardinality contextual matching, improved engagement and virality prediction beyond feature-only baselines. The results suggest that the answer depends on the task formulation. For high-engagement classification, the relational signal was useful because the objective was to rank or separate posts into higher- and lower-engagement groups. This formulation is more tolerant of noisy count variation and can benefit from neighborhood smoothing among similar posts. For regression, however, the model had to estimate exact engagement counts, which are more sensitive to unobserved exposure, platform recommendation dynamics, and temporal changes. This explains why graph-based learning improved MAE and RMSE but still failed to produce positive R^2 values. This pattern suggests that the post–post similarity graph captured enough relational structure to support ranking-based screening, but not enough information to explain full engagement-count variance. Practically, the classification outputs were therefore more actionable for content prioritization, such as for selecting a small top-ranked subset for promotion, while regression outputs were better interpreted as coarse indicators rather than precise forecasts [14], [22].

These findings align with and extend recent studies in social media engagement prediction. Recent research evaluating tabular machine learning models on pre-posting metadata consistently highlights the difficulty of explaining variance for highly skewed interactions, often requiring the extraction of complex multimodal or visual features to achieve robust point forecasting [13], [22]. The regression results are consistent with this inherent limitation of tabular metadata. However, while recent advanced graph-based applications typically rely on explicit information diffusion networks, such as temporal-spatial cascade graphs, to predict virality, the proposed approach suggests that an implicitly constructed post-similarity graph can reduce some limitations of feature-only modeling for ranking-based screening, even without explicit user-to-user sharing dynamics [15], [21]. Specifically, GraphSAGE achieved ROC-AUC of approximately 0.66 and Precision@1% up to approximately 0.75, indicating actionable ranking gains relative to the feature-only baselines evaluated in this study. Because prior studies often use different platforms, label definitions, feature sets, and evaluation windows, the comparison with earlier work should be interpreted at the level of task behavior rather than as a direct numerical benchmark. This is consistent with recent findings that framing engagement prediction as a ranking-oriented classification task rather than exact count regression provides a more robust and practical strategy for digital marketing prioritization [27].

Several limitations constrained generalization. The evaluation used a single dataset and a similarity-based constructed graph rather than an explicit diffusion or follower network, limiting direct modeling of propagation dynamics. In addition, no explainability module was included in the final pipeline. These limitations motivate extensions in future work, including temporal/dynamic graph modeling, richer content representations, and robustness analysis of graph construction choices, particularly sensitivity to edge density and similarity thresholds.

4. Conclusion

This study evaluated the prediction of social media post engagement and virality using graph neural networks combined with content-based and contextual features under a time-aware evaluation protocol. Addressing the core research questions, the empirical findings demonstrate that strictly pre-posting content and contextual features are insufficient to reliably explain the variance of exact engagement counts, yet they contain useful relational signals for identifying high-engagement outcomes. Similarity-graph modeling delivered the clearest benefit for this virality classification. Specifically, GraphSAGE consistently improved ranking-based classification performance across High_Likes, High_Comments, and High_Shares targets, effectively prioritizing high-potential posts where traditional baselines underperformed. Conversely, while GAT achieved the lowest absolute error for engagement regression, its R-squared values remained non-positive. Consequently, this study concludes that high-engagement classification represents a substantially more robust and practically informative formulation for decision-making compared to exact count prediction, which remains intrinsically difficult given heavy-tailed distributions and the deliberate exclusion of post-release signals.

The main contribution of this study is the empirical clarification that post–post similarity graphs are more valuable for early high-engagement screening than for exact engagement-count forecasting under leakage-safe and time-aware evaluation. Future work should test generalization across additional datasets and domains, examine temporal or dynamic graph formulations to better capture evolving engagement patterns, and evaluate alternative graph construction strategies to characterize when graph-based modeling provides the largest gains.

References

[1] we are social, “Digital 2025 Global Overview Report.” Accessed: Dec. 15, 2025.
[2] E. Sangiorgio, N. Di Marco, G. Etta, M. Cinelli, R. Cerqueti, and W. Quattrociocchi, “Evaluating the effect of viral posts on social media engagement,” *Scientific Reports 2025 15:1*, vol. 15, no. 1, pp. 639-, Jan. 2025. https://doi.org/10.1038/s41598-024-84960-6
[3] D. H. Kim, O. Kuru, J. Zeng, and S. Kim, “Fostering mask-wearing with virality metrics and social media literacy: evidence from the U.S. and Korea,” *Front. Psychol.*, vol. 14, p. 1151061, 2023. https://doi.org/10.3389/FPSYG.2023.1151061
[4] N. Jalli, “Viral Justice: TikTok Activism, Misinformation, and the Fight for Social Change in Southeast Asia,” *Social Media and Society*, vol. 11, no. 1, Jan. 2025. https://doi.org/10.1177/20563051251318122
[5] M. Heitmayer, “The Second Wave of Attention Economics. Attention as a Universal Symbolic Currency on Social Media and beyond,” *Interact. Comput.*, vol. 37, no. 1, pp. 18–29, Jan. 2025. https://doi.org/10.1093/IWC/IWAE035
[6] M. Tahmina Khanom, “Using social media marketing in the digital era: A necessity or a choice,” *International Journal of Research in Business and Social Science (2147- 4478)*, vol. 12, no. 3, pp. 88–98, May 2023. https://doi.org/10.20525/IJRBS.V12I3.2507
[7] S. Ren, C. Gong, C. Zhang, and C. Li, “Public opinion communication mechanism of public health emergencies in Weibo: take the COVID-19 epidemic as an example,” *Front. Public Health*, vol. 11, p. 1276083, Nov. 2023. https://doi.org/10.3389/FPUBH.2023.1276083
[8] A. Iamnitchi, L. O. Hall, S. Horawalavithana, F. Mubang, K. W. Ng, and J. Skvoretz, “Modeling information diffusion in social media: data-driven observations,” *Front. Big Data*, vol. 6, p. 1135191, May 2023. https://doi.org/10.3389/FDATA.2023.1135191
[9] E. Sangiorgio, M. Cinelli, R. Cerqueti, and W. Quattrociocchi, “Followers do not dictate the virality of news outlets on social media,” *PNAS Nexus*, vol. 3, no. 7, Jun. 2024. https://doi.org/10.1093/PNASNEXUS/PGAE257
[10] H. Metzler and D. Garcia, “Social Drivers and Algorithmic Mechanisms on Digital Media,” *Perspectives on Psychological Science*, vol. 19, no. 5, pp. 735–748, Sep. 2024. https://doi.org/10.1177/17456916231185057
[11] V. Gupta, K. Jung, and S. C. Yoo, “Exploring the Power of Multimodal Features for Predicting the Popularity of Social Media Image in a Tourist Destination,” *Multimodal Technologies and Interaction 2020, Vol. 4,* vol. 4, no. 3, pp. 1–23, Sep. 2020. https://doi.org/10.3390/MTI4030064
[12] K. R. Purba, D. Asirvatham, and R. K. Murugesan, “An analysis and prediction model of outsiders percentage as a new popularity metric on Instagram,” *ICT Express*, vol. 6, no. 3, pp. 243–248, Sep. 2020. https://doi.org/10.1016/J.ICTE.2020.07.001
[13] D. Jeong, H. Son, Y. Choi, and K. Kim, “Enhancing social media post popularity prediction with visual content,” *Journal of the Korean Statistical Society 2024 53:3*, vol. 53, no. 3, pp. 844–882, May 2024. https://doi.org/10.1007/S42952-024-00270-7
[14] S. Carta *et al.*, “Popularity Prediction of Instagram Posts,” *Information 2020, Vol. 11,* vol. 11, no. 9, Sep. 2020. https://doi.org/10.3390/INFO11090453
[15] Z. Xu, M. Qian, Z. Xu, and M. Qian, “Predicting Popularity of Viral Content in Social Media through a Temporal-Spatial Cascade Convolutional Learning Framework,” *Mathematics 2023, Vol. 11,* vol. 11, no. 14, Jul. 2023. https://doi.org/10.3390/MATH11143059

[16] M. K. Baxi, R. Sharma, and V. Mago, “Studying topic engagement and synergy among candidates for 2020 US Elections,” *Soc. Netw. Anal. Min.*, vol. 12, no. 1, p. 136, Dec. 2022. https://doi.org/10.1007/S13278-022-00959-9
[17] J. Song *et al.*, “Graph Representation-Based Deep Multi-View Semantic Similarity Learning Model for Recommendation,” *Future Internet 2022, Vol. 14,* vol. 14, no. 2, Jan. 2022. https://doi.org/10.3390/FI14020032
[18] J. Zhu and A. Yaseen, “A Recommender for Research Collaborators Using Graph Neural Networks,” *Front. Artif. Intell.*, vol. 5, p. 881704, Aug. 2022. https://doi.org/10.3389/FRAI.2022.881704
[19] Q. Tong, X. Xu, J. Zhang, and H. Xu, “Public Opinion Propagation Prediction Model Based on Dynamic Time-Weighted Rényi Entropy and Graph Neural Network,” *Entropy 2025, Vol. 27,* vol. 27, no. 5, May 2025. https://doi.org/10.3390/E27050516
[20] A. Golovin *et al.*, “Improving Recommender Systems for Fake News Detection in Social Networks with Knowledge Graphs and Graph Attention Networks,” *Mathematics 2025, Vol. 13,* vol. 13, no. 6, Mar. 2025. https://doi.org/10.3390/MATH13061011
[21] Y. Shang *et al.*, “Popularity Prediction of Online Contents via Cascade Graph and Temporal Information,” *Axioms 2021, Vol. 10,* vol. 10, no. 3, Jul. 2021. https://doi.org/10.3390/AXIOMS10030159
[22] P. K. Theodoridis, D. C. Gkikas, P. K. Theodoridis, and D. C. Gkikas, “Maximizing Social Media User Engagement Through Predictive Analytics in Retail Tourism: Identifying Key Performance Indicators That Trigger User Interactions,” *Applied Sciences 2025, Vol. 15,* vol. 15, no. 21, Nov. 2025. https://doi.org/10.3390/APP152111720
[23] “Social Media Engagement Report | Kaggle.” Accessed: Dec. 27, 2025.
[24] H. M. Lin and J. J. Lyu, “A holistic framework for intradialytic hypotension prediction using generative adversarial networks-based data balancing,” *BMC Medical Informatics and Decision Making 2025 25:1*, vol. 25, no. 1, pp. 257-, Jul. 2025. https://doi.org/10.1186/S12911-025-03094-5
[25] “Effect of transforming the targets in regression model — scikit-learn 1.8.0 documentation.” Accessed: Dec. 27, 2025.
[26] P. Koukaras and C. Tjortjis, “Data Preprocessing and Feature Engineering for Data Mining: Techniques, Tools, and Best Practices,” *AI 2025, Vol. 6, Page 257*, vol. 6, no. 10, p. 257, Oct. 2025. https://doi.org/10.3390/AI6100257
[27] P. P. Tricomi, M. Chilese, M. Conti, and A. R. Sadeghi, “Follow Us and Become Famous! Insights and Guidelines From Instagram Engagement Mechanisms,” *Web Science Conference*, pp. 346–356, Apr. 2023. https://doi.org/10.1145/3578503.3583623
[28] T. Sah and K. Jordan, “Decoding reddit memes virality,” *International Journal of Data Science and Analytics 2025 20:6*, vol. 20, no. 6, pp. 5321–5336, Apr. 2025. https://doi.org/10.1007/S41060-025-00772-5
[29] G. Marchesi, A. Ballarino, A. Brusaferri, G. Marchesi, A. Ballarino, and A. Brusaferri, “Assessing Time Series Foundation Models for Probabilistic Electricity Price Forecasting: Toward a Unified Benchmark,” *Energies 2025, Vol. 18,* vol. 18, no. 23, Nov. 2025. https://doi.org/10.3390/EN18236269
[30] N. Husin, H. Fazlurrahman, A. Safitri, R. Dhenabayu, U. A. A. Rauf, and A. M. Fitrah, “Policy Perspective on Proposed Framework of NLP AI to Bridge the Inclusive Support in Higher Education with a Mixed Methods Approach in Indonesia and Malaysia,” *International Journal of Information and Education Technology*, vol. 15, no. 12, pp. 2686–2699, 2025. https://doi.org/10.18178/ijiet.2025.15.12.2464
[31] R. Yang, A. Yu, L. Cai, and D. Meng, “Subspace clustering via graph auto-encoder network for unknown encrypted traffic recognition,” *Cybersecurity 2022 5:1*, vol. 5, no. 1, pp. 29-, Dec. 2022. https://doi.org/10.1186/S42400-022-00131-Y
[32] X. Zhang, B. Hu, S. Liu, Q. Sun, and L. Chen, “AttenFlow: Context-Aware Architecture with Consensus-Based Retrieval and Graph Attention for Automated Document Processing,” *Applied Sciences 2025, Vol. 15, Page 7517*, vol. 15, no. 13, p. 7517, Jul. 2025. https://doi.org/10.3390/APP15137517
[33] N. K. Nissa, R. T. Pusparini, A. Setiyoko, and A. M. Arymurthy, “The Implementation of Inductive Graph Neural Networks with L1 Loss for Spatiotemporal Kriging,” *AIP Conf. Proc.*, vol. 2941, no. 1, Dec. 2023. https://doi.org/10.1063/5.0184738/2929312