Research on the Collaborative Optimization of Creative Communication Models in Cross-Platform Media Ecosystems

Yiran Wang
School of Journalism and Communication, Henan University of Technology, Zhengzhou 450001, Henan, China
Corresponding author: Yiran Wang (e-mail: 15188306767@163.com)

ABSTRACT In the face of cross-platform media access in the cross-platform media environment, the problem of algorithm-pushed creative diffusion appears to be confronted with the dilemma of content fragmentation and overlap: the same platform experiences semantic fragmentation due to separate optimization on different platforms, and different user behaviors make it hard to find a single strategy to gain cross-media synergy. Firstly, this study builds a heterogeneous information network representation framework, and contents, interactions and recommendations are represented as transferable temporal graphs; Secondly, an adversarial domain adaptation strategy is introduced to generate platform-aware variants while preserving fundamental symbols to balance semantic consistency and form localization; Finally, a reinforcement learning scheduling mechanism is designed to regulate distribution sequence and resource allocation in real time according to user overlap and decay period. Experiments show that, compared with before, the mean crossdomain semantic consistency is increased by 0.86 - 0.62, duplicated reach is improved by 22.4%, and long tail discussion half-life is prolonged from 7.3 days to 15.4 days. These results confirm that hierarchical collaborative optimization could bridge the gap of dissemination discontinuity and effectiveness diffusion.

INDEX TERMS cross-platform media ecosystem, creative communication, heterogeneous graph network, reinforcement learning scheduling, network diffusion

I. INTRODUCTION

THE evolving cross-platform media ecosystem has made the dissemination path of creative content increasingly complex. Short video, text-based communities, and instant messaging platforms have each built independent algorithmic distribution systems and user behavior loops, leading to a dual dilemma for the same creative material during cross-platform circulation: on the one hand, the differentiated interpretations of content semantics by various platform recommendation engines force creative symbols to be distorted during adaptation, making it difficult for cross-screen users to form a unified brand perception; on the other hand, fragmented ad placement scheduling systems lack a global awareness of users' cross-media migration behavior, generally resulting in an efficiency gap of overexposure for high-frequency users and underreach for low-frequency users [1-3]. The aforementioned contradiction between content fragmentation and overlapping audience reach constitutes a core obstacle to cross-domain collaboration in creative dissemination, urgently requiring a systematic solution from the perspective of dissemination models.

Existing research on cross-platform content distribution mainly focuses on two directions: user cross-screen identification and channel combination optimization. For example, Unnava and Aravindakshan[4] found through analysis of brand release behavior on multiple platforms that there are significant interactive spillover effects and temporal dependencies between different social media platforms; Liu et al.[5] verified the advantages of multi-channel collaborative placement in improving reach efficiency and conversion effect from the perspective of cross-platform targeted advertising. Hase et al.[6] pointed out that the differences in content modality and expression paradigm of different platforms will significantly affect the information presentation method and user response mechanism; Nielsen and Fletcher[7] further revealed the structural reconstruction of the news dissemination system in the process of platformization from a macro perspective.

However, existing results mostly focus on local optimization of a single link, lacking a global collaborative consideration of creative semantic fidelity, variant adaptation generation and dissemination timing control. In recent years, heterogeneous information network modeling and adversarial learning technology have shown good adaptability in cross-domain transfer tasks. For example, at the level of behavioral representation, Yan et al. [8] constructed user interest representations based on heterogeneous information networks and graph neural networks, effectively integrating multi-source heterogeneous data; in terms of cross-domain migration, HassanPour Zonoozi and Seydi [9] systematically reviewed the development of adversarial domain adaptation methods, pointing out that feature alignment can be achieved between different data distributions through adversarial learning mechanisms; furthermore, Shi et al. [10] summarized the application challenges and development trends of domain adaptation in complex networks from the perspective of graph representation learning, emphasizing the key role of cross-domain representation consistency in multi-platform environments. The above work provides technical components that can be referenced for cross-platform collaborative communication, but no research has yet incorporated representation construction, adaptive adaptation, and collaborative scheduling into the same decision-making link. Based on this, this paper introduces a strategy that combines heterogeneous representation and adversarial scheduling to fill the gap in this integrated framework.

This paper seeks to establish a collaborative optimization framework for cross-platform media consumption oriented creative communication. Firstly, we construct uniform representation layer for multi-platform dissemination behavior with temporally heterogeneous graph network, mapping creative content, user interaction sequence, and platform distribution rule to transferable embedding space via a modality-specific projection layer that aligns discrete administrative parameters with high-dimensional semantic vectors using a cross-attention alignment mechanism. Secondly, we design variant adaptive generation mechanism with adversarial domain adaptation and semantic anchoring to output platform-aware expression while maintaining basic symbolic feature of creative content. Finally, we establish reinforcement learning scheduling model oriented to marginal gain of overlapping audience to dynamically adjust the delivery moment and resource allocation over platforms. To present the overall structure of cross-platform media consumption oriented creative communication collaborative optimization framework and collaborative relationship between different modules in the framework more clearly, Fig. 1 shows the schematic diagram of overall model architecture.

II. A UNIFIED CROSS-DOMAIN BEHAVIOR REPRESENTATION LAYER BASED ON TEMPORALLY HETEROGENEOUS GRAPH NETWORKS

User interaction logs from various media channels in cross-platform dissemination present high sparsity, large data structure gap, media-specific field definitions, and different update frequencies, which hinder the direct extension of these user behaviors into one quantitative analysis dataset.To address this challenge, we build a temporally heterogeneous graph network for creative dissemination in this section, which can connect user behavioral records from fragmented platforms into one transferable vector space.

The graph network contains three types of nodes: creative content nodes, user device nodes and platform context nodes. The creative content nodes hold the visual and textual feature embeddings of the material, which are extracted by pretrained CLIP encoder and finally mapped into 128 dimensions by principal component analysis. The user device nodes use salted hashed device fingerprint as unique identifier to connect users' historical interaction sequence from different platforms. The platform context nodes record the parameters of recommendation mechanism. To resolve the dimensional mismatch between discrete rules and semantic embeddings, we employ a multi-layer perceptron (MLP) to project categorical bidding strategies and traffic tags into a continuous latent space of the same dimension as the content embeddings. The three types of nodes are connected by three kinds of directed edges: the weight of edge between creative and user is the exposure time and rate of completion and type of interaction behavior; the weight of edge between user and platform is the activity level of user and his/her historical response probability on this platform; the weight of edge between creative and platform is the competition strength of material under current recommendation strategy of this platform.

Considering the temporal dependence of user cross-screen behavior, the above graph structure is divided into continuous snapshots according to the time window, and adjacent snapshots are connected by the state transition edge of the user node to maintain temporal correlation. Specifically, the day is divided into four time periods with a granularity of six hours. While the graph snapshots capture long-term state transitions, the Reinforcement Learning layer in Section 4 operates on a dual-scale mechanism, incorporating minute-level response signals to adjust budget weights within each 6-hour window to align with the real-time volatility of digital media attention. The interaction events that occur in each time period constitute a static subgraph. The same user node across time periods is connected by recursive edges, and the edge attribute is whether the user has performed platform migration behavior in that time period. The resulting temporal heterogeneous graph not only retains the interaction topology within a single platform, but also depicts the dynamic trajectory of user cross-screen migration [11].

In representation learning stage, following hierarchical aggregation strategy with meta-path guidance is conducted. For "creative-user-creative" meta-path, attention-based homogeneous neighbor aggregation is applied to learn content preference pattern of similar audience groups. For "user-platform-user" meta-path, behavioral similarity of user co-occurring between different platforms is aggregated to describe structural features of audience overlap between platforms. Update process of node embedding will introduce gated recurrent unit to model state evolution between two

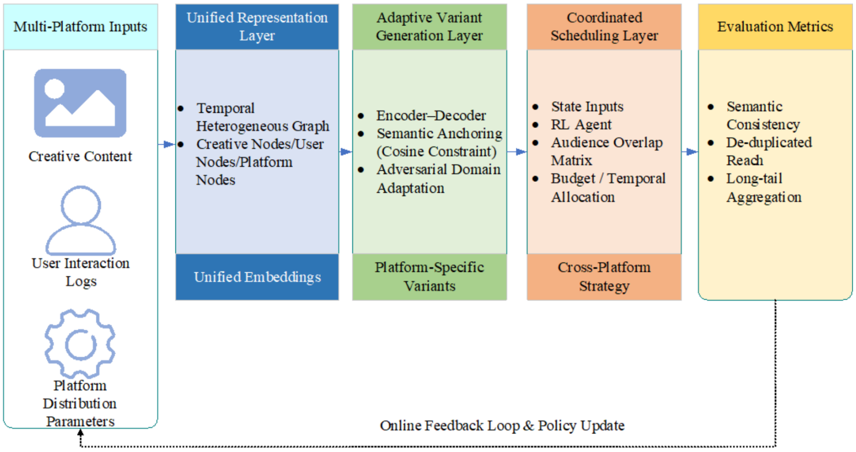

FIGURE 1. Overall Structure of the Cross-Platform Creative Communication Collaborative Optimization Framework

temporal snapshots. It is guaranteed that final output node vector is the multiplication of aggregation effect and recent behavior of dissemination.

III. AN ADAPTIVE VARIANT GENERATION LAYER INTEGRATING ADVERSARIAL DOMAIN ADAPTATION AND SEMANTIC ANCHORING

The unified representation layer constructed in the previous section solved the structuring problem of crossplatform behavioral data. However, a key contradiction remains when creative content migrates from the source platform to the target platform: directly reusing the original material often leads to decreased distribution efficiency due to incompatibility with the target platform's stylistics and recommendation logic; allowing each platform to independently adapt and optimize can easily cause semantic drift of the core creative symbols. Therefore, this section designs a variant generation mechanism that balances semantic fidelity and platform adaptation.

The input to this layer is the feature embedding vector of the source creative material, and the output is the variant feature representation adapted to the target platform. The overall architecture consists of three interconnected modules: a variant generator, a semantic anchoring constraint module, and a platform domain discriminator.

The variant generator adopts an encoder-decoder structure. The encoder reuses the creative content node embeddings output from the representation layer, while the decoder sets independent transformation matrices for different target platforms, mapping the general semantic vector to the platform-specific expression space. Taking short video platforms as an example, the variant vectors output by the decoder strengthen the weights of high saturation and fast-paced features in the visual dimension, and focus on the density of colloquial short sentences and topic tags in the text dimension; the transformation matrix of the image and text community platform improves the integrity of the composition and the structuring degree of long text descriptions [12].

The core function of the semantic anchoring constraint module is to prevent the generator from excessively catering to platform preferences and losing the original recognition features of creativity during the adaptation process. Specifically, a cosine similarity constraint term is added to the output of the generator to calculate the cosine value of the angle between the target variant vector and the source creative anchor vector. During training, this constraint term is added to the overall optimization objective as a regularization loss, forcing the cosine similarity to always remain above the preset lower bound. The source creative anchor vector is determined during model initialization, obtained by taking the mean of the semantic embedding after CLIP encoding of the source material, and remains fixed in subsequent iterations as an absolute reference for semantic consistency. The platform domain discriminator is used to drive the generator to learn the platform-specific expression style. The discriminator is a binary classification feedforward network that takes the variant vector as input and outputs the predicted probability of the platform to which the vector belongs. The generator and discriminator are connected via a gradient inversion layer: during forward propagation, gradients are transmitted normally; during backpropagation, the gradients of the discriminator are multiplied by a negative coefficient before being fed back to the generator. This adversarial training mechanism forces the variant vectors to absorb target stylistic features. To prevent stylistic drift from compromising core identity, the semantic anchoring constraint is assigned a higher optimization weight ($\omega = 2.5$) than the adversarial loss. This hierarchy ensures that the variant generator prioritizes the preservation of the source creative's invariant symbolic features even when the platform discriminator is highly confident. That is, the variants fully absorb the stylistic features of the target platform while maintaining cross-domain cognitive stability within the boundaries of semantic anchoring constraints.

We alternate in training the three modules. When updates the discriminator, the generator parameters are fixed and set to maximize the classification accuracy of platforms; When updates the generator, the following three losses will be jointly minimized: distribution difference loss between real samples of variant and target platform, semantic anchoring cosine constraint loss, adversarial loss after inverting the gradient of discriminator. After the training process converges, the generator is able to directly output the adapted variant to fit the distribution requirement of target platform, while not deviating from the core semantics of creative, based on the material and target platform identifier.

IV. REINFORCEMENT LEARNING COLLABORATIVE SCHEDULING LAYER FOR MARGINAL GAINS OF OVERLAPPING AUDIENCES

After adaptive creative variants are generated, the timing to distribute them and how to allocate resources across platforms are still open questions. If each platform independently executes its delivery plan, although it will tend to a locally optimal solution, it results in waste of budget and marginal efficiency decrease due to ignoring the audience overlap across screens.

In this section, we model the cross-platform collaborative scheduling problem as a sequential decision-making process and adopt a reinforcement learning method to dynamically update the delivery action, and the marginal gain brought by audience overlap as the optimization objective.

The essence of the scheduling problem is to balance two mutually constraining efficiency requirements: more scale of deduplicated reach, and less repeated exposure to highly overlapping users. To quantify the relationship between them, we first compute the audience overlap matrix between any pair of platforms based on user node embeddings output from representation layer. For platform A and platform B, the overlap is the proportion of device fingerprints appearing in the active user set of platform A and in the active user set of platform B at the same time to the total number of device fingerprints appearing in the active user set of platform A and in the active user set of platform B. The audience overlap matrix is updated using a 24- hour sliding window with hourly incremental refreshes and acts as dynamic background knowledge input for scheduling process. This high-frequency update mechanism allows the agent to synchronize with rapid user migration patterns and shifting media "hot spots" in real-time.

The state space of reinforcement learning agent is formed by concatenating together three types of information. The first type of information represents the cumulative exposure and unique reach of each creative variant across platforms in current time step. It reflects the absolute level of dissemination process. The second type of information represents the traffic peak index and estimated competition intensity of each platform in the past three time windows. It represents the changing trend of environment outside. The third type of information represents the actual execution of campaign action in previous time step and the real-time feedback of response rate of each platform. These three types of information are normalized to form a fixed-dimensional state vector.

The action space is defined as a continuous value vector. The dimension of vector is the number of platforms participating in scheduling. Each dimension of vector takes values in the range [0, 1], representing the proportion of budget weight allocated to the corresponding platform in the next time step. The action vector is executed after being softmax normalized to make the sum of weight of each platform equal to 1. The total budget will be divided into equal units by time period and the actual campaign amount in each time period is determined by the product of weight vector and the budget of corresponding time period.

The design of reward function adopts marginal gain orientation. After each time step, the agent will get an immediate reward signal. The immediate reward is calculated as the number of new users deduplicated in that step plus a frequency-weighted bonus for users within the "effective exposure window" (2-3 times), minus the number of times of excessive repeated exposure penalty. Following the "Three-Exposure Hypothesis," the penalty coefficient is only activated when a user's cross-platform cumulative contact frequency exceeds the saturation threshold, thereby balancing reach expansion with the necessary repetition for cognitive reinforcement. The penalty coefficient is positively correlated with the corresponding value in above mentioned audience overlap matrix. The more overlap, the more cost of repeated exposure. In addition, if the cumulative number of times a creative variant has been exposed on a single platform exceeds the maximum number of times specified by us, over-frequency suppression penalty will be applied on that platform to suppress excessive push to the same user [13].

Considering that the propagation process has obvious temporal dependence characteristics, a long short-term memory network is used as the backbone structure of the policy function and the value function. The policy network outputs the mean and variance parameters of the action vector. The actual action is sampled from the corresponding Gaussian distribution to obtain the exploration capability. The value network outputs the expected cumulative reward estimate of the current state, which is used to calculate the advantage function. The training process follows the proximal policy optimization framework and iteratively improves the scheduling policy in the interaction-sampling-update loop. In each iteration, trajectory data of a complete propagation cycle is collected from the simulation environment, the generalized advantage estimate of each time step is calculated, and then multiple rounds of gradient updates are performed on the policy network and the value network. To avoid the strategy falling into a suboptimal deterministic allocation pattern in the early stages of training, a high entropy regularization coefficient is introduced in the first twenty iterations to encourage exploration. The hyperparameter settings and decay scheduling involved in the training phase are detailed in Table 1.

TABLE 1. Training Hyperparameters and Decay Scheduling Table

Hyperparameter	Initial Value	Decay Strategy	Terminal Value / Policy
Entropy regularization coefficient	0.05	Multiply by 0.9 every 10 iterations	Minimum 0.005
Learning rate (policy network)	0.0003	Cosine annealing	Minimum 1×10^{-5}
Learning rate (value network)	0.001	Cosine annealing	Minimum 1×10^{-4}
GAE λ	0.95	Fixed	—
Discount factor γ	0.99	Fixed	—
Batch trajectory length	4×7 time slots (1 week)	Fixed	—

TABLE 2. Online A/B Evaluation of Audience Incremental Reach Efficiency

Metric	Control Group (Independent Optimization)	Experimental Group (Collaborative Scheduling)
Cross-platform deduplicated reach (10k users)	172.6	211.3
Average contact frequency per user (times)	3.84	3.21
Coefficient of variation of contact frequency	0.43	0.31
Proportion of overexposed users (>5 times)	19.2%	9.8%
Proportion of unreached users (0 times)	37.5%	23.4%

V. CROSS-PLATFORM CREATIVE COMMUNICATION COLLABORATIVE EFFICACY VERIFICATION

A. EXPERIMENTAL SETUP AND PREPARATION

The experiment selected three typical media platforms—short video platform (Douyin), text and image content community (Xiaohongshu), and instant messaging channel (WeChat Moments)—as the cross-platform observation domains. Anonymized user behavior logs were collected over four months, from September to December 2025, covering approximately 2.3 million interaction records and 47,000 unique creative material identifiers. Cross platform user identity association is probability matched using saltwater hash device fingerprints processed within a secure multi-party computing (MPC) framework to comply with the "walled garden" policy and data privacy regulations (such as PIPL). The matching recall rate is controlled at 83.6%. In order to mitigate the impact of the 16.4% matching gap on the reward function, Laplace smoothing based calibration was applied to the audience overlap matrix. Sensitivity testing confirmed that the convergence of the RL agent remained stable within a matching error range of ± 20%. This technical approach ensures that encrypted identifiers can be aggregated for overlapping analysis without leaking raw user data between competing platforms, thereby solving legal restrictions and technical silos. The baseline control group used independently operated advertising bidding systems on each platform (OCPM and eCPM native optimization strategies), while the experimental group deployed the collaborative optimization architecture constructed in this paper. The evaluation phase was set as a four-week online A/B comparison window, with both groups allocated equal budgets and a common creative asset pool, differing only in their communication strategies.

B. CROSS-DOMAIN SEMANTIC CONSISTENCY COMPARATIVE TEST EXPERIMENT

We would like to investigate how semantic fidelity of creative content in cross-platform communication. Therefore, we have applied the same source material, and generated its adapted variant by applying platform independent optimization strategy, as well as our proposed collaborative framework. We learned cosine similarity distribution between two variants and source material in CLIP semantic embedding space, and compared their mean and dispersion. We performed statistical testing on 500 randomly sampled samples.

In Figure 2, the mean cosine similarity of the semantic embedding between the creative variants and the source material under the platform-independent optimization strategies was only 0.62, with a high degree of dispersion (standard deviation 0.13). In contrast, the proposed collaborative framework increased the mean similarity to 0.86 and narrowed the standard deviation to 0.05, demonstrating a significant semantic anchoring effect. These data confirm that the introduction of adversarial domain adaptation and consistency constraints effectively suppressed semantic drift during cross-platform adaptation of creatives.

C. ONLINE A/B EXPERIMENT ON AUDIENCE INCREMENTAL REACH EFFICIENCY

This experiment was used to analyze the real effect of collaborative scheduling strategy on cross platform user reach efficiency. We ran a four week online A/B comparison among three platforms of short videos, text and image communities and instant messaging. In this experiment, the experimental group adopted collaborative scheduling method presented in this paper and control group adopted independent bidding strategy on each platform. The budget allocation and creative resource pool were the same for two groups. The marginal incremental effect of collaborative communication was measured by the scale of unique deduplicated users and dispersion of average contact frequency per user.

Table 2 shows that the collaborative scheduling strategy presented in this paper outperforms independent optimization on each platform in both audience coverage and reach balance. The experimental group reached 2.113 million users across platforms, a 22.4% increase compared to the control group's 1.726 million. The average frequency of contact per user decreased from 3.84 to 3.21, and the coefficient of variation narrowed from 0.43 to 0.31. Simultaneously, the proportion of users with excessive exposure decreased by 9.4 percentage points, and the proportion of users not reached due to low frequency decreased from 37.5% to 23.4%, a reduction of 14.1 percentage points. These results confirm that the collaborative scheduling mechanism

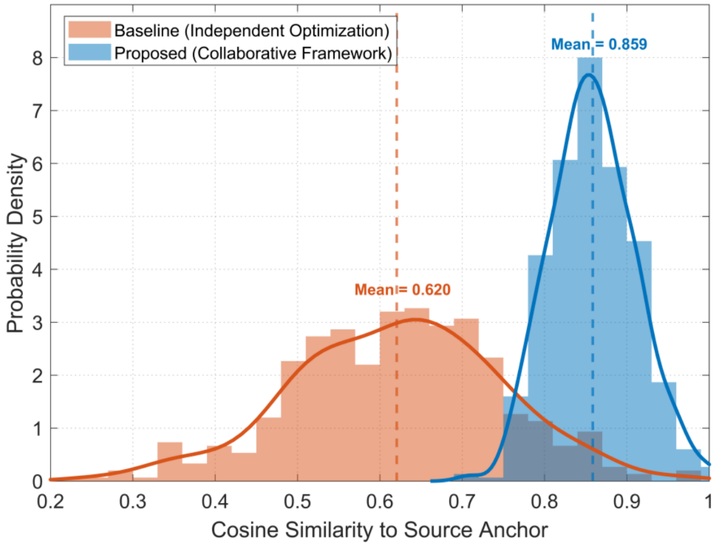

FIGURE 2. Cross-Domain Semantic Consistency Comparison Evaluation

effectively alleviates the structural imbalance in resource allocation while expanding audience coverage.

D. SIMULATION EXPERIMENT OF DISCUSSION AGGREGATION DECLINE DURING THE LONG TAIL STAGE OF SPREAD

This experiment reflects the cross-platform discussion aggregation drop after an idea enters the long tail period of dissemination. We built a simulation environment according to the historical logs and analyzed the comment interaction data from the seventh day to the thirtieth day after the release of the idea. We defined the degree of discussion aggregation based on mutual information value quantification. Compared with hard scheduling strategy, we analyzed the half-life characteristics of two scheduling strategies (collaborative scheduling framework and traditional hard scheduling strategy) in the case of hot topic.

In Figure 3, the cross-platform mutual information value decays rapidly under the traditional hard scheduling strategy, with a half-life of only 7.3 days, and the aggregation degree has dropped to 0.08 by the thirtieth day. The proposed collaborative framework effectively suppresses this decay trend, extending the half-life to 15.4 days, an improvement of 111.1% compared to the traditional method, and the mutual information value remains at 0.29 on the thirtieth day. The above data confirms that the collaborative scheduling mechanism can significantly delay the dispersion of discussion heat during the long tail stage and strengthen the resilience of cross-platform topic aggregation.

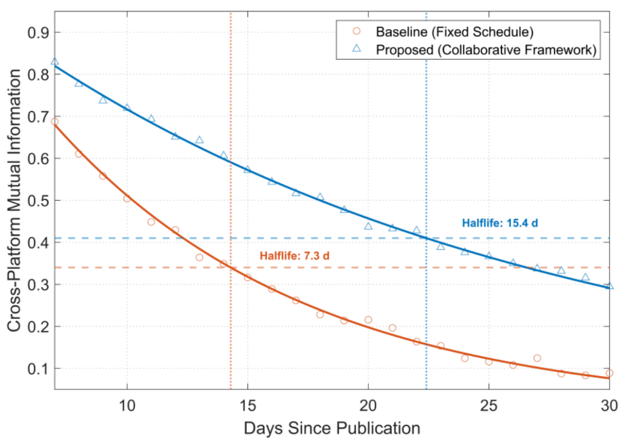

FIGURE 3. Simulation Evaluation of Discussion Aggregation Decline During the Long Tail Stage of Spread

VI. CONCLUSION

In this paper, we focus on collaborative optimization of creative communication across diverse platforms of media consumption, and build a hierarchical collaborative framework to integrate heterogeneous behavioral representations, adversarial domain adaptation and reinforcement learning scheduling. Evaluations based on the real-world dataset demonstrate that, compared with the traditional cross-domain flow of creative ideas, the hierarchical collaborative framework could reduce semantic dispersion and reach imbalance, and enhance audience reach as well as long-tail aggregation resilience while maintaining cognitive consistency. Due to the limitation of experimental period and accessible range of platform data, this paper still has not explored the robustness boundary of the model impacted by sudden trending events. It is an interesting future work to further integrate real-time public opinion signals and more multimodal semantic evolution mechanisms to enhance the adaptive ability of the hierarchical collaborative framework in different media scenarios.

AUTHOR CONTRIBUTIONS

Yiran Wang designed, collected and analyzed the data, and drafted the manuscript. Yiran Wang conducted the study, critically revised the manuscript for important intellectual content, and gave final approval of the version to be published. Yiran Wang participated fully in the work, take public responsibility for appropriate portions of the content, and agreed to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

CONFLICTS OF INTEREST

The author declares no conflict of interest.

FUNDING

This research received no external funding.

ACKNOWLEDGEMENTS

Not applicable.

REFERENCES

[1] M. Barnidge, T. Diehl, L. A. Sherrill, and J. Zhang, “Attention Centrality and Audience Fragmentation: An Approach for Bridging the Gap Between Selective Exposure and Audience Overlap,” Journal of Communication, vol. 71, no. 6, pp. 898-921, 2021, doi: 10.1093/joc/jqab023.

[2] P. Jürgens and B. Stark, “Mapping Exposure Diversity: The Divergent Effects of Algorithmic Curation on News Consumption,” Journal of Communication, vol. 72, no. 3, pp. 322-344, 2022, doi: 10.1093/joc/jqac009.

[3] S. Dvir-Gvirsman, D. Sude, and G. Raisman, “Unpacking News Engagement Through the Perceived Affordances of Social Media: A Cross-Platform, Cross-Country Approach,” New Media & Society, vol. 26, no. 11, pp. 6487-6509, 2024, doi: 10.1177/14614448231154432.

[4] V. Unnava and A. Aravindakshan, “How Does Consumer Engagement Evolve When Brands Post Across Multiple Social Media? Journal of the Academy of Marketing Science,” vol. 49, no, 2021. 5, pp. 864-881, doi: 10.1007/s11747-021-00785-z.

[5] J. Liu, W. Zhong, J. Zhang, and S. Mei, “The Effectiveness of Cross-Platform Targeted Advertising Strategy,” Electronic Commerce Research, vol. 24, no. 4, pp. 2831-2847, 2024, doi: 10.1007/s10660-022-09659-0.

[6] V. Hase, K. Boczek, and M. Scharkow, “Adapting to Affordances and Audiences? A Cross-Platform, Multi-Modal Analysis of the Platformization of News on Facebook, Instagram, TikTok, and Twitter,” Digital Journalism, vol. 11, no. 8, pp. 1499-1520, 2023, doi: 10.1080/21670811.2022.2128389.

[7] R. K. Nielsen and R. Fletcher, “Comparing the Platformization of News Media Systems: A Cross-Country Analysis,” European Journal of Communication, vol. 38, no. 5, pp. 484-499, 2023, doi: 10.1177/02673231231189043.

[8] D. Yan, W. Xie, and Y. Zhang, “Heterogeneous Information Network-Based Interest Composition with Graph Neural Network for Recommendation,” Applied Intelligence, vol. 52, no. 10, pp. 11199-11213, 2022, doi: 10.1007/s10489-021-03018-6.

[9] M. HassanPour Zonoozi and V. Seydi, “A Survey on Adversarial Domain Adaptation,” Neural Processing Letters, vol. 55, no. 3, pp. 2429-2469, 2023, doi: 10.1007/s11063-022-10977-5.

[10] B. S. Shi, Y. Q. Wang, F. D. Guo, B. B. Xu, H. W. Shen, and X. Q. Cheng, “Domain Adaptation for Graph Representation Learning: Challenges, Progress, and Prospects,” Journal of Computer Science and Technology, vol. 40, no. 2, pp. 283-300, 2025, doi: 10.1007/s11390-024-4465-x.

[11] F. Yang and D. Peng, “Spatio-Temporal Contrastive Heterogeneous Graph Attention Networks for Session-Based Recommendation,” Mathematics, vol. 12, no. 8, pp. 1-16, 2024, doi: 10.3390/math12081193.

[12] T. V. Nguyen, A. Nguyen, N. Le, and B. Le, “Semi-supervised Adversarial Discriminative Domain Adaptation,” Applied Intelligence, vol. 53, no. 13, pp. 15909-15922, 2023, doi: 10.1007/s10489-022-04288-4.

[13] C. Chen, G. Wang, B. Liu, et al., “Real-time bidding with multi-agent reinforcement learning in multi-channel display advertising,” Neural Computing and Applications, vol. 37, no. 1, pp. 499-511, 2025, doi: 10.1007/s00521-024-10649-6.