Classifying COVID-19 Related Meta Ads Using Discourse Representation Through a Hypergraph

Ujun Jeong[1(✉)], Zeyad Alghamdi[1], Kaize Ding[1], Lu Cheng[1,2], Baoxin Li[1], and Huan Liu[1]

[1] Arizona State University, Tempe, AZ, USA
{ujeong1,zalgham1,kding9,lcheng35,Baoxin.Li,huanliu}@asu.edu
[2] University of Illinois Chicago, Chicago, IL, USA

Abstract. Despite Meta's efforts to promote health information in the COVID-19 pandemic, the growing number of ads is making online content control extremely challenging. To effectively categorize the ads, this work investigates the major discourses shared across Meta ads with various categories related to COVID-19. We propose an interpretable classification model that captures common discourses in the form of keywords and phrases in ads. Particularly, we propose to use hypergraph to connect ads and discourses to capture their high-order interactions. Experiments on a curated Meta Ads dataset show that our model can provide subject-specific discourses and improve classification performance significantly.

Keywords: Meta advertisement · Text classification · Hypergraph

1 Introduction

As COVID-19 continues its global devastation, massive information related to COVID-19 has been seen on social media and other digital platforms. The unfolding of this pandemic has also demonstrated the great impact of this 'infodemic' on the sponsored content in social media. For example, Meta[1] Ads has been trying to promote health-related content by encouraging online campaigns [14]. However, this strategy becomes less effective and confuses the public even more, as advertisers often share controversial opinions such as racist theories about the origin of the virus and conspiracy theories in vaccination[2].

Regarding this issue, narrative theory suggests that identifying the repeated discourse is crucial for characterizing the content, as people tend to use specific phrases for their argument [3]. For example, in Fig. 1, *Ads 1–3* share the vaccination-related phrase "Medical Freedom" with the narrative of political need, and both *Ad 3* and *Ad 4* contain the vaccination-related phrase "Vaccine Mandate" with a negative stance. Based on these discourse relations, we

[1] Formerly known as Facebook (www.facebook.com).
[2] This claim is based on the article from Consumer Report (https://bit.ly/3I8xjeY).

R. Thomson et al. (Eds.): SBP-BRiMS 2022, LNCS 13558, pp. 35–45, 2022.
https://doi.org/10.1007/978-3-031-17114-7_4

may further infer that *Ads 1–4* are related to the policy about treatment and prevention of COVID-19.

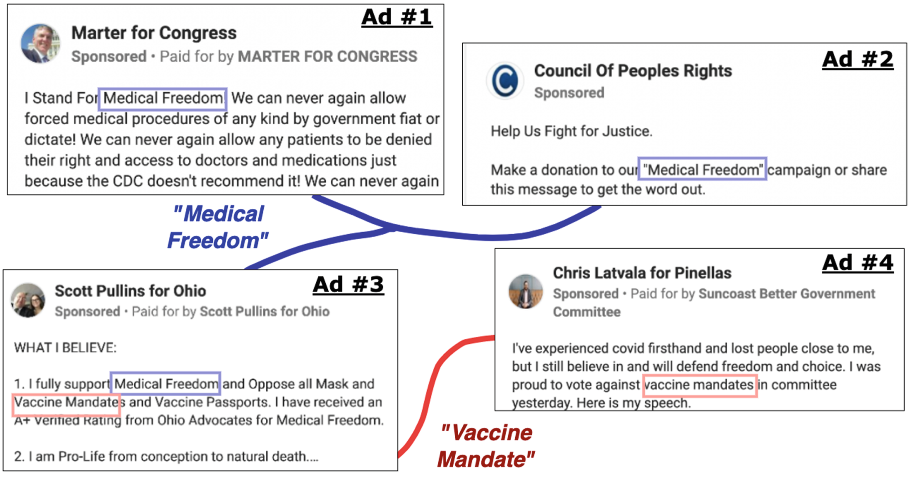

Fig. 1. Example of common discourses in COVID-19 Meta Ads. *Ads 1–3* have the shared phase about medical freedom. *Ads 3–4* are politically against vaccine mandate.

This work aims to help to efficiently classify COVID-19 related Meta ads into a set of pre-defined categories using the discourses repeatedly shared across ads. To achieve this research goal, we first define discourse as a phrase that can consist of one or multiple words shared across ads. However, using discourse for ads classification faces several major challenges: (1) Traditional text classification models are incapable of identifying important discourse during learning phase because they focus on finding an important set of relations at word-level [18]; (2) The importance of discourse cannot be flexibly measured depending on the related category, even though not all discourses are equally useful for classifying the category; and (3) As discourse is shared across the documents, the meaning of discourse should be dependent on the context of documents that contains the discourses. However, existing frameworks do not reflect this relation.

To address these challenges, in this paper, we propose a novel method coined as **DiscourseNet**. It leverages discourses shared in Meta ads for interpretable COVID-19 ads classification. To model these relations across Meta ads, we use hypergraph, an effective method to model the structure of complex relations between objects. Hypergraph is effective in modeling high-order relations [19] via connecting two or more nodes (i.e., ads) by the edge (i.e., discourses). We further integrate attention mechanism into hypergraph to capture important discourses related to specific ads categories.

Our main contributions are summarized as follows:

– Studying a novel classification problem of COVID-19 Meta ads,
– Proposing a hypergraph attention network for Meta Ads, DiscourseNet, to capture the common discourses unique to each category of COVID-19, and
– Curating a new dataset with 1,041 COVID-19 Meta ads for experiments.

2 Related Work

Meta Ad Analysis. This mainly consists of two lines of research: (1) investigating platform transparency and (2) measuring the impact of advertisement. As Meta offers marketing API to its advertisers, it has been questioned whether Meta microtargets users based on their sensitive information, such as race, religion, and sexual orientation [1]. To answer this question, many applications have been developed, such as AdAnalyst [1] and AdObservatory [9], to collect information about Meta's targeting strategies. [15] further showed that it is possible to carry out discriminatory campaigns, even without using sensitive attributes. Another work [14] proposes a classification model to identify the political content in Brazilian elections. This is for increasing the transparency of announcements made during elections with a tool external to Meta. To see the potential impact of Meta ads, Silva et al. [13] has focused on analysis on ads related to COVID-19. By using toxicity and sentiments analyses, they demonstrated the possible impact of political ads on shaping people's opinions.

Text Classification. With the superior performance of deep learning, various neural models have been developed for text classification, such as convolutional neural networks (CNNs) [12] and a series of recurrent neural networks including GRUs [4]. Hierarchical attention networks [17] further improve the model expressiveness by using attention mechanism on both word and sentence level for document classification. BERT [5] uses bidirectional encoding based on the multi-head attention mechanism. Moreover, the advent of graph neural networks has greatly helped to solve the problem of long-term dependence between words. Specifically, TextGCN [18] applies the graph convolutional networks to learn text representations on document-level graphs. HyperGAT [7] uses dual attention mechanism with hypergraphs for document classification, capturing high-order information in inductive manner.

This work aims to classify COVID-19-related Meta ads according to a set of categories, while at the same time identifying the discourses that are discussed within each category. Our framework is inherently different in that we explain the model's decision by keywords and phrases extracted from Meta ads for discourse-level interpretation. Especially, our approach represents the connection between ads and discourses through hypergraph to capture the high-order information across Meta ads while learning the importance of relations by attention networks.

3 Dataset

Data Collection. We use FBAdTracker [11] to collect the ads related to COVID-19 in Meta. FBAdTracker is an application designed to provide an integrated data collection and analysis system for researches on Meta ads. In the program, we begin by querying ads with the keyword *'covid'* for both active and inactive ads on Meta from all available countries. As collecting media content such as videos or images in Meta is not allowed without permission from each advertiser, we only use the text content of ads for our data collection. For this reason, we filter out ads without text content. i.e., we have excluded ads if the field 'ad_creative_body' is empty. Furthermore, we remove ads if the first and last ten words overlap with other ads, as there are many duplicate ads with slightly different choice of words. Finally, we only select those ads written in English. In total, we have collected 1,041 ads covering three months with the time frame from *Nov. 1st, 2021* to *Feb. 1st, 2022*. The statistics show the dataset includes 729 unique advertisers and 559 unique sponsors. Regarding the text in the ad, the average number of sentences in an ad is 6 and the average number of words in a sentence is ∼18.

Data Annotation. Following the annotation standard suggested in LitCovid [10], we divide the categories into eight different themes: *Prevention, Treatment, Diagnosis, Mechanism, Case Report, Transmission, Forecasting, and General*. We have asked two members of our team to label each sample in a multi-class fashion, which means one sample can be assigned to more or equal to one category. The instructions for the labeling process requires reading only the text content of the ad, and checking every applicable options in the suggested categories. For example, if an ad talks about vaccination and the rising number of cases, the category of both categories *'Treatment'* and *'Case Report'* will be filled with 1 and others will be filled with 0. As shown in statistics of categories in Table 1, there is imbalance between the number of classes. This is because most of the ads include content related to *'General'* as advertisers frequently talk or quote about personal stories to let audience engaged in the message. We publicly release the data by providing the list of IDs of Meta ads and the corresponding annotation in our GitHub[3] according to the terms of service of Meta[4].

Table 1. Number of collected Meta ads per category. The annotation is based on multi-class, and therefore the sum of labels is different from the total number of ads.

Prevention	Treatment	Diagnosis	Mechanism	Report	Transmission	Forecasting	General
187	311	126	18	53	32	18	634

[3] https://github.com/ujeong1/SBP22_DiscourseNet.

[4] https://www.facebook.com/terms.php.

4 Methodology

To better classify the COVID-19 related Meta ads, this work proposes to model the relations among different Meta ads based on their discourses to learn high-order representations of each advertisement. This is achieved by using hypergraph [2], a generalization of graph structure in which an edge can connect more than two nodes. In the following, we explain the details about (1) construction of hypergraph for Meta ads (2) learning relational information between nodes and hyperedge using dual attention mechanism [7], and (3) the final layer for multi-class node classification. An overview framework is depicted in Fig. 2.

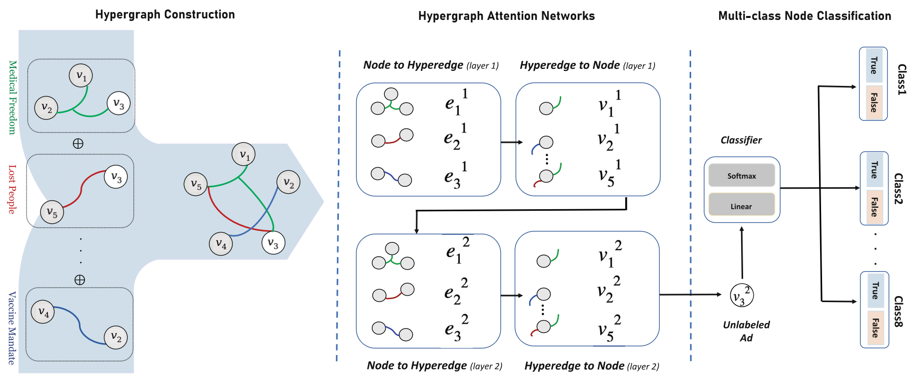

Fig. 2. The overview of the proposed model. It shows an example of node classification using hypergraph where the number of node (advertisement) is 5 and the number of hyperedge (discourse) is 3. The unlabeled ad (3rd node) is classified into eight classes.

4.1 Hypergraph Construction for Meta Ads

In real-world scenarios, relations among the objects are more complex than pair-wise relations [19]. Hypergraph has been widely used for modeling high-order relationships in different applications [7,16]. The formal definition of hypergraph is as follows:

Hyeprgraph. A graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_1, v_2, \cdots, v_n\}$ denotes the set of nodes in the graph, and $\mathcal{E} = \{e_1, e_2, \cdots, e_m\}$ represents the set of hyperedges. Different from graph, hypergraph connects two or more nodes. The hypergraph $\mathcal{G}$ can be denoted by an incidence matrix $\mathbf{H} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{E}|}$, with entries defined as:

$$h(v, e) = \begin{cases} 1 & \text{if } v \in e \\ 0 & \text{if } v \notin e \end{cases} \tag{1}$$

To apply the notion of hypergraph to Meta ads classification task, we use the Meta ads as nodes $\mathcal{V}$ and discourses as hyperedges $\mathcal{E}$. We define discourse as a form of keyword or phrase (a composition of multiple words) which is shared in two or more advertisements. We construct the hypergraph by connecting ads by the common discourse. For example, if the phrase "Lost my family" appears in three ads, we connect them by a hyperedge. To this end, we need to first extract keyword and phrases from Meta ads. We use RAKE[5], which is a tool capable of extracting discourses (keywords and phrases) in a document based on the frequency of co-occurrence. In total, we extract 1,541 of discourses from our Meta Ad dataset after selecting discourses that show at least twice among ads.

4.2 Hypergraph Attention Network

Since not all relations and advertisements in hypergraph are equally important, we introduce a hypergraph attention network well-designed for our Meta ad classification problem. As a family of graph neural networks [6,8], hypergraph attention network calculates the importance between the node (Meta advertisement) and the hyperedge (discourse) by recursively aggregating the embeddings and computing the attention coefficients using dual attention mechanism [7].

Node to Hyperedge Aggregation. Given a hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and the representations of nodes $v^{(l-1)} \in \mathbb{R}^{|\mathcal{V}| \times d}$ where l is the layer in hypergraph attention network and d is the size of hidden dimension. We use node-level attention to compute the importance of nodes to the meaning of the hyperedge. Then, nodes are aggregated to compute the hyperedge representation $e_j^{(l)} \in \mathbb{R}^d$.

$$e_j^{(l)} = ReLU\left(\sum_{v_k \in e_j} \gamma_{jk} \mathbf{W}_1 v_k^{(l-1)}\right), \tag{2}$$

The initial input of hypergraph attention network is defined as $v^{(0)} = v$, and γ_{jk} is the attention coefficient of node $v_k^{(l-1)}$. The attention coefficient is computed based on the trainable weight matrix $\mathbf{W}_1$ and weight vector a_1.

$$\gamma_{jk} = \frac{\exp\left(a_1^\intercal u_k\right)}{\sum_{v_p \in e_j} \exp\left(a_1^\intercal u_p\right)}, u_k \quad = LeakyReLU\left(\mathbf{W}_1 v_k^{(l-1)}\right), \tag{3}$$

Hyperedge to Node Aggregation. Based on the hyperedge representations $e^{(l)} \in \mathbb{R}^{|\mathcal{E}| \times d}$ in layer l and a set of hyperedge $\mathcal{E}_i$ connected to node v_i, we apply edge-level attention mechanism to discover the informative hyperedges. To this end, hyperedge representations are aggregated to compute the next-layer representation of nodes $v^{(l)}$.

$$v_i^{(l)} = ReLU\left(\sum_{e_j \in \mathcal{E}_i} \lambda_{ij} \mathbf{W}_2 e_j^{(l)}\right), \tag{4}$$

[5] https://pypi.org/project/rake-nltk/.

$v_i^{(l)}$ is the output representation of the node v_i in layer l and $\mathbf{W}_2$ is λ_{ij} is the attention coefficient of hyperedge $e_j^{(l)}$ on node $v_i^{(l)}$. The attention coefficient is computed based on a weight matrix $\mathbf{W}_2$ and a weight vector a_2 for hyperedges.

$$\lambda_{ij} = \frac{\exp\left(a_2^\intercal z_j\right)}{\sum_{v_p \in e_q, e_q \in \mathcal{E}_i} \exp\left(a_2^\intercal z_p\right)}, z_j \quad = LeakyReLU\left([\mathbf{W}_2 e_j^{(l)} \oplus \mathbf{W}_1 v_i^{(l-1)}]\right), \tag{5}$$

Node (Meta Ad) Classification. The final vector $v^{(L)} \in \mathbb{R}^{|\mathcal{V}| \times d}$ is a high-level representation resulted from node embedding through hyperedge to node aggregation, where L is the last layer of hypergraph attention network. For the node classification layer, we use a fully connected layer with weight matrix $\mathbf{W}_3$ and bias term b. Furthermore, we add the initial representation of $i-th$ node $v_i^{(0)}$ by skip-connection to maintain the original characteristics of Meta Ads. To train the model, we use BCE (Binary logit Cross Entropy) to optimize the loss using predicted probability p_i and multi-class ground-truths, $\{y_1^k, \ldots, y_C^k\} \in \{0,1\}^C$ where C is the number of the classes.

$$\begin{aligned} p_i &= Softmax(f(\mathbf{W}_3(v_i^{(L)} + v_i^{(0)}) + b)), \\ \mathcal{L}_{BCE} &= \begin{cases} -\log({p_i}^k) & if\ y_i^k = 1 \\ -\log(1 - {p_i}^k) & otherwise. \end{cases} \end{aligned} \tag{6}$$

5 Experiment

In this section, we present the experiments to evaluate the effectiveness of DiscourseNet. Our experiment design aims to answer the following questions:

- **RQ1:** Can our proposed approach improve the performance of Meta ads classification compared to the baselines in text classification task?
- **RQ2:** Can our model identify meaningful discourses related to the category?

Baseline Settings. We include the following baselines: (1) sequence-to-sequence models, including GRUs [4] and TextCNNs [12]; (2) BERT+MLP, which is a text classification model that applies MLP on BERT [5] pre-trained embeddings from Meta ads; and (3) TextGCN [18] which applies graph convolutional networks for text classification. In the experiments, we set parameters as follows: hidden dimension (128), batch size (128), optimizer (Adam), learning rate (0.001), dropout rate (0.3), and epochs (50). We use Meta ads encoded by BERT [5] with the dimension size of 768. For the evaluation, we use accuracy (exact match) and F1-score (harmonic mean of precision and recall).

5.1 RQ1: Evaluation on Text Classification Task

To answer RQ1, we compare DiscourseNet with representative baselines for text classification task based on 5-fold cross-validation. As shown in Table 2, we observe that the proposed method outperforms all the baselines on both mean accuracy and F1-score. Specifically, DiscourseNet improves accuracy over TextGCN by ∼1% point, and achieves higher F1-score over BERT+MLP around ∼2% points. It is notable TextGCN shows comparable accuracy to DiscourseNet. A possible explanation for this is that TextGCN can leverage the relational information by modeling document-word relation, which helps to find some meaningful correlation between the documents sharing the same words. On the other hand, BERT+MLP shows lower F1-score than DiscourseNet, as it only learns documents individually without any relations between them. Therefore, the experimental result demonstrates that learning relational information in discourse-level works crucially for classifying texts in Meta ads.

Table 2. Evaluation result on the curated Meta ad dataset. We run 5-fold cross-validation and average the results. The results indicate mean ± standard deviation.

Metrics	GRU	TextCNN	BERT+MLP	TextGCN	**DiscourseNet**
Accuracy	33.87 ± 4.58	52.24 ± 2.81	52.72 ± 1.77	56.48 ± 2.27	$\mathbf{57.51 \pm 3.33}$
F1-score	28.31 ± 5.28	33.16 ± 2.57	33.96 ± 2.42	32.08 ± 2.59	$\mathbf{36.31 \pm 3.53}$

5.2 RQ2: Discourse Identification in Categories

To answer RQ2, we identify the subject-specific discourses in the curated dataset by investigating the attention scores of the proposed model per category. To this end, we follow three procedures: (1) selecting nodes (Meta ads) corresponding to a certain category (2) getting hyperedges (discourses) connected to the selected nodes (3) ranking the top eight hyperedges based on attention scores.

As shown in Table 3, the top eight discourses are highly related to each category, and not associated with common terms such as stopwords. For example, in 'Prevention' category, we notice that discourses about guidelines suggested by CDC[6], such as mask mandate and wearing mask, and self-isolation. We also observe discourses predominantly used in context of vaccination in 'Treatment' category such as medical freedom, lost people, and federal vaccine. Particularly, the 'General' category shows discourses such as conspiracy theories, human rights, and school closure, which are more common yet controversial subjects.

Furthermore, we examine if the model can highlight the importance of different discourses within a document. We visualize attention scores that measure the importance of discourses in Fig. 3. As shown in the example, discourses closely associated with the category show higher attention scores, while common discourses have lower values. For instance, in the category 'Prevention', the example

[6] www.cdc.gov.

Table 3. Top 8 discourses of each category ranked based on the attention scores on hyperedges. Based on the hypergraph construction method suggested in Sect. 4.1, a discourse corresponds to a hyperedge. A discourse is expressed as a keyword or phrase.

Category	Discourses
Prevention	Wear mask, dangerous new policy, self-isolation, help prevent, surviving covid, sending asymptomatic, zero covid, mask mandate
Treatment	Lost people, medical freedom, federal vaccine, anti-vax, second dose, covid jabs, unconstitutional covid, location near
Diagnosis	Get tested, home tests, requiring proof, testing mandate, full covid, completely free, mask except, telescope health
Mechanism	Boosted person, shorter period, increased risk, experience less, unvaccinated, lost several, increase awareness, hospitalized
Case Report	Skyrocketing, infection rates, borders, covid fiscal, covid variant, nursing homes, protect americans, virus safe
Transmission	Transmit, highly transmissible, protect public, spreading, medical experts, omicron variants, guidelines, government mask
Forecasting	Inflation grow, protect small business, relyong on food bank, interest rate, exchange rate, credit card, tax cut, income
General	Covid virus, conspiracy theories, human rights, next governor, school closure, zoom broad cast, covid inquiry, stress

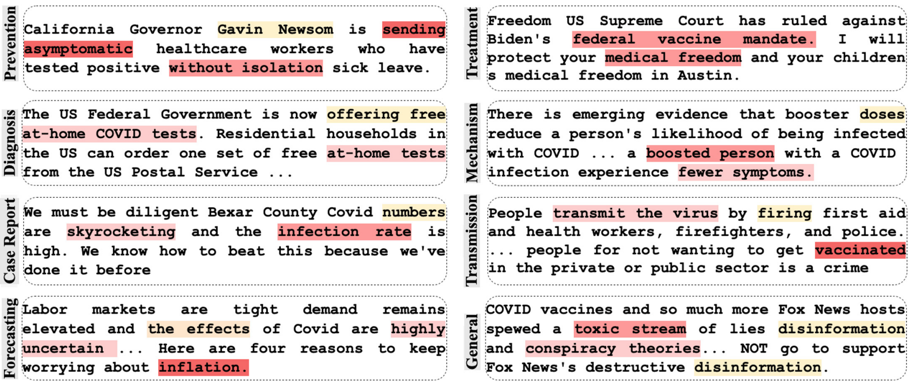

Fig. 3. Visualization of the attention score of discourses in advertisements arranged by each category. The highlighted words and phrases indicate common discourse shared across other documents, and the color represents the attention score (the color is more red when the absolute attention score is higher). The figure is best viewed in color. (Color figure online)

document shows the attention weight for the keyword ‘Gavin Newsom’ is lower than other phrases such as ‘without isolation’ and ’sending asymptomatic’.

6 Conclusion and Future Work

We propose DiscourseNet, a hypergaph-based model for text classification and discourse analysis in Meta ads. DiscourseNet outperforms representative baselines in text classification on both accuracy and F1-score by leveraging relational information through shared discourses. Through the experiments, we open a new venue of discourse analysis of Meta ads by quantitatively analyzing the discourses identified by hypergraph attention network. In future work, we focus on dynamically learning the discourses without pre-computation of keywords and phrases.

Acknowledgements. This work was supported by the Office of Naval Research under Award No. N00014-21-1-4002. Opinions, interpretations, conclusions, and recommendations are those of the authors.

References

1. Andreou, A., Silva, M., Benevenuto, F., Goga, O., Loiseau, P., Mislove, A.: Measuring the Facebook advertising ecosystem. In: Network and Distributed System Security Symposium (2019)
2. Berge, C.: Hypergraphs: Combinatorics of Finite Sets. Elsevier, Amsterdam (1984)
3. Bernardi, D.L., Cheong, P.H., Lundry, C., Ruston, S.W.: Narrative landmines. In: Narrative Landmines. Rutgers University Press (2012)
4. Cho, K., Van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: encoder-decoder approaches. arXiv (2014)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. arXiv (2018)
6. Ding, K., Wang, J., Caverlee, J., Liu, H.: Meta propagation networks for graph few-shot semi-supervised learning. In: AAAI (2022)
7. Ding, K., Wang, J., Li, J., Li, D., Liu, H.: Be more with less: hypergraph attention networks for inductive text classification. arXiv (2020)
8. Ding, K., Wang, J., Li, J., Shu, K., Liu, C., Liu, H.: Graph prototypical networks for few-shot learning on attributed networks. In: ACM International Conference on Information and Knowledge Management (2020)
9. Edelson, L., Sakhuja, S., Dey, R., McCoy, D.: An analysis of united states online political advertising transparency. arXiv (2019)
10. Gutierrez, B.J., Zeng, J., Zhang, D., Zhang, P., Su, Y.: Document classification for COVID-19 literature. arXiv (2020)
11. Jeong, U., Ding, K., Liu, H.: FbadtTacker: an interactive data collection and analysis tool for Facebook advertisements. arXiv (2021)
12. Kim, Y.: Convolutional neural networks for sentence classification. Empirical Methods Nat. Lang. Process. (2014)
13. Silva, M., Benevenuto, F.: COVID-19 ads as political weapon. In: Annual ACM Symposium on Applied Computing (2021)
14. Silva, M., Santos de Oliveira, L., Andreou, A., Vaz de Melo, P.O., Goga, O., Benevenuto, F.: Facebook ads monitor: an independent auditing system for political ads on Facebook. In: The Web Conference (2020)
15. Speicher, T., et al.: Potential for discrimination in online targeted advertising. In: Conference on Fairness, Accountability and Transparency. PMLR (2018)

16. Wang, J., Ding, K., Zhu, Z., Caverlee, J.: Session-based recommendation with hypergraph attention networks. In: SIAM International Conference on Data Mining (SDM) (2021)
17. Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., Hovy, E.: Hierarchical attention networks for document classification. In: NAACL (2016)
18. Yao, L., Mao, C., Luo, Y.: Graph convolutional networks for text classification. In: AAAI (2019)
19. Zhou, D., Huang, J., Schölkopf, B.: Learning with hypergraphs: clustering, classification, and embedding. In: NIPS (2006)