

Multimodal Post Attentive Profiling for Influencer Marketing

Seungbae Kim[†], Jyun-Yu Jiang[†], Masaki Nakada[†], Jinyoung Han[‡] and Wei Wang[†]

[†]Department of Computer Science, University of California, Los Angeles

[‡]Department of Applied Artificial Intelligence, Sungkyunkwan University

{sbkim,jyunyu,nakada,weiwang}@cs.ucla.edu,jinyoung@skku.edu

ABSTRACT

Influencer marketing has become a key marketing method for brands in recent years. Hence, brands have been increasingly utilizing influencers' social networks to reach niche markets, and researchers have been studying various aspects of influencer marketing. However, brands have often suffered from searching and hiring the right influencers with specific interests/topics for their marketing due to a lack of available influencer data and/or limited capacity of marketing agencies. This paper proposes a multimodal deep learning model that uses text and image information from social media posts (i) to classify influencers into specific interests/topics (e.g., fashion, beauty) and (ii) to classify their posts into certain categories. We use the attention mechanism to select the posts that are more relevant to the topics of influencers, thereby generating useful influencer representations. We conduct experiments on the dataset crawled from Instagram, which is the most popular social media for influencer marketing. The experimental results show that our proposed model significantly outperforms existing user profiling methods by achieving 98% and 96% accuracy in classifying influencers and their posts, respectively. We release our influencer dataset of 33,935 influencers labeled with specific topics based on 10,180,500 posts to facilitate future research.

CCS CONCEPTS

• Information systems Data mining; • Applied computing Economics.

KEYWORDS

Influencer marketing; Multimodal neural network; Multi-task learning; Influencer profiling; Social media; User profiling

ACM Reference Format:

Seungbae Kim[†], Jyun-Yu Jiang[†], Masaki Nakada[†], Jinyoung Han[‡] and Wei Wang[†]. 2020. Multimodal Post Attentive Profiling for Influencer Marketing. In *Proceedings of The Web Conference 2020 (WWW '20)*, April 20–24, 2020, Taipei, Taiwan. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3366423.3380052>

1 INTRODUCTION

Influencer marketing [15], which utilizes special individuals in social media, has gained great attention from brands. Brands expect

to promote brand awareness and advertise products to social networks of influencers [13, 17, 22, 31], who have 'influence' over a large number of followers [3, 4], since customers are often more likely to trust influencers' recommendations than brands' advertisements [5, 26, 31]. It has been reported that the global influencer market value was estimated to be 2 billion U.S. dollars as of 2017 and will increase to 10 billion U.S. dollars by 2020 [7]. The growing interest in influencer marketing has led many social media users to participate in marketing campaigns and create advertising content [19].

Due to its popularity, brands tend to increase their budgets for influencer marketing [21], and researchers have started studying various aspects of influencer marketing [8, 11, 20, 22, 29, 33]. However, big challenges remain for both brands and researchers in the rapidly growing market. So far, most brands have relied on influencer marketing agencies to hire influencers, but such agencies usually have a limited number of influencers who registered in their services, which may limit the chance to find more proper influencers, i.e., influencers who work with other agencies or did not register to any marketing agency are not considered [23]. Also, previous studies on influencer marketing mostly relied on small datasets that are acquired through surveys of influencers [8, 11, 22] or finding a few influencers on Instagram [20, 29] due to lack of available influencer data.

We, therefore, believe developing an influencer profiling model that can classify influencers with specific interests can provide valuable information for brands in their influencer hiring process or marketing strategies. Furthermore, constructing a large-scale and informative influencer dataset using such a model can foster researchers to conduct in-depth research, for example, analyzing the influencers-brands relationship, targeted audiences' responses, and influencer marketing effectiveness, as well as building efficient recommendation systems. Consequently, such research will enable brands to understand how to initialize influencer marketing campaigns, manage relationships with influencers, and effectively promote advertisements through influencers. Also, an influencer profiling model enables brands to easily find influencers for their marketing campaigns, without hiring them from marketing agencies.

To shed light on the above issue, we propose a multimodal convolutional neural network model that uses text and image information to classify influencers into specific interests such as beauty or travel. More specifically, our model takes the text and image features from social media posts published by an influencer to generate an influencer embedding by using the attention mechanism. The attention helps find posts that are more relevant to the influencer's topic, thereby obtaining a better influencer representation than existing

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '20, April 20–24, 2020, Taipei, Taiwan

© 2020 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-7023-3/20/04.

<https://doi.org/10.1145/3366423.3380052>

user profiling methods [12, 18, 27, 28, 34]. In addition to the influencer classification task, our model also classifies all posts into certain post categories (e.g., fashion, food, or interior). In our work, we use Instagram as our research context since it is the most popular social media website for influencer marketing [25]. Note that our model can be generalized and applied to any social media (e.g., Facebook, Twitter, Pinterest), where users post images with texts. We summarize our contributions as follows.

- This paper proposes a multimodal deep learning model for classifying influencers in the first place. Our model classifies users' interests by using both image and text information extracted from posts. The experimental results show that the proposed model achieves 98% accuracy in classifying influencers, and significantly outperforms the existing user profiling methods and other machine learning models with the same input features.
- We propose to use attention to make a representation for user profiling. The attention in our model assigns high scores on important posts thereby giving the robustness on our model. Our model achieves over 90% accuracy even with only 20 input posts.
- We release the Instagram influencer dataset¹ that contains 33,935 influencer information (e.g., account, biography, number of followers), and 10,180,500 posts (e.g., captions, number of likes, comments). Each influencer in the dataset includes a category label tagged by our proposed model, which is one of the eight major influencer categories, i.e., beauty, family, fashion, fitness, food, interior, pet, and travel.

2 RELATED WORK

2.1 User Profiling on Social Media

Although our work is the first attempt to profiling influencers' interests, many researchers have studied the user classification (or user profiling) on online social networks (OSNs). Pennacchiotti and Popescu [27, 28] proposed a machine learning-based user classifier on Twitter for three tasks, detecting users' political affiliation, ethnicity, and business affiliation. They employed the Gradient Boosted Decision Trees framework [14] as a classification model to profile users. Hung et al. [18] proposed a tag-based user profiling method for social media recommendation. The method finds a set of tags from users' profiles and posts and calculates weights for the tags to build the user profile. You et al. [34] focused on visual content to classify users' interests. The authors analyzed the users' individual images and then aggregated the image analysis results to obtain the user's interest distribution. These studies used various information for classifying users in OSNs, but most of the work paid little attention to use both textual and visual features to profile social media users. Farnadi et al. [12] presented a deep neural network that takes multimodal features that represent user characteristics. They introduced the power-set combination approach to aggregate multimodal features including high-level textual information, facial information, and a set of neighbors from a page like relationship network. However, the proposed features may not be applicable for identifying the main interest of a given social media user because most features represent high-level user characteristics.

¹<https://sites.google.com/site/sbkimcv/dataset>




Figure 1: The overall schema of *Influencer Profiler*.




Figure 2: The architecture of the multimodal post encoder.

2.2 Influencer Marketing

As influencer marketing has gained great attention, many researchers have studied in various topics. It has been widely recognized that the number of followers of an influencer, the size of the target audience, is one of the most important factors in influencer marketing [13, 17]. Lou and Yuan [22] presented that influencers' trustworthiness, attractiveness, and similarity to their followers positively influence on expanding brand awareness and increase purchase intentions. On the other hand, some researchers studied aspects that negatively affect to the marketing effectiveness [8, 11]. Kim et al. [20] showed that influencers are tightly connected to each other and have common followers in OSNs, especially when they have similar occupations. Yang et al. [33] analyzed social relationship between influencers and brands based on brand mentioning in social media posts. These studies have revealed valuable insights into understanding influencer marketing in OSNs. However, most of the previous work relied on survey data with a small number of influencers, or Instagram posts data uploaded by a small number of influencers. Also, little attention has been paid to identify topics (or categories) that influencers are interested in, which is crucial for brands to hire proper influencers and maximize marketing effect.

3 MULTIMODAL POST ATTENTIVE INFLUENCER PROFILER

In this section, we first state the problem to describe the objectives of our proposed model. we then present the proposed framework to classify both influencers and their posts by taking text and image features from posts.

3.1 Problem Statement

Here we formally define the goal of this paper. Given the social media posts of an influencer $P = \{p_1, \dots, p_n\}$, we aim to classify the influencer into the corresponding category of $c \in C$, where C is the classification space of categories. Moreover, each social media post p_i multimodally consists of a raw image p_i^i and a piece of texts p_i^t . More specifically, a raw image can be considered as a tensor of pixel values while texts of a post are an ordered lists of tokens.

3.2 Framework Overview

Figure 1 shows the overall schema of the proposed model, *Influencer Profiler*. The multimodal post encoder first encodes a set of posts of a given influencer to generate the post representations. The post encoder takes text and image from a post and obtains text features and image features by using pre-trained models, BERT [10] and Inception-v3 [30], respectively. The post encoder concatenates the text and image features to make a post representation. Next, the post attention layer takes post representations as an input and calculates a score for each post. The distribution of post scores is then used to output the influencer representation. Finally, our model predicts the category of a given influencer based on the influencer representation. Our proposed model is also capable of predicting the category of a social media post by learning the post representations.

3.3 Multimodal Post Encoder

To leverage the multimodal knowledge in posts, we propose the multimodal post encoder to derive a continuous representation for each post of the influencer using both text and image information. Figure 2 illustrates the architecture of the multimodal post encoder. We use the pre-trained BERT [10] to encode text features while the image features are derived by the pre-trained Inception-v3 [30].

Image Features. To generate post image features of a post p_i , we use the pre-trained Inception-v3 [30] model. We apply the transfer learning technique using the pre-trained model instead of training the model from scratch because the number of manually labeled posts is relatively small. In the transfer learning, the parameters on frozen layers are fixed and never updated. We fine-tune only the top 2 layers because those are what directly influence on the determination of the category, and we want to keep the same trained low-level feature detectors in the hidden layers, which was trained with much larger and more robustly spread dataset (ImageNet [9]). ImageNet contains 1 M images in 1,000 classes, which covers quite a variety of categories. In the image feature network, there are 49 hidden layers. We add a global spatial average pooling layer on top of the original output layer of Inception-v3, then add a fully-connected layer with the rectified linear unit (ReLU) activation function after the pooling layer to generate an image feature vector, \mathbf{x}_i^I , which has 1,024 dimensions.

Text Features. We exploit the pre-trained text model to derive text features. We use the BERT [10] model, which has 12 layers with 110 M parameters because it can capture contextualized information from text by applying the bidirectional transformers in the training procedure. We set the maximum sequence length as 128 and use only the last layer to obtain the text features. We then select the output of [CLS] token which is inserted at the beginning of an input sentence. Finally, output text feature vector \mathbf{x}_i^T for a post p_i has 768 dimensions.

After deriving the image and text features, the ultimate feature vector can be derived by concatenating the features of two resources:

$$\mathbf{x}_i = [\mathbf{x}_i^I; \mathbf{x}_i^T].$$

Finally, the continuous representation \mathbf{v}_i of the post p_i can be derived as:

$$\mathbf{v}_i = \text{ReLU}(\mathcal{F}(\mathbf{x}_i)),$$

where $\mathcal{F}(\cdot)$ is a fully-connected layer; $\text{ReLU}(\cdot)$ is the activation function.

3.4 Post Attentive Influencer Encoder

The influencer encoder generates a continuous influencer representation by taking a set of post feature vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$. We use the attention mechanism [2] to obtain the influencer embedding since all posts are not equally important to represent the category of the given influencer. For example, a small number of beauty posts published by a food influencer may not be very important to decide the influencer class. Instead, food posts should be considered as more important posts for classifying the category of food influencer. Therefore, the attention mechanism can be applied to weigh higher scores on more important posts.

For each post p_i of the influencer, a fully-connected layer with an activation function is first applied to project the post features into a hidden space as:

$$\mathbf{r}_i = \tanh(\mathcal{F}_a(\cdot)),$$

where $\mathcal{F}_a(\cdot)$ is a fully-connected layer; $\tanh(\cdot)$ is the activation function. A trainable context vector \mathbf{r}^c is then exploited to estimate the importance α_i of each post p_i with a softmax function as:

$$\alpha_i = \frac{\exp(\langle \mathbf{r}_i, \mathbf{r}^c \rangle)}{\sum_j \exp(\langle \mathbf{r}_j, \mathbf{r}^c \rangle)},$$

where $\langle \mathbf{r}_i, \mathbf{r}_j \rangle$ denotes the inner-product of \mathbf{r}_i and \mathbf{r}_j . Finally, the influencer representation \mathbf{h} can be constructed by a weighted combination of post features as:

$$\mathbf{h} = \sum_i \alpha_i \cdot \mathbf{v}_i.$$

Influencer Classification. Based on the influencer representation \mathbf{h} , the logits of influencer classification can be computed as:

$$\hat{\mathbf{y}} = \mathcal{F}_I(\text{ReLU}(\mathcal{F}_h(\mathbf{h}))),$$

where a fully-connected layer $\mathcal{F}_h(\cdot)$ and the activation function $\text{ReLU}(\cdot)$ perform a non-linear transformation while another fully-connected layer $\mathcal{F}_I(\cdot)$ infers the ultimate logits of categories for influencer classification.

3.5 Auxiliary Post Classification

In addition to influencer classification as the main task, we propose to improve the model by considering an auxiliary task. More specifically, our proposed model further classifies each post of the influencer into a post category $c_p \in C_p$. Note that the classification space of auxiliary post classification C_p is not necessary to be identical to the space of influencer classification C . If the post features are effective enough for post classification, the knowledge in the auxiliary task can be also leveraged to the main task of influencer classification.

By taking the representation \mathbf{v}_i derived from the post encoder for each post p_i , the model computes the logits of post classification as follows:

$$\hat{\mathbf{y}}_i^p = \mathcal{F}_p(\text{ReLU}(\mathcal{F}_s(\mathbf{v}_i))),$$

where we conduct a non-linear projection with a fully-connected layer $\mathcal{F}_s(\cdot)$ and an activation function $\text{ReLU}(\cdot)$; another fully-connected layer $\mathcal{F}_p(\cdot)$ generates the logits for post classification.




Figure 3: Distributions of the number of influencers in the collected dataset based on their numbers of (a) followers, (b) followees, and (c) posts.

3.6 Multi-task Learning

In this work, we learn the *Influencer Profiler* with multi-task learning for both tasks of influencer classification and post classification.

For influencer classification, we treat the task as a multi-class classification problem and utilize the cross-entropy [16] as the loss function. More precisely, the loss function of influencer classification as the main objective can be computed as:

$$\text{loss}_{\text{main}} = \sum_{c \in C} P(c | \mathbf{y}) \log P(c | \hat{\mathbf{y}}),$$

where $P(c | \mathbf{y})$ is the ground truth class distribution; $P(c | \hat{\mathbf{y}})$ is the estimated probability for the influencer category c by the logits $\hat{\mathbf{y}}$ and a softmax function.

As an auxiliary task, post classification can be also treated as a multi-class classification. Hence, the loss function of post classification can be written as:

$$\text{loss}_{\text{aux}} = \sum_{p_i \in P} \sum_{c \in C_p} P(c | \mathbf{y}_i^p) \log P(c | \hat{\mathbf{y}}_i^p),$$

where $P(c | \mathbf{y}_i^p)$ is the ground truth class distribution for the post p_i ; $P(c | \hat{\mathbf{y}}_i^p)$ can be also estimated by a softmax function.

Finally, the ultimate objective for multi-task learning can be a combination of two loss function as:

$$\text{loss} = \text{loss}_{\text{main}} + \text{loss}_{\text{aux}}.$$

4 DATA COLLECTION AND ANNOTATION

4.1 Dataset Collection

To build the influencer dataset, we employ the data collection process that includes the three steps. As the first step of our data collection, we search posts with hashtag(s) to collect user names of potential influencers. According to the FTC’s Endorsement Guides [6], influencers are required to disclose brand information by explicitly mentioning ‘paid advertisement’ and their relationships with brands if they advertise brands’ products. Therefore, we collect Instagram posts that contain the hashtag #ad which is the most commonly used hashtag for influencer marketing on Instagram [11]. We periodically query the hashtag #ad to download meta-data of newly updated posts that contain Instagram user names. We collect the data for 92 days from October 1st, 2018 to January 1st, 2019, which contains 828,045 posts by 107,656 unique Instagram users who may be potentially influencers. In the second step, we filter




Figure 4: Example images of Instagram posts of influencers in the eight influencer categories.

out users, with less than 1,000 followers or 300 posts, who are considered as non-influencers. Finally, our dataset includes 33,935 influencers. Figure 3 shows the distributions of the number of influencers based on their numbers of followers, followees, and posts. Notice that the portions of influencers who have followers less than 10,000, 25,000, and 50,000 are 37.38%, 60.22%, and 73.77%, respectively, and they are often dubbed as *micro-influencers* [24]. This result confirms that micro-influencers are actively participating in influencer marketing on Instagram, and brands are likely to work with them to reach targeted audiences. In the last step of our data collection, we download 300 recent posts of each influencer, which includes image files, captions, hashtags, usertags, numbers of likes, associated comments, etc. We finally obtain 10,180,500 posts from 33,935 influencers.

4.2 Category Identification

Influencer Categories. In this work, we define major categories of influencers based on the topic modeling. As influencers tend to introduce themselves by describing their interests in biography, we apply the Latent Dirichlet Allocation (LDA) topic model [32] to the biographies of the 33,935 influencers in the dataset. We first find ten topics, each of which is composed of ten words, and then manually select the most proper words that can represent the topics of influencers’ interests. In other words, we remove the words that cannot represent influencers’ interests such as ‘influencer’, ‘blogger’, ‘loving’, and etc. We finally identify eight major categories of influencers, *beauty*, *family*, *fashion*, *fitness*, *food*, *interior*, *pet*, and *travel*, as shown in Figure 4.

Post Categories. While we categorize topics of influencers into the eight categories, we define 10 categories of posts, which are the same categories of the eight influencer categories and two additional post categories including *product* and *other*. We add the *product* post category, which contains only products (e.g., cosmetics, fashion accessories) in the photo because many influencers often use product photos for advertising purposes. Posts that are not classified into any category, e.g., music, sports, arts, or humor posts, are labeled as *other* category.

Table 1: Model performance: influencer classification results with our proposed model and the eight baseline methods.

Model Type	Model	Accuracy	F-1 Scores on Different Influencer Categories							
			Beauty	Family	Fashion	Fitness	Food	Interior	Pet	Travel
Feature Baselines	<i>Tag</i> [18]	71.10%	0.606	0.656	0.735	0.739	0.877	0.540	0.780	0.716
	<i>Text&Social</i> [27, 28]	81.35%	0.736	0.747	0.855	0.745	0.889	0.722	0.830	0.816
	<i>User</i> [12]	84.06%	0.905	0.836	0.871	0.663	0.949	0.854	0.690	0.849
	<i>Image</i> [34]	90.02%	0.835	0.915	0.911	0.772	0.934	0.898	0.870	0.952
Model Baselines	<i>GaussianNB</i>	93.26%	0.992	0.986	0.960	0.774	0.913	0.944	0.725	0.945
	<i>KNeighbors</i>	93.87%	0.905	0.902	0.961	0.857	0.968	0.946	0.936	0.942
	<i>SVC</i>	94.13%	0.855	0.882	0.970	0.878	0.988	0.989	0.940	0.934
	<i>RandomForest</i>	94.65%	0.956	0.983	0.967	0.807	0.949	0.989	0.845	0.934
Proposed Model	<i>Influencer Profiler</i>	98.32%	0.994	0.988	0.989	0.898	0.991	0.986	0.990	0.977

4.3 Category Labeling

Influencer Labeling. We manually label influencers by examining their biographies and posts. We make two sets of labeled influencers for (i) training and validation and (ii) testing. In the first set, we label 200 influencers in each one of the eight influencer categories, which makes 1,600 labeled influencers in the first set. We then split the set into 8:2 for training and fine-tuning the model. While the first set is balanced to avoid biased results, we randomly select and label 1,142 influencers in the second set for testing purposes.

Post Labeling. To train data for the post classifier, we randomly select posts of the influencers from our dataset, and then manually label them into one of the 10 categories. Note that we do not take account of the posts that do not contain any text in the caption. Then, we obtain more than 1,000 posts for each post category and select 1,000 posts, which contain images as well as their captions, in each post category. Finally, we have 10,000 labeled posts.

5 EXPERIMENTS

5.1 Experimental Setting

We use TensorFlow [1] to implement our proposed model. We set the dimension size of the output of the post encoder as 256. For training the model, we set learning-rate as 10^{-3} and dropout probability as 0.5. We apply the heavy data augmentation to all images in our dataset by spacial rotation, spacial shifting, flipping, zooming, and channel shifting as a data preprocessing.

5.2 Baseline Methods

To evaluate the performance of our model, we consider two categories of baseline methods: (i) *Feature* and (ii) *Model*.

Feature baselines. In this category, we have four user profiling baseline methods that exploit different sets of input features from our model, to understand the importance of proposed text and image embedding to represent an influencer. We implement the baseline methods by extracting corresponding features from our dataset. (i) *Tag* method [18] uses tag information to classify social media users. (ii) *Text&Social* method [27, 28] has four feature sets: social media profile, posting behavior, linguistic content, and social network features. (iii) *User* method [12] uses textual characteristics of captions, image features from profile pictures, and neighbor information based on like relationship. (iv) *Image* method [34] classifies users' interest by using CNN based image classifier.




Figure 5: (a) Influencer classification accuracy scores on different number of posts. (b) The effects of input data modalities, fine-tuning, and attention on influencer classification.

Model baselines. In this category, baseline methods take the same input features as proposed model but use different machine learning approaches. Therefore, we can evaluate the novelty of our proposed model. We use the following four well-known learning methods: (i) Gaussian Naive Bayes, (ii) K-Nearest Neighbors, (iii) SVC, and (iv) Random Forest. Note that we aggregate all post features by averaging values since the input layer has multiple post features.

5.3 Experimental Results

5.3.1 Influencer Classification. We first evaluate the performance of the influencer classification task. Table 1 shows overall classification accuracy and F1 score of each influencer class. We find that *Tag* method shows poor performance since tags from posts are insufficient to represent the category of influencers. *Text&Social* and *User* methods improve overall classification accuracy but have poor F1 scores in specific influencer categories (e.g., fitness and pet). On the other hand, *Image* method achieves 90.02% accuracy and higher F1 scores than other baseline methods in the same category. This suggests that image information is more informative than other social media information for representing influencers' interests. The baseline methods in the Model category show better classification performance than Feature baseline methods. This implies that our proposed input features play a significant role in classifying influencer categories. Finally, Our proposed model, *Influencer Profiler*, outperforms all the baselines by achieving 98.32% accuracy. This is because our model uses the distribution of attention scores to focus on more important posts that can represent influencers.

Table 2: Post classification results of the proposed model and baseline methods across the ten post categories.

Model	Input	Accuracy	F-1 Scores on Different Post Categories									
			Beauty	Family	Fashion	Fitness	Food	Interior	Other	Pet	Product	Travel
<i>GaussianNB</i>	Text	40.70%	0.495	0.387	0.327	0.497	0.491	0.551	0.308	0.315	0.441	0.297
	Image	84.65%	0.956	0.726	0.850	0.847	0.937	0.938	0.819	0.709	0.804	0.837
	Text & Image	87.00%	0.904	0.731	0.850	0.887	0.948	0.930	0.844	0.871	0.858	0.874
<i>KNeighbors</i>	Text	38.85%	0.462	0.255	0.335	0.414	0.477	0.521	0.176	0.374	0.491	0.364
	Image	81.75%	0.946	0.690	0.826	0.812	0.931	0.935	0.793	0.623	0.773	0.794
	Text & Image	88.70%	0.923	0.706	0.839	0.886	0.954	0.939	0.840	0.922	0.891	0.927
<i>SVC</i>	Text	36.20%	0.507	0.325	0.308	0.392	0.296	0.548	0.273	0.161	0.458	0.281
	Image	67.00%	0.857	0.514	0.795	0.662	0.842	0.884	0.825	0.429	0.382	0.720
	Text & Image	71.50%	0.870	0.568	0.819	0.687	0.835	0.884	0.845	0.479	0.509	0.814
<i>RandomForest</i>	Text	31.80%	0.452	0.248	0.268	0.311	0.376	0.436	0.212	0.259	0.367	0.247
	Image	78.30%	0.866	0.624	0.770	0.772	0.859	0.902	0.736	0.757	0.714	0.811
	Text & Image	76.25%	0.909	0.624	0.738	0.791	0.830	0.871	0.701	0.770	0.624	0.724
<i>Influencer Profiler</i>	Text	60.90%	0.696	0.522	0.515	0.714	0.703	0.718	0.487	0.556	0.556	0.635
	Image	90.75%	0.970	0.752	0.874	0.928	0.955	0.958	0.884	0.897	0.905	0.931
	Text & Image	96.20%	0.982	0.891	0.932	0.975	0.990	0.985	0.955	0.968	0.955	0.979

The number of posts. We next examine how many posts are sufficient to accurately classify influencer categories. The computation cost will be remarkably reduced if we can classify influencers correctly with only a small number of posts. Figure 5(a) shows the influencer classification accuracy of the proposed model and baseline methods with a different number of input posts. The result reveals that the proposed model performs well with a small number of posts while performances of the baseline methods significantly drop. Note that the accuracy of the Influencer Profiler with 20 input posts is 92.9% and that of the baseline methods are ranged from 65% to 80%. The robustness of the model comes from the attention method which helps find more relevant posts to represent influencers even with a small number of posts.

Components importance. To understand the benefits of each component in the proposed model, we compare the performance of the full model with a model with no attention mechanism and a model with no fine-tuning. Note that we use mean values to aggregate post features in the model with no attention. We then use the following three input sets, only text, only image, and both text and image, to analyze the effect of input modalities on the models with different components. We first observe a large performance difference depending on the input sets as shown in Figure 5(b). The accuracy of the full model with only text, only image, and both text and image, are 89.04%, 95.47%, and 98.32%, respectively. This result suggests that image features are more useful than text features in identifying influencer categories. This is probably because the style of writing captions for each influencer differs more than the typical style of images in certain categories. Figure 5(b) also shows the performance gain of using attention and fine-tuning. The accuracy of the model with no fine-tuning with multi-modality is 97.41% which demonstrates that the fine-tuning the hidden layer gives a moderate gain. Attention mechanism, however, helps the model to achieve very high classification scores by weighting higher scores on more important posts which can represent the influencer category. Note that the model with no attention has 95.18% accuracy.

5.3.2 Post Classification. Table 2 shows the post classification results when the model uses (i) the text features only, (i) the image features only, and (i) both features. We first find that our model significantly outperforms other baseline methods. This demonstrates that our model can effectively capture unique characteristics for each post category. We also observe that the image features help the model to achieve higher classification performance than the text features. The accuracy of using only the text features and the image features of our proposed model are 60.9% and 90.75%, respectively. We observe that the post classification with the image features outperforms since influencers often write very short or random captions that are not related to the topic of the posts.

6 CONCLUSION

This paper proposed the multimodal deep learning-based model that uses attention to encode influencer representations to classify influencers into specific categories. Our proposed model takes multi-tasking learning to profile influencers by identifying their main interest and classifying their posts into the ten post categories. Our model achieves 98% accuracy in influencer classification and 96% accuracy in post classification, which are significantly higher than existing user profiling methods used in social media. Our model also outperforms other machine learning models that take the same input features, which demonstrates the novelty of our model. We collected 10,180,550 Instagram posts from 33,935 influencers and applied our proposed model to classify the influencers into their interests. We release the dataset which can be used by many researchers and marketers for understanding and modeling influencer marketing and help brands to effectively identify proper influencers among over 1 billion Instagram users for their marketing.

ACKNOWLEDGMENTS

We thank the reviewers for their constructive feedback. The work was partially supported by the NSF Division Of Graduate Education (DGE-1829071) and Basic Science Research Program through the NRF funded by the Ministry of Education (NRF-2018R1D1A1A020 85647). Corresponding Author: Jinyoung Han (jinyoung@skku.edu)

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*. 265–283.
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- [3] Eytan Bakshy, Jake M Hofman, Winter A Mason, and Duncan J Watts. 2011. Everyone’s an influencer: quantifying influence on twitter. In *ACM International Conference on Web Search and Data Mining (WSDM)*. ACM.
- [4] Meeyoung Cha, Hamed Haddadi, Fabricio Benevenuto, and P Krishna Gummadi. 2010. Measuring User Influence in Twitter: The Million Follower Fallacy. *International Conference on Weblogs and Social Media (ICWSM)* (2010).
- [5] Hyuk Jun Cheong and Margaret A Morrison. 2008. Consumers’ reliance on product information and recommendations found in UGC. *Journal of Interactive Advertising* 8, 2 (2008), 38–49.
- [6] Federal Trade Commission. 2017. The FTC’s endorsement guides: What people are asking.
- [7] Giordano Contestabile. 2018. Influencer Marketing in 2018: Becoming an Efficient Marketplace. *Adweek* (2018).
- [8] Marijke De Veirman, Veroline Cauberghe, and Liselot Hudders. 2017. Marketing through Instagram influencers: the impact of number of followers and product divergence on brand attitude. *International Journal of Advertising* 36, 5 (2017).
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [11] Nathaniel J Evans, Joe Phua, Jay Lim, and Hyoyeun Jun. 2017. Disclosing Instagram influencer advertising: The effects of disclosure language on advertising recognition, attitudes, and behavioral intent. *Journal of Interactive Advertising* (2017).
- [12] Golnoosh Farnadi, Jie Tang, Martine De Cock, and Marie-Francine Moens. 2018. User profiling through deep multimodal fusion. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. ACM, 171–179.
- [13] Rick Ferguson. 2008. Word of mouth and viral marketing: taking the temperature of the hottest trends in marketing. *Journal of consumer marketing* 25, 3 (2008).
- [14] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [15] Paul Gillin and Geoffrey A Moore. 2009. *The New Influencers: A Marketer’s Guide to the New Social Media*. Linden Publishing.
- [16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- [17] Stephen Guo, Mengqiu Wang, and Jure Leskovec. 2011. The role of social networks in online shopping: information passing, price of trust, and consumer choice. In *ACM conference on Electronic Commerce (EC)*. ACM.
- [18] Chia-Chuan Hung, Yi-Ching Huang, Jane Yung-jen Hsu, and David Kuan-Chun Wu. 2008. Tag-based user profiling for social media recommendation. In *Workshop on Intelligent Techniques for Web Personalization & Recommender Systems at AAAI*. AAAI.
- [19] Instagram. 2017. 2M Monthly Advertisers on Instagram. *Instagram blog* (2017).
- [20] Seungbae Kim, Jinyoung Han, Seunghyun Yoo, and Mario Gerla. 2017. How are social influencers connected in instagram?. In *International Conference on Social Informatics*. Springer, 257–264.
- [21] James Lawley. 2018. Brands to Increase Spend on Influencer Marketing Despite Fraud Concerns. *MarketingWeek* (2018).
- [22] Chen Lou and Shupej Yuan. 2019. Influencer Marketing: How Message Value and Credibility Affect Consumer Trust of Branded Content on Social Media. *Journal of Interactive Advertising* (2019), 1–16.
- [23] Mediakix. 2018. 4 Major Influencer Marketing Platform Risks Marketers Must Know. *Mediakix* (2018).
- [24] Mention. 2018. Instagram Reports 2018. <https://mention.com/en/reports/instagram/>.
- [25] Kimberlee Morrison. 2017. Instagram Dominates Influencer Marketing (Report). *Adweek* (2017).
- [26] Ac Nielsen. 2012. Nielsen: Global consumers’ trust in ‘earned’ advertising grows in importance. *Business Wire* (2012).
- [27] Marco Pennacchiotti and Ana-Maria Popescu. 2011. Democrats, republicans and starbucks aficionados: user classification in twitter. In *ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM.
- [28] Marco Pennacchiotti and Ana-Maria Popescu. 2011. A Machine Learning Approach to Twitter User Classification. In *International Conference on Weblogs and Social Media (ICWSM)*. AAAI.
- [29] Pijar Suciati, Mareta Maulidiyanti, and Amelita Lusia. 2018. Cultivation Effect of Tourism TV Program and Influencer’s Instagram Account on the Intention of Traveling. In *International Conference on Social Sciences (ICSS)*, Vol. 1.
- [30] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- [31] Ebru Uzunoğlu and Sema Misci Kip. 2014. Brand communication through digital influencers: Leveraging blogger engagement. *International Journal of Information Management* 34, 5 (2014), 592–602.
- [32] Hanna M Wallach. 2006. Topic modeling: beyond bag-of-words. In *International conference on Machine learning (ICML)*. ACM.
- [33] Xiao Yang, Seungbae Kim, and Yizhou Sun. 2019. How do influencers mention brands in social media? sponsorship prediction of Instagram posts. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. 101–104.
- [34] Quanzeng You, Sumit Bhatia, and Jiebo Luo. 2016. A picture tells a thousand words—about you! User interest profiling from user generated visual content. *Signal Processing* 124 (2016), 45–53.