Recommending Themes for Ad Creative Design via Visual-Linguistic Representations

Yichao Zhou
University of California, Los Angeles
yz@cs.ucla.edu

Shaunak Mishra
Yahoo Research
shaunakm@verizonmedia.com

Manisha Verma
Yahoo Research
manishav@verizonmedia.com

Narayan Bhamidipati
Yahoo Research
narayanb@verizonmedia.com

Wei Wang
University of California, Los Angeles
weiwang@cs.ucla.edu

ABSTRACT

There is a perennial need in the online advertising industry to refresh ad creatives, *i.e.*, images and text used for enticing online users towards a brand. Such refreshes are required to reduce the likelihood of ad fatigue among online users, and to incorporate insights from other successful campaigns in related product categories. Given a brand, to come up with themes for a new ad is a painstaking and time consuming process for creative strategists. Strategists typically draw inspiration from the images and text used for past ad campaigns, as well as world knowledge on the brands. To automatically infer ad themes via such multimodal sources of information in past ad campaigns, we propose a theme (keyphrase) recommender system for ad creative strategists. The theme recommender is based on aggregating results from a visual question answering (VQA) task, which ingests the following: (i) ad images, (ii) text associated with the ads as well as Wikipedia pages on the brands in the ads, and (iii) questions around the ad. We leverage transformer based cross-modality encoders to train visual-linguistic representations for our VQA task. We study two formulations for the VQA task along the lines of classification and ranking; via experiments on a public dataset, we show that cross-modal representations lead to significantly better classification accuracy and ranking precision-recall metrics. Cross-modal representations show better performance compared to separate image and text representations. In addition, the use of multimodal information shows a significant lift over using only textual or visual information.

CCS CONCEPTS

• Information systems → Online advertising.

KEYWORDS

Online advertising; transformers; visual-linguistic representation

ACM Reference Format:

Yichao Zhou, Shaunak Mishra, Manisha Verma, Narayan Bhamidipati, and Wei Wang. 2020. Recommending Themes for Ad Creative Design via Visual-Linguistic Representations. 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/10.1145/3366423.3380001>

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/10.1145/3366423.3380001>

Figure 1: Ad (creative) theme recommender based on a VQA approach. The inputs are derived from past ad campaigns: (i) ad images, (ii) text in ad images (OCR), (iii) inferred brands, (iv) Wikipedia pages of inferred brands, and (v) questions around the ad. The recommended themes (keyphrases) are aggregated per brand (or product category), and can be used by a creative strategist to choose images (*e.g.*, via querying stock image libraries) and text for new ads.

1 INTRODUCTION

With the widespread usage of online advertising to promote brands (advertisers), there has been a steady need to innovate upon ad formats, and associated ad creatives [2]. The image and text comprising the ad creative can have a significant influence on online users, and their thoughtful design has been the focus of creative strategy teams assisting brands, advertising platforms, and third party marketing agencies. Numerous recent studies have indicated the emergence of a phenomenon called *ad fatigue*, where online users get tired of repeatedly seeing the same ad each time they visit a particular website (*e.g.*, their personalized news stream) [25, 31]. Such an effect is common even in native ad formats where the ad creatives are supposed to be in line with the content feed they appear in [2, 31]. In this context, frequently refreshing ad creatives is emerging as an effective way to reduce ad fatigue [3, 4].

From a creative strategist's view, coming up with new themes and translating them into ad images and text is a time taking task which inherently requires human creativity. Numerous online tools have emerged to help strategists in translating raw ideas (themes)

into actual images and text, *e.g.*, via querying stock image libraries [6], and by offering generic insights on the attributes of successful ad images and text [7]. In a similar spirit, there is room to further assist strategists by automatically recommending brand specific themes which can be used with downstream tools similar to the ones described above. In the absence of human creativity, inferring such brand specific themes using the multimodal (images and text) data associated with successful past ad campaigns (spanning multiple brands) is the focus of this paper.

A key enabler in pursuing the above data driven approach for inferring themes is that of a dataset of ad creatives spanning multiple advertisers. Such a dataset [1] spanning 64,000 ad images was recently introduced in [15], and also used in the followup work [30]. The collective focus in the above works [15, 30] was on understanding ad creatives in terms of sentiment, symbolic references and VQA. In particular, no connection was made with the brands inferred in creatives, and the associated world knowledge on the inferred brands. As the first work in connecting the above dataset [1] with brands, [22] formulated a *keyword* ranking problem for a brand (represented via its Wikipedia page), and such keywords could be subsequently used as themes for ad creative design. However, the ad images were not used in [22], and recommended themes were restricted to single words (keywords) as opposed to longer keyphrases which could be more relevant. For instance, in Figure 1, the phrase *take anywhere* has much more relevance for Audi than the constituent words in isolation.

In this paper, we primarily focus on addressing both the above mentioned shortcomings by (i) ingesting ad images as well as textual information, *i.e.*, Wikipedia pages of the brands and text in ad images (OCR), and (ii) we consider keyphrases (themes) as opposed to keywords. Due to the multimodal nature of our setup, we propose a VQA formulation as exemplified in Figure 1, where the questions are around the advertised product (as in [15, 30]) and the answers are in the form of keyphrases (derived from answers in [1]). Brand specific keyphrase recommendations can be subsequently collected from the predicted outputs of brand-related VQA instances. Compared to prior VQA works involving questions around an image, the difference in our setup lies in the use of Wikipedia pages for brands, and OCR features; both of these inputs are considered to assist the task of recommending ad themes. In summary, our main contributions can be listed as follows:

- (1) we study two formulations for VQA based ad theme recommendation (classification and ranking) while using multimodal sources of information (ad image, OCR, and Wikipedia),
- (2) we show the efficacy of transformer based visual-linguistic representations for our task, with significant performance lifts versus using separate visual and text representations,
- (3) we show that using multimodal information (images and text) for our task is significantly better than using only visual or textual information, and
- (4) we report selected ad insights from the public dataset [1].

The remainder of the paper is organized as follows. Section 2 covers related work, and Section 3 describes the proposed method. The data sources, results, and insights are then described in Section 4, and we conclude the paper with Section 5.

2 RELATED WORK

In this section, we cover related work on online advertising, understanding ad creatives, and visual-linguistic representations.

2.1 Online advertising

Brands typically run online advertising campaigns in partnerships with publishers (*i.e.*, websites where ads are shown), or advertising platforms [21, 32] catering to multiple publishers. Such an ad campaign may be associated with one or more ad creatives to target relevant online users. Once deployed, the effectiveness of targeting and ad creatives is jointly gauged via metrics like click-through-rate (CTR), and conversion rate (CVR) [9]. To separate out the effect of targeting from creatives, advertisers typically try out different creatives for the same targeting segments, and efficiently explore [19] which ones have better performance. In this paper, we focus on quickly creating a pool of ad creatives for a brand (via recommended themes learnt from past ad campaigns), which can then be tested online with targeting segments.

2.2 Automatic understanding of ad creatives

The *creatives dataset* [1] is one of the key enablers of the proposed recommender system. This dataset was introduced in [15], where the authors focused on automatically understanding the content in ad images and videos from a computer vision perspective. The dataset has ad creatives with annotations including topic (category), questions and answers (*e.g.*, reasoning behind the ad, expected user response due the ad). In a followup work [30], the focus was on understanding symbolism in ads (via object recognition and image captioning) to match human-generated statements describing actions suggested in the ad. Understanding ad creatives from a brand's perspective was missing in both [15, 30], and [22] was the first to study the problem of recommending keywords for guiding a brand's creative design. However, [22] was limited to only text inputs for a brand (*e.g.*, the brand's Wikipedia page), and the recommendation was limited to single words (keywords). In this paper, we extend the setup in [22] in a non-trivial manner by including multimodal information from past ad campaigns, *e.g.*, images, text in the image (OCR), and Wikipedia pages of associated brands. We also extend the recommendations from single words to longer keyphrases.

2.3 Visual-linguistic representations and VQA

With an increasing interest in joint vision-language tasks like visual question answering (VQA) [8], and image captioning [26], there has been lot of recent work on visual-linguistic representations which are key enablers in the above mentioned tasks. In particular, there has been a surge of proposed methods using transformers [11], and we cover some of them below.

In LXMERT [28], the authors proposed a transformer based model that encodes different relationships between text and visual inputs trained using five different pre-training tasks. More specifically, they use encoders that model text, objects in images and relationship between text and images using (image,sentence) pairs as training data. They evaluate the model on two VQA datasets. More recently ViLBERT [20] was proposed, where BERT [11] architecture was extended to generate multimodal embeddings by processing both visual and textual inputs in separate streams which interact through co-attentional transformer layers. The co-attentional

transformer layers ensure that the model learns to embed the interactions between both modalities. Other similar works include VisualBERT [18], VLBERT [27], and Unicoder-VL [17].

In this paper, our goal is to focus on leveraging visual-linguistic representations to solve an ads specific VQA task formulated to infer brand specific ad creative themes. In addition, VQA tasks on ad creatives tend to be relatively challenging (*e.g.*, compared to image captioning) due to the subjective nature and hidden symbolism frequently found in ads [30]. Another difference between our work and existing VQA literature is that our task is not limited to understanding the objects in the image but also the emotions the ad creative would evoke in the reader. Our primary task is to predict different themes and sentiments that an ad creative image can invoke in its reader, and use such brand specific understanding to help creative strategists in developing new ad creatives.

3 METHOD

We first describe our formulation of ad creative theme recommendation as a classification problem in Section 3.1. This is followed by subsections on text and image representation, cross modal encoder, and optimization (Figure 2 gives an overview). Finally, in Section 3.5 we cover an alternative ranking formulation for recommendation.

3.1 Theme recommendation: classification formulation

In our setup, we are given an ad image X_i (indexed by i), and associated text denoted by S_i . Text S_i is sourced from: (i) text in ad image (OCR), (ii) questions around the ad, and (iii) Wikipedia page of the brand in the ad. Given X_i , we represent the image as a sequence of objects $x_i = \{x_{i,1}, x_{i,2}, \dots, x_{i,n}\}$ together with their corresponding regions in the image $r_i = \{r_{i,1}, r_{i,2}, \dots, r_{i,n}\}$ (details in Section 3.2). The sentence S_i is represented as a sequence of words $w_i = \{w_{i,1}, w_{i,2}, \dots, w_{i,m}\}$. Given the three sequences x_i, r_i, w_i , the objective is to recommend a keyphrase $\hat{k}_i \in \mathcal{K}$, where $\mathcal{K}$ is a pre-determined vocabulary of keyphrases. In other words, for $k \in \mathcal{K}$, the goal is to estimate the probability $\mathbb{P}(k|x_i, r_i, w_i)$, and then the top keyphrase $\hat{k}_i$ for instance i can be selected as:

$$\hat{k}_i = \arg \max_{k \in \mathcal{K}} \mathbb{P}(k|x_i, r_i, w_i). \quad (1)$$

The above classification formulation is similar to that for VQA in [30]; the difference is in the multimodal features explained below.

3.2 Text and image embeddings

Text embedding. We first use WordPiece Tokenizer [29] to convert a sentence w_i into a sequence of tokens $t_i = \{t_{i,1}, t_{i,2}, \dots, t_{i,l}\}$. Then, the tokens are projected to vectors in the embedding layer leading to $\bar{t}_i$ (as shown in (2)). Their corresponding positions p_i are also projected to vectors leading to $\bar{p}_i$ (as shown in (3)). Then, $\bar{t}_i$ and $\bar{p}_i$ are added to form $\tilde{t}_i$ as shown in (4) below:

$$\bar{t}_i = E_t * t_i, \quad (2)$$

$$\bar{p}_i = E_p * p_i, \quad (3)$$

$$\tilde{t}_i = 0.5 * (\bar{t}_i + \bar{p}_i), \quad (4)$$

where $E_t \in \mathbb{R}^{|V_t| \times D_t}$ and $E_p \in \mathbb{R}^{|V_p| \times D_p}$ are the embedding matrices. $|V_t|$ and $|V_p|$ are the vocabulary size of tokens, and token

positions. D_t and D_p are the dimensions of token and position embeddings.

Image embedding. We use bounding boxes and their region-of-interest (RoI) features to represent an image. Same as [20, 28], we leverage Faster R-CNN [24] to generate the bounding boxes and RoI features. Faster R-CNN is an object detection tool which identifies instances of objects belonging to certain classes, and then localizes them with bounding boxes. Though image regions lack a natural ordering compared to token sequences, the spatial locations can be encoded (*e.g.*, as demonstrated in [28]). The image embedding layer takes in the RoI features x_i and spatial features r_i and outputs a position-aware image embedding v_i as shown below:

$$\begin{aligned} \bar{x}_i &= W_x * x_i + b_x, \\ \bar{r}_i &= W_r * r_i + b_r, \\ v_i &= 0.5 * (\bar{x}_i + \bar{r}_i), \end{aligned} \quad (5)$$

where W_x and W_r are weights, and b_x and b_r are biases.

3.3 Transformer-based cross-modality encoder

We apply a transformer-based cross-modality encoder to learn a joint embedding from both visual and textual features. Here, without loss of generality, we follow the LXMERT architecture from [28] to encode the cross-modal features. As shown in Figure 2, the token embedding h_i is first fed into a language encoder while the image embedding v_i goes through an object-relationship encoder. The cross-modality encoder contains two unidirectional cross-attention sub-layers which attend the visual and textual embeddings to each other. We use the cross-attention sub-layers to align the entities from two modalities and to learn a joint embedding e_i of dimension D_e . We follow [11] to add a special token [CLS] to the front of the token sequence. The embedding vector learned for this special token is regarded as the cross-modal embedding¹. In terms of query (Q), key (K), and value (V), the visual cross-attention(Q, K, V) = $\text{softmax}(\frac{QK^T}{\sqrt{d}})V$ where Q, K , and V are linguistic features, visual features, and visual features, respectively; d represents the dimension of linguistic features [28]. Textual cross-attention is similar with visual and linguistic features swapped.

3.4 Learning and optimization.

Based on the joint embedding for each image and sentence pair, the keyphrase recommendation task can now be tackled with a fully connected layer. Given the cross-modal embedding, the probability distribution over all the candidate keyphrases is calculated by a fully-connected layer and the softmax function as shown below:

$$\hat{\mathbb{P}}(k|x_i, r_i, w_i) = \text{softmax}(W_f \cdot e_i + b_f) \quad (6)$$

where W_f and b_f are the weight and bias of a fully connected layer, and e_i is the cross modal embedding.

3.5 Theme recommendation: ranking formulation

We also consider solving the theme recommendation problem via a ranking model, where the model outputs a list of keyphrases in decreasing order of relevance for a given (*image, sentence*) pair,

¹Recently proposed ViLBERT [20], and VisualBERT [18] can serve as alternatives.

Figure 2: Cross modality encoder architecture, and subsequent feed forward (FF) network with softmax layer for the classification objective.

i.e., (X_i, S_i) . We use the state-of-the-art pairwise deep relevance matching model (DRMM) [13] whose architecture for our theme recommendation setup is shown in Figure 3. It is worth noting that our pairwise ranking formulation can be changed to accommodate other multi-objective or list-based loss-functions. We chose the DRMM model since it is not restricted by the length of input, as most ranking models are, but relies on capturing local interactions between query and document terms with fixed length matching histograms. Given an $(image, sentence, phrase)$ combination, the model first computes fixed length matching histogram between cross-modal embedding and the phrase embedding. Each matching histogram is passed through a multi-layer perceptron (MLP), and the overall score is aggregated with a query term gate which is a softmax function over all terms in that query.

Figure 3: DRMM for the keyphrase ranking objective.

The ranking segment of the model takes two inputs: (i) cross-modal embedding for (X_i, S_i) (as explained in Section 3.3), and (ii) the phrase embedding. It then learns to predict the relevance of the given phrase with respect to the query (*image, sentence*) pair. Given that our input documents (i.e., keyphrases) are short, we select top θ interactions in matching histogram between the cross-modal embedding, and the keyphrase embedding. Mathematically,

we denote the $(image, sentence)$ pair by just the $imgq$ below. Given a triple $(imgq, p^+, p^-)$ where p^+ is ranked higher than p^- with respect to image-question, the loss function is defined as:

$$\mathcal{L}(imgq, p^+, p^-; \theta) = \max(0, 1 - s(imgq, p^+) + s(imgq, p^-)), \quad (7)$$

where $s(imgq, p)$ denotes the predicted matching score for phrase p , and the query image-question pair.

4 EXPERIMENTS

In this section, we go over the public dataset used in our experiments, classification and ranking results, and inferred insights.

4.1 Dataset

We rely on a publicly available data set [1, 15] that consists of 64,000 advertisement creatives, spanning 700 brands across 39 categories, among which 80% is training set and 20% is test set. We select 10% data from the training set for validation. Crowdsourcing was used to gather following labels for each creative: (i) topics (39 types), (ii) questions and answers as reasons for buying from the brand depicted in the creative (~ 3 per creative). In addition to the existing annotations, we add the following annotations: (i) brand present in a creative, (ii) Wikipedia page relevant to the brand-category pair in a creative, and (iii) the set of target themes (keyphrases) associated with each image. In particular, for (i) and (ii) we follow the method in [22], and for (iii) the keyphrases (labels) were extracted from the answers using the position-rank method [10, 12] for each image. The number of keyphrases was limited to at most 5 (based on the top keyphrase scores returned by position-rank). We define a score for each keyphrase. All five keyphrases have scores of 1.0, 0.9, 0.8, 0.7, and 0.6 in order ².

The minimum, mean and maximum number of images associated with a brand are 1, 19 and 282 respectively. The top three categories of advertisements are *clothing*, *cars* and *beauty products* with 7798, 6496 and 5317 images respectively. Least number of advertisements are associated with *gambling* (32), *pet food* (37) and *security and safety services* (47) respectively. Additional statistics around the

²The annotated ads dataset can be found at <https://github.com/joey1993/ad-themes>.

dataset (i.e., keyphrase lengths, images per category, and unique keyphrases per category) are shown in Figures 4, 5, and 6.

Figure 4: Frequency and length distribution of keyphrases.

Figure 5: Distribution of images per category (top 10 categories by count on the right, bottom 10 on left).

Figure 6: Distribution of unique phrases per category (top 10 categories by count on the right, bottom 10 on left).

4.2 Evaluation Metrics

We use different evaluation metrics to measure performance of our classification and ranking models. We use three different metrics to evaluate the performance of each model.

4.2.1 Classification metrics. We rely on accuracy, text similarity and set-intersection based recall to evaluate model performance.

- **Accuracy.** We predict the keyphrase with the highest probability for each image and match it with the labels (ground truth keyphrases) for the image. We use the score of the matched phrase as the accuracy. If no labels match for a sample, the accuracy is 0. We average the accuracy scores over all the test instances to report test accuracy.
- **Similarity:** Accuracy neglects the semantic similarity between the predicted phrase and the labels. For example, a predicted keyphrase “a great offer” is similar to one of the labels, “great sale”, but will gain 0 for accuracy. So we calculate the cosine similarity [14] between the embeddings of

the predicted keyphrase and each label. Then we multiply the similarity scores with each label’s score and keep the maximum as the final similarity score for the sample.

- **VQA Recall@3:** we use Recall at 3 ($R_{VQA}@3$) as an evaluation metric for the classification task (essentially like the VQA formulation in [30]). For each test instance i , the ground truth is limited to top 3 keyphrases leading to set $\mathcal{K}_i^*$. From the classification model’s predictions the top 3 keyphrases are chosen leading to set $\hat{\mathcal{K}}_i$. $R_{VQA}@3$ is simply $\frac{|\mathcal{K}_i \cap \mathcal{K}_i^*|}{3}$.

4.2.2 Ranking metrics. We use the same evaluation metrics from prior work [22], mainly precision (P@K), recall (R@K), and NDCG@K [16] to evaluate the proposed ranking model. It is worth noting that recall is computed differently for evaluating ranking and classification models proposed in this work. Formally, given a set of queries $Q = \{q_1 \dots q_n\}$, set of phrases $\mathcal{D}_i$ labeled *relevant* for each query q_i and the set of relevant phrases $\mathcal{D}_{ik}$ retrieved by the model for q_i at position k , we define $R@K = \frac{1}{N} \cdot \sum_{i=1}^N \frac{|\mathcal{D}_{ik}|}{|\mathcal{D}_i|}$.

4.3 Implementation details

For the classification model, we set the number of object-relationship, language, and cross-modality layers as 5, 9, 5, and leverage pre-trained parameters from [28]. We fine-tune the encoders with our dataset for 4 epochs. The learning rate is $5e^{-5}$ (adam optimizer), and the batch size is 32. We also set D_t , D_p , and D_e equal to 768. For the similarity evaluation, we average the GloVe [23] embeddings of all the words in a phrase to calculate the phrase embedding. For the DRMM model (ranking formulation), we used the MatchZoo implementation [5], with 300 for batch size, 10 training epochs, 10 as the last layer’s size, and learning rate of 1.0 (adadelta optimizer). We combine textual data from different sources in a consistent order separated by a [SEP] symbol before feeding them into the encoder.

features	accuracy (%)	similarity (%)	$R_{VQA}@3$
I	10.05	58.05	0.447
I×Q	12.18	58.26	0.450
I×(Q+W)	19.01	60.12	0.467
I×(Q+O)	19.50	60.34	0.470
I×(Q+W+O)	20.40	60.95	0.473
Q+W+O	13.39	60.13	0.450
non cross-modal	18.65	60.68	0.460

Table 1: Classification performance with different features (I: image, Q: question, W: Wikipedia page of associated brand, and O: OCR, i.e., text in ad image, ×: cross-modal representation); non cross-modal denotes using an addition of separate visual and linguistic features.

4.4 Results

For different sets of multimodal features, the performance results are reported in Table 1 (for classification) and Table 2 (for ranking) respectively. The presence of Wikipedia and OCR text gives a significant lift over using only the image. Both classification and ranking metrics show the same trend in terms of feature sets.

Table 1 shows that linguistic features dramatically lift the performance by 103%, 5%, and 6% in accuracy, similarity, and $R_{VQA}@3$, compared to the performance of the model trained *only* with visual

Figure 7: Performance lifts across different categories after using text features. For accuracy, the lift is scaled (divided by 10) for better visualization with the similarity lift.

features while only using linguistic features (Q+W+O) causes a big drop in all the performances. It occurs that the OCR features bring more performance lift, compared to the Wiki. We think knowing more about the brand with the Wikipedia pages is beneficial to recommend themes to designers [22] while the written texts on the images (OCR) are sometimes more straightforward for recommendations. In addition, as reported in Table 1 (non cross-modal), using separate text and image embeddings (obtained from the model in Figure 2) is inferior in performance compared to the cross-modal embeddings. We notice that the accuracy scores are comparatively low; this reflects the difficult nature of understanding visual ads [15].

In Table 2, we observe very similar patterns: OCR features bring more benefits to ranking than Wikipedia pages. We notice that in P@10 and R@5, only using image feature (I) achieves a better score compared to adding the question features. This may indicate that the local interactions in DRMM are not effective with short questions, but favor longer textual inputs such as OCR and Wikipedia pages.

Features	Precision		Recall		NDCG	
	@5	@10	@5	@10	@5	@10
I	0.150	0.126	0.161	0.248	0.158	0.217
I×Q	0.152	0.124	0.158	0.259	0.162	0.227
I×(Q+W)	0.154	0.130	0.160	0.271	0.161	0.234
I×(Q+O)	0.174	0.137	0.182	0.287	0.185	0.254
I×(Q+W+O)	0.183	0.141	0.191	0.294	0.198	0.265

Table 2: Ranking performance with different features (I: image, Q: question, W: text from brand Wikipedia page, and O: OCR text in the ad image, ×: cross-modal representation).

4.5 Insights

Figure 7 shows the performance lifts in accuracy and similarity metrics per category (where lift is defined as ratio of improvement to baseline result without using text features in the classification task). As shown, multiple categories, *e.g.*, public service announcement (PSA) ads around domestic violence and animal rights benefit from the presence of text features; this may be related to the hidden symbolism [30] common in PSAs, where the text can help clarify the context even for humans. Also, similarity and accuracy metrics do not have the same trend in general. Along the lines of inferring

themes from past ad campaigns, and assisting strategists towards designing new creatives, we show an example based on our classification model in Figure 8. In general, a strategist can aggregate recommended keyphrases across a brand or product category, and use them to design new creatives.

Figure 8: The ad image on the left is a sample in the public dataset [1], and the ground truth keyphrases with scores are as shown. In the classification setup, using only the image has zero accuracy, while using image + text features leads to perfect accuracy. The predicted keyphrases can be used as recommended queries to a stock image library [6] (as shown on the right) to obtain new creatives.

5 CONCLUSION

In this paper, we make progress towards automating the inference of themes (keyphrases) from past ad campaigns using multimodal information (*i.e.*, images and text). In terms of model accuracy, there is room for improvement, and using generative models for keyphrase generation may be a promising direction. In terms of application, *i.e.*, automating creative design, we believe that the following are natural directions for future work: (i) automatically selecting ad images and generating ad text based on recommended themes, and (ii) launching ad campaigns with new creatives (designed via our proposed method) and learning from their performance in terms of CTR and CVR. Nevertheless, the proposed method in this paper can increase diversity in ad campaigns (and potentially reduce ad fatigue), reduce end-to-end design time, and enable faster exploratory learnings from online ad campaigns by providing multiple themes per brand (and multiple images per theme via stock image libraries).

REFERENCES

- [1] 2019. Automatic Understanding of Image and Video Advertisements. <http://people.cs.pitt.edu/~kovashka/ads>.
- [2] 2019. Banner blindness. https://en.wikipedia.org/wiki/Banner_blindness.
- [3] 2019. Facebook business: Optimize your ad results by refreshing your creative. <https://www.facebook.com/business/m/test-ads-on-facebook>.
- [4] 2019. Marketing Land: Social media ad fatigue. <https://marketingland.com/ad-fatigue-social-media-combat-224234>.
- [5] 2019. Match Zoo. <https://github.com/NTMC-Community/MatchZoo>.
- [6] 2019. Shutterstock: Search millions of royalty free stock images, photos, videos, and music. <https://www.shutterstock.com/>.
- [7] 2019. Taboola-trends. <https://trends.taboola.com/>.
- [8] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. VQA: Visual Question Answering. In *The IEEE International Conference on Computer Vision (ICCV)*.
- [9] Narayan Bhamidipati, Ravi Kant, Shaunak Mishra, and Mingzhu Zhu. 2017. A Large Scale Prediction Engine for App Install Clicks and Conversions. In *CIKM 2017*.
- [10] Florian Boudin. 2016. pke: an open source python-based keyphrase extraction toolkit. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: System Demonstrations*.
- [11] 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).
- [12] Corina Florescu and Cornelia Caragea. 2017. PositionRank: An Unsupervised Approach to Keyphrase Extraction from Scholarly Documents. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- [13] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016. A Deep Relevance Matching Model for Ad-hoc Retrieval. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*.
- [14] Jiawei Han, Jian Pei, and Micheline Kamber. 2011. *Data mining: concepts and techniques*. Elsevier.
- [15] Zaem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. 2017. Automatic Understanding of Image and Video Advertisements. In *CVPR*.
- [16] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.
- [17] Gen Li, Nan Duan, Yuejian Fang, Daxin Jiang, and Ming Zhou. 2019. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. *arXiv preprint arXiv:1908.06066* (2019).
- [18] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557* (2019).
- [19] Wei Li, Xuerui Wang, Ruofei Zhang, Ying Cui, Jianchang Mao, and Rong Jin. 2010. Exploitation and Exploration in a Performance Based Contextual Advertising System. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [20] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In *NeurIPS*.
- [21] H. Brendan McMahan, Gary Holt, D. Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, Sharat Chikkerur, Dan Liu, Martin Wattenberg, Arnar Mar Hrafnkelsson, Tom Boulos, and Jeremy Kubica. [n.d.]. Ad Click Prediction: a View from the Trenches (*KDD 2013*).
- [22] Shaunak Mishra, Manisha Verma, and Jelena Gligorijevic. 2019. Guiding Creative Design in Online Advertising (*RecSys*).
- [23] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- [24] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.
- [25] Susanne Schmidt and Martin Eisend. 2015. Advertising Repetition: A Meta-Analysis on Effective Frequency in Advertising. *Journal of Advertising* 44, 4 (2015), 415–428.
- [26] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Proceedings of ACL*.
- [27] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2019. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530* (2019).
- [28] Hao Tan and Mohit Bansal. 2019. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. In *EMNLP-IJCNLP*.
- [29] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [30] Keren Ye and Adriana Kovashka. 2018. ADVISE: Symbolism and External Knowledge for Decoding Advertisements. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XV*. 868–886.
- [31] Seounmi Youn and Seunghyun Kim. 2019. Newsfeed native advertising on Facebook: young millennials’ knowledge, pet peeves, reactance and ad avoidance. *International Journal of Advertising* 38, 5 (2019), 651–683.
- [32] Yichao Zhou, Shaunak Mishra, Jelena Gligorijevic, Tarun Bhatia, and Narayan Bhamidipati. 2019. Understanding Consumer Journey using Attention based Recurrent Neural Networks. *KDD* (2019).