VisualTextRank: Unsupervised Graph-based Content Extraction for Automating Ad Text to Image Search

Shaunak Mishra

Yahoo Research

shaunakm@verizonmedia.com

Gaurav Srivastava

Yahoo Research

gaurav.srivastava@verizonmedia.com

Mikhail Kuznetsov

Yahoo Research

kuznetsov@verizonmedia.com

Maxim Sviridenko

Yahoo Research

sviri@verizonmedia.com

ABSTRACT

Numerous online stock image libraries offer high quality yet copyright free images for use in marketing campaigns. To assist advertisers in navigating such third party libraries, we study the problem of automatically fetching relevant ad images given the ad text (via a short textual query for images). Motivated by our observations in logged data on ad image search queries (given ad text), we formulate a keyword extraction problem, where a keyword extracted from the ad text (or its augmented version) serves as the ad image query. In this context, we propose VisualTextRank: an unsupervised method to (i) augment input ad text using semantically similar ads, and (ii) extract the image query from the augmented ad text. VisualTextRank builds on prior work on graph based context extraction (biased TextRank in particular) by leveraging both the text and image of similar ads for better keyword extraction, and using advertiser category specific biasing with sentence-BERT embeddings. Using data collected from the Verizon Media Native (Yahoo Gemini) ad platform's stock image search feature for onboarding advertisers, we demonstrate the superiority of VisualTextRank compared to competitive keyword extraction baselines (including an 11% accuracy lift over biased TextRank). For the case when the stock image library is restricted to English queries, we show the effectiveness of VisualTextRank on multilingual ads (translated to English) while leveraging semantically similar English ads. Online tests with a simplified version of VisualTextRank led to a 28.7% increase in the usage of stock image search, and a 41.6% increase in the advertiser onboarding rate in the Verizon Media Native ad platform.

CCS CONCEPTS

• Information systems → Online advertising.

KEYWORDS

Online advertising; image search; multilingual; content extraction.

This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike International 4.0 License.

KDD '21, August 14–18, 2021, Virtual Event, Singapore.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8332-5/21/08.

<https://doi.org/10.1145/3447548.3467126>

ACM Reference Format:

Shaunak Mishra, Mikhail Kuznetsov, Gaurav Srivastava, and Maxim Sviridenko. 2021. VisualTextRank: Unsupervised Graph-based Content Extraction for Automating Ad Text to Image Search. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21)*, August 14–18, 2021, Virtual Event, Singapore. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3447548.3467126>

1 INTRODUCTION

Online advertising is an effective tool for both major brands, and small businesses to grow awareness about their products and influence online users towards purchases [3, 23]. Ad images are a crucial component of online ads, and small businesses with limited advertising budget may find it challenging to obtain relevant high quality images for their ad campaigns. In this context, small business are increasingly relying on online stock image libraries [2] which offer access to high quality and copyright free images available for marketing purposes. To assist such small businesses in a self-serve manner, the Verizon Media Native (Yahoo Gemini) ad platform offers onboarding advertisers the ability to query a third party stock image library for candidate ad images (as shown in Figure 1).

Although the example (an ad for bikes) shown in Figure 1 is borrowed from Verizon Media Native, it brings up a fundamental question: can we automatically figure out the search query for images given the ad text? Such automation will not only attract advertisers towards the stock image library, but will also speed up the onboarding process. However, there are several challenges in this context: (i) limited data (for supervised learning) on ad text and associated image search queries, and (ii) third party image libraries with unknown (proprietary) image indexing. The first challenge, *i.e.*, limited data stems from the fact that this is a relatively new service (*i.e.*, ad platforms offering stock image search feature), and small businesses may not be aware about such features. One can use current state-of-the-art image understanding methods (*i.e.*, object detection [10], captioning [21]) on the large set of existing ad images (which were not created by querying a stock image library) to get image annotations given ad text. But in our analysis (details in Section 3), such annotations were not specific enough to be considered as image queries in most cases, and hence the problem of limited ad text-to-image query data for supervised learning still remains. The second challenge listed above, *i.e.*, unknown indexing, makes image libraries act like black boxes allowing short (few words) queries which may often be restricted to the English language. Due to such

Figure 1: Onboarding workflow for Verizon Media Native advertisers: (1) advertiser enters their website URL with ad title and description (ad text), (2) advertiser proceeds to choose an ad image and may be shown a random selection of images from a third party image library, (3) advertiser queries the library with a relevant textual query, and (4) selects a query result to get a preview of the final ad.

a black-box nature of third party image libraries, we do not consider approaches requiring the mapping of a text query to an image.

Given the above challenges, in this paper, we focus on unsupervised approaches for figuring out the ad image search query given an ad text. As we explain later in this paper (in Section 3), available logged data on ad image search behavior shows that in a significant fraction of cases, the ad text may already contain a keyword suitable as the image search query (as in the bike ad example in Figure 1); in other words the image query is *extractive* in nature. In the remaining cases, the ideal image search query may be symbolic (abstractive) with respect to the input ad text. Citing a real example from logged data, the (anonymized) ad 'Best ETFs to buy ...' had image search queries: 'resort' and 'private jet', both of which are indicative (symbolic) of wealth after plausibly investing in the advertised financial products. Although we mostly focus on methods for extractive queries in this paper, we report preliminary findings on the usage of symbolic image queries given ad text.

Focusing on extractive image queries, we build on prior work on unsupervised graph based keyword extraction [4, 9] (none of which are specifically designed for ad text). Such graph based methods typically build a graph where nodes are tokens (e.g., words) from input text, and the weighted edges between the nodes represent similarity between the associated tokens [9]. With such a graph built from tokens, PageRank style algorithms [14] are used to infer the dominant token (keyword); the popularity of such methods spans more than a decade. Recently, biased TextRank [9] was proposed to leverage sentence-BERT (SBERT) [19] embeddings of the

Figure 2: Illustrative example where an existing similar ad has image query hints about two input ads (one in English and one in French) advertising furniture. The image tags are objects detected in the raw ad image but may not be descriptive enough to serve as image queries in many cases.

input text to adjust the bias for each node (random restart probability in Page Rank). In other words, tokens which are closer to the overall *meaning* of the entire text (as encoded by SBERT) are likely to have higher bias; this works significantly better than TextRank. In our setup focused on ad text (and ad image queries), we improve upon biased TextRank using two key ideas. Firstly, we retrieve similar existing ads to augment the (relatively short) input ad text. The augmentation is done using the text of similar ads as well as image tags (detected objects) in the images of similar ads in an unsupervised manner. Figure 2 illustrates this idea: for both the English and French ad text in the example, a semantically similar English ad already contains hints (in the image and text), that the original ad is about furniture (which is also the ground truth image query in the example for both ads). The above ideas not only make VisualTextRank ads specific, but also significantly push performance on the ad text-to-image query task. Our main contributions can be summarized as follows.

(1) *Keyword extraction for ad text-to-image query.* We formulate ad text-to-image query as a keyword extraction problem. For this task, we propose an unsupervised graph based content extraction method (VisualTextRank) which builds on biased TextRank [9] by introducing category biasing, and input text augmentation using the text and images of similar (existing) ads. For our task, we obtain ~ 11% lift in offline accuracy using VisualTextRank compared to competitive keyword extraction baselines.

(2) *Cross-lingual learning.* Leveraging VisualTextRank, we extend the ad-text-to-image task for non-English ad text (including German, French, Spanish, and Portuguese) and English-only image queries. For such a setup, we demonstrate the benefit of using semantically similar English ads to augment the non-English ad text for the ad-text-to-image query task.

(3) *Product impact.* We productionized a light-weight version of VisualTextRank for online tests with Verizon Media Native advertisers qualifying as small businesses (e.g., with relatively low advertising budgets). For advertisers with English ad text, we compared the effect of automatically showing ad text-relevant images

(via VisualTextRank queries) as initial (default) images to the advertiser versus showing random stock images. Against this baseline, we observed a 28.7% lift in the rate of advertisers selecting stock images, and a 41.6% lift in the onboarding completion rate for advertisers.

The remainder of the paper is organized as follows. Section 2 covers related work, and Section 3 covers data analysis. The proposed VisualTextRank method and its multilingual extension is covered in Sections 4 and 5. Experimental results are covered in Section 6, and we end with paper with a discussion in Section 7.

2 RELATED WORK

2.1 Keyword extraction

In this paper, we formulate the ad text-to-image query task as an unsupervised keyword extraction problem (explained in Section 4). Unsupervised keyword extraction research has a long history spanning over a decade [4]. In this paper, we focus on graph based methods which are inspired from PageRank on a graph created from the input text [4, 14]. TextRank is a widely used method in this class, with applications spanning keyword extraction, summarization, and opinion mining [9, 14]. A notable recent work is by [9], where the authors improve TextRank by introducing sentence-BERT [19] based biasing leading to the biased TextRank method. Sentence-BERT (SBERT) [19] is a sentence representation method based on BERT [5]. In particular, SBERT is a modification of pre-trained BERT that uses siamese and triplet network structures to derive semantically meaningful sentence embeddings which can be compared using cosine-similarity. This provides a computationally efficient way to achieve state-of-the-art results for sentence-pair regression tasks like semantic textual similarity. Our proposed VisualTextRank method builds on top of biased TextRank [9] with a focus on ad text, and the ad text-to-image query task.

2.2 Textual description of images: object detection and image captioning

Object detection methods (image tagging) predict the category and location details of the objects present in an image. Faster R-CNN [20] is a popular method for this task, and more recent works built on top of it [18]. Compared to image tagging which provides a bag-of-words style textual description (objects tags with confidence scores) of an image, captioning models generate a natural language description; this is a much harder task than object detection. Current state-of-the-art image captioning models use a pre-trained object detector to generate features and spatial information of the objects present in an image (e.g., as in Oscar [12]). In particular, [12] utilizes BERT-like objectives to learn cross-modal representation on different vision-language tasks (similar ideas form the basis of recent pre-trained multi-modal models, e.g., VisualBERT [11]). Prior captioning approaches have involved attention mechanisms and their variants to capture spatial relationship between objects [6] for generating captions. In this paper, we explore both object detection (image tagging) and captioning as features for assisting the ad text-to-image query task. In addition, since we focus on an unsupervised approach, we do not consider the possibility of fine-tuning the above mentioned pre-trained multi-modal models for our task.

2.3 Ad image and text understanding

Studying ad images and text using state-of-the-art deep learning models in computer vision and natural language processing (NLP) is an emerging area of research. In [8], ad image content was studied using computer vision models, and their dataset had manual annotations for: ad category, reasons to buy products advertised in the ad, and expected user response given the ad. Using this dataset, [16, 24] used ranking models to recommend themes for ad creative design using a brand’s Wikipedia page. In [17], object tag recommendations for improving an ad image was studied using data from A/B tests. Although related to ads, the above methods are not applicable in our setup due to the lack of sufficient (ad text-to-image query) data for supervised training.

3 DATA INSIGHTS

In this section, we cover data insights around ad image search behavior which guided our proposed approach for the ad text-to-image task. We first explain the data source for the analysis, followed by metadata from image understanding models (for the purpose of analysis), and the resultant insights.

3.1 Data source

To gather preliminary insights around ad text-to-image search behavior, we collected data from a sample of ~ 300 advertisers who used the stock image library feature for their ad image while launching ad campaigns (onboarding) in the Verizon Media Native (Yahoo Gemini) ad platform. For each advertiser, the data included: (i) ad text, (ii) raw ad image, and (iii) image query for the ad image. We will refer to this dataset as onboarded-sample.

3.2 Image tags and caption metadata

For the purposes of analyzing onboarded-sample for insights, we used image tags, and captions as described below. We obtained image tags using the pre-trained Inception Resnet v2 object detection model [10]. The model consists of about 5,000 classes (possible tags). For each image the model returns a list of inferred tags with confidence scores (we used tags with confidence above 0.8). For captioning, we used an object relation transformer (ORT) model [6] with a Faster R-CNN object detector. For our analysis, we trained two ORT models: one on Microsoft COCO 2014 Captions dataset [13] (COCO-captions model), and the other on Conceptual Captions dataset [21] (CC-captions model).

3.3 Insights

3.3.1 Query length, parts-of-speech, and ad text overlap. We found that 81% of image queries were a single word, and the remaining 19% of queries comprised of two or more words. For English ads in onboarded-sample, we found an overwhelming 92% of queries to be either nouns or proper nouns; verbs accounted for 3%. The parts-of-speech (POS) tags were inferred via Spacy [7]. Finally, we found that 60% of image queries were already present (as a word or as a phrase) in the ad text. This indicates majority of the image queries are extractive in nature; 40% of remaining queries are symbolic in nature (e.g., an ad about retirement investments with image queries like ‘vacation’ and ‘resort’ which are not part of the ad text).

Figure 3: Examples of ad images with captions (COCO), tags, and ground truth image query.

Supported by the above observations on query length and overlap with ad text, we formulate the ad text-to-image task as a keyword extraction task given ad text (and additional side information as explained later in Section 4.1). For our experiments, we restrict the output to be either a noun, proper noun or verb (guided by the POS tags observation above).

3.3.2 Poor correlation with image captions (but better with image tags). The idea of using image understanding models stems from the observation that there is a huge dataset of existing ad text and images which were not created via a stock image query (e.g., the public dataset of older ads in [22], or proprietary older ads in the ad platform). If it is possible to get descriptions for existing ad images via an image understanding tool (e.g., captions [21], object detectors [10]), one can treat such a description as a proxy image query, and train an end-to-end model for ad text-to-image query. To test this hypothesis, we analyzed the captions and image tags of onboarded ad images and compared them with the ground truth image queries. Checking for image query overlap with captions for the raw ad image, we found that there was an overlap in 5% of samples for the COCO-captions model and an overlap in 7.6% of samples for the CC-captions model. Figure 3 shows two examples of ad images from the onboarded sample with captions (COCO), tags and the ground truth query; clearly, captioning failed in both examples, while tagging did better in one. In comparison to captions, for the image tag model the overlap with image queries was 17%. Due to higher overlap, we retained image tags for existing ads (used in our experiments), and did not focus on captions. Note that with 17% overlap, image tags are still not expressive enough (just 5,000 labels) to create a labeled dataset for training an ad-text-to-image query predictor; but they can plausibly be used as a signal for better keyword extraction. Table 1 provides a brief summary of the insights.

feature	image query insights
POS tags	92% NOUN/PROPN, 3% VERB
query length	81% of queries are a single word
ad text	60% of queries are within the ad text
image captions (COCO)	5% query overlap
image captions (CC)	7.6% query overlap
image tags	17% query overlap

Table 1: Summary of insights from onboarded-sample data.

4 KEYWORD EXTRACTION FOR AD TEXT-TO-IMAGE QUERY

In this section, we first formulate ad text-to-image query as a keyword extraction task in Section 4.1. Next, we describe the TextRank method as necessary background in Section 4.2. Finally, we explain our proposed method VisualTextRank in Section 4.3.

4.1 Problem Formulation

Given an ad text a_{in} , our objective is to come up with a single keyword to serve as the ad image query (to a stock image library). We assume the presence of a pool of existing ads P , where each ad $a_i \in P$ has the following attributes: (i) ad text, and (ii) image tags in the raw ad image (as described in Section 3.3.2).

In this paper, we focus on a graph based unsupervised method to model the ad text (and information relevant to a_{in} in P), for the purpose of extracting a word (from the union of a and P) as a relevant ad image query given a_{in} . The key motivation behind a graph based approach as opposed to just using the image tags from the most similar ad text (i.e., a nearest neighbor approach), lies in Table 1. The intuition here is to identify the *central* entity in the ad text (which contains the query in 60% of cases as noted in Table 1), with additional help from similar ads. In the following sections, we first explain an existing graph based keyword extraction method (TextRank) in our context as necessary background, and then explain the proposed VisualTextRank method.

4.2 TextRank

At a high level, TextRank is primarily based on PageRank [14] on the graph of tokens (words) obtained from the input text. We explain below, the token graph construction, the (original) unbiased version of TextRank as well as the recently proposed biased version.

Token graph. The token-graph $\mathcal{G} = (V, E)$ is formed from tokens (words) from the input ad text a_{in} , where V denotes the set of vertices and E denotes the set of (undirected) edges. Each word w_i is mapped to a vertex v_i . An edge between two vertices has weight:

$$e_{ij} = \langle \phi(w_i), \phi(w_j) \rangle. \quad (1)$$

where $\langle \phi(w_i), \phi(w_j) \rangle$ denotes the cosine similarity between the (word) embeddings of w_i and w_j (denoted by $\phi(w_i)$ and $\phi(w_j)$ respectively). It is common to set e_{ij} to zero if it is below a similarity threshold [9].

Unbiased TextRank. The original (unbiased) TextRank method [14] iteratively computes the score v_i of each vertex $i \in V$ as:

$$v_i = (1 - d) + d \times \left(\sum_{j \in \mathcal{N}(i)} \frac{e_{ij}}{\sum_{k \in \mathcal{N}(j)} e_{kj}} \times v_j \right), \quad (2)$$

where d is the damping factor (typically 0.85), and $\mathcal{N}(i)$ is the set of vertices which share an edge (with non-zero edge weight) with vertex i .

Biased TextRank. The recently proposed biased TextRank method [9] iteratively computes the score v_i of each vertex in the token-graph in the following manner:

$$v_i = \text{bias}_i \times (1 - d) + d \times \left(\sum_{j \in \mathcal{N}(i)} \frac{e_{ij}}{\sum_{k \in \mathcal{N}(j)} e_{kj}} \times v_j \right), \quad (3)$$

where the only change with respect to the unbiased version in (2) is the addition of a bias term bias_i for vertex i . For the task of keyword extraction, in [9] bias bias_i is defined as:

$$\text{bias}_i = \langle \phi(\mathbf{a}_{in}), \phi(\mathbf{w}_i) \rangle, \quad (4)$$

where $\phi(\mathbf{a}_{in})$ is the sentence embedding for input ad text $\mathbf{a}_{in}$. Intuitively, biased TextRank tries to favor words which are closer to the overall (semantic) meaning of the input ad text. In [9], the authors demonstrate that this intuition works well and show the superiority of biased TextRank over the original TextRank method. In the remainder of the paper, we will refer to the above biasing method (as in (4)) as self-biasing since it computes the bias with respect to the input (ad text). Sentence-BERT (SBERT) [19] is an effective sentence embedding method used in biased TextRank [9], and we leverage the same in our proposed method (explained below).

4.3 VisualTextRank

We start with an overview of the proposed VisualTextRank method in Section 4.3.1, and then go over the details of important components in Sections 4.3.3 and 4.3.2.

4.3.1 Overview. The biased TextRank method (3) has the following shortcomings when it comes to understanding an ad (*i.e.*, extracting a keyword for image search):

- (1) ads are usually tied to categories (*e.g.*, IAB category taxonomy [1]) and TextRank is oblivious to this, and
- (2) ad text can be short leading to a very sparse token graph with no or negligible edges above a reasonable similarity threshold.

In addition to the above shortcomings from a keyword extraction perspective, TextRank’s original motivation was not to extract queries suitable for ad images, and hence it lacks the visual aspect needed in an image search query (in both extractive and symbolic cases). In VisualTextRank, we build on the above shortcomings as outlined below (detailed explanation after the outline).

- (1) Category biasing: we introduce an ad category specific biasing term in addition to the self-biasing term in (4).
- (2) Augmentation with similar ads’ text: for a given input ad, we fetch semantically similar ads from a pool of existing ads. We augment the token-graph with a filtered version of the text from similar ads.
- (3) Augmentation with similar ads’ image tags: we also augment the token-graph with a filtered version of the image tags (as described in Section 3.3.2) of similar ads’ images.

Steps (2) and (3) above, not only alleviate the problem of short ad text, but also offer the capability of going beyond words in the ad text while coming up with the image query. As a result of augmenting the initial token-graph $\mathcal{G}$ (as described in Section 4.2) with text and image tags from similar ads (augmentation details

described later in Section 4.3.2), we obtain an augmented token-graph $\mathcal{G}^* = (\mathbf{V}^*, \mathbf{E}^*)$.

Combining the steps above, the vertex value update for VisualTextRank can be written as follows.

$$v_i = \text{bias}_i^{(self)} \times \text{bias}_i^{(cat)} \times (1 - d) + d \times \left(\sum_{j \in \mathcal{N}^*(i)} \frac{e_{ij}}{\sum_{k \in \mathcal{N}^*(j)} e_{kj}} \times v_j \right), \quad (5)$$

where i is a vertex in the augmented graph $\mathcal{G}^*$, and $\mathcal{N}^*(i)$ denotes the set of vertices i shares an edge with in $\mathcal{G}^*$. The details behind the category biasing term $\text{bias}_i^{(cat)}$ are described in Section 4.3.3, and the details behind the graph augmentation leading to $\mathcal{G}^*$ are described in 4.3.2.

4.3.2 Token-graph augmentation with similar ads. Token graph augmentation has three steps as described below.

Retrieving similar ads. For augmentation, we assume a pool (set) of existing ads denoted by $\mathbf{P}$. An ad in $\mathbf{P}$ has the following attributes: (i) ad text, and (ii) image tags for the ad image along with confidence scores for each tag. For each ad in the pool $\mathbf{P}$, we compute the semantic similarity (relevance) of the ad with respect to the input ad $\mathbf{a}_{in}$ as follows:

$$\text{relevance}(\mathbf{a}, \mathbf{a}_{in}) = \langle \phi(\mathbf{a}), \phi(\mathbf{a}_{in}) \rangle, \quad (6)$$

where $\mathbf{a}$ is an ad in $\mathbf{P}$, and $\text{relevance}(\mathbf{a}, \mathbf{a}_{in})$ denotes its relevance with respect to the input ad $\mathbf{a}_{in}$. A sentence embedding method like SBERT [19] (trained for semantic textual similarity) can be used to obtain $\phi(\mathbf{a})$ and $\phi(\mathbf{a}_{in})$. We use the above relevance score to obtain the top- m similar ads from the pool $\mathbf{P}$ for the given input ad $\mathbf{a}_{in}$.

Augmentation with similar ads’ image tags. We augment the token-graph $\mathcal{G}$ with the image tags of similar ads’, by selecting image tags which are semantically close to a word in the ad text of the similar ad. For example, as shown in Figure 4, if the similar ad has a word like ‘furniture’, and the corresponding ad image has tags like ‘chair’ and ‘table’ (which are semantically close to furniture), we select such tags for augmenting the original token-graph $\mathcal{G}$. The details of our proposed image tag augmentation can be described as follows. We assume a bound on: (i) the maximum number of image tags that we can add ($= \text{max_tags}$), and (ii) the minimum cosine similarity between an image tag and a word in its parent ad text to select the image tag for augmentation ($= \text{min_tag_sim}$). As described in Algorithm 1, we start with a list of similar ads $\{\mathbf{a}_1, \dots, \mathbf{a}_m\}$ sorted in decreasing order of relevance to the input ad text $\mathbf{a}_{in}$. For each similar ad $\mathbf{a}_i$, $\mathcal{T}(\mathbf{a}_i)$ denotes the list of image tags associated with the corresponding ad image; this tags list is sorted in decreasing order by the confidence scores from the object detector. Given $\mathbf{a}_i$ and $\mathcal{T}(\mathbf{a}_i)$, we select a tags $t \in \mathcal{T}(\mathbf{a}_i)$ which are close to at least one word w in the $\mathbf{a}_i$ (*i.e.*, $\langle \phi(t), \phi(w) \rangle \geq \text{min_tag_sim}$). We keep iterating as shown in Algorithm 1, till we collect max_tags images tags in $\mathcal{T}^*$.

Augmentation with similar ads’ text. In addition to image tags, we also use words from similar ads to augment the token-graph $\mathcal{G}$. Our proposed word augmentation method is described in Algorithm 2. It is similar in spirit to the image tag augmentation method. We keep iterating over words in an ad text $\mathbf{a}_i$ in their order of

Figure 4: Illustrative example of VisualTextRank in action for a furniture ad. The green nodes in the token graph are augmented using the similar ad’s text and image tag, and eventually lead to ‘furniture’ being selected (as they form a distinct cluster with other furniture related words).

Data: similar ad texts $a_1, \dots, a_m$ (sorted by decreasing relevance to input), image tags list $\mathcal{T}(a_i)$ for ad a_i

Result: image tags set $\mathcal{T}^*$ to augment $\mathcal{G}$, $|\mathcal{T}^*| = \max_tags$
 $\mathcal{T}^* = \emptyset$;

```

for  $i \in \{1, \dots, m\}$  do
  for  $t \in \mathcal{T}(a_i)$  do
    if  $\max_{w \in a_i} (\langle \phi(t), \phi(w) \rangle) \geq \min\_tag\_sim$  then
       $\mathcal{T}^* = \mathcal{T}^* \cup t$ ;
      if  $|\mathcal{T}^*| = \max\_tags$  then
        return  $\mathcal{T}^*$ 
      end
    end
  end
end
end

```

Algorithm 1: Augmentation with similar ads’ image tags.

occurrence, and select words which are above the $\min_word_sim$ similarity threshold (cosine similarity between the input ad text a_{in} and the candidate word w using SBERT embeddings). The set of selected words $\mathcal{W}^*$ is used to augment $\mathcal{G}$. In the example in Figure 4, the word ‘furniture’ from the similar ad’s text was selected for augmentation using the above method.

4.3.3 Category biasing. We consider a set of *category* phrases C . Such a set can be created by using phrases from a standard (flat-tened) ad category taxonomy, e.g., IAB categories [1] (see Appendix for details). For example, the set C can contain words and phrases like ‘travel’, ‘legal services’, and ‘fashion’ which are used to denote different categories in the IAB taxonomy. Given such a set, we intend to find the closest category $\hat{c}$ in the following manner

$$\hat{c} = \arg \max_{c \in C} \langle \phi(c), \phi(a_{in}) \rangle \quad (7)$$

where $\phi(c)$ and $\phi(a_{in})$ denote (SBERT) embeddings for the category phrase and the input ad text. We do not assume the presence of category labels for a given ad text in our setup, and the above

Data: Similar ad texts $a_1, \dots, a_m$ (sorted by decreasing relevance to input)

Result: words set $\mathcal{W}^*$ to augment $\mathcal{G}$,
 namely $\mathcal{G}^* = \mathcal{G} \cup \mathcal{W}^*$, $|\mathcal{W}^*| = \max_words$

```

 $\mathcal{W}^* = \emptyset$ ;
for  $i \in \{1, \dots, m\}$  do
  for  $w \in a_i$  do
    if  $(\langle \phi(a_{in}), \phi(w) \rangle) \geq \min\_word\_sim$  then
       $\mathcal{W}^* = \mathcal{W}^* \cup w$ ;
      if  $|\mathcal{W}^*| = \max\_words$  then
        return  $\mathcal{W}^*$ 
      end
    end
  end
end
end

```

Algorithm 2: Augmentation with similar ads’ text.

method is a proxy to infer the closest category for a given ad text. Having inferred $\hat{c}$ for given input ad text, we compute the category bias $bias_i^{(cat)}$ for each vertex i in the augmented graph $\mathcal{G}^*$ as:

$$bias_i^{(cat)} = \langle \phi(\hat{c}), \phi(w_i) \rangle, \quad (8)$$

where w_i denotes the word associated with vertex i . Intuitively, category biasing tries to up-weight the words closer to the inferred category. In the final vertex value update in (5), $bias_i^{(cat)}$ is used in conjunction with the self bias as defined in (4).

5 MULTILINGUAL AD TEXT

In the multilingual setting, we assume the following: (i) ad text can be non-English, (ii) the image query is in English, and (iii) we can retrieve semantically similar (existing) English ads. Furthermore, in this paper, we assume access to a generic translation tool which can translate non-English ad text to English. Given the above assumptions, for multilingual ad text, we simply translate the ad text from its original language to English, then use the English version with VisualTextRank (i.e., fetching similar English ads, and doing category biasing). The intuition behind such an approach is that regardless of the language, visual components in the ad image will be strongly tied to the product mentions in the ad text. For example, for a furniture ad in English or an equivalent ad in French, the ad image is likely to have furniture. By borrowing from text and image tags of semantically similar English ads for a given multilingual ad, VisualTextRank facilitates a simple yet effective form of cross lingual learning. It also does not rely on having a large set of existing ads in the ad’s original language (except in the case of English ads).

In our experiments, we also explored the use of multilingual (xlm) SBERT [19] to: (i) fetch similar English ads for a given multilingual ad (without translating it), (ii) run VisualTextRank with such xlm-SBERT embeddings to get a keyword (translated to English by using word-level-translation). As we explain in Section 6.4.3, this approach was inferior to the above ad text translation approach.

6 RESULTS

6.1 Offline evaluation dataset

We collected ad text to image query data from Verizon Media Native’s stock image search feature (as illustrated in Figure 1). This feature is currently available in the onboarding workflow for small and medium businesses with relatively low advertising budgets. For a given ad text, the advertiser might make multiple search queries before converging on the final choice for onboarding (launching the ad campaign). We consider all such queries as relevant (golden set) for the given ad text. For example, for the ad text ‘Bank Foreclosing? Stop Foreclosure Now!’, if the user searches for ‘stress’, ‘eviction’, and ‘bills’, all the three queries are labeled relevant queries. We sampled the collected data to create an evaluation data set with ~ 600 ads with ad text in English, and 400 non-English (multilingual) ads spanning German, French, Spanish, Portuguese, Swedish, and Italian ad text. For the set of English ads, the image queries were editorially reviewed and corrected for spelling mistakes, and we will refer to this set as the ENG-eval set. For the multilingual set, each non-English ad was translated to English using an internal translation tool. The image queries were also editorially reviewed, translated to English (if they were not in English), and corrected for spelling mistakes as well (if they were already in English). We will refer to this as the non-ENG-eval set. While processing raw data to create the non-ENG-eval set we observed a significant number of cases where the advertiser tried only non-English queries, or had spelling mistakes in the English query (plausibly due to lack of English proficiency). Hence, editorial review was needed to correct such samples for proper evaluation with our proposed methods (and baselines). Note that both ENG-eval and non-ENG-eval datasets are of a unique nature due to the integration of the stock image search feature in the ad platform (allowing us to map ad text to corresponding ad image search queries).

6.2 Ads pool for similar ads retrieval

Since VisualTextRank leverages semantically similar ads for augmenting the input ad text, we need a pool of ads from which we can retrieve similar ads given an input ad text. For our experiments, we collected a sample of 20,000 English ads (US-only) from Verizon Media Native (Yahoo Gemini), with their ad text and the raw ad image. We obtained image tags for the raw image (as explained in Section 3.3.2); this set had 20,000 ads which were not created using the stock image search feature, and hence did not have associated image queries. The time range for the data pull was same as that for the evaluation datasets described in Section 6.1. In the remainder of this paper, we will refer to this set of existing ads as ads-pool.

6.3 Evaluation metrics

Offline metrics. For offline evaluation on the ENG-eval and non-ENG-eval sets, we consider three metrics as defined below. For a given ad text, if the predicted keyword is w^* , and the golden set of

queries is the set $Q = \{q_1, q_2, \dots\}$,

$$\text{hard-accuracy}(w^*, Q) = \begin{cases} 1 & w^* \in Q \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

$$\text{soft-accuracy}(w^*, Q) = \begin{cases} 1 & \max_{q_i \in Q} \langle \phi(w^*), \phi(q_i) \rangle \geq 0.8 \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

$$\text{w2v-similarity}(w^*, Q) = \max_{q_i \in Q} \langle \phi(w^*), \phi(q_i) \rangle \quad (11)$$

where hard-accuracy simply checks whether the predicted keyword is in the golden set. Soft-accuracy uses the word2vec [15] similarity (as implemented in Spacy [7], details in Appendix) between a query in the golden set and the predicted keyword (represented by $\langle \phi(w^*), \phi(q_i) \rangle$) to if the predicted keyword is approximately (*i.e.*, above the 0.8 cosine similarity threshold) close to any of the golden set queries. This takes care of minor differences between the golden set queries and the predicted keyword (*e.g.*, run and running would be considered similar). Finally, w2v-similarity is the absolute similarity between the predicted keyword and the closest query (in word2vec sense) in the golden set. Note that our accuracy metrics are essentially *precision@1* since we focus on a single keyword as the ad image query.

Online metrics. An advertiser session is defined as a continuous stream of interactions of the advertiser with the onboarding UI, with gaps no more than 30 minutes between consecutive events. To gauge the effectiveness of automatic ad text-to-image search, we track two events in each advertiser session: (i) selecting (clicking) stock images to get a preview of the ad, and (ii) completing the onboarding process. Tracking the image selection event, we define *image selection rate* as:

$$\text{selection rate} = \frac{\# \text{ sessions with } \geq 1 \text{ image selections}}{\# \text{ sessions}}. \quad (12)$$

Tracking the onboarding event, we define *onboarding rate* as:

$$\text{onboarding rate} = \frac{\# \text{ sessions ended with onboarding}}{\# \text{ sessions}}. \quad (13)$$

The online metrics listed above are expected to capture the effect of automatic image recommendations on the stock image selection rate and the onboarding completion rate. Note that an image selection event may or may not lead to an onboarding completion event (if the advertiser does not launch the ad campaign).

6.4 Offline results

We first describe multiple baseline methods for ad text-to-image query (most of them are existing keyword extraction methods), and then go over the offline results for VisualTextRank and the baselines (for both ENG-eval and non-ENG-eval sets).

6.4.1 Baseline methods. We use TF-IDF, existing graph based keyword extraction methods [4] (TextRank, Position Rank, Topic Rank, Topical Page Rank, Multipartite Rank, YAKE), as well as the recently proposed (self) biased TextRank [9] as baselines. For all the baselines, as well as the VisualTextRank method, we consider only nouns, proper nouns and verbs as valid output keywords, *i.e.*, we have a POS tag filter (as guided by the POS tag insight in Section 3).

model	hard accuracy	soft accuracy	avg. w2v similarity
position rank	11.5%	17.07%	0.412
YAKE	14.11%	20.03%	0.4655
topical page rank	14.81%	21.25%	0.4515
multipartite rank	15.51%	21.78%	0.4639
topic rank	16.2%	23.17%	0.4719
tf-idf	18.12%	23.69%	0.486
TextRank	21.95%	28.92%	0.5297
self-biased TextRank	26.48%	35.37%	0.5739
VisualTextRank	28.05%	39.37%	0.6147

Table 2: Offline metrics for ENG-eval.

6.4.2 *ENG-eval results.* Table 2 shows the performance of the proposed VisualTextRank method versus the baselines outlined in Section 6.4.1 for the ENG-eval set. As shown in Table 2, VisualTextRank significantly outperforms the closest baseline (self) biased TextRank by 11.3% in terms of soft-accuracy.

6.4.3 *Non-ENG-eval results.* For brevity, Table 3 just shows the performance of the proposed VisualTextRank method versus the (self) biased TextRank method (best baseline in the ENG-eval case) as well as TextRank for the non-ENG-eval set; the results for other baselines as in Table 3, are in line with the observations for the ENG-eval set. As shown in Table 3, in terms of soft-accuracy, Vi-

model	hard accuracy	soft accuracy	avg. w2v similarity
TextRank	13.19%	17.18%	0.4403
self-biased TextRank	22.39%	29.75%	0.5302
VisualTextRank	23.62%	31.60%	0.5563

Table 3: Offline metrics for non-ENG-eval.

sualTextRank outperforms (self) biased TextRank by 6.2% in the multilingual setting as well. Note that VisualTextRank leverages similar English ads for multilingual (input) ad text in this setup (as explained in Section 5). Compared to the above result for VisualTextRank, the xlm-SBERT based approach described in Section 5 led to poorer metrics (13.5% soft accuracy, 0.404 avg. w2v similarity).

6.5 Ablation study

In this section, we study the contribution of different components of VisualTextRank via offline performance on the ENG-eval set and non-ENG-eval set. Table 4 shows the impact of different components on offline metrics for the ENG-eval set. Baseline result for the ablation study is a self-biased TextRank method (line 1 in Table 4), which achieves 35.37% in terms of soft accuracy. With addition of category biasing to self biasing (as explained in Section 4.3.3), we observe a soft-accuracy improvement to 37.8% (line 3 in Table 4). The second major improvement (39.02%, line 5 in Table 4) is from adding one image tag from similar English ads (as explained in

Section 4.3.2). Putting it with similar ads’ text (as explained in Section 4.3.2), we get the best result of 39.37% soft accuracy (line 7 in Table 4). We notice that: (i) increasing number of tags or words from the similar ad does not help, because it takes away the attention from the ad text words (which are more likely to be in the golden set); (ii) adding image tags without pre-filtering by Algorithm 1 (line 6) gets lower soft accuracy than with the pre-filtering (line 5).

model	hard accuracy	soft accuracy	avg. w2v similarity
self biased TextRank	26.48%	35.37%	0.5739
cat biased TextRank	27.70%	35.71%	0.5975
self+cat biased TextRank	28.92%	37.80%	0.6061
self+cat biased + 1-nbr text	28.57%	38.15%	0.6074
self+cat biased + 1-nbr img tag	28.75%	39.02%	0.611
self+cat biased + 1-nbr img tag non-filtered	28.92%	38.5 %	0.6111
self+cat biased + 1-nbr img tag + 1-nbr text	28.05%	39.37%	0.6147

Table 4: VisualTextRank ablation study for ENG-eval.

For the non-ENG eval set, the ablation study (Table 5) shows that similar ads in English provide useful information for non-English ad text-to-image query task (and hence facilitate cross-lingual learning). In particular, the results follow the same pattern as for the ENG-eval set, except that the best result (line 5 in Table 5) is achieved by augmentation with only similar image tags, and no similar ad text words. Table 6 shows two (anonymized) examples of the data samples in the non-ENG-eval set with image search queries, and algorithm’s output for (self) biased TextRank and VisualTextRank. The first sample in Table 6 is an English ad text, and the second sample is a Spanish ad text. In both cases, VisualTextRank returns a keyword matching the (golden) query, unlike the (self) biased TextRank. For these two ad examples, we also show their (anonymized) nearest neighbours and tags of the corresponding neighbor images (Table 7). Note that in both cases neighbour image tags coincide with the ground truth user query and shift algorithm predictions towards the correct output (furniture and house, correspondingly).

6.6 Online results

We conducted an A/B test where advertisers were split into test and control buckets (50-50 split). Test advertisers saw recommended images (via automatic text-to-image search query) when they accessed the image search feature, while control advertisers saw a random set of images. Both test and control advertisers were permitted to query the stock image library further after seeing the initial set of (recommended or random) images. Due to latency constraints, a lighter version of VisualTextRank (with category biasing, averaged

model	hard accuracy	soft accuracy	avg. w2v similarity
self biased TextRank	22.39%	29.75%	0.5235
cat biased TextRank	21.78%	28.83%	0.5409
self+cat biased TextRank	24.23%	31.29%	0.5575
self+cat biased + 1-nbr text	23.93%	31.60%	0.554
self+cat biased + 1-nbr img tag	24.85%	32.21%	0.5676
self+cat biased + 1-nbr img tag non-filtered	24.54%	31.9 %	0.5608
self+cat biased + 1-nbr img tag + 1-nbr text	23.62%	31.60%	0.5563

Table 5: VisualTextRank ablation study for non-ENG-eval.

ad text	user queries	biased TextRank	Visual-TextRank
Online Store. Asian antique and vintage furniture	furniture; wooden furniture	antique	furniture
Necesitas Vender Tu Casa Rápido?	home selling; house	sell	house

Table 6: Biased TextRank vs VisualTextRank examples.

ad text	nearest nbr ad	nbr image tag
Online Store. Asian antique and vintage furniture	Furniture and Decor Sale. Up to 70% Off Top Brands And Styles!	furniture
Necesitas Vender Tu Casa Rápido?	Homeowners Could Sell Their Homes Fast. ____ realtors get the job done.	house; home

Table 7: Examples of ads nearest neighbours and image tags.

word2vec embedding for text instead of SBERT, and no augmentation via similar ads) was used. The end-to-end latency of this version was below a second. In the A/B test, we saw a 28.7% lift in the selection rate in the test bucket (compared to control). In terms of onboarding rate, we saw a 41.64% lift. The data was collected over a one month period, and was limited to US-only traffic. Note that, the final image selected from the stock image library need not be from the initial set of images shown to the advertiser. However, the selection rate lift indicates higher adoption of stock images (driven by initial recommendations which increase the advertiser’s interest to query further). The lift in onboarding rate is indicative of the positive impact on the ad platform.

7 DISCUSSION

The motivation behind VisualTextRank stemmed from insufficient data for supervised training for the ad text-to-image query task. As more advertisers explore stock images (influenced by VisualTextRank), we may collect enough data to fine-tune pre-trained multi-modal models [12] on our task. Cross-lingual learning with such multi-modal models for our task is another future direction.

REFERENCES

- [1] [n.d.]. IAB categories. <https://developers.mopub.com/publishers/ui/iab-category-blocking/>.
- [2] [n.d.]. Unsplash: Photos for everyone.
- [3] Narayan Bhamidipati, Ravi Kant, and Shaunak Mishra. 2017. A Large Scale Prediction Engine for App Install Clicks and Conversions. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM '17)*.
- [4] Florian Boudin. 2016. pke: an open source python-based keyphrase extraction toolkit. In *COLING 2016*.
- [5] 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).
- [6] Armin Kappeler Kofi Boakye Herdade, Simao and Joao Soares. 2019. Image captioning: Transforming objects into words. In *arXiv preprint arXiv:1906.05963*.
- [7] Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength Natural Language Processing in Python.
- [8] Zaeem 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*.
- [9] Ashkan Kazemi, Verónica Pérez-Rosas, and Rada Mihalcea. [n.d.]. Biased TextRank: Unsupervised Graph-Based Content Extraction (*COLING 2020*).
- [10] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. 2017. OpenImages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>* (2017).
- [11] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2020. What Does BERT with Vision Look At?. In *ACL (short)*.
- [12] Xi Yin Chunyuan Li Pengchuan Zhang Xiaowei Hu Lei Zhang Lijuan Wang et al Li, XiuJun. 2020. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV 2020*.
- [13] Michael Maire Serge Belongie James Hays Pietro Perona Deva Ramanan Piotr Dollár Lin, Tsung-Yi and C. Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *ECCV 2014*.
- [14] Rada Mihalcea and Paul Tarau. [n.d.]. TextRank: Bringing Order into Text. In *EMNLP 2004*.
- [15] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality (*NIPS '13*).
- [16] Shaunak Mishra, Manisha Verma, and Jelena Gligorijevic. 2019. Guiding Creative Design in Online Advertising. In *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys '19)*.
- [17] Shaunak Mishra, Manisha Verma, Yichao Zhou, Kapil Thadani, and Wei Wang. 2020. Learning to Create Better Ads: Generation and Ranking Approaches for Ad Creative Refinement. In *Proceedings of the 29th ACM International Conference on Information Knowledge Management (CIKM '20)*.
- [18] Liang-Chieh Chen Qiao, Siyuan and Alan Yuille. 2020. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In *arXiv preprint arXiv:2006.02334*.
- [19] Nils Reimers and Iryna Gurevych. [n.d.]. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks (*EMNLP 2020*).
- [20] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS 2015*.
- [21] Nan Ding Sebastian Goodman Sharma, Piyush and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*.
- [22] Keren Ye and Adriana Kovashka. 2018. ADVISE: Symbolism and External Knowledge for Decoding Advertisements. In *ECCV 2018*. 868–886.
- [23] Yichao Zhou, Shaunak Mishra, Jelena Gligorijevic, Tarun Bhatia, and Narayan Bhamidipati. 2019. Understanding Consumer Journey Using Attention Based Recurrent Neural Networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery Data Mining (KDD '19)*.
- [24] 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)*.

NOTES ON REPRODUCIBILITY

In this section, we list helpful information for reproducing the results in our paper.

POS tagging and offline metrics. For POS tagging based filters (*i.e.*, to focus on only nouns, proper nouns and verbs), we used Spacy’s POS tagger [7]. For, the offline evaluation method based on word2vec similarity between the ground truth query and model output we used Spacy’s word similarity function implemented using 300 dimensional word2vec embeddings (trained on Google News dataset) for English text [7].

TextRank and baselines for offline evaluation. We used the PKE [4] library for running experiments with baseline keyword extraction methods: Position Rank, Topic Rank, Topical Page Rank, Multipartite Rank, and YAKE (as reported in Section 6.4.2). For TextRank and biased TextRank we used the implementation in https://github.com/ashkankzme/biased_textrank, with damping factor $d = 0.8$, iterations = 80 (for the node importance estimate to converge), and node similarity threshold = 0.9 (using SBERT embeddings). The SBERT [19] embeddings were based on stsb-distilbert-base as listed in the pre-trained models in https://www.sbert.net/docs/pretrained_models.html. We explored the larger BERT models like BERT-base, BERT-large, RoBERTa-large but the benefits were not significant.

Category list for computing category bias. We used a list of ~ 250 categories as the set C defined in Section 4.3.3. Each entry in C is a phrase corresponding to a category in the flattened IAB ads category taxonomy [1]. For example, the ‘Arts & Entertainment’ high level IAB category has ‘Books & Literature’ and ‘Music’ as subcategories [1]. After flattening, ‘Arts & Entertainment’, ‘Books & Literature’ and ‘Music’ are listed as three separate entries in category set C used for computing category bias.

VisualTextRank. For the best results from VisualTextRank we used the following configuration (for both ENG-eval and non-ENG-eval):

- damping factor $d = 0.8$
- node similarity threshold = 0.9 (using distilbert based SBERT embeddings)
- min. tag similarity for augmentation $min_tag_sim = 0.7$
- min. word similarity for augmentation $min_word_sim = 0$
- for category biasing, we set all $bias_i^{(cat)}$ below 0.4 to zero.

Apart from using distilbert based SBERT, we also explored the larger BERT models like BERT-base, BERT-large, RoBERTa-large but the benefits were not significant. For the multilingual ad experiment with xlm-SBERT (*i.e.*, without ad text translation) described in Section 5, we used the stsb-xlm-r-multilingual model which was trained on parallel data for 50+ languages as mentioned in https://www.sbert.net/docs/pretrained_models.html.