

Visual Congruent Ads for Image Search

Yannis Kalantidis*, Ayman Farahat†, Lyndon Kennedy‡, Ricardo Baeza-Yates§ and David A. Shamma¶

*Yahoo

†Amazon

‡Futurewei Technologies Inc.

§DTIC, Universitat Pompeu Fabra & DCC, Universidad de Chile

¶Centrum Wiskunde & Informatica (CWI)

Abstract—The quality of user experience online is affected by the relevance and placement of advertisements. We propose a new system for selecting and displaying visual advertisements in image search result sets. Our method compares the visual similarity of candidate ads to the image search results and selects the most visually similar ad to be displayed. The method further selects an appropriate location in the displayed image grid to minimize the perceptual visual differences between the ad and its neighbors. We conduct an experiment with about 900 users and find that our proposed method provides significant improvement in the users’ overall satisfaction with the image search experience, without diminishing the users’ ability to see the ad or recall the advertised brand.

I. INTRODUCTION

The goal of native advertising is to present sponsored content in a way that matches the surrounding context in which the advertisements appear. Displaying ads online generally allows advertisers to target users with lower cost than in offline settings [13]. Furthermore, the additional information about site content and user interests available to online advertisers, allows them to choose ads which are particularly semantically relevant and, thus, appealing or memorable for users. For these reasons, native advertisements have rightfully taken over a large chunk of the market [15].

While newer to digital advertising, the underlying concept of native advertisements has been well established in offline domains, such as print magazines. *Vanity Fair*, a magazine covering topics in fashion, popular culture, and entertainment, has used an in-house design creative agency (Vanity Fair Agenda) since 2003 to ensure that ads match the look and feel of the surrounding content [1]. An ideal magazine ad creates a memorable experience for a brand while becoming a part of the experience of reading the magazine, rather than distracting from it.

It is estimated that 25% of search queries have image-related intent. Effective placement of native advertisements within image search results has been explored far less than in the context of traditional search. Thus, techniques that optimize the placement of native image ads present a significant opportunity. Optimizing image results in the context of photo-oriented sites or image search engines presents an interesting tradeoff. Image queries are very difficult to monetize, meaning that presenting ads that are relevant and clearly marked can lead to significant increases in revenue. On the other hand, advertisements that add clutter or distract from organic results can degrade the overall search experience. Advertisers must

find a way to make their content stand out while still matching the surrounding content.

Prior research has shown that finding the appropriate balance can be especially difficult in contexts in which users are browsing rich multimedia content. Increasing the obtrusiveness of multimedia-based advertisements has been shown to increase purchase intent [6]. However, making ads more obtrusive runs the risk of alienating users and creating a negative impression of both the site publisher and the brand being advertised [13]. The combination of relevance to website content and advertisement obtrusiveness has also been shown to raise additional privacy concerns [14], further complicating matters. The quality of user experience is highly affected by the quality, relevance and placement of advertisements in a webpage [2], [9].

We focus on the selection and placement of high-quality image-based advertisements *with an emphasis on not distracting or annoying users*. Ideally, we want to place an ad that is properly marked and that blends in with the rest of the image set, adding to the overall user experience rather than detracting from it, as in Figure 1. Although we want the advertisement to blend in with the surrounding context, it is also important that it be clearly marked. Deceiving or tricking the user with a camouflaged advertisement can degrade user trust, harming the site in the long run.




Fig. 1: An example of visually congruent ad placement in image search. The ad matches both visually and semantically with the surrounding results.

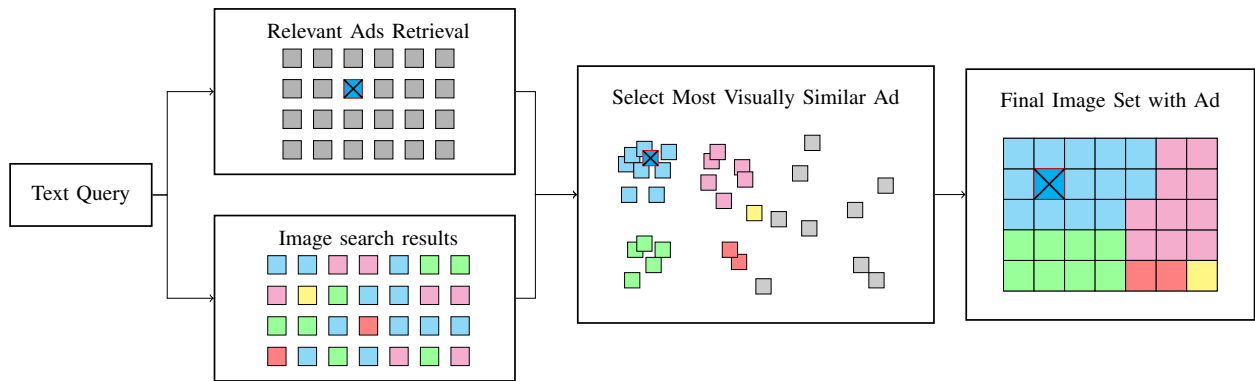


Fig. 2: Overview of the proposed approach. Advertisement images are shown in gray squares, while the one that is subsequently selected for display is shown as a blue square with a red border and a cross. Squares filled with any other color correspond to images from the same cluster in the visual space.

Therefore, we propose a specific technique that strikes a balance between making users aware of a native image ad without distracting from the overall image search experience. We accomplish this by going one step further than semantic similarity between advertisement and web content; we propose using *visual similarity* between image-based ads and surrounding images to increase their quality and effectiveness. We show that this approach makes native image ads seem more relevant and less distracting to users, without hurting clicks or recognizing the brand of the ad.

An overview of our approach is presented in Figure 2 and summarized here. Starting with a set of image search results and a set of relevant advertisements, we first select the advertisement most visually similar to the result set. We then position it within the results, either preserving the initial ordering of results or reorganizing the whole set based on visual similarity. We refer to this approach as *Visually Congruent Ad Placement* or *VCAP* for short.

We can summarize our main contributions as follows:

- 1) We propose the use of visual similarity between ad images and image search results and present a simple, automated way of selecting visually congruent native ads. We suggest different ways of positioning the ad in the result set grid in a way that the ad image is displayed among its most visually similar ones.
- 2) We use compressed image features to make our approach applicable to web-scale image search. Our algorithm runs in a few hundred milliseconds and requires storing only a few bytes per image.
- 3) We present findings from a large crowd sourced user study that shows that our algorithm has a sizable and statistically significant impact on reducing user distraction and increasing their overall ad experience while not hurting ad recognition. Given the absence of ad revenue from image search, our findings have significant financial implications.

Our experimental validation on about 900 human samples through Amazon Mechanical Turk¹ shows that our approach is able to retain brand recognition and awareness as effectively as a “distracting” ad would, while at the same time providing for a better user experience.

To the best of our knowledge, this is the first paper that investigates and evaluates visually congruent native advertisements in image search. The paper is structured as follows. § II first discusses related work and then § III presents the core of our approach, *i.e.*, selection and positioning of the ad in a web-scale image search engine. § IV presents our experimental validation while § V discusses different aspects, generalizations and implications of our approach. The paper ends in § VI with the conclusions and future work.

II. RELATED WORK

Image content representation has evolved greatly during the last years. The first wave started with the success of local SIFT features [24] and the bag-of-words model for search [29], borrowed from text retrieval [4]. As the field evolved, focus shifted on global image representations for efficiency, through aggregating local features. Most notable examples of aggregation are VLAD [20] and Fisher Vectors [27], two approaches that have influenced multiple extensions [31], [10], [16].

We are currently in a new era for image content representation. With the recent advances in GPUs and the growing number of training data available, the computer vision community has revisited Deep Convolutional Neural Networks (CNNs), *i.e.*, neural nets with many hidden layers and millions of parameters. When trained on big image databases like ImageNet [28], CNNs have been shown to “effortlessly” improve the previous state-of-the-art in many computer vision applications, *e.g.*, image classification [23] and visual search [3].

The aggregated approaches as well as the CNN-based ones, all produce a *global* image signature, *i.e.*, a high-dimensional feature vector in Euclidean space. In this case, similarity search reduces to nearest neighbor search in a feature space allowing for very scalable search. When the image database consists of billions of images, every byte counts. We therefore need to embed our visual features in a very compact image signature

¹<https://www.mturk.com/>

that reduces the storage but that also keeps nearest neighbor search quality close to the original representation space.

Large scale nearest neighbor search was traditionally based on hashing signatures [7], [26], mainly due to low memory footprints and fast search in Hamming space [25]. But even recent advances in the field [21], cannot achieve performance comparable to the slower but still memory efficient quantization-based approaches [19], [25]. What is more important, approaches that are based on Product Quantization, unlike hashing, allow us to reconstruct the original vectors from the compact signatures. The past approaches [19], [11], [22] are all good candidates for compressing CNN visual features. We choose to use the Locally Optimized Product Quantization approach [22] that currently gives the state-of-the-art results in nearest neighbor search.

III. VCAP: VISUALLY CONGRUENT AD PLACEMENT

Let $\mathcal{I} = \{I_1, \dots, I_n\}$ be the set of n images returned as search results for a text query q issued by user u . Also, let $\mathcal{A} = \{A_1, \dots, A_m\}$ be a set of m image-based *advertisements* that are *relevant* to the query and/or the user. An advertisement can be relevant to the query either in terms of semantic/topical similarity with q (e.g., a car advertisement for the query “convertible cars”) or in terms of the interests of the user u , i.e., based on any profile or browsing history data that might be available (e.g., a car advertisement for a user that frequents car-related sites).

Selecting a relevant or semantically similar ad is a subject that has been studied extensively [8]. Here, we assume that the set \mathcal{A} of such relevant ads is already available to us for any query q . What we aim for is to automatically select the advertisement from set \mathcal{A} that best matches the image set \mathcal{I} in terms of visual appearance and also present/place it in such a way that it fits seamlessly with its surrounding images.

Our approach is based on *visual similarity*, i.e., similarity based on image visual content, by choosing features extracted by Deep Convolutional Neural Networks (CNNs). Given that we aim for a scalable solution, i.e., when the image database would be the entire Web image corpus, we choose the AlexNet deep CNN architecture [23] as it offers a good trade-off between performance and speed.

We first get the $4k$ -dimensional *fc7* features, i.e., the features after the second fully connected layer. In pursue for a more compact representation, we use dimensionality reduction to get the features to only $d = 128$ dimensions. It has been shown that such embedding makes the image representation orders of magnitude more compact while minimally affecting search performance [3].

To reduce dimensionality we use PCA, learned using *fc7* features from a public 100 Million YFCC100M image set [30]. As we plan on further compressing the visual features to just a few, using Locally Optimized Product Quantization [22], we also permute the feature dimensions after PCA so that variance is balanced among all sub-vectors (see § III-C for more details).

A. Selecting the Most Visually Similar Ad

Let $\mathcal{I}_f = \{i_1, \dots, i_n\}$ and $\mathcal{A}_f = \{a_1, \dots, a_m\}$ be the sets of visual features for \mathcal{I} and \mathcal{A} , respectively, where $i_k, a_j \in \mathbb{R}^d$

and $k \in \{1, n\}, j \in \{1, m\}$. Getting the most visually similar ad requires finding the optimal ad image $\tilde{a} \in \mathcal{A}$ so that

$$\tilde{a} = \arg \min_{a \in \mathcal{A}_f} \text{dist}(a, \mathcal{I}_f), \quad (1)$$

where $\text{dist}(a, \mathcal{I}_f)$ is a function of dissimilarity between feature a and *feature set* \mathcal{I}_f of all images in the result set. There are multiple ways of defining such dissimilarity between a vector and a set. One would be to set \tilde{a} to the ad that is the closest to one of the images, i.e.,

$$\tilde{a} = \arg \min_{a \in \mathcal{A}_f} \min_{i \in \mathcal{I}_f} \|a - i\|, \quad (2)$$

where dissimilarity between two feature vectors is measured using the Euclidean distance. This way, the globally minimum dissimilarity between the ad feature set \mathcal{A}_f and image feature set \mathcal{I}_f is kept.

This is basically a nearest neighbor search in visual feature space, where we select the ad that is closest to *any* of the images. Therefore, this formulation does not take into account the set as a whole and implies the danger of selecting an ad that matches an image that is not visually consistent with the rest of the set.

Another, more robust, way of measuring this dissimilarity, would be by taking the whole set *jointly* into account and selecting the ad that is closest to the set \mathcal{I}_f as a whole, or

$$\tilde{a} = \arg \min_{a \in \mathcal{A}_f} \sum_{i \in \mathcal{I}_f} \|a - i\|. \quad (3)$$

We found that this is a better way of selecting visually congruent ads that blend better inside the result set. Therefore, unless otherwise stated we will use Equation (3) for selecting the best ad in the rest of this paper.

During ad selection, we sort all ads by proximity to the given image set and select the closest ad \tilde{a} . At no extra cost, we can keep the information on which is the most similar ad for each one of the images in set \mathcal{I} .

As we present in Section III-B, during ad placement, we also sort all images of set \mathcal{I} by proximity to the selected ad. These are dual problems, and we further require *reciprocity* in nearest neighbor search for accepting the selected ad. In other words, we want the nearest neighbor of the selected ad in the image set \mathcal{I} to also have that specific ad as its nearest neighbor in set \mathcal{A} as well.

B. Positioning the Advertisement

We investigated different advertisement positioning strategies, all of them assuming that we present results to the user in the typical way for image search, i.e., on an 2D grid.

Most of the presented strategies affect the ordering of images in different ways. For the case where the original ordering is considered very important, we also present a positioning strategy that does not affect that ordering. Figure 3 shows some of those strategies that are discussed below. Each white square is an image from set \mathcal{I} , while the selected ad image is colored in dark cyan, and its nearest images from set \mathcal{I} in lighter cyan.

Let the image indices in set \mathcal{I} also reflect the *original ordering* of the images according to relevance to the textual query q . Below we present our different ad positioning strategies ordered in terms of how much they affect the initial ordering of set \mathcal{I} .

Let \tilde{a} be the ad image selected as described in § III-A. Assuming that we want to maximize visual congruence by showing visually similar images nearby on the grid, it is important to also calculate the ordering of all images of set \mathcal{I} by visual proximity to the selected ad \tilde{a} . We therefore calculate set $\tilde{\mathcal{I}} = \{\tilde{i}_k\}$, where $\forall k \in \{1, n-1\} : \|\tilde{i}_k - \tilde{a}\| < \|\tilde{i}_{k+1} - \tilde{a}\|$.

1) *Preserving the Original Image Ordering*: If preserving the original image ordering is important, the optimal way of placing the ad is right next to its most visually similar image \tilde{i}_1 . One can choose the side of placement (left or right) based on the dissimilarity of the neighboring images of \tilde{i}_1 in the original ordering, and place it on the side of the neighbor with the largest similarity to \tilde{a} .

2) *Altering the Ordering Locally Around the Ad*: Allowing the original ordering to change, we get more freedom in placing the ad. One can place the ad image between its two nearest neighbors from set $\tilde{\mathcal{I}}$, *i.e.*, between \tilde{i}_1 and \tilde{i}_2 . This way we only change the position of \tilde{i}_2 and keep the remaining images intact in terms of order. Figure 3(a) presents a synthetic example of this strategy.

Until now, we assume that images are ordered in a 1-dimensional stream. More visually congruent positioning strategies can be achieved if we also consider the 2D nature of the displayed image grid. We assume that the image grid is static and not responsive, *i.e.*, it does not change when the browser’s window resizes. In the responsive case, one solution would be to re-compute orderings on the fly.

In this case the neighborhood of each cell can be defined in 4- or 8-way connectivity. Generalizing our previous placement, we can surround the ad image by its most similar images, either in 4 or all 8 directions. Figure 3(b) presents a synthetic example of this approach.

3) *Altering the Ordering Globally: Stochastic Neighborhood Embedding or t-SNE* [33], is an embedding that aims to reduce the dimensionality of high-dimensional data to very few dimensions, *i.e.*, 2 or 3, while at the same time preserving distances in the original space. It was shown to be very successful for visualizing image datasets in a way that visual similarity is preserved.

Such an embedding seems ideal for our case as well: Projecting all images and the ad image in that 2D space can provide us with data for placing not just the ad image among visually similar ones, but also every other image in the set.

As shown in the synthetic example in Figure 3(c), though, using this embedding does not constrain results on a grid and the output coordinates are arbitrary.

A simple way of placing images based on their 2D coordinates is to *greedily* select the closest grid position in 2D space for each image. That means that each image will take the closest grid position if it is not already occupied. Priority is given to the ad image, to ensure its positioning among

similar ones and images can be considered in ascending visual similarity from the ad.

The greedy approach may result in visually similar images being placed far apart on the grid. To avoid that, simple density based clustering can be efficiently run on the 2D image coordinates: We used a fast Mean Shift clustering [5], which is an algorithm that automatically selects the final number of modes. Given image clusters, we can now greedily place images that correspond to the same cluster in adjacent positions of the grid. A synthetic example is shown in Figure 3(d), where images of the same cluster are shown with the same color.

The advertisement image is placed among images of the same cluster, thus ensuring visual coherency. If we see that the ad image is placed in a cluster just by itself, we may discard the ad, since this is an indication that it is not consistent enough visually with at least some of the images in the set. This way we get a sense of how visually close the ad image is, relevant to the visual consistency of the image set \mathcal{I} itself, without the need for a parameter or fixed threshold.

Although slightly computationally expensive, the t-SNE embedding and clustering-based algorithm can run in just milliseconds if we use fast versions of the algorithms [32], [34] and limit their input to the first page of image search results, *i.e.*, a few dozen images. This is not restricting in practice, as in any case we would want to place our ad image within the first page of results. For our study, we chose the greedy placement to make the overall look more consistent with the other conditions (see § IV). We do not evaluate the different placements in the current experiments, leaving this for future work.

C. Scaling by Compressed Image Signatures

We want our approach to be applicable in real-time for web-scale image search. In the typical scenario, images are indexed by their surrounding textual metadata and, given a query, are ordered with relevance to those metadata through text matching [4]. We therefore only need to deal with the top- k images returned by the index.

However, visual feature extraction is a computationally heavy process that cannot run on the fly even for a few dozen photos. That means that we need to have visual features stored and available for the whole image dataset, which in a web-scale image search scenario would contain billions of images crawled from the Web. The CNN features we use are powerful, but they require storing a vector of 128 floating point numbers for each image. Although for most content-based applications this is considered a pretty compact signature, when scaling to 10^{10} photos, every byte counts.

Recent quantization approaches for large scale nearest neighbor search can provide very compact signatures with only a small drop in search performance [19], [11], [22]. Such methods perform much better than hashing, and, unlike hashing, they allow the *reconstruction* of the original vector from the compressed domain, giving a much finer way of measuring (dis-)similarity. We therefore choose to use a recently proposed extension of Product Quantization [19] for compressing visual features named *Locally Optimized Product Quantization* or *LOPQ* [22]. Keeping in mind all the different possible




Fig. 3: Different positioning strategies. The selected ad image is represented as a dark cyan colored square with a red border and a cross; all other squares correspond to images from set \mathcal{I} . White squares are images placed in the original ordering. (a): Positioning with minimal change of the original ordering; the ad is placed right next to its nearest images from set \mathcal{I} and at most one image from \mathcal{I} would change its place. (b): Altering the image order locally around the ad; the images in lighter cyan would be repositioned around the ad in ascending order of k as in set \tilde{i}_k . (c): Density based clustering in 2d space after t-SNE projection. Different colors correspond to different clusters. (d): Placing image clusters on the grid.

use cases for this content-based image signature outside this specific application of ad placement², we specifically choose to use the Multi-LOPQ approach of [22] that also enables efficient indexing and search.

Each vector is split into M sub-vectors and is quantized independently using one Byte per subvector. The codebook \mathcal{C} is considered to be composed of a Cartesian product of sub-codebooks, one for each of the sub-spaces. In the LOPQ extension, a two-level quantization is used, where codes and sub-spaces are optimally computed locally in the second quantization level in a way that variance is balanced among sub-spaces (we refer the reader to [22] for more details).

In the end, we represent each feature vector $x \in \mathbb{R}^d$ with a code $c = \{c_1, \dots, c_M\}$ of size M , where c_i is a byte-size representation for each sub-vector, typically the index of the closest sub-codebook centroid for the i -th sub-vector, together with the two coarse quantization indices from the first, coarse quantization step.

Advertisement images in the web-scale scenario are in the order of thousands. As we further restrict ourselves to *relevant* advertisements, the average number of ads to consider is even lower. In such scenario, compressing their signature is not really worth the further approximation induced by it. Of course, for the case where more than a couple thousand ads are often being considered for placement, one can also quantize the ad signatures and use the Multi-LOPQ approach for search [22]. We therefore choose not to quantize advertisement images and use *asymmetric distance computations* for approximate nearest neighbor search, *i.e.*, the *ADC* version of [19].

Given a code and all quantizers, one can project every compressed feature back to the original d -dimensional space. After projecting, advertisement selection and placement proceed as in § III. In the next section we evaluate the numerical errors due to quantization, by measuring how many times the same advertisement image is picked as closest to a given image set, with and without compression. We proceeded using the compressed signatures, which proved to be effective in selecting ad images that are perceptually consistent with the set.

²The same signature can have multiple uses, *e.g.*, deduplication, diversification, visual similarity search, *etc.*

	Topic				
	<i>Animals</i>	<i>Cars</i>	<i>Fashion</i>	<i>Movies</i>	<i>TV</i>
Ads	23	48	45	16	18
Acc.	0.74	0.69	0.67	0.81	0.81

TABLE I: Top row: Breakdown of our ads dataset showing number of ad images per topic. The total number of ads in the dataset is 150. Bottom row: Average nearest neighbor accuracy per topic using the compressed versions of the features versus the original (uncompressed) ones. The average for the whole dataset is 0.75.

IV. EXPERIMENTS

Now we present the experimental evaluation of the proposed ad placement approach. Before explicitly measuring the effectiveness of our optimal ad placement in a live field experiment, we choose to run a large-scale study on Amazon Mechanical Turk (AMT).

We devised a hard scenario to validate our approach: Given an image search result set and a set of ads that are *all relevant* to the image set *and* to the users interests, we tested three different ad placements or *conditions*: our approach, a random one, and one where we tried to make the ad to stand out of the result set as much as possible.

As mentioned before, in order to simulate targeted advertising, in all three conditions the ad was relevant to the images and the user, *e.g.*, we place a car-related ad in an image search result set for query “best car 2015” and we require people to be interested in cars in order to answer our survey.

A. Evaluation Protocol and Validation

To test our hypothesis, we curated an advertisement dataset. We selected image-based ads of high aesthetic quality from five generic topics of interest: *animals*, *cars*, *fashion*, *movies* and *TV Series*. We gathered 150 ads in total; the number of ads per topic is presented in the top row of Table I.

To simulate image search, we query the public Flickr API³ for photos with different search terms related to the 5 topics of interest. We specifically used the `photos.search` method ordered by relevance.

³<https://www.flickr.com/services/api/>

	Most similar	Random	Least similar
Workers per case	302	288	306
Saw an ad	39.7%	38.5%	38.6%
Saw the correct ad	86.7%	63.1%	79.7%
Mean time spent	54.4s	54.1s	52.2s

TABLE II: Statistics from the AMT study. Time spent carried had a standard deviation of ± 24.8 , ± 24.2 , and ± 25.6 , respectively across the Most, Random, and Least Similar categories.

To represent the visual content of ads and image search results, we follow the process presented in § III-A. We extract compact CNN features, reduce their dimensionality down to 128 dimensions and compress them using a Multi-LOPQ [22] model with $M = 16$ and coarse vocabularies of size 2^{13} . This means that we need to store in total just $128 + 26 = 154$ bits per image, *i.e.*, approximately 20 Bytes. We also L2 normalize the reduced features before compression and nearest neighbor search.

As the most interesting survey questions we asked had Likert scale answers, *i.e.*, correspond to ordinal data, we cannot rely on classic continuous-space statistical estimators like mean and standard deviation [17]. We therefore report the Mann-Whitney U test results for significance estimation, as well as percentages per condition.

For the time statistics in Table II we used robust estimates for mean and standard deviation [18] to exclude outliers.

To quantify the approximations induced by compression in nearest neighbor search, we also present some quantitative results on our ad dataset in the last row of Table II. As argued in § III-C, we compress the results image set features only and not the ad features.

For evaluation, we run 100 generic queries in the Flickr API and report the average accuracy for ad selection for each of the 5 topics, *i.e.*, how many times the ad selected as most similar was the same when using the uncompressed features and the compressed features. We see that three out of four times the same ad was picked.

These results are purely quantitative and just confirm the results reported in [22], *i.e.*, computing visual similarity in the compressed domain approximates similarity in the original domain pretty well.

However, as no ground truth can exist to tell us which ad *should* be picked, approximation results cannot give us any guaranties regarding the selected ad’s distractingness factor or relevance to the set. To measure such aspects, we run a large crowd sourced study in AMT. In this study we used the compressed image features, as this is the only viable option for a real-time and web-scale scenario. We used psiTurk⁴ to run the study.

B. Large Scale Study in AMT

We gathered result sets for 8 queries that are semantically relevant to the 5 categories of our ads dataset. For each of the 8 sets we created three conditions and each worker was presented

with *only one condition from one set*, selected randomly. We did this to avoid that the workers knew that we were testing ad placements and be biased.

The three conditions we experimented on were:

Most similar: Optimal selection and positioning of the most similar ad using our method.

Random: Random selection and positioning.

Least similar: Selection of the least similar ad and positioning it among the most dissimilar images.

To simulate targeted advertising and assume that the ad matches the user’s interests, we further required the workers to be interested in the ads topic, *e.g.*, we titled the survey that presented fashion-related ads as “A survey for people who are into fashion”.

We used the greedy positioning after 2D projection for the *Most Similar* and *Least Similar* conditions, where in the latter case we surround the ad with its most distant images from the set w.r.t. visual similarity.

We set the following 3 requirements for AMT workers: To have more than 95% of their HITs accepted, to have completed more than 1,000 HITs and, to be located in the US. Based on them we accepted the HITs from 896 workers. This implies that we showed each condition of each query (*i.e.*, in total we had 24 different versions) to about 37 workers. Of those who chose to declare their gender, 43% were female and 57% male. In terms of age distribution, 22% of the workers were under 26 years old, 49% were between 26-35 years old, 16% were between 36-45 years old, and 12% were over 45 years old.

1) *Study Protocol:* The study consisted of two pages: the *image set page* and the *questionnaire page*. On the image set page, the worker would see the result set together with two questions that would “force” them to actually pay attention to the images in the set. The first question was different for each result set, and would contain something in reference to the content of the images shown, usually in the format “How many images depict X ”, where X would be *e.g.*, “brown dogs”, “red trucks”, or “dresses”.

For example, Figure 4 shows the three conditions for the query “cats” from the topic “animals”. For all three conditions shown, the first question was “How many images depict gray cats?”.

The second and final question of the first page was shared among all queries and was the question “What is an appropriate title for this image set?”. After answering both questions the workers would proceed to the questionnaire page, *without the option of seeing the image set again*. Up to the questionnaire, nothing would reveal to the user the actual purpose of the survey. The questionnaire consisted of 11 questions in total; the complete list can be seen in Table III.

Some of those had open-ended text fields for answers, others were multiple choice and others in the 5-choice Likert scale: *Strongly Disagree*, *Disagree*, *Neutral*, *Agree*, and *Strongly Agree*. Below, we analyze the most important of the Likert scale questions independently. Of critical importance was question (Q2) “I recall seeing a sponsored ad from” that

⁴<https://psiturk.org/>

Image set page	
I1	Image set specific question (Text)
I2	What is an appropriate title for this image set? (Text)
Questionnaire page	
Q1	Did anything seem out of place? If yes, what? (Text)
Q2	I recall seeing a sponsored ad from (Multiple Choice)
Q3	Can you recall what was being advertised? (Text)
Q4	The sponsored ad was relevant to my interests (Likert)
Q5	The sponsored ad was relevant to the image set (Likert)
Q6	The sponsored ad was clearly marked (Likert)
Q7	The sponsored ad was distracting and out of place (Likert)
Q8	The ad experience was of very high quality (Likert)
Q9	I am very familiar with the advertised brand (Likert)
Q10	Gender information (Multiple Choice)
Q11	Age information (Multiple Choice)

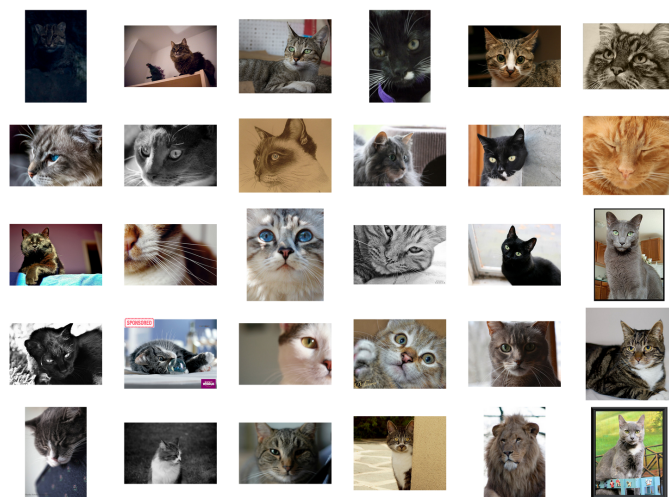
TABLE III: List of questions asked in our large scale Amazon Mechanical Turk study. We show the answer type in parenthesis after each question.

had five multiple-choice answers: The first three choices were query-dependent and were known brands from the specific topic, one of which also corresponded to the brand of the ad that the worker saw. The fourth choice was “other brand” and the fifth was “I didn’t see an ad”.

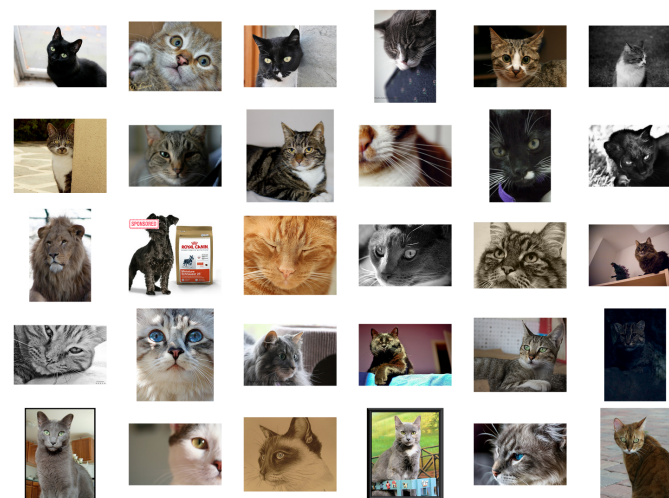
2) *Analysis of Results*: Some basic statistics are shown in Table II. In the first row, we show the number of workers per condition, which is fairly balanced for all three. The second row reports the percentage of workers that answered positively in (Q2), *i.e.*, chose one of the first four choices. It is notable that in all three conditions, about 39% of the workers claimed that they saw the ad. It is really essential for this number to be constant, as this shows that *showing an ad that is visually similar to the image set does not harm the recognition of the ad*. The third row presents the percentage of workers that answered (Q2) *correctly* per condition, *i.e.*, responded that they saw the brand that was indeed advertised in the set they saw. What is notable here is that the percentage of workers that could identify the brand correctly was *highest in the condition where the ad was selected by the proposed algorithm*. The fact that the percentage is also high in the *Least Similar* condition confirms related work that ad obtrusiveness correlates with ad recognition [6]. The final row reports the mean and standard deviation⁵ of the time spent on the image set page for the workers that claimed to have seen an ad. We see that the time is practically constant in all three cases, while the median for all conditions is 51.8 seconds.

In Figure 5 we show the percentages of each of the five possible responses per condition for the five most important questions/statements. All five are in Likert scale, so, to reject the null hypothesis [17] we also report p-values using the Mann-Whitney U test, when considering the three conditions in pairs. We always report p-values for the three pairs $p_1 =$ (Most Similar, Random), $p_2 =$ (Most Similar, Least Similar) and $p_3 =$ (Random, Least Similar) in this order. Below, we analyze the most interesting findings for some of the questions/statements independently.

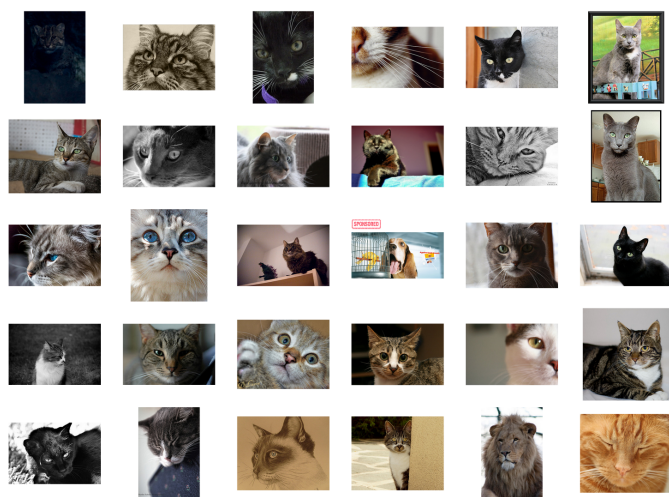
⁵To exclude outliers (*e.g.*, we saw that the slowest worker spent over 25 minutes at the image set page) we used robust estimates [18] for the mean and standard deviation reported.



(a) Most similar ad selection, optimal positioning



(b) Random ad selection and positioning



(c) Least similar ad selection and positioning

Fig. 4: The three different conditions presented to workers for query “cats”. The ad is marked with the word **sponsored**. *Best viewed magnified on a monitor.*




Fig. 5: Stacked percentages per condition for five of the questions asked in our survey. Each subplot presents percentages for the question shown in its caption. We see that our approach, *i.e.*, the condition denoted *Most Similar*, gives favorable results, while not affecting the advertisement’s recognition negatively (Q6).

(Q8) The ad experience was of very high quality. The percentages for this question presented in Figure 5 show significant signs that our basic hypothesis was right. Using visual similarity to pick the ad greatly affects the ad experience of the workers.

In fact, 44.5% of the workers that were shown the *Most Similar* condition chose Agree or Strongly Agree for the statement, while the same percentages for the *Random* and *Least Similar* conditions are 31.5% and 28.8%, respectively. Percentages reverse for Disagree and Strongly Disagree; only 18.4% of the workers shown the condition that was selected by our algorithm found their ad experience of low quality, while the percentages are 33.3% and 37.2% for the other two cases, respectively.

The Mann-Whitney test confirms the statistical significance of the results. It gives us strong evidence to reject the null hypothesis as the p -values are 0.011, 0.006, and 0.747 for the

three pairs, respectively, where the first two pairs (p_1 and p_2) compare the proposed approach to the other two conditions. Picking the *Least similar* ad was in fact not so much different than picking at *Random* in terms of the user’s ad experience.

(Q7) The sponsored ad was distracting and out of place. The Mann-Whitney test also provides us with some weak evidence to reject the null hypothesis in the distractingness question (the p -values for the three pairs are 0.0565, 0.0531, and 0.9613, respectively). This confirms the positive skew towards the *Most Similar* condition as seen in the percentages for this question in Figure 5. About 44.0% of the workers choose to Disagree or Strongly Disagree with the statement, while the same percentages are 30.6% and 32.2% for the other two conditions.

(Q5) The sponsored ad was relevant to the image set. The results for this statement were also significant and interesting. It seems that most workers (62.1%) deem the ad relevant to the

image set when it was visually consistent with its surroundings, but not for the other two conditions where there is no visual congruence. Workers disagreed on the statement 47.7% and 55.0% for the other two conditions, respectively.

This result is especially interesting, because *all ads* were chosen to be relevant or targeted to the image set in a topical sense. With the general notion of ad relevance in text-based ads, all of the ad dataset ads could have been considered “targeted”. It seems that people perceive relevance differently when it comes to images, and it comes natural to them to interpret relevance in a visual way rather than in a topical way.

The p -values show the strong statistical significance for the *Most Similar* condition versus the other two: they are $5 \cdot 10^{-5}$, $4 \cdot 10^{-7}$, and 0.328 for the three pairs, respectively. As we mention in future work, it would be very interesting to see what would happen if we take semantic relevance out of the equation and only take into account visual relevance for ad placement.

(Q4) The sponsored ad was relevant to my interests. As mentioned in the survey protocol, we required workers to be interested in the general topic of the ads that we would show them. This question was therefore more a sanity check than anything else. Percentages are in all three conditions in favor of agreement with the sentence. Still, positive agreement with the sentence is visibly higher for the *Most Similar* one over the other two. There is only weak evidence to reject the null hypothesis, however (the p -values for pairs $p1$ and $p2$ are 0.062 and 0.074, respectively).

(Q6) The sponsored ad was clearly marked. As we saw from the stats of Table II, ad recognition is not affected by the visual congruence of the ad. We had marked the ad images with the word “Sponsored” in the experiment, and as we see from the percentages shown for this question in Figure 5, most workers seem to agree with this sentence. Note that visual congruence once again does not affect how the advertisement is perceived.

V. DISCUSSION

Now we discuss related issues as well as novelty aspects of our approach. Extensions mentioned below, like multiple advertisements, diversification of results, or ad selection that takes into account other prior weights/modalities, are not investigated here; some are left out of the evaluation for clarity and consistency, and others for future work. However, we think it is important to further discuss some aspects or implications of our approach as well as clarify novelty aspects.

Multiple advertisements. All the aforementioned approaches can be extended to selecting and placing multiple advertisements. For example instead of selecting the minimum distance ad from Equation (3), one can get the k -dissimilar ones by sorting the set of dissimilarities $\{\sum_{i \in \mathcal{I}_f} \|a_j - i\|\}, a_j \in \mathcal{A}_f$ and keeping the top- k indices. This may complicate some of the naive placement strategies, but the t-SNE embedding based ones can easily cope with multiple ad placements.

Multiple modalities. The original ordering of the images is essentially a multimodal ordering⁶ of the images based on text, metadata, page-rank, and so on. In that sense, the visual appearance of the ad can be considered as another modality and incorporated in the joint ordering, together with image *aesthetics*, a modality that we chose not to investigate. Similarly, prior weights for advertisement image preference (*e.g.*, based on financial priors) can be incorporated in the selection and placement process. In such cases, though, one would need to take extra measures for ensuring that visual congruence is not heavily affected.

Monetization implications. The economic and financial implications of monetizing image search while preserving the user experience can be significant. Currently, while image intent queries account for 25% of the search queries, they go mostly unmonetized. Admittedly, image queries might not have the same commercial interest as general web queries. However, even a small improvement in monetizing image queries can lead to tens of millions of dollars in incremental revenue.

Selected advertisement rejection. In general, we do not want our approach to have a parameter (*e.g.*, a threshold) for accepting/rejecting the ad image based on visual appearance. We choose to keep it essentially *parameter free*, by only using consistency checks like reciprocity (see Section III-A) or local density (see clustering-based placement approach in § III-B) for possible rejection of the selected ad. The Mean Shift algorithm does require a scale parameter, that can however be estimated on-the-fly based on the scale of the t-SNE output. To that end, we assume that the relevant advertisement set \mathcal{A} has already been filtered by any other sense (*e.g.*, financial viability, semantic similarity, or user relevance) and only includes high-quality ad candidates. Additional parameter free conditions can be considered, *e.g.*, rejecting the ad if it is on the convex hull of the 2d t-SNE projection, but we found such further checks unnecessary in practice.

Diversification of the result set. A by-product of our clustering-based positioning approach could be a way of diversifying the result set. Given that we calculate a density-based visual clustering of the top ranked images, one may choose to limit the number of images that are shown per cluster; therefore keeping the result set more visually diverse. Of course this implies that we analyze more images (*e.g.*, the first two pages of results) and we either completely drop some of them, or just place them further down in the results.

Novelty analysis. For the ad selection part, we use simple nearest neighbor search in feature space with compact visual features as in [3] and simple aggregation for extending nearest neighbor search to vector-to-set (dis-)similarity. To the best of our knowledge, this is the first paper that uses visual similarity for advertisement retrieval given a set of images. As we discuss before, the monetization implications for such an approach are significant. Further novelty comes from the use of a two-dimensional embedding and density based clustering for presentation and ad placement, as well as from the use of state-of-the-art compact coding [22] on top of the compact feature vectors that enables our approach to run in a web-scale

⁶Image search ordering does not traditionally take visual appearance into account explicitly, although such techniques may be used for deduplication.

scenario. Finally, probably the most important contribution of this paper is the findings and implications that come from the analysis of our large crowd sourced study (see § IV-B). To the best of our knowledge, this is the first work that quantifies the effects that visual congruence has to image ad placement in web users.

Beyond image search. In this paper we present our approach using image search as a target application. However, the VCAP algorithm can be applied in a larger pool of applications, ranging from personal photo galleries, to digital magazines. In the most extreme case it could be applied in any page that contains at least one image and we also want to place an advertisement.

VI. CONCLUSIONS AND FUTURE WORK

We have presented an approach for placing native image-based ads in image search results. Our large crowd sourced user study shows that the proposed algorithm had a sizable and statistically significant impact on reducing user distraction. Our approach is able to retain brand recognition and awareness as effectively as a “distracting” ad would, while at the same time provides for a better user experience. Our method is analogous to print publications that utilized in-house design studios to ensure ads matched the magazine’s content [1], and demonstrates a similar application and utility online.

We plan to run randomized field experiments [12] to measure the impact of our algorithm on the immediate and future user engagement. Although we assume that in our approach only relevant ads are considered for placement, this is a restriction that may be relaxed if we want to further maximize visual similarity. It would be very interesting to see how the user experience changes in this setting, *i.e.*, when a user is presented with an ad that is not relevant to the images he sees semantically, but visually congruent.

In the future we also plan on investigating more into how visual congruence changes the visual perception of relevance for a user. As we found, it seems that users respond to ad relevance more if it is blending visually with the result set. Delving deep into the result set and ad positioning is another interesting future direction, as well as diversification of results using the proposed ad placement algorithm. Of high practical importance is also to extend our space model and incorporate monetization factors directly in our ad selection process.

REFERENCES

- [1] J. Abbruzzese. Vanity Fair is latest to roll out native advertising. URL, December 2013. Retrieved Oct 2, 2015 from <http://mashable.com/2013/12/13/vanity-fair-native-advertising/>.
- [2] J. Azimi, R. Zhang, Y. Zhou, V. Navalpakkam, J. Mao, and X. Fern. Visual appearance of display ads and its effect on click through rate. In *CIKM*, 2012.
- [3] A. Babenko, A. Slesarev, A. Chigorin, and V. Lempitsky. Neural codes for image retrieval. In *ECCV*, 2014.
- [4] R. A. Baeza-Yates and B. A. Ribeiro-Neto. *Modern Information Retrieval - the concepts and technology behind search, Second edition*. Pearson Education Ltd., Harlow, England, 2011.
- [5] Y. Cheng. Mean shift, mode seeking, and clustering. *PAMI*, 17(8):790–799, 1995.
- [6] S. G. Cole, L. Spalding, and A. Fayer. The brand value of rich media and video ads. Technical report, DoubleClick, 2009.
- [7] M. Datar, N. Immorlica, P. Indyk, and V. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Symposium on Computational Geometry*, 2004.
- [8] K. Dave and V. Varma. *Computational Advertising: Techniques for Targeting Relevant Ads*, volume 8 of *Foundations and Trends in Information Retrieval*. Now Publishers, 2015.
- [9] M. de Sa, V. Navalpakkam, and E. F. Churchill. Mobile advertising: Evaluating the effects of animation, user and content relevance. In *CHI*, 2013.
- [10] J. Delhumeau, P. Gosselin, H. Jegou, and P. Perez. Revisiting the VLAD image representation. In *ACM Multimedia*, 2013.
- [11] T. Ge, K. He, Q. Ke, and J. Sun. Optimized product quantization. *PAMI*, 36(4):744–755, 2014.
- [12] A. S. Gerber and D. P. Green. *Field experiments: Design, analysis, and interpretation*. W. W. Norton & Company, 2012.
- [13] A. Goldfarb. What is different about online advertising? *Review of Industrial Organization*, 44(2):115–129, 2014.
- [14] A. Goldfarb and C. Tucker. Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3):389–404, 2011.
- [15] A. Goldfarb and C. Tucker. Search engine advertising: Channel substitution when pricing ads to context. *Management Science*, 57(3):458–470, 2011.
- [16] P.-H. Gosselin, N. Murray, H. Jégou, and F. Perronnin. Revisiting the Fisher vector for fine-grained classification. *Pattern Recognition Letters*, 49:92–98, 2014.
- [17] M. Hollander, D. A. Wolfe, and E. Chicken. *Nonparametric statistical methods*. John Wiley & Sons, 2013.
- [18] P. J. Huber et al. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- [19] H. Jégou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *PAMI*, 33(1), 2011.
- [20] H. Jégou, M. Douze, C. Schmid, and P. Perez. Aggregating local descriptors into a compact image representation. In *CVPR*, 2010.
- [21] Z. Jin, Y. Hu, Y. Lin, D. Zhang, S. Lin, D. Cai, and X. Li. Complementary projection hashing. In *ICCV*, 2013.
- [22] Y. Kalantidis and Y. Avrithis. Locally optimized product quantization for approximate nearest neighbor search. In *CVPR*, 2014.
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [24] D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [25] M. Norouzi, A. Punjani, and D. J. Fleet. Fast search in Hamming space with multi-index hashing. In *CVPR*, 2012.
- [26] L. Paulevé, H. Jégou, and L. Amsaleg. Locality sensitive hashing: a comparison of hash function types and querying mechanisms. *Pattern Recognition Letters*, 31(11):1348–1358, Aug. 2010.
- [27] F. Perronnin, Y. Liu, J. Sanchez, and H. Poirier. Large-scale image retrieval with compressed Fisher vectors. In *CVPR*, 2010.
- [28] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, L. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *arXiv preprint arXiv:1409.0575*, 2014.
- [29] J. Sivic and A. Zisserman. Video Google: A text retrieval approach to object matching in videos. In *ICCV*, pages 1470–1477, 2003.
- [30] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li. YFCC100M: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.
- [31] G. Toliás, Y. Avrithis, and H. Jégou. To aggregate or not to aggregate: Selective match kernels for image search. In *ICCV*, 2013.
- [32] L. Van Der Maaten. Accelerating t-SNE using tree-based algorithms. *JMLR*, 15(1):3221–3245, 2014.
- [33] L. Van der Maaten and G. Hinton. Visualizing data using t-SNE. *JMLR*, 9(2579-2605):85, 2008.
- [34] A. Vedaldi and S. Soatto. Quick shift and kernel methods for mode seeking. In *ECCV*, 2008.