

Discovering Triangles in Portraits for Supporting Photographic Creation

Siqiong He, Zihan Zhou, Farshid Farhat, and James Z. Wang

Abstract—Incorporating the concept of triangles in photos is an effective composition technique used by professional photographers for making pictures more interesting or dynamic. Information on the locations of the embedded triangles is valuable for comparing the composition of portrait photos, which can be further leveraged by a retrieval system or used by the photographers. This paper presents a system to automatically detect embedded triangles in portrait photos. The problem is challenging because the triangles used in portraits are often not clearly defined by straight lines. The system first extracts a set of filtered line segments as candidate triangle sides, and then utilizes a modified RANSAC algorithm to fit triangles onto the set of line segments. We propose two metrics, *Continuity Ratio* and *Total Ratio*, to evaluate the fitted triangles; those with high fitting scores are taken as detected triangles. Experimental results have demonstrated high accuracy in locating preeminent triangles in portraits without dependence on the camera or lens parameters. To demonstrate the benefits of our method to digital photography, we have developed two novel applications that aim to help users compose high-quality photos. In the first application, we develop a human position and pose recommendation system by retrieving and presenting compositionally similar photos taken by competent photographers. The second application is a novel sketch-based triangle retrieval system which searches for photos containing a specific triangular configuration. User studies have been conducted to validate the effectiveness of these approaches.

Index Terms—Portrait Photography; Triangle; Photo Composition.

I. INTRODUCTION

With the widespread use of mobile cameras (e.g., smart phones), social media, and cloud storage, many more photos are being taken by consumers these days. However, ordinary users do not possess professional skills in *photo composing*, resulting in many unmemorable photos. As a result, there is a growing interest in the multimedia research community to develop intelligent software that can help amateur photographers and photo enthusiasts improve their photo-taking skills [1], [2], [3], [4]. Besides, the capability of recognizing photo composition can be important in a wide variety of applications such as photo collection summarization [5], aesthetic assessment [6], and suggesting composition improvements through image re-targeting [7], [8].

In pictorial art, good composition is viewed as a congruity or agreement among all the elements in a design [9]. The

S. He, Z. Zhou and J. Z. Wang are with College of Information Sciences and Technology, The Pennsylvania State University, University Park, Pennsylvania, USA (e-mail: hesiqiong@gmail.com; zzhou@ist.psu.edu; jwang@ist.psu.edu). F. Farhat is with the School of Electrical Engineering and Computer Science, The Pennsylvania State University, University Park, Pennsylvania, USA (e-mail: fuf111@cse.psu.edu).

Correspondence should be addressed to J. Z. Wang.

design elements appear to belong together as if there are some implicit visual connections between them. Another term to describe this form of unity is *harmony*. By reflecting this principle in photography, subjects in one scene should not aimlessly scatter around. Instead, they should unify to provide an overall impression to the viewer. To convey such unity in photos, professional photographers have designed dozens of executable techniques for composition. One universal and interesting technique is to embed basic geometrical shapes in photographic compositions [9], [10]. Human beings begin to learn about basic geometrical shapes such as circles, rectangles, and triangles since very young age. Even toddlers are capable of recognizing those basic shapes immediately. Valenzuela [10] suggests that we can explicitly or implicitly embed basic geometrical shapes in photos to attract viewers. Moreover, since these shapes are instantly recognized, subjects bounded within such shapes or an implicit construction of such shapes are perceived as a unity.




Fig. 1. Examples of the use of the triangle technique in oil paintings and photography. **First row:** Last Supper, The Birth of Venus, The Scream. **Second row:** Mona Lisa, Girl with a Pearl Earring, The Arnolfini Portrait, The Dream, Afghan Girl.

Among all basic geometric shapes, the *triangle* is arguably the most popular shape utilized by professional photographers to make a composition more interesting [11]. Such compositional technique has been called the *triangle technique*. The triangle technique has been extensively used in visual arts (Fig. 1). In particular, portrait photographer will embed “triangular” poses in photos to make an interesting composition (Figs. 7 and 8). In this work, we attempt to automatically detect triangles in portraits, which encompasses fashion, beauty, baby, family, and advertising photography, with the goal of using such capability in supporting the creative process of photographers.

Two fundamental questions are often raised when analyzing

the composition of a portrait: (1) Where are the subjects in the photo? And (2) how do the subjects pose? Traditional composition rules provide us with guidelines to answer the first question. For example, the rule of thirds suggests that putting the key aspect(s) of the subjects near one-third point(s) or line(s) of a photo is more appealing than at the center. Based on these rules, several methods have been developed to model and assess the positioning of subjects in a photo.

Nevertheless, the second question remains a challenge. To address this problem, we leverage an important observation in portrait photography [10], [11], [12], [13]: *experienced portrait photographers often use the triangle technique to create interesting, dynamic, and appealing poses for the subject*. For example, a widely-used rule for posing is to avoid 90-degree body angles, because they often look unnatural and strained. In addition, the triangle technique is also frequently used to unify multiple subjects with the surrounding environment, such as chairs and lamps, in the portrait photo.

Despite the popularity of triangle techniques in portrait photography, it is often difficult for less experienced amateurs to recognize such triangles, because most triangles do not have explicit edges and sometimes are even constructed by different objects. Moreover, triangles in portraits can be of various sizes, shapes, orientations, and appearances. Hence, our goal is to automatically detect potential triangles from professional photographers' work in order to help amateurs recognize and learn from the usage of the triangle technique.

In this paper, we adopt a line-based approach to detect triangles. But unlike photos of some man-made objects, straight lines are often incomplete or absent in portraits. To overcome this challenge, our algorithm is divided into two steps. First, a line segment detection module is used to extract candidate line segments from an image, which are subsequently filtered using the global contour information in the image. Second, the filtered line segments are fed into a triangle detection module as the candidate sides of triangles. Specifically, a RANSAC algorithm is developed to randomly pick two sides from all the candidates and fit the triangle. Two metrics, *Continuity Ratio* and *Total Ratio*, are defined to evaluate the fitness of these triangles. Only those triangles with high fitness scores will be retained.

To demonstrate the use of our triangle detection technique, we further develop two applications which aim to provide users with on-site feedback about the composition of their photos, especially in terms of the use of triangles. In the first application, we design an image retrieval system which explores the large amount of photos taken by competent photographers and available on online photo-sharing websites to provide guidance to amateur users on composing high-quality photos with a person in the scene. Given an image of the scene, we find exemplar photos in a large-scale image dataset by considering semantic similarity and the use of triangles. The retrieved photos, presented to the user, serve as an informative guide for attaining better photo composition.

In the second application, the user can provide a sketch indicating the desired shape and orientation of the triangle. Given the triangle sketch, our system will retrieve exemplar photos containing the specific triangular configuration from

a collection of photos taken by experienced or accomplished photographers. Here, the use of triangle sketches is motivated by the following two facts. First, a professional portrait photo typically contains multiple interesting triangles. The triangle sketch allows users to examine one triangle at a time, and gain a deeper understanding on such configurations across multiple images. Second, business users such as photo book or magazine editors may wish to embed certain triangular configuration in the image in order to fill in a specific page layout or a selected collection of photos about a particular subject (e.g. a celebrity).

Our research makes the following main contributions:

- *Triangle detection in portraits*: We study an important yet largely unexplored principle in portrait photography, namely the use of triangles. We propose a low-complexity and effective method to detect triangles in portraits with a variety of sizes, shapes, orientations, and appearances, with the goal of helping less experienced users to understand and learn from the usage of triangles in professional photographers' work. Our method is able to identify significant triangles in the photo even when straight lines are incomplete or absent. We have carried out a systematic experimental study on a set of portrait photos with triangles annotated by a professional photographer to validate the proposed method.
- *Composition-sensitive retrieval for on-site feedback*: We develop two new applications to demonstrate the use of our triangle detection method in portrait photography. Specifically, we design novel image retrieval systems which can provide guidance to both amateur photographers and business users, especially on the use of triangles. User studies have been conducted to validate the effectiveness of the retrieved exemplars in providing users with helpful information and suggestions on photo composition.

The remainder of the paper is organized as follows. Related work is discussed in Section II. Section III describes the triangle detection algorithm. In Section IV, experiments and results are reported. Section V describes our novel image retrieval systems for on-site feedback to photographers and validates their effectiveness in providing useful information on photo composition via user studies. We conclude and suggest future research in Section VI.

II. RELATED WORK

A. Composition Modeling

Modeling composition of a photo is challenging because it requires understanding of the relationship among visual elements. A user study indicates that composition is the most important feature related to the aesthetic quality of a photo [6], [14], [15]. Generally, viewers prefer simple and clear compositions. Hence, early work proposed high-level composition features based on locations and orientations of long dominant lines in images [16] to capture the simplicity of the composition. A more intelligent way to describe composition is to model popular composition rules such as the rule of thirds and the rule of golden ratio. For example,

researchers design features to model compositions following the rule of thirds [17], [18], [19], [20]. Further, Bhattacharya *et al.* [8] enhances the quality of amateur photos by adjusting compositions based on the rule of golden ratio, which suggests that the position of the horizon line in scenic photos should be adjusted to satisfy the golden ratio. In addition to these top-down rule-based approaches, Su *et al.* [21] proposes a bottom-up scheme to represent photo compositions by partitioning the photo into multi-size image patches and extracting features from different spaces to model color, texture, saliency, and edge information jointly.

Composition models have been applied to various real-world applications. For example, Obrador *et al.* [6] shows that by using only the composition features, one can achieve image aesthetic classification results that are comparable to the state of the art. Recently, composition rules and features have also been used to predict high-level attributes for image interest-iness classification [20] and to develop both automatic and interactive cropping and retargeting tools for image enhancement [7], [8], [22], [23], [24]. Another line of research focuses on providing onsite aesthetics guidance to the photographers at the point of photographic creation, such as choosing the best view [2], [3], recommending the locations of the subjects (*i.e.* the person or persons) in landscape photography [25], [26], and suggesting the appropriate camera parameters [4]. In addition, Yao *et al.* [1] proposes a composition-sensitive image retrieval method which classifies images into horizontal, vertical, diagonal, textured, and centered categories, and uses the classification result to retrieve exemplar images that have similar visual and compositional characteristics as the query image. However, none of them have studied the use of triangles in photography.

B. Portrait Photo Analysis

Few studies have focused on aesthetic analysis of portrait images. Jin *et al.* [27] studied the critical role of lighting in portrait photography. Varying lighting patterns shift highlights and shadows on the face, change the area ratios between them, and generate 3D perception from the 2D photograph. While the learned artistic portrait lighting templates in that work are able to capture the arrangement of low-level features, our work aims at modeling the usage of more holistic composition techniques in portrait photography. Recently, Zhang *et al.* [28] and Ma *et al.* [29] developed methods to recommend suitable poses of people in natural scenes. However, the recommendation is based on matching attentional and compositional features of the backgrounds and manually labeled human body poses, whereas our work focuses on automatic detection of triangles in portraits. Note that the use of the triangle technique in portraits is not limited to those formed by human body parts. As shown in Fig. 1, it can be used in photos containing one person or multiple persons. It can also be used to connect persons with objects in the same set up, such as chairs, blocks, etc.

III. THE METHODS

We introduce below the line segment detection algorithm (Section III-A) and discuss how triangles can be constructed

from these line segments (Section III-B).

A. Line Segment Detection

By examining high-quality portraits designed with the triangle technique (*e.g.*, Figs. 1, 7 and 8), we observe that triangles present in portraits are often composed with parts of contours or edges, such as the contours of faces, body parts, apparels, as well as edges formed by multiple subjects and other objects. In general, a triangle is geometrically defined as a polygon with three corners and three sides, where the three sides are all straight line segments. However, contours of natural objects like humans or hats are often somewhat curved. Therefore, to detect potential triangles in real-world images, our method needs to be able to identify such curved line segments, in addition to straight edges.

To achieve this, we employ the Line Segment Detector (LSD) proposed by [30] to convert gradient map of an image to a set of line segments. It first calculates a level-line angle at each pixel to produce a *level-line field*. The level line is a straight line perpendicular to the gradient at each pixel. Then, the image is partitioned into *line-support regions* by grouping connected pixels that share the same angle up to a certain tolerance. Each line-support region is treated as a candidate line segment. Next, a hypothesis testing framework is developed to test each line segment candidate. The framework approximates each line-support region with a rectangle and compare the number of “aligned points” in each rectangle in the original image with the *expected* number of aligned points in a random image. A line segment is detected if the actual number of aligned points in a rectangle is significantly larger than the expected number.




Fig. 2. LSD results with different density (as indicated).

A useful property of LSD is that, by approximating the line-support region using a rectangle of certain length, it is able to detect near-straight curves in the image. Further, it is easy to see that a larger rectangle with more unaligned points will be needed to cover a more curved line segment. Therefore, by setting a threshold on the *density* of a rectangle, which is defined as the proportion of aligned points in the rectangle, we can control the degree up to which a curved line segment is considered. Fig. 2 shows results of the line segment detector with different density values. In our experiments, the density is empirically set to 0.2.

While the line segment detector aims at extracting all potential line segments from an image using *local* image gradient cues, some line segments are more *globally* distinguishable, thus more visually attractive to viewers. To further




Fig. 3. Filtering the line segments. The parameter α is set to 0.5 and 0.7 in (d) and (e), respectively.



Fig. 4. Challenges in detecting triangles: Outliers and imperfect sides. (a) outliers: irrelevant objects, (b) outliers: multiple triangles, (c) imperfect sides: occlusion, (d) imperfect sides: bending curve.

identify such line segments, we combine the line segment detector with the ultrametric contour map (UCM) obtained by the hierarchical image segmentation algorithm [31], [32]. As shown in Fig. 3(b), the contour map has the same size as the input image. Each pixel on the contour map holds a confidence level between 0 and 1, indicating the possibility of it being on a boundary.

Next, given a line segment, we identify all the pixels falling in its support region and consider the maximum confidence level of all pixels as the confidence level of the line segment. Then, line segments with confidence levels under a certain threshold are removed, where the threshold is chosen based on the maximum confidence level present in an image. Specifically, assume the maximum confidence level of all the line segments in an image is C , where $C \in [0, 1]$, then we set the threshold as $(1 - \alpha)C$ and accept line segments whose confidence levels are within the range $[(1 - \alpha)C, C]$. The parameter α controls the number of accepted line segments. Smaller α filters out more line segments from an image, as shown in Fig. 3. In this paper, we empirically set $\alpha = 0.5$.

B. Fitting Triangles

The line segment detector described above gives us a set of candidate triangle sides. Randomly selecting three non-parallel sides from the set generates a triangle. Thus, the problem of detecting a triangle can be converted into finding three non-parallel sides. Although a triangle consists of three sides, a triangle can be uniquely defined by two sides, because the third side can be obtained by connecting the end points of the other two sides. Moreover, the presence of the third side is not as important in the practical usage of the triangle technique because viewers can easily “complete” the geometric shapes themselves. As a result, our problem is reduced to fitting a triangle using two non-parallel line segments selected from the candidate set.

Two major challenges still exist in fitting triangles using the extracted line segments: (1) there is a large number of

irrelevant line segments; and (2) the sides of a triangle are imperfect in real images. For example, as shown in Fig. 4(a), some objects in a photograph are irrelevant to the use of the triangle technique, even though they have high-contrast contours (e.g., the bird pattern on the woman’s shirt). The line segments produced by such objects are all outliers. In addition, multiple triangles often exist in one image (e.g., Fig. 4(b)). Thus, line segments from different triangles should also be considered as outliers w.r.t. each other. In Figs. 4(c) and 4(d), we further show some examples of imperfect triangle sides. As shown, these sides may be broken into several parts because of occlusion or artificial effects introduced by the line segment detector. For instance, the woman in Fig. 4(c) has her right arm occluded by her left arm. As a result, the contour of her right arm is broken into two line segments. In Fig. 4(d), the contour of the back of the subject is too curved to be approximated by a single straight line. Therefore, the line segment detector approximates it using two straight line segments with slightly different orientations.

In order to tackle these two challenges, we employ a modified RANSAC (RANDOM SAMPLE CONSENSUS) algorithm in favor of its insensitivity to outliers. RANSAC is an iterative method that robustly fits a set of observed data points (including outliers) to a pre-defined model. Our algorithm includes three steps:

- 1) Two non-parallel line segments are randomly selected from the candidate set and extended to generate two lines on which the two triangle sides lie.
- 2) All the candidate line segments within neighborhoods of the two lines are projected onto the corresponding lines, resulting in a number of projected pixels. Triangle sides are then constructed from all the projected pixels.
- 3) Once two triangle sides are constructed, two metrics *Continuity Ratio* and *Total Ratio* are calculated to measure the fitness and significance of the triangle, respectively. Triangles with high scores are accepted.

We now describe each step in details.

1) *Identifying Sides From Line Segments*: By extending the two randomly selected line segments to two lines, we obtain the shared end point of the two sides, i.e., the intersection of the lines. Moreover, two intersecting lines generate four different angles with four different opening directions: upwards, downwards, leftwards, and rightwards. Each angle corresponds to a category of triangles that contain this angle and two sides of varied lengths. Given one of the four angles, once the lengths of its two sides are determined, a unique triangle can be constructed.

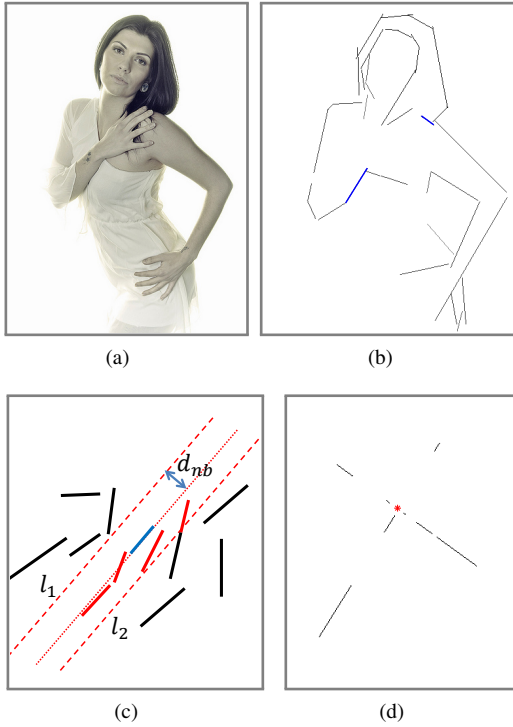


Fig. 5. Illustration of the RANSAC algorithm. (a) original image, (b) candidate sides, (c) neighborhood region, and (d) fitted pixels.

2) *Fitting All Segments on Sides*: In this step, we first mark all line segments within neighborhood of the two straight lines formed in the previous step as inliers and those falling outside neighborhood as outliers. The neighborhood region of a straight line $l: ax + by + c = 0$ is defined as

$$N(l) = \{(x, y) \in \mathbb{R}^2 \text{ and } \frac{|ax + by + c|}{\sqrt{a^2 + b^2}} \leq d_{nb}\},$$

i.e., the group of pixels whose distances to the straight line are smaller than a certain threshold d_{nb} . In our experiments, we fix $d_{nb} = 5$ pixels. Then, all the inlier line segments with respect to line l can be calculated as $I(l) = S \cap N(l)$ where S is the set of all candidate line segments. Note that, if a line segment is cut into two parts by the neighborhood boundary, the part of line segment falling within the neighborhood is included as an inlier, whereas the other part is considered as an outlier. In Fig. 5(c), we illustrate how the neighborhood of a straight line is utilized to partition all the line segments into inliers and outliers. The blue line segment is selected from the candidate line segment set and extended to a straight line l . Lines l_1 and l_2 designate boundaries of the neighborhood region. Both of them have a distance d_{nb} from l . The red line segments and black line segments represent the inliers and outliers, respectively.

Next, all the pixels on the inlier line segments are projected onto the straight line. A pixel on the straight line is called a *projected pixel* if there is at least one pixel on any inlier line segment that is projected to this pixel. We denote the set of all projected pixels as

$$P(l) = \{(x', y') \in l \mid \exists (x, y) \in I(l), (x - x', y - y') \perp l\}.$$

Fig. 5 illustrates the fitting process. Fig. 5(b) shows all the line segments extracted from the image shown in Fig. 5(a). The two line segments colored in blue are randomly selected from the candidate set. Subsequently, two straight lines, l and \tilde{l} , are generated by expanding the two line segments and their neighborhood regions are identified. Finally, inliers within their neighborhoods are projected onto the straight lines, as shown in Fig. 5(d).

Here, we note that the projected pixels typically scatter along the entire straight line. To form a triangle, we divide the line into two *half lines* at the intersection point. Formally, a half line here is defined as a straight line extending from the intersection point indefinitely in one direction only. For each pair of straight lines l and \tilde{l} , four triangles can be formed using different pairs of half lines. Therefore, when evaluating the fit of a triangle, we only consider the subsets of projected pixels which are on the two half lines that form the triangle, denoted as $P(l^h)$ and $P(\tilde{l}^h)$, as opposed to the entire sets of projected pixels $P(l)$ and $P(\tilde{l})$.

3) *Evaluating the Fitted Triangle*: In order to evaluate the quality of a fitted triangle, we first introduce a *Continuity Ratio* score to evaluate the quality of a fitted side. Here we use one half line l^h as an example. The way to calculate continuity ratio for the other half line \tilde{l}^h is exactly the same. Given any point X lying on the half line l^h , we can construct a potential side OX connecting the intersection point $O = (x_o, y_o)$ and X . We further compute the number of pixels projected onto OX divided by the length of OX as $\frac{|P(l^h) \cap OX|}{|OX|}$. Then, the *Continuity Ratio* of l^h is defined as the ratio of the best-fitted side on l^h :

$$C(l^h) = \max_{X \in P(l^h)} \frac{|P(l^h) \cap OX|}{|OX|}. \quad (1)$$

Finally, the continuity ratio for the entire triangle constructed by the two half lines l^h and \tilde{l}^h is

$$C(l^h, \tilde{l}^h) = C(l^h) \times C(\tilde{l}^h). \quad (2)$$

In addition to the continuity ratio, we define another *Total Ratio* score which is calculated as the area of the triangle divided by the area of the entire picture. Intuitively, the *Continuity Ratio* describes how well the extracted line segments fit the given side, while the *Total Ratio* represents the significance of a fitted triangle in magnitude. Because a larger triangle can be more easily recognized by viewers and has more impact on composition of the entire photo, we only retain the triangles with sufficiently high *Continuity Ratio* and *Total Ratio* scores.

IV. EXPERIMENTS

In order to evaluate the performance of our triangle detection method for portrait images, we built a dataset by collecting 4,451 professional-quality studio portrait photos from Flickr. Below we provide both qualitative and quantitative results to demonstrate the effectiveness of our method.

In the experiments, we use the publicly available LSD code¹ to detect line segments, and the image segmentation code² to

¹<http://www.ipol.im/pub/art/2012/gjmr-lsd/>

²<https://github.com/jponttuset/mcg>

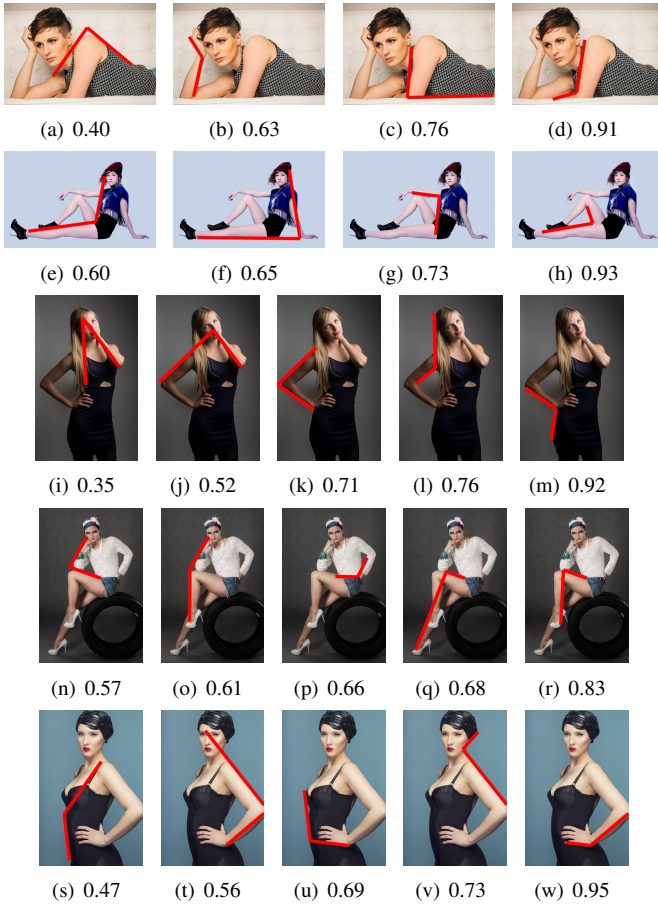


Fig. 6. Detected triangles with different *Continuity Ratios*. The red lines indicate two sides of the detected triangle. Each triangle indicates an interesting pose or composition element in the photo.

compute the ultrametric contour map (UCM). On a computer with a 2.4 GHz CPU and 8GB RAM, detecting the line segments and computing the contour map on an image take about 0.15 second and 1.5 seconds, respectively. Using our un-optimized MATLAB implementation, filtering weak line segments takes 0.2 second, and fitting the triangles takes 0.3 second. Therefore, our method takes 2 to 3 seconds in total to detect triangles in an image, thus is suitable for real-time applications.

A. Qualitative Results

Fig. 6 presents some detected triangles with different continuity ratios. Triangles with higher continuity ratios are often easier to recognize by humans. However, they do not necessarily outperform those of low continuity ratios in conveying useful compositional information about the photo. For instance, Fig. 6(i) has a much lower continuity ratio than Fig. 6(m) but it arguably conveys a more interesting pose or compositional skill. From the detected triangle, we can notice that the model slightly tilts her head to align with her left arm which constructs a charming triangle with her hair. It is common that multiple different triangles are embedded in one image, and some of them can be easily overlooked by amateurs. Here, our goal is to identify many potential triangles

from an image, in order to empower amateur photographers to better learn about composition techniques.

More detected results can be found in Fig. 7. It demonstrates that our triangle detection system can clearly identify triangles in portrait photographs despite the existence of diversions such as hair, props, and shadows/patterns/folds on cloths, etc. Moreover, triangles involving multiple people can be accurately detected as well (e.g., the fourth picture in the first row and the third picture in the second row). More interestingly, our system is able to locate triangles that may not be straightforward to human observers. For instance, considering the first photo in the last row, people typically focus on the triangle formed by two arms, while overlooking an appealing triangle constructed with one arm of the girl and the edge of her lower jaw. Another example is the fifth picture in the last row. The subject puts her arm along her dress in an intentionally designed pose so that it extends the boundary of her dress and forms a large triangle with her long hair. Such examples indicate that professional photographers design intricate poses for the subjects in order to achieve quality photo composition. However, choosing an interesting pose requires much experience and artistic inspiration. The triangles detected by a computer system can help amateurs gain deeper understanding and inspirations from high-quality photography works.

B. Quantitative Evaluation

To further study the effectiveness of our method, we involved an experienced professional portrait photographer, who has studied arts and architecture and has operated a professional portrait studio for over 30 years. Using the online annotation tool LabelMe³, we asked the photographer to manually annotate triangles that indicate interesting pose or composition in the images we collected. A total of 173 images were annotated by the photographer. We show some example triangles in Fig. 8, which illustrate a wide variety of numbers, sizes and orientations of triangles used in the portrait photography. Because professional annotations can be valuable to the research community in studying portrait composition, we will make the dataset and any future extensions freely available to researchers.

We compare the triangles detected by our method with the annotations by the professional. In this experiment, we only consider triangles whose continuity ratio and total ratio are above 0.1. Let (A, B, C) denote the set of vertices of a ground truth triangle, we consider a candidate triangle (A', B', C') a matching triangle if

$$\frac{|AA'| + |BB'| + |CC'|}{|AB| + |BC| + |CA|} \leq \delta, \quad (3)$$

where $|AB|$ is the length of the line segment connecting A and B , and δ is a threshold. We fix $\delta = 0.3$ in the experiment.

In Fig. 9, we report the *precision* and *recall* of our method as a function of the continuity ratio, the total ratio, and parameter α , respectively. Specifically, let G denote the set of ground

³<http://labelme.csail.mit.edu/Release3.0/>



Fig. 7. Examples of detected triangles in portrait photographs. Each triangle indicates an interesting pose or composition element in the photo.




Fig. 8. Example portrait photos with triangles labeled by an experienced professional photographer.

truth triangles annotated by the professional, and Q denote the set of triangles detected by our algorithm under a particular experiment setting, the precision and recall are defined as follows:

$$Precision = \frac{|G \cap Q|}{|Q|}, \quad Recall = \frac{|G \cap Q|}{|G|}. \quad (4)$$

As shown in Fig. 9(a), the precision of our method increases as the continuity ratio increases. When the continuity ratio is high, about half of the triangles detected are true positives. Meanwhile, among all manually annotated triangles, up to about half of them can be detected by our method (i.e., when continuity ratio is 0.1). In Fig. 10 (first row) we show some triangles missed by our algorithm. For many such triangles, there is a lack of explicit edges or contours in the image. For example, in the first image of Fig. 10, the triangle is formed by the pair of shoes and the subject’s knees, instead of explicit edges. Such triangles are often known as *implicit triangles* in photography. It will be an interesting future work to look for them using machine vision. In addition, Fig. 9(b) shows that our method achieves the best precision when the total ratio is about 0.2. This may suggest the most common triangle sizes in portrait photography. Lastly, Fig. 9(c) shows that, as α increases, the recall of our method increases, whereas the precision decreases. This is expected because triangles with weaker sides are also included when the parameter α gets larger.

Finally, we examine the triangles that are detected by our algorithm but not labeled by the professional. Fig. 10 (second row) shows some interesting cases where we think the triangles are actually meaningful. These examples suggest that even experienced photographers may occasionally overlook certain interesting elements, and our algorithm could potentially pro-




Fig. 9. Quantitative evaluation of triangle detection in portraits. The plots show the precision and recall of our method as a function of (a) continuity ratio, (b) total ratio, and (c) parameter α .




Fig. 10. Comparison of triangles detected by our method with professional annotations. **First row:** Triangles annotated by the professional photographer but missed by our algorithm. **Second row:** Some interesting triangles detected by our algorithm, but not labeled by the professional.

vide them with other possible interpretations of the photo.

V. PHOTOGRAPHIC APPLICATIONS

It is often difficult for untrained photographers to naturally incorporate triangles in compositions in order to form striking portraits. Our triangle technique detection method can help amateur photographers design interesting poses when taking portrait photos. Certain business users, *e.g.*, photo book or magazine editors, can also benefit from the triangle detection method. In this section, we demonstrate two applications of our method for helping amateur users take high-quality photos with a person in the scene and assisting business users select portraits to fill a specific page layout or a selected collection.

A. Position and Pose Recommendation for the Subject

Nowadays, cloud-based photo sharing services such as flickr.com, dpchallenge.com, and 500px.com allow photographers to access millions of photos taken by their peers around the world. We describe a novel algorithm which explores such resources to help photographers take high-quality photos *with a person in the scene* via composition-sensitive image retrieval. Specifically, given a sample shot of the scene (without the subject in it) taken by the user, we propose to find exemplar photos of *similar scenes* in a large collection of photos with a person in the scene. These photos

can then be used as feedback to the user for determining the position and pose of the person in the scene.

In this application, we have collected a set of 46,198 photos from the online photography community 500px.com. Each of these photos contains a person in the scene. Given a sample shot, our goal is to find a set of exemplar photos in the large collection by considering (i) the semantic similarity, and (ii) the use of triangles. Specifically, we first retrieve K images with similar semantics from the image collection, which gives us a semantic ranking $r_s(k), \forall k \in 1, \dots, K$. When the use of triangles is considered, we further compute a triangle ranking $r_t(k), \forall k \in 1, \dots, K$ based on the triangles detected by our method in the image. The final ranking is then determined as

$$r(k) = \beta r_s(k) + (1 - \beta) r_t(k). \quad (5)$$

Through this, images that incorporate more triangles are moved to the top of the ranking list. In our experiments, we fix $K = 20$ and $\beta = 0.5$.

Fig. 11 shows the top-ranked images retrieved for various query images. The retrieved images not only have similar scene semantics as the query, but also include various interesting poses. Below we describe each ranking scheme in detail. **Semantic similarity.** When computing the semantic similarity between two photos, we use the generic descriptors extracted from the convolutional neural networks (CNNs), which have recently produced state-of-the-art content-based image retrieval results [34]. We adopt the publicly available CNN model trained on the ImageNet ILSVRC challenge dataset [35], and represent each image using the ℓ_2 -normalized output of the second fully connected layer (full7 of [35]). Then, the cosine distance is used to measure the feature similarity.

The use of triangles. We apply the proposed method to detect triangles in each of the K images. Since we are focusing on the use of triangles in portrait (not the background), we first obtain the bounding box of the person in the image using the state-of-the-art object detection algorithm [36] and only detect triangles within the bounding box. Note that, in the case of multiple detections returned by the object detection algorithm, we simply keep the bounding box with the largest area. Further, based on the experiment results in Section IV, we only consider triangles with continuity ratio higher than 0.5 and total ratio higher than 0.1. Finally, the triangle ranking

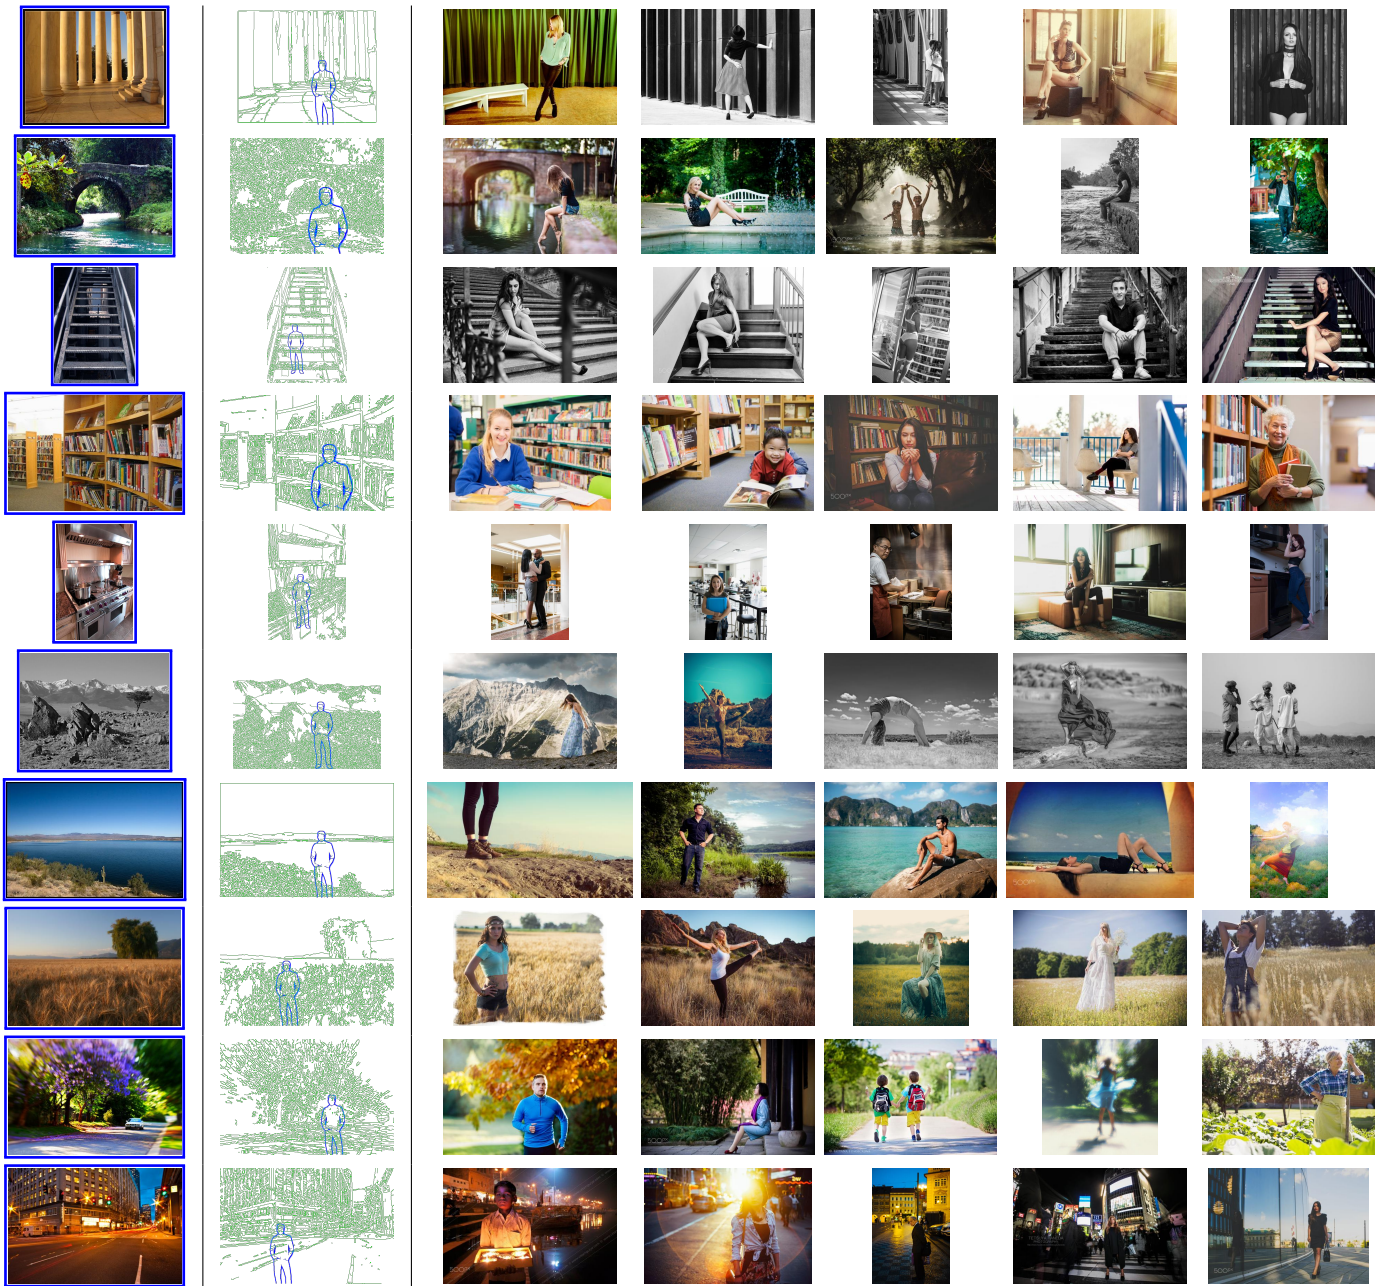


Fig. 11. Examples of human position and pose recommendation for various scene categories. Each row shows a query image, human position recommendation made by Where2Stand [33], and top-5 rank photos retrieved by our method. **Rows 1-3:** Architecture scenes. **Rows 4-5:** Interior scenes. **Rows 6-8:** Landscape scenes. **Rows 9-10:** Street scenes.

is determined by sorting the number of detected triangles in decreasing order.

It is worth mentioning that as our system focuses on recommending interesting positions and poses through the use of triangles, it could potentially complement existing systems that analyze other aspects of photo aesthetics, such as content and color. In terms of speed, both CNN feature extraction and triangle detection for the database can be done off-line. Thus, given a query image, it only takes about 0.25 second for our system to retrieve exemplar photos from the database.

Comparison to Existing Work. To quantitatively evaluate the performance of our system, we compare it to a recently

published human position recommendation system called Where2Stand [33], which is an Android app. By analyzing the content of the query image, it displays a human-shaped icon in the viewfinder that guides the user in positioning the subject in the scene as shown in Fig. 11. Note that Where2Stand does not make recommendations on the pose of the person.

We have conducted a user study which asks the participants to compare the usefulness of the two systems. The query images are chosen from the open AVA dataset [15]. The original AVA dataset contains over 250,000 images along with semantic tags describing the scene type. In this study, we randomly select 200 images as queries: from four most common scene categories in photography: “Architecture”, “Interior”,

“Landscape”, and “Street”. Since our goal is to recommend suitable position and pose for the subject, we only use images where there is a reasonable possibility to include a subject in the depicted scene. For example, we do not include images that focus on still lifes (e.g., a vase), plants and animals (e.g., a flower or a bird), or the sky and the sea.




Fig. 12. The interface in the user study. Given a query, the user is asked to rank the two systems based on how useful they are in terms of helping the user compose high-quality photos with a person in the scene. To avoid any biases, no information about the two systems was provided during the survey. For a fair comparison, we present the recommendation made by our system via an image collage containing the top-5 ranked photos.

We have recruited 16 participants to this study, mostly graduate students at Penn State with some basic photography knowledge. At our user study website, each participant is assigned with a subset of 30 randomly selected query images. For each query, we show the recommendations made by both systems and ask the user to rank the two systems based on how useful they are in terms of helping the user compose high-quality photos with a person in the scene. To avoid any biases, no information about the two systems was provided during the study (see Fig. 12). Further, for a fair comparison, we present the recommendation made by our system via an image collage containing the top-5 ranked photos.

Table I shows the user preference over recommendations made by the two systems for different scene categories. Overall, our system is ranked better for 68.1% of the time, whereas Where2Stand is ranked better for only 31.9% of the time. Although our method outperforms Where2Stand on all four categories, we note that it performs particularly well in “Interior” and “Landscape” scenes. We hypothesize that Where2Stand’s inferior performance in “Interior” scenes may be due to its training using “souvenir photos” which are typically taken outdoors. Also, the larger gaps in performance for “Interior” and “Landscape” scenes may be due to the increased importance of human poses in such scenes, for which Where2Stand is unable to assist.

TABLE I
USER PREFERENCE ON RECOMMENDATIONS FOR VARIOUS SCENE CATEGORIES.

	Architecture	Interior	Landscape	Street
Our Method	66.2%	74.3%	68.2%	63.7%
Where2Stand [33]	33.8%	25.7%	31.8%	36.3%

The retrieved examples in Fig. 11 show that the exemplar images retrieved by our method provide practical and inspiring information to the users on the human position and pose. They can also suggest different and potentially more interesting camera orientation, use of lighting, aspect ratio, and more. Meanwhile, the recommendation made by Where2Stand is restricted to the position of the subject in the scene. While such

guidance may be effective especially for “souvenir photography,” it has limitations. First, it largely ignores the interaction between the subject and the surrounding environment, which is particularly important in man-made environments (e.g., rows 1, 2, and 5 in Fig. 11). Second, since its focus is on the 2D image content, Where2Stand has difficulty inferring the correct 3D depth of the scene or the camera pose, resulting in impractical recommendations in some cases (e.g., rows 3, 5, and 8 of Fig. 11). By exploring the large-scale dataset collected from online photo-sharing platforms, our system provides a more effective solution to these issues.

B. Portrait Retrieval using Triangle Sketches

In this application, we consider another interesting scenario that photographers or magazine editors want to select portraits to fill a specific page layout or a selected collection about a subject. The shapes and orientations of embedded triangles within these photos can be critical to the overall page composition. Therefore, we develop a portrait retrieval system that can take a *triangle sketch* as a query to help users find images containing a targeted triangular configuration. Such tools can be particularly useful when a large collection of portraits are available to choose from (e.g. for a celebrity).

First, users need to provide a sketch indicating the shape and the orientation of the triangular configuration that they desire to have in the photos. In portrait photos, the third side of a triangle is often missing in most photos. Thus, the sketch query provided by users is basically an angle with two sides. Given the sketch angle, we compute the *orientations* of the two sides as well as the *opening direction* of the angle. Specifically, the orientation of a side is defined to be the angle between its extended straight line and the positive x -axis. An angle has four possible opening directions: upward, downward, leftward, or rightward. The two properties narrow down the search space to a specific type of triangles.

Recall that our triangle detection algorithm has two steps: line segment detection and fitting triangles. All line segments detected from the first step are taken as candidate triangle sides during the fitting stage. In the retrieval system, we construct two candidate sets containing line segments with similar orientations as the two sides of the sketched angle. Two sides can then be randomly selected from the two candidate sets and the combination of them generates four angles with four different opening directions (Fig. 5(d)). Only the angle that has the same opening direction as the sketched angle will be taken into consideration during the fitting process. Such a sketch-based triangle retrieval system not only assists users in searching for desired types of triangles but also reduces the search time significantly for large photo collections.

Performance Evaluation. We have conducted a human subject study to validate the effectiveness of our method. In this study, we selected 20 groups of representative queries which covered a wide range of angles in terms of magnitude and orientation. We only consider angles in the range of $[45^\circ, 135^\circ]$ because angles that are either too large or too small are often not perceived as interesting ones by humans. Each of the 20 groups of queries takes a distinct combination of orientations

for two straight lines. Twenty line combinations $\langle l_1, l_2 \rangle$ are selected in our experiment such that the angle between l_1 and l_2 falls in the closed range of $[45^\circ, 135^\circ]$ and the angle between l_1/l_2 and positive x -axis falls in $\{0^\circ, 22.5^\circ, 45^\circ, 67.5^\circ, 90^\circ, 112.5^\circ, 135^\circ\}$. Moreover, one combination of two straight lines generates four possible angles which differ in terms of their opening directions. Therefore, we use a total of 80 queries to retrieve triangles from 4,451 photos where each photo may contain many distinct triangles. For each query, we rank the results based on their continuity ratios. Higher continuity ratios represent higher quality of fitting and thus imply more accurate retrieved results.

We recruited 20 participants to this study as before. At an online website, each participant is provided with a subset of 15 randomly selected queries. For each query, we show the participant the top-20 triangles retrieved by our system. For each retrieved triangle, we ask the participant to assess whether it is “useful”, *i.e.*, whether it indicates interesting pose or composition in the image, and help the participant understand the use of triangles in the photo.

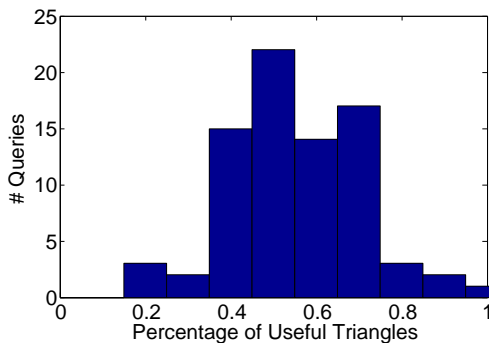


Fig. 13. The performance of the retrieval system for portraits.

Fig. 13 shows the histogram of overall percentages of “useful” triangles for all 80 queries. For most queries, between 40% and 80% of the retrieved results are considered by the participants as providing useful information and guidance on the portrait composition. Overall, 53.8% of the retrieved triangles are regarded as “useful” by the users.

In Fig. 14, we provide examples of both “useful” and “useless” triangles retrieved by our system. The first column contains twelve queries among which eight return high percentages of “useful” triangles and four return low percentage of “useful” triangles. From the retrieved examples, it can be seen that professional photographers are skillful at using different objects, such as arm, leg, shoulder, chair, wall, ground, hair, apparel, or even shadow, to construct triangles. Interestingly, very often a slight adjustment of a pose can form attractive triangles which in turn make the entire composition aesthetically appealing. For instance, the subject in row 3, column 3 slightly turns her head towards left to perfectly align with her shoulder. For the same reason, the subject in row 1, column 4 lifts up her head a little bit.

Our system also retrieves many “useless” triangles for some queries. As shown in the last four rows of Fig. 14, many of these triangles are right triangles. In fact, it is known

that professional photographers often avoid 90° body angles because they often look unnatural and strained [10]. This may partly explain why we are unable to retrieve more “useful” triangles in these cases.

VI. CONCLUSIONS AND FUTURE WORK

We have presented our work that detects and analyzes the usage of the triangle technique in portrait photography. Specifically, we extract a set of candidate line segments from a photo and then successfully fit triangles to these segments despite a large proportion of outliers. The fitted results accurately identify the presence of triangles in photographs. Based on the detected triangles, we further propose two photography support tools that incorporate portrait retrieval based on a sample shot of the scene or a user-specified angle sketch.

Several directions can be further explored to help portrait photographers design and analyze composition. The relationship between triangles and the aesthetic quality of composition can be studied. For instance, how do the number, sizes, shapes, and orientations of triangles influence the aesthetics of photos? Answering such questions can help amateur photographers learn more specific photographic techniques. Moreover, our system can be further developed to employ other design principles such as balance, contrast, viewpoint, depth, and illumination. In the future, we plan to explore these factors to provide users with more comprehensive guidance in their photo creation process.

ACKNOWLEDGMENT

Chuck S. Fong of the Studio 2 Photography provided ground truth annotations for the professional portrait dataset used in our evaluation. Jia Li of the Department of Statistics at Penn State provided valuable suggestions. Edward Chen and Sahil Mishra assisted in developing the data-collection system for the human subject studies. The authors thank Mingli Song for kindly providing the source code of Where2Stand, and the participants in the studies for their assistance. The authors would also like to acknowledge the comments and suggestions from the reviewers and the associate editor.

REFERENCES

- [1] L. Yao, P. Suryanarayan, M. Qiao, J. Z. Wang, and J. Li, “Oscar: On-site composition and aesthetics feedback through exemplars for photographers,” *International Journal of Computer Vision*, vol. 96, no. 3, pp. 353–383, 2012.
- [2] H. Su, T. Chen, C. Kao, W. H. Hsu, and S. Chien, “Preference-aware view recommendation system for scenic photos based on bag-of-aesthetics-preserving features,” *IEEE Trans. Multimedia*, vol. 14, no. 3-2, pp. 833–843, 2012.
- [3] B. Ni, M. Xu, B. Cheng, M. Wang, S. Yan, and Q. Tian, “Learning to photograph: A compositional perspective,” *IEEE Transactions on Multimedia*, vol. 15, no. 5, pp. 1138–1151, 2013. [Online]. Available: <http://dx.doi.org/10.1109/TMM.2013.2241042>
- [4] W. Yin, T. Mei, C. W. Chen, and S. Li, “Socialized mobile photography: Learning to photograph with social context via mobile devices,” *IEEE Trans. Multimedia*, vol. 16, no. 1, pp. 184–200, 2014.
- [5] P. Obrador, R. de Oliveira, and N. Oliver, “Supporting personal photo storytelling for social albums,” in *Proceedings of the ACM International Conference on Multimedia*, 2010, pp. 561–570.
- [6] P. Obrador, L. Schmidt-Hackenberg, and N. Oliver, “The role of image composition in image aesthetics,” in *Proceedings of the IEEE International Conference on Image Processing*, 2010, pp. 3185–3188.





Fig. 14. Examples of retrieved portraits. Each row shows results from a query. Red examples are considered “useless” in our evaluation.

- [7] L. Liu, R. Chen, L. Wolf, and D. Cohen-Or, "Optimizing photo composition," *Computer Graphics Forum*, vol. 29, no. 2, pp. 469–478, 2010.
- [8] S. Bhattacharya, R. Sukthankar, and M. Shah, "A framework for photo-quality assessment and enhancement based on visual aesthetics," in *Proceedings of the ACM International Conference on Multimedia*, 2010, pp. 271–280.
- [9] D. A. Lauer and S. Pentak, *Design Basics*. Cengage Learning, 2011.
- [10] R. Valenzuela, *Picture Perfect Practice: A Self-Training Guide to Mastering the Challenges of Taking World-Class Photographs*. New Riders, 2012.
- [11] M. Freeman, *The Photographer's Eye: Composition and Design for Better Digital Photos*. Focal Press, 2007.
- [12] J. Smith, *Posing for Portrait Photography: A Head-To-Toe Guide for Digital Photographers*, 2nd ed. Amherst Media, 2011.
- [13] B. Hurter, *Group Portrait Photography Handbook*, 2nd ed. Amherst Media, 2005.
- [14] R. Orendovici and J. Z. Wang, "Training data collection system for a learning-based photographic aesthetic quality inference engine," in *Proceedings of the ACM International Conference on Multimedia*, 2010, pp. 1575–1578.
- [15] N. Murray, L. Marchesotti, and F. Perronnin, "AVA: A large-scale database for aesthetic visual analysis," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 2408–2415.
- [16] W. Luo, X. Wang, and X. Tang, "Content-based photo quality assessment," in *Proceedings of IEEE International Conference on Computer Vision*, 2011, pp. 2206–2213.
- [17] R. Datta, D. Joshi, J. Li, and J. Z. Wang, "Studying aesthetics in photographic images using a computational approach," in *Proceedings of the European Conference on Computer Vision-Volume Part III*, 2006, pp. 288–301.
- [18] M. Redi and B. Merialdo, "Enhancing semantic features with compositional analysis for scene recognition," in *Proceedings of the International Conference on Computer Vision-Volume Part III*, 2012, pp. 446–455.
- [19] Y. Luo and X. Tang, "Photo and video quality evaluation: Focusing on the subject," in *Proceedings of the European Conference on Computer Vision: Part III*, 2008, pp. 386–399.
- [20] S. Dhar, V. Ordonez, and T. L. Berg, "High level describable attributes for predicting aesthetics and interestingness," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 1657–1664.
- [21] H.-H. Su, T.-W. Chen, C.-C. Kao, W. H. Hsu, and S.-Y. Chien, "Scenic photo quality assessment with bag of aesthetics-preserving features," in *Proceedings of the ACM International Conference on Multimedia*, 2011, pp. 1213–1216.
- [22] F. Zhang, M. Wang, and S. Hu, "Aesthetic image enhancement by dependence-aware object recomposition," *IEEE Trans. Multimedia*, vol. 15, no. 7, pp. 1480–1490, 2013.
- [23] C. Fang, Z. Lin, R. Mech, and X. Shen, "Automatic image cropping using visual composition, boundary simplicity and content preservation models," in *Proceedings of the ACM International Conference on Multimedia*, 2014, pp. 1105–1108.
- [24] J. Yan, S. Lin, S. B. Kang, and X. Tang, "Change-based image cropping with exclusion and compositional features," *International Journal of Computer Vision*, vol. 114, no. 1, pp. 74–87, 2015.
- [25] P. Xu, H. Yao, R. Ji, X. Liu, and X. Sun, "Where should I stand? learning based human position recommendation for mobile photographing," *Multimedia Tools Appl.*, vol. 69, no. 1, pp. 3–29, 2014.
- [26] Y. S. Rawat and M. S. Kankanhalli, "Context-aware photography learning for smart mobile devices," *TOMCCAP*, vol. 12, no. 1s, p. 19, 2015.
- [27] X. Jin, M. Zhao, X. Chen, Q. Zhao, and S.-C. Zhu, "Learning artistic lighting template from portrait photographs," in *Proceedings of the European Conference on Computer vision: Part IV*, 2010, pp. 101–114.
- [28] Y. Zhang, X. Sun, H. Yao, L. Qin, and Q. Huang, "Aesthetic composition representation for portrait photographing recommendation," in *Proceedings of the IEEE International Conference on Image Processing*, 2012, pp. 2753–2756.
- [29] S. Ma, Y. Fan, and C. W. Chen, "Pose maker: A pose recommendation system for person in the landscape photographing," in *Proceedings of the ACM International Conference on Multimedia*, 2014, pp. 1053–1056.
- [30] R. G. von Gioi, J. Jakubowicz, J. Morel, and G. Randall, "LSD: A fast line segment detector with a false detection control," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 4, pp. 722–732, 2010.
- [31] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, "Contour detection and hierarchical image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 5, pp. 898–916, 2011.
- [32] P. A. Arbeláez, J. Pont-Tuset, J. T. Barron, F. Marqués, and J. Malik, "Multiscale combinatorial grouping," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 328–335.
- [33] Y. Wang, M. Song, D. Tao, Y. Rui, J. Bu, A. C. Tsoi, S. Zhuo, and P. Tan, "Where2stand: A human position recommendation system for souvenir photography," *ACM Trans. Intelligent Systems and Technology*, vol. 7, no. 1, pp. 9:1–9:22, 2015.
- [34] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "CNN features off-the-shelf: An astounding baseline for recognition," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR Workshops*, 2014, pp. 512–519.
- [35] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional nets," in *British Machine Vision Conference*, 2014.
- [36] J. Redmon and A. Farhadi, "YOLO9000: better, faster, stronger," *CoRR*, vol. abs/1612.08242, 2016.



Siqiong He received the bachelor's degree and the ME degree in Computer Science from Zhejiang University, China, in 2008 and 2011, respectively. She obtained the MS degree in Information Sciences and Technology from The Pennsylvania State University, University Park, in 2015. She has been with Google since 2015.



Zihan Zhou received the bachelor's degree in Automation from Tsinghua University, Beijing, China, in 2007, and the master's and PhD degrees in Electrical and Computer Engineering from University of Illinois at Urbana-Champaign in 2010 and 2013, respectively. He is currently a faculty member in the College of Information Sciences and Technology at The Pennsylvania State University. His research interests lie in computer vision, machine learning, and signal processing, with a focus on developing new computational tools for efficient and robust discovery of low-dimensional data structure for 3D scene analysis and modeling. He is a member of the IEEE.



Farshid Farhat is a PhD candidate in the School of Electrical Engineering and Computer Science at The Pennsylvania State University. He received the BS and the MS degrees in Electrical Engineering from Sharif University of Technology. His current research interests include computer vision, image processing, distributed systems, and networking. He is working on image composition and aesthetics analysis on different platforms from smartphones to cluster computers.



James Z. Wang is a Professor of Information Sciences and Technology at The Pennsylvania State University. He received the bachelor's degree in mathematics and computer science *summa cum laude* from the University of Minnesota, and the MS degree in mathematics, the MS degree in computer science and the PhD degree in medical information sciences, all from Stanford University. His research interests include computational aesthetics and emotions, automatic image tagging, image retrieval, and computerized analysis of paintings. He was a visiting professor at the Robotics Institute at Carnegie Mellon University (2007–2008), a lead special section guest editor of the IEEE Transactions on Pattern Analysis and Machine Intelligence (2008), and a program manager at the Office of the Director of the National Science Foundation (2011–2012). He was a recipient of a National Science Foundation Career award (2004).