

A Multi-level Sketch-based Interface for Decorative Pattern Exploration

Yilan Chen Hongbo Fu Kin Chung Au
School of Creative Media, City University of Hong Kong




Figure 1: Our system allows users to sketch in multiple levels (Left) for decorative pattern exploration (Right: top-7 retrieved results).

Abstract

Despite the extensive usage of decorative patterns in art and design, there is a lack of intuitive ways to find a certain type of patterns. In this paper, we present a multi-level sketch-based interface that incorporates low-level geometrical features and high-level structural features, namely reflection, rotation, and translation symmetries, to support decorative pattern exploration at different levels of detail. Four brush tools are designed for users to specify any combination of such features and compose a hybrid search query. The results of a pilot study show that users are able to perform pattern retrieval tasks using our system easily and effectively.

Keywords: sketch-based image retrieval; symmetry; decorative patterns

Concepts: •Human-centered computing → Interactive systems and tools;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2016 ACM. SA '16 Technical Briefs, December 05-08, 2016, Macao ISBN: 978-1-4503-4541-5/16/12 DOI: <http://dx.doi.org/10.1145/3005358.3005362>

1 Introduction

Decorative patterns are widely used in design and art. There are rich pattern resources online, but their retrieval is still mainly limited to text-based search, which cannot intuitively represent the content, i.e., visual appearance.

Regarding the content of decorative patterns, both the structure and the detailed design units are important to the overall visual effect. There have been extensive theoretical studies [Washburn and Crowe 1988] as well as computational classification methods [Liu et al. 2004; Han and McKenna 2009] on the structure of patterns. However, the motif details are totally omitted in such categorizations, and some pattern groups are less distinctive for human perception. In contrast, sketch-based image retrieval (SBIR) [Eitz et al. 2009; Cao et al. 2013] is an intuitive way to describe the shape details. However, the structures cannot be directly specified. Detecting symmetries from vague sketches is hard, yet a near-complete sketch of complex patterns will require high drawing skills.

In this paper, we propose a multi-level sketch-based interface (Figure 1) to flexibly merge the representations of structures and details for decorative pattern exploration. Instead of the complex definition of structures [Liu et al. 2004], we utilize the basic symmetries, namely the reflection, rotation and translation symmetries, to describe a structure, since they are more readily understood for normal users. The overall distribution of a repeated pattern is depicted by translation symmetry, and the *motifs*, i.e., the units of patterns, usually contain reflection and rotation symmetries. These structural features can be specified by simple line drawings and seamlessly

integrated with the sketches of shape details to represent a hybrid search query for patterns. Besides, users can specify the desired features in any order. The flexibility of sketching and feature integration thus allows category-based exploration as well as targeted searching. A pilot study on our prototype system confirms its usability and effectiveness.

2 Related Work

Pattern Analysis. The study of pattern analysis has been focused on classifying the structure of patterned design using the geometric principles of symmetry. Washburn and Crowe [1988] have a thorough review of the theoretical analysis. Briefly speaking, pattern structures can be composed by symmetries of rotation, reflection, and translation. Based on the composition, plane patterns can be further divided into seventeen groups. Liu et al. [2004] developed a method upon the group theory and some symmetry detection methods to classify patterns automatically. However, it is less suitable for pattern searching, since some groups were found to be less discriminating to human perception [Clarke et al. 2011]. Han and McKenna [2009] used only the translation symmetries to classify and retrieve wallpapers. Some other methods for reflection, rotation and translation symmetry detection [Loy and Eklundh 2006; Park et al. 2009; Liu et al. 2013] also exist. Although we rely on these methods to extract the symmetries of pattern images, we aim to provide an easy-to-use UI which support various ways of exploration, from loosely constrained browsing to targeted retrieval.

Sketch-based Image Retrieval. Due to its simplicity and intuitiveness, many works have been done on applying a sketching interface to image retrieval. The general approach is to take a sketch of object contours as input, and represent both sketch and images as high-dimensional descriptors for matching. Global descriptors like Tensor [Eitz et al. 2009] divide images into many cells and encode each cell by a feature vector. Local descriptors like SIFT [Lowe 2004] encode essential local features and can be gathered in a bag-of-words representation. However, since these descriptors all rely on low-level geometrical features like edges, they cannot properly depict the highly structured pattern images. Some methods also used high-level features. For example, Fonseca et al. [2009] encoded the hierarchical topology of subparts in vector drawings. However, this would require the user to draw desired patterns completely and cannot describe the symmetries concisely either. Cao et al. [2013] and Lee [2013] combined the reflection symmetries of query shapes with edge features for image retrieval. These methods all require precise and complete inputs and do not support loosely-defined exploration. Different from theirs, our system lets users directly specify desired shape structures through a few simple strokes and enables the flexible composition of different features.

3 User Interface

Figure 2 shows the user interface of our prototype. Our system specifically provides four sketching tools, i.e., *Shape brush*, *Reflection brush*, *Rotation brush*, and *Distribution brush*. Each of them is associated with one type of the image features.

- The Shape brush sketches the main curves of a desired motif (e.g., the black strokes in Figure 1 (c) and (d)). Its usage is the same as the freeform sketching in traditional SBIR interfaces;
- The Reflection brush sketches the axes of reflection symmetry (see an example in Figure 1 (b));
- The Rotation brush is used to specify the order of rotational symmetry. Users can either sketch a polyline to represent the rotation angle (Figure 2), or sketch out all the axes;




Figure 2: The user interface of our prototype.

- The Distribution brush sketches one or more lines to represent the periodic distribution of motifs (Figure 1 (a)). Users can sketch a line to specify one translation direction, or sketch a lattice to simultaneously specify the relative intervals.

The strokes drawn with the last three tools are fitted into line segments once completed. When a user wants to specify a certain type of feature, s/he can simply select the corresponding tool and draw on the canvas. The user can also add any kind of sketch freely during the exploration, since there is no constraint on the sketching order. If the user combines shape sketch with any symmetry sketch, a *shadow preview* will be displayed on the canvas to visualize the effect of applying the specified symmetric transformation to the shape sketch (Figure 2). Once a stroke is finished, our system combines all the sketches for retrieval and updates the results on the right-hand pane, arranged from left to right and top to bottom.

Next we describe the workflow through two use cases. In the first case (Figure 1), the user wants to browse patterns that have similar structure but different appearances. S/he firstly sketches a rhombic lattice using the Distribution brush (Figure 1 (a)), and then specifies the desired reflection symmetry (Figure 1 (b)). The user then draws a bar across the center to define the appearance (Figure 1 (c)). S/he later undoes the sketching of shape and tries another shape (Figure 1 (d)). Note how the results are updated as the sketches accumulate or change. In the second case (Figure 2), the user wants to find floral decorative motifs. Therefore, s/he uses the Rotation brush to draw a small angle firstly, which represents near-circular symmetry, and then sketches a petal shape around the rotation center.

4 Methodology

4.1 Database Creation

We collected approximately 1,000 decorative pattern images from the Internet, including roughly 800 periodic patterns and 200 motifs. The translational symmetries of the periodic patterns were represented as one or two smallest linearly independent vectors, detected by the method of Park et al. [2009]. We also included the diagonal vectors with close magnitudes to compensate for the perceptual ambiguity of the translation directions. Each vector v_i was normalized as $\frac{v_i}{\sum_k \|v_k\|}$, where n is the number of translational vectors in the image. Hence, the translation features were invariant to the image sizes. Secondly, we manually extracted the representative motifs of the periodic patterns, while this can also be automated by the symmetry-based motif selection algorithm [Liu et al. 2004]. The reflection and rotation symmetries of all the motifs were automatically detected [Loy and Eklundh 2006] with slight refinement. Since the reflection symmetry axes always intersect at the center of the motif, the reflection symmetry can be well represented by the unit direction vectors of the axes. The center of rotation is also at the center of the motif, so we represent the rotation symmetry with

its order, i.e., the number of times a shape fits into itself in one complete rotation. The shapes of image motifs were encoded with the Tensor descriptor [Eitz et al. 2009], which captures the prominent edges well and is easy to implement, though we might resort to more advanced SBIR techniques for better performance. A similar method is used to encode the sketches of the Shape brush (4.2).

4.2 Feature Representation

Shape Feature. The sketches drawn by the Shape brush depict a part or the main shape of a motif. We use the Tensor descriptor [Eitz et al. 2009] to encode both the sketches and images for matching. This method divides each image into the same number of cells, and then finds a single vector in each cell that maximally aligns with the local image gradients. A key point is to decide the region for rasterization properly so as to preserve the relative position of the shape sketch in the desired motif. Since the region of a motif is determined by both the shape sketch and the symmetry axes, we approximate it by applying the specified reflection or rotation (if any) to the bounding box of the shape sketch.

Reflection Feature. The reflection feature is a set of unit vectors, as specified by the sketched axes of the Reflection brush. We take the sum of minimum distance approach to measure the difference between two sets of vectors. However, simply summing up the bi-directional minimum distances does not work well. For one thing, it may lead to a bias towards images with too excessive reflection symmetries, since they are more likely to contain the specified axes (i.e., having a shorter distance to the sketch query), even though such images also contain unwanted axes. Such unequal-sized matching is undesired. For another, users may sketch only one symmetry group rather than the complete set due to the hierarchy of symmetry as well as the human’s selective visual focus. It is preferred to compensate for this kind of perceptual bias. Therefore, we define the distance metric as $D_{ref} = d(S, I) + d(I, S) \cdot (1 - \alpha\beta)$, where I is the reflection feature of an image, S is that of the sketch and $d(\cdot, \cdot)$ is the distance from one to another. Note that $d(I, S)$ is weighed, where α is a small number and β is a binary number which equals to 0 when the reflection symmetries are hierarchical, so non-hierarchical and unequal-sized matching is penalized.

Rotation Feature. The rotation feature C is represented by the order of rotational symmetry. Given the fitted line segments of the Rotation sketch, our system firstly detects the center of rotation, which is the intersection of lines, and the branches. If there are only two branches, the order of rotational symmetry is computed using the rotational angle between them. Otherwise, the system treats the number of splits as the order. The similarity between two rotation features can usually be well described by their absolute difference $|C_i - C_j|$. If the rotation symmetries are hierarchical (e.g., four-fold symmetry is included in two-fold symmetry), we take it as a special case and calculate the distance as $1 - \frac{\min(C_i, C_j)}{\max(C_i, C_j)}$ to achieve a smaller, yet distinguishable value.

Translation Feature. For the translation feature, we first merge the nearly parallel sketches into one group, which denote a translational direction. The magnitudes of translation can be derived if each group contains more than one line (i.e., a lattice is drawn), otherwise they equals to 1. Similar to the pre-processing of database (4.1), we normalize the vector magnitudes so that they are invariant to the sketch size. Therefore, the similarity between two translation features can be represented as the angular distance between the vectors, plus the difference of their magnitudes. We again use the sum of minimum distance to compute the angular distance.

4.3 Joint Retrieval

The distance terms of the above features are normalized to the same range and summed up with coefficients, which is simple but still reaches plausible results. An important consideration in designing the coefficients is that images are not treated fairly in multi-level retrieval. Since some decorative patterns contain more symmetry features than others, they are more likely to show up during the retrieval. For example, since a composition of reflections over intersecting lines is equivalent to a rotation, it is possible that most of the results are also reflective when only rotation symmetry is specified. Although the outcome is reasonable, it is dominated by the images with more symmetry features and thus is less diverse. Since diversity is essential for exploration, we hope to distinguish the difference between the well-matched images.

Therefore, we associate the coefficient of an image feature with its discriminating power. If an image contains only one kind of symmetry feature, this feature is considered to be representative of the image structure. Comparatively, a similar feature vector is less discriminative when the image can also be described by other kinds of symmetry features. We thus define the weight as $e^{-\frac{1}{n}}$, where n is the number of existing symmetry features in an image. Note that the coefficients are only applied to the top matched images for better discrimination, so those with rarer symmetries among them will be ranked higher.

5 Evaluation

A pilot study was conducted to evaluate the usability and effectiveness of our system. There were two tasks: targeted retrieval, where users were asked to search for a given image or similar images, and free exploration, where users found patterns for a specific context without restriction. The usage of tools was not limited, while we encouraged users to try them freely. After the tasks, users filled a questionnaire, which included the standard System Usability Scale (SUS) [Brooke et al. 1996] and open-ended questions. We recruited 10 paid participants (P1-P10), including a professional textile designer (P1), five Design students who had experience in creating textiles (P2-P6), and four other students who had used patterns for decoration before (P7-P10).

The participants were positive about our system overall. The final SUS score was 76.5, and above the industrial average (68.0) [Sauro 2011]. P1 commented, “(it) may be a good supplement to current text-based search, since textile design usually requires uniformity”. Regarding the tools, 90% of the participants agreed that the symmetry brushes were useful and also indicated different preferences. For example, P6 liked the Distribution and Rotation brushes, because “they make the sketching process even simpler - only part of the pattern is drawn and you can find what you want”. Besides, 90% of the participants agreed that the shadow preview was useful. P4 said the preview “helps me understand the symmetry relations”.

Regarding the performance of retrieval, we sampled a few trials of the tasks here (Figure 3). Figures 4 and 5 shows some representative examples of the targeted task (Trial 1 and 2). Although users interpret the same image differently (e.g., Figure 4(a) vs. (b)), our system was able to find the target image successfully. In the failure case, Figure 4(c), the user tended to represent the given image with 4 lines of reflection symmetry and a few petal shapes. However, the actual number of axes is higher and our method rigidly found the images whose symmetry and shape are close to the user specification. Since users may have different emphasis on the specified features, like focusing more on a certain shape than an exact symmetry in this case, it may be better to provide users with a feature for interactive weight adjustment. Another failure case is Figure 5(c),




Figure 3: Sampled trials of the pilot study. (1) Find the exact reference pattern. (2) Find patterns similar to the reference. (3) Find a wallpaper to match the furniture. (4) Find a pattern to customize this grid leather case.

where the motif extracted in the pre-processing step (4.1) was not compatible with the user’s expectation and thus our method failed to find the desired images. A motif was selected so as to maximize its local symmetry, but there may be other reasonable definition of motifs, especially for hexagonal patterns. A possible solution is to extract all the symmetric motifs using Liu et al.’s method [2004], though it may require more storage. We leave this for future works.

Figure 6 presents some examples of the exploratory task. The results were not evaluated since they highly depend on personal preference. Mixing the usage of different brushes seems to be able to retrieve diverse results with just a few strokes, and P8 commented that “sometimes the unexpected ones give me lots of surprise”. Overall, users were able to find desired decorative patterns easily and effectively using our system.

Acknowledgements

We thank Pixabay and the following Flickr users for making their images available through creative common licenses: Andreas Köberle (Fig. 5(b3)), cc kk (Fig. 1(a5), Fig. 5(c5)), Crystal A Murray (Fig. 5(a1), (a4)), James Mallos (Fig. 5(b1)), Kelbv (Fig. 5(a5)), Lucy Nieto (Fig. 4(b1), (b5)), MCAD Library (Fig. 6(a4)), Mel McCarthy (Fig. 1(d1), (d2), Fig. 4(a1)), Monceau (Fig. 1(c1)). The work described in this paper was partially supported by grants from the Research Grants Council of the Hong Kong Special Administra-




Figure 4: The top-5 retrieved results of Trial 1.



Figure 5: The top-5 retrieved results of Trial 2, with the successful queries highlighted.




Figure 6: Some results of Trail 3 (a) and Trial 4 (b). Each tuple consists a sketch, the top ranked result and the user-selected result from left to right.

tive Region, China (Project No. CityU113513, CityU11204014).

References

- BROOKE, J., ET AL. 1996. Sus-a quick and dirty usability scale. *Usability evaluation in industry 189*, 194, 4–7.
- CAO, X., ZHANG, H., LIU, S., GUO, X., AND LIN, L. 2013. Sym-fish: A symmetry-aware flip invariant sketch histogram shape descriptor. In *ICCV*, 313–320.
- CLARKE, A. D., GREEN, P. R., HALLEY, F., AND CHANTLER, M. J. 2011. Similar symmetries: the role of wallpaper groups in perceptual texture similarity. *Symmetry 3*, 2, 246–264.
- EITZ, M., HILDEBRAND, K., BOUBEKEUR, T., AND ALEXA, M. 2009. A descriptor for large scale image retrieval based on sketched feature lines. In *SBM*, 29–36.
- FONSECA, M. J., FERREIRA, A., AND JORGE, J. A. 2009. Sketch-based retrieval of complex drawings using hierarchical topology and geometry. *Computer-Aided Design 41*, 12, 1067–1081.
- HAN, J., AND MCKENNA, S. J. 2009. Classifying and comparing regular textures for retrieval using texel geometry. In *VISAPP (1)*, 347–354.
- LEE, S. 2013. Symmetry-driven shape description for image retrieval. *Image and Vision Computing 31*, 4, 357–363.
- LIU, Y., COLLINS, R. T., AND TSIN, Y. 2004. A computational model for periodic pattern perception based on frieze and wallpaper groups. *IEEE TPAMI 26*, 3, 354–371.
- LIU, J., SLOTA, G., ZHENG, G., WU, Z., PARK, M., LEE, S., RAUSCHERT, I., AND LIU, Y. 2013. Symmetry detection from realworld images competition 2013: Summary and results. In *CVPRW*, IEEE, 200–205.
- LOWE, D. G. 2004. Distinctive image features from scale-invariant keypoints. *IJCV 60*, 2, 91–110.
- LOY, G., AND EKLUNDH, J.-O. 2006. Detecting symmetry and symmetric constellations of features. In *ECCV*. Springer, 508–521.
- PARK, M., BROCKLEHURST, K., COLLINS, R. T., AND LIU, Y. 2009. Deformed lattice detection in real-world images using mean-shift belief propagation. *IEEE TPAMI 31*, 10, 1804–1816.
- SAURO, J. 2011. *A practical guide to the system usability scale: Background, benchmarks & best practices*. Measuring Usability LLC.
- WASHBURN, D. K., AND CROWE, D. W. 1988. *Symmetries of culture: Theory and practice of plane pattern analysis*. University of Washington Press.