A survey of the state-of-the-art in patch-based synthesis

Connelly Barnes[1](✉), Fang-Lue Zhang[2]

Abstract This paper surveys the state-of-the-art of research in patch-based synthesis. Patch-based methods synthesize output images by copying small regions from exemplar imagery. This line of research originated from an area called "texture synthesis", which focused on creating regular or semi-regular textures from small exemplars. However, more recently, much research has focused on synthesis of larger and more diverse imagery, such as photos, photo collections, videos, and light fields. Additionally, recent research has focused on customizing the synthesis process for particular problem domains, such as synthesizing artistic or decorative brushes, synthesis of rich materials, and synthesis for 3D fabrication. This report investigates recent papers that follow these themes, with a particular emphasis on papers published since 2009, when the last survey in this area was published. This survey can serve as a tutorial for readers who are not yet familiar with these topics, as well as provide comparisons between these papers, and highlight some open problems in this area.

Keywords texture; patch; image synthesis; texture synthesis

1 Introduction

Due to the widespread adoption of digital photography and social media, digital images have enormous richness and variety. Photographers frequently have personal photo collections of thousands of images, and cameras can be used to easily capture high-definition video, stereo images, range images, and high-resolution material samples.

This deluge of image data has spurred research into algorithms that automatically remix or modify existing imagery based on high-level user goals.

One successful research thread for manipulating imagery based on user goals is patch-based synthesis. Patch-based synthesis involves a user providing one or more exemplar images to an algorithm, which is then able to automatically synthesize new output images by mixing and matching small compact regions called *patches* or *neighborhoods* from the exemplar images. Patches are frequently of fixed size, e.g., 8×8 pixels.

The area of patch-based synthesis traces its intellectual origins to an area called "texture synthesis", which focused on creating regular or semi-regular textures from small examples. See Fig. 1 for an example of texture synthesis. A comprehensive survey of texture synthesis methods up to the year 2009 is available [1]. Since then, research has focused increasingly on synthesis of larger and more diverse imagery, such as photos, photo collections, videos, and light fields. Additionally, recent research has focused on customizing the synthesis process for particular problem domains, such as synthesizing artistic or decorative brushes, synthesis of rich

1 University of Virginia, Charlottesville, VA 22904, USA. E-mail: connelly@cs.virginia.edu (✉).
2 TNList, Tsinghua University, Beijing 100084, China. E-mail: z.fanglue@gmail.com.

Fig. 1 Texture synthesis. The algorithm is given as inputting a small exemplar consisting of a regular, semi-regular, or stochastic *texture image*. The algorithm then synthesizes a large, seamless output texture based on the input exemplar. Reproduced with permission from Ref. [1], © The Eurographics Association 2009.

materials, and synthesis for 3D fabrication.

In this survey, we cover recent papers that follow these themes, with a particular emphasis on papers published since 2009. This survey also provides a gentle introduction to the state-of-the-art in this area, so that readers unfamiliar with this area can learn about these topics. We additionally provide comparisons between state-of-the-art papers and highlight open problems.

2   Overview

This survey paper is structured as follows. In Section 3, we provide a gentle introduction to patch-based synthesis by reviewing how patch-based synthesis methods work. This introduction includes a discussion of the two main stages of most patch-based synthesis methods: *matching*, which finds suitable patches to copy from exemplars, and *blending*, which composites and blends patches together on the output image. Because the matching stage tends to be inefficient, in Section 4, we next go into greater depth on accelerations to the patch matching stage. For the remainder of the paper, we then investigate different applications of patch-based synthesis. In Sections 5–7, we discuss applications to image inpainting, synthesis of whole images, image collections, and video. In Section 8, we then investigate synthesis algorithms that are tailored towards specialized problem domains, including mimicking artistic style, synthesis of artistic brushes and decorative patterns, synthesis for 3D and 3D fabrication, fluid synthesis, and synthesis of rich materials. Finally, we wrap up with discussion and a possible area for future work in Section 9.

3   Introduction to patch-based synthesis

There are two main approaches for example-based synthesis: pixel-based methods, which synthesize by copying one pixel at a time from an exemplar to an output image, and patch-based methods, which synthesize by copying entire patches from the exemplar. Pixel-based methods are discussed in detail by Wei et al. [1]. Patch-based methods have been in widespread use recently, so we focus on them.

We now provide a review of how patch-based synthesis methods work. In Fig. 2, an illustration is shown of the patch-based method of image
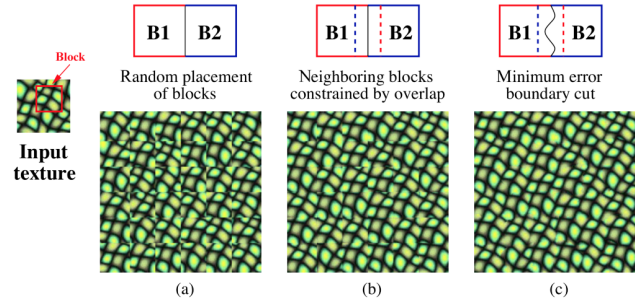

Fig. 2   The *image quilting* patch-based synthesis method [2]. Left: an input exemplar. Right: output imagery synthesis. (a) and (b) show the *matching* stage, where patches are selected by either (a) random sampling or (b) based on similarity to previously selected patches. (c) shows the *blending* stage, which composites and blends together overlapping patches. In this method, "blending" is actually done based on a minimal error boundary cut. Reproduced with permission from Ref. [2], © Association for Computing Machinery, Inc. 2001.

quilting [2]. There are two main stages of most patch-based synthesis methods: *matching* and *blending*.

The *matching* stage locates suitable patches to copy from exemplars, by establishing a correspondence between locations in the output image being synthesized, and the input exemplar image. In image quilting [2], this is done by laying down patches in raster scan order, and then selecting out of a number of candidate patches which have the best agreement with already placed patches. This matching process is straightforward for small textures but becomes more challenging for large photographs or photo collections, so we discuss in more depth different matching algorithms in Section 4.

Subsequently, the *blending* stage composites and blends patches together on the output image. See Fig. 3 for an illustration of the different patch blending methods discussed here. For sparse patches


Fig. 3   Blending methods for patch-based synthesis. Left: three methods for blending sparse patches, i.e., patches having a relatively small overlap: (a) irregularly-shaped patches can be composited [3], (b) overlapping patches can be blended in the overlap region [4], or (c) optimal seams can be found to make a hard *cut* between patches. Right: (d) shows a method for blending densely overlapping patches, where a patch is defined around every pixel: an average of all overlapping colors is computed by a voting process [6, 7].

that have a relatively small overlap region, blending can be done by simply compositing irregularly-shaped patches [3], using a blending operation in the overlap region [4], or using dynamic programming or graph cuts to find optimal seams [2, 5] (see Fig. 2(c) and Figs. 3(a)–3(c)). In other papers, dense patches are defined such that there is one patch centered at every pixel. In this case, many patches simultaneously overlap, and thus the blending operation is typically done as a weighted average of many different candidate "votes" for colors in the overlap region [6–8] (see Fig. 3(d)).

Synthesis can further be divided into *greedy* algorithms that synthesize pixels or patches only once, and *iterative optimization* algorithms that use multiple passes to repeatedly improve the texture in the output image. The image quilting [2] method shown in Fig. 2 is a greedy algorithm because it simply places patches in raster order in a single pass. A limitation of greedy algorithms is that if a mistake is made in synthesis, the algorithm cannot later recover from the mistake. In contrast, dense patch synthesis algorithms [6–9] typically repeat the matching and blending stages as an optimization. In the matching stage, patches within the current estimate for the output image are matched against the exemplar to establish potential improvements to the texture. Next, in the blending stage, these texture patches are copied and blended together by the "voting" process to give an improved estimate for the output image at the next iteration. Iterative optimization methods also typically work in a *coarse-to-fine manner*, by running the optimization on an image pyramid [10]: the optimization is repeated until convergence at a coarse resolution, and then this is repeated at successively finer resolutions until the target image resolution is reached.

4   Matching algorithms

In this section, we discuss algorithms for finding good matches between the output image being synthesized and the input exemplar image. As explained in the previous section, patch-based synthesis algorithms have two main stages: *matching* and *blending*. The matching stage locates the best patches to copy from the exemplar image to the output image that is being synthesized.

Matching is generally done by minimizing a distance term between a partially or completely synthesized patch in the output image and a same-shaped region in the exemplar image: a search is used to find the exemplar patch that minimizes this distance. We call the distance a *patch distance* or *neighborhood distance*. For example, in the image quilting [2] example of Fig. 2, this distance is defined by measuring the L^2 norm between corresponding pixel colors in the overlap region between blocks B1 and B2. For iterative optimization methods [6, 7], the patch distance is frequently defined as an L^2 norm between corresponding pixel colors of a $p \times p$ square patch in the output image and the same-sized region in the exemplar image. More generally, the patch distance could potentially operate on any image features (e.g., SIFT features computed densely on each pixel [11]), use any function to measure the error of the match (including functions not satisfying the triangle inequality), and use two or more degrees of freedom for the search over correspondences from the output image patch and the exemplar patch (for example, in addition to searching to find the (x, y) translational coordinates of an exemplar patch to match, rotation angle θ and scale s could also be searched over).

Typically, matching then proceeds by finding nearest neighbors from the synthesized output S to the exemplar E according to the patch distance. In the terminology of Barnes et al. [12], for the case of two (x, y) translational degrees of freedom, we can define a *Nearest Neighbor Field* (NNF) as a function $f : S \mapsto \mathbb{R}^2$ of offsets, defined over all possible patch coordinates (locations of patch centers) in image S, mapping to the center coordinate of the corresponding most similar patch in the exemplar.

Although search in the matching stage can be done in a brute-force manner by exhaustively sampling the parameter space, this tends to be so inefficient as to be impractical for all but the smallest exemplar images. Thus, much research has been devoted to more efficient matching algorithms. Generally, research has focused on approximation algorithms for the matching, because exact algorithms remain slower [13], and the human visual system is not sensitive to small errors in color.

We will first discuss different approximate matching algorithms for patches, followed by a discussion of how these can be generalized to apply

to correspondence finding algorithms in computer vision. In Fig. 4 are shown illustrations of the key components of a number of approximate matching algorithms. We now discuss five matching algorithms.

Matching using coherence. One simple technique to find good correspondences is to take advantage of *coherence* [14]. Typically, the nearest neighbor field (NNF) used for matching is initialized in some manner, such as by random sampling. However, random correspondences are quite poor, that is, they have high approximation error. An illustration of coherence is shown in Fig. 4(a). Suppose that during synthesis, we are examining a current patch in the output image (shown as a solid red square) and it has a poor correspondence in the exemplar (shown as a dotted red square, with the red arrow indicating the correspondence), that is, the correspondence has high patch distance. In this case, a better correspondence might be obtained from an adjacent patch such as the blue one to the left. However, the correspondence from the adjacent left patch (in blue) must be shifted to the right to obtain an appropriate correspondence for the current patch. The new (shifted) candidate correspondence is shown in green. The patch distance is evaluated for this new candidate correspondence, compared with the existing patch distance for the red correspondence, and whichever correspondence has lower patch distance is written back to the NNF. We refer to this process of choosing the best correspondence

as *improving a correspondence*. Matching using coherence is also known as *propagation*, because it has the effect of propagating good correspondences across the image. Typically, one might propagate from a larger set of adjacent patches than just the left one: for example, if one is synthesizing in raster order, one might propagate from the left patch, the above patch, and the above-left patch.

Matching using k-coherence. An effective technique for finding correspondences in small, repetitive textures is k-*coherence* [15]. This combines the previous *coherence* technique with precomputed k-nearest neighbors within the exemplar. The k-coherence method is illustrated in Fig. 4(b). A precomputation is run on the exemplar, which determines for each patch in the exemplar (illustrated by a green square), what are the k most similar patches located elsewhere in the exemplar (illustrated in green, for $k = 2$). Now suppose that during synthesis, similarly as before, we are examining a current patch, which has a poor correspondence, shown in red. We first apply the coherence rule: we look up the adjacent left patch's correspondence within the exemplar (shown in blue), and shift the exemplar patch to the right by one pixel to obtain the coherence candidate (shown in green). The coherence candidate as well as its k nearest neighbor candidates (all shown as green squares) are all considered as candidates to improve the current patch's correspondence. This process can be repeated for other coherence candidates, such as for the above patch.
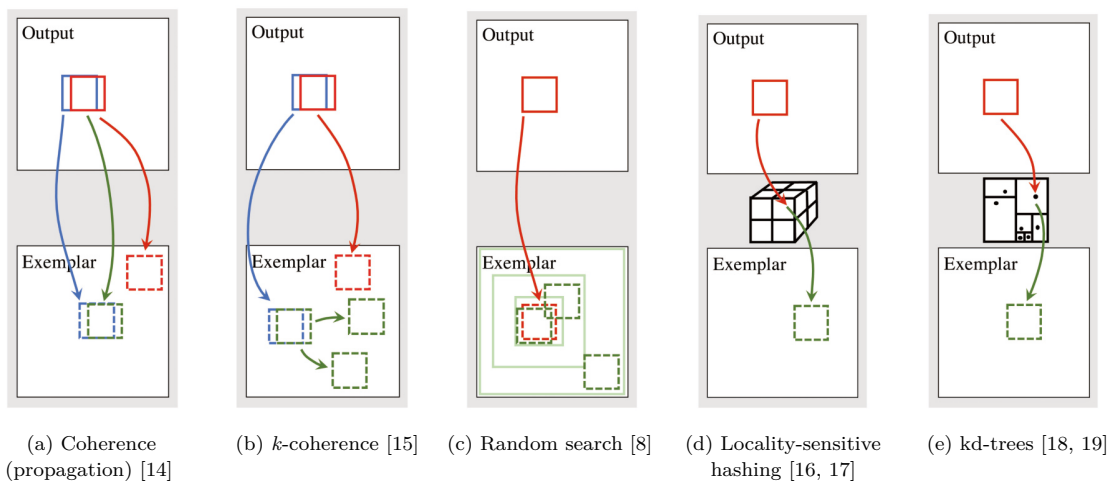

(a) Coherence (propagation) [14]   (b) k-coherence [15]   (c) Random search [8]   (d) Locality-sensitive hashing [16, 17]   (e) kd-trees [18, 19]

Fig. 4 Key components of different approximate matching algorithms. The goal of a matching algorithm is to establish correspondences between patches in the output image and patches in the exemplar. Each algorithm proposes one or more candidate matches, from which the match with minimal patch distance is selected.

Matching using PatchMatch. When real-world photographs are used, k-coherence alone is insufficient to find good patch correspondences, because it assumes that the image is a relatively small and repetitive texture. PatchMatch [8, 12, 20] allows for better global correspondences within real-world photographs. It augments the previous coherence (or propagation) stage with a random search process, which can search for good correspondences across the entire exemplar image, but places most samples in local neighborhoods surrounding the current correspondence. Specifically, the candidate correspondences are sampled uniformly within sampling windows that are centered on the best current correspondence. As each patch is visited, the sampling window initially has the same size as the exemplar image, but it then contracts in width and height by powers of two until the it reaches 1 pixel in size. This random search process is shown in Fig. 4(c). Unlike methods such as locality-sensitive hashing (LSH) or kd-trees, PatchMatch takes less memory, and it is more flexible: it can use an arbitrary patch distance function. The Generalized PatchMatch algorithm [12] can utilize arbitrary degrees of freedom, such as matching over rotations and scales, and can find approximate k-nearest neighbors instead of only a single nearest neighbor. For the specific case of matching patches using 2D translational degrees of freedom only under the L^2 norm, PatchMatch is more efficient than kd-trees when they are used naively [8]; however, it is less efficient than state-of-the-art techniques that combine LSH or kd-trees with coherence [17–19]. Recently, graph-based matching algorithms have also been developed based on PatchMatch, which operate across image collections: we discuss these next in Section 6.

Matching using locality-sensitive hashing. Locality-sensitive hashing (LSH) [21] is a dimension reduction and quantization method that maps from a high-dimensional feature space down to a lower-dimensional quantized space that is suitable to use as a "bucket" in a multidimensional hash table. For example, in the case of patch synthesis, the feature space for a $p \times p$ square patch in RGB color space could be $\mathbb{R}^{3p^2}$, because we could stack the RGB colors from each of the p^2 pixels in the patch into a large vector. The hash bucket space could be some lower-dimensional space such as $\mathbb{N}^6$. The "locality" property of LSH is that similar features map to the same hash table bucket with high probability. This allows one to store and retrieve similar patches by a simple hash-function lookup. One example of a locality-sensitive hashing function is a projection onto a random hyperplane followed by quantization [21]. Two recent works on matching patches using LSH are coherency-sensitive hashing (CSH) [17] and PatchTable [16]. These two works have a similar hashing process and both use coherence to accelerate the search. Here we discuss the hashing process of PatchTable, because it computes matches only from the output image to the exemplar, whereas CSH computes matches from one image to the other and vice versa.

The patch search process for PatchTable [16] is illustrated in Fig. 4(d). First, in a precomputation stage, a multidimensional hash table is created, which maps a hash bucket to an exemplar patch location. The exemplar patches are then inserted into the hash table. Second, as shown in Fig. 4(d), during patch synthesis, for each output patch, we use LSH to map the patch to a hash table cell, which stores the location of an exemplar patch. Thus, for each output patch, we can look up a similar exemplar patch. In practice, this hash lookup operation is done only on a sparse grid for efficiency, and coherence is used to fill in the gaps.

Matching using tree-based techniques. Tree search techniques have long been used in patch synthesis [4, 7, 18, 19, 22]. The basic idea is to first reduce the dimensionality of patches using a technique such as principal components analysis (PCA) [23], and then insert the reduced dimensionality feature vectors into a tree data structure that adaptively divides up the patch appearance space. The matching process is illustrated in Fig. 4(e). Here a kd-tree is shown as a data structure that indexes the patches. A kd-tree is an adaptive space-partitioning data structure that divides up space by axis-aligned hyperplanes to conform to the density of the points that were inserted into the tree [24]. Such kd-tree methods are the state-of-the-art for efficient patch searches for 2D translational matching with L^2 norm patch distance function [18, 19].

We now review the state-of-the-art technique of TreeCANN [19]. TreeCANN works by first inserting all exemplar patches that lie along a sparse grid into a kd-tree. The sparse grid is used to improve the efficiency of the algorithm, because kd-tree operations are not highly efficient, and are memory intensive. Next, during synthesis, output patches that lie along a sparse grid are searched against the kd-tree. Locations between sparse grid pixels are filled using coherence.

Correspondences in computer vision. The PatchMatch [12] algorithm permits arbitrary patch distances with arbitrary degrees of freedom. Many papers in computer vision have therefore adapted PatchMatch to handle challenging correspondence problems such as stereo matching and optical flow. We review a few representative papers here. Bleyer et al. [25] showed that better correspondences can be found between stereo image pairs by adding additional degrees of freedom to patches so they can tilt in 3D out of the camera plane. The PatchMatch filter work [26] showed that edge-aware filters on cost volumes can be used in combination with PatchMatch to solve labeling problems such as optical flow and stereo matching. Results for stereo matching are shown in Fig. 5. Similarly, optical flow [27] with large displacements has been addressed by computing an NNF from PatchMatch, which provides approximate correspondences, and then using robust model fitting and motion segmentation to eliminate outliers. Belief propagation techniques have also been integrated with PatchMatch [28] to regularize the correspondence fields it produces.

5 Images

The patch-based matching algorithms of Section 4 facilitate many applications in image and video manipulation. In this section, we discuss some

Fig. 5 Results from PatchMatch filter [26]. The algorithm accepts a stereo image pair as input, and estimates stereo disparity maps. The resulting images show a rendering from a new 3D viewpoint. Reproduced with permission from Ref. [26], © IEEE 2013.

of these applications. Patch-based methods can be used to inpaint targeted regions or "reshuffle" content in images. They can also be used to edit repeated elements in images. Researchers have also proposed editing and enhancement methods for image collections and videos. These applications incorporate the efficient patch query techniques from Section 4 in order to make running time be practical.

Inpainting and reshuffling. One compelling application of patch-based querying and synthesis methods is image inpainting [29]. Image inpainting removes a foreground region from an image by replacing it with background material found elsewhere in the image or in other images. In Ref. [8], PatchMatch was shown to be useful for interactive high-quality image inpainting, by using an iterative optimization method that works in a coarse-to-fine manner. Specifically, this process works by repeatedly matching partially completed features inside the hole region to better matches in the background region. Various subsequent methods took a similar approach and focused on improving the completion quality for more complex conditions, such as maintaining geometrical information that can preserve continuous structures between the hole to inpaint and existing content [30, 31]. To combine such different strategies into a uniform framework, Bugeau et al. [32] and Arias et al. [33] separately proposed variational systems that can choose different metrics when matching patches to adapt to different kinds of inpainting inputs. Kopf et al. [34] also proposed a method to predict the inpainting quality, which can help to choose the most suitable image inpainting methods for different situations. Most recently, deep learning has been introduced in Ref. [35] to estimate the missing patches and produce a feature descriptor. The patch-based approach has also been used for Internet photo applications like image inpainting for street view generation [36]. Image "reshuffling" is the problem where a user roughly grabs a region of an image and moves it to a new position. The goal is that the computer can synthesize a new image consistent with the user constraints. Reshuffling can be treated as the extension of image inpainting techniques, by initializing the regions to be synthesized by user specified contents [37]. In Barnes et al. [8], the reshuffling results were made

more controllable by adding user constraints to the PatchMatch algorithm.

Editing repeated elements in images. To perform object-level image editing, it is necessary to deform objects, and inpaint occluded parts of the objects and the background. Patch-based methods can be used to address these challenges. One scenario for object-level image manipulation involves repeated elements in textures and natural images. The idea of editing repeated elements was first proposed by the work of RepFinder [38]. In their interactive system, the repeated objects are first detected and extracted. Next, the edited objects are composited on the completed background using a PatchMatch-based inpainting method. One result is shown in Fig. 6. Huang et al. [39] improved the selection part of RepFinder and also demonstrated similar editing applications. In addition to exploring traditional image operations like moving, deleting, and deforming the repeated elements, novel editing tools were also proposed by Huang et al. In the work of ImageAdmixture [40], mixtures between groups of similar elements were created, as shown in Fig. 6. To create natural appearance in the mixed elements' boundary regions, patch-based synthesis is used to generate the appearance using the pixels from boundaries of other elements within the same group.

Denoising. The Generalized PatchMatch work [12] showed that non-local patch searches can be

Fig. 6 Editing repeated elements in an image. Above: RepFinder [38] uses patch-based methods to complete the background of the editing results before compositing different layers of foreground objects. Below: ImageAdmixture [40] uses patch-based synthesis to determine boundary appearance when generating mixtures of objects. Reproduced with permission from Ref. [38], © Association for Computing Machinery, Inc. 2010, and Ref. [40], © IEEE 2012.

integrated into the non-local means method [41] for image denoising to improve noise reduction. Specifically, this process works by finding for each image patch, the k most similar matches both globally across the image and locally within a search region. A weighted average of these patches can be taken to remove noise from the input image patches. Liu and Freeman [42] showed that a similar non-local patch search can be used for video, where optical flow guides the patch search process. The method of Deledalle et al. [43] showed that principal components analysis (PCA) can be combined with patch-based denoising to produce high-quality image denoising results. Finally, Chatterjee and Milanfar [44] showed that a patch-based Wiener filter can be used to reduce noise by exploiting redundancies at the patch level. In the closely related task of image smoothing, researchers have also investigated patch-based methods that use second order feature statistics [45, 46].

6 Image collections

In the scenario where a photographer has multiple images, there are two categories of existing works on utilizing patches as the basic operating units. In the first category, researchers extended patch-based techniques that had previously been applied to single images to multi-view images and image collections as Tong et al. did in Ref. [47]. In the second category, patches are treated as a bridge to build connections between different images. We now discuss the first category of editing tools. For stereo images, patch-based inpainting methods have been used to complete regions that have missing pixel colors caused by dis-occlusions when synthesizing novel perspective viewpoints [48–50]. In the work of Wang et al. [48], depth information is also utilized to aid the patch-based hole filling process for object removal. Morse et al. [51] proposed an extension of PatchMatch to obtain better stereo image inpainting results. Morse et al. completed the depth information first and then added this depth information to PatchMatch's propagation step when finding matches for inpainting both stereo views. Inspired by the commercial development of light field capture devices like PiCam and Lytro, researchers have also developed editing tools for light fields similar to existing 2D image editing

tools [52]. One way to look at light fields is as an image array with many different camera viewpoints. Zhang et al. [53] demonstrated a layered patch-based synthesis system which is designed to manipulate light fields as an image array. This enables users to perform inpainting, and re-arrange and re-layer the content in the light field as shown in Fig. 7. In these methods, patch querying speed is a bottleneck in the performance. Thus, Barnes et al. [16] proposed to use a fast query method to accelerate the matching process across all the image collections. The proposed patch-based applications, such as image stitching using a small album and light-field super-resolution, were reported to be significantly faster.

In the second category, patches are treated as a bridge to build connections between contents from different images. Unlike previous region matching methods which find similar shapes or global features, these methods focus on matching contiguous regions with similar appearance. This allows a dense, non-rigid correspondence to be estimated between related regions. Non-rigid dense correspondence (NRDC) [54] is a representative work. Based on Generalized PatchMatch [12], the NRDC paper proposed a method to find contiguous matching regions between two related images, by checking and merging the good matches between neighbors. NRDC demonstrates good matching results even when there is a large change in colors or lighting between input images. An example is shown in Fig. 8. The approach of NRDC was further improved and applied to matching contents across an image collection in the work of HaCohen et al. [55]. For large image datasets, Gould and Zhang [56] proposed a method to build a matching graph using PatchMatch, and optimize a conditional Markov random field to propagate pixel labels to all images from just a small subset of annotated

Fig. 8 Left, center: two input images with different color themes. Right: regions matched by NRDC [54]. Reproduced with permission from Ref. [54], © Association for Computing Machinery, Inc. 2011.

images. Patches are also used as a representative feature for matching local contiguous regions in PatchNet [57]. In this work, an image region with coherent appearance is summarized by a graph node, associated with a single representative patch, while geometric relationships between different regions are encoded by labelled graph edges giving contextual information. As shown in Fig. 9, the representative patches and the contextual information are combined to find reasonable local regions and objects for the purpose of library-driven editing.

7   Video

Here we discuss how patch-based methods can be extended for applications on videos, by incorporating temporal information into patch-based optimizations. We will start by discussing a relatively easier application to high dynamic range video [58], and then proceed to briefly discuss two works on video inpainting [59, 60], followed by video summarization by means of "video tapestries" [61]. The method of Kalantari et al. [58] reconstructs *high-dynamic range* (HDR) video. A brief definition of HDR imaging is that it extends conventional photography by using computation to achieve greater dynamic range in luminance, typically by using several photographs of the same scene with varying exposure. In Kalantari et al. [58], HDR video is reconstructed by a special video camera that can alternate the exposure at each frame of the video. For example, the exposure of frame 1 could be low, frame 2 could be medium, and frame 3 could be high, and then this pattern could repeat. The goal of the high-dynamic range reconstruction problem then is to reconstruct the missing exposure information: for example, on frame 1 we need to reconstruct the missing medium and high exposure

Fig. 7 Results of reshuffling and re-layering the content in a light field using PlenoPatch [53]. Reproduced with permission from Ref. [53], © IEEE 2016.

Fig. 9   PatchNet [57] finds a graph that connects local regions, where each local region has an internal repeating texture. PatchNet can be used for library-driven editing applications such as image composition. Reproduced with permission from Ref. [57], © Association for Computing Machinery, Inc. 2013.

information. This will allow us to reconstruct a video that has high dynamic range at every frame, and will thus allow us to take high-quality videos that simultaneously include both very dark and bright regions. One solution to this problem is to use optical flow to simply guide the missing information from past and future frames. However, optical flow is not always accurate, so better results are obtained by Kalantari et al. [58] by formulating the problem as a patch-based optimization that fills in missing pixels by minimizing both optical flow like terms and patch similarity terms. This problem is fairly well-constrained, because there is one constraint image at every frame. Results are shown in Fig. 10. Note that Sen et al. [62] also presented a similar patch-based method of HDR reconstruction where the goal is to produce a single output image rather than a video.

The problem of video inpainting, in contrast, is very unconstrained. In video inpainting, the user selects a spacetime region of the video to remove, and then the computer must synthesize an entire volume

Fig. 10   High-dynamic range (HDR) video reconstruction from Kalantari et al. [58]. Top: frames recorded by a special video camera that cyclically changes exposure for each frame (between high, low, and medium exposure). Bottom: HDR reconstruction, with light and dark areas both well-exposed. Reproduced with permission from Ref. [58], © Association for Computing Machinery, Inc. 2013.

of pixels in that region that plausibly removes the target objects. The problem is further complicated because both the camera and foreground objects may move, and introduce parallax, occlusions, disocclusions, and shadows. Granados et al. [59] developed a video inpainting method that aligns other candidate frames to the frame that is to be removed, selects among candidate pixels for the inpainting using a color-consistency term, and then removes intensity differences using gradient-domain fusion. Granados et al. assumed a piecewise planar background region. Newson et al. [60] inpainted video by using a global, patch-based optimization. Specifically, Newson et al. introduced a spatio-temporal extension to PatchMatch to accelerate the search problem, used a multi-resolution texture feature pyramid to improve texture, and estimated background movement using an affine model.

Finally, the video tapestries [61] work shows that patch-based methods can also be applied to producing pleasing summaries of video. Video tapestries are produced by selecting a hierarchical set of keyframes, with one keyframe level per zoom level, so that a user can interactively zoom in to the tapestries. An appealing summary is then produced by compacting or summarizing the resulting layout image so that it is slightly smaller along each dimension: this facilitates the joining of similar regions and the removal of repetitive features.

Deblurring. Here, we discuss the application of patch-based methods to one hard inverse problem in computer vision, deblurring. Recently, patch-based techniques have been used to deblur images that are blurred due to say camera shake. Cho et al. [63] showed video can be deblurred by observing that some frames are sharper than others. The sharp regions can be detected and used to restore blurry regions in nearby frames. This is done by a patch-

based synthesis process that ensures spatial and temporal coherence. Sun et al. [64] showed that blur kernels can be estimated by using a patch prior that is customized towards modeling corner and edge regions. The blur kernel and the deblurred image can both be estimated by an iterative process by imposing this patch prior. Sun et al. [65] later investigated whether the deblurring results could be improved by training on similar images. Sun et al. showed that deblurring results could be improved if patch priors that locally adapt based on region correspondences are used, or multi-scale patch-pyramid priors are used.

8   Synthesis for specialized domains

In this section, we investigate synthesis algorithms that are tailored towards specialized problem domains, including mimicking artistic style, synthesis of artistic brushes, decorative patterns, synthesis for 3D and 3D fabrication, fluid synthesis, and synthesis of rich materials.

Mimicking artistic style. Patch-based synthesis can mimic the style of artists, such as oil paint style, watercolor style, or even abstract style. We discuss an early work in this area, image analogies [66], followed by two recent works, Bénard et al. [67] and StyLit, which perform example-based stylization using patches. The image analogies framework [66] gave early results showing the transfer of oil paint and watercolor styles. This style transfer works by providing an exemplar photograph A, a stylized variant of the exemplar A', and an input photograph B. The image analogies framework then predicts a stylized version B' of the input photograph B. Image parsing was shown by Ref. [68] to improve results for such exemplar-based stylization. The method of Bénard et al. [67] allows artists to paint over certain keyframes in a 3D rendered animation. Other frames are automatically stylized in the target style by using a temporally coherent extension of image analogies to the video cube. The StyLit method [69] uses the "lit sphere" paradigm [70] for transfer of artistic style. In this approach, a well-lit photorealistic sphere is presented to the artist, so that the artist can produce a non-photorealistic examplar of the same sphere. In StyLit [69], the sphere is rendered using global light transport, and the rendered image is decomposed into lighting channels such as direct

diffuse illumination, direct specular illumination, first two bounces of diffuse illumination, and so forth. This is augmented by an iterative patch assignment process, which avoids producing bad matches while also avoiding excessive regularity in the produced texture. See Fig. 11 for example stylized results.

Synthesis of artistic brushes and decorative patterns. Patch-based synthesis methods can easily generate complex textures and structures. For this reason, they have been adopted for artistic design. We first discuss digital brush tools, and then discuss tools for designing decorative patterns.

Research on digital brushes includes methods for painting directly from casually-captured texture exemplars [72–74], RealBrush [75] which synthesizes from carefully captured natural media exemplars, and "autocomplete" of repetitive strokes in paintings [76]. Ritter et al. [73] initially developed a framework that allows artists to paint with textures on a target image by sampling textures directly from an exemplar image. This approach works by adapting typical texture energy functions to specially handle boundaries and layering effects. Painting by feature [72] further advanced the area of texture-based painting. They improved the sampling of patches along boundary curves, where humans are particularly sensitive to misalignments or repetitions of texture, and then filled other areas using patch-based inpainting. Results are shown in Fig. 12. Brushables [74] also improved


Fig. 11 Results from StyLit [69]. A sphere is rendered using global illumination (upper left), and decomposed into different lighting channels, e.g., direct diffuse channel (lower left) using light path expressions [71]. An artist stylizes the sphere in a desired target style (above middle, right), which is automatically transferred to a 3D rendered model using patch-based synthesis (below middle, right). Reproduced with permission from Ref. [69], © Association for Computing Machinery, Inc. 2016, and Ref. [71], © Association for Computing Machinery, Inc. 1990.

Fig. 12 Results from painting by feature [72]. Reproduced with permission from Ref. [72], © Association for Computing Machinery, Inc. 2013.

upon the example-based painting approach by allowing the artist to specify a direction field for the brush, and simultaneously synthesized both the edge and interior regions. RealBrush [75] allows for rich natural media to be captured with a camera, processed with fairly minimal user input, and then used in subsequent digital paintings. Results from RealBrush are shown in Fig. 13. The autocompletion of repetitive painting strokes [76] works by detecting repetitive painting operations, and suggesting an autocompletion to users if repetition is detected. Strokes are represented in a curve representation which is sampled so that neighborhood matching can be performed between collections of samples.

Decorative patterns are frequently used in illustrated manuscripts, formal invitations, web pages, and interior design. We discuss several recent works on synthesis of decorative patterns. DecoBrush [77] allows a designer to synthesize decorative patterns by giving the

Fig. 13 Results from RealBrush [75]. Left: physically painted strokes are captured with a camera to create a digital library of natural media. Right: an artist has created a digital painting. RealBrush synthesizes realistic brush texture in the digital painting by sampling the oil paint and plasticine exemplars in the library. Reproduced with permission from Ref. [75], © Association for Computing Machinery, Inc. 2013.

algorithm examples of the patterns, and specifying a path that the pattern should follow. See Fig. 14 for DecoBrush results, including the pattern exemplars, a path, and a resulting decorative pattern. In cases where the pattern is fairly regular and self-similar, Zhou et al. [78] showed that a simpler dynamic programming technique can be used to synthesize patterns along curves. Later, Zhou et al. [79] showed that decorative patterns can be synthesized with better controls over topology (e.g., number of holes and connected components), by first synthesizing the topology using a topological descriptor, and then synthesizing the pattern itself. This latter work also demonstrated design of 3D patterns, which we discuss next.

3D synthesis. We first discuss synthesis of 3D structures that do not need to be physically fabricated. Early texture synthesis works focused on synthesizing voxels for the purposes of synthesizing geometry on surfaces [80] or synthesizing 3D volumetric textures [81, 82]. For volumetric 3D texture synthesis, Kopf et al. [81] showed that 3D texture could be effectively synthesized from a 2D exemplar by matching 2D neighborhoods aligned along the three principal axes in 3D space. Kopf et al. [81] additionally showed that histogram matching can make the synthesized statistics more similar to the exemplar. Dong et al. [82] showed that such 3D textures can be synthesized in a lazy manner, such that if a surface is to be textured, then only voxels near the surface need be evaluated. A key observation in their paper is to synthesize the 3D volume from precomputed sets of 3D candidates, each of which is a triple of interleaved 2D neighborhoods. This reduces the search space during synthesis. See Fig. 15 for some results. In

Fig. 14 Results from DecoBrush [77]. Reproduced with permission from Ref. [77], © Association for Computing Machinery, Inc. 2014.

Fig. 15 Lazy solid texture synthesis [82] can be used to synthesize 3D texture from a 2D exemplar. This is done in a *lazy* manner, so that only voxels near the surface of a 3D model need be synthesized. Left: exemplar. Right: (a) a result for a 3D mesh, (b) a consistent texture can be generated even if the mesh is fractured. Reproduced with permission from Ref. [82], © The Author(s) Journal compilation 2008 © The Eurographics Association and Blackwell Publishing Ltd. 2008.

a similar manner, Lee et al. [83] showed that arbitrary 3D geometry could be synthesized. In their case, to reduce the computational burden, rather than synthesizing voxels directly, they synthesized an adaptive signed distance field using an octree representation.

A separate line of research focused on the synthesis of discrete vector elements in 3D [84, 85]. In Ma et al. [84], collections of discrete 3D elements are synthesized by matching neighborhoods to a reference exemplar of 3D elements. Multiple samples can be placed along each 3D element, which allows for synthesis of oblong objects such as a bowl of spaghetti or bean vegetables. Constraints can be incorporated into the synthesis process, such as physics, boundary constraints, or orientation fields. Later, Ma et al. [85] extended this synthesis method to handle dynamically animating 3D objects, such as groups of fish or animating noodles.

Synthesis for 3D fabrication. Physical artifacts can be fabricated in a computational manner by using subtractive manufacturing, such as computer-numerical control (CNC) milling and laser cutters, or using additive manufacturing, such as fused deposition modeling (FDM) printing and photopolymerization. During the printing process, it may be desirable to add fine-scale texture detail, such as ornamentation by flowers or swirls on a lampshade. Researchers have explored using patch-based and by-example methods in this direction recently. Dumas et al. [86] showed that a mechanical optimization technique called *topology optimization*

can be combined with an appearance optimization that controls patterns and textures. *Topology optimization* designs 2D or 3D shapes so as to have structural properties such as supporting force loads. Subsequently, Martínez et al. [87] showed that geometric patterns including empty regions can be synthesized along the surfaces of 3D shapes such that the printed shape is structurally sound. The method works by a joint optimization of a patch-based appearance term and a structural soundness term. The synthesized patterns follow the input exemplar, as shown in Fig. 16 (left). Recently, Chen et al. [88] demonstrated that fine filigrees can be fabricated in an example-driven manner. Their method works by reducing the filigree to a skeleton that references the base elements, and they also relax the problem by permitting partial overlap of elements. Example filigrees are shown in Fig. 16 (right).

Fluids. Texture synthesis has been used to texture fluid animations in 2D and 3D. The basic idea is to texture one frame of an animation, then advect the texture forward in time by following the fluid motion, and then re-texture the next frame if needed by starting from this initial guess. Some early works explored these ideas for synthesizing 2D textures directly on 3D fluids [89, 90]. More recently, patch-based fluid synthesis research has focused on stylizing and synthesizing fluid flows that match a target exemplar [91–93]. Ma et al. [91] focused on synthesizing high-resolution motion fields that match an exemplar, while the low-resolution flow matches a simulation or other guidance field. Browning et al. [92] synthesized stylized 2D fluid animations that match an artist's


Fig. 16 Synthesis for 3D fabrication. Left: results from Martínez et al. [87]. Right: filigree from Chen et al. [88]. Subtle patterns of material and empty space are created along the surface of the object to resemble 2D exemplars. Reproduced with permission from Ref. [87], © Association for Computing Machinery, Inc. 2015, and Ref. [88], © Association for Computing Machinery, Inc. 2016.

stylized exemplar. Jamriška et al. [93] implemented a more sophisticated system of transfer from exemplars, including support for video exemplars, encouraging uniform usage of exemplar patches, and improvement of the texture when it is advected for a long time. Figure 17 shows a result for fluid texturing with the method of Jamriška et al. [93].

Rich materials. Real-world materials tend to have complex appearance, including varying amounts of weathering, gradual or abrupt transitions between different materials, varying normals, lighting, orientation, and scale. Texture approaches can be used to factor out or manipulate these rich material properties. We discuss four recent works: the first two focus on controlling material properties [94, 95], the third focuses on controlling weathering [96], and the last focuses on capturing spatially varying BRDFs by taking two photographs of a textured material [95]. Image melding [94] permits an artist to create smooth transitions between two source images in a

way of gradual transitions between inconsistent colors, textures, and structural properties. This permits applications in object cloning, stitching of challenging panoramas, and hole filling using multiple photographs. Diamanti et al. [97] focused on a similar problem of synthesizing images or materials that follow user annotations, such as the desired scale, material, orientation, lighting, and so forth. In cases that the desired annotation is not present, Diamanti et al. [97] performed texture interpolation to synthesize a plausible guess for unobserved material properties. A result from Diamanti et al. [97] is shown in Fig. 18. The recent work by Bellini et al. [96] demonstrates that automatic control over degree of weathering can be achieved for repetitive, textured images. This method finds for each patch of the input image a measure for how weathered it is, by measuring its dissimilarity to other similar patches. Subsequently, weathering can be reduced or increased in a smooth, time-varying manner by replacing low-weathered patches with high-weathered ones, or vice versa. The method of Aittala et al. [95] uses a flash, no-flash pair of photographs of a textured material to recover a spatially varying BRDF representation for the material. This is done by leveraging a self-similarity observation that although a texture's appearance may vary spatially, for any point on the texture, there exist many other points with similar reflectance properties. Aittala et al. fit a spatially varying BRDF that models diffuse and specular, anisotropic reflectance over a detailed normal map.

Fig. 17 Results from Jamriška et al. [93]. Left: an exemplar of a fluid (in this case, fire modeled as a fluid). Right, below: the result of a fluid simulation, used to synthesize a frame of the resulting fluid animation (right, above). Reproduced with permission from Ref. [93], © Association for Computing Machinery, Inc. 2015.

9 Discussion

There have been great advances in patch-based

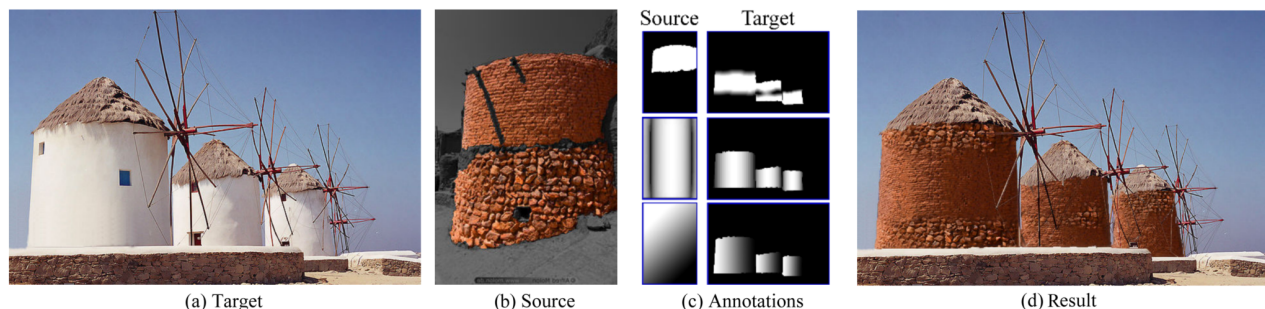(a) Target   (b) Source   (c) Annotations   (d) Result

Fig. 18 Complex material synthesis results from Diamanti et al. [97]. The target image (a) is used as a backdrop and the source (b) is used as an exemplar for the patch-based synthesis process. The artist creates annotations (c) for different properties of the material, such as large versus small bricks, the normal, and lit areas versus shadow. A result image (d) is synthesized. About 50% of the desired patch annotations were not seen in the exemplar, and were instead interpolated between the known materials. Reproduced with permission from Ref. [97], © Association for Computing Machinery, Inc. 2015.

synthesis in the last decade. Techniques that originally synthesized larger textures from small ones have been adapted to many domains, including editing of images and image collections, video, denoising and deblurring, synthesis of artistic brushes, decorative patterns, synthesis for 3D fabrication, and fluid stylization. We now discuss one alternative approach that could potentially inspire future work: deep learning for synthesis of texture.

Deep neural networks have recently been used for synthesis of texture [98–100]. Gatys et al. [98] showed that artistic style transfer can be performed from an exemplar to an arbitrary photograph by means of deep neural networks that are trained on recognition problems for millions of images. Specifically, this approach works by optimizing the image so that feature maps at the higher levels of a neural network have similar pairwise correlations as the exemplar, but the image is still not too dissimilar in its feature map to the input image. This optimization can be carried out by back-propagation. Subsequently, Ulyanov et al. [99] showed that such a texture generation process can be accelerated by pre-training a convolutional network to mimic a given exemplar texture. Denton et al. [100] showed that convolutional neural networks can be used to generate novel images, by training a different convolutional network at each level of a Laplacian pyramid, using a generative adversarial network loss.

Although neural networks are parametric models that are quite different from the non-parametric approach commonly used in patch-based synthesis, it may be interesting in future research to combine the benefits of both approaches. For example, unlike neural networks, non-parametric patch methods tend to use training-free k-nearest neighbor methods, and so they do not need to go through a training process to "learn" a given exemplar. However, neural networks have recently shown state-of-the-art performance on many hard inverse problems in computer vision such as semantic segmentation. Thus, hybrid approaches might be developed that take advantage of the benefits of both techniques.

Acknowledgements

References

[1] Wei, L.-Y.; Lefebvre, S.; Kwatra, V.; Turk, G. State of the art in example-based texture synthesis. In: Proceedings of Eurographics 2009, State of the Art Report, EG-STAR, 93–117, 2009.

[2] Efros, A. A.; Freeman, W. T. Image quilting for texture synthesis and transfer. In: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, 341–346, 2001.

[3] Praun, E.; Finkelstein, A.; Hoppe, H. Lapped textures. In: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, 465–470, 2000.

[4] Liang, L.; Liu, C.; Xu, Y.-Q.; Guo, B.; Shum, H.-Y. Real-time texture synthesis by patch-based sampling. *ACM Transactions on Graphics* Vol. 20, No. 3, 127–150, 2001.

[5] Kwatra, V.; Schödl, A.; Essa, I.; Turk, G.; Bobick, A. Graphcut textures: Image and video synthesis using graph cuts. *ACM Transactions on Graphics* Vol. 22, No. 3, 277–286, 2003.

[6] Wexler, Y.; Shechtman, E.; Irani, M. Space-time completion of video. *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 29, No. 3, 463–476, 2007.

[7] Simakov, D.; Caspi, Y.; Shechtman, E.; Irani, M. Summarizing visual data using bidirectional similarity. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1–8, 2008.

[8] Barnes, C.; Shechtman, E.; Finkelstein, A.; Goldman, D. B. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics* Vol. 28, No. 3, Article No. 24, 2009.

[9] Kaspar, A.; Neubert, B.; Lischinski, D.; Pauly, M.; Kopf, J. Self tuning texture optimization. *Computer Graphics Forum* Vol. 34, No. 2, 349–359, 2015.

[10] Burt, P. J. Fast filter transform for image processing. *Computer Graphics and Image Processing* Vol. 16, No. 1, 20–51, 1981.

[11] Lowe, D. G. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* Vol. 60, No. 2, 91–110, 2004.

[12] Barnes, C.; Shechtman, E.; Goldman, D. B.; Finkelstein, A. The generalized PatchMatch correspondence algorithm. In: *Computer Vision—ECCV 2010*. Daniilidis, K.; Maragos, P.; Paragios, N. Eds. Springer Berlin Heidelberg, 29–43, 2010.

[13] Xiao, C.; Liu, M.; Yongwei, N.; Dong, Z. Fast exact nearest patch matching for patch-based image editing

and processing. *IEEE Transactions on Visualization and Computer Graphics* Vol. 17, No. 8, 1122–1134, 2011.

[14] Ashikhmin, M. Synthesizing natural textures. In: Proceedings of the 2001 Symposium on Interactive 3D Graphics, 217–226, 2001.

[15] Tong, X.; Zhang, J.; Liu, L.; Wang, X.; Guo, B.; Shum, H.-Y. Synthesis of bidirectional texture functions on arbitrary surfaces. *ACM Transactions on Graphics* Vol. 21, No. 3, 665–672, 2002.

[16] Barnes, C.; Zhang, F.-L.; Lou, L.; Wu, X.; Hu, S.-M. PatchTable: Efficient patch queries for large datasets and applications. *ACM Transactions on Graphics* Vol. 34, No. 4, Article No. 97, 2015.

[17] Korman, S.; Avidan, S. Coherency sensitive hashing. In: Proceedings of the IEEE International Conference on Computer Vision, 1607–1614, 2011.

[18] He, K.; Sun, J. Computing nearest-neighbor fields via propagation-assisted kd-trees. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 111–118, 2012.

[19] Olonetsky, I.; Avidan, S. TreeCANN-k-d tree coherence approximate nearest neighbor algorithm. In: *Computer Vision—ECCV 2012*. Fitzgibbon, A.; Lazebnik, S.; Perona, P.; Sato, Y.; Schmid, C. Eds. Springer Berlin Heidelberg, 602–615, 2012.

[20] Barnes, C.; Goldman, D. B.; Shechtman, E.; Finkelstein, A. The PatchMatch randomized matching algorithm for image manipulation. *Communications of the ACM* Vol. 54, No. 11, 103–110, 2011.

[21] Datar, M.; Immorlica, N.; Indyk, P.; Mirrokni, V. S. Locality-sensitive hashing scheme based on p-stable distributions. In: Proceedings of the 20th Annual Symposium on Computational Geometry, 253–262, 2004.

[22] Wei, L.-Y.; Levoy, M. Fast texture synthesis using tree-structured vector quantization. In: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, 479–488, 2000.

[23] Dunteman, G. H. *Principal Components Analysis*. Sage, 1989.

[24] Muja, M.; Lowe, D. G. Fast approximate nearest neighbors with automatic algorithm configuration. In: Proceedings of VISAPP International Conference on Computer Vision Theory and Applications, 331–340, 2009.

[25] Bleyer, M.; Rhemann, C.; Rother, C. PatchMatch stereo-stereo matching with slanted support windows. In: Proceedings of the British Machine Vision Conference, 14.1–14.11, 2011.

[26] Lu, J.; Yang, H.; Min, D.; Do, M. N. Patch match filter: Efficient edge-aware filtering meets randomized search for fast correspondence field estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1854–1861, 2013.

[27] Chen, Z.; Jin, H.; Lin, Z.; Cohen, S.; Wu, Y. Large

displacement optical flow from nearest neighbor fields. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2443–2450, 2013.

[28] Besse, F.; Rother, C.; Fitzgibbon, A.; Kautz, J. PMBP: PatchMatch belief propagation for correspondence field estimation. *International Journal of Computer Vision* Vol. 110, No. 1, 2–13, 2014.

[29] Guillemot, C.; Le Meur, O. Image inpainting: Overview and recent advances. *IEEE Signal Processing Magazine* Vol. 31, No. 1, 127–144, 2014.

[30] Cao, F.; Gousseau, Y.; Masnou, S.; Pérez, P. Geometrically guided exemplar-based inpainting. *SIAM Journal on Imaging Sciences* Vol. 4, No. 4, 1143–1179, 2011.

[31] Huang, J.-B.; Kang, S. B.; Ahuja, N.; Kopf, J. Image completion using planar structure guidance. *ACM Transactions on Graphics* Vol. 33, No. 4, Article No. 129, 2014.

[32] Bugeau, A.; Bertalmio, M.; Caselles, V.; Sapiro, G. A comprehensive framework for image inpainting. *IEEE Transactions on Image Processing* Vol. 19, No. 10, 2634–2645, 2010.

[33] Arias, P.; Facciolo, G.; Caselles, V.; Sapiro, G. A variational framework for exemplar-based image inpainting. *International Journal of Computer Vision* Vol. 93, No. 3, 319–347, 2011.

[34] Kopf, J.; Kienzle, W.; Drucker, S.; Kang, S. B. Quality prediction for image completion. *ACM Transactions on Graphics* Vol. 31, No. 6, Article No. 131, 2012.

[35] Pathak, D.; Krahenbuhl, P.; Donahue, J.; Darrell, T.; Efros, A. A. Context encoders: Feature learning by inpainting. *arXiv preprint* arXiv:1604.07379, 2016.

[36] Zhu, Z.; Huang, H.-Z.; Tan, Z.-P.; Xu, K.; Hu, S.-M. Faithful completion of images of scenic landmarks using internet images. *IEEE Transactions on Visualization and Computer Graphics* Vol. 22, No. 8, 1945–1958, 2015.

[37] Cho, T. S.; Butman, M.; Avidan, S.; Freeman, W. T. The patch transform and its applications to image editing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1–8, 2008.

[38] Cheng, M.-M.; Zhang, F.-L.; Mitra, N. J.; Huang, X.; Hu, S.-M. RepFinder: Finding approximately repeated scene elements for image editing. *ACM Transactions on Graphics* Vol. 29, No. 4, Article No. 83, 2010.

[39] Huang, H.; Zhang, L.; Zhang, H.-C. RepSnapping: Efficient image cutout for repeated scene elements. *Computer Graphics Forum* Vol. 30, No. 7, 2059–2066, 2011.

[40] Zhang, F.-L.; Cheng, M.-M.; Jia, J.; Hu, S.-M. ImageAdmixture: Putting together dissimilar objects from groups. *IEEE Transactions on Visualization and Computer Graphics* Vol. 18, No. 11, 1849–1857, 2012.

[41] Buades, A.; Coll, B.; Morel, J.-M. A nonlocal algorithm for image denoising. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol. 2, 60–65, 2005.

[42] Liu, C.; Freeman, W. T. A high-quality video denoising algorithm based on reliable motion estimation. In: *Computer Vision—ECCV 2010*. Daniilidis, K.; Maragos, P.; Paragios, N. Eds. Springer Berlin Heidelberg, 706–719, 2010.

[43] Deledalle, C.-A.; Salmon, J.; Dalalyan, A. S. Image denoising with patch based PCA: Local versus global. In: Proceedings of the 22nd British Machine Vision Conference, 25.1–25.10, 2011.

[44] Chatterjee, P.; Milanfar, P. Patch-based near-optimal image denoising. *IEEE Transactions on Image Processing* Vol. 21, No. 4, 1635–1649, 2012.

[45] Karacan, L.; Erdem, E.; Erdem, A. Structure-preserving image smoothing via region covariances. *ACM Transactions on Graphics* Vol. 32, No. 6, Article No. 176, 2013.

[46] Liu, Q.; Zhang, C.; Guo, Q.; Zhou, Y. A nonlocal gradient concentration method for image smoothing. *Computational Visual Media* Vol. 1, No. 3, 197–209, 2015.

[47] Tong, R.-T.; Zhang, Y.; Cheng, K.-L. StereoPasting: Interactive composition in stereoscopic images. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 8, 1375–1385, 2013.

[48] Wang, L.; Jin, H.; Yang, R.; Gong, M. Stereoscopic inpainting: Joint color and depth completion from stereo images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1–8, 2008.

[49] Daribo, I.; Pesquet-Popescu, B. Depth-aided image inpainting for novel view synthesis. In: Proceedings of IEEE International Workshop on Multimedia Signal Processing, 167–170, 2010.

[50] Kim, C.; Hornung, A.; Heinzle, S.; Matusik, W.; Gross, M. Multi-perspective stereoscopy from light fields. *ACM Transactions on Graphics* Vol. 30, No. 6, Article No. 190, 2011.

[51] Morse, B.; Howard, J.; Cohen, S.; Price, B. PatchMatch-based content completion of stereo image pairs. In: Proceedings of the 2nd International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission, 555–562, 2012.

[52] Zhang, L.; Zhang, Y.-H.; Huang, H. Efficient variational light field view synthesis for making stereoscopic 3D images. *Computer Graphics Forum* Vol. 34, No. 7, 183–191, 2015.

[53] Zhang, F.-L.; Wang, J.; Shechtman, E.; Zhou, Z.-Y.; Shi, J.-X.; Hu, S.-M. PlenoPatch: Patch-based plenoptic image manipulation. *IEEE Transactions on Visualization and Computer Graphics* DOI: 10.1109/TVCG.2016.2532329, 2016.

[54] HaCohen, Y.; Shechtman, E.; Goldman, D. B.; Lischinski, D. Non-rigid dense correspondence with applications for image enhancement. *ACM Transactions on Graphics* Vol. 30, No. 4, Article No. 70, 2011.

[55] HaCohen, Y.; Shechtman, E.; Goldman, D. B.; Lischinski, D. Optimizing color consistency in photo collections. *ACM Transactions on Graphics* Vol. 32, No. 4, Article No. 38, 2013.

[56] Gould, S.; Zhang, Y. PatchMatchGraph: Building a graph of dense patch correspondences for label transfer. In: *Computer Vision—ECCV2012*. Fitzgibbon, A.; Lazebnik, S.; Perona, P.; Sato, Y.; Schmid, C. Eds. Springer Berlin Heidelberg, 439–452, 2012.

[57] Hu, S.-M.; Zhang, F.-L.; Wang, M.; Martin, R. R.; Wang, J. PatchNet: A patch-based image representation for interactive library-driven image editing. *ACM Transactions on Graphics* Vol. 32, No. 6, Article No. 196, 2013.

[58] Kalantari, N. K.; Shechtman, E.; Barnes, C.; Darabi, S.; Goldman, D. B.; Sen, P. Patch-based high dynamic range video. *ACM Transactions on Graphics* Vol. 32, No. 6, Article No. 202, 2013.

[59] Granados, M.; Kim, K. I.; Tompkin, J.; Kautz, J.; Theobalt, C. Background inpainting for videos with dynamic objects and a free-moving camera. In: *Computer Vision—ECCV 2012*. Fitzgibbon, A.; Lazebnik, S.; Perona, P.; Sato, Y.; Schmid, C. Eds. Springer Berlin Heidelberg, 682–695, 2012.

[60] Newson, A.; Almansa, A.; Fradet, M.; Gousseau, Y.; Pérez, P. Video inpainting of complex scenes. *SIAM Journal on Imaging Sciences* Vol. 7, No. 4, 1993–2019, 2014.

[61] Barnes, C.; Goldman, D. B.; Shechtman, E.; Finkelstein, A. Video tapestries with continuous temporal zoom. *ACM Transactions on Graphics* Vol. 29, No. 4, Article No. 89, 2010.

[62] Sen, P.; Kalantari, N. K.; Yaesoubi, M.; Darabi, S.; Goldman, D. B.; Shechtman, E. Robust patch-based hdr reconstruction of dynamic scenes. *ACM Transactions on Graphics* Vol. 31, No. 6, Article No. 203, 2012.

[63] Cho, S.; Wang, J.; Lee, S. Video deblurring for hand-held cameras using patch-based synthesis. *ACM Transactions on Graphics* Vol. 31, No. 4, Article No. 64, 2012.

[64] Sun, L.; Cho, S.; Wang, J.; Hays, J. Edge-based blur kernel estimation using patch priors. In: Proceedings of the IEEE International Conference on Computational Photography, 1–8, 2013.

[65] Sun, L.; Cho, S.; Wang, J.; Hays, J. Good image priors for non-blind deconvolution. In: *Computer Vision—ECCV 2014*. Fleet, D.; Pajdla, T.; Schiele, B.; Tuytelaars, T. Eds. Springer International Publishing, 231–246, 2014.

[66] Hertzmann, A.; Jacobs, C. E.; Oliver, N.; Curless, B.; Salesin, D. H. Image analogies. In: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, 327–340, 2001.

[67] Bénard, P.; Cole, F.; Kass, M.; Mordatch, I.; Hegarty, J.; Senn, M. S.; Fleischer, K.; Pesare, D.; Breeden, K. Stylizing animation by example. *ACM Transactions on Graphics* Vol. 32, No. 4, Article No. 119, 2013.

[68] Zeng, K.; Zhao, M.; Xiong, C.; Zhu, S.-C. From image parsing to painterly rendering. *ACM Transactions on Graphics* Vol. 29, No. 1, Article No. 2, 2009.

[69] Fišer, J.; Jamriška, O.; Lukáč, M.; Shechtman, E.; Asente, P.; Lu, J.; Sýkora, D. StyLit: Illumination-guided example-based stylization of 3D renderings. *ACM Transactions on Graphics* Vol. 35, No. 4, Article No. 92, 2016.

[70] Sloan, P.-P. J.; Martin, W.; Gooch, A.; Gooch, B. The lit sphere: A model for capturing NPR shading from art. In: Proceedings of Graphics Interface, 143–150, 2001.

[71] Heckbert, P. S. Adaptive radiosity textures for bidirectional ray tracing. *ACM SIGGRAPH Computer Graphics* Vol. 24, No. 4, 145–154, 1990.

[72] Lukáč, M.; Fišer, J.; Bazin, J.-C.; Jamriška, O.; Sorkine-Hornung, A.; Sýkora, D. Painting by feature: Texture boundaries for example-based image creation. *ACM Transactions on Graphics* Vol. 32, No. 4, Article No. 116, 2013.

[73] Ritter, L.; Li, W.; Curless, B.; Agrawala, M.; Salesin, D. Painting with texture. In: Proceedings of the Eurographics Symposium on Rendering, 371–376, 2006.

[74] Lukáč, M.; Fišer, J.; Asente, P.; Lu, J.; Shechtman, E.; Sýkora, D. Brushables: Example-based edge-aware directional texture painting. *Computer Graphics Forum* Vol. 34, No. 7, 257–267, 2015.

[75] Lu, J.; Barnes, C.; DiVerdi, S.; Finkelstein, A. RealBrush: Painting with examples of physical media. *ACM Transactions on Graphics* Vol. 32, No. 4, Article No. 117, 2013.

[76] Xing, J.; Chen, H.-T.; Wei, L.-Y. Autocomplete painting repetitions. *ACM Transactions on Graphics* Vol. 33, No. 6, Article No. 172, 2014.

[77] Lu, J.; Barnes, C.; Wan, C.; Asente, P.; Mech, R.; Finkelstein, A. DecoBrush: Drawing structured decorative patterns by example. *ACM Transactions on Graphics* Vol. 33, No. 4, Article No. 90, 2014.

[78] Zhou, S.; Lasram, A.; Lefebvre, S. By-example synthesis of curvilinear structured patterns. *Computer Graphics Forum* Vol. 32, No. 2pt3, 355–360, 2013.

[79] Zhou, S.; Jiang, C.; Lefebvre, S. Topology-constrained synthesis of vector patterns. *ACM Transactions on Graphics* Vol. 33, No. 6, Article No. 215, 2014.

[80] Bhat, P.; Ingram, S.; Turk, G. Geometric texture synthesis by example. In: Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing, 41–44, 2004.

[81] Kopf, J.; Fu, C.-W.; Cohen-Or, D.; Deussen, O.; Lischinski, D.; Wong, T.-T. Solid texture synthesis from 2D exemplars. *ACM Transactions on Graphics* Vol. 26, No. 3, Article No. 2, 2007.

[82] Dong, Y.; Lefebvre, S.; Tong, X.; Drettakis, G. Lazy solid texture synthesis. *Computer Graphics Forum* Vol. 27, No. 4, 1165–1174, 2008.

[83] Lee, S.-H.; Park, T.; Kim, J.-H.; Kim, C.-H. Adaptive synthesis of distance fields. *IEEE Transactions on Visualization and Computer Graphics* Vol. 18, No. 7, 1135–1145, 2012.

[84] Ma, C.; Wei, L.-Y.; Tong, X. Discrete element textures. *ACM Transactions on Graphics* Vol. 30, No. 4, Article No. 62, 2011.

[85] Ma, C.; Wei, L.-Y.; Lefebvre, S.; Tong, X. Dynamic element textures. *ACM Transactions on Graphics* Vol. 32, No. 4, Article No. 90, 2013.

[86] Dumas, J.; Lu, A.; Lefebvre, S.; Wu, J.; München, T. U.; Dick, C.; München, T. U. By-example synthesis of structurally sound patterns. *ACM Transactions on Graphics* Vol. 34, No. 4, Article No. 137, 2015.

[87] Martínez, J.; Dumas, J.; Lefebvre, S.; Wei, L.-Y. Structure and appearance optimization for controllable shape design. *ACM Transactions on Graphics* Vol. 34, No. 6, Article No. 229, 2015.

[88] Chen, W.; Zhang, X.; Xin, S.; Xia, Y.; Lefebvre, S.; Wang, W. Synthesis of filigrees for digital fabrication. *ACM Transactions on Graphics* Vol. 35, No. 4, Article No. 98, 2016.

[89] Kwatra, V.; Adalsteinsson, D.; Kim, T.; Kwatra, N.; Carlson, M.; Lin, M. Texturing fluids. *IEEE Transactions on Visualization and Computer Graphics* Vol. 13, No. 5, 939–952, 2007.

[90] Bargteil, A. W.; Sin, F.; Michaels, J. E.; Goktekin, T. G.; O'Brien, J. F. A texture synthesis method for liquid animations. In: Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, 345–351, 2006.

[91] Ma, C.; Wei, L.-Y.; Guo, B.; Zhou, K. Motion field texture synthesis. *ACM Transactions on Graphics* Vol. 28, No. 5, Article No. 110, 2009.

[92] Browning, M.; Barnes, C.; Ritter, S.; Finkelstein, A. Stylized keyframe animation of fluid simulations. In: Proceedings of the Workshop on Non-Photorealistic Animation and Rendering, 63–70, 2014.

[93] Jamriška, O.; Fišer, J.; Asente, P.; Lu, J.; Shechtman, E.; Sýkora, D. LazyFluids: Appearance transfer for fluid animations. *ACM Transactions on Graphics* Vol. 34, No. 4, Article No. 92, 2015.

[94] Darabi, S.; Shechtman, E.; Barnes, C.; Goldman, D. B.; Sen, P. Image melding: Combining inconsistent images using patch-based synthesis. *ACM Transactions on Graphics* Vol. 31, No. 4, Article No. 82, 2012.

[95] Aittala, M.; Weyrich, T.; Lehtinen, J. Two-shot SVBRDF capture for stationary materials. *ACM Transactions on Graphics* Vol. 34, No. 4, Article No. 110, 2015.

[96] Bellini, R.; Kleiman, Y.; Cohen-Or, D. Time-varying weathering in texture space. *ACM Transactions on Graphics* Vol. 34, No. 4, Article No. 141, 2016.

[97] Diamanti, O.; Barnes, C.; Paris, S.; Shechtman, E.; Sorkine-Hornung, O. Synthesis of complex image appearance from limited exemplars. *ACM Transactions on Graphics* Vol. 34, No. 2, Article No. 22, 2015.

[98] Gatys, L. A.; Ecker, A. S.; Bethge, M. A neural algorithm of artistic style. *arXiv preprint* arXiv:1508.06576, 2015.

[99] Ulyanov, D.; Lebedev, V.; Vedaldi, A.; Lempitsky, V. Texture networks: Feed-forward synthesis of textures and stylized images. *arXiv preprint* arXiv:1603.03417, 2016.

[100] Denton, E. L.; Chintala, S.; Fergus, R. Deep generative image models using a Laplacian pyramid of adversarial networks. In: Proceedings of Advances in Neural Information Processing Systems, 1486–1494, 2015.

Connelly Barnes is an assistant professor at the University of Virginia. He received his Ph.D. degree from Princeton University in 2011. His group develops techniques for efficiently manipulating visual data in computer graphics by using semantic information from computer vision. Applications are in computational photography, image editing, art, and hiding visual information. Many computer graphics algorithms are more useful if they are interactive; therefore, his group also has a focus on efficiency and optimization, including some compiler technologies.

Fang-Lue Zhang is a post-doctoral research associate at the Department of Computer Science and Technology in Tsinghua University. He received his doctoral degree from Tsinghua University in 2015 and bachelor degree from Zhejiang University in 2009. His research interests include image and video editing, computer vision, and computer graphics.