Research Article

Autocompletion of repetitive stroking with image guidance

Yilan Chen[1]**, Kin Chung Kwan**[2]**, and Hongbo Fu**[1] (✉)

Abstract Image-guided drawing can compensate for a lack of skill but often requires a significant number of repetitive strokes to create textures. Existing automatic stroke synthesis methods are usually limited to predefined styles or require indirect manipulation that may break the spontaneous flow of drawing. We present an assisted drawing system to autocomplete repetitive short strokes during a user's normal drawing process. Users draw over a reference image as usual; at the same time, our system silently analyzes the input strokes and the reference to infer strokes that follow the user's input style when certain repetition is detected. Users can accept, modify, or ignore the system's predictions and continue drawing, thus maintaining fluid control over drawing. Our key idea is to jointly analyze image regions and user input history to detect and predict repetition. The proposed system can effectively reduce the user's workload when drawing repetitive short strokes, helping users to create results with rich patterns.

Keywords interaction; autocompletion; digital drawing; prediction; texture synthesis

1 Introduction

Drawing is a common form of artistic expression. By varying the strokes, texture, and shading, artists can create drawings in various styles [1]. However, it remains a largely manual process that may require significant artistic expertise and repetitive manual labor.

1 School of Creative Media, City University of Hong Kong, Tat Chee Avenue, Kowloon, Hong Kong SAR, China. E-mail: Y. Chen, yil.ellen.chan@gmail.com; H. Fu, hongbofu@cityu.edu.hk (✉).

2 University of Konstanz, University Street 10, 78464 Konstanz, Germany. E-mail: kckwan@ieee.org.

To reduce repetitive workload, various methods have been proposed to automatically synthesize strokes from user-provided examples [4–6] or through procedural steps [7]. However, these methods usually perform in batches, reducing user participation in the creative artistic process. Furthermore, since many methods have predefined styles and only allow users to modify a few global parameters, the final results may look monotonous and lack originality (see Fig. 1). Other interactive systems [8, 9] preserve the normal drawing flow while automating significant stroke synthesis. We share a similar goal to theirs. However, they typically target experts, requiring artistic expertise for high-level picture composition and fine-grained control.

One common way to overcome this skill barrier is to use a reference photo as a scaffold for drawing, by tracing a reference photo physically using transparent paper, or digitally via layers in digital drawing applications. Prior research [10] shows that even when a reference image is used as a scaffold, people still enjoy the freedom of individual expression. We thus propose to enhance drawing using an image


(a) our result   (b) produced with Ref. [2]   (c) produced with Ref. [3]

Fig. 1 Style comparison. (a) Our work is designed to reduce the workload of completing repetitive patterns during the manual drawing process. Full control over the drawing process leads to more dynamic results than (b) Photoshop's art history brush tool [2] and (c) StippleShop [3].

scaffold by automating tedious repetitions. Our idea is to bridge the two extremes: *manual drawing*, which allows full control but can be tedious, and *image-based algorithmic synthesis*, which saves effort but provides limited user control and interactivity. As the first attempt towards this goal, we focus on autocompleting repetitive *short* strokes, which are very common in pen-and-ink drawing (see Fig. 2), under the guidance of a reference image. As in typical digital drawing applications, users can draw freely on a reference image with our system. Meanwhile, our system analyzes the relationships between user inputs and the reference image, detects potential repetitions, and suggests what users might want to draw next. Users can accept, reject, or ignore the suggestions and continue drawing, thus maintaining fluid control of drawing. See Fig. 3 for an example.

The major contribution of this paper is the technical design of an image-guided autocompletion drawing tool that can preserve the natural drawing process and individual user styles. Our approach is inspired by image analogy [4] and operation history analysis and synthesis [9] while leveraging two key insights. Firstly, since the act of drawing repetitive strokes usually indicates specific intentions (e.g., filling an object or hatching a shaded region), we use common image features shared by the coherent repetitive strokes to infer the intended region. Secondly, the drawing is usually related to the underlying reference image (e.g., the density of strokes depends on image brightness). Therefore, we analyze the properties of both the drawing and the reference image to infer possible relationships as contextual constraints for stroke prediction.

We have implemented a prototype and conducted a pilot study with participants from different backgrounds to evaluate its utility and usability. The quantitative analysis and qualitative feedback, as well as various drawing results created by the users, suggest that our system effectively reduces user's workload when drawing repetitive short strokes, helping users to create results with rich patterns.

2 Related work

2.1 Image-assisted drawing

Many drawing support tools adopt reference images and provide intelligent assistance to novices, e.g., beautifying user's sketches with extracted image features [10–13], or providing educational guidance to novice users [14–16]. We share a similar goal to Refs. [17–19] of reducing the user's workload. However, these works use predefined algorithms to


(a) by Alphonso Dunn   (b) by Vincent van Gogh

Fig. 2 Inspiring manual drawings by artists.

(a) user input   (b) suggestion   (c) accept   (d) type visualization   (e) result

Fig. 3 Example of our system workflow. (a) A user stipples over a leaf region of a reference image while our system predicts what she might draw next (b) (blue strokes: inferred exemplars; pale red region: inferred target region; semi-transparent strokes: system suggestions), (c) which is then accepted by the user (green strokes: user inputs or accepted suggestions). (d) shows the manually drawn content (black, 261 strokes) and autocompleted content (red, 3510 strokes). (e) shows the final result; note the different repetitive stroke patterns in different regions. Our autocompletion system can reduce tedious repetitive input, while providing full user control.

generate strokes guided by cursor movement and only take the user's input as an indicator of where to render, thus greatly limiting the user's artistic freedom. In contrast, we aim to provide more flexibility between automatic synthesis and manual artistic control by autocompleting tedious repetitions during the user's normal drawing processes.

2.2 Image-based artistic rendering

Our work is related to image-based artistic rendering (IB-AR) [20], especially stroke-based methods and example-based methods.

Stroke-based methods create artistic results from images by strategically generating brushstrokes whose properties (e.g., position, density, orientation, color, size) are related to image properties (e.g., gradient, edges, color, salience) [7]. Among those methods, the closest to ours are the early image-based pen-and-ink rendering methods [21, 22], which allow users to input sample elements for distribution. However, users have to prepare the sample elements separately (usually as a standalone file) and then adjust parameters to produce the rendered output. In contrast, our system lets users directly provide exemplars on a reference image while silently inferring the distribution properties.

Example-based methods aim to model the visual features of example images for transferal. There are two major modeling approaches: the parametric approach [6, 23, 24] that is based on statistical analysis of stroke characteristics, thus preserving global textures better, and the non-parametric approach [4, 5, 25] based on patch-wise mapping, thus capturing local structures better. We combine both methods to generate strokes: the parametric approach is used to infer statistical relationships between stroke properties and image features, and the patch-wise matching method is used to preserve the local arrangement of strokes. Stylit [5] allows users to stylize a rendered ball and simultaneously propagates the style to arbitrary 3D shapes. Our method shares a similar idea of interactive style propagation, but with two main differences. Firstly, instead of propagating a style globally, we propagate a style to perceptually similar local areas so that users can conveniently define different styles in different areas. Secondly, we represent drawings as discrete stroke operations instead of raster textures to better preserve their structure, e.g., changing the color or size of the drawn strokes.

2.3 Operation history-assisted authoring

Operation histories [26] have been leveraged for various authoring tasks, such as sketching [9], animation [27, 28], modeling [29, 30], beautification of freehand drawings [31], and handwriting [32]. Our work is most closely related to that of Xing et al. [9], which autocompletes repetitive sketching by analyzing dynamic operations recorded during authoring. Our method extends their work to consider additional information from a reference image and thus enables the propagation of strokes to regions with similar image attributes such as color or semantic meaning.

In our use case, an operation is an input stroke, so our work is also related to stroke pattern analysis and synthesis [8, 33–36]. These works disregard the temporal relationships between past strokes, and do not use image guidance, so are different from ours.

We summarise the major differences between our work and the discussed closely related work in Table 1.

Table 1 Differences between our tool and closely related work. B means generation is performed in batches, based on predefined attributes, while Dy means generation is based on dynamic operation history. S indicates operation on strokes, P on pixels. Di means users can specify a style by directly operating on the output, I indicates indirect interaction Y and N represent yes and no, respectively, for using image references

Method	Ref. [22]	Ref. [4]	Ref. [6]	Ref. [8]	Ref. [9]	Ours
Reference	Y	Y	Y	N	N	Y
Process	B	B	B	B	Dy	Dy
Format	S	P	S	S	S	S
Operate	I	I	I	Di	Di	Di

3 User interface

Our system prototype follows a standard digital drawing interface, with addition of our autocompletion feature, as shown in Fig. 4. The user draws on top of the reference image displayed semi-transparently on the main canvas, while our system analyzes the input strokes and the reference image in the background.

3.1 Autocompletion

In autocompletion mode, our system automatically performs analysis whenever the user finishes a new stroke. When a potential repetition is detected, our system highlights the current repetitive strokes and an inferred propagation region, updates the inferred parameters in the filling property panel, and generates
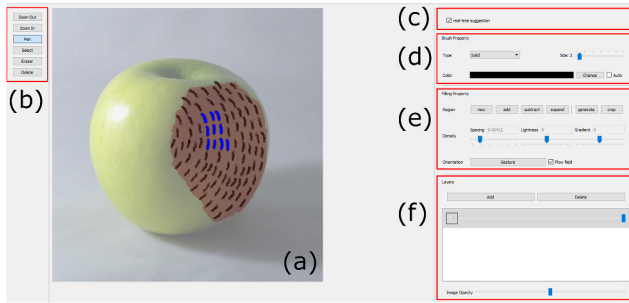
Fig. 4 User interface, comprising (a) a central drawing canvas, (b) a toolbar for drawing and selection, (c) a toggle-switch for autocompletion mode, (d) a brush property toolbar, (e) a filling property toolbar, and (f) a layer panel.

an autocompletion suggestion. Users can accept or reject the suggestion using hotkeys, accept part of it using lasso selection, or ignore it and continue to draw (Fig. 5). The suggestion will continuously update according to user input.

3.2 Interactive editing

Our system provides a set of tools to refine the autocompleted results.

Propagation region editing. Users can create, add, or subtract a region using the intelligent scissors tool [37], or expand an existing region by a fixed width (see Fig. 4(e)) for stroke autocompletion. Figure 6 shows an example of creating a new region for stroke regeneration.

Density editing. Users can modify three parameters to adjust the density of the generated strokes: the average *spacing*, the *lightness* coefficient, and the *gradient* coefficient. The latter two define the relationships between density and image lightness and gradient, respectively. Our system automatically updates these parameters upon prediction, and the updated parameters provide a starting point for user manipulation. Figure 7 shows an example.

Orientation editing. Our system automatically predicts whether the input exemplar is correlated with the image flow; orientation can also be adjusted by the user manually. The user can also modify the image flow field using the gesture brush, and the touched strokes will be rotated to be aligned with the gesture. See Fig. 8 for an example.
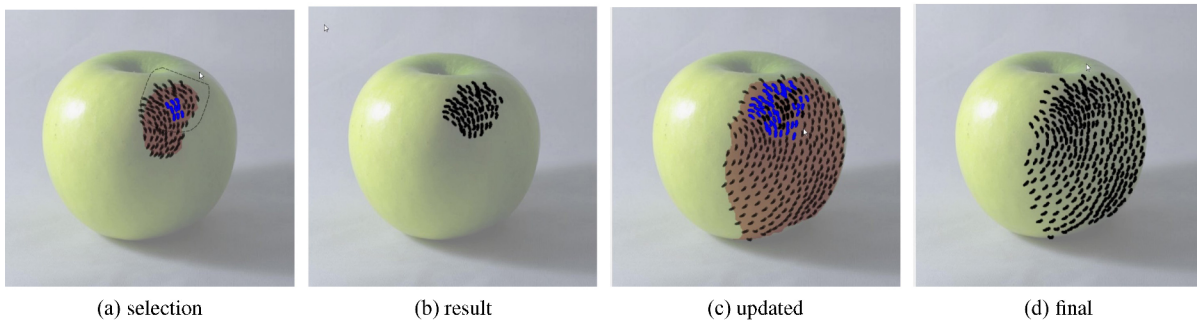

(a) selection   (b) result   (c) updated   (d) final

Fig. 5 An example of autocompletion. The user selects part of the suggestion using the lasso tool (a) with the result shown in (b), then continues to draw leading to the updated suggestion (c), and accepts all the suggestions using a hotkey (d). The blue strokes in (a) and (c) indicate inferred exemplars from the user's input strokes.
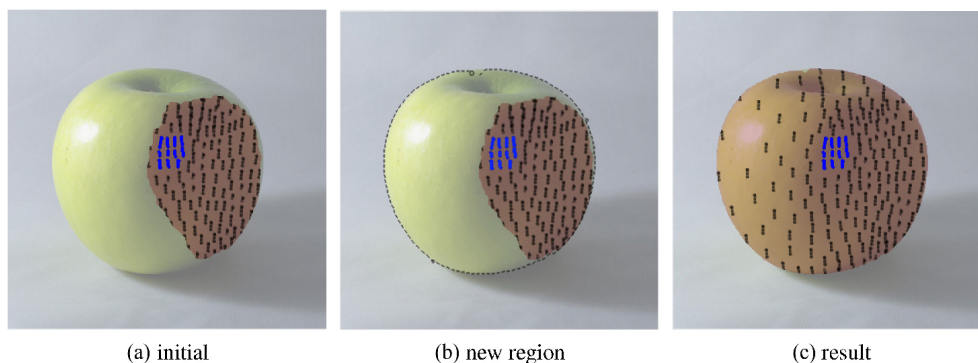

(a) initial   (b) new region   (c) result

Fig. 6 Region editing example. The initial prediction (a) contains only the brown region. The user-specified region (b) contains the entire apple, with the corresponding synthesis result in (c).

清华大学出版社
Tsinghua University Press   ⚡ Springer

(a) (8, 0, 0)   (b) (15, 0, 0)   (c) (8, 0.2, 0)   (d) (8, 0, 0.6)

Fig. 7 Density editing example with different values of *spacing*, *lightness*, and *gradient* parameters. Larger spacings lead to sparser strokes, while greater lightness and gradient lead to larger stroke density variations.


(a)   (b)   (c)   (d)

Fig. 8 Orientation editing example. (a) User gesture. (b) Orientation field updated based on the user gesture and the original image flow field. (c) Updated result. (d) Result without any orientation field.

3.3 Auxiliary functions

Our prototype also includes the auxiliary functions below. These are not unique to our system but can facilitate the usual drawing processes.

Post-editing stroke properties. Users can select existing strokes and edit their properties, such as size and color.

Auto-coloring. This function, when used, can automatically colorize strokes with color from the reference image.

View switching. Users can press the space key to switch between the canvas view, reference view, and pure drawing view.

4 Approach

Our system involves two key steps: (i) inferring the input exemplar, the output region, and the contextual constraints from the stroke history and the reference image, and (ii) synthesizing suggestive strokes accordingly. Section 4.1 first describes how to *synthesize* strokes, assuming that all the information is available, and then Section 4.2 explains how to *infer* the necessary information for synthesis.

4.1 Stroke synthesis

4.1.1 Problem statement

The inputs to our stroke synthesis method include an exemplar E consisting of repetitive strokes, the reference image I, a target region mask M, an orientation map O, and a radius map R. Pixel values of R determine the stroke spacing: a smaller value leads to a denser distribution. Our goal is to compute an output set of strokes X over the output region M, such that X is similar to E with respect to I. We describe how to infer E, M, O, and R from user interaction with I in Section 4.2.

4.1.2 Idea

To support autocompletion using the reference image, we extend the discrete element texture synthesis method [9, 38], which represents strokes as point samples and iteratively improves the sample distribution by minimizing the neighborhood difference between the exemplar and the output, using an additional reference image. Firstly, we combine sample neighborhoods [38] with image features [4] to measure neighborhood differences. Secondly, the range and orientation of each sample neighborhood

is determined by the radius and orientation maps inferred from the reference image. Figure 9 shows our key idea.

4.1.3 Stroke representation

A stroke s is an ordered list of sample points, each with a timestamp and appearance attributes such as thickness and color. Here we focus on autocompleting short strokes, so we represent each stroke by its centroid p and the average direction v (see Fig. 10) for efficiency of synthesis, without considering any other information about the original stroke. To take drawing order into consideration, we obtain the dominant direction by averaging the vectors from the start point to each subsequent point. After synthesis, we reconstruct all sample points according to the updated centroid and direction.

4.1.4 Initialization

We pre-process the target region mask M by removing the area occupied by existing strokes in the same layer to avoid clutter, and then initialize the output X by generating sample positions with Poisson-disk sampling based on the radius map R. For each sampled position, we copy the input stroke with the


Fig. 9 Synthesis algorithm. We synthesize the predicted strokes (green) from previously drawn strokes (gray) by matching their neighborhoods as well as image features.


Fig. 10 Left: A stroke, with centroid p and dominant direction v. Right: The neighborhood of the black stroke includes the n ($n = 1$ in this example) closest strokes (green) from each quadrant and the middle image patch (blue pixel grid).

smallest image feature distance d_I, which will be explained in Eq. (2). We then optimize the output for several objectives, as detailed below.

4.1.5 Neighborhood term

We define the neighborhood of a stroke s by both its neighboring strokes as well as an $R(s) \times R(s)$ image patch around its centroid, where $R(s)$ is the radius value at s. Prior methods (e.g., Ref. [38]) determine the neighboring strokes by spatial distance. Thus, the neighborhood radius should be large enough to capture any underlying pattern. However, this might include redundant strokes and thus decrease performance. Therefore, we adopt Zhao and Zhu's method [39] to automatically find a minimal representative neighborhood, considering not only the distances between strokes but also their locations. As Fig. 10(b) shows, we set the neighborhood radius of the center stroke s to $2R(s)$. We then divide all strokes within the neighborhood radius into four quadrants with respect to the local frame defined by the orientation at $O(s)$, and collect the n nearest strokes from each quadrant as the representative neighborhood $\boldsymbol{N}(s)$. In our implementation, we set $n = 4$ for the input exemplar and $n = 1$ for the output strokes to ensure that each output neighborhood can be maximally matched.

For a stroke s and a neighboring stroke $s' \in \boldsymbol{N}(s)$, we compute their offsets in position and direction to be

$$\hat{u}(s', s) = \left(O(s)^{-1}\left(p(s') - p(s)\right)/R(s)\,,\right.$$
$$\left.O(s)^{-1}\left(v(s') - v(s)\right)\right) \tag{1}$$

where $O(s)^{-1}$ indicates rotating the vector inversely to $O(s)$. Note that the position and direction difference is computed in the local frame defined by the density map and orientation map. For an output stroke s_o and an input stroke s_i, we first find their best matching pairs $\{(s'_\mathrm{o}, s'_\mathrm{i})\}$ in the neighborhoods $\boldsymbol{N}(s_\mathrm{o})$ and $\boldsymbol{N}(s_\mathrm{i})$ using the Hungarian algorithm [38, 40]. We use the norm-2 distance of the offsets from s_o or s_i in Eq. (1) as the matching cost. The neighborhood distance is then defined as

$$d_\mathrm{neigh}(s_\mathrm{o}, s_\mathrm{i}) = \sum_{s'_\mathrm{o} \in \boldsymbol{N}(s_\mathrm{o})} \left|\hat{u}(s'_\mathrm{o}, s_\mathrm{o}) - \hat{u}(s'_\mathrm{i}, s_\mathrm{i})\right|^2$$
$$+ \mu \underbrace{\left|I(s_\mathrm{o}) - I(s_\mathrm{i})\right|^2}_{d_I} \tag{2}$$

The second term measures the image feature distance d_I; μ ($= 0.1$ in our implementation) controls its

relative weight. We use the mean Lab* color of an $r \times r$ patch at the stroke centroid as the image feature vector. The overall neighborhood term to minimize is

$$\phi_{\text{neigh}}(X, E) = \sum_{s_{\text{o}} \in X} \min_{s_{\text{i}} \in E} d_{\text{neigh}}(s_{\text{o}}, s_{\text{i}}) \qquad (3)$$

4.1.6 Correction term

Since the neighborhood term is a one-way matching from output neighborhoods to input neighborhoods, sometimes optimization may tend to leave out some void regions. Furthermore, the neighborhood term does not preserve the alignment of strokes to the image (e.g., see Fig. 11(e)). To address these issues, we apply a correction term. We compute a weighted centroidal Voronoi diagram from all the strokes' center points, using $1/R$ as weight; we denote the computed region centroids as $\{\bar{p}\}$. Then we can minimize the total distance between each output stroke centroid and the region centroid, defined as Eq. (4):

$$\phi_{\text{corr}}(X) = \sum_{s_{\text{o}} \in X} |p(s_{\text{o}}) - \bar{p}(s_{\text{o}})|^2 \qquad (4)$$

4.1.7 Solver

The energy function we aim to minimize is defined as

$$\phi(X, E) = (1 - w)\phi_{\text{neigh}} + w\phi_{\text{corr}} \qquad (5)$$

We iteratively minimize the energy function following the EM methodology in Ref. [38]. In each iteration, for each output stroke s_{o}, we search for the closest matching input stroke s_{i} to minimize ϕ_{neigh}, compute the Voronoi diagram centroid $\bar{p}$ to minimize ϕ_{corr}, and solve a least-squares system combining both terms. Let m be the total number of iterations. For the i-th iteration, we set $w = (i/m)^2$, so that more weight is given to ϕ_{neigh} in earlier iterations, to optimize the neighborhood distribution first before making corrections: this leads to better results.

Figures 11(b)–11(d) show iterative optimization of both objectives. For comparison, Fig. 11(e) shows the result without the correction term and Fig. 11(f) shows the result without using the image neighborhood in both initialization and optimization.
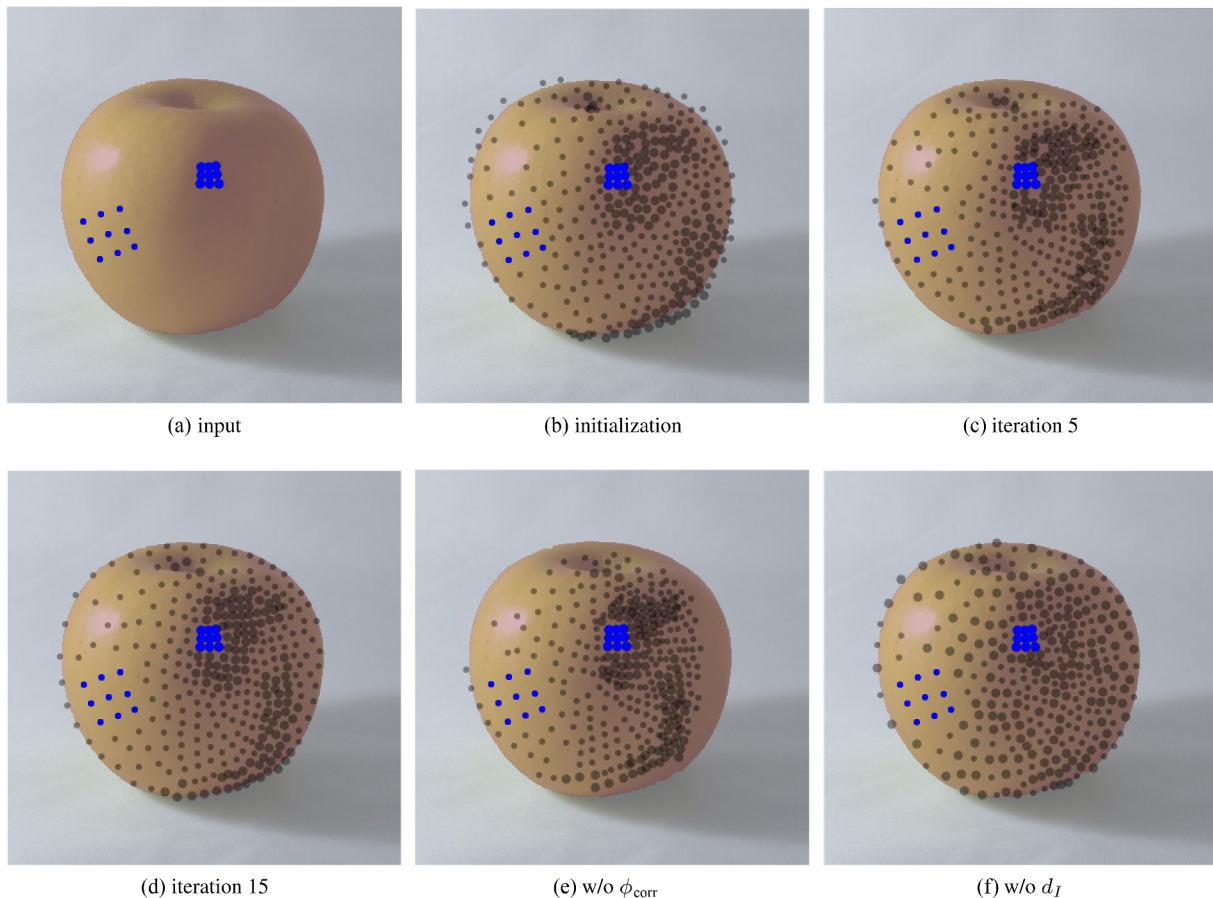

(a) input   (b) initialization   (c) iteration 5

(d) iteration 15   (e) w/o ϕ_{corr} (f) w/o d_I

Fig. 11 (a) Input. (b–d) Iteration process. (e, f) Ablation studies. Without the correction term ϕ_{corr} the predicted strokes tend to cluster together (e). Without the image term d_I the predicted strokes may not follow the reference sufficiently (f).

4.2 Inference

In this section, we describe how we infer E, M, O, and R for synthesis from user interactions with I.

4.2.1 Input exemplar E

In this step we aim to detect whether stroke repetitions exist and obtain the repetitive group as an exemplar for the synthesis process. Since people usually draw strokes in a coherent manner [9] and they usually have specific intentions when drawing repetitive strokes, we assume the example strokes to be temporally consecutive and to have certain similar properties.

We start from the last stroke input by the user and search backward in the stroke sequence to incrementally find strokes with similar shape and image features to the last stroke. Stroke shape similarity is measured using the Fréchet distance, while the image features include Lab* color (weighted by 0.12, 0.44, and 0.44 to suppress the impact of lightness) and precomputed semantic segmentation [41] at a stroke's center. Alternatively, one could use different image features to capture different drawing intentions. We compare the standard deviation of a feature in the traversed k strokes against a threshold (15/255 for the color feature, 1 for the segmentation feature) for similarity measurement. Back-traversal stops when the next

stroke does not contain a similar feature or $k > 50$. These k strokes serve as the input exemplar for the synthesis process. See Fig. 12 for an example of the incremental search process.

4.2.2 Output region M

The shared features of the obtained stroke exemplar also indicate the intended region. For instance, if all exemplar strokes lie inside the same object segmentation region, it is very likely that the user intends to fill that region. Therefore, we use the shared features obtained in the exemplar grouping process to find a similar region for output.

Since we have only two features in our implementation, we simply obtain the region by GrabCut [42] if the Lab* color feature is shared by the exemplar strokes, we directly take the corresponding segmentation if the semantic feature is shared, and we take the intersection if both features are shared. See Fig. 12 for an example. When there are multiple disconnected regions, we retain the nearest region to the user's last stroke and discard the remainder, as it is less natural to propagate to distant regions.

4.2.3 Contextual constraints

Since the drawing is usually related to the underlying reference image, we analyze the properties of both the drawn strokes and the reference image to infer possible relationships to control the global
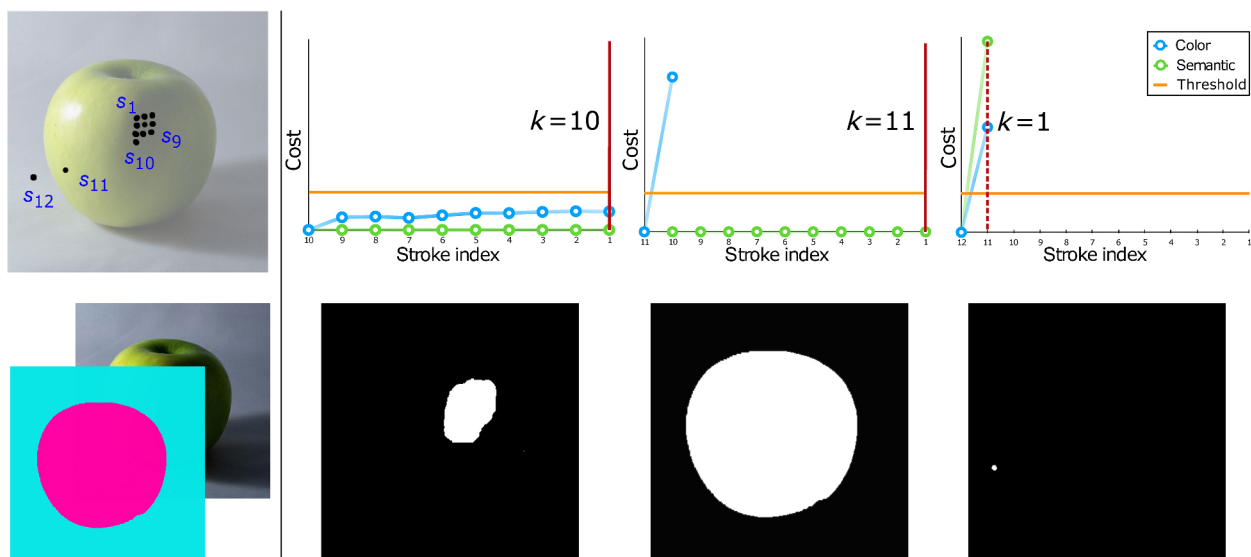

Fig. 12 Predicting the input exemplar and output region. Left, above: the input stroke sequence (black dots, only a few indices are shown for clarity) on the reference image. Left, below: image features. Right, above: threshold lines and the image feature cost curves for s_{10}, s_{11}, s_{12} respectively. Right, below: corresponding predicted output regions. The cumulative number k is determined when both cost curves exceed the threshold. Note that the third region prediction result is only for demonstration: since the exemplar only contains one stroke (i.e., $k = 1$), it is not considered a valid exemplar and would not be used for synthesis.

distribution of strokes. The constraints we consider are orientation O and radius R.

Artists usually adjust stroke directions to convey curvature, but may sometimes randomize or fix stroke orientation regardless of the depicted objects to create different visual effects. Therefore, the problem is to decide which case the input exemplar implies. We first compute the edge tangent field (ETF) [43] for the reference image and then calculate the angles between the exemplar strokes and the ETF directions at their centroids. If the standard deviation of the angles is small (less than $15°$), we consider the stroke orientations to be related to the ETF and take the ETF as the orientation field; otherwise, we set a default global coordinate frame at each point of the orientation field.

Since density is inversely proportional to the spacing between strokes, we reframe the spacing problem as predicting a radius map that controls the extents of stroke neighborhoods. First, we compute the distance from each exemplar stroke to its nearest neighbor. We assume a linear relationship between these minimum distances r and the image features, including image lightness l and gradient strength g at a stroke's centroid, represented as

$$r = \begin{pmatrix} l & g & 1 \end{pmatrix} \cdot \boldsymbol{t} \tag{6}$$

where $\boldsymbol{t}$ denotes the coefficients to be found. Using the fitted linear model, if the squared correlation value is lower than 0.5 (the closer to 1, the better explanation), we use the model to compute a radius map. Otherwise, we consider the density to be uniform and create a constant radius map using the average spatial distance of the exemplar. We then update the user interface with the computed coefficients.

5 Evaluation

5.1 Approach

We conducted a pilot study to evaluate the utility and usability of our approach. We compared three modes through quantitative analysis and qualitative feedback.

In *autocompletion mode*, users had full access to our prototype, including autocompletion and interactive editing.

In *interactive batch filling mode* (batch mode), users were required to create a texture example first and then manually specify properties for batch filling. It simulates the sequential procedure in many IB-AR methods (e.g., Ref. [21]), although they rarely allow users to directly define examples on target images. This mode was performed using our system with autocompletion turned off.

In *fully manual drawing mode* (manual mode), users had to manually draw each stroke without any automatic synthesis.

We also tested the expressiveness of our system through an open creation session and obtained comments for future improvements.

5.2 Target session

The goal of this session was to compare the three interaction modes in terms of utility and usability. Since we aim to facilitate drawing using an image scaffold, we included general users with different backgrounds but focusing on less skillful users, who are more likely to want to use reference images. We thus recruited 12 participants, including nine novices with little drawing experience, two amateurs with some experience (P3, P4), and a student majoring in illustration (P5). Most of the studies were conducted on a Lenovo Miix 520 tablet with a stylus, in a lab environment, except for two studies conducted remotely using a mouse (due to the covid pandemic).

The study procedure consisted of a tutorial followed by target tasks, and took each participant about two hours in total.

Each participant was first given a brief introduction to our system and then asked to fill the apple in Fig. 4 with short hatches as a training task. They were encouraged to vary the density and orientation of input strokes and to become familiar with the features of our system.

We used a within-subject design, in which each participant was asked to reproduce two target drawings (see Fig. 13) in all three modes: autocompletion, batch mode, and manual mode. The target drawings contained an object and a landscape, common illustration topics (e.g., see Fig. 2). The assigned order of modes was randomised over all participants. Since we focus on region filling, we asked the participants to draw the outlines of both images in advance, so that they could focus on drawing the textures during the study. We encouraged the participants to finish each drawing as soon as possible, preferably within about a dozen minutes, but without

(a) bear   (b) drawing   (c) segmentation   (d) orientation

(e) beach   (f) drawing   (g) segmentation   (h) orientation

Fig. 13 Target session tasks. (a, e) Reference photos. (b, f) Corresponding sample outputs.

any hard time limit. After completing the two drawings in each mode, each participant filled in a NASA-TLX questionnaire [44]. Finally, we asked the participants about their preferred mode and usage experience, and for other comments.

5.3 Open session

The goal of this session was to observe how users interact with our system and to learn about user's subjective experiences. We invited seven participants (one professional artist, two amateurs, and four novices) for this session. They were asked to create a drawing freely from the same reference image (Fig. 15(a)) using our system. The reference image was a portrait photo, also common in illustrations. The only requirement was that the drawings should contain some repetitive content. We again commenced with a tutorial and conducted the task on a Lenovo Miix 520 tablet with stylus. The participants were encouraged to think aloud and describe their thought processes and interaction during this session. After this task, participants could optionally create further drawings from any images they wanted. Since our prototype does not contain all common functionality of commercial drawing tools, we allowed the participants to retouch the resulting drawings, without adding more strokes, in Photoshop.

5.4 Results and observations

5.4.1 Workload

Figure 14(a) shows the perceived workload scores from the target session. Generally, the autocompletion mode received the lowest (best) scores for almost all factors. One-way ANOVA showed the three modes have significant differences in physical demand ($F = 10.69$, $p < 0.001$) while having no significant difference in other factors. Regarding physical demand, post-hoc pairwise tests showed that the autocompletion mode and batch mode both rated significantly lower than manual mode, but had no significant difference from each other. This matches our expectation, since automatic synthesis should only reduce physical load and not cause extra pressure over manual work.

5.4.2 Efficiency

We calculated the average completion time (Fig. 14(b)) and stroke count (Fig. 14(c)) for each mode and task. Generally, the system synthesized about 82% of the strokes in the autocompletion mode and about 92% of the strokes in batch mode. Although manual mode took the shortest time for participants to complete, it also resulted in the fewest total strokes. We thus calculated the strokes per minute for each mode: autocompletion (111.03, SD = 38.76), batch (101.98, SD = 45.13), manual (115.95, SD = 46.73). It
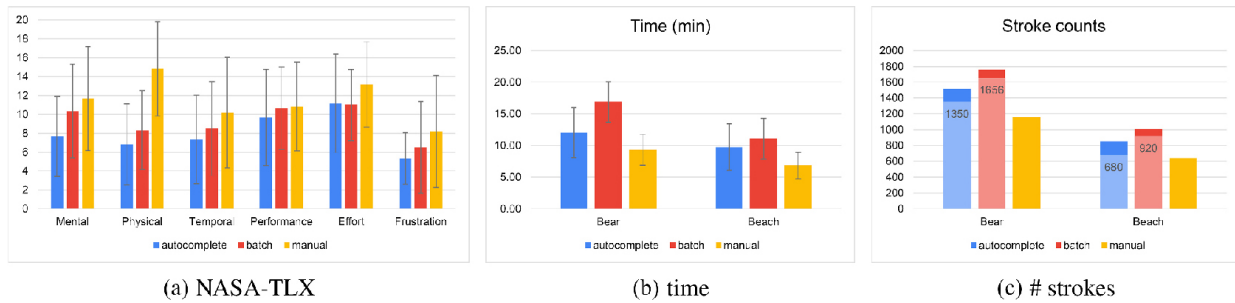
(a) NASA-TLX   (b) time   (c) # strokes

Fig. 14 Target session results. (a) Average NASA-TLX scores from 12 participants. Lower scores are better. (b) Average completion time. (c) Average stroke counts. The number of system-generated strokes is labeled in each column.

(a) reference   (b) 81/1563   (c) 428/4593   (d) 272/1266

(e) 68/8356   (f) 165/17111   (g) 443/2931   (h) 261/6018

Fig. 15 Example drawing results from the open session. Each case indicates the number of manual/autocompleted strokes.

turns out that automatic generation did not improve efficiency, probably because the users spent extra time adjusting and experimenting with the generated effects instead of just drawing strokes. It should be noted that such directed tasks omit the time taken to explore alternative patterns, which, however, might be high in a fully manual case.

5.4.3 Quality

We asked 30 external volunteers to evaluate the quality of participants' drawings. We randomized all drawings created by the participants, showed each output drawing alongside the target drawing, and asked volunteers to rate the resemblance of the output drawing to the target drawing, on a scale from 1 (very dissimilar) to 5 (very similar). The volunteers were instructed to focus more on the overall stroke distributions and flows instead of individual stroke thickness and detailed shapes. We calculated average scores for each mode: autocompletion (3.10, SD = 1.24), batch (3.09, SD = 1.21), manual (2.98, SD = 1.20). The quality of the drawings created with automatic synthesis was slightly better than for the fully manual drawings, but without a significant difference. From the participants' perspective, three novices commented that the automated strokes were better than their manual strokes, because they tend to become impatient when manually drawing all strokes, lowering quality.

5.5 Preferred mode

Seven participants preferred autocompletion mode while the other five participants preferred batch mode. Generally, autocomplete mode was considered more convenient, but less precise; batch mode was considered more precise, but requires too much

interaction. P12 commented, "*autocompletion mode is more straightforward, because you can see the filling effects instantly without doing a lot of manipulation beforehand; while in batch mode, you have to remember the meanings of parameters and adjust them in order to create strokes.*" P10 also said, "*compared to batch filling, autocompletion mode provides a quick guess for filled regions and allows me to get results more quickly with less work.*" However, autocompletion mode is "*less accurate in some vague and detailed regions, such as the shadows of the boat, where it tends to include some unwanted regions, so I had to manually subtract those*

regions, which is a bit tedious", according to P3. The professional, P5, also preferred batch mode for being able to precisely select regions. Therefore, we consider the autocompletion function and the interactive editing function to be complementary in usability.

5.6 Creative results and experience

Figure 15 shows outcomes from the open session. Although from the same reference image and widely using repetitive short strokes, the study participants were able to create different results by varying the stroke shapes and arrangement. Figure 16
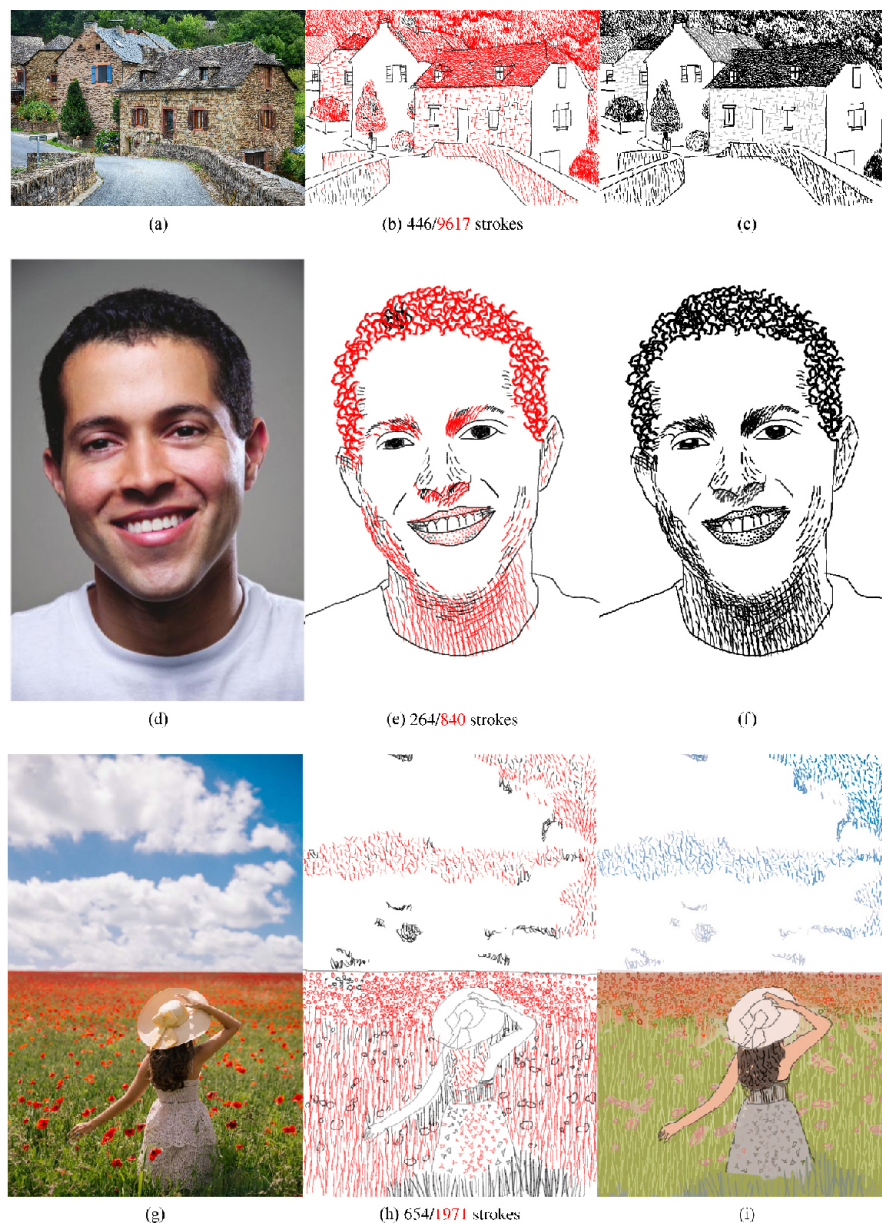

Fig. 16 Sample results. For each example: left: reference image, center: manual (black) and autocompleted (red) strokes, right: final drawings. In the last example, the strokes were created with our system first and then imported into Photoshop for background coloring.

demonstrates further results. Regarding the creation experience, one user said "it is playful, and the final result is also good", two users described it as "encouraging", because the system allows beginners to quickly create stylistic drawings, and one user commented that she "felt creative when drawing with this system", because she could try out patterns over image regions conveniently and she was more comfortable with drawing from a reference image than from scratch. The professional suggested that the tool itself was somewhat limited to pointillism and hatching styles, but could be helpful in adding interesting textures to color paintings (e.g., see Fig. 16(i)). Two users commented that the reduction in workload is useful, but they also complained about some inaccurate inferences of the autocompletion. We further discuss this problem in Section 6.

6 Limitations and future work

From our observation and users' feedback, we identified several opportunities for improvement.

6.1 Accuracy of autocompletion

We rely on simple Lab* color and semantic segmentation for region inference. While color features suffice for most cases, regions with similar colors but different semantics require sufficient segmentation accuracy for region inference (e.g., see Fig. 13(c)). More advanced semantic selection methods (e.g., Ref. [45]) might help to infer more accurate regions. However, granularity of selection requires further study. For example, when users draw on a bear's limb, is the intended region the whole bear, or all limbs? We leave this as future work.

6.2 Visual blocking

Since the drawing and the system's suggestions are overlaid on the reference image, it can be difficult for users to see the image when selecting parts of the suggestions (see Fig. 17) or adding a new layer of strokes. Although users can switch views using a hotkey, it might be helpful to provide some reference information, like image darkness or boundaries, through additional visual hints [10, 16].

6.3 Higher-level image features

We only consider relationships between strokes and low-level image features, like colors and flows, over regions. By considering higher-level image features,

Fig. 17 Example of visual blocking. Left: reference image. Right: canvas view.

such as elements and edges, it may be possible to extend the scope of autocompletion, such as autocompleting the sparse flowers in the foreground of Fig. 16(i) through the correspondences between strokes and elements.

6.4 Stroke types

Our method only supports short strokes, while artists frequently also use long repetitive strokes [1]. It is worth investigating the possibility of incorporating continuous strokes [46] in our analysis and synthesis framework and extending support to different input strokes.

7 Conclusions

We have presented a new drawing concept and designed an assisted drawing system to help users autocomplete repetitive short strokes with guidance from reference images while maintaining the flexible control of manual drawing. By extending operation history analysis and synthesis with image analysis, our system is able to generate results adapted to reference images and users' prior inputs. We conducted a pilot study to validate the usefulness of our approach and show various drawing results from the users.

Acknowledgements

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article.

References

[1] Dunn, A. *Pen and Ink Drawing: A Simple Guide.* New Jersey: Three Minds Press, 2015.

[2] Adobe. Paint stylized strokes with the Art History Brush. 2017. Available at https://helpx.adobe.com/photoshop/using/painting-stylized-strokes-art-history.html.

[3] Martín, D.; Arroyo, G.; Rodríguez, A.; Isenberg, T. A survey of digital stippling. *Computers & Graphics* Vol. 67, 24–44, 2017.

[4] 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.

[5] 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.

[6] Gerl, M.; Isenberg, T. Interactive example-based hatching. *Computers & Graphics* Vol. 37, Nos. 1–2, 65–80, 2013.

[7] Hegde, S.; Gatzidis, C.; Tian, F. Painterly rendering techniques: A state-of-the-art review of current approaches. *Computer Animation and Virtual Worlds* Vol. 24, No. 1, 43–64, 2013.

[8] Kazi, R. H.; Igarashi, T.; Zhao, S. D.; Davis, R. Vignette: Interactive texture design and manipulation with freeform gestures for pen-and-ink illustration. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1727–1736, 2012.

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

[10] Xie, J.; Hertzmann, A.; Li, W.; Winnemöller, H. PortraitSketch: Face sketching assistance for novices. In: Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology, 407–417, 2014.

[11] Kang, H. W.; He, W. J.; Chui, C. K.; Chakraborty, U. K. Interactive sketch generation. *The Visual Computer* Vol. 21, Nos. 8–10, 821–830, 2005.

[12] Su, Q.; Li, W. H. A.; Wang, J.; Fu, H. EZ-sketching: Three-level optimization for error-tolerant image tracing. *ACM Transactions on Graphics* Vol. 33, No. 4, Article No. 54, 2014.

[13] Li, G. B.; Bi, S.; Wang, J.; Xu, Y. Q.; Yu, Y. Z. ColorSketch: A drawing assistant for generating color sketches from photos. *IEEE Computer Graphics and Applications* Vol. 37, No. 3, 70–81, 2017.

[14] Iarussi, E.; Bousseau, A.; Tsandilas, T. The drawing assistant: Automated drawing guidance and feedback from photographs. In: Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology, 183–192, 2013.

[15] Matsui, Y.; Shiratori, T.; Aizawa, K. DrawFromDrawings: 2D drawing assistance via stroke interpolation with a sketch database. *IEEE Transactions on Visualization and Computer Graphics* Vol. 23, No. 7, 1852–1862, 2017.

[16] Williford, B.; Doke, A.; Pahud, M.; Hinckley, K.; Hammond, T. DrawMyPhoto: Assisting novices in drawing from photographs. In: Proceedings of the 2019 on Creativity and Cognition, 198–209, 2019.

[17] Haeberli, P. Paint by numbers: Abstract image representations. *ACM SIGGRAPH Computer Graphics* Vol. 24, No. 4, 207–214, 1990.

[18] Benedetti, L.; Winnemöller, H.; Corsini, M.; Scopigno, R. Painting with Bob: Assisted creativity for novices. In: Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology, 419–428, 2014.

[19] Tsai, H. C.; Lee, Y. H.; Lee, R. R.; Chu, H. K. User-guided line abstraction using coherence and structure analysis. *Computational Visual Media* Vol. 3, No. 2, 177–188, 2017.

[20] Kyprianidis, J. E.; Collomosse, J.; Wang, T. H.; Isenberg, T. State of the "art": A taxonomy of artistic stylization techniques for images and video. *IEEE Transactions on Visualization and Computer Graphics* Vol. 19, No. 5, 866–885, 2013.

[21] Salisbury, M. P.; Wong, M. T.; Hughes, J. F.; Salesin, D. H. Orientable textures for image-based pen-and-ink illustration. In: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, 401–406, 1997.

[22] Hiller, S.; Hellwig, H.; Deussen, O. Beyond stippling—Methods for distributing objects on the plane. *Computer Graphics Forum* Vol. 22, No. 3, 515–522, 2003.

[23] Kalogerakis, E.; Nowrouzezahrai, D.; Breslav, S.; Hertzmann, A. Learning hatching for pen-and-ink illustration of surfaces. *ACM Transactions on Graphics* Vol. 31, No. 1, Article No. 1, 2012.

[24] Gatys, L. A.; Ecker, A. S.; Bethge, M.; Hertzmann, A.; Shechtman, E. Controlling perceptual factors in neural style transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 3730–3738, 2017.

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

[26] Nancel, M.; Cockburn, A. Causality: A conceptual model of interaction history. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1777–1786, 2014.

[27] Xing, J.; Wei, L. Y.; Shiratori, T.; Yatani, K. Autocomplete hand-drawn animations. *ACM Transactions on Graphics* Vol. 34, No. 6, Article No. 169, 2015.

[28] Peng, M.; Wei, L. Y.; Kazi, R. H.; Kim, V. G. Autocomplete animated sculpting. In: Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology, 760–777, 2020.

[29] Peng, M. Q.; Xing, J.; Wei, L. Y. Autocomplete 3D sculpting. *ACM Transactions on Graphics* Vol. 37, No. 4, Article No. 132, 2018.

[30] Suzuki, R.; Yatani, K.; Gross, M. D.; Yeh, T. Tabby: Explorable design for 3D printing textures. In: Proceedings of the 26th Pacific Conference on Computer Graphics and Applications: Short Papers, 29–32, 2018.

[31] Fišer, J.; Asente, P.; Sýkora, D. ShipShape: A drawing beautification assistant. In: Proceedings of the International Symposium on Sketch-based Interfaces and Modeling, 49–57, 2015.

[32] Zitnick, C. L. Handwriting beautification using token means. *ACM Transactions on Graphics* Vol. 32, No. 4, Article No. 53, 2013.

[33] Barla, P.; Breslav, S.; Markosian, L.; Thollot, J. Interactive hatching and stippling by example. INRIA Research Report RR-6461, 2006.

[34] Ijiri, T.; Mêch, R.; Igarashi, T.; Miller, G. An example-based procedural system for element arrangement. *Computer Graphics Forum* Vol. 27, No. 2, 429–436, 2008.

[35] Alves dos Passos, V.; Walter, M.; Sousa, M. C. Sample-based synthesis of illustrative patterns. In: Proceedings of the 18th Pacific Conference on Computer Graphics and Applications, 109–116, 2010.

[36] Hsu, C. Y.; Wei, L. Y.; You, L.; Zhang, J. J. Autocomplete element fields. In: Proceedings of the CHI Conference on Human Factors in Computing Systems, 1–13, 2020.

[37] Mortensen, E. N.; Barrett, W. A. Intelligent scissors for image composition. In: Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques, 191–198, 1995.

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

[39] Zhao, M.; Zhu, S. C. Customizing painterly rendering styles using stroke processes. In: Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Non-Photorealistic Animation and Rendering, 137–146, 2011.

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

[41] Zhao, H. S.; Shi, J. P.; Qi, X. J.; Wang, X. G.; Jia, J. Y. Pyramid scene parsing network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 6230–6239, 2017.

[42] Rother, C.; Kolmogorov, V.; Blake, A. "GrabCut": Interactive foreground extraction using iterated graph cuts. In: Proceedings of the ACM SIGGRAPH 2004 Papers, 309–314, 2004.

[43] Kyprianidis, J. E.; Kang, H. Image and video abstraction by coherence-enhancing filtering. *Computer Graphics Forum* Vol. 30, No. 2, 593–602, 2011.

[44] Hart, S. G.; Staveland, L. E. Development of NASA-TLX (task load index): Results of empirical and theoretical research. In: *Advances in Psychology*. Amsterdam: Elsevier, 139–183, 1988.

[45] Chen, X.; Zhao, Z. Y.; Yu, F. W.; Zhang, Y. L.; Duan, M. N. Conditional diffusion for interactive segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 7325–7334, 2021.

[46] Tu, P. H.; Wei, L. Y.; Yatani, K.; Igarashi, T.; Zwicker, M. Continuous curve textures. *ACM Transactions on Graphics* Vol. 39, No. 6, Article No. 168, 2020.

Yilan Chen received her Ph.D. degree from the School of Creative Media, City University of Hong Kong, Hong Kong, China. Her research interests include computer graphics and human–computer interaction.

Kin Chung Kwan received his B.Sc. and Ph.D. degrees from the Chinese University of Hong Kong in 2009 and 2015, respectively. He is now a postdoctoral researcher at the University of Konstanz. His research interests include computer graphics and human–computer interaction.

Hongbo Fu is a professor in the School of Creative Media, City University of Hong Kong. Previously, he had postdoctoral research training at the Imager Lab, University of British Columbia, Canada, and the Department of Computer Graphics, Max-Planck-Institut Informatik, Germany. He received his Ph.D. degree in computer science from Hong Kong University of Science and Technology in 2007 and his B.S. degree in information sciences from Peking University, China, in 2002. His primary research interests fall in the fields of computer graphics and human–computer interaction. He has served as an associate editor of *The*

Visual Computer, Computers & Graphics, and *Computer Graphics Forum.*

Other papers from this open access journal are available free of charge from http://www.springer.com/journal/41095. To submit a manuscript, please go to https://www.editorialmanager.com/cvmj.

清華大学出版社
Tsinghua University Press   Springer