

Autocomplete Repetitive Stroking with Image Guidance

Yilan Chen
City University of Hong Kong
Hong Kong
yil.ellen.chan@gmail.com

Kin Chung Kwan
University of Konstanz
Germany
kckwan@ieee.org

Hongbo Fu*
City University of Hong Kong
Hong Kong
hongbofu@cityu.edu.hk

Abstract

Image-guided drawing can compensate for the lack of skills 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 users’ normal drawing process. Users can 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 users’ input style when certain repetition is detected. Users can accept, modify, or ignore the system predictions and continue drawing, thus maintaining the fluid control of drawing. Our key idea is to jointly analyze image regions and operation history for detecting and predicting repetitions. The proposed system can effectively reduce users’ workload in drawing repetitive short strokes and facilitates users in creating results with rich patterns.

1. Introduction

Drawing is a common form of artistic expression. By varying the stroke, texture, and shading, artists can create drawings with various styles [6]. Yet, it remains a largely manual process that may require significant artistic expertise and repetitive manual labor.

To reduce repetitive workload, various methods have been proposed to automatically synthesize strokes from user-provided examples [14, 8, 10] or through procedural steps [13]. However, these methods usually perform in a batch and less involve users in the art creation process. Besides, since many methods have predefined the styles and only allow users to tweak a few global parameters, the final results may look monotonous and lack originality (Figure 3). There are also interactive systems [22, 43] which preserve the normal drawing flow while automating significant stroke synthesis. We share a similar goal to theirs. However, they are more

targeted at experts and require sufficient 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, i.e., tracing a reference photo physically via transparent papers or digitally via layers in digital drawing applications. Prior research [42] shows that even when scaffolded with a reference image, people still enjoy freedom for individual expressing. We thus propose to enhance such image-scaffolded drawing process 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 efforts 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 (Figure 2), under the guidance of a reference image. Like 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 the fluid control of drawing. See Figure 1 for an example scenario.

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 [14] and operation history analysis and synthesis [43] while leveraging two key insights. First, since the act of drawing repetitive strokes usually indicates specific intentions (e.g., filling an object or hatching a shading region), we use the common image features among the coherent repetitive strokes to infer the intended regions. Second, the drawing usually relates to the underlying reference image (e.g., the density of strokes with respect to the image lightness). Therefore, we analyze the properties of both the drawing and the reference image to infer possible relationships as contextual constraints for stroke prediction.

We implemented a prototype and conducted a pilot study with participants in different backgrounds to evaluate its

*Corresponding author.

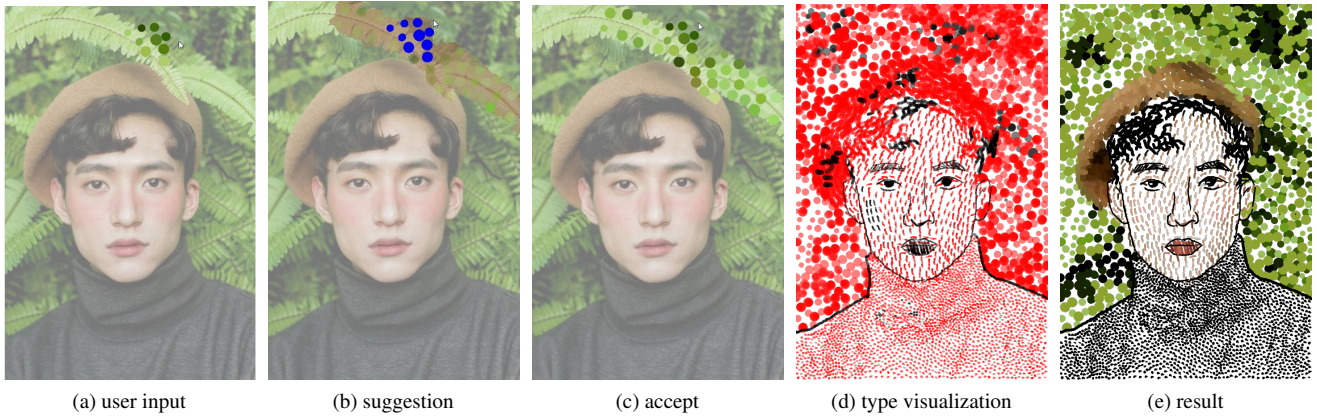


Figure 1: *Example of our system workflow.* A user stipples over a leaf region of a reference image (a) while our system predicts what she might draw next (b) (blue strokes: inferred exemplars; pale red region: inferred target region; semi-transparent strokes: system suggestions), which is then accepted by the user (c) (green strokes: user inputs or accepted suggestions in this scene). (d) visualizes all the manually drawn content in black (261 strokes) and autocompleted content in red (3510 strokes). (e) shows the final result with different repetitive stroke patterns over different regions. Our autocomplete system can reduce tedious repetitive inputs, while being fully under user control.




Figure 2: *Inspiring manual drawings by artists.*




Figure 3: Our work is designed to reduce the workload of completing repetitive patterns during the manual drawing process. The full control of the drawing process leads to more dynamic results than (b) Photoshop’s Art History Brush Tool [1] and (c) StippleShop [28].

utility and usability. The quantitative analysis and qualitative feedback, as well as various drawing results created by users, suggest that our system effectively reduces users’ workload in drawing repetitive short strokes and facilitates users in creating 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 users’ sketches with extracted image features [20, 37, 42, 25], or providing educational guidance to novice users [17, 29, 41]. We share a similar goal to [11, 3, 39] so as to reduce the user workload. However, these works use predefined algorithms to generate strokes along cursor movement and only take users’ input as an indicator of where to render, thus greatly limiting users’ artistic freedom. In contrast, we aim to provide more flexibility between automatic synthesis and manual artistic control by autocompleting tedious repetitions during users’ normal drawing processes.

2.2. Image-based Artistic Rendering

Our work is related to image-based artistic rendering (IB-AR) [23], 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 the image properties (e.g., gradient, edge, color, salience) [13]. Among those methods, the closest to ours are the early

image-based pen-and-ink rendering methods [35, 15], 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 tweak parameters to view the rendered output. In contrast, our system lets users directly specify exemplars on a reference image while silently inferring the distribution properties.

Example-based methods aim to model the visual features of example images for transferring. There are two major modeling approaches: the parametric approach [19, 10, 9] that is based on the summary statistics of stroke characteristics and thus preserves the global textures better, and the non-parametric approach [14, 21, 8] that is based on patch-wise mapping and thus captures the local structures better. We combine both methods for generating strokes: the parametric approach to infer statistical relationships between stroke properties and image features, and the patch-wise matching method to preserve the local arrangements of strokes. Stylit [8] allows users to stylize a rendered ball and simultaneously propagates the style to arbitrary 3D shapes. Our method shares a similar idea in interactive style propagation but with two main differences. First, instead of propagating a style globally, we propagate a style to its perceptually similar local areas so that users can conveniently define different styles in different areas. Second, we represent drawings as discrete stroke operations instead of raster textures for better preserving their structures and enabling procedural editing [36], such as changing the color or size of the drawn strokes.

2.3. Operation History-assisted Authoring

Operation histories [31] have been leveraged in different authoring tasks, such as sketching [43], animation [44, 32], modeling [33, 38], beautification of freehand drawings [7], and handwritings [47]. Our work is most closely related to that by Xing et al.’s [43], which autocompletes repetitive sketching by analyzing the 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 scenario, an operation is an input stroke, so our work is also related to stroke pattern analysis and synthesis [2, 18, 5, 22, 16]. These works disregard the temporal relationship among past strokes and do not use image guidances and thus are different from ours.

To sum up, we list our major differences from the discussed closely related works in Table 1.

3. User Interface

Our system prototype follows a standard digital drawing interface, with the added autocomplete feature, as shown in Figure 4. A user draws on top of the reference image displayed semi-transparently on the main canvas, while our

Table 1: *The differences between our tool and closely related works.* “batch” means the generation is performed in a batch, based on predefined attributes; “dynamic” means the generation is performed based on dynamic operation history. “direct” means users can specify a style by directly operating on the output. “Y” and “N” represent yes and no, respectively, for using image references.

Method	[15]	[14]	[10]	[22]	[43]	Ours
Reference	Y	Y	Y	N	N	Y
Process	batch	batch	batch	batch	dynamic	dynamic
Format	stroke	pixel	stroke	stroke	stroke	stroke
Operate	indirect	indirect	indirect	direct	direct	direct

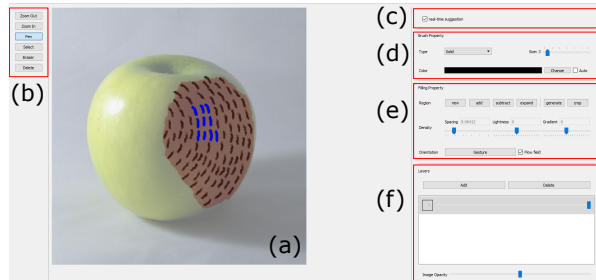


Figure 4: *User interface*, consisting of a central drawing canvas (a), a toolbar for drawing and selection (b), a toggle-switch of the autocomplete mode (c), a brush property toolbar (d), a filling property toolbar (e), and a layers panel (f).




Figure 5: *An example of autocompletion.* The user selects part of the suggestions via the lasso selection tool (a) with the result in (b), continues to draw leading to the updated suggestions (c), and accepts all the suggestions via a hotkey (d). The blue strokes in (a) and (c) indicate inferred exemplars from user-input strokes.

system analyzes the input strokes and the reference image in the background.

3.1. Autocomplete

In the autocomplete mode, our system automatically analyzes whenever the user finishes a new stroke. When a potential repetition is detected, our system highlights the currently repetitive strokes and an inferred propagation region, updates the inferred parameters in the filling property panel, and generates autocompletion suggestions. Users can

accept or reject all the suggestions via hotkeys, accept part of them via lasso selection, or ignore them and continue to draw (Figure 5). The suggestions will keep updating according to user inputs.

3.2. Interactive Editing




Figure 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).




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




Figure 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) A result without any orientation field.

Our system provides a set of tools to refine the auto-completed results.

Propagation region editing. Users can create/add/subtract a new region using the intelligent scissors tool [30] or expand an existing region by a fixed width (Figure 4e) for stroke autocompletion. Figure 6 shows an example of creating a new region for stroke regeneration.

Density editing. Users can tweak 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/gradient, respectively. Our system automatically updates these parameters upon prediction, and the updated parameters provide a starting point for users to manipulate. Figure 7 shows an example.

Orientation editing. Our system automatically predicts whether the input exemplar correlates with the image flow, which can also be tweaked by users manually. Users can also modify the image flow field via the gesture brush, and the touched strokes will be rotated to align with the gesture direction. See Figure 8 for an example.

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

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

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

4. Our Approach

Our system involves two key steps: (1) inferring the input exemplar, the output region, and the contextual constraints from the stroke history and the reference image; (2) synthesizing suggestive strokes accordingly. This section first describes how to *synthesize* (Section 4.1) strokes, assuming all the information is available, and then explains how to *infer* (Section 4.2) the necessary information for synthesis.

4.1. Stroke Synthesis

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 denote the extents of 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 interactions with I in Section 4.2.




Figure 9: (a) A stroke, with centroid p and dominant direction v . (b) The neighborhood of the black stroke includes the n ($n = 1$ in this example) closest strokes (in green) from each quadrant and the middle image patch (blue pixel grid).




Figure 10: *Illustration of our synthesis algorithm.* We synthesize the predicted strokes (in green) from previously drawn strokes (in gray) by matching their neighborhoods as well as the image features.

Idea. To support the autocomplete function with the reference image, we extend the discrete element texture synthesis method [27, 43], which represents strokes as point samples and iteratively improves the sample distribution by minimizing the neighborhood difference between the exemplar and the output, with an additional reference image. First, we combine sample neighborhoods [27] with image features [14] for measuring neighborhood difference. Second, the range and orientation of sample neighborhoods are determined by the radius and orientation maps inferred from the reference image. Figure 10 shows our key idea.

Stroke representation As shown in Figure 9a, 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 auto-completing short strokes, so we represent each stroke by its centroid p and the average direction v for efficiency during synthesis, without considering any other information of the original stroke. To take the 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 the sample points according to the updated centroid and direction.

Initialization. We pre-process the target region mask M by removing the area occupied by existing strokes in the same layer to avoid cluttering, 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 smallest image feature distance d_I , which will be explained in Equation (2). We then optimize the output for a few objectives, as detailed below.

Neighborhood term. We define the neighborhood of a stroke s as 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. [27]) determine the neighboring strokes by spatial distances. Thus, the neighborhood radius should be large enough in order to capture an underlying pattern. However, this might include redundant strokes and thus decrease the performance. Therefore, we adopt Zhao et al.’s method [46] to automatically find a minimum representative neighborhood, considering not only the spatial distance between strokes but also their locations. As depicted in Figure 9b, we set the neighborhood radius of the center stroke s to $2R(s)$. We then divide all the 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, denoted as $\mathbf{N}(s)$. 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 $s' \in \mathbf{N}(s)$, we compute their offsets in position and direction as:

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

where $O(s)^{-1}$ means to rotate the vector by the inverse orientation of $O(s)$. That is, 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 $\{(s'_o, s'_i)\}$ in the neighborhoods $\mathbf{N}(s_o)$ and $\mathbf{N}(s_i)$ using the Hungarian algorithm [27, 26]. We use the norm-2 distance of the offsets from s_o or s_i in Equation (1) as the matching cost. The neighborhood distance is then defined as:

$$d_{neigh}(s_o, s_i) = \sum_{s'_o \in \mathbf{N}(s_o)} |\hat{u}(s'_o, s_o) - \hat{u}(s'_i, s_i)|^2 + \mu \underbrace{|I(s_o) - I(s_i)|^2}_{d_I}, \quad (2)$$

The second term measures the image feature distance d_I , and μ ($= 0.1$ in our implementation) controls the relative weighting. 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:

$$\phi_{neigh}(X, E) = \sum_{s_o \in X} \min_{s_i \in E} d_{neigh}(s_o, s_i). \quad (3)$$

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

$$\phi_{corr}(X) = \sum_{s_o \in X} |p(s_o) - \bar{p}(s_o)|^2. \quad (4)$$

Solver. The energy function we aim to minimize is defined as:

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

We iteratively minimize the energy function following the EM methodology in [27]. In each iteration, for each output stroke s_o , we search for the most matched 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 $w = (i/m)^2$, which means that more weight is given to ϕ_{neigh} in the beginning of iterations, so that we can optimize the neighborhood distribution first before doing corrections, which leads to better results.

Figure 11b - Figure 11d show the iterative optimization process of both the objectives. In comparison, Figure 11e shows the result without the correction term and Figure 11f shows the result without using the image neighborhood in both initialization and optimization.

4.2. Inference

In this section, we describe how to infer E , M , O , and R used for our synthesis method in Section 4.1 from user interactions with I .

4.2.1 Input exemplar E

This step aims 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 [43] and they usually have specific intentions when drawing repetitive strokes, we assume the example strokes to be temporally consecutive and 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 that have similar shape and image features to the last stroke. Specifically, the stroke shape similarity is measured with the Fréchet distance, and 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 [45] at a stroke's center. Alternatively, one can also use different image features for capturing 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. The back-traversal stops when the next stroke does not contain any similar feature or $k > 50$. These k strokes serve as the input exemplar for the synthesis process. See Figure 12 for an example of the incremental searching process.

4.2.2 Output region M

The shared features of the obtained stroke exemplar also indicate the intended region. For instance, if all of the exemplar strokes are 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 there are only two features in our implementation, we simply obtain the region by GrabCut [34] if the *Lab** color feature is shared among the exemplar strokes, directly take the corresponding segmentation if the semantic feature is shared, and take the intersection if both features are shared. See Figure 12 for an example. When there are multiple disconnected regions, we retain the nearest region to the user's last stroke and discard the rest, because it is less natural to propagate to distant regions.

4.2.3 Contextual constraints

Since the drawing usually relates to the underlying reference image, we analyze the properties of both the drawn strokes and the reference image to infer possible relationships that control the global distribution of strokes.

Orientation O . Artists usually adjust the stroke directions to convey curvatures, but they may sometimes randomize or fix the 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) [24] for the reference image and then calculate the angles between the exemplar




Figure 11: *Iteration process in (b) to (d) and ablation studies in (e) and (f). Without the correction term ϕ_{corr} the predicted strokes tend to clutter together as in (e). Without the image term d_I the predicted strokes might not follow the reference sufficiently as in (f).*




Figure 12: *An example of predicting the input exemplar and output region. The left column shows the input stroke sequence visualized in black dots (only a few indices are shown for clarity) on the reference image (top) and the image features (bottom). The right columns show the threshold lines and the image feature cost curves for s_{10} , s_{11} , s_{12} respectively (top), and the corresponding predicted output regions (bottom). 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 will not be used for synthesis.*

strokes and the ETF directions at their centroids. If the standard deviation of the angles is small (less than 15 degrees), 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 to each point of the orientation field.

Radius R . Since density is inversely proportional to the spacing between strokes, we reframe the problem as predicting a radius map that controls the extent 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:

$$r = \begin{pmatrix} l & g & 1 \end{pmatrix} \cdot \mathbf{t}, \quad (6)$$

where \mathbf{t} denotes the coefficients to solve. With 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 as uniform and create a constant radius map with the average spatial distance of the exemplar. We then update the UI with the computed coefficients.

5. Evaluation

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

Autocomplete. Users have full access to our prototype, including autocomplete and interactive editing.

Interactive batch filling. (aka *batch mode*) Users are required to create a texture example first and then manually specify the properties for batch filling. It simulates the sequential procedure in many IB-AR methods (e.g., [35]), although they rarely allow users to directly define examples on target images. This mode is performed on our system with the autocomplete function off.

Fully manual drawing. (aka *manual mode*) Users have 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.

5.1. Target Session

The goal of this session is to compare the three interaction modes in utility and usability. Since we aim to facilitate image-scaffolded drawing, we hope to include general users from different background while focusing more on less skillful users, who are more likely to use reference images. We

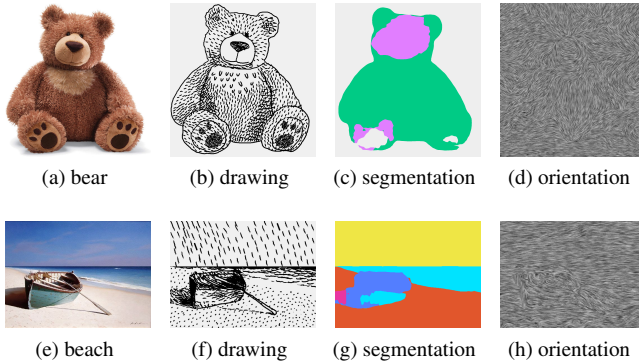


Figure 13: *Target session tasks*. Reference photos in (a) and (e), and the corresponding sample outputs in (b) and (f).

thus recruited 12 participants, including nine novices with little drawing experiences, two amateurs with some experiences (P3, P4), and a student majored in illustration (P5). Most of the studies were conducted on a Lenovo Miix 520 tablet with stylus in a lab environment, except two studies conducted remotely with mouse due to the pandemic.

The study procedure consisted of the following parts and took each participant about two hours in total.

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

Target tasks. We used a within-subjects design, where each participant was asked to reproduce two target drawings (Figure 13) in all the three modes: autocomplete, interactive batch filling, and fully manual drawing. The target drawings include an object and a landscape, which are common illustration topics (e.g., Figure 2). The assigned order of modes was counter-balanced among all the 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 in a dozen of minutes, but without any hard time limit. After completing the two drawings in each mode, each participant filled in a NASA-TLX questionnaire [12]. At the end, we asked the participants about their preferred mode, usage experience and other comments.

5.2. Open session

The goal of this session is to observe users' interaction with our system and learn about users' subjective experience. 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

(Figure 15a) with our system. The reference image was a portrait photo, which is also common in illustrations. The only requirement was that the drawings should contain some repetitive content. We again gave a tutorial in the beginning and conducted the task on a Lenovo Miix 520 tablet with stylus. The participants were encouraged to think aloud and describe their thought process and interactions during this session. After this task, participants could optionally create more drawings with any images they want. Since our prototype does not contain all common functions in commercial drawing tools, we allow the participants to retouch the result drawings without adding more strokes in Photoshop.

5.3. Results and Observations

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

Efficiency. We calculate the average completion time (Figure 14b) and stroke count (Figure 14c) in each mode and each task. Generally, the system synthesized about 82% strokes in the autocomplete mode and about 92% strokes in the batch mode. Although the manual mode took the shortest time for the participants to complete, it also resulted in the fewest total number of strokes. We thus calculated the strokes per minute for each mode: autocomplete (111.03, $SD=38.76$), batch (101.98, $SD=45.13$), manual (115.95, $SD=46.73$). It turns out automatic generation did not improve the 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 for exploring alternative patterns, which, however, might be high in a fully manual case.

Quality. We asked 30 external volunteers to evaluate the quality of participants' drawings, as shown in the supplementary material. We randomized all the 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

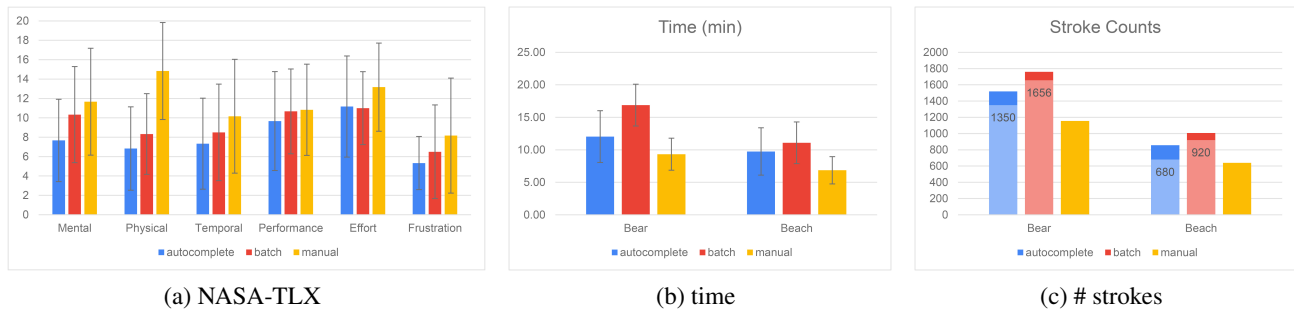


Figure 14: *Target sessions results.* (a) Average NASA-TLX scores from 12 participants. The lower the better. (b) Average completion time. (c) Average stroke counts. The number of system-generated strokes is labeled in each column.

shapes. We calculated the average scores for each mode: autocomplete (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 is slightly better than the fully manual drawings, but without significant difference. From the participants’ perspective, three novices commented the automated strokes were better than their manual strokes, because they tend to lose patience when manually drawing all strokes, which results in worse quality.

Preferred mode. Seven participants preferred the autocomplete mode while the rest five participants preferred the batch mode. Generally, the autocomplete mode is considered more convenient, yet less precise; the batch mode is considered more precise, but requires too many interactions. P12 commented, “*the autocomplete mode is more straightforward, because you can see the filled effects instantly without doing a lot of manipulation beforehand; while in the batch mode, you have to remember the meaning of parameters and tweak them in order to create strokes.*” P10 also said, “*Compared with batch filling, the autocomplete mode provides a quick guess of filled regions and allows me to get the results more quickly with less work.*” However, the autocomplete mode is “*less accurate at some vague and detailed regions, such as the shadows of the boat, where it tends to include some unwanted regions, so I have to manually subtract those regions, which is a bit tedious*”, according to P3. The professional, P5, also preferred the batch mode for being able to precisely select the regions. Therefore, we consider the autocomplete function and the interactive editing function are complementary in usability.

Creation results and experience. Figure 15 shows the 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 demonstrates some sample results. Regarding the creation experience, one user said “*it is playful, the final result is*




Figure 15: *Example drawing results from the open session.* Each case is marked with the # of manual and **autocompleted** strokes.

also good”; two users described it as “*encouraging*”, because the system allows beginners to quickly create stylistic drawings; one user commented that she “*felt creative when drawing with this system*”, because she could test 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 can be helpful in adding interesting textures into color paintings (e.g., Figure 16i). Two users commented that the reduction of workload is useful, but they also complained about some inaccurate inference of autocomplete. We will discuss about this problem in Section 7.

6. Conclusion

We have presented a new drawing concept and designed an assisted drawing system to help users autocomplete repeti-




Figure 16: *Sample results.* In each example, the left column shows the reference images, the middle column visualizes the manual strokes (black) and auto-completed results (red) of the final drawings on the right column. In the last example, the strokes are created with our system first and then imported into Photoshop for background coloring. Please check the supplementary material for more results.

tive 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 that adapt 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.

7. Limitations and Future Work

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

Improve accuracy of autocompletion. We rely on simple *Lab** color and semantic segmentation for region inference. While color feature is sufficient for most cases, regions with similar colors but different semantics will require sufficient segmentation accuracy for region inference (e.g., Figure 13c). More advanced semantic selection methods (e.g., [4]) might help inferring more accurate regions. However, the 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 the limbs? We left this as a future work.

Resolve visual blocking.

Since the drawing and the system suggestions are overlaid on the reference image, it might be difficult for users to refer to the image when selecting parts of the suggestions (e.g., Figure 17) or adding a new layer of strokes. Although users can switch the views via a hotkey, it might be helpful to provide some reference information, like image darkness or boundaries, through additional visual hints [42, 41].




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

Consider relationships with higher-level image features.

We only consider the relationships between strokes and low-level image features, like colors and flows, over regions. By considering higher-level image features, such as elements and edges, it is possible to extend the scope of autocompletion, such as autocompleting the sparse flowers in the foreground of Figure 16i through the correspondences between strokes and elements.

Support more stroke types. Our method only supports short strokes, while artists also use long repetitive strokes frequently [6]. It is worth investigating the possibility of

incorporating continuous strokes [40] in our framework and extending the support for different input strokes.

Acknowledgement

We thank the anonymous reviewers for feedback, and funding from Adobe Research and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)-Project-ID 251654672-TRR 161.

References

- [1] Adobe. Paint stylized strokes with the art history brush. <https://helpx.adobe.com/photoshop/using/painting-stylized-strokes-art-history.html>, 2017. 2
- [2] P. Barla, S. Breslav, L. Markosian, and J. Thollot. Interactive Hatching and Stippling by Example. Research Report RR-6461, INRIA, 2006. 3
- [3] L. Benedetti, H. Winnemöller, M. Corsini, and R. Scopigno. Painting with bob: Assisted creativity for novices. In *UIST '14*, pages 419–428, New York, NY, USA, 2014. ACM. 2
- [4] X. Chen, Z. Zhao, F. Yu, Y. Zhang, and M. Duan. Conditional diffusion for interactive segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7345–7354, 2021. 11
- [5] V. A. dos Passos, M. Walter, and M. C. Sousa. Sample-based synthesis of illustrative patterns. In *2010 18th Pacific Conference on Computer Graphics and Applications*, pages 109–116, Sep. 2010. 3
- [6] A. Dunn. *Pen and Ink Drawing: A Simple Guide*. Three Minds Press, 2015. 1, 11
- [7] J. Fišer, P. Asente, and D. Šýkora. Shipshape: A drawing beautification assistant. In *Proceedings of the Workshop on Sketch-Based Interfaces and Modeling*, SBIM '15, pages 49–57, Goslar Germany, Germany, 2015. Eurographics Association. 3
- [8] J. Fišer, O. Jamriška, M. Lukáč, E. Shechtman, P. Asente, J. Lu, and D. Šýkora. Stylist: Illumination-guided example-based stylization of 3d renderings. *ACM Trans. Graph.*, 35(4):92:1–92:11, July 2016. 1, 3
- [9] L. A. Gatys, A. S. Ecker, M. Bethge, A. Hertzmann, and E. Shechtman. Controlling Perceptual Factors in Neural Style Transfer. *ArXiv e-prints*, Nov. 2016. 3
- [10] M. Gerl and T. Isenberg. Interactive example-based hatching. *Comput. Graph.*, 37(1-2):65–80, Feb. 2013. 1, 3
- [11] P. Haeberli. Paint by numbers: Abstract image representations. In *SIGGRAPH '90*, pages 207–214, 1990. 2
- [12] S. G. Hart and L. E. Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. In *Advances in psychology*, volume 52, pages 139–183. Elsevier, 1988. 8
- [13] S. Hegde, C. Gatzidis, and F. Tian. Painterly rendering techniques: a state-of-the-art review of current approaches. *Computer Animation and Virtual Worlds*, 24(1):43–64, 2013. 1, 2
- [14] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin. Image analogies. In *SIGGRAPH '01*, pages 327–340, 2001. 1, 3, 5

- [15] S. Hiller, H. Hellwig, and O. Deussen. Beyond stippling—methods for distributing objects on the plane. In *Computer Graphics Forum*, volume 22, pages 515–522. Wiley Online Library, 2003. 3
- [16] C.-Y. Hsu, L.-Y. Wei, L. You, and J. J. Zhang. Autocomplete element fields. In *CHI '20*, pages 1–13, 2020. 3
- [17] E. Iarussi, A. Bousseau, and T. Tsandilas. The drawing assistant: Automated drawing guidance and feedback from photographs. In *UIST '13*, pages 183–192, 2013. 2
- [18] T. Ijiri, R. M ech, T. Igarashi, and G. Miller. An example-based procedural system for element arrangement. *Computer Graphics Forum*, 27(2):429–436, 2008. 3
- [19] E. Kalogerakis, D. Nowrouzezahrai, S. Breslav, and A. Hertzmann. Learning hatching for pen-and-ink illustration of surfaces. *ACM Trans. Graph.*, 31(1):1:1–1:17, Feb. 2012. 3
- [20] H. W. Kang, W. He, C. K. Chui, and U. K. Chakraborty. Interactive sketch generation. *The Visual Computer*, 21(8):821–830, Sep 2005. 2
- [21] A. Kaspar, B. Neubert, D. Lischinski, M. Pauly, and J. Kopf. Self Tuning Texture Optimization. *Computer Graphics Forum*, 2015. 3
- [22] R. H. Kazi, T. Igarashi, S. Zhao, and R. Davis. 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, CHI '12*, pages 1727–1736, 2012. 1, 3
- [23] J. E. Kyprianidis, J. Collomosse, T. Wang, and T. Isenberg. State of the “Art”: A Taxonomy of Artistic Stylization Techniques for Images and Video. *IEEE Transactions on Visualization and Computer Graphics*, 19(5):866–885, 2013. 2
- [24] J. E. Kyprianidis and H. Kang. Image and video abstraction by coherence-enhancing filtering. *Computer Graphics Forum*, 30(2):593–602, 2011. Proceedings Eurographics 2011. 6
- [25] G. Li, S. Bi, J. Wang, Y. Xu, and Y. Yu. Colorsketch: A drawing assistant for generating color sketches from photos. *IEEE Computer Graphics and Applications*, 37(3):70–81, May 2017. 2
- [26] C. Ma, L.-Y. Wei, S. Lefebvre, and X. Tong. Dynamic element textures. *ACM Trans. Graph.*, 32(4):90:1–90:10, July 2013. 5
- [27] C. Ma, L.-Y. Wei, and X. Tong. Discrete element textures. *ACM Trans. Graph.*, 30(4):62:1–62:10, July 2011. 5, 6
- [28] D. Mart n, G. Arroyo, A. Rodr guez, and T. Isenberg. A survey of digital stippling. *Computers & Graphics*, 67:24–44, Oct. 2017. 2
- [29] Y. Matsui, T. Shiratori, and K. Aizawa. Drawfromdrawings: 2d drawing assistance via stroke interpolation with a sketch database. *IEEE transactions on visualization and computer graphics*, 23(7):1852–1862, 2017. 2
- [30] E. N. Mortensen and W. A. Barrett. Intelligent scissors for image composition. In *SIGGRAPH '95*, pages 191–198, 1995. 4
- [31] M. Nancel and A. Cockburn. Causality: A conceptual model of interaction history. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems, CHI '14*, pages 1777–1786, New York, NY, USA, 2014. ACM. 3
- [32] M. Peng, L.-Y. Wei, R. H. Kazi, and V. G. Kim. Autocomplete animated sculpting. In *UIST '20*, 2020. 3
- [33] M. Peng, J. Xing, and L.-Y. Wei. Autocomplete 3d sculpting. *ACM Trans. Graph.*, 37(4), Aug. 2018. 3
- [34] C. Rother, V. Kolmogorov, and A. Blake. “grabcut”: Interactive foreground extraction using iterated graph cuts. In *ACM SIGGRAPH 2004 Papers, SIGGRAPH '04*, page 309–314, New York, NY, USA, 2004. Association for Computing Machinery. 6
- [35] M. P. Salisbury, M. T. Wong, J. F. Hughes, and D. H. Salesin. Orientable textures for image-based pen-and-ink illustration. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97*, pages 401–406, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co. 3, 7
- [36] M. Schwarz, T. Isenberg, K. Mason, and S. Carpendale. Modeling with rendering primitives: An interactive non-photorealistic canvas. In *Proceedings of the 5th International Symposium on Non-photorealistic Animation and Rendering, NPAR '07*, pages 15–22, New York, NY, USA, 2007. ACM. 3
- [37] Q. Su, W. H. A. Li, J. Wang, and H. Fu. Ez-sketching: Three-level optimization for error-tolerant image tracing. *ACM Trans. Graph.*, 33(4):54:1–54:9, July 2014. 2
- [38] R. Suzuki, K. Yatani, M. D. Gross, and T. Yeh. Tabby: Explorable design for 3d printing textures. In *PG '18 Short Papers*, pages 29–32, Goslar, DEU, 2018. Eurographics Association. 3
- [39] H.-C. Tsai, Y.-H. Lee, R.-R. Lee, and H.-K. Chu. User-guided line abstraction using coherence and structure analysis. *Computational Visual Media*, 3(2):177–188, 2017. 2
- [40] P. Tu, L.-Y. Wei, K. Yatani, T. Igarashi, and M. Zwicker. Continuous curve textures. *ACM Trans. Graph.*, 39(6), Nov. 2020. 11
- [41] B. Williford, A. Doke, M. Pahud, K. Hinckley, and T. Hammond. Drawmyphoto: Assisting novices in drawing from photographs. In *Proceedings of the 2019 on Creativity and Cognition, C&C '19*, pages 198–209, 2019. 2, 11
- [42] J. Xie, A. Hertzmann, W. Li, and H. Winnem ller. Portraitsketch: Face sketching assistance for novices. In *UIST '14*, pages 407–417, 2014. 1, 2, 11
- [43] J. Xing, H.-T. Chen, and L.-Y. Wei. Autocomplete painting repetitions. *ACM Trans. Graph.*, 33(6):172:1–172:11, Nov. 2014. 1, 3, 5, 6
- [44] J. Xing, L.-Y. Wei, T. Shiratori, and K. Yatani. Autocomplete hand-drawn animations. *ACM Trans. Graph.*, 34(6):169:1–169:11, Oct. 2015. 3
- [45] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *CVPR*, 2017. 6
- [46] M. Zhao and S.-C. Zhu. Customizing painterly rendering styles using stroke processes. In *NPAR '11*, pages 137–146, 2011. 5
- [47] C. L. Zitnick. Handwriting beautification using token means. *ACM Trans. Graph.*, 32(4), July 2013. 3