

Fast Accurate and Automatic Brushstroke Extraction

YUNFEI FU, National Centre for Computer Animation, Bournemouth University, UK

HONGCHUAN YU, National Centre for Computer Animation, Bournemouth University, UK

CHIH-KUO YEH, Department of Computer Science and Information Engineering, National Cheng-Kung University, Taiwan

TONG-YEE LEE, Department of Computer Science and Information Engineering, National Cheng-Kung University, Taiwan

JIAN J. ZHANG, National Centre for Computer Animation, Bournemouth University, UK

Brushstrokes are viewed as the artist's "handwriting" in a painting. In many applications such as style learning and transfer, mimicking painting, painting authentication, it is highly desired to quantitatively and accurately identify brushstroke characteristics from old masters' pieces using computer programs. However, due to the nature of hundreds or thousands of intermingling brushstrokes in the painting, it still remains challenging. This paper proposes an efficient algorithm for brush Stroke extraction based on a Deep neural network, i.e. DStroke. Compared to the state of the art research, the main merit of the proposed DStroke is to automatically and rapidly extract brushstrokes from a painting without manual annotation, while accurately approximating the real brushstrokes with high reliability. Herein, recovering the faithful soft transitions between brushstrokes is often ignored by the other methods. In fact, the details of brushstrokes in a master piece of painting, e.g. shapes, colors, texture, overlaps, are highly desired by artists since they hold promise to enhance and extend the artists' powers, just like microscopes extend biologists' powers. To demonstrate the high efficiency of the proposed DStroke, we perform it on a set of real scans of paintings and a set of synthetic paintings respectively. Experiments show that the proposed DStroke is noticeably faster and more accurate at identifying and extracting brushstrokes, outperforming the other methods.

CCS Concepts: • **Computing methodologies** → **Image manipulation**.

Additional Key Words and Phrases: Brushstroke extraction, Painting authentication, Hard and Soft segmentation, Pix2Pix network

ACM Reference Format:

Yunfei Fu, Hongchuan Yu, Chih-Kuo Yeh, Tong-Yee Lee, and Jian J. Zhang. 2018. Fast Accurate and Automatic Brushstroke Extraction. 1, 1 (October 2018), 26 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

In the recent years, paintings processing has received wide attentions in computer vision and graphics fields, such as animating paintings [37], style transfer [18, 38], mimicking paintings [47, 39]. Particularly, deep learning technology is

Authors' addresses: Yunfei Fu, fuyunfei1991@gmail.com, National Centre for Computer Animation, Bournemouth University, Fern Barrow, Poole, UK; Hongchuan Yu, hyu@bournemouth.ac.uk, National Centre for Computer Animation, Bournemouth University, Fern Barrow, Poole, UK; Chih-Kuo Yeh, Department of Computer Science and Information Engineering, National Cheng-Kung University, Taiwan, City 701, Tainan, Taiwan, simpson.ycg@gmail.com; Tong-Yee Lee, Department of Computer Science and Information Engineering, National Cheng-Kung University, Taiwan, City 701, Tainan, Taiwan; Jian J. Zhang, National Centre for Computer Animation, Bournemouth University, Fern Barrow, Poole, UK.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1




Fig. 1. Paintings with highly mixed brush stroke in similar colors (left); Extracted brushstrokes by [7] (middle, many strokes are missed and are classified in dark). Extracted brushstrokes by our DStroke (right, the soft transitions between strokes can be noted).

boosting this research. There have been several emerging AI created paintings available, which falls in the category of stroke-based rendering [9, 10, 36, 12], i.e. creating non-photorealistic imagery through placing discrete elements such as paint strokes or stipples. However, for learning styles and mimicking paintings, it usually requires strokes to be extracted from masters' pieces in advance. These applications indeed demand an old masters' brushstroke database for rendering purposes. A rising challenge is to automatically and accurately extract brushstrokes from a brush painting, which inspires the research work of this paper.

In some circumstances, computer programs can extract certain patterns from scans more thoroughly than manual attempts, process a much larger number of paintings, and be less subjective. However, despite encouraging results from research groups, none are perfect. For instance, current methods cannot automatically extract all the brushstrokes from a painting, and therefore manual input is required, which is a tedious and time-consuming task due to the nature of hundreds or thousands of brushstrokes intermingling with each other in the painting.

To tackle the rising challenge, this paper aims to develop an efficient algorithm that can automatically and correctly detect close to all the brushstrokes on a brush painting and accurately approximate them. The proposed Deep neural network based brush Stroke extraction (DStroke) method employs deep neural network and image matting techniques (i.e. foreground and background extraction problems) to brushstroke extraction and accurately approximates the real brushstrokes. The challenging problem we encounter is to deal with the scenario of brushstrokes overlapping in a painting. Compared to the existing representative work [7, 25, 41], the proposed DStroke can clearly identify close to all the brushstrokes, whilst accurately acknowledging the faithful soft transitions between brushstrokes (see the right of Fig.1). This is significant for image segmentation because brushstrokes similar in color due to either overlapping or adjacent to each other (see the left of Fig.1) tend to be incorrectly classified as one (see the middle of Fig.1), thus requiring human correction. This also sums up the main limitations of [7, 25, 41]. To the best of our knowledge, the other stroke extraction methods don't take into account overlapped strokes. Additionally, for visualization purpose, we label the extracted brushstrokes in colors (see the middle and right of Fig.1). Particularly, to highlight the soft transitions of brushstrokes, the color values depend on the individual alpha values, which results in the blurred boundaries of strokes as the right of Figure 1. The blurred boundaries indeed represent soft transitions.

Our research is inspired by the state-of-the-art work, i.e. Pix2Pix network [17] and Semantic Soft Segmentation [2], both of which facilitate semantic segmentation. Likewise, the other GAN neural network based semantic segmentation methods [40, 6, 11, 15], Pix2Pix [17] cannot also be directly applied to brushstroke extraction due to the following facts:

- There is no brushstroke training dataset available for deep neural network training currently. Manually annotating strokes on a painting is a tedious task making it difficult to build up a large training dataset for deep learning purposes.
- The current semantic segmentation methods always fail to segment overlapping strokes on a painting since all the strokes may share the same class.
- The current segmentation methods are not suited to deal with a dense labelling problem since labels are discrete valued. Unfortunately, extracting hundreds or thousands of brushstrokes from a painting is a dense labelling problem.

We modify the Pix2Pix network to cope with these difficulties in this paper. Aksoy et al. [2] presented the other implementation of semantic segmentation on soft transitions between image regions, which embeds texture and color features from the image as well as high-level semantic information generated by a neural network. The soft segments are generated through eigen-decomposition of the Laplacian matrix. Nevertheless, this method suffers from at least three limitations which results in the failure of soft segmentation on the application of brushstroke extraction:

- It only supports a small number of segments, and the segment number is usually fixed. This is not suitable for segmenting hundreds or thousands of strokes from a painting.
- It does not support instance-level segmentation, since there is no instance-aware semantic information available. This is not suitable in dealing with the scenario of overlapping strokes of a similar color.
- The spectral decomposition of the matting Laplacian is very expensive, decreasing the efficiency of the algorithm.

In contrast, the proposed DStroke overcomes these deficiencies. Our research focuses on the efficiency of the algorithm, i.e. computational complexity, accuracy and automation. The main contributions include:

- A large, automatically generated painting training dataset. To the best of our knowledge, we are the first to provide an automatic method to build up a large painting sample dataset for deep learning purpose;
- The proposed DStroke method supports soft segmentation on instance level to identify every brushstroke and recover the faithful soft transitions between them;
- The modified Pix2Pix network can output the labels of segmentation without limitation on segment number, which are exactly discrete values.

Moreover, the proposed DStroke method can fulfil soft segmentation for fuzzy boundaries by involving the guided filter [13], which is a linear time algorithm. In fact, the efficiency and accuracy of guided filter entirely depend on a given guidance image. A merit of DStroke is to rapidly generate accurate guidance images required by the guided filter.

Additionally, an increasing number of problems in the history of art, particularly authenticating and dating brush paintings by the masters, have received wide attention, and many rigorous computer methods have been developed through significant interdisciplinary efforts across computer vision, graphics and art history in recent year. Herein, brushstroke extraction plays an important role. This is because brushstrokes can be employed to assess the level of distinction between categories of paintings and identifying attributes that differ significantly on average [25, 14, 16, 31, 35, 4, 41]. Moreover, identifying the order of brushstrokes and their individual colors is important as well for authentication purpose. In fact, it is not only a request from painting authentication, but also is the basis of many

existing layer decomposition methods [7, 30, 37, 34, 33]. However it still remains challenging since overlapped strokes result in mixed color and transparency change leads to blur. This paper demonstrates that the DStroke method can automatically detect close to all the brushstrokes whilst accurately recovering the fragile soft transitions between strokes, has high efficiency, and also work well on most brush paintings such as oil brush paintings, Chinese paintings, watercolor and acrylic paintings.

2 RELATED WORK

Layer Decomposition: In digital image editing systems, artists deposit colors throughout a painting via a set of strokes, which are classified into different layers in terms of opacity values. Skilled artists commonly blend multiple layers, each of which is composed of simple colors and transparency gradients, to represent an object. However, the scans of paintings and photographs have no such layer information. Even for digital paintings, the available image files usually lack layer information as well. Without layer information, even simple editing may become very challenging. Our previous work [7] attempts to employ layer decomposition and boundary constraints to stroke extraction. Richardt et al. [30] present an interactive approach for decomposing bitmap drawings and photographs into opaque and semi-transparent vector layers. Xu et al. [37] aim to decompose Chinese paintings into a collection of layered brushstrokes with an assumption that an overlapping region contains at most two strokes and has minimal variation in transparency. Moreover, their approach requires the knowledge of the order of strokes and a brushstroke library for recognition, which is built by professional artists. McCann et al. [26, 27] present two generalized layer decomposition methods, in which pixels have individual layers and partially overlap with each other, the layer orderings of which may be manipulated. Aharoni-Mack et al. [3] propose a recoloring painting method for watercolor paintings based on color palette estimation and layer decomposition. Tan et al. [34] present a layer decomposition method based on RGB-space geometry. An assumption is that all possible image colors are convex combinations of the palette colors. Computing the convex hull of image colors and per-pixel layer opacities is converted into a convex optimization problem. Thus, their method can work well without prior knowledge of shape and even with some overlapping strokes. Furthermore, Tan et al. [33] proposed a palette-based layer decomposition algorithm. The distinct advantage is to require no numerical optimization and allow users to interactively edit the palette to adjust the layers. Koyama et al. [19] argued that the previous methods typically only support linear color-blend modes, and further proposed a layer decomposition method to support any user-specified color-blend modes. However, these methods still cannot accurately extract all the strokes from a painting. The proposed DStroke method will tackle this challenge.

Soft Segmentation: For oil paintings such as van Gogh’s artefacts, most of brushstrokes tend to be opaque, which benefits edge detection. However the overlapped parts tend to be missed. Li et al. [25] and Lamberti et al. [41] employ the seed growing-based brushstroke extraction scheme to painting authentication and artist identification. Firstly, some pixels are selected as seeds. Then, neighboring pixels are exploited through region growing. Region growing can be controlled through a shape validation method. In their implementation, the metric of shape validation is defined based on 10 examples of brushstroke regions that are manually extracted from van Gogh’s paintings. When the boundaries of objects are vague, or if objects are translucent, e.g. water, hair, brushstrokes with transparency, traditional image segmentation methods begin to fail. Even if manual segmentation is taken, it is no longer reliable. Our proposed DStroke method uses soft segmentation technology to tackle this challenge.

Compared to usual segmentation [46], the key feature of soft segmentation is that pixels may be assigned to the foreground or background in terms of transparency values. Thus, it can capture the fuzzy boundary between objects. The core is image matting. Matting aims to estimate the per-pixel transparency of the foreground region based on the

indicators that users provide [23, 5, 1]. The indicators are typically represented using trimap, in which the foreground, background, and unknown transparent regions are distinguished using different colors. Instead of identifying the foreground objects from the image, soft segmentation decomposes an image into multiple layers. As a result, each pixel is likely to be classified into multiple layers. Singaraju et al.[32] presented an approach of segmenting an image into multiple layers by the estimation of alpha mattes, which need to be optimized iteratively. The spectral matting technique in [24] extended spectral segmentation techniques [23], from the extraction of only hard segments, to include the extraction of soft matting components.

However, these soft segmentation methods cannot generate semantically meaningful segmented regions without user indications. Recently, Aksoy et al.[2] leveraged deep network for semantic soft segmentation, using the high-level information (semantic information) from a deep network to define affinities between different regions to generate soft segmentation corresponding to semantic regions. On the other hand, low-level information such as color and matting affinity, are also calculated to construct the matting Laplacian matrix to handle soft boundaries. Like spectral matting, regions with local soft transitions are generated using matting Laplacian and spectral decomposition. However, when applied to brushstroke extraction, their methods suffer at least two limitations. Firstly, a brush painting normally contains hundreds or even thousands of brushstrokes. Unfortunately the methods in [24, 2] only support a limited number of segmented regions. Secondly, despite using high-level information from a deep network, the semantic soft segmentation in [2] cannot do instance-level segmentation. This will bring about serious errors for brush paintings since hundreds, or thousands of instances (brushstrokes) in one painting may share the same semantic label. In contrast, our proposed DStroke can handle thousands of brushstrokes in a given brush painting and can effectively generate instance-level segmentations corresponding to semantic regions.

Deep learning based segmentation: Krizhevsky et al.[20] introduce a deep convolution network called AlexNet consisting of 8 layers and millions of parameters which was trained on the ImageNet dataset with 1 million images. Since then, even larger convolutional networks have been designed and competed with the state of the art in image segmentation, including [40, 6, 17, 11, 45]. U-Net [29] was proposed for biomedical image segmentation, which provides a pixel-wise accuracy for dense cells segmentation. Deep convolution neural networks were used to learn through minimizing a loss function. It is crucial to design an appropriate loss function. For instance, if simply using the Euclidean distance between the predicted and ground truth pixels as loss, the results may tend to blur the boundaries. Unceasingly involving expert knowledge into loss functions can improve the performance of deep networks. The discriminator network is combined with U-Net in the Pix2Pix network [17], which can automatically learn a loss function appropriate for satisfying this goal. It dramatically improves the accuracy of segmentation which is employed in our method.

The main challenge currently is the lack of a training dataset for brushstrokes extraction. Deep neural networks need a large dataset for effective training. Manually extracting brushstrokes is impractical since a painting usually contains thousands of brushstrokes. To deal with this challenge, we propose an automatic method to build a large training dataset, in which brush paintings consist of a given set of brushstrokes and also contain information of the brushstrokes' edge maps. This addresses the challenges of manually constructing a large dataset for machine learning purpose.

3 DEEP NETWORK BASED BRUSH STROKE EXTRACTION (DSTROKE) METHOD

We focus on the efficiency of the proposed DStroke method, which is illustrated in Figure 2. The first problem is to automatically create a large painting training dataset for deep neural network training purpose. To the best of our knowledge, we are the first to propose an automatic method of building up a large training dataset of brush paintings.

The modified version of the Pix2Pix network [17] is applied to brushstroke extraction for edge map detection (also called as hard segmentation). We prefer instance-level segmentation on paintings and therefore employ the edge maps of sample paintings as instance-level information to neural network training. After that, the guided filter [13] is employed to soft segmentation, i.e. identifying soft transitions between brushstrokes to refine the boundaries of strokes, which is a linear time algorithm and drastically improve the efficiency of DStroke.

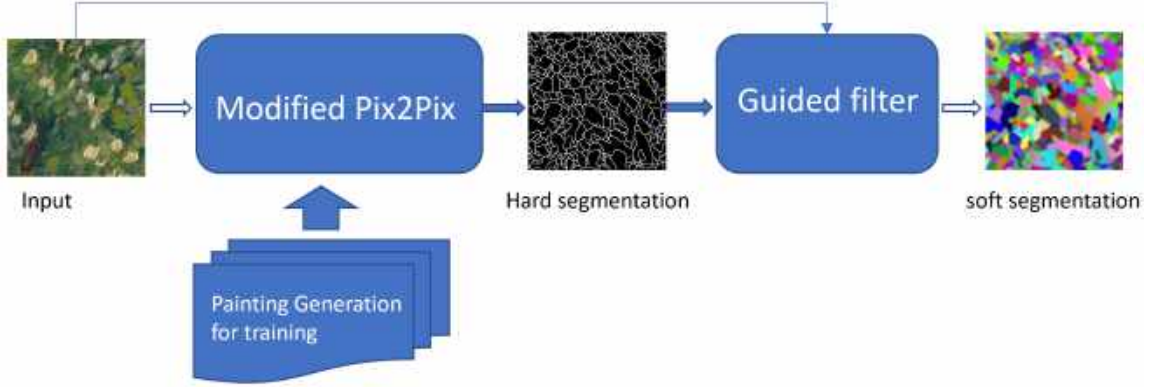


Fig. 2. Overview of the proposed DStroke.

3.1 Training Database Generation

We address the automatic painting generation method at first and then address how to build up a large painting training dataset. Brushstrokes on a painting are usually similar. It tends to classify multiple brushstrokes into one class. To identify every brushstroke from a painting, the training sample dataset is required to contain a set of paintings associated with the individual edge maps of brushstrokes in order to provide instance-level information. The edge map is the collection of the brushstrokes' boundaries on the painting.

Our painting generation method is to convert photos to paintings through a set of given brushstrokes. Although there have been many methods for painting production, they do not usually provide brushstroke information. In our implementation, when painting brushstrokes, the boundaries are cumulated to form the edge map of brushstrokes on the resulting painting.

Moreover, for reality purpose, the brushstroke samples are required to maintain the illumination effect of paintings. Usually, real brush paintings have the individual height maps since the pigment associated with each brushstroke on a canvas may be layered to have differing thicknesses. The photo or image of a painting is the realistic appearance of the canvas surface with plausible lighting. Thus, every brushstroke sample is assigned to the individual height map (e.g. let alpha map as height map), and the height field of the created painting is generated by rendering the brushstrokes textured with the height maps. The final painting is rendered by the input of image colors alongside the height maps. Algorithm 1 addresses how to create a painting by mimicking the physical appearance of brushstrokes.

In our implementation, the canvas C is firstly initialized as a blank image plane, and the region of interest (ROI) is selected in terms of the pixel with the largest color difference between the sample image and the canvas. The gradient within this pixel's neighborhood on the sample image is computed as the ROI's gradient. Secondly, a stroke is randomly picked up from the small stroke library-kylebrush [44] and put over the ROI on the canvas. The orthogonal direction

Algorithm 1: Painting Generation

```

Input:  $S$  sample image;  $SA$  stroke alpha map;  $ST$  stroke thickness/height map;  $C$  blank canvas;
Output:  $C$  painted canvas;  $E$  edge map;
Initialising:  $Diff = \infty$  color difference;  $H=0$  height field;  $R=0$  the ratio of painted area over the canvas area;
while  $R < 1$  do          /* may change  $R$  to threshold  $\epsilon < \|S - C\|$  for a pleasing visual effect */
   $Diff_1 = S - C$ 
   $ROI = \arg \max_{i,j} |Diff_1(i, j)|$  /* region of interest ( $ROI$ ) is a neighborhood of the pixel  $(i, j)$  */
  Select a stroke from library;
   $color = avg(ROI)$  /* average color of  $ROI$  on  $S$  is viewed as the selected stroke's color */
   $C = compose(C, Stroke, Color)$  /* paint the selected stroke on canvas with color */
   $Diff_2 = S - C$  /* update color difference */
  if  $\|Diff_1\| > \|Diff_2\|$  then
    Keep the selected stroke on  $C$ ;
    Save the stroke boundary on  $E$ ;
    Remove the overlapped edge from  $E$ ;
     $H = H + ST$  /* add stroke's height map to height field */
     $R = \frac{\text{painted region}}{\text{canvas}}$  /*  $R \in [0, 1]$ , the painted regions will gradually spread over the whole canvas */
  else
    Remove the selected stroke from  $C$ ;
  endif
end while
 $C = Rendering(C, H)$ ;

```

to the gradient usually indicates edges. We compose the stroke along with the orthogonal direction on the canvas. After that, the new ROI is detected over the updated canvas and is overlapped by a new selected stroke from the stroke library until the original canvas is thoroughly covered by strokes. The height map ST of a brushstroke sample is set equal to its alpha map (see Fig.4). The height field H of the canvas is cumulative and updated each time a stroke is painted onto the canvas. The normal of pixels on the canvas is calculated by the directional derivative of the height field. The illumination of each pixel is then calculated under the different illumination models in the rendering step, e.g. the reflection model may be Phong model or Cook-Torrance model [43].

Figure 3 shows the process of creating a brush painting through a set of given brushstrokes, particularly Fig.3c which shows the effect from the height field of the painting. Herein the goals are to make the painting look like the sample image and limit the number of strokes in some way to make the result look like a painting.

To this end, the terminal criteria include the ratio R of the painted area over the whole canvas area and the differences $Diff$ between the sample image and the rendering. $Diff$ is minimized in a trail-and-error way, i.e. if the change reduces the cost function, the change is incorporated; otherwise, it is discarded. Using R can limit the number of strokes. If pursuing a pleasing visual effect (see Fig.3d), we can simply change the ratio R to the threshold ϵ of the $Diff$'s error (see Algorithm 1). However this will require a longer running time.

We further address how to build up a large training dataset using our automatic painting generation method. In deep learning applications, the training sample set is acquired in a limited set of conditions. But the application may exist in a variety of conditions. It is natural to add synthetically modified data into the training dataset to account for these new situations, which is called data augmentation. To this end, the created paintings in our training dataset are randomly flipped, rotated, and slightly distorted; noise is added (e.g. Gaussian noise, salt and pepper noise) to generate new training samples. We also shift the colors of the created paintings by randomly changing hues for new training




Fig. 3. Process of painting production. (a) Scene image. (b) Iteratively painting strokes on canvas (Top row); height map of canvas (bottom row); (c) Illumination of canvas; (d) Effect image using the threshold ϵ , which looks more like (a).



Fig. 4. The alpha maps of different brushstrokes from the kylebrush Library[44]

samples. Figure 5 shows such synthetic paintings. As the ratio R is employed rather than the threshold ϵ , the running time is acceptable; that is, Algorithm 1 spends around 2 hours on producing 1000 paintings for our training dataset. Regarding the core of the training dataset generation, i.e. Algorithm 1, we hope to point out the following distinct advantages when compared with the existing “painterly rendering methods” such as [38, 8, 9, 10, 18],




Fig. 5. Illustration of data augmentation. (a) Input painting, (b) adding noise, (c) distortion, (d) changing hue, (e) flipping

- Algorithm 1 can better avoid over-fitting issues in training by utilizing many stroke types for training instead of only a few, e.g. using more than 500 stroke types in our implementation. Moreover, the data augmentation is also applied to increasing the variety of data.
- Utilizing the height maps of brushstrokes, output paintings are rendered under different light settings to further avoid over-fitting issue.

3.2 Deep Network for Hard Segmentation

We apply the deep neural network, i.e. Pix2Pix network [17], to paintings for hard segmentation. The modified Pix2Pix network is end-to-end trainable, i.e. the output is a binary segmentation map instead of intensity image. The network consists of two parts: the generator is trained to generate the edge maps of brushstrokes from input paintings, and the discriminator is trained to detect the generator output's fakes. For most of the paintings, Pix2Pix can output good edge maps of paintings. However, in some scenarios, edge detection still remains challenging (see Figure 6). It can be noted that there are still some small gaps or blurred edges on the outputs, that is, the boundaries of brushstrokes are not closed. This is unacceptable since these gaps usually result in wrong segmentation. The following soft segmentation step always relies on accurate hard segmentation.




Fig. 6. An example of blurred edges and gaps within strokes. (b) edge map by Pix2Pix from (a) input image. (c) a small gap from (b). (d) wrongly merging the adjacent regions due to a gap on their boundaries. (e) correct segmentation.

To tackle this challenge, we add region information into the Pix2Pix network. There are two distinct advantages, (1) improving convergence of deep network training; (2) obtaining a binary edge map without gaps. The results are very encouraging, i.e. the boundaries of segments are closed. For ease of use, we keep the same mathematical symbols and terminology as in [17]. Figure 7 shows the modified Pix2Pix network structure, i.e. the generator G 's inputs are

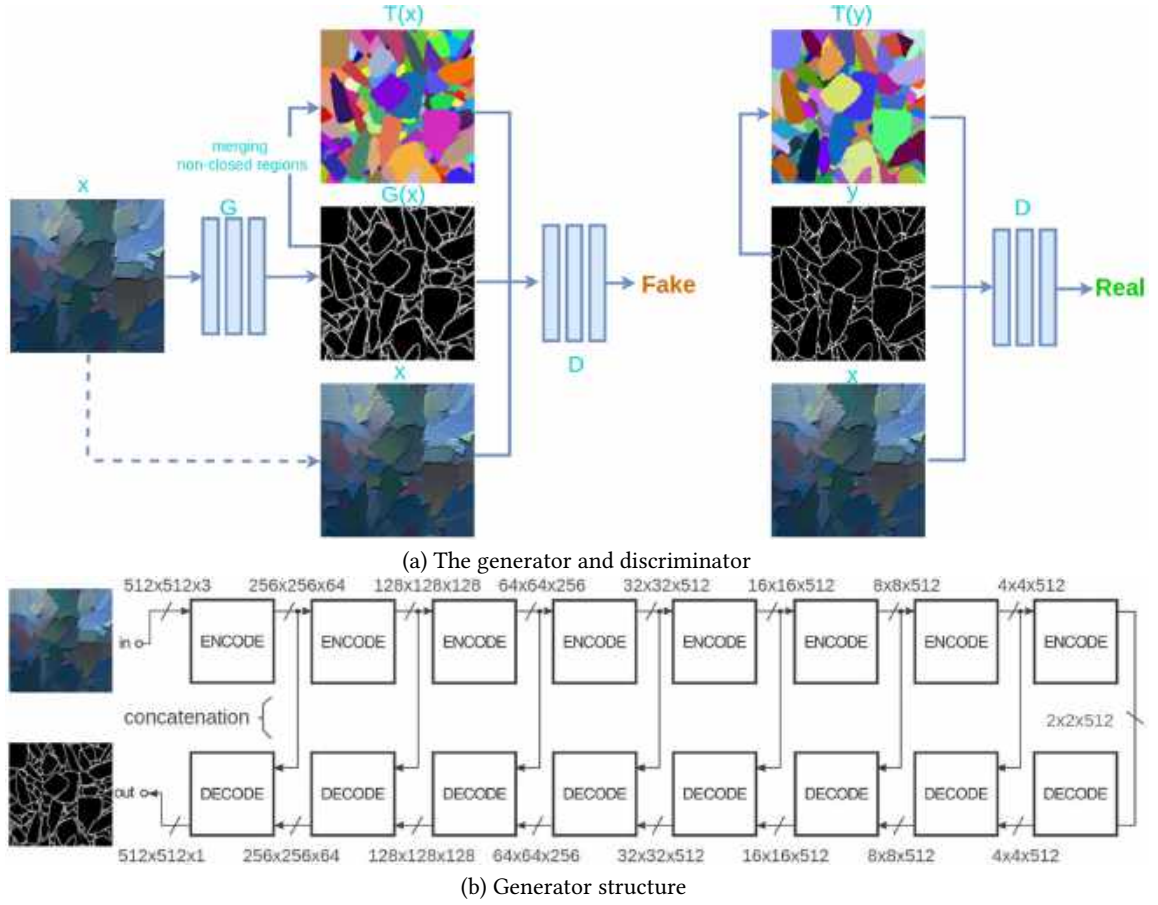


Fig. 7. The modified Pix2Pix network structure. Note that the colors in $T(x)$ and $T(y)$ are used for visualization, but not labels. $T(\cdot)$ denotes edge map that is involved in Pix2Pix network computation. Note that $T(x)$ and $T(y)$ should be of binary image edge maps without colors. Herein we color regions in order to visualize the regions closed.

paintings x while the outputs being the brushstrokes' edge maps $G(x)$. Our modification is to add a reference edge map $T(x)$ as one of the G 's outputs, which can be obtained by simple merging operation such as trapped-ball method [42]. Herein, the $T(x)$ is a binary image, whereas the $G(x)$ is an intensity image. The merging strategy is to apply the region growing to $G(x)$ (see Fig.6b) so that the regions' boundaries are closed in the resulting $T(x)$. There is no gap or blurred edge on $T(x)$. It is likely to wrongly merge adjacent regions (see Fig.6d). Compared to the output $G(x)$, the difference between the $T(x)$ and the ground truth will be likely enlarged. Thus it cannot pass the discriminator check unless wrong segmentation is corrected (see Fig.6e). The distinct advantage is that the region growing ensures the region boundaries closed and outputs the binary edge map $T(x)$ rather than ambiguous intensity image. The loss of the modified Pix2Pix network can be rewritten as,

$$L_{cGAN}(G, T, D) = E_{x, y \sim p_{data}(x, y)} [\log D(x, y, T(y))] + E_{x \sim p_{data}(x), z \sim p_z(z)} [\log (1 - D(x, G(x, z), T(G(x, z))))] \quad (1)$$

where $T(\cdot)$ denotes the reference edge map. Note that for the ground truth y , obviously $T(y) = y$ when y is binarized. The regularizer with the l_1 -norm metric is written as,

$$L_{l_1}(G) = E_{x,y \sim p_{data}(x,y), z \sim p_z(z)} [\|y - G(x, z)\|_1] \quad (2)$$

We do not add another regularizer for the reference edge map T since both $T(x)$ and $G(x)$ are edge maps and $T(x)$ is derived from $G(x)$. The discriminator is trained by,

$$\begin{cases} D[x, G(x, z), T(G(x, z))] = fake \\ D[x, y, T(y)] = real \end{cases} \quad (3)$$

The resulting generator is,

$$G^*, T^* = \arg \min_{G, T} \max_D L_{cGAN}(G, T, D) + \lambda L_{l_1}(G) \quad (4)$$

The network architecture uses the form of convolution-BatchNorm-ReLu. The generator in Fig.7b shows the skip connections. Regarding the new added term, i.e. reference edge map T , it may be regarded as a post-process of the generator G . The resulting $T(x)$ is still an edge map as well as the output $G(x)$. In fact, the binarization of the output $G(x)$ likely results in gaps from the blurred edges of the $G(x)$. The $T(x)$ has no gaps due to the merging operation. This is indeed to amplify the difference between the discriminator input and the ground truth if there is wrong segmentation, which is beneficial for deep network training. The input painting x associated with its reference edge map $T(x)$ and edge map $G(x)$ is regarded as a tensor, which is used for training the discriminator patchGAN. Moreover, we test the deep network with/without the reference edge map $T(x)$. It can be noted that the final output, $G^*(x)$, may still contain blurred edges due to the property of intensity image. This will be no longer a big deal here since the final reference edge map, $T^*(x)$, is a binary image and has no such deficiency. Thus, $T^*(x)$ is the desired hard segmentation.

In our implementation, 5000 training paintings generated by Algorithm 1 are used for 200 epochs, batch size 1, with random jitter and mirroring. To visualize the regions closed, we color every region in the edge maps $T(x)$ and $T(y)$ in Figure 6 and 7. The color values are labels which are randomly assigned and not involved in the computation. It is worthy mentioned that the final reference $T^*(x)$ is indeed a label map since the boundaries of all the regions are closed. $T^*(x)$ can be expediently labelled in any form. Currently the applications of GAN models need to convert the labels from discrete values to continuous-valued variation [17]. Our modified Pix2Pix network can successfully generate “labels” without such troubles.

3.3 Soft Segmentation

Brushstrokes overlapping usually results in fuzzy boundaries. A rising issue is to identify the boundary of a stroke on the overlapping regions. The proposed DStroke employs the guided filter [13] to the scenario of strokes overlapping in a painting and extracts the alpha mattes of brushstrokes as soft segmentation. The challenge is to identify the soft transitions between brushstrokes in terms of transparency. The guided filter has been proved to be a good explicit image matting method, particularly capturing the thin structures in a composite image [13]. This is because with the help of the guidance image, it can make the filtering output more structured and less smoothed than the input. Accurate guidance image (or called as trimap) plays a vital role in such soft segmentation. We employ the hard segmentation of a painting as its guidance image for brushstroke soft segmentation and show it below by briefly addressing the guided filter. The guided filter outputs the alpha mattes of every brushstroke in a painting which is used to further refine their boundaries. Additionally, we also apply the resulting alpha mattes to stroke ordering and coloring issue, which is required by stroke analysis for painting authentication. In terms of the pigment colors, stroke orientation and the order

of strokes from master pieces, artists hope to have an insight to the individual old masters' traits [22]. Moreover, it also is the basis of many existing layer decomposition methods [7, 30, 37, 34, 33].

Relating to the matting Laplacian, the guided filter involves an input image I , a guidance image p (or called as a trimap) and an output alpha matte α , and minimizes the cost function,

$$E(\alpha) = (\alpha - p)^T \Sigma (\alpha - p) + \alpha^T L \alpha \quad (5)$$

where L denotes an $N \times N$ matting Laplacian matrix, and Σ denotes a diagonal matrix encoded with the weights of the constraints. We view hard segmentation as p . The solution to this optimization problem can be approximated [13] by,

$$\alpha_i \approx \sum_{j \in \omega_i} W_{ij}(I) p_j \quad (6)$$

which has an $O(n)$ linear algorithm. The weight of guided filter kernel, $W_{ij}(I)$, is defined on a small sized window ω_i . Obviously, once trimap p is available, DStroke can take soft segmentation in linear time.




Fig. 8. Illustration of soft segmentation. (a) Input painting, (b) Hard segmentation, (c) Top: A segmented brushstroke mask from (b); Bottom: The selected brushstroke mask from (d), (d) Soft segmentation.

We further analyse the errors of the alpha matte α . Let the offset of p be Δp in a vector form. Substituting Δp into Eq.(6) in a matrix form yields,

$$\Delta \alpha \approx W \Delta p \quad (7)$$

Note that $\sum_{j \in \omega_i} W_{ij}(I) = 1$. If the elements of error vector Δp share the same value, this value is simply added onto the alpha matte vector α . This implies that the error Δp is transferred linearly to the alpha matte α . As a result, the guided filter performance relies on quickly providing the accurate trimaps. In our implementation, each brushstroke is processed independently, i.e. one stroke is regarded as the foreground object to be segmented from the others. The resulting alpha matte of each stroke is independent of the others. Figure 8 shows the effect of soft segmentation. It can be noted that the mask of the alpha matte is closer to the original brushstroke than the hard segmentation in Fig.8c.

Additionally, we apply the resulting alpha mattes of brushstrokes to the stroke ordering and coloring issue. The standard Porter-Duff's "A over B" compositing and blending model [28] is used,

$$(A \cdot B)_{rgb} = \frac{\alpha_A A_{rgb} + (1 - \alpha_A) \alpha_B B_{rgb}}{\alpha_A + (1 - \alpha_A) \alpha_B} \quad (8)$$

Algorithm 2: Brushstroke Ordering and Coloring

Input: a pair of overlapped strokes, A and B , with the observed colors of each pixel $(AB)_{rgb}(i)$, $i \in A \cup B$;
Output: the order of A and B ; the colors A_{rgb}, B_{rgb} ;
/ computing the colors of A and B under two supposed scenarios respectively */*

Suppose the order $A \cdot B$;
 Computing A_{rgb1}, B_{rgb1} by Eq.(8)

Suppose the order $B \cdot A$;
 Computing A_{rgb2}, B_{rgb2} by Eq.(8)
/ reconstructing the color of each pixel within A and B */*

Using A_{rgb1}, B_{rgb1} with the assumption of the order $A \cdot B$;
 Computing each pixel's color $(A \cdot B)_{rgb}(i)$ by Eq.(8) within $i \in A \cup B$

Using A_{rgb2}, B_{rgb2} with the assumption of the order $B \cdot A$;
 Computing each pixel's color $(B \cdot A)_{rgb}(i)$ by Eq.(8) within $i \in A \cup B$
/ computing order */*

$$order = \operatorname{argmin} \left(\sum_{i \in A \cup B} ((AB)_{rgb}(i) - (A \cdot B)_{rgb}(i))^2, \sum_{i \in A \cup B} ((AB)_{rgb}(i) - (B \cdot A)_{rgb}(i))^2 \right)$$

where pixel A with color A_{rgb} and alpha α_A overlays pixel B with color B_{rgb} and alpha α_B and the observed color is $(A \cdot B)_{rgb}$. The ordering is unchangeable since the “ A over B ” operation is not commutative. For a pair of overlapped brushstrokes, we can apply Eq.(8) within the union region of these two brushstrokes to each pixel, and compute the two stroke colors, A_{rgb}, B_{rgb} . In terms of our observation, it is plausible to assume that every brushstroke contains only one color in a painting. The color change within one brushstroke (e.g. from light to dark) is most likely caused by transparency, i.e. alpha change. For a given painting, the alpha mattes may have a variation within a brushstroke while the brushstroke color remains the same. If the order of stroke “ A over B ” is correct, reconstructing the color $(A \cdot B)_{rgb}$ of each pixel by Eq.(8) using the resulting colors, A_{rgb}, B_{rgb} , should result in a small error. Otherwise, the order needs to be reversed. Solving the stroke colors A_{rgb} and B_{rgb} will result in an over-determining linear system depending on the number of pixels within these two strokes. This is a least square solution. If the order of stroke “ A over B ” is correct, the linear system is compatible and reconstructing the color $(A \cdot B)_{rgb}$ of each pixel will finally result in a small error. Otherwise, the linear system is incompatible, which will accumulate a big error in reconstructing each pixel color. Thus, simply reconstructing pixel colors may discriminate which order and brushstroke colors are acceptable. The brushstroke ordering and coloring is summarized in Algorithm 2.

Remark Algorithm 2 is useful for the layer decomposition methods [7, 30, 37, 34, 33]. These methods usually solve the opacities for each layer through minimizing a polynomial cost function, which is an optimization problem with the initial guesses of layer number, order and color values. It is likely to apply Algorithm 2 to a specified region in a painting to estimate layer colors and order of two successive layers, which are utilized as the estimations of layer color and order to solving this optimization problem.

4 RESULTS AND ANALYSIS

The experiments focus on the running time and accuracy of our DStroke. In our implementation, we applied Algorithm 1 to create 5000 paintings associated with the individual set of brushstroke edge maps as our training dataset. We didn't apply the scans of real paintings to train deep network. This is because (1) for the synthesized paintings, the available edge maps are exact without any error; (2) for the scans of real paintings, the edge maps must be marked by manual. If applying the edge maps marked by manual to deep network training, it will result in a big error. Moreover, to avoid over-fitting issue, several training strategies are employed such as Data argumentation, i.e. the samples of paintings and edge maps are randomly flipped, rotated, slightly distorted, and added noise. Moreover, the samples of paintings and edge maps are randomly cropped into different sizes and resized into 512×512 pixels. Training the modified Pix2Pix network took less than two hours on a single GTX 1080 GPU.

In this section, we compare the proposed DStroke method with several existing approaches, including [7], and Semantic Soft Segmentation [2]. The test painting set is composed of two groups, one for the scans of 15 real brush paintings and the other for the synthetic paintings produced by Algorithm 1. For numerical comparison, we have to employ some synthesized paintings here. This is because the ground truth of the brushstrokes in the synthesized paintings is exact without any errors, whereas the ground truth of the scans is delineated by manual. Additionally, for fairness, the synthetic paintings in the test painting set are reproduced rather than taken from the training dataset. Additionally, we also show several examples of applying brushstrokes to image editing, such as recoloring and inserting objects, which herald a great potential in image processing.

4.1 Comparing DStroke with [7,41]

We selected the real brush painting group in our test painting set for comparison, include acrylic paintings, watercolor paintings, and oil paintings. To demonstrate the robustness of our DStroke, paintings are carefully picked, in which brushstrokes in similar colors are heavily employed or there are many blending areas, as shown in Figure 9a. Due to the brushstrokes in similar colors and the complexity of the paintings, some of brushstrokes are undetected by [7] (note that the undetected regions are labelled in black in Figure 9b). In contrast, our DStroke successfully extracts close to all the brushstrokes, even in similar color regions as shown in Figure 9c. Moreover, we selected three well known van Gogh's painting scans and performed our DStroke on the whole paintings rather than patches. To better visually match the hard segmentation to their counterpart in the source paintings, we used the dominant color of the brushstrokes and boundaries in the hard segmentation as shown in Figure 10.

To numerically evaluate our results, likewise [7, 25], we created the ground truth through manually marking brushstrokes by the experienced artists, and applied the same accuracy metrics, i.e. valid ratio and detection ratio, to tests (the reader is referred to [25] for the details of these two ratios). Moreover, we introduced a new metric, i.e. Intersection over Union (IoU), to evaluate the overlap ratio between the target masks (manually masked strokes) and the detected masks (detected strokes) as below,

$$IoU = \frac{Target \cap Detected}{Target \cup Detected}$$

The mean of Intersection over Unions (meanIoU) are evaluated and shown in Table 1. Computing the valid ratio and detection ratio, the valid covering percentage is given as 80% in [25], whereas it is set to 85% in our experiments. Although this change results in the low valid ratios and low detection ratios in Table 1, the quality of the matched brushstrokes is satisfactory. We also tried to set the percentage to 60%. All the ratios reached 100%, which means that all the brushstrokes can be detected and extracted though the accuracy is low.




Fig. 9. Highly mixed brushstrokes in similar colors are extracted. (a) Input paintings, (b) Brushstrokes extracted by [7]. (c) Brushstrokes extracted by DStroke (only hard segmentation).

Moreover, we further compared the proposed DStroke with [25, 41] using the metrics proposed by [41], including the Mean Square angular Distance of orientation (MSD_α), the Mean Square Percentage Distance of length ($MSPD_\lambda$), Mean Square Percentage Distance of width ($MSPD_\phi$), fill rate and representativeness rate, and showed the results in Table 2 and 3 respectively. These quantitative comparisons show that the proposed method can correctly detect on an average 84% of the brushstrokes from a painting (see Table 1), and accurately approximate the real brushstrokes with high reliability (see Table 2 and 3). Thus, we only claim that the proposed DStroke can extract close to all the brushstrokes from a painting. Figure 14 shows 15 real painting scans, the effect images of segmentation and manual



Fig. 10. Illustration of hard segmentation on whole paintings. The 2nd row shows the hard segmentation.

marked strokes mentioned in Table 1. Although the modified Pix2Pix network is trained by the synthetic paintings rather than real ones, these experiments show that DStroke can achieve results comparable to manual labelling on the real paintings. To justify it, we hope to further point out the following three facts, (1) Figure 14 shows the comparison of the brushstrokes extracted by DStroke and the manual segmentations on the real paintings rather than synthetic ones; (2) the brushstroke boundaries extracted by DStroke are closed as shown in the edge maps of the column-b; (3) DStroke does not miss any strokes compared to the manual segmentations as shown in the column-e. Moreover, our DStroke can work well on a large class of brush paintings, including oil paintings and Chinese paintings.

Table 1. Evaluation and Comparison of DStroke and [7] (Higher value is better)

Painting ID	[7]			DStroke Method		
	meanIoU(%)	Valid Rate(%)	Detection Rate(%)	meanIoU(%)	Valid Rate(%)	Detection Rate(%)
Mixed strokes1 (Fig.14, row 1, oil painting)	35.01	36.15	11.11	80.42	77.00	82.17
Beach (Fig.14, row 2, acrylic painting)	33.25	40.91	8.37	82.11	80.92	87.07
Mixed strokes2 (Fig.14, row 3, oil painting)	28.30	46.88	17.36	78.16	74.52	80.99
Mixed strokes3 (Fig.14, row 4, oil painting)	32.88	39.06	14.08	81.99	79.27	82.99
Cloud (Fig.14, row 5, acrylic painting)	31.86	38.64	12.86	86.66	88.43	90.95
Dusk1 (Fig.14, row 6, watercolor painting)	35.39	40.54	12.84	76.13	74.81	89.91
Dusk2 (Fig.14, row 7, watercolor painting)	22.56	38.71	12.12	75.68	68.97	79.70
Trees (Fig.14, row 8, acrylic painting)	33.19	48.72	21.09	76.72	74.47	82.03
Mixed strokes4 (Fig.14, row 9, oil painting)	35.23	34.13	11.37	76.04	70.26	77.84
Road (Fig.14, row 10, oil painting)	26.79	46.15	19.59	76.69	74.44	85.57
Landscape (Fig.14, row 11, watercolor painting)	27.53	35.00	11.79	79.22	78.00	88.97
White lotus (Fig.14, row 12, Chinese painting)	17.22	49.57	8.45	65.40	77.19	71.01
Yellow lotus (Fig.14, row 13, Chinese painting)	23.48	45.16	12.14	73.39	88.90	95.71
Flying bird (Fig.14, row 14, Chinese painting)	19.78	42.48	10.11	71.02	88.39	93.09
Peony (Fig.14, row 15, Chinese painting)	33.43	43.33	8.13	64.07	75.49	70.03

Table 2. Comparison with Table 5 and 6 of [41] (Paintings are from [25] and the higher value is better.)

Painting ID	Valid rate (%)				Detection rate (%)				Fill rate (%)			Representative rate (%)		
	[25]	[41]	[7]	DStroke	[25]	[41]	[7]	DStroke	[25]	[41]	DStroke	[25]	[41]	DStroke
F218	42.7	20.2	88.1	86.3	21.6	20.5	48.5	52.3	29.1	55.2	62.2	2.1	22.3	27.3
F386	73.7	41.6	82.4	83.1	68.4	59.7	90.2	90.6	15.0	24.6	35.9	0.7	13.2	22.4
F518	60.7	35.3	71.2	75.2	75.2	60.1	78.9	81.9	24.2	40.3	51.5	8.4	29.8	28.0
F538	49.1	23.5	84.2	87.1	44.9	32.2	81.3	85.1	12.9	31.8	41.8	3.2	23.7	34.6

Table 3. Comparison of brushstroke's length, width and orientation (We ask experts to label these 4 paintings and compare our results with Table 7 of [41])

Painting ID	MSD_{α}			$MSPD_{\lambda}$			$MSPD_{\phi}$		
	[25]	[41]	DStroke	[25]	[41]	DStroke	[25]	[41]	DStroke
F218	144.00	137.00	95.00	1.29	1.06	0.54	2.40	0.86	0.45
F386	112.00	94.00	65.00	3.89	3.66	1.25	2.75	0.74	0.42
F518	47.00	68.00	24.00	1.44	0.79	0.26	2.76	0.45	0.15
F538	97.00	83.00	42.00	1.12	0.38	0.12	2.50	0.43	0.28




Fig. 11. Comparison of semantic segmentation and instance-level segmentation. (a) Input paintings; (b) The results of Semantic Soft Segmentation. This method accidentally merges many strokes into one segment. (c) Soft segmentations and (d) Hard segmentation by DStroke. DStroke can do segmentation on instance level to identify all the strokes.

4.2 Comparing DStroke with Semantic Soft Segmentation




Fig. 12. Comparison of the alpha map of a segmented stroke and the ground truth. (a) Input painting, (b) Alpha map of a stroke segmented by Eq.(7), (c) Alpha map of the ground truth, (d) MSE with varying overlap ratio.




Fig. 13. Comparison of alpha MSE and running time at different size levels.

Our DStroke is also compared with the most relevant work: Semantic Soft Segmentation in [2]. The advantages of DStroke are, for hard segmentation, (1) there is no limitation of segment number; (2) it can do segmentation on instance level as shown in Figure 11d. In contrast, due to lack of appropriate hard segmentations, Semantic Soft Segmentation [2] wrongly classifies many strokes into one stroke as shown in Figure 11b. For soft segmentation, our DStroke works in a $O(n)$ complexity based on the guided filter [13]. For numerical evaluation, we used the synthetic painting group in our test painting set for comparison since all the brushstroke information is available in advance, which can be regarded as the ground truth. We select 500 synthetic paintings with different sizes (from 0.1K to 1M) for test. The accuracy of brushstroke alpha mattes is computed by MSE (mean squared error),

$$MSE = \frac{1}{M} \sum_{j=1}^M \left(\frac{1}{N} \sum_{i=1}^N (\alpha_{ij} - \hat{\alpha}_{ij})^2 \right)$$

where M denotes the stroke number, N denotes the pixel number within one stroke, and the alpha of the i -th pixel on the j -th brushstroke α_{ij} is generated by Eq.(7), and $\hat{\alpha}_{ij}$ is the corresponding ground truth. Figure 12 illustrates the alpha mattes of a stroke segmented by Eq.(7) and the ground truth. Moreover, Eq.(7) indicates that the errors of hard segmentations are linearly transferred to the alpha mattes. Figure 12d shows the example of MSE with varying the ratio of the detected hard segmentation over the ground truth hard segmentation, which is approximately linear.

Thus, the soft segmentation entirely depends on the hard segmentation. Figure 13 further shows the comparison of the proposed DStroke method and Semantic Soft Segmentation [2] on both running time and alpha matte accuracy. It can be noted that the MSE of our DStroke is obviously less than [2]. Whether overlapping or changing image size, a big error is introduced by [2] into hard segmentation. In contrast, DStroke can continually provide a more accurate hard segmentation despite an increasing image size, hence reducing MSE.

In summary, compared to Semantic Soft Segmentation, our DStroke method has a better performance in the following aspects:

- DStroke is able to do soft segmentation on instance level;
- DStroke has no limitation on the segment number enabling soft segmentation for thousands of strokes whereas Semantic Soft Segmentation can only segment a small number of regions;
- The accuracy of DStroke’s alpha mattes is noticeably higher than Semantic Soft Segmentation;
- The running time of DStroke method is much less than Semantic Soft Segmentation. For a 640×480 image, Semantic Soft Segmentation takes around 3 mins on segmentation, while DStroke taking less than 1 second.

4.3 Applications in Image Editing

When almost all the brushstrokes in a painting are extracted, it is likely to recompose the painting through manipulating the strokes. Figure 15 shows two applications in image editing, recoloring paintings through changing brushstrokes’ colors and inserting objects into paintings. To make inserted objects and recolored colors look reasonable and vivid, we performed the DStroke for brushstroke extraction on three whole real paintings and then carried out inserting and recoloring separately.

5 CONCLUSION

In this paper, we propose a Deep learning based brush Stroke extraction method-DStroke, which consists of two parts, (1) edge detection (or hard segmentation) through the modified Pix2Pix network; (2) soft segmentation by the guided filter. Compared to the state of the art methods, the main merit of the proposed DStroke is its high efficiency, i.e. to automatically and rapidly extract close to all the brushstrokes from a brush painting and accurately recover the faithful soft transitions between brushstrokes. The numerical results show that our DStroke can accurately approximate brushstrokes with high reliability.

Our main contributions include (1) proposing a painting production method which automatically builds up a large brush painting training dataset without the need for manual annotation. To the best of our knowledge, we are the first to provide an automatic method of building up a large painting sample dataset for deep learning purpose. (2) The modified Pix2Pix network can do segmentation on instance level and output “labels” for segmentation. In addition we also apply the estimated alpha mattes of brushstrokes to identifying stroke ordering and coloring issues for painting authentication purposes. Experimental results further demonstrate that our DStroke method outperforms the current state of the art methods.

Limitations Due to the diversity of drawing art, our DStroke is unsuitable for traditional (or classic) western paintings, whose style was developed in the Renaissance and emphasized realism such as "Mona Lisa" by Da vinci, and abstract paintings like "Composition 8" from Vasily Kandinsky. Additionally, it cannot deal with the mixed color strokes that appear in watercolor and oil paintings with brush mixing and pickup. These limitations direct our further research.

We plan to build up an old masters' brushstroke database using the proposed DStroke method in the near future. Moreover we are also interested in the recent work, alphaGAN [21], and aim to integrate hard segmentation and soft segmentation into the GAN model soon.

ACKNOWLEDGEMENT

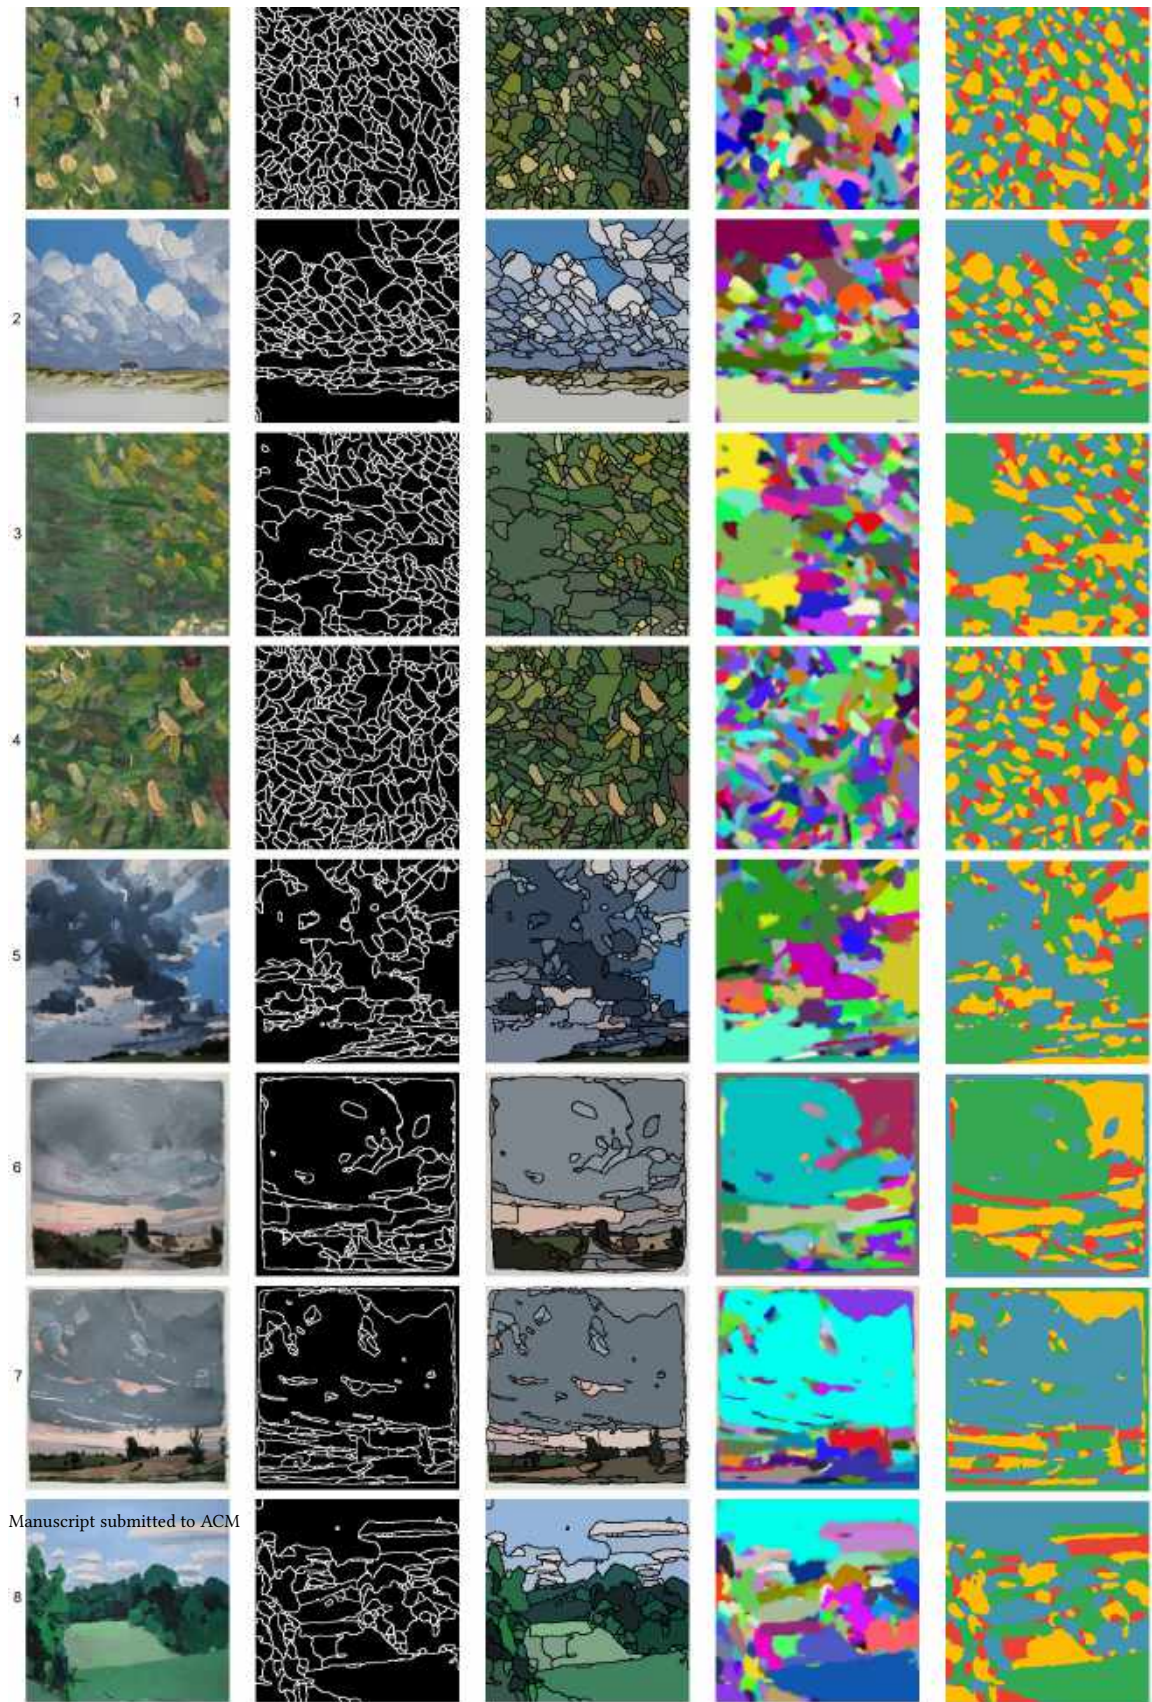
This research was supported by EU-H2020-MSCA project: AniAge (No.691215) and the grant of the Ministry of Science and Technology, Taiwan (No.107-2221-E-006-196-MY3 and 108-2221-E-006 -038 -MY3).

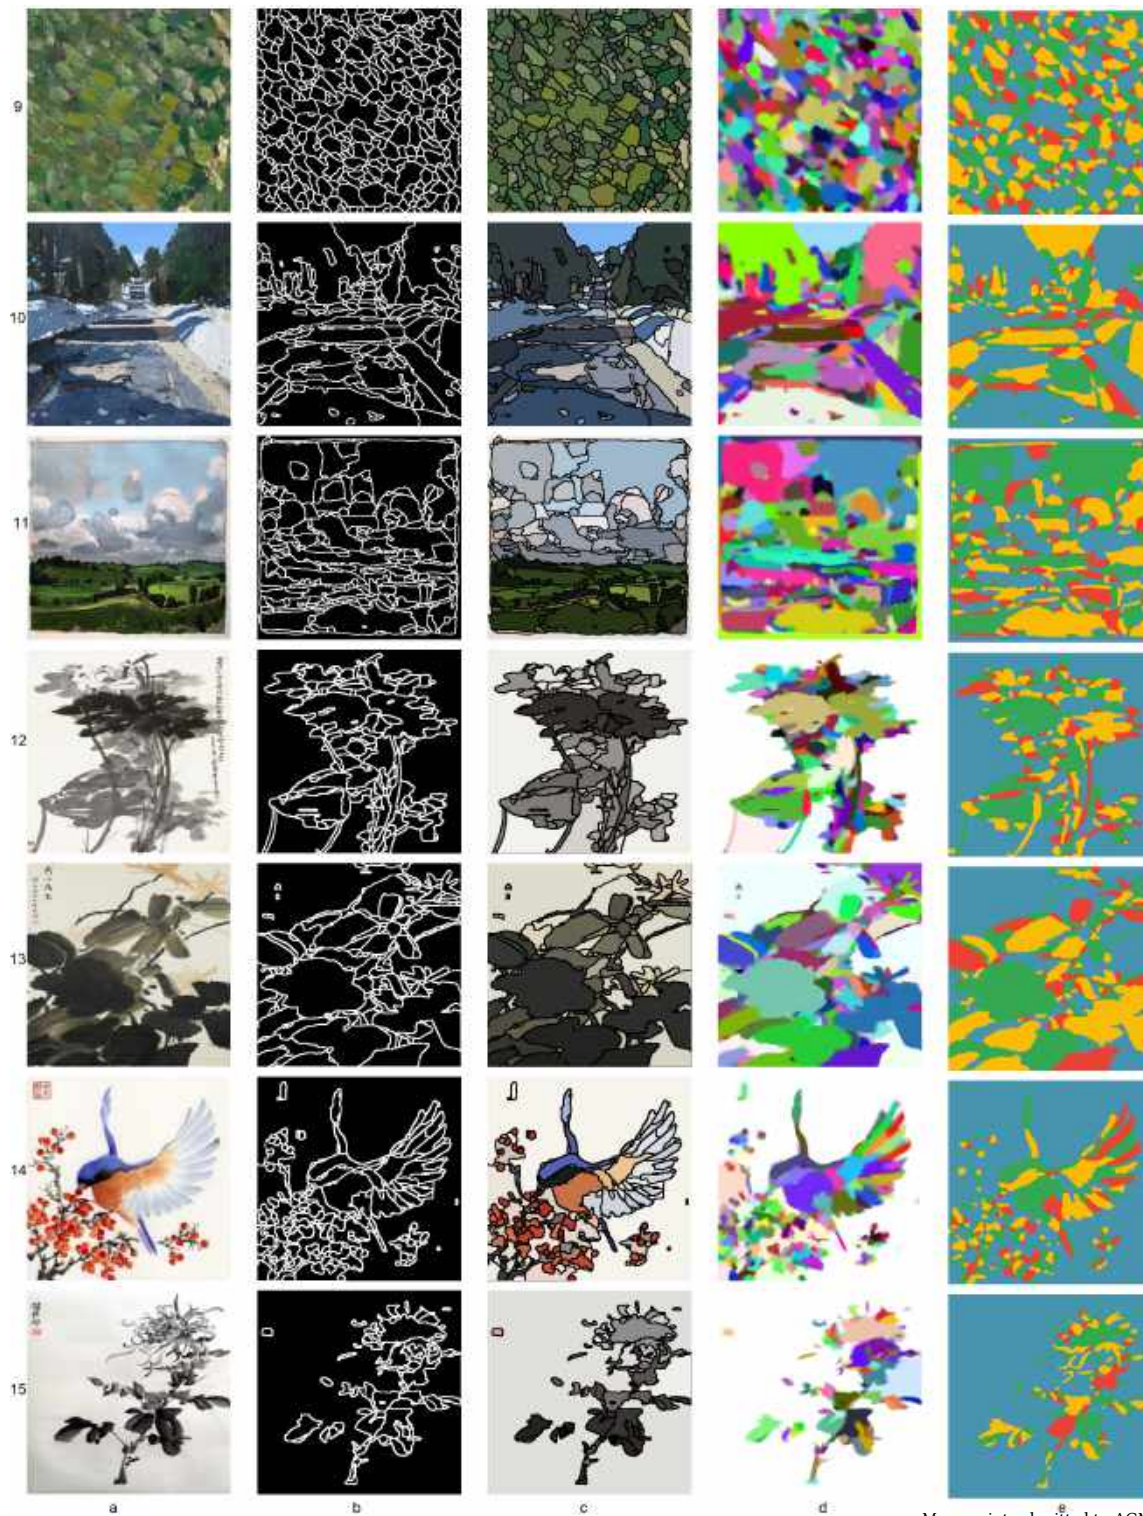
REFERENCES

- [1] Yagiz Aksoy, Tunc Ozan Aydin, Marc Pollefeys, 2017, Designing Effective Inter-Pixel Information Flow for Natural Image Matting, in Proc. of CVPR'17
- [2] Aksoy, Y., Oh, T.-H., Paris, S., Pollefeys, M., and Matusik, W., 2018, Semantic soft segmentation, ACM Trans. on Graphics, 37(4):72
- [3] Elad Aharoni-Mack, Yakov Shambik, and Dani Lischinski, 2017, Pigment-based recoloring of watercolor paintings, in Proc. of ACM Symposium on Non-Photorealistic Animation and Rendering (NPAR'17), Article 1
- [4] Berezhnoy, I.E., Postma, E.O. van den Herik, H.J., 2009, Automatic extraction of brushstroke orientation from paintings, Machine Vision and Applications, Vol.20, Issue 1.
- [5] Chen, Q., Li, D., and Tang, C.-K., 2013, KNN matting, IEEE Trans. on PAMI, 35(9):2175–2188.
- [6] Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L., 2017, Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE Trans. on PAMI, 40(4):834–848.
- [7] Y. Fu, H. Yu, C. Yeh, J. J. Zhang and T. Lee, 2018, High Relief from Brush Painting, IEEE Trans. on Visualization and Computer Graphics, doi: 10.1109/TVCG.2018.2860004
- [8] Gooch, B., Coombe, G., and Shirley, P., 2002, Artistic vision: painterly rendering using computer vision techniques, in Proc. of ACM 2nd Int'l symposium on Non-photorealistic animation and rendering, pp.83–ff.
- [9] Hertzmann, A., 2003, A survey of stroke-based rendering, IEEE Computer Graphics and Applications, Vol.23, No.4, pp.70-81
- [10] Hertzmann A., 2002, Fast paint texture, in Proc. of ACM 2nd Int'l symposium on Non-photorealistic animation and rendering (NPAR'02), pp.91-ff.
- [11] He, K., Gkioxari, G., Dollár, P., and Girshick, R., 2017, Mask r-CNN, in Proc. of IEEE ICCV'17, pp.2980–2988.
- [12] Hegde, S., Gatzidis, C., and Tian, F., 2013, Painterly rendering techniques: a state-of-the-art review of current approaches. Computer Animation and Virtual Worlds, 24(1):43–64.
- [13] He, K., Sun, J., and Tang, X., 2013, Guided image filtering, IEEE Trans. on PAMI, (6):1397–1409
- [14] Hurtut, T., 2D Artistic Images Analysis, a Content-based Survey, 2010, http://hal.archivesouvertes.fr/hal-00459401_v2/
- [15] Zhongyi Han, Benzhenq Wei, Ashley Mercado, Stephanie Leung, Shuo Li, 2018, Spine-GAN: Semantic segmentation of multiple spinal structures, Medical Image Analysis, Vol.50, pp.23-35
- [16] Ivanova, Krassimira; Stanchev, Peter; Vanhoof, Koen; and Ein-Dor, Phillip, 2010, Semantic and abstraction content of art images, in Proc. of Mediterranean Conf. on Information Systems 2010.
- [17] Isola, Phillip, Zhu, Jun-Yan, Zhou, Tinghui and Efros, Alexei A., 2017, Image-to-Image Translation with Conditional Adversarial Networks, in Proc. of CVPR2017
- [18] Yongcheng Jing, Yang Liu, Yezhou Yang, Zunlei Feng, Yizhou Yu, Dacheng Tao, and Mingli Song, Stroke Controllable

- Fast Style Transfer with Adaptive Receptive Fields, in Proc. of European Conf. on Computer Vision 2018, pp.238-254
- [19] Koyama, Y. and Goto, M., 2018, Decomposing Images into Layers with Advanced Color Blending, *Computer Graphics Forum*, 37: 397-407. doi:10.1111/cgf.13577
- [20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton, 2017, ImageNet classification with deep convolutional neural networks, *Commun. of ACM*, Vol.60, Issue 6, pp.84-90.
- [21] Lutz, S., Amliantitis, K. and Smolic, A., 2018, AlphaGAN: Generative adversarial networks for natural image matting. in Proc. of British Machine Vision Conf. 2018.
- [22] C. Li and T. Chen, 2009, Aesthetic Visual Quality Assessment of Paintings, *IEEE Journal of Selected Topics in Signal Processing*, Vol.3, No.2, pp.236-252
- [23] Levin, A., Lischinski, D., and Weiss, Y., 2008, A closed-form solution to natural image matting. *IEEE Trans. on PAMI*, 30(2):228–242.
- [24] Levin, A., Rav-Acha, A., and Lischinski, D., 2008, Spectral matting. *IEEE Trans. on PAMI*, 30(10):1699–1712.
- [25] Li, J., Yao, L., Hendriks, E., and Wang, J. Z., 2012, Rhythmic brush-strokes distinguish van Gogh from his contemporaries: findings via automated brushstroke extraction, *IEEE Trans. on PAMI*, 34(6):1159–1176.
- [26] McCann, J. and Pollard, N., 2009, Local layering. *ACM Trans. on Graphics*, 28(3):84.
- [27] McCann, J. and Pollard, N. S., 2012, Soft stacking. *Computer Graphics Forum*, Vol.31, pp.469–478.
- [28] T. Porter and T. Duff, 1984, Compositing digital images. in Proc. of ACM Siggraph84-Computer Graphics, Vol.18, pp.253–259.
- [29] Ronneberger O., Fischer P., Brox T., 2015, U-Net: Convolutional Networks for Biomedical Image Segmentation. in Proc. of Medical Image Computing and Computer-Assisted Intervention–MICCAI2015, LNCS Vol.9351, Springer.
- [30] Richardt, C., Lopez-Moreno, J., Bousseau, A., Agrawala, M., and Drettakis, G., 2014, Vectorising bitmaps into semi-transparent gradient layers, *Computer Graphics Forum*, Vol.33, pp.11–19.
- [31] Shamir, Lior, 2015, What makes a Pollock Pollock: a machine vision approach, *Int'l J. of Arts and Technology*, Vol.8, No.1, pp.1-10
- [32] Singaraju, D. and Vidal, R., 2011, Estimation of alpha mattes for multiple image layers, *IEEE Trans. on PAMI*, 33(7):1295–1309.
- [33] Jianchao Tan, Jose Echevarria, and Yotam Gingold, 2018, Efficient palette-based decomposition and recoloring of images via RGBXY-space geometry, *ACM Trans. Graph.*, 37(6), Article 262, doi.org/10.1145/3272127.3275054
- [34] Tan, J., Lien, J.-M., and Gingold, Y., 2016, Decomposing images into layers via rgb-space geometry. *ACM Trans. on Graphics*, 36(1):7
- [35] van der Maaten, Laurens J. P.; Postma, Eric O., 2010, Texton-based analysis of paintings, in Proc. of the SPIE, Vol.7798, id.77980H
- [36] Ning XIE and Hirotaka HACHIYA and Masashi SUGIYAMA, 2013, Artist Agent: A Reinforcement Learning Approach to Automatic Stroke Generation in Oriental Ink Painting, *IEICE Trans. on Information and Systems*, Vol.E96.D, No.5, pp.1134-1144
- [37] Xu, S., Xu, Y., Kang, S. B., Salesin, D. H., Pan, Y., and Shum, H.-Y., 2006, Animating Chinese paintings through stroke-based decomposition, *ACM Trans. on Graphics*, 25(2):239–267
- [38] Xie, N., Zhao, T., Tian, F., Zhang, X., and Sugiyama, M., 2015, Stroke-based stylization learning and rendering with inverse reinforcement learning, in Proc. of the 24th Int'l Conf. on Artificial Intelligence (IJCAI'15), pp.2531-2537.
- [39] Ningyuan Zheng and Yifan Jiang and Dingjiang Huang, StrokeNet: A Neural Painting Environment, in Proc. of Int'l Conf. on Learning Representations 2019

- [40] Zhao, H., Shi, J., Qi, X., Wang, X., and Jia, J., 2017, Pyramid scene parsing network, in Proc. of IEEE CVPR'17, pp.6230–6239.
- [41] Lamberti, F., Sanna, A. Paravati, G., 2014, Computer-assisted analysis of painting brushstrokes: digital image processing for unsupervised extraction of visible features from van Gogh's works, EURASIP Journal on Image and Video Processing, 2014:53.
- [42] S. Zhang, T. Chen, Y. Zhang, S. Hu and R. R. Martin, 2009, Vectorizing Cartoon Animations, IEEE Trans. on Visualization and Computer Graphics, 15(4), pp. 618-629.
- [43] Robert L. Cook and Kenneth E. Torrance. 1981. A reflectance model for computer graphics. SIGGRAPH Comput. Graph. 15, 3 (August 1981), 307–316. DOI:<https://doi.org/10.1145/965161.806819>
- [44] <https://www.kylebrush.com/>
- [45] Liqiang Nie; Meng Liu; Xuemeng Song, 2019, Multimodal Learning toward Micro-Video Understanding, Multimodal Learning toward Micro-Video Understanding, Morgan-Claypool.
- [46] P. Arbelaez, M. Maire, C. Fowlkes and J. Malik, 2009, From contours to regions: An empirical evaluation, IEEE CVPR09, Miami, FL, pp.2294-2301, doi: 10.1109/CVPR.2009.5206707.
- [47] Jingwan Lu, Connelly Barnes, Stephen DiVerdi, and Adam Finkelstein. 2013. RealBrush: painting with examples of physical media. ACM Trans. Graph. 32(4), Article 117, 12 pages. DOI:<https://doi.org/10.1145/2461912.2461998>





Manuscript submitted to ACM

Fig. 14. Comparison of the results by DStroke and manual segmentation. (a) Input paintings, (b) Detected brushstroke edges, (c) Hard segmentation, (d) Soft segmentation, (e) Manual segmentation. We use the dominant color with segmentation boundaries in (c) to easily compare the hard segmentation with the brushstrokes of the paintings in (a). However, to highlight transition areas in soft segmentation, we still use random colorization in (d) to ensure a big color difference between any two adjacent brushstrokes.




Fig. 15. Image editing. Row 1 shows the whole real paintings. Row 2 shows the results of inserting objects. A tree in the painting 3 is inserted in the painting 1. The farmer in the painting 1 is inserted into the painting 2. Row 3 shows the results of recoloring paintings.