Journal of Visual Communication and Image Representation 51 (2018) 29–39

Contents lists available at ScienceDirect

Journal of Visual Communication and Image Representation

journal homepage: www.elsevier.com/locate/jvci

Stroke-based stylization by learning sequential drawing examples☆

Ning Xie[a], Yang Yang[a], Heng Tao Shen[a], Ting Ting Zhao[b,*]

[a] Center for Future Media, School of Computer Science and Engineering, University of Electronic Science and Technology of China, Chengdu 611731, China
[b] School of Computer Science and Information Engineering, Tianjin University of Science and Technology, Tianjin 300457, China

ARTICLE INFO

Keywords:
Stroke-based stylization
Reinforcement learning
Inverse reinforcement learning

ABSTRACT

Among various traditional art forms, *brush stroke drawing* is one of the widely used styles in modern computer graphic tools such as GIMP, Photoshop and Painter. In this paper, we develop an *AI-aided art authoring* (*A*4) system of *non-photorealistic rendering* that allows users to automatically generate brush stroke paintings in a specific artist's style. Within the reinforcement learning framework of brush stroke generation proposed by Xie et al. (2012), the first contribution in this paper is the application of regularized policy gradient method, which is more suitable for the stroke generation task; the other contribution is to learn artists' drawing styles from video-captured stroke data by *inverse reinforcement learning*. Through experiments, we demonstrate that our system can successfully learn artists' styles and render pictures with consistent and smooth brush strokes.

1. Introduction

Artistic stylization in non-photorealistic rendering enables users to stylize pictures with the appearance of traditional art forms, such as pointillism painting, line sketching, or brush stroke drawing. Among them, brush stroke drawing is one of the widely used art styles across different cultures in history. In computer-generated painterly rendering, *stroke placement* is a big challenge and significant efforts have been made to investigate how to draw a stroke with realistic brush texture in a desired shape and how to organize multiple strokes [10].

The goal of this paper is to develop an AI-aided art authoring system for artistic brush stroke generation. In this section, we first review backgrounds in computer graphics and artificial intelligence, then give an overview of our proposed system, and finally give an introduction of a new policy learning method.

1.1. Background in computer graphics

The most straightforward approach for painterly rendering would be *physics-based painting*, i.e., giving users an intuitive feeling just like drawing with a real brush. Some works modeled physical virtual brushes including its 3D structure, dynamics, interaction with paper surface [6] and simulating the physical effect of the ink dispersion [7]. These virtual brushes can be used to draw various styles of strokes with a digital pen or mouse. However, it is very complex to control a virtual brush. Furthermore, since the computational cost is often very high to achieve satisfactory visual effects to human eyes, some physics-based painting approaches rely on graphics processing units (GPUs) for obtaining reasonable performance [5].

To address these issues associated with physics-based painting, the *stroke-based rendering* approach was proposed to directly simulate rendering marks (such as lines, brush strokes, or even larger primitives such as tiles) on a 2D canvas. This stroke-based rendering [13] underpins many artistic rendering algorithms, especially on those emulating traditional brush-based artistic styles such as oil painting and watercolor.

Although physics-based painting and stroke-based rendering are useful for (semi-) professional usage, often users who have no painting expertise are only interested in final results rather than the painting process itself [16]. To make the painterly rendering system more accessible to novice users, several researchers investigated *beautification*. The early work [28] explored automatic techniques to beautifying geometric drawings by enforcing various relations such as the collinearity of lines and the similarity of their lengths. The approach by Igarashi et al. [29] offered users several choices in the beautifying process.

Filter-based methods are also widely used for building artistic rendering algorithms applied in the image manipulation software such as Photoshop and Gimp. The main task is to find out the nice or beautiful outlines/edges based on novel filters, such as bilateral filter [24], DoG filter [27], morphological filter [4], shock filter [17] and Kuwahara filters [18]. Recently, the deep network is applied for artistic filter [2,11]. These techniques are usually based on heuristics developed through hands-on experience, showing that certain combinations of

☆ This paper has been recommended for acceptance by Meng Wang.

* Corresponding author.
E-mail addresses: seanxiening@gmail.com (N. Xie), dlyyang@gmail.com (Y. Yang), shenhengtao@hotmail.com (H.T. Shen), tingting@tust.edu.cn (T.T. Zhao).

https://doi.org/10.1016/j.jvcir.2017.12.012
Received 28 February 2017; Received in revised form 7 November 2017; Accepted 29 December 2017
Available online 04 January 2018
1047-3203/

filters produce an artistic look, more precisely, called stylized cartoon rendering, pen-and-ink illustration, or watercolor painting.However, the connection between the edge-preserving image simplification and the artistic rendering is less obvious, because the significant artistic look is often achieved or further reinforced by taking the local image structure and brush stroke details into account, rather than the global image abstraction. In practice, the designers usually firstly apply the painting style filter on the real photo in order to imagine the whole art authoring in terms of the entire layout. Then, another layer is created on the top to emphasize the local parts of the image that are important to the users by hand or stroke-based methods.

More recently, methods that attempt not only to beautify generated artistic images, but also to maintain users' personal styles are pursued. Studies of style imitation in artistic rendering focused on ink sketching. Baran et al. [3] proposed a method to draw smooth curves while maintaining the details. The sketch beautification approach by Orbay et al. [21] used the model that automatically learns how to parse a drawing. Zitnick [41] proposed a general purpose approach to handwriting beautification using online input from a stylus. Since techniques of line sketching style imitation are not suitable to synthesizing brush strokes, quite a few previous works [36,37,19] tried to reproduce brush stroke texture as reference samples.

1.2. Background in artificial intelligence

Differently from the above approaches developed in computer graphics, the system we propose in this paper trains a virtual brush agent to learn the stroke drawing model according to particular artists' styles using their stroke drawing videos. The problem on truncated fault texture can be solved by using our learned stroke drawing behavior model. Since the brush agent is trained locally with the data set of basic stroke shapes, we can create strokes in new shapes even when they are quite different from an artist's examples. This is eminently suitable for the artistic stylization of images when non-expert users try to render their favorite photos into a particular artist's style with just a few button clicks.

Our proposed system is based on the *reinforcement learning* (RL) method to artistic brush stroke generation [33], which allows users to automatically produce consistent and smooth brush strokes. In this RL approach, the task of synthesizing the texture of each individual stroke is formulated as a sequential decision making problem based on the Markov decision process, where a soft-tuft brush is regarded as an RL agent. Then, to sweep over the shape closed by the contour, the agent is trained by a policy learning method to learn which direction to move and how to keep the stable posture while sweeping over various stroke shapes provided as training data. Finally, in the test phase, the agent chooses actions to draw strokes by moving a virtual inked brush within a newly given shape represented by a closed contour.

In this paper, we extend this RL-based approach to be able to incorporate personal artistic stylization. More specifically, we propose the method of *inverse RL* [1] for apprenticeship learning to acquire skilled stylized stroke drawing behaviour. Our system learns to characterize the reward function from stroke data video-captured from artists: we first invite artists to draw strokes using our handcrafted device for recording the brush movement. The brush footprint in each key frame of the captured stroke-drawing video is then extracted, and time series data are obtained by assembling the extracted posture configuration of each footprint including the motion attitude, pose, and locomotion of the brush. The data are used to mimic the artist's stroke drawing style through the reward function learned by inverse RL (IRL).

1.3. Overview of our proposed system

An overview of our system, called *AI-aided art authoring* ($A4$) system for artistic brush stroke generation, is illustrated in Fig. 1. Our system consists of two phases: an *online synthesis phase* and an *offline training phase*.

In the *online synthesis phase*, $A4$ provides a fast and easy-to-use graphical-user interface so that users can focus on developing art work concepts just by sketching the position and attitude of desired strokes. Given an input picture or photo, even non-expert users can sketch the shapes of desired strokes using either closed contours or simple curves.

In the *offline training phase*, the main task is to train the virtual agent so as to synthesize strokes in an artist's drawing style. Instead of the classical policy gradient method [31], we use the state-of-the-art policy gradient algorithm called *importance-weighted policy gradients with parameter-based exploration* [40], which allows stable policy update and efficiently reuse of previously collected data.

Through experiments, we demonstrate that the proposed system is promising in producing stroke placement with a personalized style.

1.4. Policy learning algorithm

Among reinforcement learning methods, the *policy gradient* (PG) method [32,25] demonstrated remarkable successes in complex systems [22,9,26].

Nevertheless, the PG method has a weakness that estimation of policy gradients can be unreliable in practice. In order to reduce the variance of gradient estimation, useful techniques have been proposed, including, the natural gradient [15,23], parameter-based exploration [25,20], and the optimal baseline [30,12,39]. While all of these methods were shown to stabilize the policy update to some extent, none of them directly take the variance of gradient estimates into account in the objective. Thus, further stabilization is necessary to efficiently solve challenging RL problems.

In this paper, we adopt a more explicit way for further variance reduction, by directly employing the variance of policy gradients as a regularizer. More specifically, we design a new framework for PG algorithms by directly incorporating the variance of policy gradients in the objective function. The proposed variance-regularized framework can naturally increase the expected return, but also reduce the variance of gradient estimates.

In practice, we combine our variance-regularization technique with parameter-based exploration [25,20] and the optimal baseline [30,12] for further variance reduction.More specifically, we implement a state-of-the-art PG method, policy gradient with parameter based exploration (PGPE) with optimal baseline subtraction [39], in our proposed variance-regularized framework.

2. Reinforcement learning formulation of brush agent

In order to synthesize the painting imagery of an artist's personal style, we construct our brush agent equipped with the style learning ability by extending the existing RL-based approach [33] as illustrated in Fig. 2.

We assume that our stroke drawing problem is a discrete-time Markov decision process. At each time step t, the agent observes a state $\boldsymbol{s}_t \in \mathscr{S}$, selects an action $a_t \in \mathscr{A}$, and then receives an immediate reward r_t resulting from a state transition. The state space $\mathscr{S}$ and action space $\mathscr{A}$ are both defined as continuous spaces in this paper. The dynamics of the environment is characterized by unknown conditional density $p(\boldsymbol{s}_{t+1}|\boldsymbol{s}_t,a_t)$, which represents the transition probability density from current state $\boldsymbol{s}_t$ to next state $\boldsymbol{s}_{t+1}$ when action a_t is taken. The initial state of the agent is determined following unknown probability density $p(\boldsymbol{s}_1)$. The immediate reward r_t is given according to the reward function $R(\boldsymbol{s}_t,a_t,\boldsymbol{s}_{t+1})$. The agent's decision making procedure at each time step t is characterized by a parameterized policy $p(a_t|\boldsymbol{s}_t,\theta)$ with parameter θ, which represents the conditional probability density of taking action a_t in state $\boldsymbol{s}_t$.

A sequence of states and actions forms a *trajectory* denoted by

$$h := [\boldsymbol{s}_1,a_1,\ldots,\boldsymbol{s}_T,a_T],$$

where T denotes the number of steps called the horizon length. Given