Neural Style Transfer: A Paradigm Shift for Image-based Artistic Rendering?

Amir Semmo, Tobias Isenberg, Jürgen Döllner

► To cite this version:

Amir Semmo, Tobias Isenberg, Jürgen Döllner. Neural Style Transfer: A Paradigm Shift for Image-based Artistic Rendering?. NPAR 2017 - Proceedings of the International Symposium on Non-Photorealistic Animation and Rendering (NPAR as part of Expressive), ACM, Jul 2017, Los Angeles, United States. pp.5:1-5:13, 10.1145/3092919.3092920 . hal-01527495v3

HAL Id: hal-01527495

<https://hal.inria.fr/hal-01527495v3>

Submitted on 7 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Neural Style Transfer: A Paradigm Shift for Image-based Artistic Rendering?

Amir Semmo
Hasso-Plattner-Institut,
Faculty of Digital Engineering,
University of Potsdam, Germany
amir.semmo@hpi.uni-potsdam.de

Tobias Isenberg
Inria & Université Paris-Saclay
France
tobias.isenberg@inria.fr

Jürgen Döllner
Hasso-Plattner-Institut,
Faculty of Digital Engineering,
University of Potsdam, Germany
doellner@hpi.uni-potsdam.de

Figure 1: Outputs of a feed-forward NST using CNNs for image processing (Johnson et al. 2016a,b). The potentials and impact of NST on IB-AR and its combinations with paradigms such as image filtering (here: oil paint, watercolor) are discussed in this paper. “Brooklyn Bridge” by Curtis MacNewton is licensed under CC BY-SA 2.0;¹ derivatives implicate style variants.

ABSTRACT

In this meta paper we discuss *image-based artistic rendering* (IB-AR) based on *neural style transfer* (NST) and argue, while NST may represent a paradigm shift for IB-AR, that it also has to evolve as an interactive tool that considers the design aspects and mechanisms of artwork production. IB-AR received significant attention in the past decades for visual communication, covering a plethora of techniques to mimic the appeal of artistic media. *Example-based rendering* represents one of the most promising paradigms in IB-AR to (semi-)automatically simulate artistic media with high fidelity, but so far has been limited because it relies on pre-defined image pairs for training or informs only low-level image features for texture transfers. Advancements in deep learning showed to alleviate these limitations by matching content and style statistics via activations of neural network layers, thus making a generalized style transfer practicable. We categorize style transfers within the taxonomy of IB-AR, then propose a semiotic structure to derive a technical research agenda for NSTs with respect to the grand challenges of NPAR. We finally discuss the potentials of NSTs, thereby identifying applications such as casual creativity and art production.

This is the authors' version of the work. It is posted here for your personal use. Not for redistribution. The definitive version will be published in Proceedings of Non-Photorealistic Animation and Rendering (NPAR '17).

<http://dx.doi.org/10.1145/3092919.3092920>.

NPAR'17, Los Angeles, CA, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. 978-1-4503-5081-5/17/07...\$15.00

DOI: 10.1145/3092919.3092920

CCS CONCEPTS

• Computing methodologies → Non-photorealistic rendering; Image processing;

KEYWORDS

style transfer, stylization, convolutional neural networks, image-based artistic rendering, image processing, semiotics

ACM Reference format:

Amir Semmo, Tobias Isenberg, and Jürgen Döllner. 2017. Neural Style Transfer: A Paradigm Shift for Image-based Artistic Rendering? In *Proceedings of Expressive, Los Angeles, CA, USA, July 28-29, 2017 (NPAR'17)*, 13 pages.

DOI: 10.1145/3092919.3092920

1 INTRODUCTION

Non-photorealistic rendering (NPR) constitutes a highly active research domain of computer graphics that deals with the expression, recognition, and communication of complex image contents by means of information abstraction and highlighting (DeCarlo and Santella 2002; Gooch 2010; Hertzmann 2010; Lansdown and Schofield 1995). In particular, image-based artistic rendering (IB-AR) enjoys a growing popularity in mobile expressive rendering (Dev 2013; Winnemöller 2013) to simulate the appeal of traditional artistic styles and media for visual communication (Kyprianidis et al. 2013; Rosin and Collomosse 2013) such as pencil, pen-and-ink, oil paint, and watercolor. Classical IB-AR techniques typically model the design aspects that are involved with these artistic styles, i. e., to direct the smoothing and contour highlighting of image filtering,

¹CC BY-SA 2.0 license: <http://creativecommons.org/licenses/by-sa/2.0>

Figure 2: Overview of NST techniques and applications. Previous works used implementations of NSTs to perform generalized color and texture transfers, stylize videos, and provide means for casual creativity in mobile expressive rendering. Images © Risser et al. (2017), from Gatys et al. (2016b) © IEEE, from Selim et al. (2016) © ACM, all used with permission.

the approximation of image contents via rendering primitives (e. g., brush strokes, stipples), or an image segmentation. A more generalized approach has been introduced by example-based rendering (EBR), which employs machine learning or statistical models to emulate characteristics of artistic styles from visual examples (Kyprianidis et al. 2013). Previous techniques in EBR, however, typically require analogous style and content pairs for training (Hertzmann et al. 2001) or only inform low-level image features for texture transfers, thus limiting its application and creative control over the design aspects. Advancements in deep learning and convolutional neural networks (CNNs) demonstrated that these technical limitations can be alleviated as follows:

- (1) Deep CNNs are able to accurately classify high-level image contents across generalized data sets (Simonyan and Zisserman 2015).
- (2) Layers of pre-trained deep CNNs can be activated to match content and style statistics, and thus perform a *neural style transfer* (NST) between *arbitrary* images (Gatys et al. 2016b) (Figure 1).

To this end, we argue that deep learning denotes a key technique in the chronology of IB-AR (Kyprianidis et al. 2013), as it makes—for the first time—a generalized style transfer practicable. First applications demonstrate this process using the example of color and texture transfers as well as casual creativity systems and services (Figure 2). To provide a sophisticated paradigm shift for IB-AR, however, we believe that NSTs need to mature from color and texture transfers to interactive tools that consider the design aspects and mechanisms involved in artwork production, i. e., to ease the visual expression of artists, non-artists (i. e., general public), and scientists (Gooch et al. 2010; Isenberg 2016; Salesin 2002).

In this paper we discuss the potentials and challenges of NST for IB-AR. In the following section, we first provide a conceptual overview for (neural) style transfer and show how the design process differs from classical IB-AR paradigms (Section 2). Next, we provide a semiotic structure for IB-AR that combines design aspects and mechanisms of artwork production with well-established design principles of NPAR (Section 3). We then use this structure to categorize current (neural) style-transfer techniques (Section 4) and derive a technical research agenda for NST (Section 5) including potential mutual inclusions with other IB-AR paradigms such as image filtering (Figure 1). With this research agenda we shed light on how NSTs may contribute to deal with the grand challenges of NPAR put forth by Salesin (2002) and revisited by Gooch et al. (2010), and how they can be evolved as interactive tools that consider mechanisms of artwork production. Finally, we identify potential future applications such as casual creativity (Section 6).

2 ARTISTIC STYLE TRANSFER WITHIN THE TAXONOMY OF IB-AR TECHNIQUES

IB-AR is related to the processes of visual abstraction that are involved in the creation of general artworks (Hertzmann 2010; Ma 2002) and used to express uncertainty, communicate abstract ideas, and evoke the imagination (Gooch et al. 2010) by addressing the rational, emotional, and cognitive qualities of the human mind (Halper et al. 2003; Hertzmann 2010). For an effective visual abstraction, the separation of *content* from *style* is thus considered to be a key factor to allow us to distinguish between the mechanisms used for capturing the essence of an image, on the one side, and the design aspects that drive the aesthetic appeal to stimulate human senses (Gooch et al. 2010; Salesin 2002), on the other side. To this end, research in

Figure 3: Overview of style transfer concepts, which differ in the way artistic styles are modeled or transferred: heuristics-based algorithms (left) and style transfers based on image statistics or analogies (middle) require explicit modeling or training phases prior to application, whereas NSTs (right) combine both aspects in a single phase.

the field of IB-AR has been devoted to deduce the design aspects of an *artistic style* that are involved in artwork production:

Definition “Artistic Style”: The constant form—and sometimes the constant elements, qualities, and expression—in the art of an individual or a group.

— Meyer Schapiro (Schapiro 1994)

IB-AR implementations typically require programmers to model the design space as well as the defining and distinguishing characteristics of an artistic style. Here, we see two general approaches which align with Kyprianidis et al.’s (2013) taxonomy as follows:

- (1) **Heuristics-based Algorithms:** Paradigms that are based on rendering functions, which are implemented by a domain expert who explicitly models individual artistic styles and its correspondent design aspects or mechanisms. This group basically comprises stroke-based rendering, region-based techniques, image processing and filtering, and may also account for physically-based simulations.
- (2) **Style Transfer Algorithms:** Example-based rendering which is directed to learn or reproduce artistic styles from visual examples (ground-truth data sets). This type often comprises statistical models and optimization schemes to balance aspects of *content* and *style* in the stylized output.

Prominent examples of heuristics-based algorithms are the stroke-based rendering approach of Hertzmann (1998), the cartoon pipeline of Winnemöller et al. (2006), and the watercolor system of Bousseau et al. (2006). For style transfer algorithms, by contrast, the literature primarily distinguishes between EBR techniques that transfer color or texture (Kyprianidis et al. 2013). However—with the maturation of machine learning—we believe that this strict separation is no longer practicable because color and texture represent only two out of many variables to define the composition of artistic styles, and deep learning enables NSTs to abstract from applications (e. g., color/texture transfer). To this end, we conjecture that it is worthwhile to provide a process-oriented taxonomy for EBR that reflects how artistic style transfers are modeled or technically implemented. Given artistic works as ground-truth data, we argue that three concepts may distinguish current and future EBR techniques (Figure 3):

- I. **Style Transfer using Image Statistics:** Techniques that balance content and style of two separate inputs using statistical models. Prominent examples are histogram-based

color transfers that equalize the mean and variance between content and style images (Neumann and Neumann 2005; Reinhard et al. 2001).

- II. **Style Transfer using Image Analogies:** Techniques that use image pairs for training—a source image and an artistic depiction of this image—i. e., to learn an analogous transformation such that content images can be transformed into an artistic rendering of similar visual style (Hertzmann et al. 2001).
- III. **Style Transfer using Neural Networks:** Techniques that employ neural networks to separate and recombine the content and style of arbitrary inputs. Typically, loss functions are minimized iteratively to balance the components of *style* and *content* in the output (Gatys et al. 2016b), or train feed-forward neural networks for linear image transformation (Johnson et al. 2016a,b).

We believe this classification helps to organize EBR techniques by their technical foundation and underpins the maturation from application-specific (e. g., color transfers) towards generalized style transfers. In the following section, we define design aspects and mechanisms important for implementing these three concepts.

3 A SEMIOTIC STRUCTURE FOR ARTISTIC STYLE TRANSFER

Semiotics deals with the study of symbols and how they communicate image contents or information in a meaningful way (Bertin 2010). In artwork production, elements of design are considered to be fundamental aspects of pictorial semiotics (Rudner 1951), whose mutual impact define the “composition” of an artwork, and thus its artistic style. Therefore, we believe that the transfer of proven design aspects and mechanisms of artwork production to modern media and imaging technologies, and the development of new artistic styles are key challenges for current and future research. In IB-AR theory (Hertzmann 2010), a semiotic structure that considers these design aspects and the mechanisms of interactive NPAR has not been formulated yet. We believe, however, that such a structure is essential to provide developers of NPAR techniques with the conceptual means to help them compose and extend artistic styles as well as evolve (neural) style transfers as interactive tools that ease the visual expression of artists, non-artists and scientists for illustrative visualization (Gooch et al. 2010; Isenberg 2016; Salesin 2002). We thus formulate a semiotic structure that is based on graphic

Figure 4: Semiotic structure comprising graphical core variables and mechanisms that may be considered for style transfers.

semiology principles of Bertin (2010) and MacEachren et al. (2012) that provide a theoretical foundation to visualization (Figure 4).

The visual variables described by Bertin (2010) and MacEachren et al. (2012), however, cannot fully express the unique requirements of interactive media and systems (e. g., animation, video, interactive parameterizations). We thus extend the classification by the concepts of filtering and perception to consider interactivity, level of abstraction, and coherence/continuity issues of NPAR as well. This way, user involvement can be considered as a key mechanism for maintaining an iterative feedback loop between a system—as design instance implementing NPAR techniques—and the user’s requirements—as consumer/artist. In particular, it is directed to interactively adjust the semiotic structure that defines aspects of modeling, filtering, composition and perception (Figure 4):

- 1. Modeling Aspects:** They deal with encoding real-world phenomena as *color maps*, and complementary information as *feature maps* (e. g., results of an image segmentation, saliency analysis, optical flow estimation) and *geometry maps* (e. g., depth).
- 2. Filtering Aspects:** They are used to select and apply different configurations of composition variables according to image *location*, *color*, or *feature*. Filtering aspects should provide effective control to globally and locally adjust the *level of abstraction*. Examples are the luminance-based placement of stipples (Martin et al. 2015), the location-dependent placement of contour lines (Cole et al. 2008), and feature-guided image filtering using orientation information (Kyprianidis and Döllner 2008).
- 3. Graphical Elements:** These elements comprise rendering primitives such as *points*, *lines*, *areas*, and generalized *2D elements*. They may also define rendering paths or locations for texturing, e. g., stippling, contour-lining, and the decoration of image segments.
- 4. Graphical Variables:** They refer to the illusion of physical mass and density (*form*), image regions with well-defined boundaries (*shape*), the *size* of graphical elements, and *color* including *brightness* as phenomena of light and human visual perception. Prominent examples refer to rendering with reduced color palettes and at multiple scales (Kyprianidis et al. 2013).
- 5. Design Mechanisms:** They deal with the surface character and relationships among image features with respect to position and direction (*space/texture*), *transparency* to infer

Table 1: Overview of image-based artistic style transfer techniques and how they relate to semiotic aspects. Current NST techniques apparently lack to model graphical elements/variables and provide interactive (creative) control.

Publication	Color Maps	Feature Maps	Geometry Maps	Location-based Filtering	Color-based Filtering	Feature-based Filtering	Point / Line / Area	Color / Brightness	Form / Shape / Size	Space / Texture	Transparency	Orientation	Shading / Shadows	Crispness / Resolution	Coherence / Continuity	Pictorial Cues	User Interaction
Image Statistics																	
Arbelot et al. (2016)	x	x			x			x		x						x	x
Chang et al. (2015)	x	x			x			x		x							x
Kim et al. (2009)	x			x			x	x	x	x			x	x			
Maciejewski et al. (2008)	x			x			x	x	x	x			x	x			
Martin et al. (2011)	x			x			x	x	x	x			x	x			x
Neumann Broth. (2005)	x						x	x									x
Pouli & Reinhard (2011)	x																
Reinhard et al. (2001)	x							x									
Wu et al. (2013)	x	x						x									
Xiao & Ma (2009)	x	x						x									
Yang et al. (2017)	x	x						x									
Image Analogies																	
Ashikhmin (2003)	x	x								x							x
Bénard et al. (2013)	x	x	x		x				x	x	x		x			x	x
Berger et al. (2013)	x			x			x										
Efros & Freeman (2001)	x									x							
Fiser et al. (2016)	x		x							x			x			x	x
Hashimoto et al. (2003)	x	x								x							
Hertzmann (2001)	x									x							
Hertzmann et al. (2002)	x			x	x	x											
Lee et al. (2011)	x	x	x							x							
Wang et al. (2013)	x						x				x						
Zhao & Zhu (2011)	x			x	x	x											
Neural Networks																	
Anderson et al. (2016)	x	x								x						x	
Champandard (2016)	x	x		x	x					x							
Chen & Schmidt (2016)	x									x							
Dumoulin et al. (2017)	x									x							
Gatys et al. (2016a)	x				x					x							
Gatys et al. (2016b)	x				x					x							
Gatys et al. (2016c; 2017)	x	x		x	x				x								
Gupta et al. (2017)	x	x								x						x	
Huang & Belongie (2017)	x									x				x			
Iizuka et al. (2016)	x							x		x							
Johnson et al. (2016a)	x									x				x			
Li & Wand (2016)	x									x							
Liu et al. (2017)	x	x	x	x						x							x
Risser et al. (2017)	x	x		x	x					x							
Ruder et al. (2016)	x	x								x						x	
Selim et al. (2016)	x	x			x				x	x						x	
Taigman et al. (2016)	x									x							
Ulyanov et al. (2016a)	x									x							
Ulyanov et al. (2017a)	x	x			x					x							
Ulyanov et al. (2016b)	x									x

Figure 5: Adjusting the level of abstraction: (top) using different relative weightings between content/style reconstruction, (bottom) matching the style representation with layer subsets of the VGG-Network. Images from Gatys et al. (2016b), © IEEE, used with permission. Style image by Wassily Kandinsky is in the public domain (source: Google Art Project).

color blending via overdraw or layering, the *orientation* of graphical elements, the *shading* and *lighting* conditions, and the *crispness/resolution* of image features. Previous works deal with mechanisms for stylized shadows (DeCoro et al. 2007), the orientation and layering of curved brush strokes (Hertzmann 1998), and low-pass image filters.

6. **Perceptual Aspects:** IB-AR typically aims to reproduce a hand-drawn look, where “distracting flickering and sliding artifacts” for animated scenes (e. g., virtual environments, video) should be minimized (Bénard et al. 2011). Bénard et al. (2011) propose this challenge to be a concurrent fulfillment of three goals: *flatness*, *motion coherence*, and *temporal continuity*. In addition, we conjecture that *pictorial cues* are important perceptual aspects because artists often carefully consider linear perspective, occlusion, and texture gradients to infer depth in their artworks.

The mutual impact of these aspects define the individual artistic style and composition, and thus should be considered when designing and implementing style transfers. In particular, we argue that *color* and *texture* are only two semiotic aspects most techniques currently serve. By contrast, a “successful” modeling approach should consider the distinctive design aspects and mechanisms involved in a particular artistic style, i. e., with respect to the rendering functions, optimization functions for image statistics and analogies, or loss functions for neural networks (Section 2).

4 SEMIOTICS-ORIENTED OVERVIEW OF ARTISTIC STYLE TRANSFER TECHNIQUES

In this section we now provide an overview² of existing techniques with respect to the three concepts of style transfer and show how they consider aspects of the semiotic structure (Figure 4). We provide a summary of this discussion in Table 1.

²The overview gives a non-exhaustive general picture of how semiotics are considered in current research, we expect it to be extended with future research.

4.1 Style Transfer using Image Statistics

Most techniques using image statistics are designed to perform *color transfers*. Here we can only mention the most representative works and refer to Faridul et al.’s (2014) survey for a comprehensive overview. The majority of techniques equalizes the mean and variance of a style and content image to control *color distributions* via luminance-based (Reinhard et al. 2001) or HSL-based (Neumann and Neumann 2005) histograms. Extensions integrate *feature maps* to consider local information as well, such as image segmentation (Wu et al. 2013; Xiao and Ma 2009), edge-aware texture descriptors (Arbelot et al. 2016), and semantics (Yang et al. 2017) to colorize grayscale images. With *interactive* methods it is also possible to maintain control over a set of colors that is involved in palette-based color transfers (Chang et al. 2015; Pouli and Reinhard 2011).

Another classical application for image statistics can be found in image stippling (Martin et al. 2017). Here, patterns are learned and applied through example using statistical texture measures (Kim et al. 2009; Maciejewski et al. 2008), modeling aspects such as the location of *points* (stipples), *texture*, *shading*, and *resolution*, which should depend on the spatial *size* of the output image. Martín et al. (2011) evolve these methods towards a “*scale-dependent*, example-based stippling technique that supports both low-level stipple *placement* and high-level *interaction* with the stipple illustration.” These methods are prime examples for how style transfers can be implemented on a primitive level, considering graphical elements explicitly rather than texture patches.

4.2 Style Transfer using Image Analogies

Most style transfer techniques defined by image analogies are based on *texture transfers*. Its basic idea is to copy image patches from a style image to a content image in a way that locally shares and minimizes pixel differences in the content image, thereby using a smoothness constraint to provide similarity with adjacent textures (Efros and Freeman 2001). Hertzmann (2001) defines this as an optimization problem by learning the analogous transformation of a style/ground-truth image pair (A, A') and applying it to a content image B to obtain a stylized output B' such that $A : A' :: B : B'$.

Ashikhmin (2003) provides conditions for how to integrate user-defined *feature maps* to adjust parameter values of the texture transfer. The approach can also be used to learn stroke placements for contour-lining (Hertzmann et al. 2002) in domains such as portrait sketches using templates (Zhao and Zhu 2011) and modeling image features at multiple scales for *level of abstraction* rendering (Berger et al. 2013). Further extensions use edge and orientation information encoded in *feature maps* to control the placement of texture patches (Lee et al. 2011) and individual brush strokes (Wang et al. 2013), learn multiple styles and stroke patterns for portrait sketching and painting (Berger et al. 2013; Zhao and Zhu 2011), and estimate motion using flow fields to stabilize *temporal coherence* (Hashimoto et al. 2003). Bénard et al. (2013) propose a sophisticated system for artists that performs style transfers for animations using orientation, velocity, and geometry information of 3D models to direct the transfer with *shading* and *lighting* conditions, and to ensure *temporal and style continuity*. In addition, they support overdraw and partial *transparency* using a layering approach explicitly defined by the artist. Most of these works, however, typically consider only luminance- or color-guidance texture transfers, yet other information may be considered as well such as *illumination* as shown by Fišer et al. (2016) for stylized 3D models.

4.3 Style Transfer using Neural Networks

To discuss this subfield, we draw on Gatys et al. (2016b) definition of NSTs. Given a style image, a content image and a loss network, e. g., VVG-16 (Simonyan and Zisserman 2015), that is used to define several loss functions to measure the difference between the output image and a target image, one can compute an output image by minimizing a weighted combination of the loss functions. Gatys et al. (2016b) initially define perceptual loss functions that control feature and style reconstructions to balance the components of *content* and *style*, and control spatial smoothness by regularizing the total variation, then solve the optimization problem using L-BFGS (Figure 5). Besides texture transfers, this approach can be employed to perform sophisticated color transfers as well, e. g., to colorize grayscale images (Iizuka et al. 2016). Because this generalized style transfer employs back-propagation and combines learning and application in a single phase, we denote it as an iterative approach and distinguish it from the approach that separates learning from application to train a feed-forward neural network.

Iterative Approaches. Extensions of Gatys et al.'s (2016b) work primarily define additional loss functions to control the output's composition. MRFs loss functions, for instance, can be used as a local constraint to provide a more accurate *texture* patch matching and blending (Li and Wand 2016), histogram losses may produce outputs that statistically match style images more accurately (Risser et al. 2017), and a depth loss function to consider the spatial distribution of image features (Liu et al. 2017). Further, a temporal loss function based on optical flow can be used to stabilize *temporal coherence* when applied on a per-frame basis to video (Anderson et al. 2016; Gupta et al. 2017; Ruder et al. 2016; Selim et al. 2016). A few works controlled perceptual factors locally by considering *feature maps* using semantics-based image segmentation, such as to subdivide the optimization problem of NST to local image regions (Champandard 2016) or facial regions of portrait images (Selim et al. 2016) to

Figure 6: Adjustments of loss functions for NST to control graphical variables of the semiotic structure. Images © Gatys et al. (2016c), used with permission.

provide semantically more accurate transfers. Some enhancements also considered composition variables of the semiotic structure such as *color*, *size*, and *location-based filtering* by introducing control measures (Gatys et al. 2016a,c, 2017) (Figure 6). We see these works as a starting point to evolve NSTs as interactive tools for IB-AR that facilitate creative expression, which we discuss below.

Feed-forward Approaches. The solving of NSTs optimization problems is computationally extensive. Some approaches thus provide approximations by computing the weights of a feed-forward neural network. Here, test images sets, e. g., ImageNet (Krizhevsky et al. 2012) or MS-COCO (Lin et al. 2014), are often used in a training phase performed once per artistic style, after which the obtained generative convolutional networks are used for linear image transformation (Johnson et al. 2016a,b; Ulyanov et al. 2016a, 2017a,b). Johnson et al. (2016a; 2016b) and Ulyanov et al. (2016a) showed that these networks can be three orders of magnitude faster than the iterative approach. The output quality of these approaches can be further improved by employing network layers for (adaptive) instance normalization (Huang and Belongie 2017; Ulyanov et al. 2016b) that align the mean and variance of features of the content and style images. Conceptual limitations of these approaches, however, lie in the limited level of detail: style characteristics are generalized and not balanced for a unique style/content image pair (Figure 7). Alternative approaches either employ simpler loss functions with only local matching constraints, e. g., using a single layer of a pre-trained loss network (Chen and Schmidt 2016), or learn multiple styles or generative networks at once (Dumoulin et al. 2017; Zhang and Dana 2017) to improve versatility.

5 A TECHNICAL RESEARCH AGENDA FOR NEURAL STYLE TRANSFER

NST is a relatively new field of research but has already shown promising results for generalized style transfers. We believe its future directions can be defined in the context of some of the grand challenges of NPAR (Gooch et al. 2010; Isenberg 2016; Salesin 2002),

Figure 7: Comparison of iterative approaches for NST (deepart.io and Pikazo) with feed-forward approaches (BN: batch normalization, IN: instance normalization). Images © Ulyanov et al. (2017b), used with permission.

i. e., its combination with other IB-AR paradigms for providing algorithmic aesthetics, improving the fidelity in reproducing and extending artistic styles towards new forms of art, and its parameterization to evolve as interactive tools that “support full design cycle” (Salesin 2002) and ease visualization tasks. With these challenges and semiotics-oriented overview of Section 4 in mind, we thus propose the following technical research agenda.

Proposal 1: Semiotics-based Loss Functions

Current NST techniques primarily depend on color statistics for style transfer, but model color as a mutual inclusion and effect of multiple composition variables. However, we believe that loss functions need to be defined for individual composition variables and controlled filtering-wise by providing modeling information that, e. g., encode how the size, shape, orientation, transparency, shading, and shadows are aligned with the contents of a style image. For instance, stroke-based rendering models the image composition by placing, orienting, and layering individual brush strokes as graphical elements (Kyprianidis et al. 2013). Typically, techniques estimate image flow (Wang et al. 2004; Yan et al. 2008; Zeng et al. 2009) or derive local surface properties (Sloan et al. 2001) to guide brush strokes with the orientation of image features or the shading and lighting conditions (Fišer et al. 2016). Together with texture layering, e. g., of painterly art maps or dictionaries (Yan et al. 2008; Zeng et al. 2009), they provide better quality in preserving fine texture details and modeling style characteristics induced by *form*, *shape*, and *orientation*. For the latter, we believe a similar loss used for temporal consistency (Gupta et al. 2017; Ruder et al. 2016)—but based on image orientation information—could help guide the texture transfer. There is also demand to explicitly model semiotic aspects that consider feature semantics. Here, Figure 8 exemplifies some limitations that NSTs currently face for three artistic styles:

- **Divisionism** represents images by regularly aligned rendering primitives, e. g., brush strokes that optically compose image features when viewed from distance. Because

of its analogy to patch-based texturing, divisionism can be modeled quite accurately by current loss functions.

- **Cubism** depicts subjects using simplified shapes and forms for composition, which are often portrayed using multiple perspectives. Here, NST techniques would need to infer geometric transformations and match geometric representations, e. g., as practiced by Mital et al. (2013), in correspondence with the color similarity.
- **Pop Art** typically composes images by thick outlines, bold solid colors and Ben-Day dots. Here, current NST techniques face multiple limitations in reproducing shape, preserving the semantic composition, and style characteristics such as the regularity and color inversion of halftoning.

The examples of cubism and pop art demonstrate that the coupling of individual semiotic aspects with the semantics of content and style images requires sophisticated rule-based algorithms. Eventually, this would lead to couple feature-level engineering with the architecture engineering approach of deep learning.

Proposal 2: Combination with IB-AR Paradigms

Local effects and phenomena of traditional artistic media such as oilpaint, pencil, or watercolor at high-fidelity and resolution are still hard to reproduce by NSTs. Here, we believe that NSTs may be used as one of multiple processing stages in IB-AR, and combined with the knowledge and algorithms of other paradigms. NSTs would thus not operate at the lowest level of detail, but as a first stage that introduces higher-level abstractions—to be followed by a low-level, established technique to simulate drawing media and their interplay with substrates. For instance, specialized line drawing algorithms can be used to detect and stylize (salient) edges, e. g., via difference-of-Gaussians (Winnemöller et al. 2012), edge-preserving filtering for noise reduction (Kyprianidis et al. 2013), and the constraints of stroke-based rendering to control the placement of graphical elements, e. g. based on luminance to direct (tonal) art maps for pencil rendering (Lee et al. 2006; Praun et al. 2001) or structure grids for feature-guided stippling (Son et al. 2011) to

Figure 8: Limitations of current NSTs for simulating artistic styles. Style images of Robert Delaunay and Juan Gris are in the public domain (source: Wikimedia Commons). Content image from Son et al. (2011) © Elsevier Inc, used with permission.

Figure 9: Comparison of NSTs with heuristics-based algorithms for stippling and hatching. Content image and stippling from Son et al. (2011) © Elsevier Inc, style images © Randy Glass and © Igor Lukyanov, all used with permission.

avoid the artifacts from pure NSTs shown in Figure 9. In Figure 10 we show results of a case study, where image filtering is employed in a post-processing stage to NST to simulate local effects such as edge darkening, pigment density variation, and wet-in-wet of watercolors quite accurately (Bousseau et al. 2006; Wang et al. 2014), whereas flow-based Gaussian filtering with Phong shading is used to filter low-level noise and create smooth continuous oilpaint-like texture effects (Hertzmann 2002; Semmo et al. 2016b) In both cases we used the abstract style of Pablo Picasso’s “La Muse” to generate an effect of higher-level abstraction, before adding mentioned filters to simulate the respective low-level, local paint characteristics.

Proposal 3: New Forms of Styles

Gooch et al. (Gooch et al. 2010) provided an overview of NPAR research through Heinlein’s maturation model, and argue that NPAR has left the first stage—emulating and imitating artistic styles—, evolved towards the second stage by optimizing the performance of the (used) technology, and is about to move towards the last stage, where the technology becomes seamless and almost transparent. In this respect, we believe that NST provides new opportunities for the first two stages, but needs to “incorporate elements such as interaction, collaboration, human perception and cognition” (Gooch et al. 2010) to approach the third stage. In particular, here we see two potential use cases for NST. First, modifying learned artistic styles by providing mechanisms to specify transfer or loss functions that change particular design aspects or variables. Second, performing a style transfer by taking rule-based algorithms into account, i. e., to learn styles not only from style images but also a set of descriptions how an artistic style should look like, which makes new forms of styles—that have never been *seen* before—practicable.

Proposal 4: Providing Interactivity

Recently, Isenberg (2016) argued that EBR approaches have the potential to enable users to provide “both higher-level interaction and low-level control”—suggesting that this allows us to create both interaction environments for artists who need a wide range of low-level to high-level control and for non-artists whose interaction needs are likely easier satisfied with high-level interactions such as the application of filters. Many traditional EBR approaches, however, have relied on a close relationship between input style and input context, e. g., for hatching (Gerl and Isenberg 2013). NSTs have the potential to address this very problem: styles are more easy to capture and thus the interactive application of style becomes easier. So far, however, NST are typically treated like a “black box”, supporting only the high-level application of a captured style. To enable the interaction spectrum that Isenberg (2016) calls for, it would be necessary to integrate more local control. Artists need to be able to affect the result on a semantic level: controlling how larger regions are treated, change groups of marks, and even adjust a single mark. One approach could be to provide loss functions that operate on primitive-level and single design aspects as well, e. g., graphical elements such as brush strokes in a style image. For example, Figure 9 demonstrates how a purely global NST approach fails in several regions, and local control such as the change of an underlying directional field, e. g., as practiced by (Salisbury et al. 1994), seems to be missing.

Moreover, it is important to consider the input from several style images, which is technically demonstrated by Johnson et al. (2016a; 2016b) for blending multiple styles. This could be extended to either learn a particular technique/style or even an artist’s design principles more reliably, or it could be used to combine two different

Figure 10: Post-process image filtering to reduce low-level noise and inject paint characteristics. NST results are combined with watercolor rendering (Bousseau et al. 2006; Wang et al. 2014) and oilpaint filtering (Semmo et al. 2016b). Content image by Frank Köhntopp is in the public domain.

styles in the same target image. For example for the latter, illustrations that combine different depiction styles to steer attention and create focus and context view would be an important application domain. Such an approach, however, would need local control or a semantic/semiotic processing of the content image by the NST algorithm, e. g., as is partially practiced by Gatys et al. (2016c; 2017) using feature maps, and interactive performance for immediate visual feedback, but which is currently a strong limitation of iterative NST techniques.

Proposal 5: Supporting Visualization Tasks

Semiotics are inherently linked with the theory of (information) visualization (Bertin 2010). In particular, style transfers have been commonly used in illustrative visualization (Rautek et al. 2008), e. g., for the stylization of lines to depict flow (Everts et al. 2015), to make phenomena—hidden in complex data sets—visible to the human mind. However, effective visualization must also “enable analysis of the supplied information, while easing the cognitive burden of a user” (Gooch et al. 2010). NSTs based on deep CNNs emulate functionalities of the visual cortex by solving tasks through hierarchical processing (DiCarlo et al. 2012), but need to be performed in a context-dependent manner, e. g., with respect to a user’s task and data domain, for effective visualization. Here, we imagine the development of toolboxes or palettes of illustration styles that can be interactively applied by professional illustrators, in a way that considers an interaction spectrum from low-level to high-level controls (Isenberg 2016). For example, a palette for computer-supported hatching and stippling could be provided that alleviates some of the tediousness of manual processes, but that includes support for higher-level illustration processes, e. g., (Martin et al. 2011), where NSTs could suggest regions to be filtered or regions to be contrast-adjusted. The layers of deep CNNs that capture multiple levels of abstraction could be interactively used for this purpose to direct the interactive visualization/illustration process. Finally, we believe that, with the generalized application of NSTs, more complex artistic styles of several visualization domains could be served, such as medical imaging or cartography, but which requires NSTs to consider the semantics of style and content images (e. g., as shown

for portrait images (Selim et al. 2016)), and data-domain specific design mechanisms such as generalization (MacEachren 1995).

Proposal 6: Evaluation

The evaluation of aesthetics and practical benefits for illustration or visualization tasks remains an important issue in IB-AR (Gooch 2010; Hall and Lehmann 2013; Hertzmann 2010; Isenberg 2013). For effective comparison of NST techniques, we believe there is demand for a standardized benchmark image set such as the general NPAR set provided by Mould and Rosin (2016).

With respect to aesthetic evaluation, Salesin (2002) and Gooch et al. (2010) raised the issue of a “Turing Test” that determines if CG imagery can be indistinguishable from imagery produced by humans. While the utility of such a test is being debated (Hall and Lehmann 2013), some authors have included respective questions in their evaluations (Gatys et al. 2016b; Isenberg et al. 2006). Gatys et al. for instance, evaluated their NST technique (2016b) in a preliminary choice experiment, asking participants to find the hand-painted images in a set of 10 hand-painted/NST image pairs. The average of their 45,000 participants answered 6.1 image pairs correctly.³ With the further consideration of semiotic aspects, in particular filtering that includes semantics to resolve incoherences in color transfers, it would be great to gather more information such as response time and eye fixations to determine apparent locations or aspects of style incoherence—information that may be injected into the learning phase for improving a style transfer.

With respect to task efficiency, studies are required to determine if NSTs only copy low-level style aspects or if they also maintain higher-level semantics of image contents. These studies could also be used to determine to what degree NSTs introduce abstraction, whether the degree of abstraction can be intentionally controlled, and how it can be seamlessly interpolated for an interactive application as discussed above. In particular, the meaningful interaction with NSTs as tools for artists or scientists (e. g., with respect to illustrative visualization) requires investigation.

³According to Leon Gatys in his talk at CVPR 2016 on “Image Style Transfer Using Convolutional Neural Networks” (Gatys et al. 2016b).

6 APPLICATIONS

The shift from feature engineering towards architecture engineering⁴ of deep learning enables IB-AR to abstract from input data, and thus increase the general applicability in highly dynamic environments. Here, we see the following potentials for using NSTs.

Casual Creativity

NSTs have particularly enriched casual creativity applications (Winnemöller 2013) in ubiquitous environments such as mobile computing. This domain has largely been devoted to image filtering and processing to date, providing only constrained effects (Dev 2013). Prominent examples are the web service *deepart.io* and the iOS app *Prisma*—attracting 60 million users in three weeks—, which also started to establish their own social media communities for sharing and commenting on stylized outputs. We believe, however, that these apps have to evolve from “black box” solutions towards user-centric tools (Winnemöller 2013) to further promote visual expression. Here, a metaphor for on-screen parameter painting (Semmo et al. 2016a) may be used to tune hyperparameters of neural networks, while hiding the computational complexity.

Art Production

Salesin (2002) had envisioned the support of artists to be a major goal of NPAR, i. e., developing tools that make their life easier but that do not constrain their capabilities in visual expression (Isenberg 2016). We discussed in Section 5 that this requires NSTs to evolve as interactive tools. One example is the system by Fišer et al. (2016) in which artists are able to draw over a printed stencil, while their individual style is transferred in real-time onto 3D models, dealing with proper light propagation and auto-completion. Another example is the system for watercolor rendering with art-directed control of Montesdeoca et al. (2016; 2017), where the effects shown in Figure 10 (among others) can be controlled via on-screen painting. Here, a long-term goal would be to integrate NSTs in the production pipeline of feature films, e. g., as evaluated by Joshi et al. (2017) for *Come Swim*, reaching a quality level to assist the laborious production of fully painted animated films such as *Loving Vincent* (Mackiewicz and Melendez 2016) (Figure 11), e. g., with respect to temporal coherence and the placement of graphical elements such as brush strokes.

Teaching Art Classes

We see potentials to use NSTs for teaching purposes, i. e., to help study and explore artistic styles of famous artists or epochs. In particular, we consider semiotics-oriented loss functions (Section 5) as a key goal for providing algorithmic support at a high-level (e. g., texture transfer) and low-level (e. g., primitive-level transfer). This way, interactive art explorations could be feasible for children using (semi-)automatic transfers, e. g., using the semantics of two-bit doodles (Champandard 2016). A similar scenario can also be created for adults who could explore, e. g., the modeling, painting, and mixing of style invariances (e. g., brush size, pattern, etc.).

Figure 11: Comparison between emulating an artistic style via painting (oil on canvas) and a NST. Results from “Loving Vincent” © BreakThru Films, used with permission. Style image by Vincent van Gogh is in the public domain.

Exhibitions and Art Installations

Machine learning has gathered particular interest as an interactive component of exhibition and art installations, e. g., Tate Modern’s IK Prize 2016 winner *Recognition*⁵ uses pattern recognition to compare art to photojournalism. For instance, Adobe’s *Artistic Eye*⁶ uses NSTs to enable children transform their self-portraits into artistic renditions in the style of a museum’s exhibits, while Becattini et al. (2016) combined NSTs with art explorations, allowing users to scan exhibits and transfer their style to user-defined images.

7 CONCLUSION

Deep learning has opened new possibilities for IB-AR to make a generalized style transfer practicable. On the one hand, NSTs provide new potentials for using IB-AR in context-sensitive and creative application domains, such as casual creativity apps for mobile expressive rendering and production tools for feature films. On the other hand, NSTs currently provide only “black box” solutions from a HCI point-of-view: research (so far) has mainly focused on tuning hyperparameters of deep neural networks. To this end, we propose a semiotic structure to provide developers of NST techniques with the conceptual means of artworks production to help them compose and extend artistic styles, as well as consider design aspects and mechanisms for evolving NSTs as interactive tools. In particular, we hope that this structure helps researchers to identify requirements for semiotics-based loss functions, combine NSTs with the knowledge of other IB-AR paradigms, promote completely new artistic styles, and assist applications in illustrative visualization.

Finally, we argue that semiotics can be considered for defining artistic style and used to systematically evaluate NST techniques. Eventually, this evaluation should also account for the application space, level of interactivity, and audience including the user’s context and environment, skills and competence, and the purpose of

⁴Stephen Merity. 2016. In deep learning, architecture engineering is the new feature engineering. http://smerity.com/articles/2016/architectures_are_the_new_feature_engineering.html. Last followed: 04/09/2017.

⁵Tate IK Prize 2016. <http://www.tate.org.uk/about/projects/ik-prize-2016>. Last followed: 04/09/2017.

⁶Adobe Artistic Eye. <http://blogs.adobe.com/conversations/2017/03/de-youngsters-photos-get-the-look-of-masterpieces.html>. Last followed: 04/09/2017.

artistic rendering, e. g., the user's task—conditions that affect the "success" of a NST. For example, while a real-time feed-forward NST of texture and color as semiotic aspects may provide hallucination results of sufficient quality in mobile expressive rendering, artists typically wish to have full control over each individual semiotic aspect involved in the composition and transfer of artistic styles.

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their comments and suggestions on how to improve the paper. This work was partly funded by the Federal Ministry of Education and Research (BMBF), Germany, for the AVA project 01IS15041B.

REFERENCES

Meta/Surveys

- Pierre Bnard, Adrien Bousseau, and Jolle Thollot. 2011. State-of-the-Art Report on Temporal Coherence for Stylized Animations. *Computer Graphics Forum* 30, 8 (Dec. 2011), 2367–2386. doi: 10.1111/j.1467-8659.2011.02075.x
- Jacques Bertin. 2010. *Semiology of Graphics: Diagrams, Networks, Maps*. Esri Press, Redlands, California. <http://esripress.esri.com/display/index.cfm?fuseaction=display&websiteID=190>
- Doug DeCarlo and Anthony Santella. 2002. Stylization and Abstraction of Photographs. *ACM Transactions on Graphics* 21, 3 (July 2002), 769–776. doi: 10.1145/566654.566650
- Kapil Dev. 2013. Mobile Expressive Renderings: The State of the Art. *IEEE Computer Graphics and Applications* 33, 3 (May/June 2013), 22–31. doi: 10.1109/MCG.2013.20
- James J. DiCarlo, Davide Zoccolan, and Nicole C. Rust. 2012. How Does the Brain Solve Visual Object Recognition? *Neuron* 73, 3 (Feb. 2012), 415–434. doi: 10.1016/j.neuron.2012.01.010
- Hasan Sheikh Faridul, Tania Pouli, Christel Chamaret, Jrgen Stauder, Alain Trmeau, and Erik Reinhard. 2014. A Survey of Color Mapping and its Applications. In *Eurographics State of the Art Reports*. Eurographics Association, Goslar, Germany, 43–67. doi: 10.2312/egst.20141035
- Amy A. Gooch. 2010. Towards Mapping the Field of Non-Photorealistic Rendering. In *Proc. NPAR*. ACM, New York, 159–164. doi: 10.1145/1809939.1809958
- Amy A. Gooch, Jeremy Long, Li Ji, Anthony Estey, and Bruce S. Gooch. 2010. Viewing Progress in Non-Photorealistic Rendering Through Heinelein's Lens. In *Proc. NPAR*. ACM, New York, 165–171. doi: 10.1145/1809939.1809959
- Peter Hall and Ann-Sophie Lehmann. 2013. Don't Measure—Appreciate! NPR Seen Through the Prism of Art History. In *Image and Video based Artistic Stylisation*, Paul Rosin and John Collomosse (Eds.). Computational Imaging and Vision, Vol. 42. Springer, London/Heidelberg, Chapter 16, 333–351. doi: 10.1007/978-1-4471-4519-6_16
- Nick Halper, Mara Mellin, Christoph S. Herrmann, Volker Linneweber, and Thomas Strothotte. 2003. Psychology and Non-Photorealistic Rendering: The Beginning of a Beautiful Relationship. In *Proc. Mensch & Computer*. Teubner Verlag, Wiesbaden, 277–286. doi: 10.1007/978-3-322-80058-9_28
- Aaron Hertzmann. 2010. Non-Photorealistic Rendering and the Science of Art. In *Proc. NPAR*. ACM, New York, 147–157. doi: 10.1145/1809939.1809957
- Tobias Isenberg. 2013. Evaluating and Validating Non-Photorealistic and Illustrative Rendering. In *Image and Video based Artistic Stylisation*, Paul Rosin and John Collomosse (Eds.). Computational Imaging and Vision, Vol. 42. Springer, London/Heidelberg, Chapter 15, 311–331. doi: 10.1007/978-1-4471-4519-6_15
- Tobias Isenberg. 2016. Interactive NPAR: What Type of Tools Should We Create?. In *Proc. NPAR*. Eurographics Association, Goslar, Germany, 89–96. doi: 10.2312/exp.20161067
- Jan Eric Kyprianidis, John Collomosse, Tinghuai Wang, and Tobias Isenberg. 2013. State of the "Art": A Taxonomy of Artistic Stylization Techniques for Images and Video. *IEEE Transactions on Visualization and Computer Graphics* 19, 5 (May 2013), 866–885. doi: 10.1109/TVCG.2012.160
- Jan Eric Kyprianidis and Jrgen Dllner. 2008. Image Abstraction by Structure Adaptive Filtering. In *Proc. TP.CG*. Eurographics Association, Goslar, Germany, 51–58. doi: 10.2312/LocalChapterEvents/TPCG/TPCG08/051-058
- John Lansdown and Simon Schofield. 1995. Expressive Rendering: A Review of Non-photorealistic Techniques. *IEEE Computer Graphics and Applications* 15, 3 (May 1995), 29–37. doi: 10.1109/38.376610
- Kwan-Liu Ma (Ed.). 2002. *Recent Advances in Non-Photorealistic Rendering for Art and Visualization*. SIGGRAPH Course Notes, Vol. 23. ACM, New York. <http://www.cs.ucdavis.edu/~ma/SIGGRAPH02/course23/>
- Alan M. MacEachren. 1995. *How Maps Work: Representation, Visualization, and Design*. Guilford Press, New York. <https://www.abebooks.com/products/isbn/>
- 9780898625899
- Alan M. MacEachren, Robert E. Roth, James O'Brien, Bonan Li, Derek Swingle, and Mark Gahegan. 2012. Visual Semiotics & Uncertainty Visualization: An Empirical Study. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec. 2012), 2496–2505. doi: 10.1109/TVCG.2012.279
- Domingo Martn, Germn Arroyo, Alejandro Rodrguez, and Tobias Isenberg. 2017. A Survey of Digital Stippling. *Computers & Graphics* (2017). To appear.
- David Mould and Paul L. Rosin. 2016. A Benchmark Image Set for Evaluating Stylization. In *Proc. NPAR*. Eurographics Association, Goslar, Germany, 11–20. doi: 10.2312/exp.20161059
- Peter Rauterk, Stefan Bruckner, Eduard Grller, and Ivan Viola. 2008. Illustrative Visualization: New Technology or Useless Tautology? *ACM SIGGRAPH Computer Graphics* 42, 3, Article 4 (Aug. 2008), 8 pages. doi: 10.1145/1408626.1408633
- Paul Rosin and John Collomosse (Eds.). 2013. *Image and Video based Artistic Stylisation*. Computational Imaging and Vision, Vol. 42. Springer, London/Heidelberg. doi: 10.1007/978-1-4471-4519-6
- Richard Rudner. 1951. On Semiotic Aesthetics. *The Journal of Aesthetics and Art Criticism* 10, 1 (Sept. 1951), 67–77.
- David H. Salesin. 2002. Non-Photorealistic Animation & Rendering: 7 Grand Challenges. Keynote talk at NPAR. (June 2002).
- Meyer Schapiro. 1994. *Theory and Philosophy of Art: Style, Artist, and Society*. Vol. 4. George Braziller, New York.
- Holger Winnemller. 2013. NPR in the Wild. In *Image and Video based Artistic Stylisation*, Paul Rosin and John Collomosse (Eds.). Computational Imaging and Vision, Vol. 42. Springer, London/Heidelberg, Chapter 17, 353–374. doi: 10.1007/978-1-4471-4519-6_17

Style Transfer using Image Statistics

- Benot Arbelot, Romain Vergne, Thomas Hurtut, and Jolle Thollot. 2016. Automatic Texture Guided Color Transfer and Colorization. In *Proc. NPAR*. Eurographics Association, Goslar, Germany, 21–32. doi: 10.2312/exp.20161060
- Huiwen Chang, Ohad Fried, Yiming Liu, Stephen DiVerdi, and Adam Finkelstein. 2015. Palette-based Photo Recoloring. *ACM Transactions on Graphics* 34, 4, Article 139 (July 2015), 11 pages. doi: 10.1145/2766978
- SungYe Kim, Ross Maciejewski, Tobias Isenberg, William M. Andrews, Wei Chen, Mario Costa Sousa, and David S. Ebert. 2009. Stippling by Example. In *Proc. NPAR*. ACM, New York, 41–50. doi: 10.1145/1572614.1572622
- Ross Maciejewski, Tobias Isenberg, William M. Andrews, David S. Ebert, Mario Costa Sousa, and Wei Chen. 2008. Measuring Stipple Aesthetics in Hand-Drawn and Computer-Generated Images. *IEEE Computer Graphics and Applications* 28, 2 (March/April 2008), 62–74. doi: 10.1109/MCG.2008.35
- Domingo Martn, Germn Arroyo, M. Victoria Luzn, and Tobias Isenberg. 2011. Scale-Dependent and Example-Based Stippling. *Computers & Graphics* 35, 1 (Feb. 2011), 160–174. doi: 10.1016/j.cag.2010.11.006
- Attila Neumann and Lszl Neumann. 2005. Color Style Transfer Techniques using Hue, Lightness and Saturation Histogram Matching. In *Proc. CAe*. Eurographics Association, Goslar, Germany, 111–122. doi: 10.2312/COMPAESTH/COMPAESTH05/111-122
- Tania Pouli and Erik Reinhard. 2011. Progressive Color Transfer for Images of Arbitrary Dynamic Range. *Computers & Graphics* 35, 1 (Feb. 2011), 67–80. doi: 10.1016/j.cag.2010.11.003
- Erik Reinhard, Michael Adhikmin, Bruce Gooch, and Peter Shirley. 2001. Color Transfer Between Images. *IEEE Computer Graphics and Applications* 21, 5 (Sept./Oct. 2001), 34–41. doi: 10.1109/38.946629
- Yi-Chian Wu, Yu-Ting Tsai, Wen-Chieh Lin, and Wen-Hsin Li. 2013. Generating Pointillism Paintings Based on Seurat's Color Composition. *Computer Graphics Forum* 32, 4 (July 2013), 153–162. doi: 10.1111/cgf.12161
- Xuezhong Xiao and Lizhuang Ma. 2009. Gradient-Preserving Color Transfer. *Computer Graphics Forum* 28, 7 (Oct. 2009), 1879–1886. doi: 10.1111/j.1467-8659.2009.01566.x
- Yue Yang, Hanli Zhao, Lihua You, Renlong Tu, Xueyi Wu, and Xiaogang Jin. 2017. Semantic Portrait Color Transfer with Internet Images. *Multimedia Tools and Applications* 76, 1 (Jan. 2017), 523–541. doi: 10.1007/s11042-015-3063-x

Style Transfer using Image Analogies

- Michael Ashikhmin. 2003. Fast Texture Transfer. *IEEE Computer Graphics and Applications* 23, 4 (July 2003), 38–43. doi: 10.1109/MCG.2003.1210863
- Pierre Bnard, Forrester Cole, Michael Kass, Igor Mordatch, James Hegarty, Martin Sebastian Senn, Kurt Fleischer, Davide Pesare, and Katherine Breeden. 2013. Stylizing Animation by Example. *ACM Transactions on Graphics* 32, 4, Article 119 (July 2013), 12 pages. doi: 10.1145/2461912.2461929
- Itamar Berger, Ariel Shamir, Moshe Mahler, Elizabeth Carter, and Jessica Hodgins. 2013. Style and Abstraction in Portrait Sketching. *ACM Transactions on Graphics* 32, 4, Article 55 (July 2013), 12 pages. doi: 10.1145/2461912.2461964
- Alexei A. Efros and William T. Freeman. 2001. Image Quilting for Texture Synthesis and Transfer. In *Proc. SIGGRAPH*. ACM, New York, 341–346. doi: 10.1145/383259.383296

- Jakub Fišer, Ondřej Jamriška, Michal Lukáč, Eli Shechtman, Paul Asente, Jingwan Lu, and Daniel Sýkora. 2016. *StyLit: Illumination-guided Example-based Stylization of 3D Renderings*. *ACM Transactions on Graphics* 35, 4, Article 92 (July 2016), 11 pages. doi: 10.1145/2897824.2925948
- Ryota Hashimoto, Henry Johan, and Tomoyuki Nishita. 2003. *Creating Various Styles of Animations Using Example-Based Filtering*. In *Proc. CGI*. IEEE Computer Society, Los Alamitos, 312–317. doi: 10.1109/CGI.2003.1214488
- Aaron Hertzmann, Charles E. Jacobs, Nuria Oliver, Brian Curless, and David H. Salesin. 2001. *Image Analogies*. In *Proc. SIGGRAPH*. ACM, New York, 327–340. doi: 10.1145/383259.383295
- Aaron Hertzmann, Nuria Oliver, Brian Curless, and Steven M. Seitz. 2002. *Curve Analogies*. In *Proc. EGWR*. Eurographics Association, Goslar, Germany, 233–246. doi: 10.2312/EGWR/EGWR02/233-246
- Bhautik Joshi, Kristen Stewart, and David Shapiro. 2017. *Bringing Impressionism to Life with Neural Style Transfer in Come Swim*. arXiv.org report 1701.04928. <https://arxiv.org/abs/1701.04928>
- Hochang Lee, Sanghyun Seo, and Kyunghyun Yoon. 2011. *Directional Texture Transfer with Edge Enhancement*. *Computers & Graphics* 35, 1 (Feb. 2011), 81–91. doi: 10.1016/j.cag.2010.11.008
- Tinghui Wang, John Collomosse, Andrew Hunter, and Darryl Greig. 2013. *Learnable Stroke Models for Example-based Portrait Painting*. In *Proc. British Machine Vision Conference*. BMVA Press, Durham, UK, Article 36, 11 pages. doi: 10.5244/C.27.36
- Mingtian Zhao and Song-Chun Zhu. 2011. *Portrait Painting using Active Templates*. In *Proc. NPAR*. ACM, New York, 117–124. doi: 10.1145/2024676.2024696

Style Transfer using Neural Networks

- Alexander G. Anderson, Cory P. Berg, Daniel P. Mossing, and Bruno A. Olshausen. 2016. *DeepMovie: Using Optical Flow and Deep Neural Networks to Stylize Movies*. arXiv.org report 1607.08022. <https://arxiv.org/abs/1605.08153>
- Federico Becattini, Andrea Ferracani, Lea Landucci, Daniele Pezzatini, Tiberio Uricchio, and Alberto Del Bimbo. 2016. *Imaging Novecento. A Mobile App for Automatic Recognition of Artworks and Transfer of Artistic Styles*. In *Proc. EuroMed*. Springer International, Cham, Switzerland, 781–791. doi: 10.1007/978-3-319-48496-9_62
- Alex J. Champandard. 2016. *Semantic Style Transfer and Turning Two-Bit Doodles into Fine Artworks*. arXiv.org report 1612.04337. <https://arxiv.org/abs/1603.01768>
- Tian Q. Chen and Mark Schmidt. 2016. *Fast Patch-based Style Transfer of Arbitrary Style*. arXiv.org report 1612.04337. <https://arxiv.org/abs/1612.04337>
- Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. 2017. *A Learned Representation For Artistic Style*. arXiv.org report 1610.07629. <https://arxiv.org/abs/1610.07629>
- Leon A. Gatys, Matthias Bethge, Aaron Hertzmann, and Eli Shechtman. 2016a. *Preserving Color in Neural Artistic Style Transfer*. arXiv.org report 1606.05897. <https://arxiv.org/abs/1606.05897>
- Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. 2016b. *Image Style Transfer Using Convolutional Neural Networks*. In *Proc. CVPR*. IEEE Computer Society, Los Alamitos, 2414–2423. doi: 10.1109/CVPR.2016.265
- Leon A. Gatys, Alexander S. Ecker, Matthias Bethge, Aaron Hertzmann, and Eli Shechtman. 2016c. *Controlling Perceptual Factors in Neural Style Transfer*. arXiv.org report 1611.07865. <https://arxiv.org/abs/1611.07865>
- L. A. Gatys, A. S. Ecker, M. Bethge, A. Hertzmann, and E. Shechtman. 2017. *Controlling Perceptual Factors in Neural Style Transfer*. In *Proc. CVPR*. IEEE Computer Society, Los Alamitos. To appear.
- Agrim Gupta, Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2017. *Characterizing and Improving Stability in Neural Style Transfer*. arXiv.org report 1705.02092. <https://arxiv.org/abs/1705.02092>
- Xun Huang and Serge Belongie. 2017. *Arbitrary Style Transfer in Real-time with Adaptive Instance Normalization*. arXiv.org report 1703.06868. <https://arxiv.org/abs/1703.06868>
- Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. 2016. *Let There Be Color! Joint End-to-end Learning of Global and Local Image Priors for Automatic Image Colorization with Simultaneous Classification*. *ACM Transactions on Graphics* 35, 4, Article 110 (July 2016), 11 pages. doi: 10.1145/2897824.2925974
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016a. *Perceptual Losses for Real-Time Style Transfer and Super-Resolution*. In *Proc. ECCV*. Springer International, Cham, Switzerland, 694–711. doi: 10.1007/978-3-319-46475-6_43
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016b. *Perceptual Losses for Real-Time Style Transfer and Super-Resolution*. arXiv.org report 1603.08155. <https://arxiv.org/abs/1603.08155>
- Chuan Li and Michael Wand. 2016. *Combining Markov Random Fields and Convolutional Neural Networks for Image Synthesis*. In *Proc. CVPR*. IEEE Computer Society, Los Alamitos, 2479–2486. doi: 10.1109/CVPR.2016.272
- X. Liu, Paul L. Rosin, Yu-Kun Lai, and M.-M. Cheng. 2017. *Depth-Aware Neural Style Transfer*. In *Proc. NPAR*. ACM, New York. To appear (this proceedings).
- Eric Risser, Pierre Wilmot, and Connelly Barnes. 2017. *Stable and Controllable Neural Texture Synthesis and Style Transfer Using Histogram Losses*. arXiv.org report 1701.08893. <https://arxiv.org/abs/1607.08022>

- Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. 2016. *Artistic Style Transfer for Videos*. In *Proc. GCPR*. Springer International, Cham, Switzerland, 26–36. doi: 10.1007/978-3-319-45886-1_3
- Ahmed Selim, Mohamed Elgharib, and Linda Doyle. 2016. *Painting Style Transfer for Head Portraits Using Convolutional Neural Networks*. *ACM Transactions on Graphics* 35, 4, Article 129 (July 2016), 18 pages. doi: 10.1145/2897824.2925968
- Yaniv Taigman, Adam Polyak, and Lior Wolf. 2016. *Unsupervised Cross-Domain Image Generation*. arXiv.org report 1611.02200. <https://arxiv.org/abs/1611.02200>
- Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S. Lempitsky. 2016a. *Texture Networks: Feed-forward Synthesis of Textures and Stylized Images*. In *Proc. International Conference on Machine Learning*. JMLR.org, New York, 1349–1357. <http://jmlr.org/proceedings/papers/v48/ulyanov16.html>
- Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. 2017a. *Improved Texture Networks: Maximizing Quality and Diversity in Feed-forward Stylization and Texture Synthesis*. In *Proc. CVPR*. IEEE Computer Society, Los Alamitos. To appear.
- Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. 2017b. *Improved Texture Networks: Maximizing Quality and Diversity in Feed-forward Stylization and Texture Synthesis*. arXiv.org report 1701.02096. <https://arxiv.org/abs/1701.02096>
- Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. 2016b. *Instance Normalization: The Missing Ingredient for Fast Stylization*. arXiv.org report 1607.08022. <https://arxiv.org/abs/1607.08022>
- Hang Zhang and Kristin Dana. 2017. *Multi-style Generative Network for Real-time Transfer*. arXiv.org report 1703.06953. <https://arxiv.org/abs/1703.06953>

Others

- Adrien Bousseau, Matt Kaplan, Joëlle Thollot, and François X. Sillion. 2006. *Interactive Watercolor Rendering with Temporal Coherence and Abstraction*. In *Proc. NPAR*. ACM, New York, 141–149. doi: 10.1145/1124728.1124751
- Forrester Cole, Aleksey Golovinskiy, Alex Limpaecher, Heather Stoddart Barros, Adam Finkelstein, Thomas Funkhouser, and Szymon Rusinkiewicz. 2008. *Where Do People Draw Lines?* *ACM Transactions on Graphics* 27, 3, Article 88 (Aug. 2008), 11 pages. doi: 10.1145/1360612.1360687
- Christopher DeCoro, Forrester Cole, Adam Finkelstein, and Szymon Rusinkiewicz. 2007. *Stylized Shadows*. In *Proc. NPAR*. ACM, New York, 77–83. doi: 10.1145/1274871.1274884
- Maarten H. Everts, Henk Bekker, Jos B. T. M. Roerdink, and Tobias Isenberg. 2015. *Interactive Illustrative Line Styles and Line Style Transfer Functions for Flow Visualization*. arXiv.org report 1503.05787. <https://www.arxiv.org/abs/1503.05787>
- Moritz Gerl and Tobias Isenberg. 2013. *Interactive Example-based Hatching*. *Computers & Graphics* 37, 1–2 (Feb.–April 2013), 65–80. doi: 10.1016/j.cag.2012.11.003
- Aaron Hertzmann. 1998. *Painterly Rendering with Curved Brush Strokes of Multiple Sizes*. In *Proc. SIGGRAPH*. ACM, New York, 453–460. doi: 10.1145/280814.280951
- Aaron Hertzmann. 2002. *Fast Paint Texture*. In *Proc. NPAR*. ACM, New York, 91–96. doi: 10.1145/508530.508546
- Tobias Isenberg, Petra Neumann, Sheelagh Carpendale, Mario Costa Sousa, and Joaquim A. Jorge. 2006. *Non-Photorealistic Rendering in Context: An Observational Study*. In *Proc. NPAR*. ACM, New York, 115–126. doi: 10.1145/1124728.1124747
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. *ImageNet Classification with Deep Convolutional Neural Networks*. In *Advances in Neural Information Processing Systems (NIPS)*. Curran Associates, Inc., Red Hook, NY, USA, 1097–1105. <https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks>
- Hyunjun Lee, Sungtae Kwon, and Seungyong Lee. 2006. *Real-Time Pencil Rendering*. In *Proc. NPAR*. ACM, New York, 37–45. doi: 10.1145/1124728.1124735
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. *Microsoft COCO: Common Objects in Context*. In *Proc. ECCV*. Springer International, Cham, Switzerland, 740–755. doi: 10.1007/978-3-319-10602-1_48
- Lukasz Mackiewicz and Francho Melendez. 2016. *Loving Vincent: Guiding Painters Through 64,000 Frames*. In *SIGGRAPH Talks*. ACM, New York, Article 6, 2 pages. doi: 10.1145/2897839.2927394
- Domingo Martín, Vicente del Sol, Celia Romo, and Tobias Isenberg. 2015. *Drawing Characteristics for Reproducing Traditional Hand-Made Stippling*. In *Proc. NPAR*. Eurographics Association, Goslar, Germany, 103–115. doi: 10.2312/exp.20151183
- Parag K. Mital, Mick Grierson, and Tim J. Smith. 2013. *Corpus-based Visual Synthesis: An Approach for Artistic Stylization*. In *Proc. SAP*. ACM, New York, 51–58. doi: 10.1145/2492494.2492505
- Santiago E. Montesdeoca, Hock-Soon Seah, and Hans-Martin Rall. 2016. *Art-directed Watercolor Rendered Animation*. In *Proc. NPAR*. Eurographics Association, Goslar, Germany, 51–58. doi: 10.2312/exp.20161063
- Santiago E. Montesdeoca, Hock-Soon Seah, Hans-Martin Rall, and Davide Benvenuti. 2017. *Art-directed Watercolor Stylization of 3D Animations in Real-time*. *Computers & Graphics* (2017). doi: 10.1016/j.cag.2017.03.002 To appear.
- Emil Praun, Hughes Hoppe, Matthew Webb, and Adam Finkelstein. 2001. *Real-Time Hatching*. In *Proc. SIGGRAPH*. ACM, New York, 581–586. doi: 10.1145/383259.383328

- Michael P. Salisbury, Sean E. Anderson, Ronen Barzel, and David H. Salesin. 1994. Interactive Pen-and-Ink Illustration. In *Proc. SIGGRAPH*. ACM, New York, 101–108. doi: 10.1145/192161.192185
- Amir Semmo, Tobias Dürschmid, Matthias Trapp, Mandy Klingbeil, Jürgen Döllner, and Sebastian Pasewaldt. 2016a. Interactive Image Filtering with Multiple Levels-of-control on Mobile Devices. In *Proc. MGIA*. ACM, New York, Article 2, 8 pages. doi: 10.1145/2999508.2999521
- Amir Semmo, Daniel Limberger, Jan Eric Kyprianidis, and Jürgen Döllner. 2016b. Image Stylization by Interactive Oil Paint Filtering. *Computers & Graphics* 55 (April 2016), 157–171. doi: 10.1016/j.cag.2015.12.001
- Karen Simonyan and Andrew Zisserman. 2015. *Very Deep Convolutional Networks for Large-Scale Image Recognition*. arXiv.org report 1409.1556. <https://arxiv.org/abs/1409.1556>
- Peter-Pike Sloan, William Martin, Amy Gooch, and Bruce Gooch. 2001. The Lit Sphere: A Model for Capturing NPR Shading from Art. In *Proc. Graphics Interface*. Morgan Kaufmann, San Francisco, 143–150. doi: 10.20380/GI2001.17
- Minjung Son, Yunjin Lee, Henry Kang, and Seungyong Lee. 2011. Structure Grid for Directional Stippling. *Graphical Models* 73, 3 (May 2011), 74–87. doi: 10.1016/j.gmod.2010.12.001
- Bin Wang, Wenping Wang, Huaiping Yang, and Jianguang Sun. 2004. Efficient Example-Based Painting and Synthesis of 2D Directional Texture. *IEEE Transactions on Visualization and Computer Graphics* 10, 3 (May/June 2004), 266–277. doi: 10.1109/TVCG.2004.1272726
- Miaoyi Wang, Bin Wang, Yun Fei, Kanglai Qian, Wenping Wang, Jiating Chen, and Jun-Hai Yong. 2014. Towards Photo Watercolorization with Artistic Verisimilitude. *IEEE Transactions on Visualization and Computer Graphics* 20, 10 (Feb. 2014), 1451–1460. doi: 10.1109/TVCG.2014.2303984
- Holger Winnemöller, Jan Eric Kyprianidis, and Sven Olsen. 2012. XDoG: An eXtended Difference-of-Gaussians Compendium including Advanced Image Stylization. *Computers & Graphics* 36, 6 (Oct. 2012), 740–753. doi: 10.1016/j.cag.2012.03.004
- Holger Winnemöller, Sven C. Olsen, and Bruce Gooch. 2006. Real-Time Video Abstraction. *ACM Transactions on Graphics* 25, 3 (July 2006), 1221–1226. doi: 10.1145/1141911.1142018
- Chung-Ren Yan, Ming-Te Chi, Tong-Yee Lee, and Wen-Chieh Lin. 2008. Stylized Rendering Using Samples of a Painted Image. *IEEE Transactions on Visualization and Computer Graphics* 14, 2 (March 2008), 468–480. doi: 10.1109/TVCG.2007.70440
- Kun Zeng, Mingtian Zhao, Caiming Xiong, and Song-Chun Zhu. 2009. From Image Parsing to Painterly Rendering. *ACM Transactions on Graphics* 29, 1, Article 2 (Dec. 2009), 11 pages. doi: 10.1145/1640443.1640445