

Non-photorealistic Rendering of Images: from Image Processing to Machine Learning

Dr. Yu-Kun Lai, Cardiff University GAMES webinar

4/7/2019

Non-Photorealistic Rendering



A well researched topic that mimics artistic styles using computer algorithms





Painting

Pen & Ink

Technical illustrations



Scientific visualisation





Our Work: Traditional Methods

CARDIFF UNIVERSITY PRIFYSGOL CAERDYD

Dedicated algorithms are developed to create specific styles.



Rosin & Lai, Graphical Models, 2013.



Lai & Rosin, IEEE Trans. Image Processing, 2014.



Rosin & Lai, Computational Aesthetics, 2013.



Rosin & Lai, Computational Aesthetics, 2015.

Artistic Minimal Rendering

CARDIFF UNIVERSITY PRIFYSGOL CAERDY

Key observations:

- Using a small number of tones
- Simple primitives: lines and tonal blocks



Portrait by Andy Warhol

Input

Single-scale

Multi-scale

Artistical minimal rendering with lines and blocks, Graphical Models, 2013.

Artistic Minimal Rendering: Pipeline





Artistic Minimal Rendering: 3-tone lines





Input

Lines only (2-tone)

3-tone

posterised into 3 levels

Artistic Minimal Rendering: Tonal Blocks



(f)

(e)



(g)



Artistic Minimal Rendering: Tone Overlay



Rules for determining the output tone (B = black, G = gray, W = white) to be rendered according to the tones of the lines and blocks (if present).

Lines	В	×	×	×	×	×	×	×	×	\checkmark	\checkmark	\checkmark	\checkmark		\checkmark	\checkmark	\checkmark
	W	×	×	×	×	\checkmark	\checkmark	\checkmark	\checkmark	×	×	×	×	\checkmark	\checkmark	\checkmark	\checkmark
Block	В	×	×	\checkmark	\checkmark	×	×	\checkmark	\checkmark	×	×	\checkmark	\checkmark	×	×	\checkmark	\checkmark
	W	×	\checkmark	×	\checkmark	×	\checkmark	×	\checkmark	×	\checkmark	×	\checkmark	×	\checkmark	×	\checkmark
Output		G	W	В	G	W	G	W	W	В	В	G	В	G	W	В	G





Results and Extensions





Results and Extensions







GAN-based Image Synthesis





CartoonGAN



We address cartoon stylisation of images

- Using unpaired content and style image sets
- Training data easy to obtain
 - Content: normal photos from Flickr
 - Style: key frames from cartoon films



(a) Original scene

(b) Our result

CartoonGAN: Generative Adversarial Networks for Photo Cartoonization. CVPR, 2018.



Network architecture

- Generator: similar to [Johnson et al. 2016]
- Discriminator: differentiate real and synthesised images





Content loss:

$$\mathcal{L}_{con}(G, D) = \mathbb{E}_{p_i \sim S_{data}(p)}[||VGG_l(G(p_i)) - VGG_l(p_i)||_1]$$

Use L1 instead of L2 to cope with local large differences (recover flat shading)





Adversarial loss $\mathcal{L}_{adv}(G, D) = \mathbb{E}_{c_i \sim S_{data}(c)}[\log D(c_i)] + \mathbb{E}_{e_j \sim S_{data}(e)}[\log(1 - D(e_j))]$

$$+ \mathbb{E}_{p_k \sim S_{data}(p)} [\log(1 - D(G(p_k)))].$$

- Edges are often lost since they only cover a small number of pixels
- We add an edge promoting term in adversarial loss by penalising cartoon images with edges smoothed



(a) A cartoon image c_i



(b) The edge-smoothed version e_i



Initialisation

- Traditional random initialisation does not give good results
- To avoid the GAN model stuck at poor local minima, we start generator learning that only aims to reconstruct the content of the input images.



(a) Original photo



(b) Image after initialization



CartoonGAN: Results





(a) input photo

(b) Shinkai style



CartoonGAN: Results





APDrawingGAN



Portrait drawings: a longstanding and distinct art form, which typically use a sparse set of continuous graphical elements (e.g., lines) to capture the distinctive appearance of a person



APDrawingGAN: Generating Artistic Portrait Drawings from Face Photos with Hierarchical GANs. CVPR, 2019 (oral).

APDrawingGAN: Challenges



Artistic portrait drawings (APDrawings) are substantially different from painting styles studied in previous work:

Highly abstract, sparse but continuous graphical elements

Stronger semantic constraints for portrait

Different rendering for different facial parts

Elements not located precisely by artists

Conceptual lines not directly related to low level features







APDrawingGAN: Challenges



(b) Our method

in the training set

(c) Existing methods



APDrawingGAN Highlights:



Novel distance transform (DT) loss

A new APDrawing Dataset

APDrawingGAN: Hierarchical Structure





each of which includes a global network and six local networks.

APDrawingGAN: Hierarchical Generator G





APDrawingGAN: Hierarchical Discriminator D





APDrawingGAN: Loss Function



There are four terms in the loss function:



 $L(G,D) = L_{adv}(G,D) + \lambda_1 L_{\mathcal{L}1}(G,D) + \lambda_2 L_{DT}(G,D) + \lambda_3 L_{local}(G,D)$



L1 loss:

Penalize minor misalignments. Treat small misalignments and big misalignments as the same... DT loss: Based on distance. Tolerate small misalignments Penalize big misalignments

APDrawingGAN: Dataset





We build an artistic portrait drawing dataset containing 140 pairs of high-resolution portrait photos and corresponding professional artistic drawings.

APDrawingGAN: Pre-training





6655 pairs of face photos and NPR results for pre-training.

























APDrawingGAN: Ablation study





APDrawingGAN: Comparison with Gatys, CycleGAN, Pix2Pix



Input face photo

Ground truth

Gatys

CycleGAN

Pix2Pix

APDrawingGAN

APDrawingGAN: Comparison with CNNMRF, Deep Image Analogy and Headshot Portrait







APDrawingGAN: User study





Conclusions



- We show how traditional image processing techniques and machine learning can be used for non-photorealistic rendering of images
- Machine learning is more flexible to cope with different styles
- We demonstrate that they can be combined to perform effective non-photorealistic rendering
- Many challenges remain
 - Quality of results
 - Robustness
 - Evaluation

Conclusions



Quality evaluation

- Subjective, typically using a few examples to demonstrate the method works
- One approach is to create benchmark datasets



Mould & Rosin, An image stylization benchmark, Expressive 2016

NPRportrait Benchmark

Image and ancillary data from the NPRportrait Benchmark described in the paper below can be downloaded.

PL. Rosin, D. Mould,, I. Berger, J. Collomosse, Y.K. Lai, C. Li, H. Li, A. Shamir, M. Wand, T. Wang, H. Winnemöller, "Benchmarking Non-Photorealistic Rendering of Portraits"

Level 1

Thumbnails for the level 1 images from the above paper are shown below, and the (full resolution) set is zipped up for downloading here.



Level 2

Thumbnails for the level 2 images from the above paper are shown below, and the (full resolution) set is zipped up for downloading here.



Rosin et al. Benchmarking Non-Photorealistic Rendering of Portraits, Expressive, 2017



Thank you!