Computational Visual Media
DOI 10.1007/s41095-xxx-xxxx-x Vol. x, No. x, month year, xx–xx

Article Type (Research/Review)

NPRportrait 1.0: A three-level benchmark for non-photorealistic rendering of portraits [Electronic supplementary material]

Paul L. Rosin[1]✉, Yu-Kun Lai[1], David Mould[2], Ran Yi[3], Itamar Berger[4], Lars Doyle[2], Seungyong Lee[5], Chuan Li[6], Yong-Jin Liu[3], Amir Semmo[7], Ariel Shamir[4], Minjung Son[8], and Holger Winnemöller[9]

1 Description of the Evaluated NPR Algorithms

Li and Wand's method [8] treats styles as textures, and forces the synthesised image and the reference style image to have the same Markovian texture statistics. Non-parametric sampling is first used to capture patches from the style image; patch matching and blending are then used to transfer the style to the synthesised image. For portrait stylisation, they include an additional content constraint that minimises the L_2 distance between the CNN encoding of the portrait photo and the synthesised image.

Berger *et al.* [1] mimic the style of specific artists' line-drawings in a data-driven manner. Sample drawings of artists are collected and their statistics are analysed. Then, given a new portrait photograph and an artist style, the algorithm first creates a contour image by using a variant of the XDoG method [18]. Using the detected facial features, the face geometry is modified to follow the specific artist's geometric style. Lastly, the face contours are drawn using strokes from the artist's stroke database following the artist's drawing statistics.

Yi *et al.* [19] proposed APDrawingGAN, a hierarchical system of generative adversarial networks (GANs) that transforms face photographs into high-quality artistic portrait drawings. Since artists usually use different drawing styles for different facial regions, this hierarchical GAN model combines a global network (for fusing local parts) and six local networks (for individual facial regions). Finally, to train this model, a novel line-promoting distance transform loss was proposed to capture the fact that an artist's drawing is usually not perfectly aligned with image features.

Rosin and Lai's algorithm [10] first stylises the image with abstracted regions of flat colours plus black and white lines [7], then fits a partial face model to the input image and attempts to detect the skin region. Shading and line rendering is stylised in the skin region, and in addition, the face model helps inform portrait-specific enhancements: reducing line clutter; improving eye detail; colouring the lips and teeth; and inserting synthesised highlights. It is straightforward to modify this pipeline to render, in place of this "puppet" style, a more abstracted version, inspired by the artist Julian Opie.

Note that some subsequent modifications have been made to the published description of the algorithm. The first change was introduced, and described in [13], to better cope with different ethnicities. The portraits are first classified as containing dark or light skin, which then selected a slight variation in the pipeline. In the original version of the algorithm [10] the three classes of pixels in HSI space have their intensity values quantised to {0,200,255}. However, subsequently it was found that dark skin is better rendered if the mean intensity value for each class was used instead. Previously [13] the determination of dark versus light skin was done by testing whether the mean intensity of a central region in the face lies below the mid intensity range value (128). This has been replaced in this paper by applying a threshold of 5 to the individual typology angle (ITA) [17] calculated from the facial region. Another change is that orientation of the cylinder used to provide the shading effect for the puppet style is estimated more accurately; see [12].

Winnemöller *et al.*'s XDoG filter [18] can be conceptualised as the weighted sum of a blurred

source image and a scaled difference-of-Gaussians (DoG) response of the same image, effectively applying unsharp masking to the DoG response. Combined with subsequent soft thresholding, this computationally simple filter allows a wide range of stylistic and artistic effects, including cartoon shading, black-and-white thresholding, and charcoal shading. If required, local modification of filter parameters, according to facial features, would be trivial to implement.

Rosin and Lai [12] create an engraving style rendering of an image using a dither matrix, which is a spatially-varying threshold. The dither matrix has been designed so that it generates a pattern of black and white lines forming cross hatching. The method is enhanced by using a simple cylindrical model of the face to warp the dither matrix so that the lines curve around the face, providing a pseudo-3D effect.

Son *et al.* [16] proposes a novel method for hedcut, where the varying sizes of dots and hatching lines are regularly spaced along the local feature orientations. A smooth grid curved along the feature vector field, named the structure grid, is synthesized to contain the tangential and normal distances to the nearest grid intersection at each pixel. Given a structure grid, the appropriate positions and attributes of primitives are determined via rapid pixel-based primitive rendering. The method works well for human faces even though it is not specially designed for portraits.

Semmo *et al.*'s [15] oil paint filter is based on non-linear image smoothing to obtain painterly looks with a soft color blending. The method uses Gaussian-based filter kernels that are aligned to the main feature contours of an image for structure-adaptive filtering. By using the construct of the smoothed structure tensor and principles of line integral convolution to synthesize paint textures in real-time, the filter responses are locally controllable. In particular, the level of abstraction can be easily adjusted by interactive painting or could be based on facial feature masks.

Doyle *et al.*'s [3] pebble mosaic stylisation process begins with a superpixel segmentation of the image, guided by an orientation field derived from the structure tensor. Each superpixel is converted into a pebble by first smoothing the exterior boundary and then computing a height field for the tile interior, determined by harmonic interpolation between the tile boundary and an interior contour placed at a set height. The resulting 3D geometry can be conventionally rendered and textured, using a tile color that is the average color of the pixels within the image segment.

Rosin and Lai [11] use a filter based approach to generate a watercolour stylisation. In order to achieve the multiple characteristics of watercolour – namely brightening, abstraction, edge darkening, wobbling, granulation, glazing, pigment and paper variations – they employ various steps such as smoothing, morphological opening and closing, contrast-limited local histogram equalisation, edge detection, overlay blend, local geometric distortion, superpixel segmentation, and level of detail masks controlled by face detection and saliency masks.

2 Full Results of 11 NPR Algorithms

Figures 1–11 show the results of applying the 11 NPR algorithms to the full *NPRportrait1.0* benchmark dataset, while in figures 12–22 the results are shown ranked as described in Experiment 2.

3 Examples of some Additional Stylisations

Figure 23 shows some further examples, covering a wide range of stylisations, on two images from level 1 of the *NPRportrait1.0* benchmark. All the algorithms are general, i.e. not portrait-specific.

4 Photo credits

Level 1:

1. Nuclear Regulatory Commission (Wikimedia) Creative Commons Attribution 2.0 Generic license;
2. Rod Waddington (Flickr) Attribution-ShareAlike 2.0 Generic (CC BY-SA 2.0);
3. Paul Rosin (no copyright);
4. Paul Rosin (no copyright);
5. United States Air Force (Wikimedia) public domain in the United States;
6. Kaydiddy35 (Wikimedia) Creative Commons Attribution-Share Alike 3.0 Unported license;
7. Charles Mims (Flickr) Attribution 2.0 Generic (CC BY 2.0);
8. Moid Rasheedi (Wikimedia) Creative Commons CC0 1.0 Universal Public Domain Dedication;
9. Rahman Uddin (Wikimedia) Creative Commons Attribution-Share Alike 4.0 International licence;
10. Paul attick (Wikimedia) Creative Commons Attribution-Share Alike 4.0 International licence;
11. Department of Foreign Affairs and Trade (Wikimedia) Creative Commons Attribution 3.0 Australia;
12. UC Davis College of Engineering (Flickr) Attribution 2.0 Generic (CC BY 2.0);

13. icannphotos (Flickr) Attribution-ShareAlike 2.0 Generic licence;
14. Paul Rosin (no copyright);
15. MARRON1119 (Wikimedia) Creative Commons Attribution-Share Alike 4.0 International licence;
16. Choi Kwang-mo (Wikimedia) Creative Commons Attribution-Share Alike 4.0 International license;
17. Fred Bchx (Flickr) Attribution 2.0 Generic;
18. Iolo Penri (Wikimedia) Creative Commons Attribution-Share Alike 4.0 International license;
19. Zboralski (Wikimedia) Creative Commons Attribution-Share Alike 3.0 Unported license;
20. Julia Miele (Flickr) Attribution 2.0 Generic (CC BY 2.0).

Level 2:

1. Sgt. Matthew Callahan (Wikimedia) public domain;
2. BBC World Service (Flickr) Attribution-NonCommercial 2.0 Generic (CC BY-NC 2.0) but use granted by author for NPRportrait;
3. Rod Waddington (Flickr) Attribution-ShareAlike 2.0 Generic (CC BY-SA 2.0);
4. Adam McGuffie (Flickr) Attribution 2.0 Generic;
5. Dominick D (Wikimedia) Creative Commons Attribution-Share Alike 2.0 Generic license;
6. www.j-pics.info (Flickr) Attribution-NonCommercial-ShareAlike 2.0 Generic (CC BY-NC-SA 2.0) but use granted by author for NPRportrait;
7. certified_su (Flickr) Attribution 2.0 Generic;
8. Nando.uy (Flickr) Attribution-NonCommercial-ShareAlike 2.0 Generic (CC BY-NC-SA 2.0) but use granted by author for NPRportrait;
9. Christopher Blizzard (Flickr) Attribution-ShareAlike 2.0 Generic (CC BY-SA 2.0);
10. Andrew Braithwaite (Flickr) Attribution 2.0 Generic (CC BY 2.0);
11. y5RZouZwNsH6MI (Flickr) Attribution-ShareAlike 2.0 Generic (CC BY-SA 2.0);
12. shankar s. (Flickr) Attribution 2.0 Generic (CC BY 2.0);
13. Noah Buscher (Unsplash);
14. Sparky (Flickr) Attribution-NonCommercial 2.0 Generic (CC BY-NC 2.0) but use granted by author for NPRportrait;
15. unknown (Pexels);
16. Martin Sharman (Flickr) Attribution-NonCommercial-ShareAlike 2.0 Generic (CC BY-NC-SA 2.0) but use granted by author for NPRportrait;
17. Wilfredor (Wikimedia) Creative Commons CC0 1.0 Universal Public Domain Dedication;
18. Einar Aslaksen (Wikimedia) Creative Commons Attribution 2.0 Generic license;
19. Hamish Irvine (Flickr) Attribution-NonCommercial 2.0 Generic (CC BY-NC 2.0) but use granted by author for NPRportrait;
20. Greg Peverill-Conti (Flickr) Attribution-NonCommercial-ShareAlike 2.0 Generic (CC BY-NC-SA 2.0) but use granted by author for NPRportrait.

Level 3:

1. Johanna Buguet (Unsplash);
2. Commander, U.S. Naval Forces Europe-Africa/U.S. 6th Fleet (Flickr) public domain;
3. Felipe Sagn (Unsplash);
4. Nathan Dumlao (Unsplash);
5. Old Youth (Unsplash);
6. Olesya Yemets (Unsplash);
7. Azamat Zhanisov (Unsplash);
8. Claudia (Unsplash);
9. Lachlan Hardy (Flickr) Attribution 2.0 Generic (CC BY 2.0);
10. Calvin Lupiya (Unsplash);
11. Artyom Kim (Unsplash);
12. Jordan Bauer (Unsplash);
13. Alex Iby (Unsplash);
14. Andrew Robinson (Unsplash);
15. David in Lisburn (Flickr) Attribution 2.0 Generic (CC BY 2.0);
16. Gabriel Silvério (Unsplash);
17. Lachlan Hardy (Flickr) Attribution 2.0 Generic (CC BY 2.0);
18. Michael Vadon (Flickr) Attribution-ShareAlike 2.0 Generic (CC BY-SA 2.0);
19. Orin Zebest (Flickr) Attribution 2.0 Generic (CC BY 2.0) ;
20. irumge (Flickr) Attribution 2.0 Generic (CC BY 2.0).

Note: The Unsplash license allows images to be used for free for commercial and non-commercial purposes, and no permission is needed. Unsplash images cannot be sold without significant modification.

1 School of Computer Science and Informatics, Cardiff University, Cardiff, UK. E-mail: PaulRosin@cs.cf.ac.uk, Yukun.Lai@cs.cardiff.ac.uk.

2 School of Computer Science, Carleton University, Canada. E-mail: mould@scs.carleton.ca, larsdoyle@cmail.carleton.ca.

3 Department of Computer Science and Technology, Tsinghua University, Beijing, China. E-mail: yr16@mails.tsinghua.edu.cn,liuyongjin@tsinghua.edu.cn.

4 The Interdisciplinary Center Herzliya, Israel. E-mail: berger.itamar@gmail.com, arik@idc.ac.il.

5 Department of Computer Science and Engineering, Pohang University of Science and Technology, South Korea. E-mail: leesy@postech.ac.kr.

6 Lambda Labs, Inc., USA. E-mail: c@lambdal.com.

7 Hasso Plattner Institute, University of Potsdam, Germany. E-mail: Amir.Semmo@hpi.de.

8 Multimedia Processing Laboratory, Samsung Advanced Institute of Technology, South Korea. E-mail: minjungs.son@samsung.com.

9 Adobe Systems, Inc. USA. E-mail: hwinnemo@adobe.com.

Manuscript received: 2014-12-31; accepted: 2015-01-30.

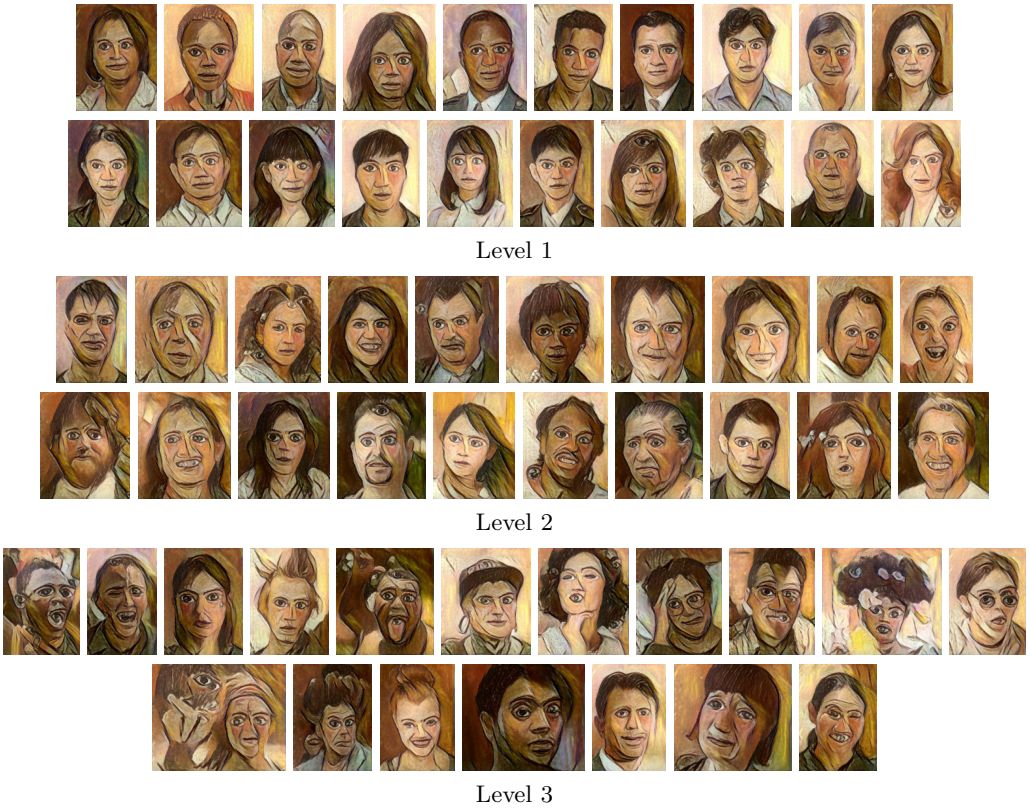

Fig. 1 *NPRportrait1.0* benchmark stylised using neural style transfer: Li and Wand [8]

References

[1] I. Berger, A. Shamir, M. Mahler, E. Carter, and J. Hodgins. Style and abstraction in portrait sketching. *ACM Transactions on Graphics (TOG)*, 32(4):55, 2013.

[2] C.-C. Chiu, Y.-H. Lo, R.-R. Lee, and H.-K. Chu. Tone- and feature-aware circular scribble art. In *Computer Graphics Forum*, volume 34, pages 225–234, 2015.

[3] L. Doyle, F. Anderson, E. Choy, and D. Mould. Automated pebble mosaic stylization of images. *Computational Visual Media*, 5(1):33–44, 2019.

[4] N. Faraj, G.-S. Xia, J. Delon, and Y. Gousseau. A generic framework for the structured abstraction of images. In *Proceedings of the International Symposium on Non-Photorealistic Animation and Rendering*, pages 1–11, 2017.

[5] A. S. Javid, L. Doyle, and D. Mould. Irregular pebble mosaics with sub-pebble detail. In C. S. Kaplan, A. Forbes, and S. DiVerdi, editors, *ACM/EG Expressive Symposium*, 2019.

[6] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711, 2016.

[7] Y.-K. Lai and P. L. Rosin. Efficient circular thresholding. *IEEE Trans. Image Processing*, 23(3):992–1001, 2014.

[8] C. Li and M. Wand. Combining Markov Random Fields and convolutional neural networks for image synthesis. In *Proc. Computer Vision and Pattern Recognition*, pages 2479–2486, 2016.

[9] P. L. Rosin and Y.-K. Lai. Artistic minimal rendering with lines and blocks. *Graphical Models*, 75(4):208–229, 2013.

[10] P. L. Rosin and Y.-K. Lai. Non-photorealistic rendering of portraits. In *Proceedings of the Workshop on Computational Aesthetics*, pages 159–170. Eurographics Association, 2015.

[11] P. L. Rosin and Y.-K. Lai. Watercolour rendering of portraits. In *Pacific-Rim Symposium on Image and Video Technology*, pages 268–282, 2017.

[12] P. L. Rosin and Y.-K. Lai. Image-based portrait engraving. *arXiv preprint arXiv:2008.05336*, 2020.

[13] P. L. Rosin, D. Mould, I. Berger, J. P. Collomosse, Y. Lai, C. Li, H. Li, A. Shamir, M. Wand, T. Wang, and H. Winnemöller. Benchmarking non-photorealistic rendering of portraits. In *Proceedings of the International Symposium on Non-Photorealistic Animation and Rendering*, pages 11:1–11:12, 2017.

[14] A. Secord. Weighted voronoi stippling. In *Proceedings of the International Symposium on Non-Photorealistic Animation and Rendering*, pages 37–43, 2002.

[15] A. Semmo, D. Limberger, J. E. Kyprianidis, and J. Döllner. Image stylization by interactive oil paint

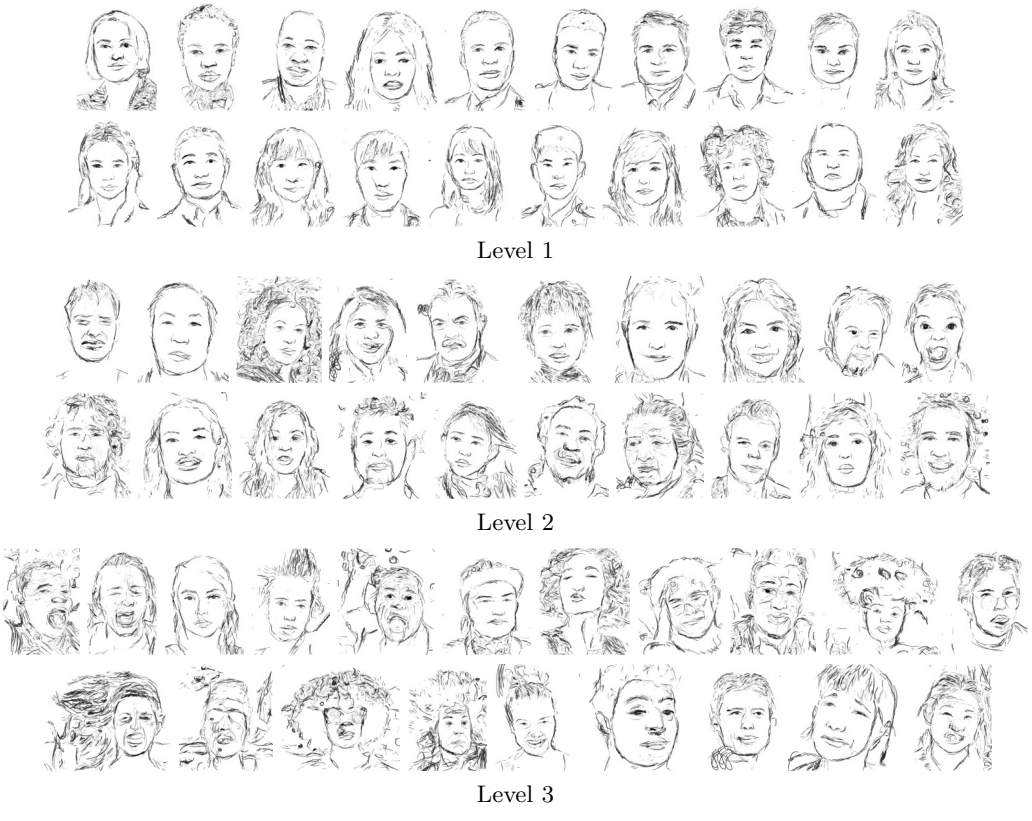

Fig. 2 *NPRportrait1.0* benchmark stylised by the artistic sketch method: Berger *et al.* [1]

filtering. *Computers & Graphics*, 55:157–171, 2016.

[16] M. Son, Y. Lee, H. Kang, and S. Lee. Structure grid for directional stippling. *Graphical Models*, 73(3):74–87, 2011.

[17] M. Wilkes, C. Y. Wright, J. L. du Plessis, and A. Reeder. Fitzpatrick skin type, individual typology angle, and melanin index in an African population: steps toward universally applicable skin photosensitivity assessments. *JAMA Dermatology*, 151(8):902–903, 2015.

[18] H. Winnemöller, J. E. Kyprianidis, and S. C. Olsen. XDoG: An eXtended difference-of-gaussians compendium including advanced image stylization. *Computers & Graphics*, 36(6):740–753, 2012.

[19] R. Yi, Y.-J. Liu, Y.-K. Lai, and P. L. Rosin. APDrawingGAN: Generating artistic portrait drawings from face photos with hierarchical GANs. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10743–10752, 2019.

Fig. 3 *NPRportrait1.0* benchmark stylised by APDrawingGAN: Yi *et al.* [19]

Fig. 4 *NPRportrait1.0* benchmark stylised as puppets: Rosin and Lai [10]

Level 1

Level 2

Level 3

Fig. 5 *NPRportrait1.0* benchmark stylised by XDoG: Winnemöller *et al.* [18]

Level 1

Level 2

Level 3

Fig. 6 *NPRportrait1.0* benchmark stylised as engravings: Rosin and Lai [12]

Fig. 7 *NPRportrait1.0* benchmark stylised as hedcuts: Son *et al.* [16]

Fig. 8 *NPRportrait1.0* benchmark stylised as oil paintings: Semmo *et al.* [15]

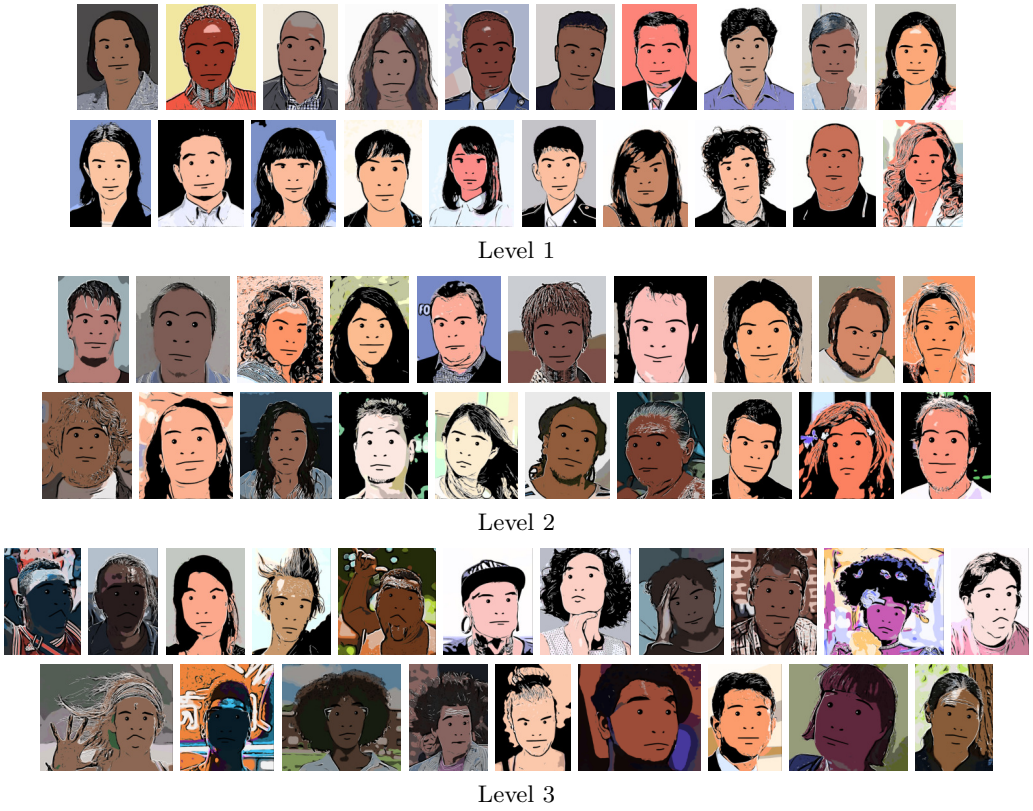

Fig. 9 *NPRportrait1.0* benchmark stylised in the Julian Opie style: Rosin and Lai [10]

Fig. 10 *NPRportrait1.0* benchmark stylised as pebble mosaics: Doyle *et al.* [3]

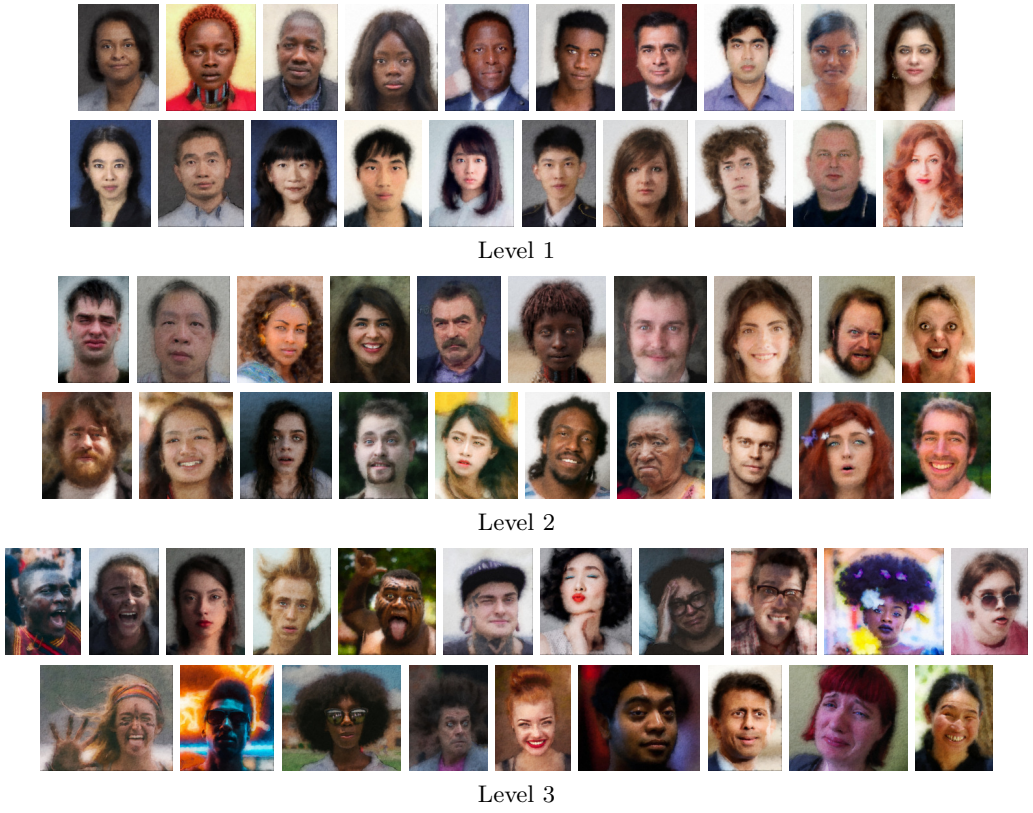

Fig. 11 *NPRportrait1.0* benchmark stylised as watercolours: Rosin and Lai [11]

Fig. 12 Images from *NPRportrait1.0* benchmark stylised using neural style transfer: Chuan and Wand [8] and ranked according to Experiment2.

Fig. 13 Images from *NPRportrait1.0* benchmark stylised by the artistic sketch method: Berger *et al.* [1] and ranked according to Experiment2.

Fig. 14 Images from *NPRportrait1.0* benchmark stylised by APDrawingGAN: Yi *et al.* [19] and ranked according to Experiment2.

Fig. 15 Images from *NPRportrait1.0* benchmark stylised as puppets: Rosin and Lai [10] and ranked according to Experiment2.

Fig. 16 Images from *NPRportrait1.0* benchmark stylised by XDoG: Winnemöller *et al.* [18] and ranked according to Experiment2.

Fig. 17 Images from *NPRportrait1.0* benchmark stylised as engravings: Rosin and Lai [12] and ranked according to Experiment2.

Fig. 18 Images from *NPRportrait1.0* benchmark stylised as hedcuts: Son *et al.* [16] and ranked according to Experiment2.

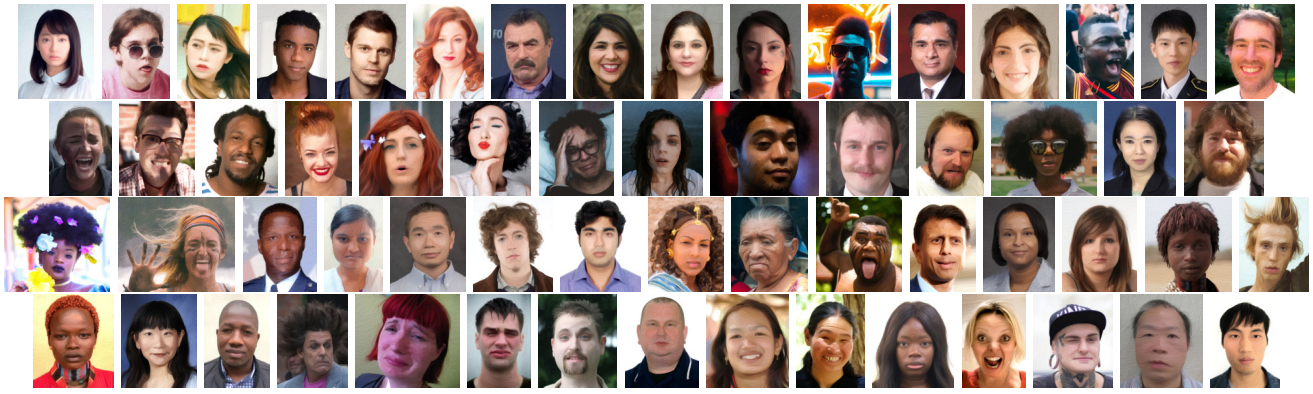

Fig. 19 Images from *NPRportrait1.0* benchmark stylised as oil paintings: Semmo *et al.* [15] and ranked according to Experiment2.

Fig. 20 Images from *NPRportrait1.0* benchmark stylised in the Julian Opie style: Rosin and Lai [10] and ranked according to Experiment2.

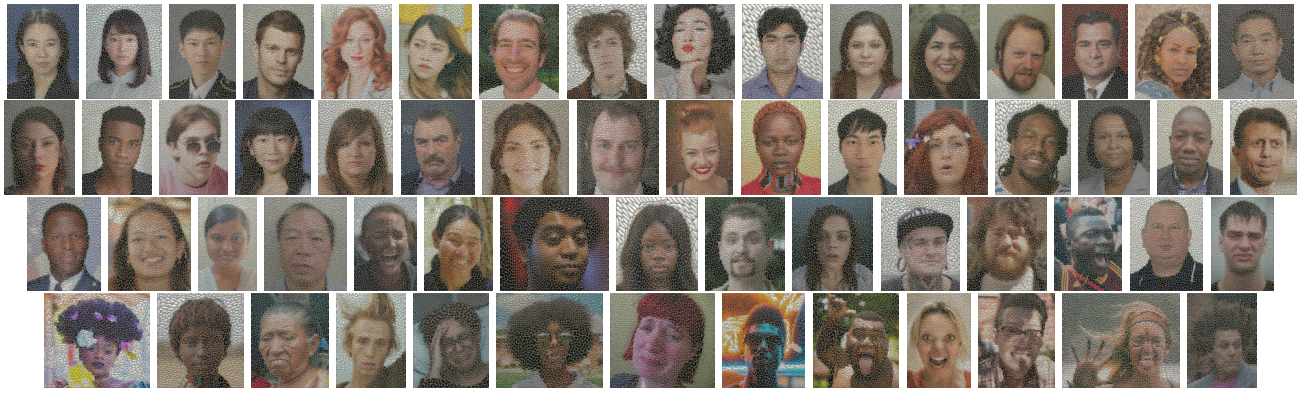

Fig. 21 Images from *NPRportrait1.0* benchmark stylised as pebble mosaics: Doyle *et al.* [3] and ranked according to Experiment2.

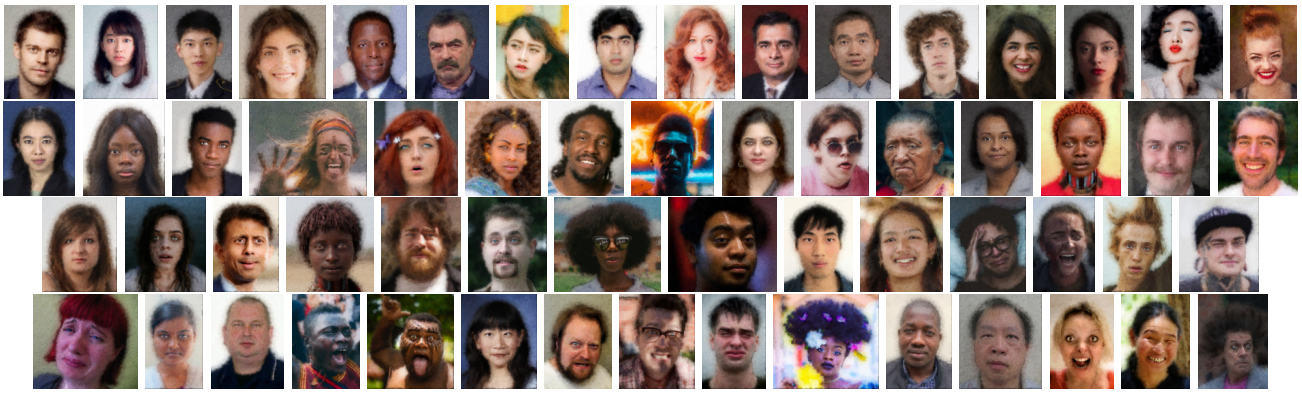

Fig. 22 Images from *NPRportrait1.0* benchmark stylised as watercolours: Rosin and Lai [11] and ranked according to Experiment2.

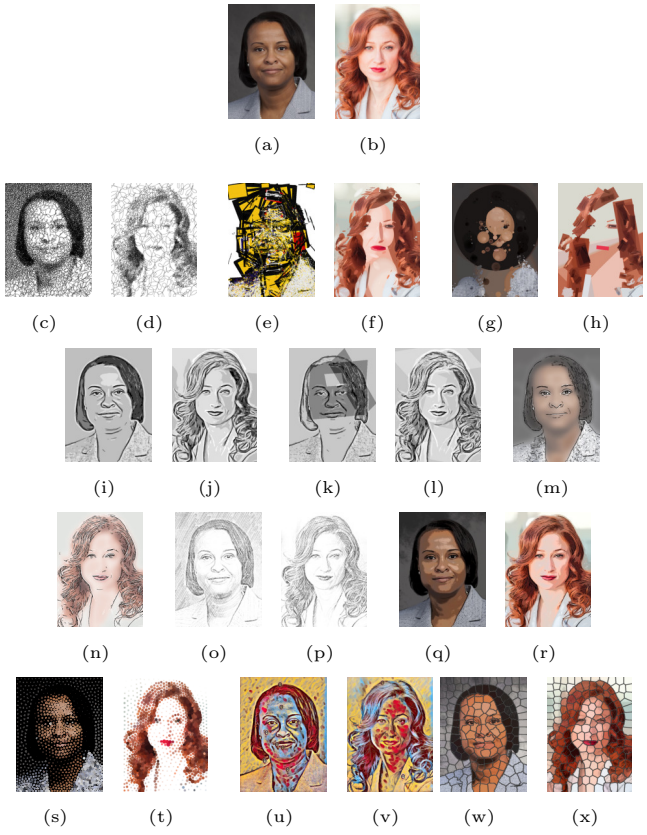

Fig. 23 Additional stylisations. (a) & (b) source images from level 1, (c) & (d) circular scribble art [2], (e)–(h) structured abstraction [4], (i) & (j) 3-tone, 3 level pyramid, line and region-based stylization [9], (k) & (l) as previous example, but with parallelogram region stylization [9], (m) & (n) line drawing overlaid on simplified and diffused colour image [9], (o) & (p) hatching effect using XDoG [18], (q) & (r) cartoon effect using XDoG [18], (s) & (t) stippling with coloured, overlapping stipples [14], (u) & (v) neural style transfer [6], (w) & (x) pebble mosaic [5].