Applying Medical Imaging Tractography Techniques to Painterly Rendering of Images

Alberto Di Biase[1]

[1]National Heart and Lung Institute, Faculty of Medicine, Imperial College London, UK
adibiase@ic.ac.uk

November 4, 2025

Abstract

Doctors and researchers routinely use diffusion tensor imaging (DTI) and tractography to visualize the fibrous structure of tissues in the human body. This paper explores the connection of these techniques to the painterly rendering of images. Using a tractography algorithm the presented method can place brush strokes that mimic the painting process of human artists, analogously to how fibres are tracked in DTI. The analogue to the diffusion tensor for image orientation is the structural tensor, which can provide better local orientation information than the gradient alone. I demonstrate this technique in portraits and general images, and discuss the parallels between fibre tracking and brush stroke placement, and frame it in the language of tractography. This work presents an exploratory investigation into the cross-domain application of diffusion tensor imaging techniques to painterly rendering of images. All the code is available at `https://github.com/tito21/st-python`.

Figure 1: Example of stylization using tractography.

1 Introduction

An invaluable tool for visualizing the fibrous structure of tissue in the human body is tractography. Tractography algorithms use the diffusion tensor to determine the local orientation of fibres in tissue, and then place streamlines along these orientations to create a 3D representation of the fibre pathways. This problem is surprisingly similar to the placing coherent brush strokes in painterly rendering of images.

This paper explores this similarity in a cross-domain framing, applying tractography techniques to place brush strokes in images. I present the parallels between the two problems and introduce the computer graphics community to the mathematical framework of tractography highlighting the potential for cross-disciplinary collaboration.

The strokes of a painter follow the contours and edges of the image subject. This motivates the need for the extraction of gradient based quantities. In my algorithm, this is done by the eigen decomposition of the structural tensor (analogue to the diffusion tensor). Next, the extracted vectors are followed with a simple tractography algorithm and then rendered with an artistic brush stroke. An example of an image stylized with the proposed method is shown in figure 1.

The article is divided as follows: first, an overview of the diffusion tensor imaging including a review of the tractography algorithm is given, followed by a brief discussion of existing painterly rendering techniques. Next, a full description of the proposed algorithm is presented. After that, results from the NPRportrait and NPRgeneral are presented followed by discussions and conclusions.

1.1 Primer on Diffusion Tensor Imaging and Tractography

DTI is an image modality in Magnetic Resonance Imaging (MRI) that measures the diffusion of water across the body [15, 2]. One of the main applications of this technique is to visually represent the nerve tracts of the brain (and other organs), helping doctors and researchers to study neurological disease [11].

In DTI the acquired images are later used to fit the diffusion tensor of water for that voxel (pixel). This is a 3×3 symmetric matrix that quantifies how much water can diffuse in each direction. Because water in biological tissue is restricted by cell membranes the shape of the tensor is elongated in areas were a predominant orientation is present such as in the fibre tracts. In this case the tensor is said to be anisotropic. In the case of free water, there is no preferred orientation and the tensor is isotropic. To measure the amount of anisotropy in the tensor several metrics exist. The most common being the fractional anisotropy (FA) which is a scalar value ranging from 0 (completely isotropic) to 1 (completely anisotropic).

An eigenvalue decomposition is performed on the tensor to find the predominant orientation. Larger eigenvalues indicate more diffusion in the direction of the associated eigenvector. The eigensystem is generally sorted in descending order and the first eigenvector is followed to create streamlines that represent fibres in the tissue. This process is called tractography [15].

First introduced by Basser, et al. in 2000, [1] tractography consists of following the primary orientation of the diffusion tensor forming tracts that can be visualized in 3D. The most basic form of this is done by solving this differential equation:

$$\frac{d\vec{r}(s)}{ds} = \vec{u_1}(\vec{r}(s))$$

With $\vec{r}(s)$ being the tract and $\vec{u_1}$ the primary eigenvector (in diffusion the one associated with the largest eigenvalue, and in structural tensor the smallest). Solving this equation can be done by the well established Runge–Kutta methods [8] (RK2(3) is used in this work).

Solving the presented differential equation was the first method proposed for tractography but more advanced techniques exist. For example, probabilistic tractography methods that take into account uncertainty in the tensor estimation [3].

1.2 Related Works in image stylization

In the past other methods for image stylization have used the normal to the gradient to orient the brush strokes [13, 9]. These methods also achieve a painterly effect by following the direction normal to the gradient. The algorithms to place the brush strokes are determined by a simple procedure that can be viewed as solving the tractography differential equation with Euler's method.

The structural tensor as a measure of orientation produces smoother brush strokes without the need to do additional post-processing. Additionally, it is more robust to noise and can provide more information about the image structure. Metrics such as the coherence [10] of the tensor can be used to control the brush stroke width and length. In animation the structural tensor can also be used in 3D to orient the brush strokes in a coherent manner [7]. To the best of my knowledge the structural tensor has not been used for image stylization in the literature.

2 Methods

2.1 Structural Tensor

For non diffusion images an analogue to the diffusion tensor is the structural tensor [12, 4]. This tensor captures the local changes in intensity of the image. Areas with flat texture will exhibit an isotropic tensor, while areas with edges or linear structures will exhibit an anisotropic tensor oriented along the structure. As with the diffusion tensor an eigenvalue decomposition is performed to extract the local orientations.

The structural tensor is defined by the tensor that minimizes the changes in intensities (V) when moving in an arbitrary direction. The general solution is $S = \sum \nabla V (\nabla V)^T$, where ∇V is the gradient of the image and the summation is over a

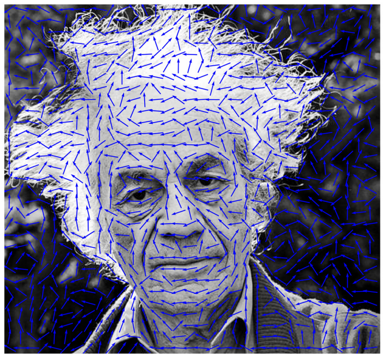

(a)

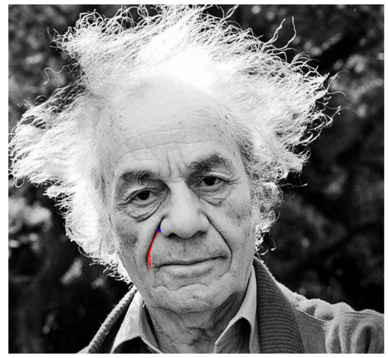

(b)

Figure 2: Example of the structure tensor. The arrows are the primary eigenvector of the structure tensor (a). Example of a tractography tract. The tract starts from the blue point (b).

neighbourhood of the image. In practice the structure tensor is calculated using the following procedure:

1. Calculate the derivatives I_x and I_y using a Sobel filter or a derivative of Gaussian filter. (in this work a derivative of Gaussian with standard deviation $\sigma = 1.0$ is used).

2. Average the products of the gradients with a Gaussian filter with a standard deviation ρ (in this work $\rho = 1.0$ is used): $(\overline{I_{ij}} = K_\rho * (I_i I_j))$.

3. Build the tensor as:

$$S = \begin{bmatrix} \overline{I_{xx}} & \overline{I_{xy}} \\ \overline{I_{xy}} & \overline{I_{yy}} \end{bmatrix}$$

Again, if there are areas of uniform intensities in one direction, the tensor will be oriented in that direction and is anisotropic. The other case, an isotropic tensor would be an area with uniform intensities in all directions. The orientation of the structures corresponds to the direction of constant intensities (small gradients). To find the predominant orientation, the eigenvalue decomposition is performed. However, the eigenvalues are sorted in ascending order (reverse of the diffusion tensor) because a smaller eigenvalue indicates the predominant orientation. This is because smaller changes in intensity indicate that structure is present along that direction. See [7] for more discussion on the structure tensor. The predominant orientation is visualized in figure 2a. It can be seen that the arrows are aligned with the features of the image.

2.2 Structural Tensor vs Gradient

To demonstrate the improvement of the structural tensor over the gradient for defining the orientation of the brush strokes I implemented the same filter using the gradient. The gradient is calculated using a derivative of Gaussian filter (with standard deviation $\sigma = 1.0$) and the tracts are followed in the same way as with the structural tensor.

2.3 Image Stylization Filter

For the generation of painterly images I use a multilayer approach [9] where the image is first downsampled and lower-resolution (thicker) strokes are first rendered followed by higher-resolution (thinner) strokes rendered on top. This follows the painting process of most artists where the low-frequency content is painted first, and later the details are painted on top. The following paragraphs describe the steps for each of the layers.

First, the image converted to greyscale and is downsampled to a lower resolution by a factor of σ. Then the structure tensor and the associated eigensystem is calculated. The image is divided into a grid of square cells with the size of the brush stroke width. If the mean Euclidean distance between the colours of the original image and the current image is above a user defined threshold (colour threshold) in the cell a stroke is followed by the tractography algorithm. The strokes are followed for a set length or until the tract reaches an image border or until the coherence [10], a measure of anisotropy of the local orientation, of the pixel is below a user defined parameter (0.5 is used throughout). Coherence for the structural tensor in 2D is defined as $\alpha = ((\lambda_1 - \lambda_2)/(\lambda_1 + \lambda_2))^2$ with λ_1 and λ_2 the largest and smallest eigenvalues of the tensor, respectively. The order in which the cells are drawn is set at random.

Once the tract is obtained it is simplified using the Ramer–Douglas–Peucker [16] algorithm and converted to a Bézier curve using the algorithm from Schneider [19]. The colour of the stroke is determined by the middle point of the tract. The width of the stroke is controlled by the user. In this implementation a simple circular brush stroke rendering is used, but any brush rendering algorithm can be used.

These steps are repeated for 4 layers with decreasing values of σ and brush stroke length and width. The parameters where selected empirically for a visually appealing result and are shown for each layer in table 1. The layers are rendered one on top of the other. Each layer provides more detail to the image. The first layer uses large brush strokes to render the low frequency content of the image. The number of strokes and resolutions are increased in each layer to provide more detail.

2.4 Implementation Details

A reference implementation of the algorithm is available at `https://github.com/tito21/st-python` with examples of usage and the data and code necessary to reproduce the results in this article. The code is written in Python and uses Pycairo for rendering. To sample the vector field during the tractography process bilinear inter-

	Layer 1	Layer 2	Layer 3	Layer 4
Scale factor (σ)	10	5	1	0.5
Length of strokes	1000	500	100	100
Width of strokes	50	25	5	2.5
Colour threshold	100	100	50	50

Table 1: Algorithm parameters. Each layer is rendered on top of the previous one. The scale factor is the blurring standard deviation for each level. The length of the strokes is the maximum length of the tract to be followed in pixels. The width of the strokes is the width of the brush stroke in pixels. The colour threshold is the maximum error (Euclidean distance in RGB space) allowed to render a stroke in that cell.

polation is used. Execution time for a 2000×1300 image is approximately 15 minutes on a standard laptop computer, though no optimization for speed was attempted.

3 Results

3.1 Structural Tensor and Gradient Comparison

In figure 3 the results of the stylization using the gradient and the structural tensor are shown. Both images are generated using the same parameters. The image on the left corresponds to the gradient and the image on the right corresponds to the structural tensor. The gradient based method produces jaggier brush strokes and results in less clear definition of the image features. Close tracts in the structural tensor are more aligned and produce a smoother texture.

3.2 Evaluation Datasets

To evaluate the proposed filter I use the NPRportrait 1.0 [17] and NPRgeneral [14] datasets. The first dataset consists of portraits at three levels of difficulty. The first level the subject is looking straight into the camera with a neutral expression and with a simple background. At the second level the images might have some background content and the expression are not necessary neutral. In the final third level the images have more complex lighting and backgrounds. All levels are matched by gender and ethnicity. The NPRgeneral dataset consists of a variety of images with different subjects and backgrounds.

For this exploratory investigation I focus on a qualitative description of the results on the standard datasets. A formal evaluation of the performance against other stylization methods and a more extensive ablation study of the parameters should be conducted in future work.

The stylized images from NPRportrait level 3 are shown in figure 4, and in figure 5 the stylized images from NPRgeneral are shown. The reader is encouraged to view

Figure 3: Stylized images using the gradient (left) and the structural tensor (right)

the full resolution images in the project repository. In appendix B the original images are shown for comparison, and in appendix A the stylized images from the other NPRportrait levels are shown.

Images with textures in a linear direction (for example the hair in the cat from NPRgeneral) give the most interesting effect because each stroke follows the hair.

In images with high frequency content or texture (for example arch in NPRgeneral) the stylization somewhat destroys the details (however gives an interesting effect) because they are smaller than the textures. For images with smooth gradients (for example headlight in NPRgeneral) the stylization is not as effective, given flat images similar to the original.

Because the colour is taken directly from the underlying image the filter performs equally under any illumination conditions. The same applies for low contrast situations.

Small details such as eyes are sometimes not rendered properly either because the brush stroke is too thick or because there are no tracts in that area. The latter problem can be solved by increasing the number of strokes in that layer or by using a smarter algorithm to select the starting points for the tracts.

I emphasize that this work is an exploratory investigation into the cross-domain application of diffusion tensor imaging techniques to painterly rendering and the proposed method has not been optimized for quality or speed. Other computer graphics experts could improve on the proposed algorithm with their expertise in the field of non-photorealistic rendering.

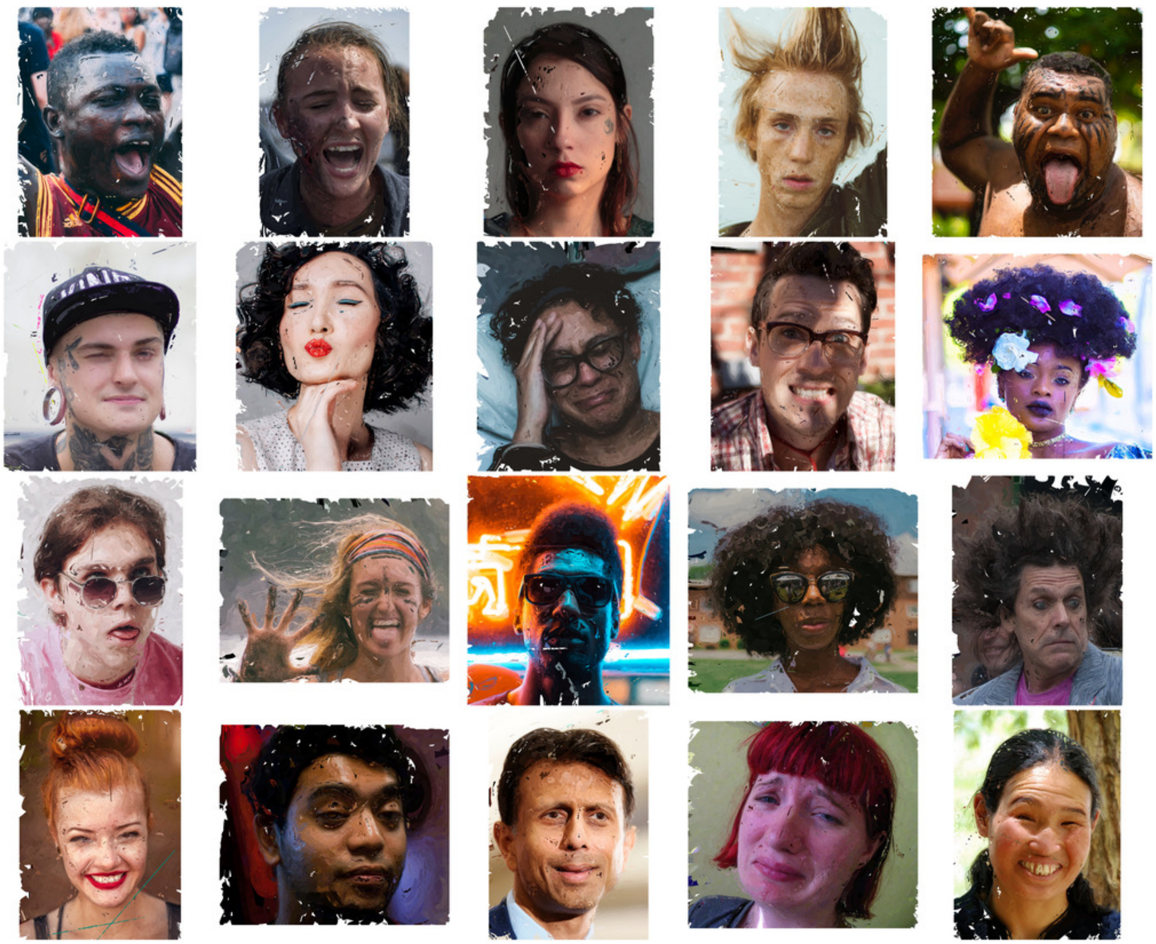

Figure 4: Stylized images from NPRportrait 1.0 level 3

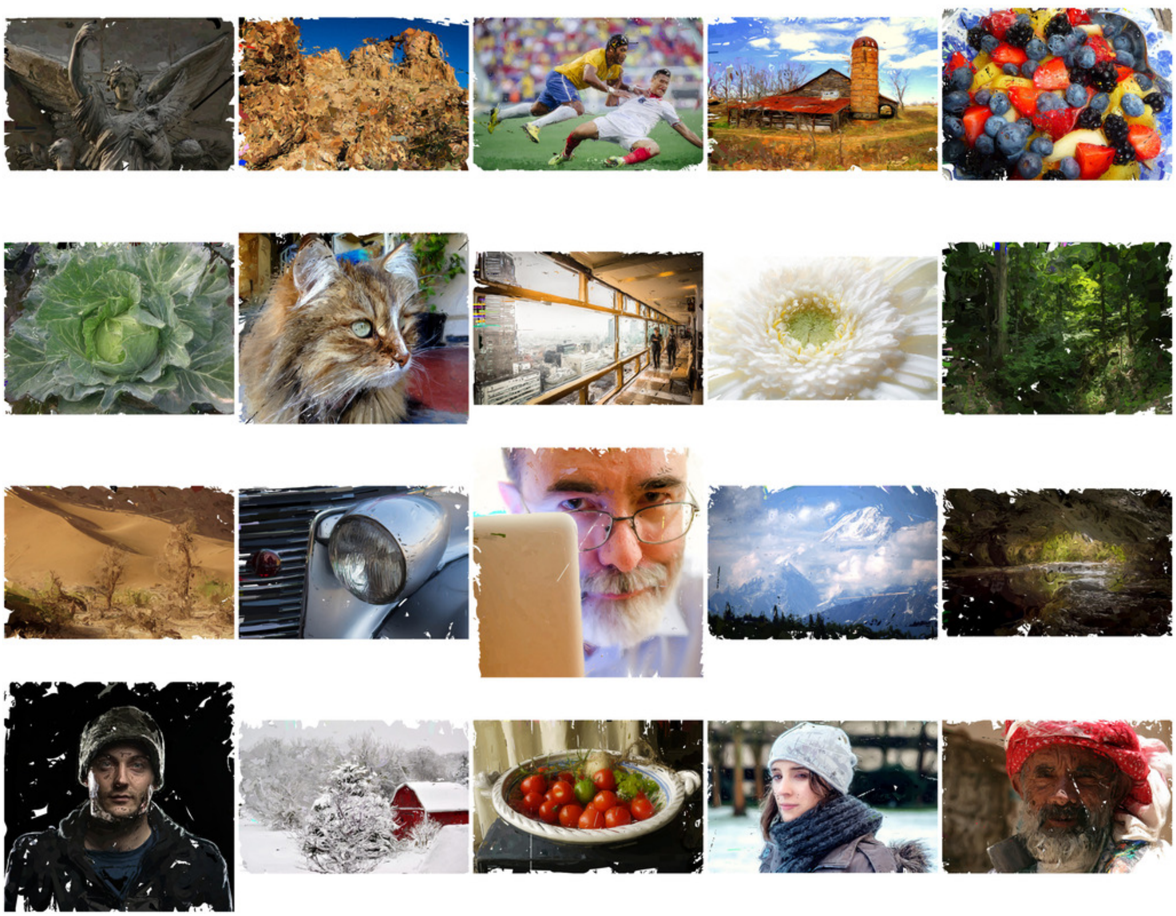

Figure 5: Stylized images from NPRgeneral

4 Discussion

4.1 Limitations and Future Work

This article focused only on the selection of brush strokes from an image. The method can select the lines to be drawn to replicate a target image by following its contours. Once the brushes have been selected, other parameters, such as the colour and width could be adjusted. The selection of these parameters can be based on the target style and requires further research.

The basic implementation presented in this article preserves the original colours of the image, but it is well known that artists play with colour for effect. A simple randomization in the colour selection could give the desired effect, or a more complex colour-picking selection can be used to give a particular effect.

This implementation uses a simple brush stroke rendering that consists of drawing simple Bézier curves between the points of the tract. More sophisticated brush rendering algorithms can be used to give a more realistic effect [20, 5, 6]. Additional data can be used to render the strokes such as angle and pressure can be estimated from the tract and used to aid the brush rendering.

As in more advanced tractography algorithms, the selected strokes can be further processed to avoid similar or overlapping strokes by removing strokes that are too close to each other. Other post-processing of the strokes could merge crossing strokes for an effect that would be more difficult to represent by a real artist.

The tractography algorithm presented here corresponds with the first presented in the literature during the early 2000s. More robust techniques have been developed since then that improve the quality of the tracts in the presence of noise [18]. Also, it is common to use a segmented approach to tractography where the image is segmented into regions and the tracts are followed in each region. For this filter, face features can be extracted first and then followed by tractography, ensuring that there is appropriate detail in the most important areas.

A limitation of this work is that the strokes are selected in a greyscale version of the image. The interaction of colours in how the tracts are selected was not explored. Future work could include following tracts based on other colour properties, such as hue instead of brightness, or in other colour spaces.

A final limitation is a lack of extensive algorithm evaluation. Future work should include a formal evaluation of the performance against other stylization methods and a more extensive ablation study of the parameters of the algorithm.

4.2 Differences and Similarities between Medical Imaging and Painterly Rendering

Medical imaging tractography and stroke placement for painterly rendering are surprisingly similar problems. In both cases the local structure of the object under study is extracted using the eigen decomposition of a tensor (diffusion tensor or structural tensor). Then, the primary eigenvector is followed to create either fibre tracts or brush strokes.

However, there are some differences in the two problems. In medical imaging the goal is to reconstruct the 3D structure of the fibres in the body as accurately as possible, while in painterly rendering the goal is to create visually appealing brush strokes that convey the essence of the original image. This means that in painterly rendering there is more freedom to deviate from the original image for artistic effect.

Interestingly the diffusion imaging community and the computer graphics community developed similar measures of anisotropy (FA in DTI and coherence in structural tensor analysis). The medical imaging community routinely takes advances from computer graphics to improve their algorithms, for example when selecting interpolation methods for vector fields, however the opposite is not common. This work hopes to inspire researchers from both fields to collaborate and further improve the results.

5 Conclusion

This work presents an exploratory investigation into applying medical imaging concepts to the painterly rendering of images. I present the problem of brush placement in the language of tractography, and demonstrate the use of the structural tensor to guide the placement of brush strokes. This exploration will hopefully inspire researchers from both fields to collaborate and further improve the results.

Acknowledgement

AI Chat bots were used to help with the writing of this article and in assisting with writing the code. This work was done independently and does not represent the views of Imperial College London.

References

[1] P. J. Basser et al. “In Vivo Fiber Tractography Using DT-MRI Data”. In: *Magnetic Resonance in Medicine* 44.4 (Oct. 2000), pp. 625–632. ISSN: 0740-3194. DOI: `10.1002/1522-2594(200010)44:4<625::aid-mrm17>3.0.co;2-o`. PMID: `11025519`.

[2] Peter J Basser and Denis LeBihant. “Fiber Orientation Mapping in an Anisotropic Medium with NMR Diffusion Spectroscopy”. In: ().

[3] T. E. J. Behrens et al. "Probabilistic Diffusion Tractography with Multiple Fibre Orientations: What Can We Gain?" In: *NeuroImage* 34.1 (Jan. 1, 2007), pp. 144–155. ISSN: 1053-8119. DOI: 10.1016/j.neuroimage.2006.09.018. URL: https://www.sciencedirect.com/science/article/pii/S1053811906009360 (visited on 10/31/2025).

[4] Josef Bigun and Gosta H. Grandlund. "Optimal Orientation Detection of Linear Symmetry". In: *Proceedings of the IEEE First International Conference on Computer Vision*. London: IEEE Computer Society Press, 1987, pp. 433–438. URL: https://www2.hh.se/staff/josef/publ/publications/bigun87london.pdf.

[5] N.S.H. Chu and Chiew-Lan Tai. "Real-Time Painting with an Expressive Virtual Chinese Brush". In: *IEEE Computer Graphics and Applications* 24.5 (Sept. 2004), pp. 76–85. ISSN: 1558-1756. DOI: 10.1109/MCG.2004.37. URL: https://ieeexplore.ieee.org/abstract/document/1333630/authors (visited on 08/01/2025).

[6] Cassidy J. Curtis et al. "Computer-Generated Watercolor". In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '97*. The 24th Annual Conference. Not Known: ACM Press, 1997, pp. 421–430. ISBN: 978-0-89791-896-1. DOI: 10.1145/258734.258896. URL: http://portal.acm.org/citation.cfm?doid=258734.258896 (visited on 08/05/2025).

[7] Vedrana Andersen Dahl. *Structure Tensor Computation, Visualization, and Application*. June 8, 2019. URL: https://people.compute.dtu.dk/vand/notes/ST_intro.pdf (visited on 07/31/2025).

[8] Paul DeVries. *A First Course in Computational Physics*.

[9] Aaron Hertzmann. "Painterly Rendering with Curved Brush Strokes of Multiple Sizes". In: *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '98*. The 25th Annual Conference. Not Known: ACM Press, 1998, pp. 453–460. ISBN: 978-0-89791-999-9. DOI: 10.1145/280814.280951. URL: http://portal.acm.org/citation.cfm?doid=280814.280951 (visited on 08/02/2025).

[10] "Tensor Methods". In: *Spatio-Temporal Image Processing: Theory and Scientific Applications*. Ed. by Bernd Jähne. Berlin, Heidelberg: Springer, 1993, pp. 143–157. ISBN: 978-3-540-48145-4. DOI: 10.1007/3-540-57418-2_20. URL: https://doi.org/10.1007/3-540-57418-2_20 (visited on 08/03/2025).

[11] Ben Jeurissen et al. "Diffusion MRI Fiber Tractography of the Brain". In: *NMR in Biomedicine* 32.4 (2019), e3785. ISSN: 1099-1492. DOI: 10.1002/nbm.3785. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/nbm.3785 (visited on 07/31/2025).

[12] Hans Knutsson, Carl-Fredrik Westin, and Mats Andersson. "Representing Local Structure Using Tensors II". In: *Image Analysis*. Ed. by Anders Heyden and Fredrik Kahl. Vol. 6688. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 545–556. ISBN: 978-3-642-21226-0 978-3-642-21227-7. DOI: 10.1007/978-3-642-21227-7_51. URL: http://link.springer.com/10.1007/978-3-642-21227-7_51 (visited on 08/01/2025).

[13] Peter Litwinowicz. "Processing Images and Video for an Impressionist Effect". In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques - SIGGRAPH '97*. The 24th Annual Conference. Not Known: ACM Press, 1997, pp. 407–414. ISBN: 978-0-89791-896-1. DOI: 10.1145/258734.258893. URL: http://portal.acm.org/citation.cfm?doid=258734.258893 (visited on 08/02/2025).

[14] David Mould and Paul L. Rosin. *A Benchmark Image Set for Evaluating Stylization*. The Eurographics Association, 2016. ISBN: 978-3-03868-002-4. URL: https://doi.org/10.2312/exp.20161059 (visited on 07/31/2025).

[15] Lauren J. O'Donnell and Carl-Fredrik Westin. "An Introduction to Diffusion Tensor Image Analysis". In: *Neurosurgery clinics of North America* 22.2 (Apr. 2011), pp. 185–viii. ISSN: 1042-3680. DOI: 10.1016/j.nec.2010.12.004. PMID: 21435570. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3163395/ (visited on 07/31/2025).

[16] Urs Ramer. "An Iterative Procedure for the Polygonal Approximation of Plane Curves". In: *Computer Graphics and Image Processing* 1.3 (Nov. 1, 1972), pp. 244–256. ISSN: 0146-664X. DOI: 10.1016/S0146-664X(72)80017-0. URL: https://www.sciencedirect.com/science/article/pii/S0146664X72800170 (visited on 10/19/2025).

[17] Paul L. Rosin et al. "NPRportrait 1.0: A Three-Level Benchmark for Non-Photorealistic Rendering of Portraits". In: *Computational Visual Media* 8.3 (Sept. 2022), pp. 445–465. ISSN: 2096-0662. DOI: 10.1007/s41095-021-0255-3. URL: https://ieeexplore.ieee.org/document/10897536/ (visited on 07/27/2025).

[18] Tabinda Sarwar, Kotagiri Ramamohanarao, and Andrew Zalesky. "Mapping Connectomes with Diffusion MRI: Deterministic or Probabilistic Tractography?" In: *Magnetic Resonance in Medicine* 81.2 (2019), pp. 1368–1384. ISSN: 1522-2594. DOI: 10.1002/mrm.27471. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27471 (visited on 08/01/2025).

[19] Philip J. Schneider. "AN ALGORITHM FOR AUTOMATICALLY FITTING DIGITIZED CURVES". In: *Graphics Gems*. Elsevier, 1990, pp. 612–626. ISBN: 978-0-08-050753-8. DOI: 10.1016/B978-0-08-050753-8.50132-7. URL: https://linkinghub.elsevier.com/retrieve/pii/B9780080507538501327 (visited on 10/19/2025).

[20] Songhua Xu et al. "Advanced Design for a Realistic Virtual Brush". In: *Computer Graphics Forum* 22.3 (2003), pp. 533–542. ISSN: 1467-8659. DOI: 10.1111/1467-8659.t01-2-00701. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.t01-2-00701 (visited on 08/01/2025).

Appendices

A Stylized images from the NPRportrait 1.0 dataset

Figure 6: Stylized images from NPRportrait 1.0 level 1

Figure 7: Stylized images from NPRportrait 1.0 level 2

B Original images

Figure 8: Original images from NPRportrait 1.0 level 1

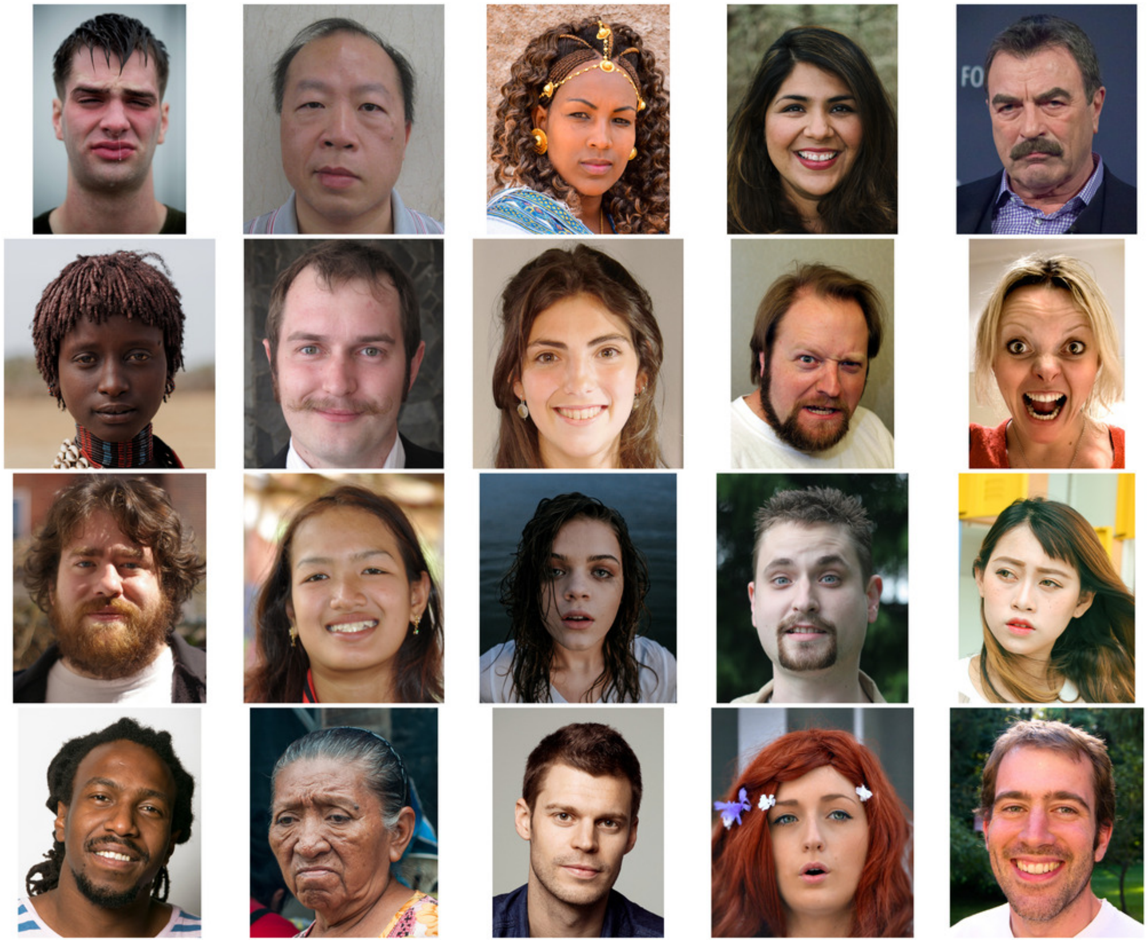

Figure 9: Original images from NPRportrait 1.0 level 2

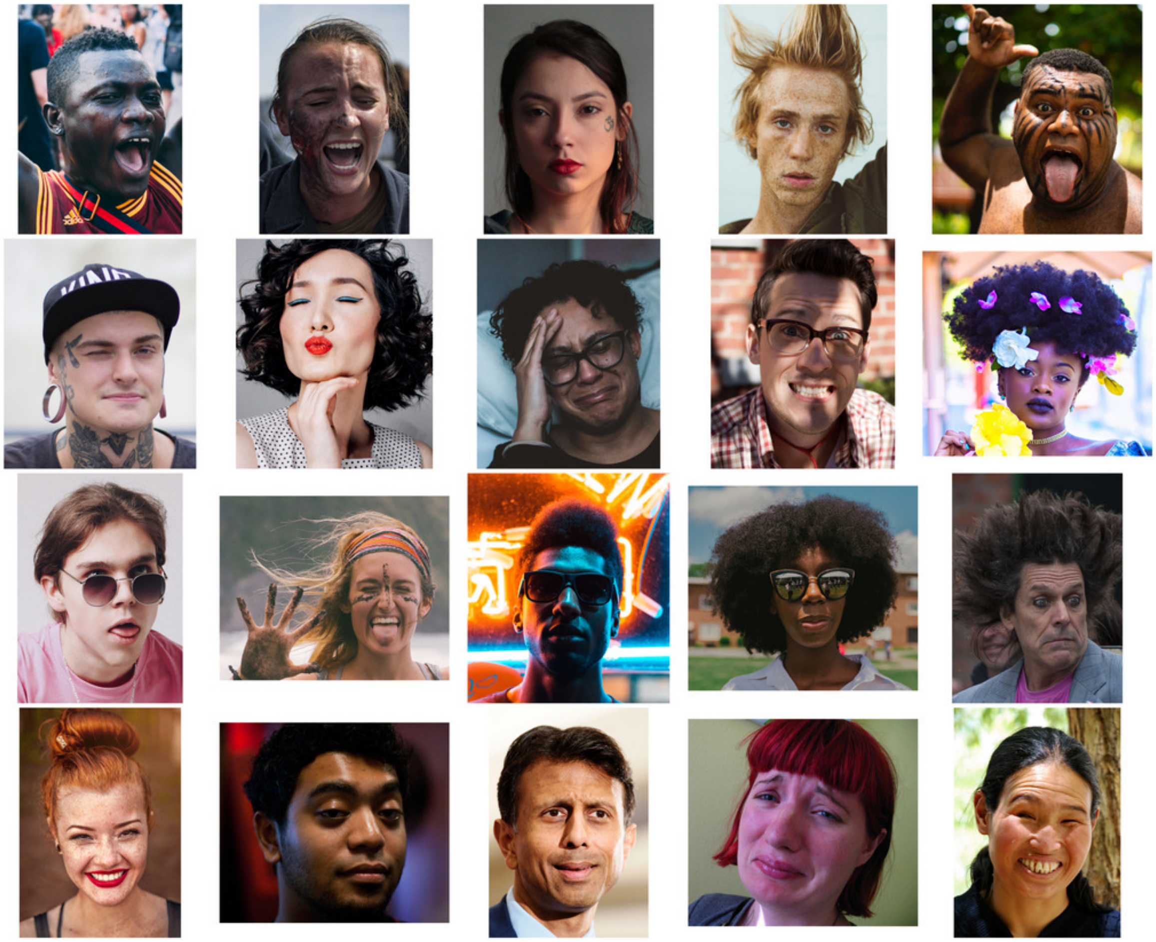

Figure 10: Original images from NPRportrait 1.0 level 3

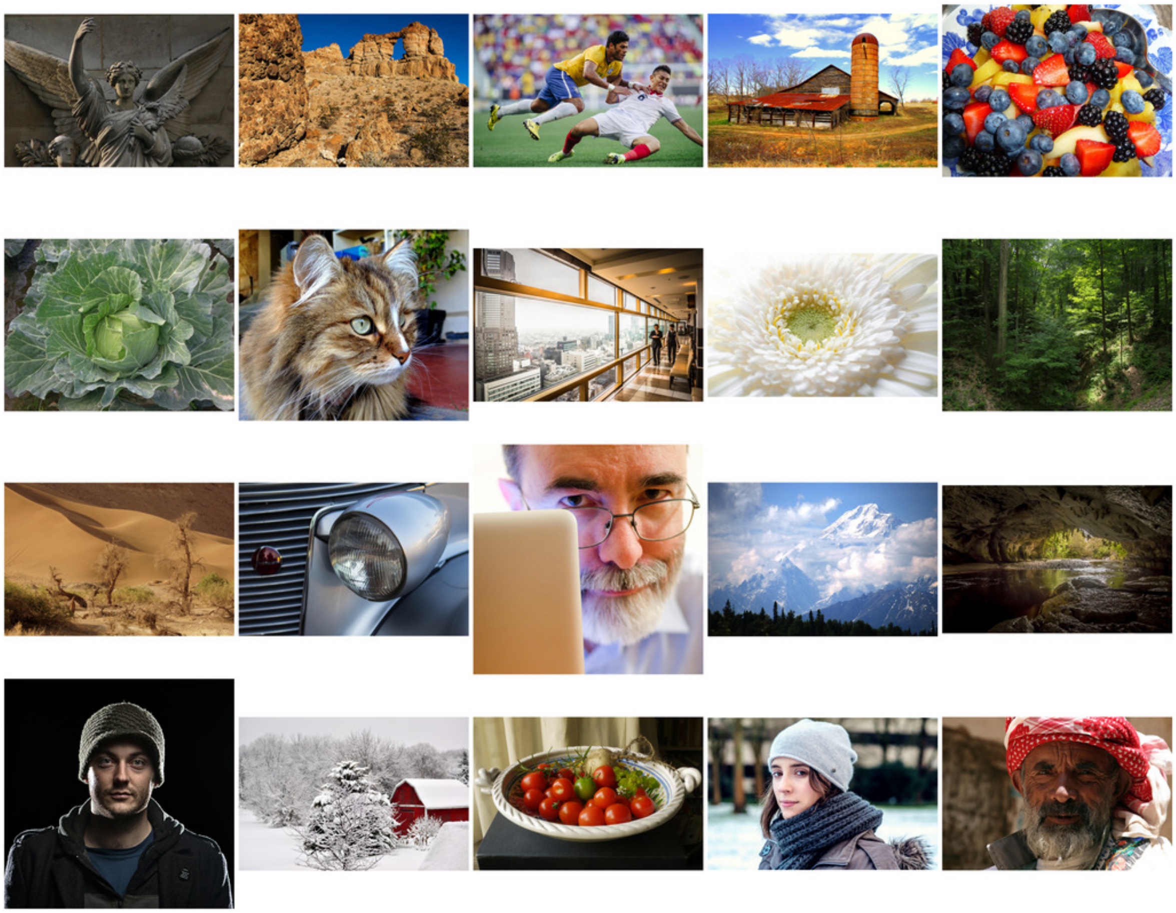

Figure 11: Original images from NPRgeneral