https://doi.org/10.2352/CIC.2022.30.1.22

Digital Restoration of Lost Art: Applying the Colorization Transformer to the Ghent Altarpiece panels

Milan Kresović and Jon Yngve Hardeberg, Norwegian University of Science and Technology, Gjøvik, Norway

Abstract

In recent years, image processing has proven to be a great tool to help document, preserve, and restore art pieces, especially visual art. One example is using colorization techniques when the original color information of the image has been lost or is unavailable. To expand on that, we used Ghent Altarpiece as the study case. One of the panels of this polyptych has been lost, and only an archival photograph exists. Using the state-of-the-art colorization method, we want to digitally restore what has been lost in its original form. In this work, we proposed a pipeline that consists of a colorization transformer (ColTran) trained on the captured patch-based imaging data of the Ghent Altarpiece. We evaluated the proposed pipeline and addressed its strengths and weaknesses. Moreover, we presented planned future steps and improvements for this project.

Index Terms—Lost art, Colorization, Ghent Altarpiece

Introduction

In recent years, image processing has proven to be a great tool that can help in the documentation, preservation, and restoration of art pieces, especially in visual art. The task of digital image restoration can be described as a computer vision task where the main objective is the reconstruction or estimation of the digital image that has been somehow damaged. Sometimes, the original color information of the image has been lost or is not available. In that case, the colorization technique is usually applied. This technique tries to solve an ill-posed problem of estimating color information from grayscale image values. The use cases of this technique in art reconstruction are characterized by the fact that the original art piece is lost, but the archival black-and-white mediums, like photographs, are preserved.

The Ghent Altarpiece by The Van Eyck brothers is considered one of the seminal works of the 15th-century [6]. Presented in Figure 1, this polyptych consists of multiple panels. Two of the lower panels were stolen. Fortunately, one was returned but the other one - *"The Just Judges"* - remains lost. In 2020 a major restoration process was completed by Royal Institute for Cultural Heritage (KIK-IRPA) [4]. The captured image data from this project motivated many different image processing researchers to contribute [5, 16, 18].

As previously mentioned, "The Just Judges" panel is still lost, but one archival black-and-white photograph obtained before the theft has been preserved. Our motivation stems from here: by using the state-of-the-art colorization method, we want to digitally restore what has been lost in its original form. Unfortunately, as acquiring the original black-and-white photographs proved challenging in this time frame, we decided to focus on developing the colorization pipeline that could be used on the black-and-white photographs once they are acquired. Therefore, instead of using black-and-white photographs of the original and stolen artifacts, we use color and grayscale images of the restored Ghent Altarpiece panels to develop and validate our methodology. In this paper, we propose a pipeline that consists of a colorization transformer [13] trained on the captured patch-based image data of the Ghent Altarpiece. Additionally, we propose extending the proposed pipeline with color correction and image stitching in the future. In summary, this paper presents these key points:

- We present a novel pipeline for colorization of the grayscale captured image data of the Ghent Altarpiece
- We evaluate the pipeline and address its strengths and weaknesses
- We propose future steps and improvements

The paper is organized as follows. First, we provide an overview of the existing colorization methods. The following section defines the overall proposed pipeline, and it gives details of the colorization transformer. Furthermore, we describe the used dataset, implementation, and training details. Next, the visual inspection and evaluation metrics obtained using the proposed pipeline are presented. The results are discussed, and the limitations of the pipeline are laid out. Finally, the last section concludes this paper.

Related works

One inspiration for this paper was the work done by Demetriou et al. in [7, 8]. In their paper, they focus on reclaiming the lost art from art books in which the records of lost artworks are only preserved. As these art books were produced in the past century, printed captures of art pieces in the books are of low resolution with the additional problem of halftone patterns. To correct these and reclaim the lost art, example-based super-resolution [24], inverse halftoning [20], and color correction [22] have been applied.

Image colorization is an ill-posed stochastic problem that needs semantic information to work well [13]. Before the prevalence of the deep learning methods, traditional colorization algorithms usually needed human input, which would provide valuable hints [2, 14, 23]. With this idea in mind, exemplar-based techniques started to emerge. Here, a reference image is used from which the colors are copied to the colorized image [15, 19, 21]. In recent years, more advanced automated deep learning methods have been developed. Some networks condition color information on the pixel intensity [10, 11, 25]. Others use unconditional generative colorization networks (like GANs, VAEs) and modify them to condition on the grayscale images [1, 12]. Besides these, there are also autoregressive approaches that we used for this project. Specifically, PixColor [9], and ColTran [13]

Figure 1. The Ghent Altarpiece inner panels with the stolen panel framed in red. Sint-Baafskathedraal Gent - www.artinflanders.be - Dominique Provost en Hugo Maertens - CC-BY-NC-ND 4.0.

networks use an autoregressive approach to do colorization and upsampling at the same time to obtain the finalized colorized image.

Methodology

This paper presents the proposed colorization pipeline framed in red in Figure 2. However, in the same figure, we give the whole workflow we plan to create in the future and how the colorization pipeline fits into it. The idea would be to estimate color from grayscale image patches using the proposed colorization pipeline, then color-correct the obtained patches, and stitch them back together to derive the final image.

In the case of the colorization pipeline, the first step is to crop the original captured grayscale images of the panels in small equally-sized patches. These patches can then be used to train the colorization network. For the colorization network, we used Colorization Transformer [13] developed by a team at Google.

Colorization transformer

Colorization transformer (ColTran) [13] is a three-stage probabilistic model that uses self-attention to colorize grayscale images, or in our case patches. The idea behind this network is that by using axial transformers, we first map low-resolution grayscale images to the 3-bit colored images. In order for this to work, axial attention is used. That implies applying self-attention alternatingly on each row and column independently, thus capturing the global receptive field. After obtaining a colorized low-resolution image, color and spatial upsamplers are used where all pixels are modified simultaneously. Color upsampler generates 8-bit colored images, whereas spatial upsampler increases the colorized image resolution to the original grayscale image resolution. The loss function of the ColTran is described as follows:

$$\mathscr{L} = (1-\lambda)\log p_c + \lambda \log \widetilde{p}_c + \log \widetilde{p}_{c\uparrow} + \log \widetilde{p}_{s\uparrow} \quad (1)$$

It is a simple negative log likelihood-based loss where the network tries to maximize per pixel-channel distributions outputs p_c, $\widetilde{p}_c$, $\widetilde{p}_{c\uparrow}$, $\widetilde{p}_{s\uparrow}$ after colorization, color, and spatial upsampling, respectively. λ is a hyperparameter that controls the contribution of the low-resolution colorization outputs in the overall loss.

Experiments

This section gives an overview of the acquired captured data of the Ghent Altarpiece. Additionally, information about training and implementation is given.

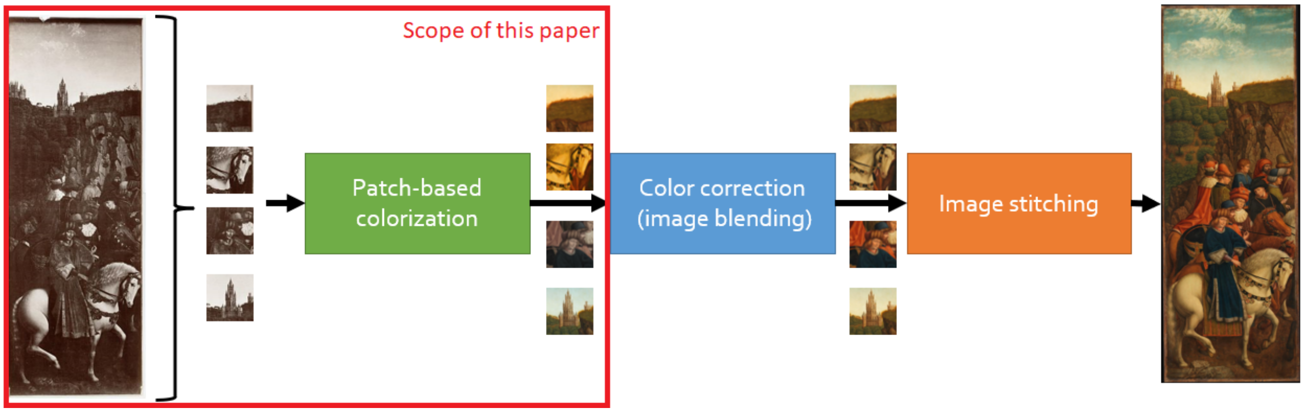

Figure 2. *The workflow for digital restoration of lost panels with the colorization pipeline framed in red*

Dataset

The dataset was acquired courtesy of Art in Flanders, an institution that supported restoration efforts made between 2012 and 2020. From this project, all lower panels of the Ghent Altarpiece have been restored. We received captures of the inner panels both before the restoration process and after. We have obtained in total 55 *GB* of data. Side panels sizes were 6500×16000px, whereas the central panel was of size 26000×16000px.

As the capture images were very large, they had to be border-trimmed and cropped to patches of size 250×250px with 100px step for easier image stitching that we intend to include in the future. The Just Judges panel was excluded from the training set, which, in the end, totaled 5475 patches. Figure 3 shows training set patches examples.

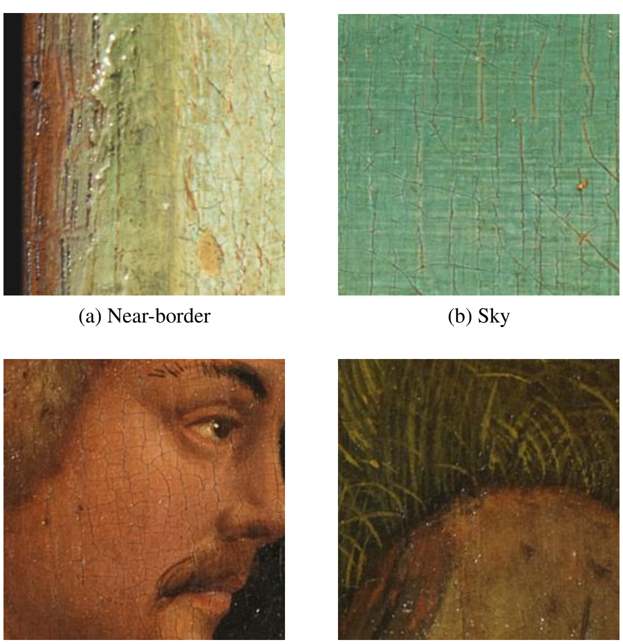

(a) Near-border (b) Sky

(c) Human face (d) Nature

Figure 3. *Different types of patches that were present in the dataset*

Implementation

ColTran was trained on this custom dataset in three steps. First, we trained the low-resolution colorization model - *ColTran core*. The second step is to train a model that upsamples color information - *Color upsampler*. Lastly, we trained a model that increases the resolution of the colorized patch - *Spatial upsampler*. After this, the inference step is run, from which we obtain the colorized patch. Each training step took around 2.5 days on a 4 GPUs NVIDIA Tesla V100 16Gb machine. Inference mode took 3 hours on 10 GPUs. The training loss functions for the ColTran network can be seen in Figure 4. Additional information on training can be found in Table 1.

Table 1: Training information

Hyperparameter	Value
Optimizer	RMSProp
Learning rate	3e-4
Num. of Steps	13k
Batch size	32

Results and discussion

Figure 5 shows one example of patch outputs in different steps of the colorization process. Figure 5b shows a colorized low-dimension patch that uses 3-bit colors. In the next step, low color resolution is upsampled, followed by image resolution upsampling to match the original patch. Here, we can see that the network struggles with red tones, but it manages to reconstruct the overall color.

We can see some differences if we compare original patches with colorized patches, but the overall color is restored. This is further illustrated in Figure 6. The blue colors of the sky are, in most cases, well reconstructed. That is not surprising if we have in mind that most of the panel images, and thus the dataset, contains sky in one-quarter of the full image. Therefore, the network is well ”aware” of this case. On visual inspection, we can also see that neutral colors are well preserved, as in the case of the castle patch. Furthermore, the spatial upsampler managed to maintain the spatial variation of the original patch. Even though the network learned the general color information of the panels, it still makes mistakes with some colors. That is especially noticeable in the last example, where blue shades are mistaken for brown.

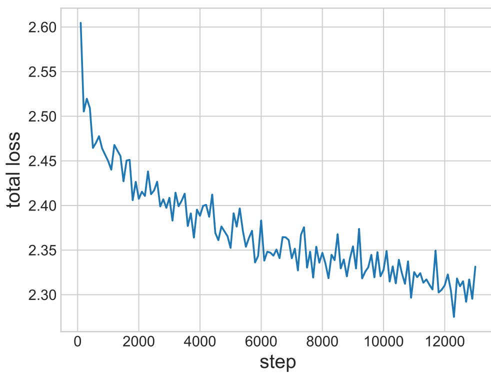

(a) Color upsampler loss

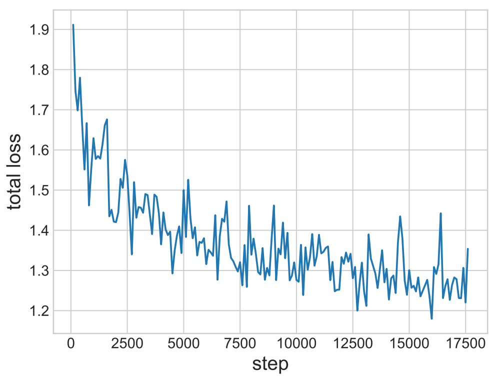

(b) Spatial upsampler loss

Figure 4. *The loss functions for different steps during the training*

To measure the performance of the ColTran network, we have calculated SSIM, PSNR, and ΔE^*_{ab} difference. We calculated these metrics for pairs of original and colorized patches in the whole dataset, and the results are presented in Table 2. According to SSIM, we can see that the spatial variation is well preserved; therefore, the spatial upsampler shows good performance. Additionally, if we observe the ΔE^*_{ab} difference value, we can see that overall, the color information is acceptably reconstructed.

Table 2: Metrics for colorized patches

Metric	Mean	Median	Min	Max	95th
SSIM	0.944	0.995	0.119	0.997	0.996
PSNR	27.388	27.638	5.571	39.247	35.672
ΔE^*_{ab}	10.687	9.661	2.624	99.145	18.805

In Figure 7, we plotted the a*b* values of colors found in original (red points) and colorized (blue points) patches in the whole dataset. We can see that the set of colors is mostly overlapping, with colorized patches having some extreme values due to the colorization error.

We combined multiple connected colorized patches as seen in Figure 8. We notice the purple pixels representing the buildup of the colorization error, which follows the observations made for data in Figure 7.

(a) Original patch

(b) Grayscale

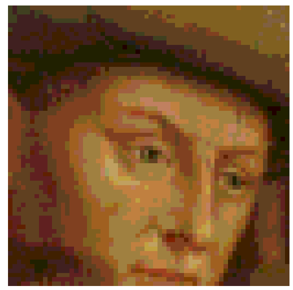

(c) Low-resolution colorization

(d) Color upsampled

(e) Spatially upsampled (final)

Figure 5. *Evolution of patch for different steps in the colorization network*

Conclusion

In this paper, we wanted to address the problem of digital restoration of art pieces that have been lost where some archival photographs have been preserved. Our study case includes the *Ghent Altarpiece* and, more specifically, *The Just Judges* panel. We have modified and trained the ColTran network on the Ghent Altarpiece imaging data. We have shown that our model is capable of colorizing the grayscale captures of this Van Eyck's masterpiece. Nonetheless, the colorization network could be improved by focusing on the color artifacts that the network generates. These artifacts are known as color biases, and simple $L^*a^*b^*$ Euclidean distance in the loss function should be expanded by including priors for the color palette to constrain the colors generated by the network as suggested in [3].

In the future, the proposed workflow containing the colorization pipeline, color correction, and image stitching step from Figure 2 should be fully implemented and evaluated. Moreover, this project should be expanded by using the pre-theft black-and-white archival photographs of the Ghent Altarpiece to dgitally restore the original panel fully. Additionally, [17] showed that by analyzing the pearls in the panels of the Ghent Altarpiece, they could

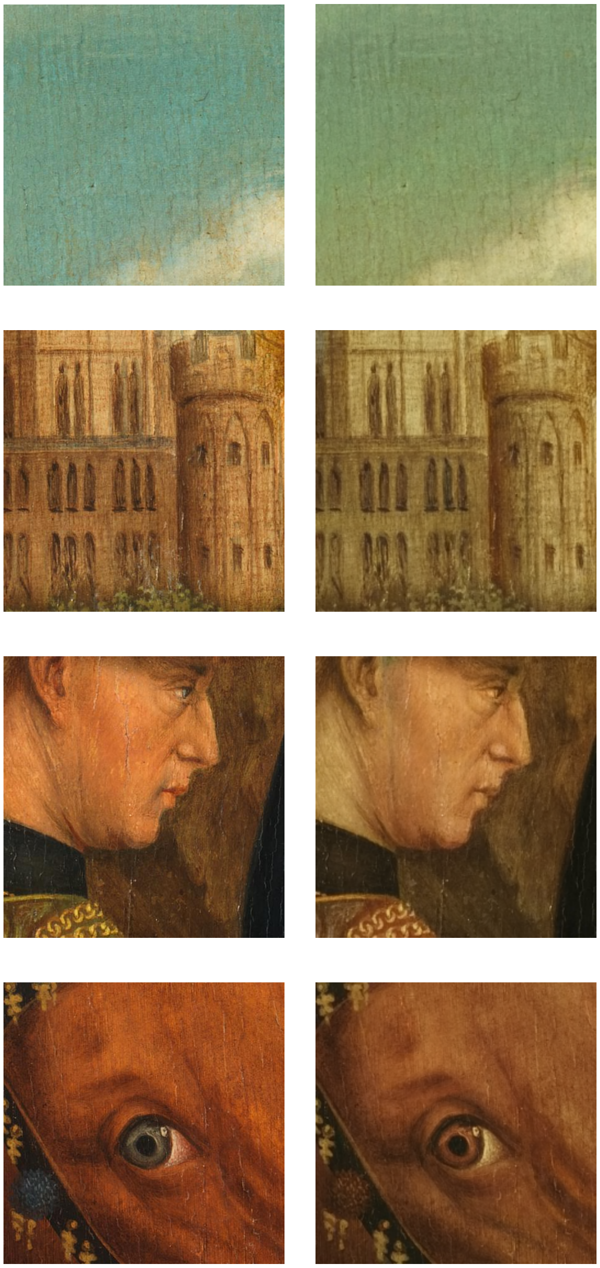

(a) Original (b) Colorized

Figure 6. Original vs. colorized patches

indicate which panels were painted by which Van Eyck brother. In other words, each pearl has an individual painter's signature. If we apply the same idea to the colorized pre-theft black-and-white photograph of the original panel, we should see the difference in the signature of the pearls in the original and copied panel.

Acknowledgments

We would like to thank Art in Flanders for providing the captures of the Ghent Altarpiece needed for this project. We also thank Gabriel Adelmann for proposing the original idea of applying advanced imaging methods to the field of lost art, and for his continued support, fruitful discussions, and insight into the topic of art preservation and restoration.

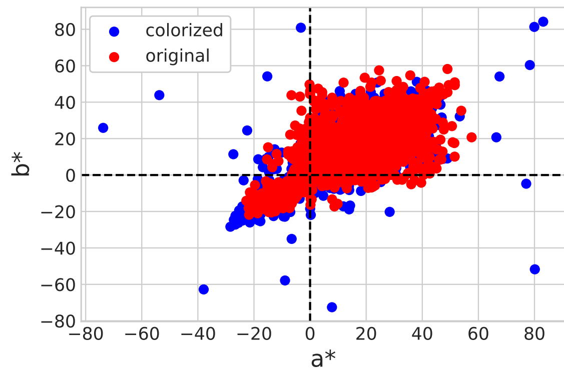

Figure 7. a*b* pixel values of original and colorized patches

Figure 8. Connected colorized patches.

References

[1] Yun Cao, Zhiming Zhou, Weinan Zhang, and Yong Yu. "Unsupervised Diverse Colorization via Generative Adversarial Networks." In: *ECML/PKDD (1)*. Ed. by Michelangelo Ceci, Jaakko Hollmén, Ljupco Todorovski, Celine Vens, and Saso Dzeroski. Vol. 10534. Lecture Notes in Computer Science. Springer, 2017, pp. 151–166.

[2] Yuanzheng Ci, Xinzhu Ma, Zhihui Wang, Haojie Li, and Zhongxuan Luo. "User-Guided Deep Anime Line Art Colorization with Conditional Adversarial Networks". In: *Proceedings of the 26th ACM International Conference on Multimedia*. New York, NY, USA: Association for Computing Machinery, 2018, pp. 1536–1544.

[3] Irina-Mihaela Ciortan, Sony George, and Jon Yngve Hardeberg. "Colour-Balanced Edge-Guided Digital Inpainting: Applications on Artworks". In: *Sensors* 21.6 (2021). ISSN: 1424-8220.

[4] *Closer to Van Eyck*. (Accessed on 12/05/2021). URL: `http://closertovaneyck.kikirpa.be/`.

[5] Bruno Cornelis, Tijana Ružić, E. Gezels, Ann Dooms, Aleksandra Pižurica, Ljiljana Platiša, Jan Cornelis, Maximiliaan Martens, Marc De Mey, and Ingrid Daubechies. "Crack detection and inpainting for virtual restoration of

paintings: The case of the Ghent Altarpiece”. In: *Signal Processing* 93.3 (Mar. 2013), pp. 605–619.

[6] De Mey, Marc and Martens, Maximiliaan and Stroo, Cyriel, ed. *Vision & material: interaction between art and science in Jan van Eyck's time*. KVAB Press, 2012, p. 330.

[7] Maria Lena Demetriou, Jon Yngve Hardeberg, and Gabriel Adelmann. “Computer-Aided Reclamation of Lost Art”. In: *Computer Vision – ECCV 2012. Workshops and Demonstrations*. Springer Berlin Heidelberg, 2012, pp. 551–560.

[8] Maria-Lena Demetriou. “Computer-Aided Reclamation of Lost Art”. MA thesis. Norwegian University of Science and Technology, 2012, p. 94.

[9] Sergio Guadarrama, Ryan Dahl, David Bieber, Jonathon Shlens, Mohammad Norouzi, and Kevin Murphy. “Pix-Color: Pixel Recursive Colorization”. In: *British Machine Vision Conference 2017, BMVC 2017, London, UK, September 4-7, 2017*. BMVA Press, 2017.

[10] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. “Let there be color!” In: *ACM Transactions on Graphics* 35.4 (July 2016), pp. 1–11.

[11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. “Image-to-Image Translation with Conditional Adversarial Networks”. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, July 2017.

[12] Diederik P. Kingma and Max Welling. “Auto-Encoding Variational Bayes”. In: *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*. 2014.

[13] Manoj Kumar, Dirk Weissenborn, and Nal Kalchbrenner. “Colorization Transformer”. In: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. 2021.

[14] Anat Levin, Dani Lischinski, and Yair Weiss. “Colorization using optimization”. In: *ACM SIGGRAPH 2004 Papers on - SIGGRAPH '04*. ACM Press, 2004.

[15] Yuji Morimoto, Yuichi Taguchi, and Takeshi Naemura. “Automatic colorization of grayscale images using multiple images on the web”. In: *SIGGRAPH 2009: Talks on - SIGGRAPH '09*. ACM Press, 2009.

[16] Aleksandra Pižurica, Ljiljana Platiša, Tijana Ružić, Bruno Cornelis, Ann Dooms, Maximiliaan Martens, Helene Dubois, Bart Devolder, Marc De Mey, and Ingrid Daubechies. “Digital Image Processing of The Ghent Altarpiece: Supporting the painting's study and conservation treatment”. In: *IEEE Signal Processing Magazine* 32.4 (July 2015), pp. 112–122.

[17] Ljiljana Platiša, Bruno Cornelis, Tijana Ružić, Aleksandra Pižurica, Ann Dooms, Maximiliaan Martens, Marc De Mey, and Ingrid Daubechies. “Pearls and beads in Jan van Eyck's paintings”. In: *Vision and material : interaction between art and science in Jan van Eyck's time, Abstracts*. Brussels, Belgium, 2010.

[18] Zahra Sabetsarvestani, Barak Sober, Catherine Higgitt, Ingrid Daubechies, and Miguel R. D. Rodrigues. “Artificial intelligence for art investigation: Meeting the challenge of separating x-ray images of the Ghent Altarpiece”. In: *Science Advances* 5.8 (Aug. 2019).

[19] Tomihisa Welsh, Michael Ashikhmin, and Klaus Mueller. “Transferring color to greyscale images”. In: *ACM Transactions on Graphics* 21.3 (July 2002), pp. 277–280.

[20] Ping Wah Wong. “Inverse halftoning and kernel estimation for error diffusion”. In: *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society* 4 4 (1995), pp. 486–98.

[21] Chufeng Xiao, Chu Han, Zhuming Zhang, Jing Qin, Tien-Tsin Wong, Guoqiang Han, and Shengfeng He. “Example-Based Colourization Via Dense Encoding Pyramids”. In: *Computer Graphics Forum* 39.1 (Apr. 2019), pp. 20–33.

[22] Wei Xu and Jane Mulligan. “Performance Evaluation of Color Correction Approaches for Automatic Multi-view Image and Video Stitching”. In: June 2010, pp. 263–270.

[23] Liron Yatziv and Guillermo Sapiro. “Fast image and video colorization using chrominance blending”. In: *IEEE Transactions on Image Processing* 15.5 (May 2006), pp. 1120–1129.

[24] Roman Zeyde, Michael Elad, and Matan Protter. “On Single Image Scale-Up Using Sparse-Representations”. In: *Curves and Surfaces*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 711–730.

[25] Richard Zhang, Phillip Isola, and Alexei A. Efros. “Colorful Image Colorization”. In: *Computer Vision – ECCV 2016*. Springer International Publishing, 2016, pp. 649–666.

Author Biography

Milan Kresović completed his master's degree in the Computational Color and Spectral imaging program in 2022 at the Norwegian University of Science and Technology, Norway. He obtained his BSc in Computer Science & Control Systems in 2020 from the Faculty of technical sciences, University of Novi Sad, Serbia.

Jon Y. Hardeberg received his sivilingeniør (MSc) degree in signal processing from the Norwegian Institute of Technology (now NTNU) in Trondheim, Norway in 1995, and his PhD from Ecole Nationale Supérieure des Télécommunications in Paris, France in 1999. He is now Professor of Colour Imaging at NTNU, Department of Computer Science, The Norwegian Colour and Visual Computing Laboratory, Gjøvik, Norway. His current research interests include spectral imaging, image quality, colour management, material appearance, cultural heritage imaging, and medical imaging, and he has co-authored more than 300 publications within the field. He has led several research projects funded by the Research Council of Norway, been NTNU's representative in two Erasmus Mundus Joint Master Degrees (CIMET and COSI), and the coordinator of three Marie Sklodowska Curie ITN projects (CP7.0, ApPEARS, CHANGE).