Preprint

A Deep Active Learning Framework for Crack Detection in Digital Images of Paintings

Nicolas Nadisic[a,b,*], Yoann Arhant[a,c], Niels Vyncke[a], Sebastiaan Verplancke[a], Srdan Lazendić[a], Aleksandra Pižurica[a]

[a] *Ghent University, St.-Pietersnieuwstraat 41, 9000 Ghent, Belgium*
[b] *Royal Institute for Cultural Heritage (KIK-IRPA), Jubelpark 1, 1000 Brussels, Belgium*
[c] *Royal Military Academy, Renaissancelaan 30, 1000 Brussels, Belgium*

Abstract

Paintings deteriorate over time due to aging and storage conditions, with cracks being a common form of degradation. Detecting and mapping these cracks is crucial for art analysis and restoration but it presents challenges. Traditional methods often require tedious manual effort, while deep learning (DL) relies on large annotated datasets, which are expensive to produce. Also, DL does not generalize well, in the sense that it is conditioned by the properties of the training data and often performs poorly on unseen data with slightly different properties. To address these issues, we developed a deep active learning (DAL) method called DAL4ART. DAL methods start with minimal annotated data, perform their task, and then retrain iteratively on newly annotated samples to improve efficiency. This iterative learning process makes our method require less data, learn progressively from human input, handle partially annotated data, and perform better on previously unseen paintings. Additionally, our method can integrate various imaging modalities and is equipped with a user-friendly web interface. We demonstrate the application of the proposed crack detection tool in a concrete use case as a means of supporting the restoration of old master paintings.

Keywords: Digital painting analysis; crack detection; deep active learning.

1. Introduction

Paint cracking is the most common type of deterioration encountered in old paintings. Cracks, also called craquelures, appear on layers of paint due to several factors such as the aging of the underlying material (wooden panel or support, canvas), oxidation of the varnish layer, or inadequate storage conditions including fluctuations of humidity and temperature. See fig. 1 for an example.

* Corresponding author.
E-mail address: nicolas.nadisic@ugent.be

In the last two decades, the digitization of paintings and the increasingly high resolution of digital images have created opportunities for the automated analysis of paintings (Cornelis et al., 2011; Pizurica et al., 2015; Sizyakin et al., 2020; Sober et al., 2022). The detection and mapping of cracks contribute to many tasks in art investigation and restoration. For example, how cracks develop and spread depends greatly on the choice of raw materials and the techniques used by the artist. As such, a thorough analysis of cracks can help to date a painting, to judge its authenticity, to understand the technical aspects of its realization and factors of its degradation (Bucklow, 1997). A precise study of these factors can then support preventive measures and bring insight to prepare a restoration (Abas, 2004). An advantage of this technique over physical or chemical methods is that it is not invasive and does not damage the painting (Bucklow, 1997).

Apart from standard digital photography, art investigation also benefits from other imaging technologies, such as infrared macrophotography, infrared reflectography, or X-radiography; see fig. 1 for an illustration. Those modalities are complementary, in the sense that a type of crack that is hardly distinguishable with one imaging modality might stand out much clearer with another modality.

(a) Visual macrophotography (b) Infrared reflectography (c) X-radiography

Fig. 1: Illustration of a portion of a painting showing cracks, in several imaging modalities. These images are from the panel *Virgin Annunciate* of the Ghent Altarpiece, publicly available on the website of the *Closer to Van Eyck* project[1]. Image copyright: Ghent, Kathedrale Kerkfabriek, Lukasweb.

Existing methods for automatic crack detection are based either on filtering or on machine learning, see section 2 for a detailed review. Filtering-based methods, see for example Gupta et al. (2008), use different kinds of gray-scale morphological filters to increase the contrast between cracks and background, and then apply thresholding to extract the cracks. Machine learning methods learn from previously annotated data how to perform the crack detection task, and they have shown good performance in previous studies (Giakoumis et al., 2005; Spagnolo and Somma, 2010; Cornelis et al., 2013b). However, they still require tedious manual effort in feature engineering and hyperparameter tuning. Methods based on deep learning can learn features from data and show promising results (Sizyakin et al., 2020, 2022a,b). However, deep learning requires a substantial amount of previously annotated data to perform well. For crack detection in paintings, such annotated data is extremely scarce because the annotation needs to be done manually by expert art historians, thus limiting the practical applicability of deep learning and supervised learning in general. Furthermore, the crack detection task is complex because of the variety in the types of painting and in the shapes and patterns of cracks, which makes applying deep learning even more challenging, especially on previously unseen paintings.

The active learning paradigm can help tackle the challenge of data scarcity. Deep active learning (DAL) methods start with a deep learning model trained on little annotated data, perform their task, and then suggest to the oracle (for example a human annotator) new data to annotate, before retraining the model. By retraining the model iteratively in an efficient way, active learning typically needs much less data than traditional deep learning, learns continuously from new annotations by a human-in-the-loop, enables learning from partially annotated data, and performs better on

[1] https://closertovaneyck.kikirpa.be

previously unseen data. To the best of our knowledge, DAL has never been used for crack detection in paintings; the main objective of this work is therefore to develop a DAL framework for this purpose.

Our contributions are as follows:

- We make significant additions to existing deep learning models, in order to enable their use in a deep active learning framework. This includes fast re-learning, continuous learning, and transfer learning techniques.
- We introduce DAL4ART, a deep active learning framework for crack detection in multimodal digital images of paintings. It is designed to integrate deep learning models for crack detection into an active learning process, with iterations of the loop "annotation-training-prediction".
- We present the web interface of DAL4ART, designed to be easily usable by art historians who are not necessarily experts in computer vision tools and to facilitate efficient annotation.

The rest of this article is organized as follows. We introduce the state of the art in crack detection in section 2 and review the active learning paradigm in section 3. Section 4 describes DAL4ART, the deep active learning framework we developed. We present experimental results in section 5, and section 6 concludes the article.

DAL4ART is accessible at the following link, `https://dal4art.ugent.be`.

2. Crack detection

In the last two decades, many machine learning algorithms have been proposed to detect cracks in digital images of paintings. This includes methods based on vector classification (Giakoumis et al., 2005; Spagnolo and Somma, 2010; Cornelis et al., 2013b; Pizurica et al., 2015). A few methods combined filtering techniques with machine learning (Cornelis et al., 2013a), and others such as Huang et al. (2016) leveraged sparse representations. Those machine-learning-based methods still suffered from the need for hand-crafted features, requiring complex manual tuning. Methods based on deep learning (DL) learn relevant features from data and have led to tremendous improvements. In particular, convolutional neural networks (CNN) (Sizyakin et al., 2020) can capture patterns in images such as edges, textures, or shapes at various levels of abstraction.

A related task in digital painting analysis is inpainting or visual restoration. It consists in not only detecting the cracks and paint losses but also recovering the missing content, using machine learning (Ruzic et al., 2010; Ruzic, 2013; Pizurica et al., 2014) and more recently deep learning (Meeus et al., 2020; Sizyakin et al., 2022a,b). It has proven useful in preparing restoration works and also to analyze the paintings, for example by making more readable areas of text affected by paint loss.

Most of the methods discussed above apply only to single-modality images, and as such they do not leverage the rich information given by multiple modalities, that are now relatively common for art study and restoration. A few recent works leveraged this multimodality to improve crack detection, see for example Huang et al. (2020) and the references therein. The approach introduced by Sizyakin et al. (2020) is based on a convolutional neural network (CNN), handles multimodal data, and is to the best of our knowledge the current state of the art in crack detection in paintings.

One important limitation of existing works is the impossibility of performing online learning, that is, the update of existing weights of the neural network using new data. Above all, deep neural networks require large quantities of labeled data to achieve satisfying performance. In crack detection in paintings, this data is scarce as the annotation needs to be done by human experts. Above all, the crack detection task is complex because cracks have different shapes, thickness, and irregular patterns, they can be either lighter than the background or darker, and they can be visually similar to elements of the painting with elongated shapes such as hair, eyelashes, or writing. This diversity requires even more training data for the models to be accurate and robust. To tackle these issues, we developed an active learning framework. First, let us introduce the active learning paradigm.

3. Active learning

Active learning is a machine learning paradigm where an algorithm learns through an iterative process, by actively querying an oracle to label samples (Settles, 2008; Ren et al., 2021). This oracle is usually a human annotator, but it can sometimes be an automated system. The algorithm is retrained at each step with the newly labeled samples. By strategically selecting the samples to label or annotate, active learning models are more efficient and typically need much less data than traditional methods to reach the same accuracy. They also improve the labeling speed, by enabling the oracle to validate or reject the algorithm's results instead of labeling from scratch. Active learning was introduced to tackle some limitations of these traditional machine learning algorithms, notably their reliance on large datasets labeled beforehand. Active learning is especially useful when labeled data is scarce and annotation is complex and time-consuming, for example, when it needs to be done by human experts, as in crack detection in images of paintings.

One key idea of active learning is to focus on the most uncertain or informative samples. Instead of randomly selecting unlabeled data to send to the oracle, active learning algorithms employ various strategies to select samples that are expected to provide the most learning benefit. Common strategies include uncertainty sampling, query-by-committee, and diversity sampling. Uncertainty sampling selects samples where the model is least confident in its predictions, aiming to reduce uncertainty in the model. Query-by-committee methods maintain multiple models trained on the data and select samples where the models disagree the most, thus requiring more information. Diversity sampling prioritizes samples that are the least similar to the currently labeled data. The goal is to cover the widest possible range of features to make sure that the algorithm learns from diverse samples and is then more robust.

Another key property of active learning algorithms is their ability to retrain the underlying machine learning model efficiently, instead of relearning from scratch when new labeled samples are added. Several training strategies are possible, and they imply updating the weights of the neural network continuously, each time a new sample or batch of samples is added; see Shui et al. (2020) and the references therein for a more exhaustive study. Re-learning is the most naive strategy, as it implies retraining the model from scratch after adding the newly annotated samples to the existing training data. In practice, this technique is computationally expensive and is not practical for large models. Continuous learning considers the existing weights in the model and updates them finely using the new data. Last, transfer learning is a strategy where part of the neural network is frozen while the last few layers are retrained. Transfer learning is common for larger models such as U-Net as it ensures more computationally efficient learning and prevents "catastrophic forgetting", that is the forgetting of previously learned information upon learning new information.

When the algorithm being trained actively is a deep neural network, it is referred to as deep active learning (DAL). While there are numerous deep active learning approaches discussed in the literature (Ren et al., 2021), none of them have been applied or validated for problems similar to the one addressed in our manuscript. Our focus is on crack detection in very high-resolution and multimodal images, which presents unique challenges and necessitates tailored solutions.

4. Our proposed method: DAL4ART

In this section, we present DAL4ART, an active learning framework for crack detection in multimodal images of paintings. DAL4ART consists of a web interface to upload, manage, and annotate images, see fig. 2 for an illustration. In the backend, deep learning models run to perform the crack detection task.

We designed the web interface based on the existing library CVAT[2]. The choice of a web-based tool ensures that it suits the majority of users, independently from their hardware or software specifications. The interface supports different image formats for the input data and, above all, can handle any number of modalities in the input. An additional key feature of the user interface is the capability to fine-tune models specifically on selected areas of paintings. This functionality holds great significance for art historians who wish to apply the already-trained models in practical scenarios. By enabling fine-tuning on specific parts of the paintings, the user interface empowers art historians to refine the models according to their expertise and domain-specific requirements. The possibility to switch easily between modalities, for example, from standard photography to infrared and X-ray, is also a game changer. It is especially helpful when cracks and painting elements are very similar in one modality.

From the point of view of a user, crack detection with DAL4ART works as follows, see fig. 3 for an illustration:

1. The user uploads an image (or several in case of multiple modalities) of a painting.

[2] https://www.cvat.ai/

Fig. 2: Illustration of the DAL4ART web interface. The red filter represents annotated cracks and the green filter represents annotated non-cracks.

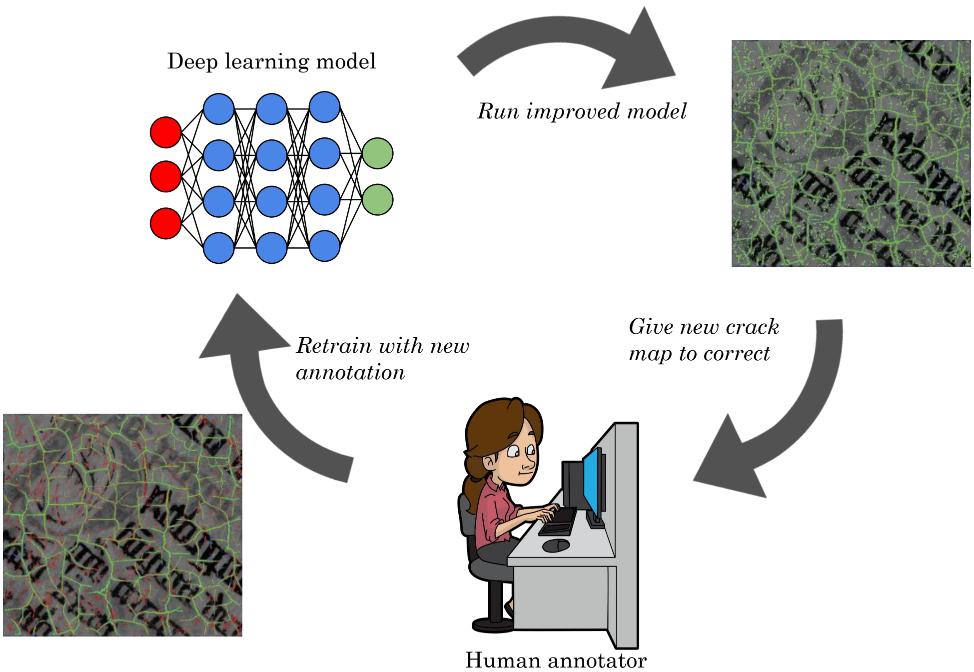

Fig. 3: Illustration of the deep active learning cycle.

2. The user annotates some of the cracks with a paintbrush-like tool on the web interface.
3. The deep learning model is updated using this new annotation.
4. The deep learning model runs and returns a crack map.
5. The user corrects the annotation partly by adding or removing zones to the crack map.
6. If the user is satisfied with the crack map, stop the process, otherwise do another iteration by returning to step 3.

As discussed in section 3, in standard active learning approaches, a key element is the querying or sampling strategy. In DAL4ART, rather than relying on predefined sampling strategies, we emphasize the control of the expert user and let them select the parts of the crack map to be corrected, by adding or removing zones in this crack map. This user-centric approach allows for more flexibility and customization in the annotation process and empowers the users to leverage their expertise effectively.

A significant contribution of our work is the adaptation of existing deep learning algorithms for crack detection so that they are easily integrated into an active learning framework, through efficient training strategies. We modified the convolutional neural network from Sizyakin et al. (2020), the current state-of-the-art deep learning model for crack detection, to allow continuous learning, that is the efficient update of the existing weights using little newly annotated data. The challenge of this adaptation is to integrate new knowledge without forgetting the previous knowledge and without overfitting the model to the task at hand.

5. Experiments

To measure the gain in annotation time provided by DAL4ART, we measured the time spent by a human operator to annotate cracks in 3 patches of 256×256 pixels taken from the Ghent Altarpiece, first completely manually, then using our deep active learning framework DAL4ART. See table 1 for numerical results. This experiment shows that

Table 1: Measured annotation times for three patches of 256×256 pixels of the Ghent Altarpiece.

	Patch 1	Patch 2	Patch 3	Average
Manual annotation	16min 20s	18min 31s	20min 15s	18min 22s
DAL4ART	6min 41s	8min 55s	7min 9s	7min 35s
Time reduction	59%	52%	65%	59%

DAL4ART reduces the annotation time by more than 50% in all 3 test cases, with an average reduction of 59%.

In fig. 4 we present the results of a crack detection task on a large patch of dimensions 2000×2000 pixels, which would take hours for an expert to annotate. This task is challenging because the patch features both cracks that are darker than the background (on the faces in the upper left-hand corner and the lower right-hand corner, and on the light beige zone in the upper right-hand corner) and cracks that are lighter than the background (on the hair and the pearls). Also, parts of the hair intrinsically resemble cracks because of an elongated shape. We observe that, without active learning, the pre-trained CNN produces a very noisy crack map. The beige zone is relatively well extracted, although with a lot of pixels wrongly classified as cracks. The faces are very noisy, and the cracks in the hair are not recognized at all. After a few iterations of active learning and a few manual annotations, we observe that the crack map produced improves fast. The noise in the beige zone and the faces is gradually removed, and the cracks in the hair are partially detected, while the hair itself is no longer classified as cracks. For the 2nd iteration, the new annotations focus on the most problematic zone (corresponding to hair) and the improvement is immediate. For the 3rd iteration, the new annotations focus on the left face. Afterward, the cracks in that zone are almost perfectly detected. This experiment shows that, by annotating only a fraction of the input image, we can obtain a crack map of high quality.

The accuracy of the annotation is difficult to measure given the absence of ground truth but, qualitatively speaking, we observed that the crack map produced with DAL4ART is more accurate than manual annotation, notably at the crack-paint borders. Also, the results of DAL4ART are qualitatively finer and more accurate than the crack maps produced by the pre-trained CNN model.

6. Conclusion

In this paper, we presented DAL4ART, a deep active learning framework equipped with a web interface for the detection of cracks in digital images of paintings. We showed empirically that it can reduce significantly the time required for an expert to annotate cracks while improving the accuracy. We hope this tool can be useful to the community of art historians and restorers, and that DAL4ART will self-improve through continuous learning with the input of its successive users. Crack detection is a key problem in other fields such as civil engineering, for example to monitor and rehabilitate roads and buildings, and adapting DAL4ART to these tasks is a promising research direction. We also hope DAL4ART can be a foundation to develop new active learning tools, in the world of art investigation and beyond.

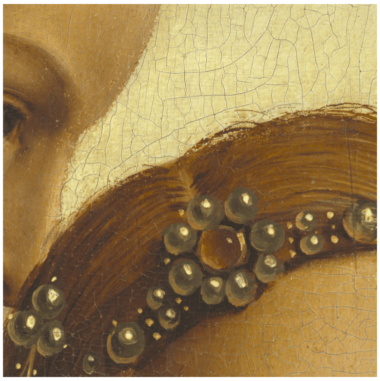

(a) Original patch (visual macrophotography)

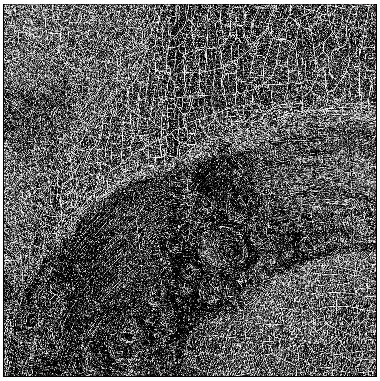

(b) Crack map produced by the pretrained CNN

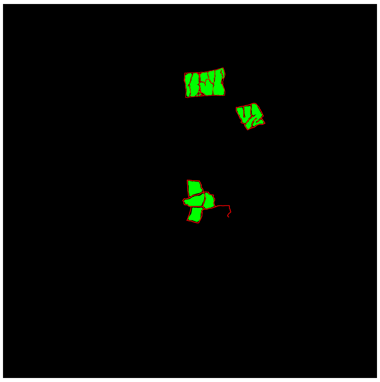

(c) Manual annotations for the 1st iteration of active learning.

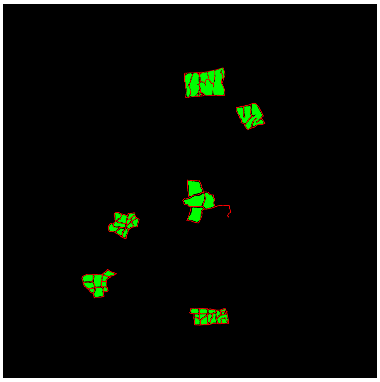

(d) Manual annotations for the 2nd iteration of active learning.

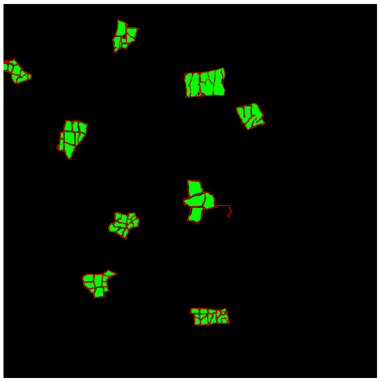

(e) Manual annotations for the 3rd iteration of active learning.

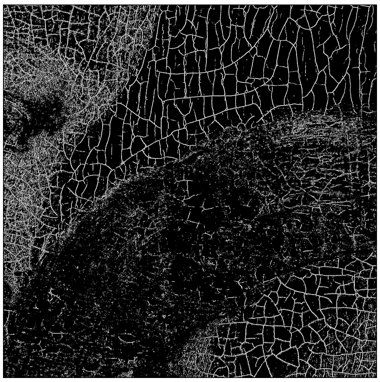

(f) Crack map produced after the 1st iteration of active learning

(g) Crack map produced after the 2nd iteration of active learning

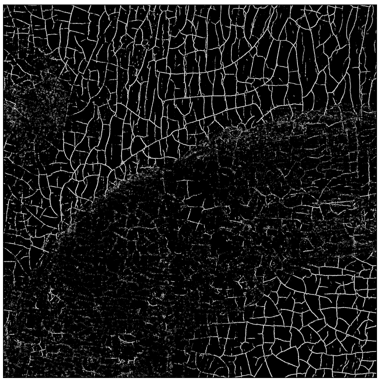

(h) Crack map produced after the 3rd iteration of active learning

Fig. 4: Results of crack detection using DAL4ART on a patch from the panel *Singing Angels* of the Ghent Altarpiece, publicly available on the website of the *Closer to Van Eyck* project. Image copyright: Ghent, Kathedrale Kerkfabriek, Lukasweb. For the annotations, we show in red the pixels annotated as cracks and in green the pixels annotated as non-cracks.

Acknowledgements

The authors thank Philippe Serbruyns for his valuable input and advice in developing the web application, and for enabling the web hosting of DAL4ART. They also thank Roman Sizyakin, Hélène Dubois, and Maximiliaan Martens for the useful discussions and for their expert feedback from an art history and conservation perspective.

The authors acknowledge the partial support of the Flanders AI Research Program (FAIR) under grant number 174B09119. NN acknowledges the support of the Belgian Federal Science Policy (BELSPO) through the FED-tWIN project Prf-2022-050_BALaTAI.

References

Abas, F., 2004. Analysis of craquelure patterns for content-based retrieval. Ph.D. thesis. University of Southampton.

Bucklow, S.L., 1997. A stylometric analysis of craquelure. Computers and the Humanities 31, 503–521.

Cornelis, B., Dooms, A., Cornelis, J., Leen, F., Schelkens, P., 2011. Digital painting analysis, at the cross section of engineering, mathematics and culture, in: 2011 19th European Signal Processing Conference, IEEE. pp. 1254–1258.

Cornelis, B., Ružić, T., Gezels, E., Dooms, A., Pižurica, A., Platiša, L., Cornelis, J., Martens, M., De Mey, M., Daubechies, I., 2013a. Crack detection and inpainting for virtual restoration of paintings: The case of the ghent altarpiece. Signal Processing 93, 605–619.

Cornelis, B., Yang, Y., Vogelstein, J.T., Dooms, A., Daubechies, I., Dunson, D., 2013b. Bayesian crack detection in ultra high resolution multimodal images of paintings, in: 2013 18th International Conference on Digital Signal Processing (DSP), IEEE. pp. 1–8.

Giakoumis, I., Nikolaidis, N., Pitas, I., 2005. Digital image processing techniques for the detection and removal of cracks in digitized paintings. IEEE Transactions on Image Processing 15, 178–188.

Gupta, A., Khandelwal, V., Gupta, A., Srivastava, M., 2008. Image processing methods for the restoration of digitized paintings. Science & Technology Asia , 66–72.

Huang, S., Cornelis, B., Devolder, B., Martens, M., Pizurica, A., 2020. Multimodal target detection by sparse coding: Application to paint loss detection in paintings. IEEE Transactions on Image Processing 29, 7681–7696.

Huang, S., Liao, W., Zhang, H., Pizurica, A., 2016. Paint loss detection in old paintings by sparse representation classification, in: Proceedings of the third" international Traveling Workshop on Interactions between Sparse models and Technology"(iTWIST'16), pp. 62–64.

Meeus, L., Huang, S., Zizakic, N., Xie, X., Devolder, B., Dubois, H., Martens, M., Pizurica, A., 2020. Assisting classical paintings restoration: efficient paint loss detection and descriptor-based inpainting using shared pretraining, in: Optics, Photonics and Digital Technologies for Imaging Applications VI, SPIE. pp. 99–110.

Pizurica, A., Platisa, L., Ruzic, T., Cornelis, B., Dooms, A., Martens, M., De Mey, M., Daubechies, I., 2014. Virtual restoration and mathematical analysis of pearls in the adoration of the mystic lamb, in: Het Lam Gods Series of Lectures.

Pizurica, A., Platisa, L., Ruzic, T., Cornelis, B., Dooms, A., Martens, M., Dubois, H., Devolder, B., De Mey, M., Daubechies, I., 2015. Digital image processing of the ghent altarpiece: Supporting the painting's study and conservation treatment. IEEE Signal Processing Magazine 32, 112–122.

Ren, P., Xiao, Y., Chang, X., Huang, P.Y., Li, Z., Gupta, B.B., Chen, X., Wang, X., 2021. A survey of deep active learning. ACM computing surveys (CSUR) 54, 1–40.

Ruzic, T., 2013. Context aware image inpainting with application to virtual restoration of old paintings. IEICE Proceedings Series 16.

Ruzic, T., Cornelis, B., Platisa, L., Dooms, A., Pizurica, A., Martens, M., De Mey, M., Daubechies, I., 2010. Craquelure inpainting in art work, in: Vision and Material-Interaction between Art and Science in Jan van Eyck's Time,(24-26 November 2010),(VLAC), Brussels (Belgium).

Settles, B., 2008. Curious machines: Active learning with structured instances. Ph.D. thesis. University of Wisconsin–Madison.

Shui, C., Zhou, F., Gagné, C., Wang, B., 2020. Deep active learning: Unified and principled method for query and training, in: International Conference on Artificial Intelligence and Statistics, PMLR. pp. 1308–1318.

Sizyakin, R., Cornelis, B., Meeus, L., Dubois, H., Martens, M., Voronin, V., Pižurica, A., 2020. Crack detection in paintings using convolutional neural networks. IEEE Access 8, 74535–74552.

Sizyakin, R., Voronin, V., Pižurica, A., 2022a. Virtual restoration of paintings based on deep learning, in: Fourteenth International Conference on Machine Vision (ICMV 2021), SPIE. pp. 422–432.

Sizyakin, R., Voronin, V., Zelensky, A., Pižurica, A., 2022b. Virtual restoration of paintings using adaptive adversarial neural network. Journal of Electronic Imaging 31.

Sober, B., Bucklow, S., Daly, N., Daubechies, I., Dragotti, P.L., Higgitt, C., Huang, J.J., Pižurica, A., Pu, W., Reynolds, S., et al., 2022. Revealing and reconstructing hidden or lost features in art investigation. IEEE BITS the Information Theory Magazine 2, 4–19. URL: https://ieeexplore.ieee.org/abstract/document/9900245.

Spagnolo, G.S., Somma, F., 2010. Virtual restoration of cracks in digitized image of paintings, in: Journal of Physics: Conference Series, IOP Publishing. p. 012059.