



Universiteit Antwerpen
| Faculteit Ontwerpwetenschappen

On Possible Roles of Machine Learning Applications for Painting Conservation: a Review of the State of the Art

Aster Van Vijle



University of Antwerp
Faculty Design Sciences
Course Heritage Studies
Academic year 2023-2024

2051POWERF: Masterscriptie
Prof. dr. Piraye Hacigüzeller and Prof. dr. Geert Van der Snickt
23/08/2024

I declare in good conscience that this master's thesis is an original product written solely by the author.
References were always correctly made to the literature from which information or ideas were derived.

Foreword

With this thesis I aim to give an overview of the state-of-the-art of machine learning applications that were created to facilitate the conservation of paintings. The hope is to bridge the gap between the two research domains of artificial intelligence and artwork conservation. This thesis should help conservators gain more knowledge about machine learning and serve as inspiration for future applications.

By writing this thesis I have finished the Master of Heritage Studies. The past two years were filled with gaining new knowledge both theoretical and practical about the heritage sector. I am very thankful for the opportunity to get a look behind the scenes of several heritage institutions. This made me realise that I wanted to be of aid in the conservation of heritage objects. Furthermore, I enjoyed the interdisciplinarity of the master. I gained interest in both digital heritage and conservation practices. This master's thesis was the perfect chance to unite both of those and it led to the subject under study "Possible Roles of Machine Learning Applications for Painting Conservation".

First and foremost, I would like to thank my promotors, Prof. Dr. Piraye Hacigüzeller and Prof. Dr. Geert Van der Snickt. They provided me with guidance each from their respective study domains, to bring this thesis to a good end. Additionally, I would like to thank my friend Lisa for providing feedback and proofreading my thesis and most importantly for supporting me when I was stressed.

Aster Van Vijle

Antwerp, 23 August 2024

Abstract

In this thesis an overview of state-of-the-art applications of machine learning for the conservation of paintings is given. Machine learning and especially deep learning have made a significant impact in various research domains, including artwork conservation. However, for conservators not familiar with the technical terminology of these techniques, it can be hard to understand the processes and see the opportunities. The other way around engineers working with machine learning and deep learning might not always understand the needs of conservators. With this thesis an attempt is made to bridge that communication gap. To ensure a good understanding of the discussed algorithms, a glossary and an extended glossary on machine learning are provided. Currently, the focus with researchers within this broad topic lies with four major subjects: enhancing the readability of technical images, pigment analysis, damage detection, and damage prediction. Various types of input data are employed, however the lack of data and especially data from historical sources remains an issue. Some researchers have found solutions by working with artificial and modern data and/or by employing algorithms or techniques which facilitate working with little data. Despite these solutions, some challenges remain and call for further research.

Keywords: machine learning, painting conservation, enhancing technical images, pigment analysis, damage detection, damage prediction

Table of Contents

Foreword	3
Abstract.....	4
Glossary	7
1. Introduction	15
2. Machine Learning: An Extended Glossary	16
2.1. What is Machine Learning?	16
2.2. What is Deep Learning?	16
2.2.1. Neural Networks	16
2.2.2. Deep Learning	17
2.3. Deep learning vs. Traditional Machine Learning	17
2.4. Types of Machine Learning	18
2.4.1. Supervised Learning.....	18
2.4.2. Unsupervised Learning	18
2.4.3. Semi-Supervised Learning.....	18
2.4.4. Reinforcement Learning.....	18
2.4.5. Self-Supervised Learning.....	19
2.5. Algorithms Mentioned in the Text	19
2.5.1. Machine Learning.....	19
2.5.2. Deep Learning	21
3. Enhancing the Readability of Technical Images	27
3.1. Machine Learning Applications	27
3.1.1. Visualising Underdrawings and Pentimenti	27
3.1.2. Separating Mixed X-ray Images	28
3.2. Discussion and Future Directions	29
4. Pigment Analysis.....	31
4.1. Machine Learning Applications.....	31
4.1.1. Pigment Identification.....	32
4.1.2. Pigment Unmixing	35
4.1.3. Pigment Layer Thickness	39
4.2. Discussion and Future Directions	41
5. Damage Detection	43
5.1. Machine Learning Applications.....	44
5.1.1. Craquelure Detection.....	44
5.1.2. Paint Loss Detection.....	46

5.1.3. General Damage Detection.....	47
5.1.4. Damage Detection with Infrared Thermography	48
5.2. Discussion and Future Directions	50
6. Damage Prediction	52
6.1. Machine Learning Applications	52
6.2. Discussion and Future Directions	54
7. Conclusion.....	56
8. Bibliography	58
9. Bibliography Illustrations	65

Glossary

Active learning: Active learning is a technique to improve the performance of supervised machine learning algorithms. By actively choosing the important data points in a dataset, this method builds a high-performance classifier with a small quantity of training data. The algorithm first selects a set of examples to be labelled from the pool of unlabelled data. A user can then be instructed by the learning algorithm to label this data with the correct results. The underlying theory of the active learning technique is that a machine learning algorithm might attain a greater accuracy with fewer training labels if it could choose which data it wanted to learn from.¹

Algorithm unrolling: Interpretability remains a challenge in deep neural networks, due to their black-box nature and the requirement for sizable training sets. Algorithm unrolling holds promise for resolving these difficulties. Initially, unrolling techniques were put out to create quick neural network approximations for sparse coding. However, the technique can also be applied to general iterative algorithms. Analytical procedures are repeatedly carried out by these algorithms. By mapping each iteration unique onto a single network layer and stacking a finite number of layers on top of each other, the algorithm can be unrolled into a deep network. The parameters in each iteration step, for example the model parameters, are represented as vectors. After unrolling, the parameters match those of the deep network, and by training the network end-to-end with real datasets, they can be optimised for real-world applications. Since unrolling encodes domain knowledge, unrolled networks typically have considerably less parameters than standard generic neural networks. Furthermore, their architectures are more precisely designed with specified applications in mind. Along with ensuring increased efficiency, these advantages also improve generalisability, particularly in the case of limited training data.²

Backbone network: A backbone network is commonly used in object detection model architectures. The backbone extracts and encodes features from the input data. It serves as the primary feature extractor, extracting both low- and high-level features from the input data.³

CIELab: Finding methods to define a colour's characteristics and calculate the numerical difference between shades is essential to consistently reproducing the exact desired colour. The $L^*a^*b^*$ colour model was therefore established by the Commission Internationale de l'Eclairage (CIE) with the aim of establishing a standard for colour communication. CIELab or CIE $L^*a^*b^*$, is a 3D colour space which allows precise measurement and comparison of all colours that may be perceived using just three colour values. The amount of variation that people perceive in colours is approximately reflected in this colour space by numerical disparities between values. The three values that the CIELab colour space uses to measure objective colour and compute colour differences are represented by the letters L^* , a^* , and b^* . On a scale of zero to one hundred, L^* denotes lightness from black to white, whereas a^* and b^* denote chromaticity without numerical bounds. Green is associated with negative a^* , red with positive a^* , blue with negative b^* , and yellow with positive b^* .⁴

¹ 'ML | Active Learning', GeeksforGeeks, 23 July 2019, <https://www.geeksforgeeks.org/ml-active-learning/>.

² Vishal Monga, Yuelong Li, and Yonina C. Eldar, 'Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing', *IEEE Signal Processing Magazine* 38, no. 2 (2021): 18–44, <https://doi.org/10.1109/MSP.2020.3016905>.

³ Megha Shroff, 'Know Your Neural Network Architecture More by Understanding These Terms', *Medium* (blog), 16 May 2023, <https://medium.com/@shroffmegha6695/know-your-neural-network-architecture-more-by-understanding-these-terms-67faf4ea0efb>.

⁴ 'What Is CIELAB Color Space?', accessed 5 August 2024, <https://www.hunterlab.com/blog/what-is-cielab-color-space/>.

Continuous learning: Traditional algorithms are normally trained on static datasets and their learning process is fixed as soon as the training is finished. However, they typically become less effective at tasks they have previously learnt when faced with new data. Catastrophic forgetting is the main problem that continuous learning attempts to solve. Continuous learning thus addresses the need for algorithms to integrate new data without fully retraining the model. This way older models can stay relevant as new data becomes available.⁵

Convolutional sparse coding: Convolutional sparse coding (CSC) is a feature extraction method in computer vision and machine learning. CSC can learn image features and use these for classification and reconstruction applications. Unlike patch-based techniques, convolutional sparse coding works on entire images, which allows it to capture the correlation between local neighbourhoods.⁶

Data augmentation: Massive volumes of data are required for deep learning. However, there are situations where gathering this much data is not practical, in which case data augmentation is helpful. Next to aiding in expanding the dataset's size, it also adds variability. When using images as input, common operations include rotation, shearing, zooming, cropping, flipping, and adjusting the brightness level.⁷

Ensemble Learning: By combining numerous models, ensemble learning enhances the results of machine learning. When compared to a single model, this approach produces greater predictive performance. The fundamental concept is to train a group of classifiers and give them the capability to vote, which then leads to a more accurate result.⁸

Features: A feature, also referred to as a variable or attribute, is a distinct measurable property or characteristic of a data point in a data set that serves as input for a machine learning algorithm. Features can be text-based, categorical, or numerical, and they describe various aspects of the data that are significant to the present challenge.⁹

Fibre Optic Reflectance Spectroscopy: The Fibre Optic Reflectance Spectroscopy (FORS) analytical method makes use of portable spectrophotometers that have been fitted with optical fibres accessories. By transporting light electromagnetic radiation from the instruments to the object under study and back again via optical fibres, in-situ measurements without the need for samples are made feasible. The flexibility of optical fibres allows for non-invasive measurement and easy access to any area of the object, regardless of the artwork's size or shape.¹⁰

Finetuning: Further training a pre-trained model with a fresh dataset is known as finetuning. When applying finetuning, the complete model including the initial layers is retrained. To avoid large changes,

⁵ 'Continual Learning in Machine Learning', GeeksforGeeks, 13 December 2023, <https://www.geeksforgeeks.org/continual-learning-in-machine-learning/>.

⁶ Felix Heide, Wolfgang Heidrich, and Gordon Wetzstein, 'Fast and Flexible Convolutional Sparse Coding', in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA: IEEE, 2015), 5135–43, <https://doi.org/10.1109/CVPR.2015.7299149>.

⁷ 'Python | Data Augmentation', GeeksforGeeks, 5 September 2019, <https://www.geeksforgeeks.org/python-data-augmentation/>.

⁸ 'Ensemble Classifier | Data Mining', GeeksforGeeks, 14 May 2019, <https://www.geeksforgeeks.org/ensemble-classifier-data-mining/>.

⁹ 'What Is Feature Engineering?', GeeksforGeeks, 20 March 2023, <https://www.geeksforgeeks.org/what-is-feature-engineering/>.

¹⁰ 'Fiber Optics Reflectance Spectroscopy | Popart', accessed 5 August 2024, <https://popart-highlights.mnhn.fr/identification/fiber-optics-reflectance-spectroscopy/index.html>.

the learning rate for the first few layers is set to a low number. In contrast, following layers use a higher learning rate in order to adjust to the updated dataset.¹¹

Fundamental parameters approach: Measured intensities are used in X-ray fluorescence analysis to assess a sample's composition. Either a computation using Fundamental Parameters (FP) or a direct comparison with calibration standards can be used to accomplish this. Calibration standards are always required because the spectrometer's illumination strength needs to be determined. However, as long as accurate intensity calculations can be performed, their composition does not need to closely resemble that of the unidentified sample. When analysing thin or multilayer materials, where it is challenging to get precise standards, these computations are quite useful. The absorption coefficients, absorption edges, characteristic wavelengths, fluorescence yields, relative emission rates, and absorption jump factors are the required parameters for these computations, which are also referred to as fundamental parameters.¹²

Hyperparameter: A model hyperparameter is a parameter whose value is established prior to the training of the model. A hyperparameter cannot be learned through fitting the model to the data. Examples of hyperparameters are the number of layers in a neural network or the number of neurons per layer in a neural network.¹³

Hyperspectral Imaging: By taking pictures at hundreds of consecutive narrow spectral bands, hyperspectral imaging (HIS), also known as Reflectance Imaging Spectroscopy (RIS), is able to provide an extensive spectral signature for every pixel in the picture. It outputs a three-dimensional data cube (hypercube), in which each pixel contains spatial and spectral information. Because of its great spectral resolution, materials may be precisely identified and categorised using their own spectral fingerprints.¹⁴

Infrared: Infrared (IR) radiation is an electromagnetic radiation type that is located slightly beyond the red portion of the visible light spectrum. The sun, lightbulbs, and even organisms all release infrared radiation, which has the ability to transmit heat. The wavelength range of infrared is around 700 nanometres (nm) to 1 millimetre (mm). This range is commonly separated into Near Infrared (NIR), Short-Wave Infrared (SWIR), Mid-Wave Infrared (MWIR), Long-Wave Infrared (LWIR), and Far-Infrared (FIR). The most used for conservation purposes are NIR, which extends from 700 to around 1400 nm and SWIR, which ranges from 1400 to 3000 nm.¹⁵

Infrared Photography: IR (Infrared Photography) is an imaging method that uses reflected infrared radiation with wavelengths between approximately 700 and 1100 nm. Reflected infrared photography was previously carried out on film formats, but it is now typically done using an adjusted commercial camera (e.g. DSLR) that has a silicon-based sensor. When an object is subjected to infrared radiation, an image known as a reflected infrared photograph is created that captures the object's response in the infrared region. IR can be used to separate pigments that behave similarly in the visible spectrum

¹¹ 'Transfer Learning with Fine-Tuning', GeeksforGeeks, 11 July 2023, <https://www.geeksforgeeks.org/transfer-learning-with-fine-tuning-in-nlp/>.

¹² D.K.G. De Boer, 'Fundamental Parameters for X-Ray Fluorescence Analysis', *Spectrochimica Acta Part B: Atomic Spectroscopy* 44, no. 11 (1989): 1171–90, [https://doi.org/10.1016/0584-8547\(89\)80114-4](https://doi.org/10.1016/0584-8547(89)80114-4).

¹³ 'Difference Between Model Parameters VS HyperParameters', GeeksforGeeks, 11 May 2020, <https://www.geeksforgeeks.org/difference-between-model-parameters-vs-hyperparameters/>.

¹⁴ 'What Is Hyperspectral Imaging? Where It Is Used?', GeeksforGeeks, 14 May 2024, <https://www.geeksforgeeks.org/what-is-hyperspectral-imaging-where-it-is-used/>; 'What Is Spectral Imaging in Image Processing?', GeeksforGeeks, 24 April 2024, <https://www.geeksforgeeks.org/what-is-spectral-imaging-in-image-processing/>.

¹⁵ 'What Is Infrared (IR)?', GeeksforGeeks, 30 April 2024, <https://www.geeksforgeeks.org/what-is-infrared/>.

(i.e., appear the same colour) from one another, given that they show distinct behaviours in the infrared region.¹⁶ If underdrawings are present and they are done in a, infrared-absorbing material and applied over a material that reflects infrared radiation so that the drawing lines appear darker than the underlying layer(s) of the painting, it can often be detected below overlaying paint layers using IR imaging.¹⁷

Infrared Reflectography: Layers beneath the surface can be explored and infrared-absorbing materials become visible using Infrared Reflectography (IRR). By illuminating a broadband infrared light source on a painting's surface and measuring the reflected light across roughly the wavelength range of 900–1700 nm, an infrared reflectogram can be generated. IRR can be used to detect underdrawings, just like IR.¹⁸

Keypoint detection: In computer vision, keypoint detection is the process of locating and detecting distinctive points or features in an image that are resistant to changes in viewpoint, scale, rotation, and illumination. By acting as anchor points, these keypoints help algorithms recognise objects, track motion, align images, and carry out a variety of other computer vision tasks. A keypoint detection head and a keypoint description head are involved in keypoint detection. The detection head is an algorithm that identifies keypoints in an image. The descriptor head contains feature vectors that describe the local appearance surrounding each keypoint.¹⁹

Kubelka-Munk theory: The Kubelka–Munk (KM) theory is used for modelling the optical properties of light-scattering materials. The theory illustrates the link between the absorption (K) and scattering (S) coefficients of incident light in materials. To implement the KM model, knowledge of the optical properties of the materials under study is needed.²⁰

Model parameter: A model parameter is a variable of the chosen model which may be determined by fitting the provided data to the model. An example of a model parameter is the weights in a neural network.²¹

Morphological filtering: To extract image elements that are helpful in representing and describing the shape of a region, morphological filtering can be applied. Morphological filtering includes a few basic operations that rely on the shape of the image. Usually, it is applied to binary images. The input image and the structuring component are the only data sources required. Morphological filters use set operators to combine these two data sources. The structuring component encodes shape characteristics, which determine how the objects in the input image are processed.²²

Multispectral imaging: Typically employing a small number of spectral bands, multispectral imaging entails taking pictures at distinct discrete wavelengths throughout the electromagnetic spectrum.

¹⁶ 'Reflected Infrared Digital Photography (IR)', MediaWiki, 28 February 2022, [https://www.conservation-wiki.com/wiki/Reflected_infrared_digital_photography_\(IR\)](https://www.conservation-wiki.com/wiki/Reflected_infrared_digital_photography_(IR)).

¹⁷ Barak Sober et al., 'Revealing and Reconstructing Hidden or Lost Features in Art Investigation', *IEEE BITS the Information Theory Magazine*, 2022, 1–16, <https://doi.org/10.1109/MBITS.2022.3207125>.

¹⁸ Sober et al., 'Revealing and Reconstructing Hidden or Lost Features in Art Investigation', 3.

¹⁹ 'What Is KeyPoint Detection?', GeeksforGeeks, 11 June 2024, <https://www.geeksforgeeks.org/what-is-keypoint-detection/>.

²⁰ Kleynhans et al., 'An Alternative Approach to Mapping Pigments in Paintings with Hyperspectral Reflectance Image Cubes Using Artificial Intelligence', *Heritage Science* 8, no. 1 (2020): 84, <https://doi.org/10.1186/s40494-020-00427-7>.

²¹ 'Difference Between Model Parameters VS HyperParameters'.

²² 'Python | Morphological Operations in Image Processing (Opening) | Set-1', GeeksforGeeks, 27 May 2019, <https://www.geeksforgeeks.org/python-morphological-operations-in-image-processing-opening-set-1/>.

Differentiating distinct materials and characteristics based on their spectral fingerprints is made possible by the fact that each band corresponds to a certain range of wavelengths.²³

Overfitting: When a statistical model fails to produce reliable predictions on test data, it is referred to as overfitted. An overabundance of data causes a model to start learning from noise and incorrect data entries in the training set. The model's sensitivity to changes in the training data means that testing with test data may lead to mistakes in the output. As a result of noise and excessive details, the model fails to classify the input accurately. Non-parametric non-linear techniques (i.e., no particular mathematical function can adequately capture the relationship between the dependent and independent variables) can easily cause overfitting since they give machine learning algorithms more freedom in constructing a model from the dataset. Too much freedom may lead to the creation of flawed models.²⁴

Principal Component Analysis: In order to decrease the dimensionality of data, Principal Component Analysis (PCA) which is a popular multivariate statistical analysis technique can be utilised. Primary characteristics of high-dimensional data are extracted by PCA and projected onto a lower-dimensional space via linear transformations. Feature extraction, noise reduction, and easy implementation are just a few of the benefits this technique offers.²⁵

Pulsed Phase Thermography: The method known as pulse phase thermography (PPT) combines the benefits of pulsed thermography (PT) and modulated thermography (MT). In MT, a single frequency is inserted into the specimen and examined in the stationary domain; in PT, however, all of the frequency responses are combined into a single transient signal. PPT makes use of the MT data analysis method and the PT experimental process.²⁶

RGB: The most widely used colour model in digital image processing is the RGB colour model. There are three channels in the colour image. and each channel has one colour. This model's primary colour components are red, green, and blue. The proportionate ratio of these three colours can make any other colour. Black is represented by 0, and the colour intensity increases as the value increases.²⁷

Semantic image segmentation: During the process of image segmentation, each object in an image is assigned a pixel-wise mask that indicates its presence. By using segmentation to identify the individual pixels that make up each object in the image, this technique can assist in determining the shape of each object contained in the image. In semantic image segmentation, every object belonging to the same class forms a single classification; as a result, every object in the same class receives the same colour.²⁸

Sparse representation: Sparse representation, also known as sparse approximation, is an approach to represent most or all of the data in a dataset by employing a weighted linear combination of a small

²³ 'What Is Spectral Imaging in Image Processing?'

²⁴ 'ML | Underfitting and Overfitting', GeeksforGeeks, 23 November 2017, <https://www.geeksforgeeks.org/underfitting-and-overfitting-in-machine-learning/>.

²⁵ Yi Liu et al., 'Generative Deep Learning-Based Thermographic Inspection of Artwork', *Sensors* 23, no. 14 (13 July 2023): 6362, <https://doi.org/10.3390/s23146362>.

²⁶ Sergio Marinetti et al., 'Pulse Phase Thermography for Defect Detection and Visualization', ed. Ajit K. Mal (Nondestructive Evaluation Techniques for Aging Infrastructures & Manufacturing, Newport Beach, CA, 1999), 230–38, <https://doi.org/10.1117/12.339890>.

²⁷ 'Difference Between RGB, CMYK, HSV, and YIQ Color Models', GeeksforGeeks, 16 November 2021, <https://www.geeksforgeeks.org/difference-between-rgb-cmyk-hsv-and-yiq-color-models/>.

²⁸ 'Object Detection vs Object Recognition vs Image Segmentation', GeeksforGeeks, 27 February 2020, <https://www.geeksforgeeks.org/object-detection-vs-object-recognition-vs-image-segmentation/>.

number of elements from a repetitive basis or dictionary. By reducing the amount of non-zero elements in the representation, this technique seeks to generate a concise representation of the input data.²⁹

Spatial median filtering: Spatial median filtering is a non-linear approach to minimise impulsive, or salt-and-pepper noise. The filter does this by replacing the pixel value with the median value of the neighbouring pixel. The median filter can also protect edge properties while lowering noise.³⁰

Spectral Angle Mapper: Spectral Angle Mapper (SAM) is one of the most often utilised spectral classification techniques. Mapping spectral similarity is done quickly and easily with this technique. When two spectra (test and reference) are compared, SAM calculates how similar their spectrums are. By considering the two spectra as N-dimensional vectors in space—where N is the number of spectral bands—arccosine angles between them can be calculated. The smaller the angle, the more decisive the match between the spectra.³¹

Spectral Correlation Mapper: By using Spectral Correlation Mapper (SCM), the Pearson correlation coefficient between two spectra can be determined. Pearson Correlation Coefficient (PCC) measures linear correlation between two sets of data. By centralising in the mean of the test and reference spectra, SCM standardises the data. It can be represented in angles by using arccosine. This approach excludes negative correlation.³²

Spectral imaging: Spectral imaging is the process of acquiring images at different wavelengths, usually ranging from Ultraviolet (UV) to NIR, but sometimes much farther. Each pixel in the image possesses a spectral signature that indicates its characteristics, whether they are reflecting or emitting, across the entire spectrum. By analysing these spectral signatures, professionals and experts can identify unique characteristics and patterns that are invisible to the human eye. Spectral imaging systems usually consist of a light source, sensors or optical filters, and data analysis software. Light from the light source illuminates the target object, and light emitted or reflected at different wavelengths is detected by filters or sensors. Software is used to process the collected data in order to generate spectral images and extract useful information. Spectral imaging is typically divided into two types, namely multispectral and hyperspectral.³³

Spectral Normalisation: Training traditional Generative Adversarial Networks (GANs) is often difficult because of the intricacy of the discriminator (D) and generator (G). To solve these problems, Spectral Normalised Generative Adversarial Network (SNGAN) uses spectral normalisation to stabilise the training of the discriminator. This means that the weight vectors of the GAN model are simply reparametrized. It also eliminates the need for intricate hyperparameter tuning and thereby increases simplicity. This technique improves the training stability of GANs and aids in producing higher-quality images.³⁴

²⁹ Yigit Oktar and Mehmet Turkan, 'A Review of Sparsity-Based Clustering Methods', *Signal Processing* 148 (2018): 20–30, <https://doi.org/10.1016/j.sigpro.2018.02.010>.

³⁰ 'Spatial Filters - Averaging Filter and Median Filter in Image Processing', GeeksforGeeks, 2 May 2020, <https://www.geeksforgeeks.org/spatial-filters-averaging-filter-and-median-filter-in-image-processing/>; 'Noise Removal Using Median Filter in C++', GeeksforGeeks, 17 June 2019, <https://www.geeksforgeeks.org/noise-removal-using-median-filter-in-c/>.

³¹ Dipendra J. Mandal et al., 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging', *Journal of Imaging Science and Technology* 67, no. 3 (1 May 2023): 030403-1-030403–18, <https://doi.org/10.2352/J.ImagingSci.Technol.2023.67.3.030403>.

³² Mandal et al., 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging,' 3.

³³ 'What Is Spectral Imaging in Image Processing?'

³⁴ Liu et al., 'Generative Deep Learning-Based Thermographic Inspection of Artwork,' 5.

Strain Energy Density: When a force is applied to a solid, it will deform and store the work done by the applied force as potential energy. This potential energy is a balance between the work delivered by external loads (such as body forces, surface tractions, and point loads) and the volume integral of the work delivered by internal stresses. The former is known as strain energy, whereas strain energy density (SED) is then the strain energy per unit volume.³⁵

Thermographic Signal Reconstruction: Thermographic Signal Reconstruction (TSR) is a data processing technique for thermographic images. To employ TSR, the best derivative images corresponding to each depth range must be chosen. This allows to assess the flaws' depths qualitatively by identifying them, or numerically using characteristic timings.³⁶

Transfer learning: In the process of transfer learning, the final layers of a pre-trained model are usually removed and replaced with new layers. To maintain the learnt representations, a low learning rate is used to finetune the pre-trained model's initial layers. These layers aid to capture the general features. After that, the new dataset is used to train the newly added layers.³⁷

Translation invariance: When something is translation-invariant, it signifies that the result is unaffected by changes made to the input data. Any translation of the input has no effect on the algorithms' response, which concentrates on characteristics that remain consistent during shifts.³⁸

X-radiography: By capturing variations in the way X-ray radiation is transmitted through a painting, X-radiography can be used to picture pentimenti and concealed compositions that have been altered or painted over. Furthermore, X-rays can shed light on structural elements like the composition and wood grain of panel paintings or the canvas weave and the positioning of battens, stretcher bars, and nails in canvas paintings. This is due to the fact that everything in the path of the incident X-ray radiation will contribute to the final image: the areas where the X-ray beam is most attenuated appear brightest in an X-ray image, and the degree of attenuation depends on the energy of the X-rays (which is also affected by the depth of a specific material within the structure of a painting) as well as the composition, density, and thickness of paint layers and other materials in the beam path.³⁹

X-ray: X-rays are a type of electromagnetic radiation which are found before the UV region of the electromagnetic spectrum and are extremely valuable in art conservation. X-ray radiation has a wavelength range of 1 to 10^{-3} nm and can be generated by blasting high-energy electrons at a metal target.⁴⁰

X-ray fluorescence: XRF scanning is a commonly used spectroscopic imaging technique in the cultural heritage field. The atoms in a painting are excited by X-ray radiation, which causes them to emit or fluoresce X-rays at particular energies that are typical of the elements found in both the surface and underlying layers. The XRF scanner creates a datacube by scanning the object's whole surface pixel by pixel and capturing the spectra of the distinctive X-rays emitted at each pixel. Using this method, one

³⁵ 'Strain Energy Density: What, Why and How?', *OnScale* (blog), 3 January 2022, <https://https%253A%252F%252Fonscale.com%252Fstrain-energy-density-what-why-and-how%252F>.

³⁶ Daniel L. Balageas et al., 'The Thermographic Signal Reconstruction Method: A Powerful Tool for the Enhancement of Transient Thermographic Images', *Biocybernetics and Biomedical Engineering* 35, no. 1 (2015): 1–9, <https://doi.org/10.1016/j.bbe.2014.07.002>.

³⁷ 'Transfer Learning with Fine-Tuning'.

³⁸ 'What Is the Difference between "Equivariant to Translation" and "Invariant to Translation"', *GeeksforGeeks*, 19 February 2024, <https://www.geeksforgeeks.org/what-is-the-difference-between-equivariant-to-translation-and-invariant-to-translation/>.

³⁹ Sober et al., 'Revealing and Reconstructing Hidden or Lost Features in Art Investigation' 3.

⁴⁰ 'Electromagnetic Spectrum | Wavelengths, Diagram and Uses', *GeeksforGeeks*, 20 December 2021, <https://www.geeksforgeeks.org/electromagnetic-spectrum/>.

can map and identify the distribution of various chemical elements at and below the surface of a painting, which can be linked to various pigments and materials.⁴¹

⁴¹ Sober et al., 'Revealing and Reconstructing Hidden or Lost Features in Art Investigation' 4.

1. Introduction

All artworks are susceptible to physical and visual degradation caused by a variety of natural or human factors. The risks that these works encounter, begins with the artist's technique and materials and is impacted by age and preservation conditions. Conservation measures attempt to prevent artworks from deteriorating due to age or environmental factors by keeping track of storage and exhibition conditions. Periodic imaging, as well as physicochemical documenting and research, are important conservation tasks.⁴² Most heritage institutions collect and save information on their art collections digitally. This data in combination with machine learning can serve as an essential source of information and as a non-invasive method for in-depth analyses of works of art.⁴³ Most conservators do not have sufficient knowledge about machine learning to implement this themselves, meaning that working with engineers is necessary. However, they might not always understand or see the challenges within conservation research. This thesis is meant as a starting point for both conservators and engineers who want to learn more about the applications of machine learning in art conservation.

Therefore, this thesis aims to give an answer to the question: what machine learning methods are currently being applied to facilitate the conservation of paintings? To answer this question, a review of the state-of-the-art is performed. Different from previous state-of-the-arts on this subject a comprehensive overview is given of all applications on paintings, instead of focussing on one topic. Additionally, it is considered that the thesis must be understandable and usable for both engineers and conservators. To make sure that all the technical terminology is not an obstacle to read this thesis, a glossary is provided.

This thesis is structured as follows: in chapter two a short overview of what machine learning is, will be given. In addition, the relevant algorithms are discussed to aid those who lack this knowledge in comprehending the mentioned methods. Related studies are then divided into larger topics, which each get discussed in a chapter. In chapter three two methods are discussed where machine learning is applied to enhance the readability of technical images. More specifically the enhancement of a Near-Infrared reflectogram and the separation of mixed X-rays is discussed. Chapter four deals with utilising machine learning to interpret spectral data in order to analyse pigments. The chapter is divided in methods that aim to identify single pigments, unmix pigment mixtures, and estimate the paint layer thickness. Chapter five focusses on using machine learning with the aim of automatically detecting damage in paintings. The identification of several types of damages, such as craquelure, paint loss and blisters are discussed in this chapter. In chapter six a study on the application of machine learning for predicting damage in paintings is discussed. Each chapter is structured as follows: First, there is a short explanation on the subject and why applying machine learning is an improvement compared to previous methods. Next, an overview is given of the state-of-the-art research and the methodologies that have been tested. Finally, a comprehensive overview of the general workflow is given and directions for future research are discussed.

⁴² Laurentiu Marian Angheluta and Alecsandru Chiroasca, "Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models," *Romanian Reports in Physics* 72, no. 805 (2020).

⁴³ Roman Sizyakin et al., "A Two-Stream Neural Network Architecture for the Detection and Analysis of Cracks in Panel Paintings," in *Optics, Photonics and Digital Technologies for Imaging Applications VI*, ed. Peter Schelkens and Tomasz Kozacki (Presented at the Optics, Photonics and Digital Technologies for Imaging Applications VI, Online Only, France: SPIE, 2020), 10, accessed November 23, 2023, <https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11353/2555857/A-two-stream-neural-network-architecture-for-the-detection-and/10.1117/12.2555857.full>.

2. Machine Learning: An Extended Glossary

2.1. What is Machine Learning?

Machine learning is a branch of computer science focused on teaching computers to learn from data without explicit programming (Fig. 1). The idea of self-learning is central to machine learning. Self-learning is the use of statistical modelling to find patterns and enhance performance based on data information, rather than following predefined commands. It is crucial to remember that although machines do not need direct commands to carry out certain tasks in machine learning, they do require input data. To put it briefly, the machine is given data, an algorithm is chosen, hyperparameters are




Fig. 1: Relationship between computer science, artificial intelligence, machine learning, and deep learning.

established and adapted, and then the algorithm starts analysing the data. Through a process of trial and error, the algorithm learns to recognise patterns within the data. The final model can be descriptive to learn from data to interpret underlying patterns or predictive, using learned patterns to predict future events, or it can be both.⁴⁴ Some traditional machine learning models include Support Vector Machine, Decision Tree, Random Forest, K-Means Clustering. Important to note is that deep neural networks, or more commonly called deep learning is regarded as a subset of machine learning (Fig. 1). Considering that throughout this paper deep learning methods are employed the most, a separate explanation is given in the next section.

2.2. What is Deep Learning?

2.2.1. Neural Networks

Artificial neural networks' (ANN) similarity to the workings of neural networks in the human brain led to the term "artificial neural networks". ANN's, similar to neurons in the human brain, are composed of interconnected neurons, or nodes, that engage in mutual interactions. These nodes are arranged in layers within the neural network. Raw data, such as text, audio, images, and numbers, serves as input for the nodes in the first layer. Then, along the network's edges, each node transmits data to the nodes in the next layer. A numerical weight is assigned to each edge and can be modified and formulated depending on experience. A neuron in the next layer will become active if the total sum of the connected edges meets the activation function, which is a predetermined threshold. It will not, however, activate if the total sum of the connected edges is less than the predetermined threshold. There are three layers in a neural network: input, hidden, and output. The input layer receives data first and uses it to discover broad features. The data is subsequently processed and analysed by the hidden layer(s). The data is simplified based on earlier calculations as each hidden layer is passed. The output layer displays the final result.⁴⁵

One of the drawbacks of neural networks is that, although they can approximate precise results, they function as "black boxes". This means that, although the network's structure can be straced, little to no information about the variables influencing the output is revealed.⁴⁶

⁴⁴ Oliver Theobald, *Machine Learning for Absolute Beginners: A Plain English Introduction*, Second edition. ([United States]: The author, 2017), 8-9.; Andriy Burkov, *The Hundred-Page Machine Learning Book* (Polen: Andriy Burkov, 2019), 3.

⁴⁵ Theobald, *Machine Learning for Absolute Beginners*, 78-81.

⁴⁶ Theobald, *Machine Learning for Absolute Beginners*, 78-81.

Neural networks are generally most effective at handling complex pattern recognition problems, particularly ones that are challenging for computers to solve yet intuitive for humans. For example, in the case of self-driving cars, it can be difficult for the system to decide whether a person is going to cross the road, while for a human this is an easy situation to judge.⁴⁷

2.2.2. Deep Learning

Basic neural networks or other alternative classification tools, such as logistic regression and k-nearest neighbours, are usually adequate for the examination of simple patterns. However, when the patterns in the data are more complex, a simple or shallow model can no longer be relied upon or used for analysis, especially when the data comes in the form of a high number of inputs, such as the total amount of pixels in pictures. This is because when the number of inputs increases, the model gets exponentially more complex. For neural networks, this implies adding more layers to handle more input nodes. Nonetheless, a neural network with a deep number of layers, also referred to as a deep neural network, is able to decompose complicated patterns into less complex ones. A shallow learning algorithm uses the features of the training data to directly learn the model's parameters. In contrast, in deep learning the algorithm derives the majority of its model parameters from the outputs of earlier layers.⁴⁸

2.3. Deep learning vs. Traditional Machine Learning

Deep Learning	Traditional Machine Learning
Applies ANN architectures to learn hidden patterns and relationships in a dataset.	Employs statistical algorithms to learn hidden patterns and relationships in a dataset.
Requires larger datasets compared to other machine learning algorithms.	Can work on smaller datasets.
Relevant features are automatically extracted from a dataset. It is an end-to-end learning process.	Relevant features are manually extracted from a dataset.
Better for complex tasks like image processing, natural language processing, etc.	Better for low-label tasks.
Takes more time to train a model.	Takes less time to train a model.
More complex because it works like a black box. Interpretations of the result are not always easy.	Less complex and the results are easy to interpret.

Fig. 2: Comparison of deep learning and traditional machine learning methods

⁴⁷ Theobald, *Machine Learning for Absolute Beginners*, 85-6.

⁴⁸ Barkov, *The Hundred-Page Machine Learning Book*, 26., Theobald, *Machine Learning for Absolute Beginners*, 85-6.

2.4. Types of Machine Learning

There are various types of learning in machine learning such as supervised, unsupervised, semi-supervised, reinforcement learning, and self-supervised learning. Each type has its benefits and shortcomings and excels in addressing different research questions.

2.4.1. Supervised Learning

Working with labelled datasets and drawing connections between input variables ("X") and known outcomes ("y"), are key components of supervised learning. This type of machine learning focuses on discovering patterns and making predictions based on these patterns. The way supervised learning operates is by providing the algorithm with sample data that includes different features ("X") and the data's corresponding correct value output ("y"). The dataset is considered "labelled" since both the input features and the output values are known. After identifying patterns in the data, the algorithm builds a model that can replicate the same patterns using new data. After the model is ready, it is tested for accuracy by applying it to new data. The model is prepared for application and real-world use once it has successfully completed both these training and test data phases. Examples of supervised learning algorithms include linear regression, logistic regression, Decision Tree, Random Forest, K-nearest neighbours, Support Vector Machine, Neural Networks (see section 2.5.).⁴⁹

2.4.2. Unsupervised Learning

Contrary to supervised learning, no data is labelled when it comes to unsupervised learning. Instead, by employing unsupervised learning techniques, the machine must find hidden patterns and generate labels. One of the benefits of unsupervised learning is finding patterns and relationships in the data that were previously unknown. Unsupervised learning thus offers the potential to extract new insights from unlabelled data that would not have been achievable in a different way. Some common unsupervised learning techniques are K-means clustering, autoencoder, etc. (see section 2.5.).⁵⁰

2.4.3. Semi-Supervised Learning

The dataset in semi-supervised learning includes both labelled and unlabelled data. Generally, there is a lot more unlabelled data than labelled data. The aim of a supervised learning algorithm and a semi-supervised learning algorithm are the same. However, by providing a large number of unlabelled examples along with a smaller set of labelled examples, the algorithm should be able to build a more superior model, that results in a more accurate output than a model only trained on labelled datasets. An example of an algorithm that can be employed for semi-supervised learning is a Generative Adversarial Network (see section 2.5.).⁵¹

2.4.4. Reinforcement Learning

In contrast to supervised and unsupervised learning, reinforcement learning uses input from earlier iterations to constantly improve its model. In reinforcement learning, algorithms are configured to train the model continuously through a trial-and-error process. Standard reinforcement learning models

⁴⁹ Theobald, *Machine Learning for Absolute Beginners*, 15-16.; Barkov, *The Hundred-Page Machine Learning Book*, 3.

⁵⁰ Theobald, *Machine Learning for Absolute Beginners*, 16-19.; Barkov, *The Hundred-Page Machine Learning Book*, 4.; 'Supervised and Unsupervised Learning', GeeksforGeeks, 1 October 2017, <https://www.geeksforgeeks.org/supervised-unsupervised-learning/>.

⁵¹ Barkov, *The Hundred-Page Machine Learning Book*, 4.

have quantifiable performance criteria and assign grade rather than tag outputs. An example of a reinforcement learning algorithm is Q-learning (see section 2.5.).⁵²

2.4.5. Self-Supervised Learning

Self-supervised learning is a deep learning approach in which unlabelled data is used to pre-train a model. The model then autonomously generates labels for the data, which are then utilised as ground truths in following rounds. The basic principle behind self-supervised learning is to interpret the unlabelled data that is fed to it in an unsupervised manner on the first iteration, in order to generate supervisory signals. Subsequently, the model is trained in following iterations using the high-confidence data labels among those created. Every iteration involves a change to the data labels that serve as the ground truth. An algorithm that can be employed for self-supervised learning is a Convolutional Neural Network (see section 2.5.).⁵³

2.5. Algorithms Mentioned in the Text

2.5.1. Machine Learning

Decision Tree: Decision Trees (DT) are tree structures that resemble flowcharts, with core nodes representing features, branches representing rules, and leaf nodes representing the algorithms output. It is a flexible supervised machine-learning approach that may be applied to regression and classification issues. Decision trees are powerful algorithms due to their simplicity and interpretability. Additionally, they are utilised by Random Forest (RF), one of the most powerful machine learning algorithms, which trains on various subsets of training data to create an ensemble of “trees” for better accuracy and robustness.⁵⁴

A DT was employed to identify pure pigments by Chen, Jesus, and Vilarigues (see section 4.1.2.).⁵⁵

eXtreme Gradient Boosting: Extreme Gradient Boosting (XGBoost) is an advanced machine learning method that has gained recognition for its remarkable prediction capabilities. In the context of ensemble learning, particularly with regard to gradient-boosting algorithms, it is considered the gold standard. To create a dependable and accurate predictive model, XGBoost gradually creates a number of “weak learners”, i.e. learning algorithms that, because they are unable to learn complex models, are usually quick in both training and prediction time. Essentially, it combines the predictions of multiple weak learners, e.g. decision trees, to create a strong predictive model. It employs a boosting strategy, whereby each weak learner corrects the errors of its predecessors, to produce an incredibly precise ensemble model.⁵⁶

The XGBoost algorithm was employed by Califano et al. to predict future damage in paintings (see section 6.1.).⁵⁷

K-Means Clustering: Taking an unlabelled dataset, K-Means Clustering aims to separate data points into a given number of clusters. Each group's data points are more similar to each other than they are

⁵² Theobald, *Machine Learning for Absolute Beginners*, 19-21.; Barkov, *The Hundred-Page Machine Learning Book*, 4.

⁵³ ‘Self-Supervised Learning (SSL)’, GeeksforGeeks, 27 January 2023, <https://www.geeksforgeeks.org/self-supervised-learning-ssl/>.

⁵⁴ ‘Decision Tree’, GeeksforGeeks, 16 October 2017, <https://www.geeksforgeeks.org/decision-tree/>.

⁵⁵ Ailin Chen, Rui Jesus, and Márcia Vilarigues, ‘Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms’, *SN Computer Science* 4, no. 2 (21 December 2022): 115, <https://doi.org/10.1007/s42979-022-01529-8>.

⁵⁶ ‘ML | XGBoost (EXtreme Gradient Boosting)’, GeeksforGeeks, 19 August 2019, <https://www.geeksforgeeks.org/ml-xgboost-extreme-gradient-boosting/>.

⁵⁷ America Califano et al., ‘Predicting Damage Evolution in Panel Paintings with Machine Learning’, *Procedia Structural Integrity* 41 (2022): 145–57, <https://doi.org/10.1016/j.prostr.2022.05.017>.

to those of the other groups. Depending on how far the data points are from the cluster centres, they are allocated to one of the K clusters. First, the centres of the clusters are randomly assigned to the available space. After the assignment of each point to a cluster, the points are then each assigned to another cluster. This procedure is performed repeatedly until every data point finds a good cluster.⁵⁸

K-Nearest Neighbours: The K-Nearest Neighbours (K-NN) algorithm is a supervised machine learning algorithm. It can handle both numerical and categorical data and does not require any assumptions about the underlying data distribution. This non-parametric technique bases its predictions on how similar the data points in a particular dataset are to one another. Using a distance measure, the K-NN algorithm locates the K neighbours who are closest to a particular data point. The majority vote or the average of the K neighbours are then used to establish the class or value of the data point. Using this approach enables the algorithm to adjust to various patterns and predict results based on the local structure of the data.⁵⁹

Linear regression: Linear regression is a supervised machine learning algorithm which predicts continuous output variables. The algorithm establishes the linear relationship between a dependent variable and one or more independent features. It does this by fitting a linear equation to the data.⁶⁰

Logistic regression: Logistic regression is a supervised machine learning algorithm that can be employed for classification tasks where the aim is to estimate the probability that something falls beneath a given class or not. Logistic regression is thus used for binary classification and where independent variables are taken as input and a probability value between 0 and 1 is generated as output.⁶¹

Random Forest: Random Forest (RF) algorithm is a supervised machine learning algorithm that works by generating a numerous Decision Trees during the training phase. Each tree is generated by taking an arbitrary subset of the data set to measure an arbitrary subset of features in each section. Due to each tree being more diverse as a result of the randomness, there is less chance of overfitting and the overall prediction performance is enhanced. When making predictions, the algorithm averages (for regression tasks) or votes (for classification tasks) the output of each tree. The findings of this collective decision-making process, which is aided by the insights of several trees, are reliable and accurate.⁶²

Support Vector Machine: Utilised for both regression and classification, Support Vector Machine (SVM) is a supervised machine learning technique. While SVM can also be used to handle regression problems, it works best for classification tasks. The SVM algorithm's primary goal is to locate the best hyperplane in an N-dimensional space that can be used to divide data points into different classes. This hyperplane is chosen to maximise the margin between the nearest points of the different classes, known as support vectors. The number of features determines the hyperplane's dimension. For example, the hyperplane is a line if there are two input features. The hyperplane transforms into a 2-D plane if there are three input features. For higher-dimensional spaces, the hyperplane becomes more complex, but the principle remains the same.⁶³

⁵⁸ 'K Means Clustering - Introduction', GeeksforGeeks, 2 May 2017, <https://www.geeksforgeeks.org/k-means-clustering-introduction/>.

⁵⁹ 'K-Nearest Neighbor(KNN) Algorithm', GeeksforGeeks, 14 April 2017, <https://www.geeksforgeeks.org/k-nearest-neighbours/>.

⁶⁰ 'Linear Regression in Machine Learning', GeeksforGeeks, 13 September 2018, <https://www.geeksforgeeks.org/ml-linear-regression/>.

⁶¹ 'Logistic Regression in Machine Learning', GeeksforGeeks, 9 May 2017, <https://www.geeksforgeeks.org/understanding-logistic-regression/>.

⁶² 'Random Forest Algorithm in Machine Learning', GeeksforGeeks, 22 February 2024, <https://www.geeksforgeeks.org/random-forest-algorithm-in-machine-learning/>.

⁶³ 'Support Vector Machine (SVM) Algorithm', GeeksforGeeks, 20 January 2021, <https://www.geeksforgeeks.org/support-vector-machine-algorithm/>.

Mandal et al. carried out two comparative investigations of machine learning models to identify pigments and study the effect of surface elevation on pigment identification. Among the machine learning models tested, was an SVM (see section 4.1.1.).⁶⁴ An SVM algorithm was also employed to identify pure pigments by Chen, Jesus, and Vilarigues (see section 4.1.2.).⁶⁵ Furthermore, Sidorov and Hardeberg used an SVM model to classify the crack maps outputted by a GNN into geographical origin classes (see section 5.1.1.).⁶⁶

Q-Learning: Q-learning is a popular machine learning model-free reinforcement learning algorithm. This kind of learning is known as temporal difference learning, where a system learns new information by seeing outcomes, engaging with the surroundings, and receiving feedback in the form of rewards. To train the model, different components work together iteratively in Q-learning models. The agent explores the environment as part of this iterative process, and the model is updated on a continuous basis.⁶⁷

2.5.2. Deep Learning

Autoencoder: A special type of encoder-decoder algorithms known as autoencoders are capable of learning effective representations of input data without the requirement for labels. As such, unsupervised learning is applied by this type of ANN. The fundamental idea behind an autoencoder is to learn how to efficiently represent and compress input data. Two components, an encoder and a decoder, make up the two-fold structure used to do this. The encoder, also known as "latent space" or "encoding," converts the input data into a reduced-dimensional representation. The decoder reconstructs the original input from that representation.⁶⁸

Employing an autoencoder, Fukumoto, Tsumura, and Berns presented an efficient and fast approach to estimate pigment concentrations (see section 4.1.2.).⁶⁹ Furthermore, Shitomi et al. developed a workflow employing an autoencoder for estimating the pigment mixing ratio and thickness (see section 4.1.3.).⁷⁰

Convolutional Autoencoder: A Convolutional Autoencoder (CAE) is a type of autoencoder that is applied often for unsupervised learning tasks, like image compression and denoising. Building further on the architecture of a traditional autoencoder, in CAE convolutional layers are added to both the encoder and decoder parts of the network. The encoder is made up of several layers that compress the input (an image or grid) by passing it through several convolutional layers. The decoder breaks down the compressed representation in an attempt to reconstruct the original image.⁷¹

⁶⁴ Mandal et al., 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging,' 2.; Dipendra J. Mandal et al., 'Comparison of Pigment Classification Algorithms on Non-Flat Surfaces Using Hyperspectral Imaging,' *Journal of Imaging Science and Technology* 67, no. 5 (1 September 2023): 1–25, <https://doi.org/10.2352/J.ImagingSci.Technol.2023.67.5.050405>.

⁶⁵ Chen, Jesus, and Vilarigues, 'Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms,' 8.

⁶⁶ Oleksii Sidorov and Jon Yngve Hardeberg, 'Craquelure as a Graph: Application of Image Processing and Graph Neural Networks to the Description of Fracture Patterns', 2019, <https://doi.org/10.48550/ARXIV.1905.05010>.

⁶⁷ 'Q-Learning in Python', GeeksforGeeks, 6 February 2019, <https://www.geeksforgeeks.org/q-learning-in-python/>.

⁶⁸ 'Autoencoders -Machine Learning', GeeksforGeeks, 21 June 2019, <https://www.geeksforgeeks.org/auto-encoders/>.

⁶⁹ Kensuke Fukumoto, Norimichi Tsumura, and Roy Berns, 'Estimating Concentrations of Pigments Using Encoder-Decoder Type of Neural Network.', *Color and Imaging Conference* 27, no. 1 (21 October 2019): 149–52, <https://doi.org/10.2352/issn.2169-2629.2019.27.28>.

⁷⁰ Ryuta Shitomi et al., 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments', *Journal of the Optical Society of America A* 40, no. 1 (1 January 2023): 116, <https://doi.org/10.1364/JOSAA.472775>.

⁷¹ 'Implement Convolutional Autoencoder in PyTorch with CUDA', GeeksforGeeks, 31 July 2023, <https://www.geeksforgeeks.org/implement-convolutional-autoencoder-in-pytorch-with-cuda/>.

Meeus et al. utilised a CAE as an unsupervised pre-training step to detect paint losses (see section 5.1.2.).⁷² Furthermore, Liu et al. utilised a CAE to enhance PCT's performance and detect damages in paintings (see section 5.1.4.).⁷³

Convolutional Neural Network: A typical deep learning neural network architecture commonly used in computer vision is the Convolutional Neural Network (CNN). The main use of a CNN is feature extraction from grid-like matrix datasets. For instance, visual datasets with a lot of data patterns, such as images or videos. Multiple layers make up a CNN, including the input layer, pooling layers, convolutional layers, and fully connected layers. The input image is processed by the convolutional layer to extract features, the pooling layer reduces computation and prevent overfitting by downsampling (i.e. reducing the spatial dimensions of) the image, and the fully connected layer generates the final output.⁷⁴

Karella, Blažek, and Striová processed NIR data using a CNN in order to improve the visualisation of underdrawings and pentimenti in a painting (see section 3.1.1.).⁷⁵ Furthermore, Pu et al. proposed a separation network for X-ray image separation utilising an CNN (see section 3.1.2.).⁷⁶ Chen, Jesus, and Vilarigues, and Jones et al. employed a CNN for automatic pigment identification (see section 4.1.1.) and Xu et al. used a CNN for pigment unmixing (see section 4.1.2.).⁷⁷ Finally, Nadisic et al. proposed an active learning framework employing a CNN to detect cracks (see section 5.1.1.).⁷⁸

Encoder-Decoder: Sequence-to-sequence learning is accomplished with neural network architectures such as encoder-decoders (ENDEC). A sequence, which is an ordered list of elements (e.g. words or characters in a text, frames in a video), is fed into the encoder, which converts it into a set of “context vectors”, an encoded summary representation, that the decoder can use. Taking this encoded representation, the decoder's task is to rebuild it into it a similar sequence (e.g. translating a sentence to another language), or a transformed version (e.g. summarising a text). This is not possible by speculation; instead, there must be some sort of relationship between what was encoded and what needs to be reconstructed, ensuring meaningful reconstruction.⁷⁹

Feedforward Neural Network: One of the most fundamental types of ANN's is the Feedforward Neural Network (FNN). The data or input supplied to this ANN only moves in one direction, namely forward.

⁷² Laurens Meeus et al., '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*, ed. Peter Schelkens and Tomasz Kozacki (Optics, Photonics and Digital Technologies for Imaging Applications VI, Online Only, France: SPIE, 2020), 16, <https://doi.org/10.1117/12.2556000>.

⁷³ Liu et al., 'Generative Deep Learning-Based Thermographic Inspection of Artwork', 1.

⁷⁴ 'Introduction to Convolution Neural Network', GeeksforGeeks, 21 August 2017, <https://www.geeksforgeeks.org/introduction-convolution-neural-network/>.

⁷⁵ Tomáš Karella, Jan Blažek, and Jana Striová, 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms', *Journal of Cultural Heritage* 58 (2022): 186–98, <https://doi.org/10.1016/j.culher.2022.09.022>.

⁷⁶ Wei Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs', *IEEE Transactions on Image Processing* 31 (2022): 4458–73, <https://doi.org/10.1109/TIP.2022.3185488>.

⁷⁷ Ailin Chen, Rui Jesus, and Marcia Vilarigues, 'Convolutional Neural Network-Based Pure Paint Pigment Identification Using Hyperspectral Images', in *ACM Multimedia Asia (MMAsia '21: ACM Multimedia Asia, Gold Coast Australia: ACM, 2021)*, 1–7, <https://doi.org/10.1145/3469877.3495641>; Bingjie Jenny Xu et al., 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?', *Journal of Analytical Atomic Spectrometry* 37, no. 12 (2022): 2672–82, <https://doi.org/10.1039/D2JA00246A>; Cerys Jones et al., 'Neural Network-Based Classification of X-Ray Fluorescence Spectra of Artists' Pigments: An Approach Leveraging a Synthetic Dataset Created Using the Fundamental Parameters Method', *Heritage Science* 10, no. 1 (2022): 88, <https://doi.org/10.1186/s40494-022-00716-3>.

⁷⁸ Nicolas Nadisic et al., 'A Deep Active Learning Framework for Crack Detection in Digital Images of Paintings', in *7th International Conference on Smart Monitoring, Assessment and Rehabilitation of Civil Structures, SMAR 2024 (7th International Conference on Smart Monitoring, Assessment and Rehabilitation of Civil Structures (SMAR 2024), Elsevier, 2024)*, 1–8, <http://hdl.handle.net/1854/LU-01J18DKNDKGSNKM5HVB9PYS7C>.

⁷⁹ Ahmadsabry, 'A Perfect Guide to Understand Encoder Decoders in Depth with Visuals', *Medium* (blog), 24 June 2023, <https://medium.com/@ahmadsabry678/a-perfect-guide-to-understand-encoder-decoders-in-depth-with-visuals-30805c23659b>.

The data enters the ANN through the input layer and exits the ANN through the output layer. Hidden layers can be, but are not necessarily present between the input and output layers.⁸⁰

This is different from how bidirectional neural networks operate e.g. a Recurrent Neural Network (RNN). An RNN takes the output of a layer and re-uses this output as new input in order to obtain a more accurate prediction of the layer.⁸¹

Pu et al. proposed a separation network for X-ray image separation utilising an FNN (see section 3.1.2.).⁸² Furthermore, to unmix pigments Rohani et al. suggested using a two-stream model where both the pigment identification and unmixing stages are FNN's that are trained simultaneously (see section 4.1.2.).⁸³ Pouyet et al. utilises the same methodology as Rohani et al. for pigment unmixing (see section 4.1.2.).⁸⁴

Fully Connected Neural Network: A Fully Connected Neural Network (FC-NN) consists of an input layer, one or more hidden layers, and an output layer. Input to each neuron comes from either the input layer or the neurons in the layer that precedes it. One neuron's output thus serves as the input for neurons in the network's subsequent layer, and so on, until the last layer generates the network's output. The neural network's layers apply several non-linear transformations to the input data, enabling the network to learn complicated representations of the original data.⁸⁵

Mandal et al. carried out a comparative investigation of machine learning models to identify pigments. Among the machine learning models tested, was a FC-NN (see section 4.1.1.).⁸⁶

Fully Convolutional Network: A CNN that has been adjusted for picture segmentation as opposed to classification is called a Fully Convolutional Network (FCN). FCNs produce a corresponding segmentation map with the same spatial resolution as the input image, in contrast to conventional CNNs that output a single class prediction. The absence of dense layers is the main distinction between CNNs and FCNs. Since FCNs do not need to make predictions, they do not require dense layers that might destroy the segmentation map's spatial resolution. Rather, according to their name, FCNs are constructed only with the help of upsampling (increasing the resolution as well as the size of an image) and convolutional layers.⁸⁷

Sindel, Maier, and Christlein developed a method for multi-modal image registration and artwork visualisation that is based on a deep learning craquelure detection method, employing an FCN (see section 5.1.1.).⁸⁸

Generative Adversarial Network: A Generative Adversarial Network (GAN) is a powerful kind of neural network for unsupervised learning. A GAN consist of two neural networks, namely a discriminator and

⁸⁰ 'Artificial Neural Networks and Its Applications', GeeksforGeeks, 24 June 2020, <https://www.geeksforgeeks.org/artificial-neural-networks-and-its-applications/>.

⁸¹ 'Difference Between Feed-Forward Neural Networks and Recurrent Neural Networks', GeeksforGeeks, 3 January 2023, <https://www.geeksforgeeks.org/difference-between-feed-forward-neural-networks-and-recurrent-neural-networks/>.

⁸² Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 6.

⁸³ Neda Rohani et al., 'Pigment Unmixing of Hyperspectral Images of Paintings Using Deep Neural Networks', in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Brighton, United Kingdom: IEEE, 2019), 3217–21, <https://doi.org/10.1109/ICASSP.2019.8682838>.

⁸⁴ Emeline Pouyet et al., 'Artificial Intelligence for Pigment Classification Task in the Short-Wave Infrared Range', *Sensors* 21, no. 18 (13 September 2021): 6150, <https://doi.org/10.3390/s21186150>.

⁸⁵ 'Introduction to Deep Learning', GeeksforGeeks, 1 June 2018, <https://www.geeksforgeeks.org/introduction-deep-learning/>.

⁸⁶ Mandal et al., 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging,' 2.

⁸⁷ Akanksha Chokshi, 'Supporting Fully Convolutional Networks (and U-Net) for Image Segmentation', accessed 5 June 2024, <https://www.datature.io/blog/supporting-fully-convolutional-networks-and-u-net-for-image-segmentation>.

⁸⁸ Aline Sindel, Andreas Maier, and Vincent Christlein, 'A Multi-Modal Registration and Visualization Software Tool for Artworks Using CraquelureNet', 2022, <https://doi.org/10.48550/ARXIV.2208.08836>.

a generator. The latter creates artificial data, and the discriminator evaluates them against real data samples to determine their authenticity. That is, the generator aims to produce realistic but synthetic data to fool the discriminator, and the discriminator strives to correctly identify real versus artificial data. As a result of this adversarial training process both networks improve over time, propel towards progress and deliver realistic, high-quality data samples.⁸⁹

Liu et al. utilised a Spectral Normalised GAN to enhance Principle Component Thermography's performance and detect damages in paintings (see section 5.1.4.).⁹⁰

Graph Neural Network: A Graph Neural Network (GNN) is used to manipulate data represented as graphs in different ways. Although traditional neural networks are already available, they are mainly capable of processing and interpreting textual and image data. Graphs (e.g. craquelure networks) are powerful data structures that store links between nodes (i.e. data points), and GNN's allow to investigate these connections in new and innovative ways.⁹¹

They enable tasks such as data point classification, link prediction, and graph classification. Sidorov and Hardeberg proposed a new method utilising a GNN to describe craquelure patterns and on the basis of this classify painting origins (see section 5.1.1.).⁹²

Multi-layer Perceptron: A Multi-layer Perceptron (MLP) is an ANN with three or more layers of “perceptrons”, which are basic units of performing weighted sums and activation functions to introduce non-linearity. These layers consist of one input layer, one or more hidden layers, and one perceptron output layer. From the input layers to the hidden layer or layers and then to the output layer, the data only moves forward, meaning MLP is a type of FNN.⁹³

To perform image segmentation and identify areas of an image with pure pigments MLP was used by Chen, Jesus, and Vilarigues (see section 4.1.1.).⁹⁴ Furthermore, Moradi et al. proposed a workflow utilising a MLP to categorise the pixels of a painting into defective and healthy regions on the basis of IRT (see section 5.1.4.).⁹⁵

One-Dimensional CNN: The purpose of a One-Dimensional CNN (1D-CNN) is to process One-Dimensional (1D) sequence data, such as time-series data or genomic sequences, which are one-dimensional in the sense that they consist of a sequence of values ordered in a single dimension. In a 1D convolutional layer, moving a convolutional filter over the input sequence is the primary task, which works like a window or magnifying glass sliding across a row of data points. During training, the network modifies a set of learnable weights that serve as this filter. These weights can be thought of as adjustable settings that the network fine-tunes to perform better. In each section of the input sequence, the filter moves along the data and the convolution process multiplies the filter values by the original input values in that section. This multiplication helps the network to highlight important patterns in the data. The results of these multiplications are then added together, resulting in a single

⁸⁹ ‘Generative Adversarial Network (GAN)’, GeeksforGeeks, 15 January 2019, <https://www.geeksforgeeks.org/generative-adversarial-network-gan/>.

⁹⁰ Liu et al., ‘Generative Deep Learning-Based Thermographic Inspection of Artwork,’ 1.

⁹¹ ‘What Are Graph Neural Networks?’, GeeksforGeeks, 2 July 2023, <https://www.geeksforgeeks.org/what-are-graph-neural-networks/>.

⁹² Sidorov and Hardeberg, ‘Craquelure as a Graph,’ 1.

⁹³ ‘Difference between Multilayer Perceptron and Linear Regression’, GeeksforGeeks, 4 March 2022, <https://www.geeksforgeeks.org/difference-between-multilayer-perceptron-and-linear-regression/>.

⁹⁴ Chen, Jesus, and Vilarigues, ‘Convolutional Neural Network-Based Pure Paint Pigment Identification Using Hyperspectral Images,’ 2.

⁹⁵ Morteza Moradi et al., ‘A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography’, *Sensors* 22, no. 23 (1 December 2022): 9361, <https://doi.org/10.3390/s22239361>.

output point for that section. The output is a transformed sequence, which is produced by repeating this operation across the whole sequence.⁹⁶

Mandal et al. carried out two comparative investigations of machine learning models to identify pigments and study the effect of surface elevation on pigment identification. Among the machine learning models tested, was a 1D-CNN (see section 4.1.1.).⁹⁷ Furthermore, Kleynhans et al. utilised a 1D-CNN to unmix pigments (See section 4.1.2.).⁹⁸

Residual Networks: It is a common problem in deep learning that when the number of layers in a neural network are increased, the training and test error rate also increases. In order to solve this problem Residual Networks (ResNet) were introduced. ResNets are made by stacking residual blocks which utilises skip connections. The skip connection connects one layer to a further layer by skipping some layers in between. By introducing this kind of skip connection, the architecture's performance can be preserved by skipping any layers which negatively impact it. As a result, a deep neural network is trained without overfitting.⁹⁹

Sindel, Maier, and Christlein developed a method for multi-modal image registration and artwork visualisation that is based on a deep learning craquelure detection method, employing a ResNet (see section 5.1.1.).¹⁰⁰

Self Organizing Map: A Self Organizing Map (SOM), also known as a Kohonen map, is a type of ANN that uses a competitive learning algorithm to train their network in an unsupervised manner. SOM is used to map multidimensional data onto lower-dimensional spaces using clustering and mapping (also known as dimensionality reduction) approaches, making complicated problems easier to comprehend. SOMs are shallow networks that are only made up of two layers, namely the input layer and the output layer.¹⁰¹

Kogou et al. utilized a SOM to effectively cluster an entire X-ray fluorescence (XRF) hypercube collected from a painting into separate pigment groups (see section 4.1.1.).¹⁰²

U-Net: A U-Net's main goal is to solve the problem of the minimal amount of annotated data. This network is made to perform efficiently with less data while being precise and fast. A U-Net's architecture is distinct because it has two paths: a contracting path and an expansive path. The contracting path consists of encoder layers that collect contextual information and lower the spatial resolution of the input. The expansive path is made up of decoder layers that decode the encoded data and utilize the information from the contracting path to produce a segmentation map.¹⁰³

⁹⁶ 'What Is a 1D Convolutional Layer in Deep Learning?', GeeksforGeeks, 19 February 2024, <https://www.geeksforgeeks.org/what-is-a-1d-convolutional-layer-in-deep-learning/>.

⁹⁷ Mandal et al., 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging,' 2.; Mandal et al., 'Comparison of Pigment Classification Algorithms on Non-Flat Surfaces Using Hyperspectral Imaging,' 3.

⁹⁸ Kleynhans et al., 'An Alternative Approach to Mapping Pigments in Paintings with Hyperspectral Reflectance Image Cubes Using Artificial Intelligence,' 3.

⁹⁹ 'Residual Networks (ResNet) - Deep Learning', GeeksforGeeks, 3 June 2020, <https://www.geeksforgeeks.org/residual-networks-resnet-deep-learning/>.

¹⁰⁰ Sindel, Maier, and Christlein, 'A Multi-Modal Registration and Visualization Software Tool for Artworks Using CraquelureNet,' 3.

¹⁰¹ 'Self Organizing Maps - Kohonen Maps', GeeksforGeeks, 1 July 2020, <https://www.geeksforgeeks.org/self-organising-maps-kohonen-maps/>.

¹⁰² Sotiria Kogou et al., 'A New Approach to the Interpretation of XRF Spectral Imaging Data Using Neural Networks', *X-Ray Spectrometry* 50, no. 4 (7 August 2021): 310–19, <https://doi.org/10.1002/xrs.3188>.

¹⁰³ 'U-Net Architecture Explained', GeeksforGeeks, 8 June 2023, <https://www.geeksforgeeks.org/u-net-architecture-explained/>.

Meeus et al. utilised an adapted U-Net to detect paint loss in paintings (see section 5.1.2.).¹⁰⁴ Furthermore, Moradi et al. proposed a workflow utilising a U-Net to categorise the pixels of a painting into defective and healthy regions on the basis of IRT (see section 5.1.4.).¹⁰⁵

VGG-16: The Visual Geometry Group (VGG) at the University of Oxford developed the convolutional neural network (CNN) architecture known as the VGG-16 model. With 16 layers total—13 convolutional and three fully connected—it differentiates itself by its depth.¹⁰⁶

Chirosca and Angheluta employed a VGG-16 to examine works of art that show various kinds of physical deteriorations in the painted layer (see section 5.1.3.).¹⁰⁷

¹⁰⁴ Meeus et al., 'Assisting Classical Paintings Restoration,' 2.

¹⁰⁵ Moradi et al., 'A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography,' 3.

¹⁰⁶ 'VGG-16 | CNN Model', GeeksforGeeks, 26 February 2020, <https://www.geeksforgeeks.org/vgg-16-cnn-model/>.

¹⁰⁷ Angheluta and Chirosca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 3.

3. Enhancing the Readability of Technical Images

Technical images can be used to visualize what is hidden under the top painting layer in an artwork and thus not visible with the naked eye. However, interpreting these images can be challenging due to the diversity of materials and the stratigraphy of an artwork. When it comes to non-invasively studying historical artworks, one of the most used methods is Near-Infrared (NIR) spectroscopy. NIR scanning, however, produces a mixture of the layers, such as the underdrawing, visible below the painting surface and the visible light (VIS) information of the painting surface itself. Separating these two layers and displaying the underlying layers as they appear without those features coming from the VIS layer is a difficult task.¹⁰⁸

X-ray imaging is another much used technique to obtain technical images. In X-ray imaging the material's physical thickness as well as its atomic number affect the attenuation of X-rays and, consequently, the brightness of the resulting region. The fact that X-ray images are two-dimensional (2D) representations of three-dimensional (3D) objects presents another difficulty. The radiograph will show elements in the front, rear, or even inside the painting. Thus, cradles or stretcher bars, for instance, as well as the support's structural elements (such as the wood grain, canvas weave, or fixings) will be visible. Each image (or phase of an image's development) is clearly layered or "blended" together if the support is painted on both sides, if the artist has modified the design, or if a support has been utilised more than once.¹⁰⁹

3.1. Machine Learning Applications

The non-linear transfer function from VIS to NIR for painting materials is hard to describe, since no precise relation exists to characterise it. However, machine learning algorithms can be utilized to estimate the relation representing this transition and unmix both signals.¹¹⁰

To separate X-ray images several methods including machine learning have already been explored. However, the separated X-ray images in previous research results still had elements belonging to the two different paintings recorded in the mixed x-ray image. Furthermore, these methods were only suited to separate X-rays belonging to artworks with paintings on both sides. This is because RGB images of both sides were required as input, however for paintings with concealed designs this is not possible.¹¹¹

3.1.1. Visualising Underdrawings and Pentimenti

Karella, Blažek, and Striová processed VIS-NIR data using a Convolutional Neural Network (CNN) that operates by suppressing VIS information found in NIR in order to improve the visualisation of underdrawings and pentimenti in a painting. Mock-ups were used as input data and as a basis to create virtual phantoms. These are artificial samples with changeable size, materials and underdrawing ratio that were created using known pigments and underdrawing coverage in accordance with historical painting methods. The phantoms were divided into training and validation sets for the


¹⁰⁸ Karella, Blažek, and Striová, 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms,' 186-87.

¹⁰⁹ Z. Sabetsarvestani et al., 'Artificial Intelligence for Art Investigation: Meeting the Challenge of Separating x-Ray Images of the *Ghent Altarpiece*', *Science Advances* 5, no. 8 (2 August 2019): eaaw7416, <https://doi.org/10.1126/sciadv.aaw7416>; Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 1-2.

¹¹⁰ Karella, Blažek, and Striová, 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms,' 187.

¹¹¹ Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 1-2.

CNN.¹¹² The trained CNN model estimates the non-linear transfer function and predicts the visible contribution (visible cover) in the NIR reflectogram. This visible cover is then subtracted from the NIR



reflectogram to obtain the enhanced reflectogram. The enhanced reflectogram shows only the NIR regions that are below the top painting layer, eliminating any remaining VIS information. This way it presents a more precise estimate of the painting's hidden features. To show the network's applicability in real-world scenarios, it was further applied to historical paintings (Fig. 3). The performed tests validate the technique as a method for NIR estimation both for virtual phantoms and historical paintings.¹¹³

Fig. 3: (a) RGB image of *Madonna dei Fusi* attributed to Leonardo Da Vinci; (b) Visible cover (the output of the neural network); (c) NIR reflectogram centered at 950 nm; (d) Enhanced reflectogram (the subtraction image c-b).

3.1.2. Separating Mixed X-ray Images

Pu et al. proposed a separation network for X-ray image separation which works for both double-sided paintings as well as for paintings with concealed designs. The proposed network is entirely self-supervised, with a training dataset consisting of sections of the mixed X-ray and an RGB image of the painting visible at the surface. In the case of a double-sided painting, an RGB image of only one side of the painting is used as input. Pu et al. utilised a convolutional sparse coding approach to model each individual image patch. The justification for adopting convolutional sparse coding is that it produces state-of-the-art performance while also allowing it to be unrolled into a CNN architecture. Furthermore, because the algorithm is based on CNN's it can deal with the RGB version of the visual image patches rather than the greyscale version of the same image patches. This means that it can capture colour information that may be useful in improving mixed X-ray image separation accuracy. Next, the mixed X-ray picture patch needs to be split into its constituent image patches. For this task a deep neural network structure with two components was used: an analysis and synthesis component. The analysis component unrolls the sparse coding algorithm into a Feedforward Neural Network (FNN) to generate sparse representations of the mixed X-ray image and RGB image of the surface painting. The synthesis component employs a linear CNN to generate a reconstruction of the mixed X-ray picture and RGB images of the surface painting.¹¹⁴ The researchers carried out several tests to evaluate the efficiency of their suggested X-ray separation method. These included evaluating the approach's efficiency in comparison to previous techniques. These experiments were performed with an artificially generated mixed X-ray image and



Fig. 4: Results of proposed approach on *Dona Isabel de Porcel* by de Goya: Top row from left to right: RGB image of the surface painting. Mixed X-ray image. Manually modified image of the surface painting. Bottom row from left to right: the reconstructed X-ray image of the surface painting, the reconstructed X-ray image of the concealed design, the synthetically mixed X-ray from the separated results.

¹¹² Karella, Blažek, and Striová, 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms,' 187-90.

¹¹³ Karella, Blažek, and Striová, 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms,' 192-93.

¹¹⁴ Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 4-5.

data from three historical paintings (Fig. 4). From the results of the experiments, it could be deduced that the proposed method outperforms the previous separation algorithms and addresses many of the problems associated with them.¹¹⁵

3.2. Discussion and Future Directions




Fig. 5: Summary of the general workflow for enhancing the readability of technical images.

¹¹⁵ Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 7-14.

In both techniques discussed above, the input data is derived from technical images, that are either within the VIS, X-ray, or the VIS-NIR range. The technical images were collected from historical data such as real paintings, or from modern data such as mock-ups. In addition, artificial data such as artificial mixed X-rays or virtual phantoms, derived from mock-ups were also used.¹¹⁶ In the future, focussing on enhancing underdrawings with less commonly used materials such as iron gall ink, earth pigments, vermilion, metalpoints or lake pigments could lead to interesting results. Furthermore, for enhancing underdrawings, experimenting with more complex training data, artificial (such as the virtual phantoms) as well as modern data (mock-ups), that considers the multi-layered paint stratigraphy in historical paintings, could deliver more accurate results when testing the model on historical data.

To enhance the NIR reflectogram, the virtual phantoms were manually labelled with the known pigments and underdrawing coverage. The algorithm was trained in a supervised manner.¹¹⁷ When trying to separate mixed X-ray images however, manually labelling the data to train a model is not possible. Therefore, a self-supervised approach was chosen, where labels are generated by the algorithm itself.¹¹⁸

To enhance the visibility of an underdrawing in a NIR reflectogram a CNN was used. In future research the benefits of more advanced neural network models such as ResNet to prevent overfitting can be explored.¹¹⁹ To separate a mixed X-ray image two algorithms, namely a FNN and a CNN were used in a sequential manner.¹²⁰

Since the method developed to separate the mixed X-rays treats the mixed X-rays as images it can also be applied to other images of paintings with mixed compositions, such as infrared reflectograms, XRF element distribution maps, or outputs from RIS. However, it is important to consider the increased complexity of light interaction with mixtures and layers of material in the resulting images from these techniques.¹²¹

¹¹⁶ Karella, Blažek, and Striová, 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms,' 187-90.; Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 4-5.

¹¹⁷ Karella, Blažek, and Striová, 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms,' 187-90.

¹¹⁸ Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 4-5.

¹¹⁹ Karella, Blažek, and Striová, 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms,' 193.

¹²⁰ Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 7-14

¹²¹ Pu et al., 'Mixed X-Ray Image Separation for Artworks With Concealed Designs,' 15.

4. Pigment Analysis

Knowing the pigments used by an artist is essential for dealing with a number of conservation-related issues, such as getting rid of old varnish and dirt without damaging the paint layer, stabilising paint by consolidating loose paint and removing consolidant residue from prior treatments, deciding on retouching materials that prevent metamerism and age similarly to the original paint layer, and determining appropriate storage conditions with regard to humidity, temperature, and light exposure. To safely preserve artworks, it is necessary to have a full understanding of the materials employed by the artist and materials used in later interventions, as well as their chemical structure and condition of preservation.¹²²

Many analytical and image analysis techniques have been developed for this purpose. Spectral imaging is one such approach. It is a quick, non-invasive, in situ method that allows to examine an artwork by measuring the light it absorbs and/or reflects at various wavelengths. This provides a wealth of information, such as chemical data related to the pigments used in an artwork.¹²³ Large datasets are then typically analysed using chemometric techniques. These are techniques based on the principles of statistics and mathematics. For example, in multivariate statistical analysis, every variable is considered simultaneously. This is particularly interesting to analyse spectra derived from image pixels.¹²⁴ However, due to the intricate material mixes and layered structures in paintings, the interpretation of measured spectral features requires expert knowledge and can be difficult as well as time-consuming.¹²⁵

4.1. Machine Learning Applications

Machine learning techniques can assist and reduce the manual work in interpreting spectral data to study the pigments in a painting. Currently, the primary goals of applying machine learning to the analysis of spectral and chemical data in cultural heritage are pigment identification and pigment unmixing. Pigment identification is a matter of classification that establishes whether a pigment is present or absent based on the given spectrum. Pigment unmixing is the quantitative decomposition of a given spectrum into its endmembers (EM) or base constituents. The EMs represent the pure materials utilised to create the mixtures, whereas the concentrations show how much of each EM is present in each pixel of the image. The process of unmixing is then split into two tasks: endmember identification and concentration estimation.¹²⁶ The unmixing process has its basis in statistical mixing models that describe the physical reactions that take place when various pigments are mixed. The mixing models fall into one of two categories: linear or non-linear.¹²⁷ If the spectra are linearly mixed, then the EM spectrum originates from a pixel that represents a part of the imaged area that is made up of multiple spatially separated materials. In this case, the EM spectrum can be fitted by an algorithm using a linear combination of the pure materials that are included in the pigment library. The measured spectrum is generally not a linear sum, though, if the materials are

¹²² Bartosz Grabowski et al., 'Automatic Pigment Identification from Hyperspectral Data,' *Journal of Cultural Heritage* 31 (2018): 1–12, <https://doi.org/10.1016/j.culher.2018.01.003>.

¹²³ Grabowski et al., 'Automatic Pigment Identification from Hyperspectral Data,' 1-2.; Chen, Jesus, and Vilarigues, 'Convolutional Neural Network-Based Pure Paint Pigment Identification Using Hyperspectral Images,' 1.

¹²⁴ Grabowski et al., 'Automatic Pigment Identification from Hyperspectral Data,' 1-2.

¹²⁵ Liu, Lingxi, Tsveta Miteva, Giovanni Delnevo, Silvia Mirri, Philippe Walter, Laurence De Viguierie, and Emeline Pouyet. 'Neural Networks for Hyperspectral Imaging of Historical Paintings: A Practical Review.' *Sensors* 23, no. 5 (February 22, 2023): 2419. <https://doi.org/10.3390/s23052419>.

¹²⁶ Liu et al., 'Neural Networks for Hyperspectral Imaging of Historical Paintings,' 5.

¹²⁷ Eva M. Valero et al., 'Unmixing and Pigment Identification Using Visible and Short-Wavelength Infrared: Reflectance vs Logarithm Reflectance Hyperspaces', *Journal of Cultural Heritage* 64 (2023): 290–300, <https://doi.org/10.1016/j.culher.2023.10.016>.

intimately mixed and light is reflected and/or absorbed by the other nearby materials (i.e., scattering) rather than simply reflecting off one material and into the hyperspectral camera. For this reason, a non-linear model will give better results. Since machine learning models are able to simulate non-linear functions, they have emerged in pigment analysis to automate the unmixing process.¹²⁸ Next to pigment identification and unmixing, estimating the pigment layer thickness has recently also been explored by researchers. Being able to estimate the pigment layer thickness in layered objects such as paintings, can aid in studying the stratigraphy of a painting and the possible deterioration of various layers in more detail.¹²⁹

Ultimately, applying machine learning for pigment analysis may reduce the number of destructive tests while raising the number of artworks that is examined, making the entire process more cost-effective.¹³⁰

4.1.1. Pigment Identification

An unsupervised deep learning technique to effectively cluster an entire X-ray Fluorescence (XRF) hypercube into separate chemical element groups was presented by Kogou et al. The XRF data was collected from a watercolour painting. The dataset is pre-treated by reducing the number of spectral bands and using spatial median filtering. This is required not only to minimise the amount of data and thereby increase processing efficiency, but also to elevate the signal-to-noise ratio and eliminate bands that have excessive noise or undesired signals to prevent alteration of the clustering results.¹³¹ The standard method for




Fig. 6: (a) Colour image of 'Peru type' #478 by Francisco Pancho Fierro; (b) Cluster map of the XRF hypercube with each cluster coded by a unique colour; (c) Corresponding single-cluster maps with each cluster presented individually for clarity.

interpreting XRF data is generally based on single-element maps, co-location maps, or spatial correlation maps between pairs of elements. However, because the clustering analysis in this research is carried out directly on the pre-treated XRF hypercube, there is no necessity for single-element maps. Procedures like spectral deconvolution (decomposing overlapping elemental peaks) are also not required until the interpretation step, where each mean cluster spectrum is analysed. The researchers utilised the Self Organizing Map (SOM) algorithm to automatically group pixels that have similar spectra. SOM works by reducing the total number of spectra that need to be examined for material identification to a small number of mean cluster spectra. Important to note is that the clustering method accounts for fluctuations in peak intensity, making it possible to distinguish between variations in the paint's density and the relative concentration of materials. The results show that an overview of the material changes throughout an object can be acquired by visualising the clusters' spatial distribution in the cluster map (Fig. 6b). Furthermore, the single cluster maps make it possible to

¹²⁸ Kleynhans et al., 'An Alternative Approach to Mapping Pigments in Paintings with Hyperspectral Reflectance Image Cubes Using Artificial Intelligence,' 2.

¹²⁹ Shitomi et al., 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments,' 116-28.

¹³⁰ Adam Polak et al., 'Hyperspectral Imaging Combined with Data Classification Techniques as an Aid for Artwork Authentication', Journal of Cultural Heritage 26 (2017): 1-11, <https://doi.org/10.1016/j.culher.2017.01.013>.

¹³¹ Kogou et al., 'A New Approach to the Interpretation of XRF Spectral Imaging Data Using Neural Networks,' 1-5.

analyse the mean spectra for each cluster and based on this give an indication of the pigments in each cluster (Fig. 6c).¹³²



Fig. 7: Results using the proposed model on a painting by Amadeo de Souza-Cardoso: RGB representation of a hyperspectral image (left); Predicted pure pigments (right).

To perform image segmentation and identify areas of an image with pure pigments, a deep learning technique was proposed by Chen, Jesus, and Vilarigues. The input data consisted of hyperspectral cubes related to images of 11 artworks and a database of spectral reference signatures (i.e. the reflectance of a material as a function of the wavelength) of 16 pigments. The reflectance data was collected within the VIS range. Additionally, the main chemical elements attributed to each pigment obtained through XRF were included in the input data. To train the model

a dataset of artificial hyperspectral samples was created using the pigment reference database. These samples included both pure pigment samples, generated with data augmentation, and mixed pigment samples. A hyperspectral training image was produced from these pure and mixed pigment samples.¹³³ A deep learning model consisting of three separate branches was employed by the researchers. The model's first branch consists of a convolutional neural network (CNN) to analyse the reflectance hypercube corresponding to a pixel in the hyperspectral image. This branch makes it possible to identify the feature maps that characterize the reflectance of pigments at various wavelengths. The spatial component of the hypercube analysis concentrates on identifying colour transitions. This leads to the eventual colour segmentation. The second branch of the model consists of a CNN that analyses the pixel's derivative of the reflectance hypercube. To eliminate high-frequency noise, a filter is utilised during the pre-processing stage which smooth the input hypercube data for this branch. The third branch examines the error present between the examined pixel's reflectance and every spectral signature present in the reference pigment database utilising the Spectral Correlation Mapper (SCM) value. The output of the three branches is combined into a linear array, which is then run through a Multi-layer Perceptron (MLP) to examine the detected features. To calculate the likelihood that the input data belongs to a pure pigment signature in the dictionary or a blend of these signatures (a pigment mixture), a multi-class output is utilised. By applying a threshold function to the model's output, which distinguishes between marginally pure predictions, the classification is further enhanced. When the probability falls below an established threshold, it indicates that the central pixel under analysis is not pure enough, and it is classified as a mixture.¹³⁴ The resulting image displays that the proposed model is able to effectively predict pure pigments. This was also verified via XRF analysis (Fig. 7).¹³⁵

To train a CNN for automatic classification of XRF spectra, Jones et al.'s research involved creating an extensive dataset of artificial XRF spectra of pigment often found in Renaissance paintings. By altering the concentrations of the characteristic chemical elements in each pigment present in the 15 chosen pigment classes, numerous synthetic XRF spectra (construed as single layers consisting of individual

¹³² Kogou et al., 'A New Approach to the Interpretation of XRF Spectral Imaging Data Using Neural Networks,' 5-9.

¹³³ Chen, Jesus, and Vilarigues, 'Convolutional Neural Network-Based Pure Paint Pigment Identification Using Hyperspectral Images,' 3-4.

¹³⁴ Chen, Jesus, and Vilarigues, 'Convolutional Neural Network-Based Pure Paint Pigment Identification Using Hyperspectral Images,' 2-3.

¹³⁵ Chen, Jesus, and Vilarigues, 'Convolutional Neural Network-Based Pure Paint Pigment Identification Using Hyperspectral Images,' 5-4.

pigments) were produced utilising the Fundamental Parameters (FP) approach. To generate artificial XRF spectra and accommodate for the fluctuating elemental compositions of historical pigments, consideration was given to the mean and standard deviation of the concentration of the chemical elements in the studied pigments.¹³⁶ Important to note is that the broad continuum (bremsstrahlung radiation), Rayleigh and Compton scatter peaks originating from the rhodium X-ray source, and the argon K lines attributed to atmospheric XRF measurements found in spectra from historical paintings were not included in the artificial data.¹³⁷ A CNN was trained with the data generated through FP. After testing the model for a first time with both artificial and real data, transfer learning was applied to finetune the model utilising a dataset of XRF spectra from historical paintings. The model was tested once again on real data. The results of these tests show that even though the accuracy of the model was significantly increased due to transfer learning, the model still failed when classifying spectra that consist of the same major chemical elements. Furthermore, it also struggled with spectra where minor chemical element peaks were required for differentiation. Finally, the model also failed when signals from pigments in the layers beneath the surface were present in the spectra (Fig. 8).¹³⁸




Fig. 8: Examples of XRF spectra of pigments correctly and misclassified using the CNN with transfer learning. The relative fluorescence intensities are plotted on a square-root scale with major characteristic X-ray lines labelled for each relevant element.




Fig. 9: Classification results: Euclidean Distance (ED), Spectral Angle Mapper (SAM), Spectral Correlation Mapper (SCM), Spectral Information Divergence (SID), Spectral Similarity Scale (SSS), Jeffries Matusita-Spectral Angle Mapper (JMSAM), Support Vector Machine (SVM), Fully Connected Neural Network (FC-NN) and One-dimensional Convolutional Neural Network (1D-CNN).

Mandal et al. carried out a study where they compared several popular non-machine learning methods and machine learning models to identify pigments. The data that they used for their research was collected in the Visible Near-Infrared (VNIR) region from a mock-up consisting of 10 pigments.¹³⁹ The machine learning models tested were Support Vector Machine (SVM), a Fully Connected Neural Network (FC-NN), and a one-dimensional CNN (1D-CNN). From the images showing the classification results, the researchers established that a FC-NN has the highest accuracy, followed by 1D-CNN and SVM (Fig. 9). All three algorithms outweigh the non-machine learning methods that were tested. This confirms the superiority of machine learning for an accurate identification of pigments.¹⁴⁰

In a subsequent study, Mandal et al. attempted to explore the impact of surface elevation and data augmentation on the accuracy of pure pigment classification. The researchers again performed a comparative study of non-machine learning methods and machine learning models. The data was collected withing the VNIR region from a 3D mock-up consisting of 10 pigments. The mock-up was

¹³⁶ Jones et al., 'Neural Network-Based Classification of X-Ray Fluorescence Spectra of Artists' Pigments,' 2-3.


¹³⁷ Jones et al., 'Neural Network-Based Classification of X-Ray Fluorescence Spectra of Artists' Pigments,' 7.

¹³⁸ Jones et al., 'Neural Network-Based Classification of X-Ray Fluorescence Spectra of Artists' Pigments,' 6-12.

¹³⁹ Mandal et al., 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging,' 1-7.

¹⁴⁰ Mandal et al., 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging,' 8-11.

made with various surface elevation levels (0 mm, 2.5 mm, 5 mm, and 10 mm).¹⁴¹ During this research two machine learning models were tested, namely SVM and 1D-CNN. The images demonstrate that the machine learning models deliver better classification results than the other tested methods (Fig. 10). The 1D-CNN outperforms the rest, followed by the SVM which struggles at the 10 mm elevation.




However, the researchers were able to improve the classification results of SVM with data augmentation. In general, it can be said that with increased elevation, the classification accuracy decreased. This could be due to the fact that elevation can cause shadows to appear, which in turn may influence the algorithms' classification accuracy. Furthermore, the researchers could conclude that data augmentation has a positive effect on the classification accuracy.¹⁴²

Fig. 10: Classification results: Euclidean Distance (ED), Spectral Angle Mapper (SAM), Spectral Correlation Mapper (SCM), Spectral Information Divergence (SID), Spectral Similarity Scale (SSS), Jeffries Matusita-Spectral Angle Mapper (JMSAM), Support Vector Machine (SVM) without data augmentation, and One-dimensional Convolutional Neural Network (1D-CNN).

4.1.2. Pigment Unmixing

Fig. 11: Semi-quantitative maps of pigment distribution by different networks. Each bar refers to a mixture, for which the dark blue, light blue and green show the average estimated coefficient values of the first, second, and third (if present in the mixtures) pigments in the mixtures. Yellow refers to the average estimated coefficient values of other pigments (not present in the mixture) that the model has predicted wrongly. (a) baseline network with the spectrum as the input, (b) a network with the spectrum and labels concatenated in the input layer (early fusion), (c) a network with the spectrum as the input to the first layer and labels concatenated to the output of the third layer (intermediate fusion), (d) a network with the spectrum as the input to the first layer and the output of the fifth layer multiplied by labels (late fusion).



An unmixing problem was addressed by Rohani et al. by employing a deep learning model to generate semi-quantitative maps of pigment distribution. To test the final model a hyperspectral image was made, consisting of reflectance spectra collected within the visible range (VIS) from a set of mock-up paintings. The mock-up paintings were composed of 11 pure colours and 41 mixtures of either two or three colours. To first train the identification and unmixing network a dataset with known concentration values is required. However, it is not always easy to obtain precise quantitative information for real and experimental datasets. Thus, to produce simulated data Kubelka-Munk (KM) theory was used. The simulated pigment mixtures were constructed using the 11 pure pigments found in the set of the mock-up paintings. Once more, combinations of two or three colours are taken into consideration, and every possible combination of two or three pure pigments was modelled. The KM theory was used to describe random non-linear mixtures, and for each mixture, 500 randomly chosen combinations of absorption(K)/scattering(S) coefficients and concentration values were chosen.¹⁴³ To solve the unmixing challenge the researchers suggested using a two-stream model where both the pigment identification and unmixing stages are deep Feedforward Neural Networks (FNN) that are trained simultaneously. Rohani et al. deliberately decided to work with a two-stream model where the first

¹⁴¹ Mandal et al., 'Comparison of Pigment Classification Algorithms on Non-Flat Surfaces Using Hyperspectral Imaging,' 1-6.

¹⁴² Mandal et al., 'Comparison of Pigment Classification Algorithms on Non-Flat Surfaces Using Hyperspectral Imaging,' 9-17.

¹⁴³ Rohani et al., 'Pigment Unmixing of Hyperspectral Images of Paintings Using Deep Neural Networks,' 3218-219.

stream identifies the pigments for any unknown spectrum. The second stream then utilises this information to deconstruct the provided spectrum into the concentration values of the pigments that had been identified. By first identifying the pigments a more precise result can be obtained. The pigment identification was modelled as a supervised multi-label classification, because the presence of various pigments in a measured spectrum allows for the possibility of numerous binary labels. The identified labels have two possible uses: they can be added as an additional input to the network, as "auxiliary input" (Fig. 11a) or they can be combined with other network layers (Figs. 11b (early fusion), 11c (intermediate fusion), and 11d (late fusion)). Experiments have confirmed that architecture d has the most accurate results (Fig. 11d). This is because the network performs better when the labels are included in the outputs of the network's higher layers.¹⁴⁴




Fig. 12: (a) Visible picture of a Buddhist painting (Tibetan thangka); (b) SAM classification of VIS-RIS dataset with analysed sites and corresponding reflectance spectra presented in (c, f); (d) DNN model results for minium, vermilion, and orpiment (in red, green and blue, respectively) with analysed sites and corresponding reflectance spectra presented in (c, f); (e) DNN model results for azurite and malachite (in blue and green, respectively) with analysed sites and corresponding reflectance spectra presented in (c, f).

Using the same methodology as Rohani et al., Pouyet et al. used Short-Wavelength Infrared (SWIR) hyperspectral datasets as input for the two FNN's to obtain semi-quantitative maps of pigment distribution. An expanded training dataset was created to increase the model's robustness. The researchers investigated the application of diverse input data, such as SWIR data from 8 different pure pigment pellets, from mock-ups, and from historical paintings. The mock-ups were made from the same 8 pigments as the pellets in a similar style as the historical painting and consisted of single layers and mixtures of up to three different pigments. Ensuring so much variability in the input data is an essential step in enhancing deep neural network (DNN) prediction.¹⁴⁵ Just like in the method proposed by

Rohani et al., the first FNN then identifies the pigments and the second FNN estimates the concentrations of the pigments in the mixtures.¹⁴⁶ Two methods were used to evaluate the acquired results from the two-stream network. To assess the model's performance, the accuracy of the model was first computed on a test dataset made up of spectral data from one mock-up and one historical painting. Subsequently, the results from the DNN were contrasted with those acquired using Spectral Angle Mapper (SAM), after which the detected pigments were labelled using supplementary VIS-RIS (Reflectance Imaging Spectroscopy) and XRF data. It was demonstrated that the DNN model performs better than SAM, when applied to both mock-ups and historical paintings (Fig. 12). Pouyet et al. could conclude that their multi-labelling approach and the multiple inputs from historical paintings and mock-ups used as a training dataset, enables the deep learning method to have a more representative mapping of the pigments.¹⁴⁷

Fukumoto, Tsumura, and Berns presented an efficient approach to estimate pigment concentrations by employing an autoencoder model. The input data consisted of the reflectance spectra of 19 different pigments measured within the VIS range. Next to that, utilising KM theory the corresponding K/S

¹⁴⁴ Rohani et al., 'Pigment Unmixing of Hyperspectral Images of Paintings Using Deep Neural Networks,' 3219-220.

¹⁴⁵ Pouyet et al., 'Artificial Intelligence for Pigment Classification Task in the Short-Wave Infrared Range,' 1-3.

¹⁴⁶ Pouyet et al., 'Artificial Intelligence for Pigment Classification Task in the Short-Wave Infrared Range,' 3-5.

¹⁴⁷ Pouyet et al., 'Artificial Intelligence for Pigment Classification Task in the Short-Wave Infrared Range,' 5-9.

coefficients were measured for each pigment. To train the autoencoder spectral data was obtained from mixtures that were made of the same four pigments out of the 19 different pigments.¹⁴⁸ The encoder will take the spectral data and encode it into a middle output. In this case, the middle output is represented by the pigment concentrations. The decoder then takes this middle output and decodes it back into spectral data. To test the model, spectral data was used from mixtures consisting of four different pigments which could be any of the 19 pigments. The resulting concentration estimations of the autoencoder were compared with the ground truth (Fig. 13). The comparison shows that the model can accurately unmix pigment mixtures and estimate the concentration of the pigments.¹⁴⁹




Fig. 13: The results of concentration estimation compared to the ground truth (GT) with error rates of best (left) and worst (right) accuracy.




Fig. 14: Comparison of pigment labelled maps: (a) Colour image of *Pentecost* by the Master of the Dominican Effigies; (b) Table of pigment labels for the truth map and the 1D-CNN map; (c) Truth pigment map (SAM); (d) 1D-CNN map.

mixtures containing white pigments (such as chalk, lead white, etc.). Most of the hue and intensity variability in pigments needed for reliable classification was thus present in this training set.¹⁵¹ Two methods were used to evaluate the final 1D-CNN. First, the mean-per-class accuracy was calculated; second, the accuracy of the model was confirmed by comparing its output with the resulting truth maps of three other classification methods (MLP, SVM, and SAM). The 1D-CNN was able to outperform all other methods (Fig. 14).¹⁵²

After constructing a DNN methodology for the identification of pure pigments, Chen, Jesus, and Vilarigues studied the identification and unmixing of pure and mixed pigments using data

¹⁴⁸ Fukumoto, Tsumura, and Berns, 'Estimating Concentrations of Pigments Using Encoder-Decoder Type of Neural Network,' 150.

¹⁴⁹ Fukumoto, Tsumura, and Berns, 'Estimating Concentrations of Pigments Using Encoder-Decoder Type of Neural Network,' 150-52.

¹⁵⁰ Kleynhans et al., 'An Alternative Approach to Mapping Pigments in Paintings with Hyperspectral Reflectance Image Cubes Using Artificial Intelligence,' 1-4.

¹⁵¹ Kleynhans et al., 'An Alternative Approach to Mapping Pigments in Paintings with Hyperspectral Reflectance Image Cubes Using Artificial Intelligence,' 4-6.

¹⁵² Kleynhans et al., 'An Alternative Approach to Mapping Pigments in Paintings with Hyperspectral Reflectance Image Cubes Using Artificial Intelligence,' 6-11.

augmentation and a hierarchical approach of machine learning techniques. This hierarchical technique restricts the classification of pixels or areas at levels with similar categories to reduce the complexity of the training and testing procedure. For each pigment that was taken from a real artist paint tube, the database the researchers used contained the following information: the name of the pigment, its RGB and CIELab colour components, a list of the most important chemical elements for each pigment that was obtained using



Fig. 15: Painting Amadeo de Souza-Cardoso (left); Reconstructed image from the resulting concentration percentages of the DNN (right).

XRF, and reflectance spectra that were obtained using both a hyperspectral camera and Fibre Optic Reflectance Spectroscopy (FORS) for comparison purposes. In addition, they employed spectral data from the VIS range and XRF-data from 12 paintings. The RGB and CIELab colour components of the assessed points were also provided for reference. To get over the problem of data limitation and expand the number of samples, the researchers employed data augmentation. This minimises overfitting of the assessed model while producing enough training samples within the reference signal's standard deviation.¹⁵³ The researchers used a classification tree consisting of four phases to identify and/or unmix pigments. First, a shallow neural network was utilised to identify black and white pigments. The data used consisted of reflectance spectra of black, white, and coloured pigments as three separate groups.¹⁵⁴ Next, Decision Trees (DT) and SVM, were compared to identify pure pigments. The DT was trained and tested with reflectance spectra of pixels of a single image for classification of reflectances only. The SVM was modelled to identify any of the 17 pure pigments (including black and white) from the database. Testing was done with data from a historical painting.¹⁵⁵ Lastly DNNs were used to estimate pigment concentrations. The training data for the DNN consisted of artificial samples created with data augmentation applied to mixtures consisting of 3 pigments of any of the 17 pigments from the database.¹⁵⁶ Reflectance spectra from a historical painting were used to test the model and the resulting concentration percentages were then employed to reconstruct the painting (Fig. 15). The reconstructed image clearly shows a high correlation with the original painting.¹⁵⁷

A framework for pigment identification and unmixing based on XRF data was presented by Xu et al. However, compared to previous methods, they specifically focused on mixed pigments and the intricate layered structure of paintings. They focused on this layered structure because one of the paintings they studied was a special case with an underlying painting. To create, train, and test their framework, they collected XRF data from three separate sources: XRF spectra from historical paintings, XRF spectra from oil-paint mock-ups, and simulated XRF spectra of multi-layered pigments. The XRF spectra from the historical paintings had already been compared during previous research with the ground truth results from RIS and cross-section analysis. The mock-ups were simplified to a three-layer structure using 11

¹⁵³ Chen, Jesus, and Vilarigues, 'Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms,' 1-4.

¹⁵⁴ Chen, Jesus, and Vilarigues, 'Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms,' 6-7.

¹⁵⁵ Chen, Jesus, and Vilarigues, 'Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms,' 8.

¹⁵⁶ Chen, Jesus, and Vilarigues, 'Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms,' 8-10.

¹⁵⁷ Chen, Jesus, and Vilarigues, 'Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms,' 10-13.





Fig. 16: Pigment identification results of *The Bathers* by Paul Cézanne from the finetuned model, including cobalt blue, vermilion, and emerald green from top to bottom. The first, second and third columns display the pigment map, element map(s) and the comparison figure and scatter plots, respectively.

distinct pigments which together make up 64 different layer structures. 16224 simulated XRF spectra also consisting of a three-layer structure were generated with the FP approach. With this simulation dataset, the researchers were able to minimise the need for mock-ups and still have a sufficient amount of data.¹⁵⁸ The simulation dataset was used to train a CNN, which was subsequently finetuned using the experimental dataset consisting of spectra from the mock-ups, to produce the final pigment identification model. The finetuning process allows to extract characteristics from

specific pigment mixtures in a painter's palette which aids in the identification of pigments, particularly in artworks that have similar painting styles.¹⁵⁹ The trained model was then tested on XRF data obtained from two historical paintings. The pigment maps, which are the visualization of the probability output of the model, and matching element maps were compared. The results demonstrated that the model is capable of accurately identifying pigments, even when the pigments are present in low concentrations or covered by overpaint (Fig. 16).¹⁶⁰ Nevertheless, there are still several issues with the model. Pigments with comparable elemental profiles are sometimes difficult for the model to differentiate. In addition, when highly absorbing pigments obstruct the radiation released from hidden paint layers, the model proves inefficient. Finally, to avoid model overfitting and inaccurate predictions due to limited pigment mixtures in the mock-ups, the refined model requires experimental data containing a minimum of a few ground truth measurements.¹⁶¹

4.1.3. Pigment Layer Thickness

Fig. 17: (a) RGB image converted from an original measured spectral image; (b) RGB image converted from a reconstructed spectral image; (c) Mean squared errors in (a) and (b) spectra; (d) Estimation results of pigment thickness.



Shitomi et al. modelled an unsupervised physics-based autoencoder for estimating the mixing ratio and thickness of pigments in layered surface objects such as mural paintings and watercolours. The reason the researchers went with an unsupervised framework is that it doesn't need a big dataset, which makes it possible to estimate measured data without the need for annotated thickness data. They trained and tested the model with both artificial spectral data based on KM theory and real data from mock-ups obtained within the VIS range.¹⁶² In essence solving the inverse problem of KM theory is the aim of the physics-based autoencoder. First the spectrum of an arbitrary pixel is processed by the encoder to estimate the thickness and mixing ratio of the pigments present in that pixel. Subsequently, the decoder, which is a

¹⁵⁸ Xu et al., 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?' 1-4.

¹⁵⁹ Xu et al., 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?' 4-5.

¹⁶⁰ Xu et al., 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?' 5-7.

¹⁶¹ Xu et al., 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?' 8-10.

¹⁶² Shitomi et al., 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments,' 119-126.



Fig. 18: Estimated mixing ratios including pigments other than those used for the mock-up.

model grounded on KM theory, reconstructs these latent values received from the encoder. The latent variables in the middle layer, in this case the mixing ratio and thickness of pigments, are interpretable since the decoder is based on a physics-based model (Fig. 17 and Fig. 18). Lastly, the error between the input and output spectra is used to optimise the encoder.¹⁶³ The resulting images show that the model can correctly estimate the mixing ratio and thickness of the pigments most of the time.

However, the model still has some shortcomings: the model has a hard time estimating the thickness of

the pigment layer if it is too large; errors in determining mixing ratios arise when various pigment mixtures have similar spectra. Regardless of these challenges, the researchers could confirm that the results of their unsupervised approach was superior to those of a supervised model.¹⁶⁴

¹⁶³ Shitomi et al., 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments,' 116-19.

¹⁶⁴ Shitomi et al., 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments,' 127-28.

4.2. Discussion and Future Directions




Fig. 19: Summary of the general workflow for pigment analysis.

In all the described methods the training and testing data consisted of spectral data collected with one or more of the following techniques: FORS, SWIR, RIS, VIS, VNIR, XRF. The spectral data was collected from numerous different sources: artificial data generated with data augmentation or with physical models such as KM or FP; modern data such as mock-ups, pigment pellets, single pigments; or historical data such as real paintings or original paint tubes. Several studies only used historical data to test the model and not to train it. However, artificial and modern data are oftentimes less complex than data from real paintings. Additionally, these types of data lack certain characteristics that spectral data collected from historical sources does have. First, weathering, ageing, and fading influence the reflectance of pigments. Ageing can be modelled to some extent; however multiple variables are involved. The change in the reflectance of a pigment depends on composition, light exposure, and environmental factors. In many cases these specific factors are unknown.¹⁶⁵ Second, the material composition of the whole painting and especially the variety in pigments and mixtures should be considered. The base (e.g. canvas, wood, ...), preparation layer, sequence of paint layers (e.g. highly absorbing pigments obstruct the radiation coming from the concealed layer), binder, thickness and depth of pigments (e.g. thin painting layers like glazing), pigments consisting of minor chemical components, pigments with comparable elemental features, rare pigments, the mixing ratios of pigments, and varnish are all parameters that will complicate the correct identification and unmixing of pigments.¹⁶⁶

The training data consisted of both manually labelled data, such as characterised paintings and manual identification of spectral features for supervised methods and unlabelled training data, such as chemical elements, pigments, concentrations and pigment layer thickness for unsupervised methods. As mentioned before, some variables that influence the identification and unmixing of pigments are hard to determine. Unsupervised learning methods can offer a solution due to their capability to extract latent features of signals without any labels. When considering unsupervised learning increasing the amount of unannotated data as input will only enhance a model's performance. Therefore, in future research the use of unsupervised techniques should be further explored.¹⁶⁷

To identify pigments the following algorithms were used: 1D-CNN, CNN, DT, FC-NN, MLP, SOM and SVM. For the pigment unmixing task a 1D-CNN, autoencoder, CNN, and FNN were employed. To estimate the thickness of pigments an autoencoder was used. Some researchers utilised multiple algorithms in a sequential manner to perform the challenge at hand.

To improve the suggested algorithm's performance, finetuning and transfer learning were utilised by some researchers. These techniques operate by adding unseen data to an already trained model to expand the model's usability.

¹⁶⁵ Valero et al., 'Unmixing and Pigment Identification Using Visible and Short-Wavelength Infrared,' 298-99.

¹⁶⁶ Mandal et al., 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging,' 11.; Polak et al., 'Hyperspectral Imaging Combined with Data Classification Techniques as an Aid for Artwork Authentication,' 10.; Valero et al., 'Unmixing and Pigment Identification Using Visible and Short-Wavelength Infrared,' 298-99.; Xu et al., 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?,' 10.; Jones et al., 'Neural Network-Based Classification of X-Ray Fluorescence Spectra of Artists' Pigments,' 12.

¹⁶⁷ Xu et al., 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?,' 10.; Kogou et al., 'A New Approach to the Interpretation of XRF Spectral Imaging Data Using Neural Networks,' 9-10.; Shitomi et al., 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments,' 127-28.

5. Damage Detection

Polychrome surfaces in artworks are susceptible to both physical and optical degradation as a result of numerous natural or human-caused factors. Problems with the painted layers of the artworks can be attributed to the materials employed by an artist, as well as to the technique for blending pigments and treating the support. In addition, the history of conservation affects the types of damages an artwork suffers. To preserve artworks for the future it is necessary to document the various types of damages an artwork may have.¹⁶⁸

One of the most frequent deteriorations seen in paintings is cracking, also known as craquelure. In general, pressures that build up within or on paint layers due to a variety of reasons typically result in the material breaking, leading to cracks in paint layers. Age cracks are usually caused by the substrate expanding or contracting and can be found in every layer of a painting. These changes in the substrate are often due to fluctuations in temperature and relative humidity. It's important to differentiate premature cracks from age cracks. The former originate in only one layer of paint and have a duller edge than those due to ageing. They typically indicate poor technical execution during the painting process, such as applying a layer that dries faster than the underlying one or not allowing enough time for a layer to dry.¹⁶⁹ The entire paint layer structure, including the preparation and paint layers on top of it, might also be subjected to mechanical cracks. This type of cracking may be generated by transport vibrations, framing stress or other types of physical impact.¹⁷⁰ Varnish cracks are another type of cracking. They appear exclusively in the varnish layer when oxidation causes it to become brittle. Cracks can be visually classified as either dark cracks on a bright background or bright cracks on a dark background. Additionally, they differ in thickness from fine hair-like lines to very thick lines extending into areas of missing paint. They create a pattern that might be entirely random or in the shape of trees, webs, rectangles, circles, or unidirectional.¹⁷¹

Another form of damage is paint loss. They typically result from cracks combined with ground layer separations. The regions with paint loss, also known as lacunae, may result from physical impacts such as abrasions, tears, etc., or from the combination of other forms of damages.¹⁷²

Blisters, often called bubbling, are also a typical form of degradation. They are the result of a paint film's localised loss of adhesion. This film grows increasingly brittle as it ages, and its lack of adhesion turns into flaking. This effect can be seen between a painted layer and the support or between two painted layers. The moisture found under the paint layer, which primarily results from new woodwork that was not adequately protected before painting, is the most frequent source of this kind of damage on wood supports.¹⁷³

Damages may be detected by visually assessing an artwork in real life, or through assessing images under Visible Light (VIS), Infrared Photography (IR), Infrared Reflectography (IRR), X-radiography or

¹⁶⁸ Angheluta and Chiroasca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 1.

¹⁶⁹ Aleksandra Pizurica et al., 'Digital Image Processing of The Ghent Altarpiece: Supporting the Painting's Study and Conservation Treatment', *IEEE Signal Processing Magazine* 32, no. 4 (2015): 112–22, <https://doi.org/10.1109/MSP.2015.2411753>; B. Cornelis et al., 'Crack Detection and Inpainting for Virtual Restoration of Paintings: The Case of the Ghent Altarpiece', *Signal Processing* 93, no. 3 (2013): 605–19, <https://doi.org/10.1016/j.sigpro.2012.07.022>.

¹⁷⁰ Angheluta and Chiroasca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 2.; I. Giakoumis, N. Nikolaidis, and I. Pitas, 'Digital Image Processing Techniques for the Detection and Removal of Cracks in Digitized Paintings', *IEEE Transactions on Image Processing* 15, no. 1 (2006): 178–88, <https://doi.org/10.1109/TIP.2005.860311>.

¹⁷¹ Pizurica et al., 'Digital Image Processing of The Ghent Altarpiece,' 114.; Cornelis et al., 'Crack Detection and Inpainting for Virtual Restoration of Paintings,'.

¹⁷² Angheluta and Chiroasca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 2.; Laurens Meeus et al., 'Deep Learning for Paint Loss Detection: A Case Study on the Ghent Altarpiece', *Image Processing for Art Investigation (IP4AI)*, 2018, 30–2, <http://hdl.handle.net/1854/LU-8577702>.

¹⁷³ Angheluta and Chiroasca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 2.

Infrared Thermography (IRT). However, numerous factors might cause human mistakes during a visual assessment of the physical state of an artwork. The volume of data that needs to be analysed can make even photographic evidence, that will also be subjected to human examination, difficult. Next to that, annotating data to document all damages can be a time-consuming task.¹⁷⁴

5.1. Machine Learning Applications

Considering that human examination is time consuming and can cause many mistakes, damage identification, classification, and documentation may be completed more accurately and quicker with the use of automatic interpretation techniques.¹⁷⁵

Research on damage detection utilising machine learning has focused mostly on craquelure detection. Some studies have also focussed on paint loss detection and one study tried to create a model to detect multiple types of damage at the same time. Furthermore, Infrared Thermography (IRT) in combination with machine learning has also been explored to detect damages in paintings.

5.1.1. Craquelure Detection

There are several reasons why a crack pattern recognition algorithm is desirable. The state of the structural components can be determined from crack patterns.¹⁷⁶ This can in turn be used to track the paint layer's deterioration or to conduct a more thorough investigation into the elements that lead to crack formation so that preventive measures can be implemented.¹⁷⁷ Assessing cracks is also useful for determining the geographical origin and authenticity because the artist's choice of materials and techniques affects how cracks grow and spread. A painting's conservation history may also be deduced from the presence of cracks. Overpainting, for instance, can be seen in places where crack patterns diverge, suggesting that these regions may have previously undergone restoration.¹⁷⁸ Finally, the digital restoration of artworks also uses crack detection as a pre-processing phase.¹⁷⁹




Fig. 20: Schematic representation of the proposed framework.

Most studies on crack detection solely consider fundamental statistical features. They merely retrieve the length and/or histogram of orientation of the cracks, leaving out the majority of the shape-related information. Sidorov and Hardeberg proposed a new method to deal with these shortcomings. By graphically describing craquelure patterns, their model complements the basic linear approximation with additional information. Furthermore, they are also the first researchers to try and classify the geographical origins of a painting

through crack detection based on deep learning.¹⁸⁰ The dataset they used consisted of high-resolution greyscale photographs of Italian, French, Flemish, German, and Dutch paintings of XIV–XVIII

¹⁷⁴ Angheluta and Chiroasca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 2-3.

¹⁷⁵ Angheluta and Chiroasca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 2-3.

¹⁷⁶ Roman Sizyakin et al., 'A Deep Learning Approach to Crack Detection in Panel Paintings', *Image Processing for Art Investigation (IP4AI)*, 2018, 40–2, <http://hdl.handle.net/1854/LU-8577805>.

¹⁷⁷ Cornelis et al., 'Crack Detection and inpainting for Virtual Restoration of Paintings,' 606.

¹⁷⁸ F.S. Abas and K. Martinez, 'Craquelure Analysis for Content-Based Retrieval', in *2002 14th International Conference on Digital Signal Processing Proceedings. DSP 2002 (Cat. No. 02TH8628)* (14th International Conference on Digital Signal Processing (DSP2002), Santorini, Greece: IEEE, 2002), 111–14, <https://doi.org/10.1109/ICDSP.2002.1027828>.

¹⁷⁹ Sizyakin et al., 'A Deep Learning Approach to Crack Detection in Panel Paintings,' 1.

¹⁸⁰ Sidorov and Hardeberg, 'Craquelure as a Graph,' 1.

centuries.¹⁸¹ The images were pre-processed with a morphological filter to extract a graph-like structure of the cracks whereby the topography of the nodes, the shape of the edges and the structure of the network were all considered during labelling. Notably, graph representation is robust to all types of geometrical distortion. According to Sidorov and Hardeberg, the extracted graph's geometrical structure provides meaningful information that is suitable to characterize the crack pattern.¹⁸² The researchers employed a Graph Neural Network (GNN) as a feature extractor and fed it unweighted non-directional graphs (node values and the connections between them) to demonstrate this. This process is completely independent of statistical processing, however when employed together, they allow for the most complete and thorough characterisation of a craquelure (Fig. 20). An SVM model was then used to classify the crack maps generated by the GNN into geographical origin classes. The test results showed that the GNN can accurately detect cracks and on the basis of this classify the geographical origin of paintings with an SVM.¹⁸³

Technical imaging has become an established method for research and conservation of paintings. However, since technical images are obtained with diverse imaging modalities, the varying image resolutions and different viewpoints of the devices should be considered. Image registration is therefore essential to align these technical images and compare them at pixel level.¹⁸⁴ Consequently, Sindel, Maier, and Christlein developed a program for automatic multi-modal image registration and artwork visualisation that is based on a deep

learning craquelure detection method. The registration method of the application operates with a software called CraquelureNet. This software aligns images from various modalities (VIS-IRR, VIS-UV, VIS-XR) based on the detected crack patterns. CraquelureNet utilises a Residual Network (ResNet) as a backbone and a Fully Convolutional Network (FCN) for keypoint detection and keypoint description. The network was trained on small sized image patches of multi-modal images collected from historical paintings and prints. The keypoint detection head classifies the small patches into “craquelure” and “background” and the keypoint description head is trained to learn cross-modal descriptors. By integrating CraquelureNet, multi-modal image registration can be carried out entirely automatically.¹⁸⁵ The Graphical User Interface (GUI) that was created by the researchers displays the registration results in an interactive manner, including the ability to superimpose and blend the registered image pair as well as individually and synchronously zoom in and out of the image views (Fig. 21).¹⁸⁶




Fig. 21: VIS-UV registration using CraquelureNet App of *Martin Luther* by the workshop of Lucas Cranach the Elder or Circle. In (a) the keypoint matches of VIS and UV are densely concentrated at the craquelure in the facial region; (b) and (c) show the registered pair and image fusions of the complete images and synchronised details.

In the past, a variety of (machine learning) techniques have been used to identify cracks. The downside of these models compared to deep learning is that they still involve manual labour in feature engineering and hyperparameter tuning. That being said, deep learning can only function well with a large volume of pre-annotated data. Such annotated data is hard to come by for crack detection in paintings since it requires human annotation by art historians. This limits the practical application of

¹⁸¹ Sidorov and Hardeberg, 'Craquelure as a Graph,' 6.

¹⁸² Sidorov and Hardeberg, 'Craquelure as a Graph,' 2-4.

¹⁸³ Sidorov and Hardeberg, 'Craquelure as a Graph,' 5-7.

¹⁸⁴ Sindel, Maier, and Christlein, 'A Multi-Modal Registration and Visualization Software Tool for Artworks Using CraquelureNet,' 1.

¹⁸⁵ Sindel, Maier, and Christlein, 'A Multi-Modal Registration and Visualization Software Tool for Artworks Using CraquelureNet,' 3.

¹⁸⁶ Sindel, Maier, and Christlein, 'A Multi-Modal Registration and Visualization Software Tool for Artworks Using CraquelureNet,' 5-6.




Fig. 22: First row: Original patch from the *Ghent Altarpiece* by Jan and Hubert van Eyck (left) and crack map produced by the pretrained CNN (right). Second row: Manual annotations for the first, second and third iteration of active learning, respectively from left to right. Third row: Crack map produced for the first, second and third iteration, respectively from left to right.

supervised learning and deep learning in general. The diversity in painting styles as well as the shapes and patterns of cracks adds to the complexity of the crack detection task and increases the difficulty of applying deep learning. Deep learning models struggle especially when previously unseen data from paintings is used.¹⁸⁷ As a solution to these challenges, Nadisic et al. proposed DAL4ART, an active learning framework utilising deep learning models for crack detection. The researchers decided to address the problem of data scarcity with the use of active learning. This technique usually requires significantly less data than standard deep learning since it retrain the model in an effective manner. Furthermore, it is able to learn continuously from new annotations (provided by humans), allows learning from partially annotated data, and reaches a better performance on previously unseen data. For the deep learning model, the researchers built further on the CNN proposed by Sizyakin et al.¹⁸⁸ Compared to previous methods this CNN model is an improvement because it operates with multi-modal images as

input data, it doesn't need full retraining when adding new data, and it doesn't suffer from false thickening of the crack boundaries.¹⁸⁹ The input data consisted of multi-modal images acquired with infrared photography, infrared reflectography, or X-radiography during the restoration campaign of the Ghent Altarpiece and can be accessed via <https://clostovaneck.kikirpa.be>. To enable fast re-learning, continuous learning, and transfer learning, the CNN from Sizyakin et al. was modified. Integrating new information without forgetting previous information or overfitting the model to the assignment at hand can be overcome with this adaptation. The researchers could conclude that the crack map generated by DAL4ART is more accurate than manual annotation, particularly when it comes to the boundaries of the cracks (Fig. 22). Moreover, DAL4ART produces results that are qualitatively finer and more accurate than the crack maps generated by the CNN model from Sizyakin et al. The proposed framework was integrated in a web interface where researchers can upload, manage and annotate images. The interface accommodates an infinite number of input modalities and accepts various picture formats for the input data. The interface's ability to finetune models on specific regions of paintings is another important feature.¹⁹⁰

5.1.2. Paint Loss Detection

Mapping lacunae and larger paint loss areas is one of the documentation tasks essential to painting conservation and restoration. Since non-machine-learning methods can only provide an approximate estimate of paint loss, mapping requires a significant amount of manual labour. The procedure becomes quite laborious and slow as a result. Smarter image processing methods, such as machine

¹⁸⁷ Nadisic et al., 'A Deep Active Learning Framework for Crack Detection in Digital Images of Paintings,' 2-3.

¹⁸⁸ Nadisic et al., 'A Deep Active Learning Framework for Crack Detection in Digital Images of Paintings,' 2-4.

¹⁸⁹ Roman Sizyakin et al., 'Crack Detection in Paintings Using Convolutional Neural Networks', *IEEE Access* 8 (2020): 74535–52, <https://doi.org/10.1109/ACCESS.2020.2988856>.

¹⁹⁰ Nadisic et al., 'A Deep Active Learning Framework for Crack Detection in Digital Images of Paintings,' 4-7.

learning, can be applied to improve automatic mapping. Furthermore, just as with crack detection, paint loss detection is used as a pre-processing phase for the digital restoration of artworks.¹⁹¹

Meeus et al. suggested a deep learning method that is semi-supervised and does not rely on a big number of labelled data to detect paint loss. The aim of this study was to maximise the performance of the model with little user intervention and thus increase the usefulness of automatic paint loss detection. As input multi-modal images taken of the Ghent Altarpiece were used.¹⁹² The proposed method utilises a translation invariant U-Net (TI-U-Net) to enhance the prediction quality with little data. The researchers added an unsupervised pretraining step with a convolutional autoencoder




Fig. 23: Paint loss learning with continuous training from a pretrained model with added annotations: (a) Close-up during treatment of the *Ghent Altarpiece* by Jan and Hubert van Eyck; (b) Detected paint loss based on the network trained on a different painting; (c) Added minimal annotations where red marks paint loss and blue marks regions which are not to be restored; (d) Paint loss detection after retraining the network.

to facilitate transfer learning. This step consists of integrating new annotations with preliminary prediction results and thus expanding the model continually, thereby limiting user interference and improving classification accuracy.¹⁹³ The performance of the trained TI-U-Net is significantly improved if the encoder is initialised by the encoder of a pretrained autoencoder (T TI-U-Net) in the case of limited annotations. Finally, a last training step to finetune the model (F-T TI-U-Net) was undertaken. The resulting images revealed that the detected paint loss areas were consistent with those treated during the actual restoration of the painting. Directly applying the model to new images could generate below satisfactory results. However, with transfer learning only a small number of annotations have to be added to the earlier paint loss detection results and then the model can be applied to new images with confidence (Fig. 23).¹⁹⁴

5.1.3. General Damage Detection




Fig. 24: Results of the proposed algorithm on a XIXth century icon: Activation map for several degradation image segments (upper row) vs the same segments without the activation map (lower row).

Contrary to other studies, Chiroasca and Angheluta decided to not solely focus on one type of damage, but to examine works of art that showed various kinds of physical deteriorations in the painted layer. The three forms of damage that the researchers specifically concentrated on were lacunae, cracks, and blisters. The set of images that they utilised is part of a photogrammetric documentation of an icon from the 19th century. The proposed method is based on the use of encoder patterns with a specialised CNN called VGG-16, created at Oxford University's Visual Geometry Group. This network is pre-trained for traditional computer vision object identification applications, and it further refines its training for application-specific use.¹⁹⁵ From the

¹⁹¹ Meeus et al., 'Assisting Classical Paintings Restoration,' 1.

¹⁹² Meeus et al., 'Assisting Classical Paintings Restoration,' 2.

¹⁹³ Meeus et al., 'Assisting Classical Paintings Restoration,' 3-4.

¹⁹⁴ Meeus et al., 'Assisting Classical Paintings Restoration,' 8-11.

¹⁹⁵ Angheluta and Chiroasca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 3.

resulting activation maps it can be stated that the trained model is able to assign the correct label, namely clean, crack, blister, paint loss (Fig. 24).¹⁹⁶

5.1.4. Damage Detection with Infrared Thermography

When it comes to inspecting cultural heritage objects, Infrared Thermography (IRT) is a convenient technique since it can detect surface defects. The main criterion for IRT to detect damage in artworks is the temperature difference between the damaged area and the surrounding regions visible in thermal images. However, the thermographic data produced using IRT frequently contains noise due to a variety of circumstances, such as the experimental environment, the equipment for image acquisition and an uneven background which may cause damages to be overlooked. For instance, a painting surface presents a difficulty for IRT because the various pigments alter the surface's emissivity. As such, it becomes difficult to recognise and differentiate damages based alone on visual inspection of thermal images. Analysing data while considering the heat equation can help obtain important insights from thermograms captured by infrared cameras. Yet, because an artwork may have complex damage configurations (various sizes, depths, materials, and types), analysing IRT data using only physical models is extremely difficult, especially when access to thermal data is limited to the external surface.¹⁹⁷ Consequently, researchers have concentrated on studying machine learning methods to enhance thermographic data and their ability to detect damage.¹⁹⁸

Principal Component Thermography (PCT) is a widely utilised technique in thermography where Principal Component Analysis (PCA) is employed for improving damage detection and studying damage distributions. However, the limited availability of thermal images makes it more difficult to determine the best projection directions for differentiating between damage and background elements, which limits the potential of PCT for feature extraction. The development of methods that improve PCT's capabilities and allow for more precise and effective damage detection has made this area a key research focus in the field. Liu et al. utilised two distinct approaches to enhance PCT's performance: Spectral Normalised Generative Adversarial Network (SNGAN) and Convolutional Autoencoder (CAE). The aim was to evaluate their relative impacts on PCT outcomes.¹⁹⁹ The input data was collected from a panel painting sample.²⁰⁰ Through a complete assessment, the performance improvements realised with these two techniques were evaluated by comparing the PCT analysis results generated from the two image enhancement methodologies with those acquired from the original data (Fig. 25). It could be concluded that SNGAN optimises and refines the generated images, thereby overcoming the problem of minimal data availability. CAE on the other hand, lessens the effect of noise on damage visibility which enhances the




Fig. 25: Results based on deep learning strategies for the Madonna sample: Top row: Original image (left) with damage distribution and PCT (right); Bottom row: PCT with SNGAN (left) and PCT with CAE (right).

¹⁹⁶ Angheluta and Chiroasca, 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models,' 7-9.

¹⁹⁷ Moradi et al., 'A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography,' 1-2.; Liu et al., 'Generative Deep Learning-Based Thermographic Inspection of Artwork,' 1-2.

¹⁹⁸ Liu et al., 'Generative Deep Learning-Based Thermographic Inspection of Artwork,' 2.

¹⁹⁹ Liu et al., 'Generative Deep Learning-Based Thermographic Inspection of Artwork,' 2-3.

²⁰⁰ Liu et al., 'Generative Deep Learning-Based Thermographic Inspection of Artwork,' 8.

images and in turn leads to better PCT results. The resulting images show that both SNGAN and CAE outperform regular PCT, but CAE delivers the best results overall.²⁰¹




Fig. 26: Results of MLP based on only temporal information (2nd row), MLP-U-Net based on both temporal and spatial information (3rd row), and the well-known algorithms of TSR, PPT, and PCT (4th and 5th rows) compared to the ground truth image (the 1st row) for the first dataset and PCT (the 4th and 5th rows) compared to the ground truth image (the 1st row) where the white spots represent damages. The red colour represents the unused pixels during training or testing.

The method proposed by Moradi et al. is based on the spatiotemporal physical features of the IRT image sequence. The analysis of IRT image sequences (potentially from a thermal video) can be performed using two types of input: The spatial information aspect, that involves the temperature variation within each image (at one moment) over all pixels and the temporal information aspect, that involves the temperature variation of each pixel over time. The researchers created a supervised method utilising a spatiotemporal deep neural network (STDNN) to categorise pixels into defective and healthy regions, thereby displaying the relevant intact and damaged regions. Moradi et al. deliberately chose a supervised approach because previous studies with unsupervised learning were not able to deal with defects in signal-to-noise ratio (SNR) and complex interference.²⁰² In this study a photograph of a replica of a mural painting was used as input data. There are two sub-models in the proposed framework which operate sequentially: the first being an MLP, which is used to classify the temporal signal correlated to every pixel from the input image as either healthy or defective; the second being a U-Net that uses the MLP's output as input and segments images based on the associations between surrounding pixels in spatial information into healthy and faulty regions. To reconstruct the entire final image, the training and test sets from the U-Net model outputs are concatenated.²⁰³ The resulting images were compared with non-machine-learning methods such as PCT, Pulsed Phase

Thermography (PPT), and Thermographic Signal Reconstruction (TSR). This comparison shows that the deep learning model is able to perform more accurate damage detection (Fig. 26). Furthermore, the surface emissivity issue brought on by pigments has no negative effect on the accuracy of the proposed method.²⁰⁴

²⁰¹ Liu et al., 'Generative Deep Learning-Based Thermographic Inspection of Artwork,' 10-3.

²⁰² Moradi et al., 'A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography,' 2.

²⁰³ Moradi et al., 'A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography,' 3-4.

²⁰⁴ Moradi et al., 'A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography,' 7-10.

5.2. Discussion and Future Directions




Fig. 27: Summary of the general workflow for damage detection.

The input data for damage detection always consisted of technical images such as IR, IRR, IRT, VIS, X-ray. Sometimes multi-modal datasets consisting of multiple types of technical images were used. The technical images were either derived from historical data or from modern data; no artificial data was used. Some research was conducted on a dataset with little variation, sometimes even without data from historical paintings. This leaves room to experiment with the input data to improve models. For example, in the model proposed by Moradi et al. to detect damages with IRT they only tested the model on a mock-up.²⁰⁵ When it comes to enhancing infrared thermograms for damage detection, it may be interesting to investigate how deep learning techniques might be combined with other feature extraction techniques than PCT.

Most of the research described above works with supervised methods and manually labelled data according to the type(s) of damage present. However, this can be very time-consuming therefore it could be interesting to explore unsupervised methods like Liu et al. tested to enhance infrared thermograms.²⁰⁶

For crack detection a CNN, GNN, and a combination of a ResNet and FCN were employed. Investigating the use of state-of-the-art U-Net models could deliver interesting results in crack detection. To detect paint losses an adapted U-Net model was utilised. Furthermore, an adapted CNN, namely a VGG-16 was used for general damage detection, in this case focussing on cracks, paint losses, and blisters. Further improvement could be made by training a model for more types of damages including abrasion, flaking, water damage. To detect damage with IRT a SNGAN, CAE, MLP and U-Net were chosen.

Some researchers opted to improve their proposed model by utilising various techniques such as active learning, continuous learning, finetuning, and transfer learning. These techniques are all focused on adding new input data after the model was already trained in order to expand the usefulness of the model.

²⁰⁵ Moradi et al., 'A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography,' 2.

²⁰⁶

6. Damage Prediction

Panel paintings are artworks distinguished by their hygroscopic and heterogeneous properties. The hygroscopic nature is primarily caused by the presence of a wooden support (panel), which is very prone to water, particularly when changes in temperature (T) and relative humidity (RH) are at play. The heterogeneous nature is caused by the presence of layers of various materials, such as the wooden support, gesso (a mixture of animal glue and chalk) as preparatory layer, paint, and varnishes. Furthermore, there is a possibility of damage in the pictorial layer due to the mismatch between the wooden panel's moisture response and that of a stiff gesso layer. As a result, the mechanical integrity and qualities of paintings may be negatively impacted when they are, for instance, loaned from one exhibition to another and their microclimate naturally shifts. A solid understanding of the mechanical behaviour of panel paintings under microclimatic fluctuations is essential for their conservation.²⁰⁷

Countless research endeavours have relied on monitoring, experimental, and numerical projects to collect the vast amount of data regarding strategies for minimising damage. Yet, because they must adhere to the timelines of microclimatic variations and fluctuations as well as the characteristic response time of multi-material objects during diffusive phenomena, such projects are often laborious.²⁰⁸

6.1. Machine Learning Applications

As a solution to current time-consuming research projects, the generalisation abilities of machine learning, along with established failure criteria from the fields of mechanical engineering and fracture mechanics, can be combined to obtain reliable findings and guidance on how to move forward with conservation practices. It could speed up the decision-making process regarding preventive conservation of paintings and collections and add value to the decision-making process.²⁰⁹




Fig. 28: Classification results: Separation Lines (SL) in plane $z = 6$ selected within the Border Volume (BV) for each configuration. The blue and red points follow the logic of the SED criterion.

Califano et al. have devised a new method of applying machine learning to predict damage in panel paintings with an already existing crack. Their approach is novel in that it has the potential to lower the computational cost of the thousands of numerical analyses required to evaluate the potential for damage development in a large number of geometric configurations.²¹⁰ The researchers started by generating a parametric 3D model of a panel painting composed of a layer of gesso with a crack on a wooden framework. They decided to use the strain energy density (SED) failure criterion as a decision rule for predicting the

future probability of safe or unsafe conditions, meaning whether the crack in the gesso would grow or not. Several Finite Element (FE) simulations were performed with the 3D model using different

²⁰⁷ Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 146.

²⁰⁸ Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 146.

²⁰⁹ Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 146.

²¹⁰ Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 146.

geometrical parameters to determine with the SED failure criterion whether additional damage could arise.²¹¹ Simplifying the problem and concentrating on 2D models retrieved from the simulated 3D configurations was done based on preliminary geometrical considerations. An eXtreme Gradient Boosting (XGBoost) machine, which is capable of predicting the emergence of safe or harmful conditions through basic classification tasks, was trained and tested using the extracted data (Fig. 28). The generated graphs demonstrate that the proposed method can accurately predict whether further damage will occur in the future.²¹²

²¹¹ Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 147-49.

²¹² Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 150-52.

6.2. Discussion and Future Directions




Fig. 29: Summary of the general workflow for damage prediction.

The technique explained above utilised a digital 3D-model consisting of a layer of gesso with a crack on top of a wooden panel as a sample to perform FE simulations where SED was used as failure criterion. The FE simulations were subsequently reduced to 2D sub-models.²¹³ In future research, when choosing to base the input data on a digital 3D-model, more complex configurations that consider the layered structure of a painting could be employed. Investigating the feasibility of using machine learning approaches on a complete 3D model as opposed to just 2D sub-models, could offer another avenue for further research. Furthermore, digital models cannot always capture the history of a painting (e.g. climatic history, earlier repaired damages). For example, paintings get used to their climatic circumstances, even if they are not ideal. Therefore, when a painting is brought into a different climatic environment it might not react well, even though those new circumstances are theoretically more ideal. Further possibilities thus lie in utilising input data based on previously measured deformations on real objects. This will be more time-consuming but might deliver more realistic results.

The 2D-sub models of the FE simulations that were used as input data were manually labelled as safe or unsafe to train a supervised algorithm.²¹⁴ Since a lot of factors such as the materials used, structural components (e.g. cradling) the surrounding climatic environment, but also the history of a painting that play a role in the degradation of a painting cannot always be modelled, using unlabelled data and unsupervised methods might also deliver interesting results.

Here, the researchers used a traditional machine learning method, namely XGBoost. Future research could focus on optimising the XGBoost model's hyperparameters. Furthermore, comparing an XGBoost model's output against those of other algorithms, such as deep neural networks, may result in damage predictions that are more accurate.²¹⁵

Finally, the study discussed here only focusses on the future declining of a crack in the gesso. However, a painting can undergo various types of damages. Therefore, in future studies the focus should also lie on other types of damages, e.g. based on images of a painting surface future paint losses might be predicted.

²¹³ Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 146.

²¹⁴ Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 150-52

²¹⁵ Califano et al., 'Predicting Damage Evolution in Panel Paintings with Machine Learning,' 155-56.

7. Conclusion

An overview of machine learning approaches for painting conservation has been provided in this thesis. It is clear that these applications have been able to facilitate the conservation process. Though in terms of input data and used algorithms further research is possible. The studies discussed in this thesis were divided into four primary topics: enhancing the readability of technical images, pigment analysis, damage detection, and damage prediction.

The methods discussed to enhance the readability of technical images used input data that consisted of technical images collected from historical data or modern data. Artificial data created on the basis of the historical or modern data, was also used. In the future, focussing on enhancing underdrawings with less commonly used materials could lead to interesting results. Additionally, for enhancing underdrawings, experimenting with more complex training data, artificial as well as modern data, that considers the multi-layered paint stratigraphy in historical paintings, could deliver more accurate results. Both a supervised method with manual labelling and a self-supervised method without labelling the input data were employed. Deep learning models such as CNN, and FNN were used to enhance the visibility of an underdrawing in a NIR reflectogram and to separate a mixed X-ray image. In future work the benefits of more advanced neural network models should be explored. The method developed to separate the mixed X-rays can also be applied to other images of paintings with mixed compositions, such as infrared reflectograms, XRF element distribution maps, or outputs from RIS.

In all the mentioned methods for pigment analysis the training and testing data consisted of spectral data collected with one or more complementary techniques. The spectral data was collected from artificial data, modern data, or historical data. Some researchers only used historical data to test the model and not to train it. However, artificial and modern data are oftentimes less complex than historical data. In addition, spectral data collected from historical sources contains certain characteristics that artificial and modern data lack. For example, weathering, ageing, and fading influence the reflectance of pigments and are hard to model. Furthermore, the material composition of the whole painting and especially the variety in pigments and mixtures should be considered. The training data consisted of both manually labelled data for supervised methods and unlabelled training data for unsupervised methods. As mentioned, some variables that influence the identification and unmixing of pigments can be difficult to determine. Unsupervised learning methods offer a solution because they can extract latent features of signals without any labels. Therefore, in future research the use of unsupervised techniques should be further explored. Pigment analysis has been performed with the following traditional machine learning models and deep learning models: 1D-CNN, autoencoder, CNN, DT, FC-NN, FNN, MLP, SOM and SVM. Some researchers used multiple algorithms to split up the tasks. Finetuning and transfer learning were employed by in several studies to improve the suggested algorithms' performance.

The input data for damage detection always consisted of technical images that were either derived from historical data or from modern data. Certain studies were performed on a dataset with little variation, and sometimes no data historical data was used. To further improve models, working with historical data to train the model should be explored. Most of the researchers worked with supervised methods and manually labelled data. However, manually labelling data can be very time-consuming, therefore it could be interesting to explore unsupervised methods. A variety of deep learning models were used for damage detection: CAE, CNN, FCN, GNN, MLP, ResNet, SNGAN, U-Net, and VGG-16. Currently the focus is only on certain types of damage. Therefore, in the future other types of damages such as abrasion, flaking, water damage should also be included in the input data. Some researchers

were able to improve their proposed model by utilising active learning, continuous learning, finetuning, and transfer learning.

The method for damage prediction utilised a digital 3D-model as a sample to perform FE simulations where SED was used as failure criterion. The FE simulations were reduced to 2D sub-models. In future research, a more complex digital 3D-model, that consider the stratigraphy of a painting could be employed. Utilising a complete 3D model as input opposed to just 2D sub-models, could offer another avenue for further research. Additionally, digital models cannot always capture the history of a painting (e.g. climatic history, earlier repaired damages). Further possibilities thus lie in utilising input data based on deformations measured on real objects. The training data was manually labelled as safe or unsafe. Since a lot of factors such cannot always be modelled, using unlabelled data and unsupervised methods might also deliver interesting results. The researchers used XGBoost, a traditional machine learning method. Focusing on optimising the XGBoost model's hyperparameters, could lead to more accurate results in the future. Employing a type of deep neural network instead of a traditional machine learning algorithm may also result in better damage predictions. Finally, in future studies the focus should also lie on the prediction of other types of damages than a crack in the gesso-layer.

In summary, most research can benefit from working with a variety of data sources preferably including historical data to produce more accurate results. However, labelled data and large datasets in general can be hard to collect. To create a model with better generalisation capabilities, improvement methods can be used or unsupervised learning methods can be further explored.

8. Bibliography

- 'Fiber Optics Reflectance Spectroscopy | Popart'. Accessed 5 August 2024. <https://popart-highlights.mnhn.fr/identification/fiber-optics-reflectance-spectroscopy/index.html>.
- 'What Is CIELAB Color Space?' Accessed 5 August 2024. <https://www.hunterlab.com/blog/what-is-cielab-color-space/>.
- Abas, F.S., and K. Martinez. 'Craquelure Analysis for Content-Based Retrieval.' In *2002 14th International Conference on Digital Signal Processing Proceedings. DSP 2002 (Cat. No.02TH8628)*, 111–14. Santorini, Greece: IEEE, 2002. <https://doi.org/10.1109/ICDSP.2002.1027828>.
- Ahmadsabry. 'A Perfect Guide to Understand Encoder Decoders in Depth with Visuals'. *Medium* (blog), 24 June 2023. <https://medium.com/@ahmadsabry678/a-perfect-guide-to-understand-encoder-decoders-in-depth-with-visuals-30805c23659b>.
- Angheluta, Laurentiu Marian, and Alecsandru Chiroasca. 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models.' *Romanian Reports in Physics* 72, no. 805 (2020).
- Balageas, Daniel L., Jean-Michel Roche, François-Henri Leroy, Wei-Min Liu, and Alexander M. Gorbach. 'The Thermographic Signal Reconstruction Method: A Powerful Tool for the Enhancement of Transient Thermographic Images'. *Biocybernetics and Biomedical Engineering* 35, no. 1 (2015): 1–9. <https://doi.org/10.1016/j.bbe.2014.07.002>.
- Burkov, Andriy. *The Hundred-Page Machine Learning Book*. Polen: Andriy Burkov, 2019.
- Califano, America, Pietro Foti, Filippo Berto, Marco Baiesi, and Chiara Bertolin. 'Predicting Damage Evolution in Panel Paintings with Machine Learning.' *Procedia Structural Integrity* 41 (2022): 145–57. <https://doi.org/10.1016/j.prostr.2022.05.017>.
- Chen, Ailin, Rui Jesus, and Marcia Vilarigues. 'Convolutional Neural Network-Based Pure Paint Pigment Identification Using Hyperspectral Images'. In *ACM Multimedia Asia*, 1–7. Gold Coast Australia: ACM, 2021. <https://doi.org/10.1145/3469877.3495641>.
- Chen, Ailin, Rui Jesus, and Márcia Vilarigues. 'Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms'. *SN Computer Science* 4, no. 2 (21 December 2022): 115. <https://doi.org/10.1007/s42979-022-01529-8>.
- Chokshi, Akanksha. 'Supporting Fully Convolutional Networks (and U-Net) for Image Segmentation'. Accessed 5 June 2024. <https://www.datature.io/blog/supporting-fully-convolutional-networks-and-u-net-for-image-segmentation>.
- Cornelis, B., T. Ružić, E. Gezels, A. Doms, A. Pižurica, L. Platiša, J. Cornelis, M. Martens, M. De Mey, and I. Daubechies. 'Crack Detection and Inpainting for Virtual Restoration of Paintings: The Case of the Ghent Altarpiece.' *Signal Processing* 93, no. 3 (2013): 605–19. <https://doi.org/10.1016/j.sigpro.2012.07.022>.
- De Boer, D.K.G. 'Fundamental Parameters for X-Ray Fluorescence Analysis'. *Spectrochimica Acta Part B: Atomic Spectroscopy* 44, no. 11 (1989): 1171–90. [https://doi.org/10.1016/0584-8547\(89\)80114-4](https://doi.org/10.1016/0584-8547(89)80114-4).
- Fukumoto, Kensuke, Norimichi Tsumura, and Roy Berns. 'Estimating Concentrations of Pigments Using Encoder-Decoder Type of Neural Network.' *Color and Imaging Conference* 27, no. 1 (21 October 2019): 149–52. <https://doi.org/10.2352/issn.2169-2629.2019.27.28>.

GeeksforGeeks. 'Artificial Neural Networks and Its Applications', 24 June 2020.
<https://www.geeksforgeeks.org/artificial-neural-networks-and-its-applications/>.

GeeksforGeeks. 'Autoencoders -Machine Learning', 21 June 2019.
<https://www.geeksforgeeks.org/auto-encoders/>.

GeeksforGeeks. 'Continual Learning in Machine Learning', 13 December 2023.
<https://www.geeksforgeeks.org/continual-learning-in-machine-learning/>.

GeeksforGeeks. 'Decision Tree', 16 October 2017. <https://www.geeksforgeeks.org/decision-tree/>.

GeeksforGeeks. 'Difference Between Feed-Forward Neural Networks and Recurrent Neural Networks', 3 January 2023. <https://www.geeksforgeeks.org/difference-between-feed-forward-neural-networks-and-recurrent-neural-networks/>.

GeeksforGeeks. 'Difference Between Model Parameters VS HyperParameters', 11 May 2020.
<https://www.geeksforgeeks.org/difference-between-model-parameters-vs-hyperparameters/>.

GeeksforGeeks. 'Difference between Multilayer Perceptron and Linear Regression', 4 March 2022.
<https://www.geeksforgeeks.org/difference-between-multilayer-perceptron-and-linear-regression/>.

GeeksforGeeks. 'Difference Between RGB, CMYK, HSV, and YIQ Color Models', 16 November 2021.
<https://www.geeksforgeeks.org/difference-between-rgb-cmyk-hsv-and-yiq-color-models/>.

GeeksforGeeks. 'Electromagnetic Spectrum | Wavelengths, Diagram and Uses', 20 December 2021.
<https://www.geeksforgeeks.org/electromagnetic-spectrum/>.

GeeksforGeeks. 'Ensemble Classifier | Data Mining', 14 May 2019.
<https://www.geeksforgeeks.org/ensemble-classifier-data-mining/>.

GeeksforGeeks. 'Generative Adversarial Network (GAN)', 15 January 2019.
<https://www.geeksforgeeks.org/generative-adversarial-network-gan/>.

GeeksforGeeks. 'Implement Convolutional Autoencoder in PyTorch with CUDA', 31 July 2023.
<https://www.geeksforgeeks.org/implement-convolutional-autoencoder-in-pytorch-with-cuda/>.

GeeksforGeeks. 'Introduction to Convolution Neural Network', 21 August 2017.
<https://www.geeksforgeeks.org/introduction-convolution-neural-network/>.

GeeksforGeeks. 'Introduction to Deep Learning', 1 June 2018.
<https://www.geeksforgeeks.org/introduction-deep-learning/>.

GeeksforGeeks. 'K-Nearest Neighbor(KNN) Algorithm', 14 April 2017. <https://www.geeksforgeeks.org/k-nearest-neighbours/>.

GeeksforGeeks. 'K Means Clustering - Introduction', 2 May 2017. <https://www.geeksforgeeks.org/k-means-clustering-introduction/>.

GeeksforGeeks. 'Linear Regression in Machine Learning', 13 September 2018.
<https://www.geeksforgeeks.org/ml-linear-regression/>.

GeeksforGeeks. 'ML | Active Learning', 23 July 2019. <https://www.geeksforgeeks.org/ml-active-learning/>.

GeeksforGeeks. 'ML | Underfitting and Overfitting', 23 November 2017.
<https://www.geeksforgeeks.org/underfitting-and-overfitting-in-machine-learning/>.

GeeksforGeeks. 'ML | XGBoost (EXtreme Gradient Boosting)', 19 August 2019.
<https://www.geeksforgeeks.org/ml-xgboost-extreme-gradient-boosting/>.

GeeksforGeeks. 'Noise Removal Using Median Filter in C++', 17 June 2019.
<https://www.geeksforgeeks.org/noise-removal-using-median-filter-in-c/>.

GeeksforGeeks. 'Object Detection vs Object Recognition vs Image Segmentation', 27 February 2020.
<https://www.geeksforgeeks.org/object-detection-vs-object-recognition-vs-image-segmentation/>.

GeeksforGeeks. 'Python | Data Augmentation', 5 September 2019.
<https://www.geeksforgeeks.org/python-data-augmentation/>.

GeeksforGeeks. 'Python | Morphological Operations in Image Processing (Opening) | Set-1', 27 May 2019. <https://www.geeksforgeeks.org/python-morphological-operations-in-image-processing-opening-set-1/>.

GeeksforGeeks. 'Q-Learning in Python', 6 February 2019. <https://www.geeksforgeeks.org/q-learning-in-python/>.

GeeksforGeeks. 'Random Forest Algorithm in Machine Learning', 22 February 2024.
<https://www.geeksforgeeks.org/random-forest-algorithm-in-machine-learning/>.

GeeksforGeeks. 'Residual Networks (ResNet) - Deep Learning', 3 June 2020.
<https://www.geeksforgeeks.org/residual-networks-resnet-deep-learning/>.

GeeksforGeeks. 'Self Organizing Maps - Kohonen Maps', 1 July 2020.
<https://www.geeksforgeeks.org/self-organising-maps-kohonen-maps/>.

GeeksforGeeks. 'Self-Supervised Learning (SSL)', 27 January 2023. <https://www.geeksforgeeks.org/self-supervised-learning-ssl/>.

GeeksforGeeks. 'Spatial Filters - Averaging Filter and Median Filter in Image Processing', 2 May 2020.
<https://www.geeksforgeeks.org/spatial-filters-averaging-filter-and-median-filter-in-image-processing/>.

GeeksforGeeks. 'Supervised and Unsupervised Learning', 1 October 2017.
<https://www.geeksforgeeks.org/supervised-unsupervised-learning/>.

GeeksforGeeks. 'Support Vector Machine (SVM) Algorithm', 20 January 2021.
<https://www.geeksforgeeks.org/support-vector-machine-algorithm/>.

GeeksforGeeks. 'Transfer Learning with Fine-Tuning', 11 July 2023.
<https://www.geeksforgeeks.org/transfer-learning-with-fine-tuning-in-nlp/>.

GeeksforGeeks. 'U-Net Architecture Explained', 8 June 2023. <https://www.geeksforgeeks.org/u-net-architecture-explained/>.

GeeksforGeeks. 'VGG-16 | CNN Model', 26 February 2020. <https://www.geeksforgeeks.org/vgg-16-cnn-model/>.

GeeksforGeeks. 'What Are Graph Neural Networks?', 2 July 2023.
<https://www.geeksforgeeks.org/what-are-graph-neural-networks/>.

GeeksforGeeks. 'What Is a 1D Convolutional Layer in Deep Learning?', 19 February 2024.
<https://www.geeksforgeeks.org/what-is-a-1d-convolutional-layer-in-deep-learning/>.

GeeksforGeeks. 'What Is Feature Engineering?', 20 March 2023. <https://www.geeksforgeeks.org/what-is-feature-engineering/>.

GeeksforGeeks. 'What Is Hyperspectral Imaging? Where It Is Used?', 14 May 2024.
<https://www.geeksforgeeks.org/what-is-hyperspectral-imaging-where-it-is-used/>.

- GeeksforGeeks. 'What Is Infrared (IR)?', 30 April 2024. <https://www.geeksforgeeks.org/what-is-infrared/>.
- GeeksforGeeks. 'What Is KeyPoint Detection?', 11 June 2024. <https://www.geeksforgeeks.org/what-is-keypoint-detection/>.
- GeeksforGeeks. 'What Is Spectral Imaging in Image Processing?', 24 April 2024. <https://www.geeksforgeeks.org/what-is-spectral-imaging-in-image-processing/>.
- GeeksforGeeks. 'What Is the Difference between "Equivariant to Translation" and "Invariant to Translation"', 19 February 2024. <https://www.geeksforgeeks.org/what-is-the-difference-between-equivariant-to-translation-and-invariant-to-translation/>.
- Giakoumis, I., N. Nikolaidis, and I. Pitas. 'Digital Image Processing Techniques for the Detection and Removal of Cracks in Digitized Paintings'. *IEEE Transactions on Image Processing* 15, no. 1 (2006): 178–88. <https://doi.org/10.1109/TIP.2005.860311>.
- Grabowski, Bartosz, Wojciech Masarczyk, Przemysław Głomb, and Agata Mendys. 'Automatic Pigment Identification from Hyperspectral Data'. *Journal of Cultural Heritage* 31 (2018): 1–12. <https://doi.org/10.1016/j.culher.2018.01.003>.
- Heide, Felix, Wolfgang Heidrich, and Gordon Wetzstein. 'Fast and Flexible Convolutional Sparse Coding'. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5135–43. Boston, MA, USA: IEEE, 2015. <https://doi.org/10.1109/CVPR.2015.7299149>.
- Jones, Cerys, Nathan S. Daly, Catherine Higgitt, and Miguel R. D. Rodrigues. 'Neural Network-Based Classification of X-Ray Fluorescence Spectra of Artists' Pigments: An Approach Leveraging a Synthetic Dataset Created Using the Fundamental Parameters Method'. *Heritage Science* 10, no. 1 (2022): 88. <https://doi.org/10.1186/s40494-022-00716-3>.
- Karella, Tomáš, Jan Blažek, and Jana Striová. 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms'. *Journal of Cultural Heritage* 58 (2022): 186–98. <https://doi.org/10.1016/j.culher.2022.09.022>.
- Kleynhans, Tania, Catherine M. Schmidt Patterson, Kathryn A. Dooley, David W. Messinger, and John K. Delaney. 'An Alternative Approach to Mapping Pigments in Paintings with Hyperspectral Reflectance Image Cubes Using Artificial Intelligence'. *Heritage Science* 8, no. 1 (2020): 84. <https://doi.org/10.1186/s40494-020-00427-7>.
- Kogou, Sotiria, Lynn Lee, Golnaz Shahtahmassebi, and Haida Liang. 'A New Approach to the Interpretation of XRF Spectral Imaging Data Using Neural Networks'. *X-Ray Spectrometry* 50, no. 4 (7 August 2021): 310–19. <https://doi.org/10.1002/xrs.3188>.
- Liu, Lingxi, Tsveta Miteva, Giovanni Delnevo, Silvia Mirri, Philippe Walter, Laurence De Viguerie, and Emeline Pouyet. 'Neural Networks for Hyperspectral Imaging of Historical Paintings: A Practical Review'. *Sensors* 23, no. 5 (22 February 2023): 2419. <https://doi.org/10.3390/s23052419>.
- Liu, Yi, Fumin Wang, Zhili Jiang, Stefano Sfarra, Kaixin Liu, and Yuan Yao. 'Generative Deep Learning-Based Thermographic Inspection of Artwork'. *Sensors* 23, no. 14 (13 July 2023): 6362. <https://doi.org/10.3390/s23146362>.
- Mandal, Dipendra J., Marius Pedersen, Sony George, and Clotilde Boust. 'Comparison of Pigment Classification Algorithms on Non-Flat Surfaces Using Hyperspectral Imaging'. *Journal of Imaging*

Science and Technology 67, no. 5 (1 September 2023): 1–25.
<https://doi.org/10.2352/J.ImagingSci.Technol.2023.67.5.050405>.

Mandal, Dipendra J., Marius Pedersen, Sony George, Hilda Deborah, and Clotilde Boust. 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging'. *Journal of Imaging Science and Technology* 67, no. 3 (1 May 2023): 030403-1-030403–18.
<https://doi.org/10.2352/J.ImagingSci.Technol.2023.67.3.030403>.

Marinetti, Sergio, Yuri A. Plotnikov, William P. Winfree, and Alberto Braggiotti. 'Pulse Phase Thermography for Defect Detection and Visualization'. edited by Ajit K. Mal, 230–38. Newport Beach, CA, 1999. <https://doi.org/10.1117/12.339890>.

MediaWiki. 'Reflected Infrared Digital Photography (IR)', 28 February 2022. [https://www.conservation-wiki.com/wiki/Reflected_infrared_digital_photography_\(IR\)](https://www.conservation-wiki.com/wiki/Reflected_infrared_digital_photography_(IR)).

Meeus, Laurens, Shaoguang Huang, Bart Devolder, Maximiliaan Martens, and Aleksandra Pizurica. 'Deep Learning for Paint Loss Detection: A Case Study on the Ghent Altarpiece.' *Image Processing for Art Investigation (IP4AI)*, 2018, 30–32. <http://hdl.handle.net/1854/LU-8577702>.

Meeus, Laurens, Shaoguang Huang, Nina Zizakic, Xianghui Xie, Bart Devolder, Hélène Dubois, Maximiliaan Martens, and Aleksandra Pizurica. '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*, edited by Peter Schelkens and Tomasz Kozacki, 16. Online Only, France: SPIE, 2020. <https://doi.org/10.1117/12.2556000>.

Monga, Vishal, Yuelong Li, and Yonina C. Eldar. 'Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing'. *IEEE Signal Processing Magazine* 38, no. 2 (2021): 18–44.
<https://doi.org/10.1109/MSP.2020.3016905>.

Moradi, Morteza, Ramin Ghorbani, Stefano Sfarra, David M.J. Tax, and Dimitrios Zarouchas. 'A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography'. *Sensors* 22, no. 23 (1 December 2022): 9361.
<https://doi.org/10.3390/s22239361>.

Nadisic, Nicolas, Yoann Arhant, Niels Vyncke, Sebastiaan Verplancke, Srdan Lazendic, and Aleksandra Pizurica. 'A Deep Active Learning Framework for Crack Detection in Digital Images of Paintings'. In *7th International Conference on Smart Monitoring, Assessment and Rehabilitation of Civil Structures, SMAR 2024*, 1–8. Elsevier, 2024. <http://hdl.handle.net/1854/LU-01J18DKNDKGSNKM5HVB9PYS7C>.

Oktar, Yigit, and Mehmet Turkan. 'A Review of Sparsity-Based Clustering Methods'. *Signal Processing* 148 (2018): 20–30. <https://doi.org/10.1016/j.sigpro.2018.02.010>.

OnScale. 'Strain Energy Density: What, Why and How?', 3 January 2022.
<http://https%253A%252F%252Fonscale.com%252Fstrain-energy-density-what-why-and-how%252F>.

Pizurica, Aleksandra, Ljiljana Platisa, Tijana Ruzic, Bruno Cornelis, Ann Doms, 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.' *IEEE Signal Processing Magazine* 32, no. 4 (2015): 112–22. <https://doi.org/10.1109/MSP.2015.2411753>.

Polak, Adam, Timothy Kelman, Paul Murray, Stephen Marshall, David J.M. Stothard, Nicholas Eastaugh, and Francis Eastaugh. 'Hyperspectral Imaging Combined with Data Classification Techniques as an Aid

- for Artwork Authentication'. *Journal of Cultural Heritage* 26 (2017): 1–11. <https://doi.org/10.1016/j.culher.2017.01.013>.
- Pouyet, Emeline, Tsveta Miteva, Neda Rohani, and Laurence De Viguier. 'Artificial Intelligence for Pigment Classification Task in the Short-Wave Infrared Range'. *Sensors* 21, no. 18 (13 September 2021): 6150. <https://doi.org/10.3390/s21186150>.
- Pu, Wei, Jun-Jie Huang, Barak Sober, Nathan Daly, Catherine Higgitt, Ingrid Daubechies, Pier Luigi Dragotti, and Miguel R. D. Rodrigues. 'Mixed X-Ray Image Separation for Artworks With Concealed Designs'. *IEEE Transactions on Image Processing* 31 (2022): 4458–73. <https://doi.org/10.1109/TIP.2022.3185488>.
- Rohani, Neda, Emeline Pouyet, Marc Walton, Oliver Cossairt, and Aggelos K. Katsaggelos. 'Pigment Unmixing of Hyperspectral Images of Paintings Using Deep Neural Networks'. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3217–21. Brighton, United Kingdom: IEEE, 2019. <https://doi.org/10.1109/ICASSP.2019.8682838>.
- Sabetsarvestani, Z., B. Sober, C. Higgitt, I. Daubechies, and M. R. D. Rodrigues. 'Artificial Intelligence for Art Investigation: Meeting the Challenge of Separating x-Ray Images of the *Ghent Altarpiece*'. *Science Advances* 5, no. 8 (2 August 2019): eaaw7416. <https://doi.org/10.1126/sciadv.aaw7416>.
- Shitomi, Ryuta, Mayuka Tsuji, Yuki Fujimura, Takuya Funatomi, Yasuhiro Mukaigawa, Tetsuro Morimoto, Takeshi Oishi, Jun Takamatsu, and Katsushi Ikeuchi. 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments'. *Journal of the Optical Society of America A* 40, no. 1 (1 January 2023): 116. <https://doi.org/10.1364/JOSAA.472775>.
- Shroff, Megha. 'Know Your Neural Network Architecture More by Understanding These Terms'. *Medium* (blog), 16 May 2023. <https://medium.com/@shroffmegha6695/know-your-neural-network-architecture-more-by-understanding-these-terms-67faf4ea0efb>.
- Sidorov, Oleksii, and Jon Yngve Hardeberg. 'Craquelure as a Graph: Application of Image Processing and Graph Neural Networks to the Description of Fracture Patterns,' 2019. <https://doi.org/10.48550/ARXIV.1905.05010>.
- Sindel, Aline, Andreas Maier, and Vincent Christlein. 'A Multi-Modal Registration and Visualization Software Tool for Artworks Using CraquelureNet,' 2022. <https://doi.org/10.48550/ARXIV.2208.08836>.
- Sizyakin, Roman, Bruno Cornelis, Laurens Meeus, Helene Dubois, Maximiliaan Martens, Viacheslav Voronin, and Aleksandra Pizurica. 'Crack Detection in Paintings Using Convolutional Neural Networks'. *IEEE Access* 8 (2020): 74535–52. <https://doi.org/10.1109/ACCESS.2020.2988856>.
- Sizyakin, Roman, Bruno Cornelis, Laurens Meeus, Maximiliaan Martens, Viacheslav Voronin, and Aleksandra Pizurica. 'A Deep Learning Approach to Crack Detection in Panel Paintings.' *Image Processing for Art Investigation (IP4AI)*, 2018, 40–42. <http://hdl.handle.net/1854/LU-8577805>.
- Sizyakin, Roman, Bruno Cornelis, Laurens Meeus, Viacheslav Voronin, and Aleksandra Pizurica. 'A Two-Stream Neural Network Architecture for the Detection and Analysis of Cracks in Panel Paintings.' In *Optics, Photonics and Digital Technologies for Imaging Applications VI*, edited by Peter Schelkens and Tomasz Kozacki, 10. Online Only, France: SPIE, 2020. <https://doi.org/10.1117/12.2555857>.
- Sober, Barak, Spike Bucklow, Nathan Daly, Ingrid Daubechies, Pier Luigi Dragotti, Catherine Higgitt, Jun-Jie Huang, et al. 'Revealing and Reconstructing Hidden or Lost Features in Art Investigation'. *IEEE BITS the Information Theory Magazine*, 2022, 1–16. <https://doi.org/10.1109/MBITS.2022.3207125>.

- Theobald, Oliver. *Machine Learning for Absolute Beginners: A Plain English Introduction*. Second edition. [United States]: The author, 2017.
- Valero, Eva M., Miguel A. Martínez-Domingo, Ana B. López-Bal-domero, Ana López-Montes, David Abad-Muñoz, and José L. Vílchez-Quero. 'Unmixing and Pigment Identification Using Visible and Short-Wavelength Infrared: Reflectance vs Logarithm Reflectance Hyperspaces'. *Journal of Cultural Heritage* 64 (2023): 290–300. <https://doi.org/10.1016/j.culher.2023.10.016>.
- Verplancke, Sebastiaan. 'DAL-ART: Deep Active Learning Tool for Multimodal Image Analysis Applied to Crack Detection in Paintings.' Master thesis, University of Ghent, 2023.
- Xu, Bingjie Jenny, Yunan Wu, Pengxiao Hao, Marc Vermeulen, Alicia McGeachy, Kate Smith, Katherine Eremin, et al. 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?' *Journal of Analytical Atomic Spectrometry* 37, no. 12 (2022): 2672–82. <https://doi.org/10.1039/D2JA00246A>.

9. Bibliography Illustrations

Figure 1: Image, Aster Van Vijle, 2024.

Figure 2: GeeksforGeeks. 'Introduction to Deep Learning', 1 June 2018.
<https://www.geeksforgeeks.org/introduction-deep-learning/>.

Figure 3: Karella, Tomáš, Jan Blažek, and Jana Striová. 'Convolutional Neural Network Exploiting Pixel Surroundings to Reveal Hidden Features in Artwork NIR Reflectograms'. *Journal of Cultural Heritage* 58 (2022): 186–98. <https://doi.org/10.1016/j.culher.2022.09.022>.

Figure 4: Pu, Wei, Jun-Jie Huang, Barak Sober, Nathan Daly, Catherine Higgitt, Ingrid Daubechies, Pier Luigi Dragotti, and Miguel R. D. Rodrigues. 'Mixed X-Ray Image Separation for Artworks With Concealed Designs'. *IEEE Transactions on Image Processing* 31 (2022): 4458–73.
<https://doi.org/10.1109/TIP.2022.3185488>.

Figure 5: Image, Aster Van Vijle, 2024.

Figure 6: Kogou, Sotiria, Lynn Lee, Golnaz Shahtahmassebi, and Haida Liang. 'A New Approach to the Interpretation of XRF Spectral Imaging Data Using Neural Networks'. *X-Ray Spectrometry* 50, no. 4 (7 August 2021): 310–19. <https://doi.org/10.1002/xrs.3188>.

Figure 7: Chen, Ailin, Rui Jesus, and Marcia Vilarigues. 'Convolutional Neural Network-Based Pure Paint Pigment Identification Using Hyperspectral Images'. In *ACM Multimedia Asia*, 1–7. Gold Coast Australia: ACM, 2021. <https://doi.org/10.1145/3469877.3495641>.

Figure 8: Jones, Cerys, Nathan S. Daly, Catherine Higgitt, and Miguel R. D. Rodrigues. 'Neural Network-Based Classification of X-Ray Fluorescence Spectra of Artists' Pigments: An Approach Leveraging a Synthetic Dataset Created Using the Fundamental Parameters Method'. *Heritage Science* 10, no. 1 (2022): 88. <https://doi.org/10.1186/s40494-022-00716-3>.

Figure 9: Mandal, Dipendra J., Marius Pedersen, Sony George, Hilda Deborah, and Clotilde Boust. 'An Experiment-Based Comparative Analysis of Pigment Classification Algorithms Using Hyperspectral Imaging'. *Journal of Imaging Science and Technology* 67, no. 3 (1 May 2023): 030403-1-030403–18. <https://doi.org/10.2352/J.ImagingSci.Technol.2023.67.3.030403>.

Figure 10: Mandal, Dipendra J., Marius Pedersen, Sony George, and Clotilde Boust. 'Comparison of Pigment Classification Algorithms on Non-Flat Surfaces Using Hyperspectral Imaging'. *Journal of Imaging Science and Technology* 67, no. 5 (1 September 2023): 1–25. <https://doi.org/10.2352/J.ImagingSci.Technol.2023.67.5.050405>.

Figure 11: Rohani, Neda, Emeline Pouyet, Marc Walton, Oliver Cossairt, and Aggelos K. Katsaggelos. 'Pigment Unmixing of Hyperspectral Images of Paintings Using Deep Neural Networks'. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3217–21. Brighton, United Kingdom: IEEE, 2019. <https://doi.org/10.1109/ICASSP.2019.8682838>.

Figure 12: Pouyet, Emeline, Tsveta Miteva, Neda Rohani, and Laurence De Viguerie. 'Artificial Intelligence for Pigment Classification Task in the Short-Wave Infrared Range'. *Sensors* 21, no. 18 (13 September 2021): 6150. <https://doi.org/10.3390/s21186150>.

Figure 13: Fukumoto, Kensuke, Norimichi Tsumura, and Roy Berns. 'Estimating Concentrations of Pigments Using Encoder-Decoder Type of Neural Network.' *Color and Imaging Conference* 27, no. 1 (21 October 2019): 149–52. <https://doi.org/10.2352/issn.2169-2629.2019.27.28>.

Figure 14: Kleynhans, Tania, Catherine M. Schmidt Patterson, Kathryn A. Dooley, David W. Messinger, and John K. Delaney. 'An Alternative Approach to Mapping Pigments in Paintings with Hyperspectral Reflectance Image Cubes Using Artificial Intelligence'. *Heritage Science* 8, no. 1 (2020): 84. <https://doi.org/10.1186/s40494-020-00427-7>.

Figure 15: Chen, Ailin, Rui Jesus, and Márcia Vilarigues. 'Identification and Visualization of Pure and Mixed Paint Pigments in Heritage Artwork Using Machine Learning Algorithms'. *SN Computer Science* 4, no. 2 (21 December 2022): 115. <https://doi.org/10.1007/s42979-022-01529-8>.

Figure 16: Xu, Bingjie Jenny, Yunan Wu, Pengxiao Hao, Marc Vermeulen, Alicia McGeachy, Kate Smith, Katherine Eremin, et al. 'Can Deep Learning Assist Automatic Identification of Layered Pigments from XRF Data?' *Journal of Analytical Atomic Spectrometry* 37, no. 12 (2022): 2672–82. <https://doi.org/10.1039/D2JA00246A>.

Figure 17: Shitomi, Ryuta, Mayuka Tsuji, Yuki Fujimura, Takuya Funatomi, Yasuhiro Mukaigawa, Tetsuro Morimoto, Takeshi Oishi, Jun Takamatsu, and Katsushi Ikeuchi. 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments'. *Journal of the Optical Society of America A* 40, no. 1 (1 January 2023): 116. <https://doi.org/10.1364/JOSAA.472775>.

Figure 18: Shitomi, Ryuta, Mayuka Tsuji, Yuki Fujimura, Takuya Funatomi, Yasuhiro Mukaigawa, Tetsuro Morimoto, Takeshi Oishi, Jun Takamatsu, and Katsushi Ikeuchi. 'Unsupervised Learning with a Physics-Based Autoencoder for Estimating the Thickness and Mixing Ratio of Pigments'. *Journal of the Optical Society of America A* 40, no. 1 (1 January 2023): 116. <https://doi.org/10.1364/JOSAA.472775>.

Figure 19: Image, Aster Van Vijle, 2024.

Figure 20: Sidorov, Oleksii, and Jon Yngve Hardeberg. 'Craquelure as a Graph: Application of Image Processing and Graph Neural Networks to the Description of Fracture Patterns', 2019. <https://doi.org/10.48550/ARXIV.1905.05010>.

Figure 21: Sindel, Aline, Andreas Maier, and Vincent Christlein. 'A Multi-Modal Registration and Visualization Software Tool for Artworks Using CraquelureNet', 2022. <https://doi.org/10.48550/ARXIV.2208.08836>.

Figure 22: Nadisic, Nicolas, Yoann Arhant, Niels Vyncke, Sebastiaan Verplancke, Srdan Lazendic, and Aleksandra Pizurica. 'A Deep Active Learning Framework for Crack Detection in Digital Images of Paintings'. In *7th International Conference on Smart Monitoring, Assessment and Rehabilitation of Civil Structures, SMAR 2024*, 1–8. Elsevier, 2024. <http://hdl.handle.net/1854/LU-01J18DKNDKGSNKMK5HVB9PYS7C>.

Figure 23: Meeus, Laurens, Shaoguang Huang, Nina Zizakic, Xianghui Xie, Bart Devolder, Hélène Dubois, Maximiliaan Martens, and Aleksandra Pizurica. '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*, edited by Peter Schelkens and Tomasz Kozacki, 16. Online Only, France: SPIE, 2020. <https://doi.org/10.1117/12.2556000>.

Figure 24: Angheluta, Laurentiu Marian, and Alecsandru Chiroasca. 'Physical Degradation Detection on Artwork Surface Polychromies Using Deep Learning Models'. *Romanian Reports in Physics* 72, no. 805 (2020).

Figure 25: Liu, Yi, Fumin Wang, Zhili Jiang, Stefano Sfarra, Kaixin Liu, and Yuan Yao. 'Generative Deep Learning-Based Thermographic Inspection of Artwork'. *Sensors* 23, no. 14 (13 July 2023): 6362. <https://doi.org/10.3390/s23146362>.

Figure 26: Moradi, Morteza, Ramin Ghorbani, Stefano Sfarra, David M.J. Tax, and Dimitrios Zarouchas. 'A Spatiotemporal Deep Neural Network Useful for Defect Identification and Reconstruction of Artworks Using Infrared Thermography'. *Sensors* 22, no. 23 (1 December 2022): 9361. <https://doi.org/10.3390/s22239361>.

Figure 27: Image, Aster Van Vijle, 2024.

Figure 28: Califano, America, Pietro Foti, Filippo Berto, Marco Baiesi, and Chiara Bertolin. 'Predicting Damage Evolution in Panel Paintings with Machine Learning'. *Procedia Structural Integrity* 41 (2022): 145–57. <https://doi.org/10.1016/j.prostr.2022.05.017>.

Figure 29: Image, Aster Van Vijle, 2024.