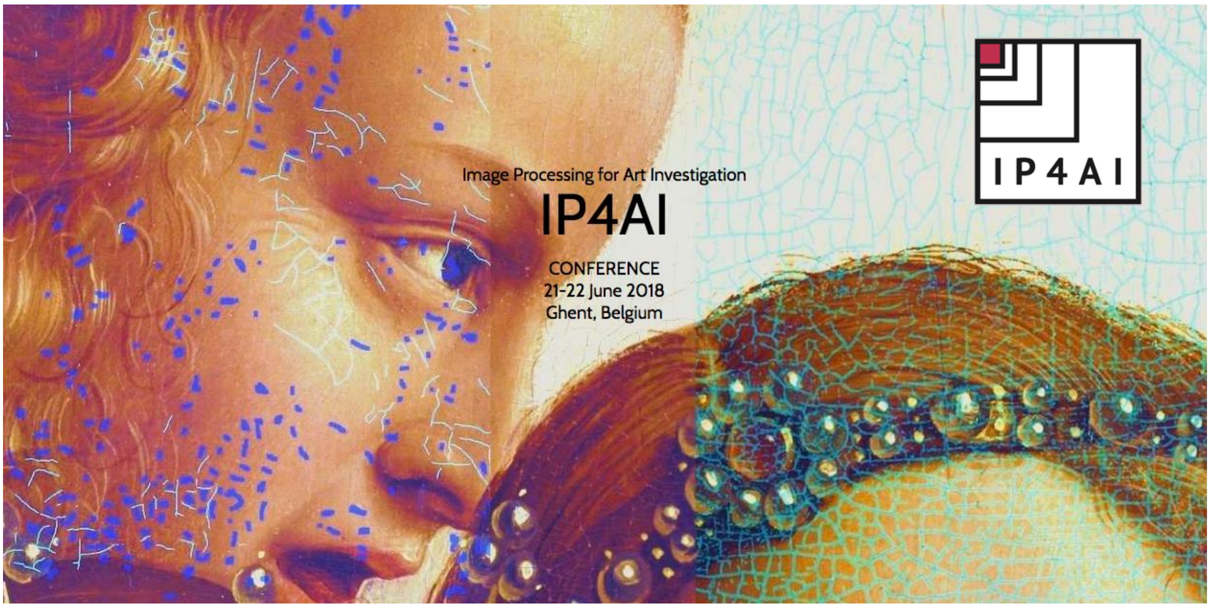

PROCEEDINGS

of the sixth International workshop on Image Processing for Art Investigation

IP4AI 2018

www.ip4ai.ugent.be

21-22 June 2018

Museum of Fine Arts (MSK)

Ghent, Belgium

GHENT UNIVERSITY

Supported by Ghent University

Table of Contents

IP4AI 2018 Organizing Committee

Aleksandra Pižurica, Ghent University
Ljiljana Platiša, imec-IPI, Ghent University
Maximiliaan Martens, Ghent University
Bruno Cornelis, Vrije Universiteit Brussel
Nikos Deligiannis, Vrije Universiteit Brussel
Miguel Rodrigues, University College London
Ingrid Daubechies, Duke University

Scientific Committee

Patrice Abry, École Normale Supérieure De Lyon, CNRS
Laura Alba Carcelén, Museo del Prado, Madrid
Bruno Cornelis, Vrije Universiteit Brussel
Ingrid Daubechies, Duke University
Marc De Mey, Royal Flemish Academy of Belgium (KVAB)
Nikos Deligiannis, Vrije Universiteit Brussel
Ann Dooms, Vrije Universiteit Brussel
Ella Hendriks, University of Amsterdam
Catherine Higgitt, The National Gallery, London
Stephane Jaffard, University Paris-Est Créteil (UPEC)
Maximiliaan Martens, Ghent University
Aleksandra Pižurica, Ghent University
Ljiljana Platiša, imec-IPI, Ghent University
Eric Postma,Tilburg University
Miguel Rodrigues, University College London
Carola-Bibiane Schönlieb, University of Cambridge

Image Processing for Art Investigation Workshop 2018 Program
June 21-22, 2018
Auditorium of the Museum of Fine Arts (MSK) in Ghent
Fernand Scribedreef 1, 9000 Ghent, Belgium

Thursday June 21, 2018

Start time	*End time*		*Authors*	*Title*
8:30	9:00	**Registration**		
9:00	9:10	**Conference welcome and logistics**		
		Plenary		**Chair: I. Daubechies**
9:10	10:00		John Delaney	Multi-modal imaging spectroscopy of paintings and illuminated manuscripts
		Session 1	**Combining different imaging modalities to uncover hidden features in paintings**	**Chair: A. Pizurica**
10:00	10:20		Zahra Sabetsarvestani, Francesco Renna, Franz Kiraly, and Miguel Rodrigues	Un-mixing X-ray images with the application in art investigation
10:20	10:40		Jan Blažek, and Barbara Zitová	Information separation in art investigation: A survey
10:40	11:00		Shaoguang Huang, Bruno Cornelis, Bart Devolder, and Aleksandra Pizurica	Paint loss detection via kernel sparse representation
11:00	11:20	Coffee break		
		Session 2	**Image processing for Photography**	**Chair: S. Jaffard**
11:20	11:50		Patrice Abry	Anisotropic multiscale representations for an automated and reproducible analysis and classification of photographic paper
11:50	12:20		Henri Maître	Automatic appreciation of aesthetics in photography: Where are we going?
12:20	13:40	Lunch		
13:40	13:50	Group picture		
		Session 3	**Poster pitches**	**Chair: L. Platisa**
13:50	13:55		Marie d'Autume, and Enric Meinhardt-Llopis	Disrobing Adam and Eve with the linear osmosis model
13:55	14:00		Ana Martins	"NO CHAOS, DAMN IT!" Extracting paint maps from Macro-X-Ray Fluorescence scanning data to deconstruct Jackson Pollock's "action painting" in Number 1A, 1948
14:00	14:05		Rasha Ahmed Shaheen, Mona Fouad Ali, Medhat El-Dabaa	Documentation and digitalizing of royal albumen photographic collection of King Farouk, dating from the 19th Century
14:05	14:10		Gjorgji Strezoski, and Marcel Worring	OmniArt: A large scale artistic benchmark
		Session 4	**Imaging tools for artwork diagnostics**	**Chair: M. Martens**
14:10	14:40		Hélène Dubois	The conservation of the brothers van Eyck's Ghent Altarpiece
14:40	15:10		Ella Hendriks	Transforming conservation
		Session 5	**New initiatives and collaborations**	**Chair: M. Rodrigues**
15:10	15:40		Daniele Zavagno	The Visual Science of Art conference: History and aims
15:40	16:00		Collaborative projects	
			Massimo Fornasier	Mantegna Project 4.0: Some novel directions over the work done in the Mantegna Project
			Miguel Rodrigues	Transforming art study, conservation, preservation, and presentation via digital technology
16:00		Coffee		
16:00	17:00	**Session 6**	**Posters & Demonstrators**	**Chair: L. Platisa**
			Poster presentations	*Corresponding to the pitch presentations from Session 3*
			Demonstrators	
			Laurens Meeus	Fast and accurate paint loss detection using deep learning
			Gjorgji Strezoski	ArtSight: A visual artistic data exploration engine
			Roman Sizyakin	Detection of cracks in paintings using deep learning
17:00	17:30	Visit to the Ghent Altarpiece restoration studio at the MSK		
			Evening program	
17:30	19:00	Boat & walking tour through the heart of Ghent		
19:00	22:00	Conference dinner		
		Evening presentation by Ingrid Daubechies: Reunited: a digital and art-historical adventure		

Friday June 22, 2018			
Start time	*End time*	*Authors*	*Title*
	Session 7	**Imaging as a tool in understanding creation of the artwork**	**Chair: P. Abry**
9:00	9:30	Koen Janssens	Paintings' alternations in the past and in the future: Non-invasive X-Ray based imaging of subsurface information and how to improve upon it
9:30	9:50	Barak Sober	Iron age Hebrew epigraphy in the silicon age - An algorithmic approach to study Paleo-Hebrew inscriptions
9:50	10:10	R. van Liere, K.J. Batenburg, A. Kostenko, C-L. Wang, and I. Garachon	Imaging ancient Chinese ivory puzzle balls: deducing the make Process
10:10	10:30	Ellen van Bork, and Bruno Cornelis	The Monitoring of Cracks in Historical Silver with Image Processing Techniques
10:30	10:50	Coffee break	
	Session 8	**Frescoes**	**Chair: E. Hendriks**
10:50	11:20	Dubravka Đukanović and Siniša Zeković	The renovation and protection of the rich cultural and artistic heritage of the Hilandar monastery
11:20	11:50	Massimo Fornasier	The Mategna frescoes in Padua: Computer assisted puzzle solving and recolorization
11:50	12:10	Simone Parisotto, Luca Calatroni, and Claudia Daffara	Mathematical osmosis imaging for multi-modal and multi-spectral applications in cultural heritage conservation
12:10	13:30	Lunch	
	Session 9	**Virtual restoration**	**Chair: M. Fornasier**
13:30	14:00	Carola-Bibiane Schönlieb	Unveiling the invisible - mathematical approaches for virtual image restoration
14:00	14:20	Nanne van Noord and Eric Postma	Pixel context encoders for painting region inpainting
	Session 10	**Deep learning techniques for digital artwork analysis**	**Chair: K. Janssens**
14:20	14:40	Laurens Meeus, Shaoguang Huang, Bart Devolder, Maximiliaan Martens, and Aleksandra Pizurica	Deep learning for paint loss detection: A case study on the Ghent Altarpiece
14:40	15:00	Roman Sizyakin, Bruno Cornelis, Laurens Meeus, Maximiliaan Martens, Viacheslav Voronin, and Aleksandra Pizurica	A Deep Learning Approach to Crack Detection in Panel Paintings
15:00	15:30	Eric Postma	Improving the reliability of CNNs for digital artwork analysis
15:30	15:40	**Closing remarks**	

John K. Delaney is senior imaging scientist in the scientific research department of the conservation division of the National Gallery of Art, Washington. His research focuses on the adaptation of remote sensing sensors and processing methods for the study of paintings and works on paper. Before joining the Gallery, Delaney was chief scientist for the Advanced Sensors Business Unit of the ISR Airborne Systems Division of the Goodrich Corporation. He received his BS from the Worcester Polytechnic Institute and his PhD from the Rockefeller University, and he completed postdoctoral studies at the University of Arizona and the Johns Hopkins University School of Medicine. Delaney has published more than 75 papers in the areas of imaging and spectroscopy.

When Delaney joined the Gallery in 2007 thanks to funding provided by the Andrew W. Mellon Foundation, the Gallery became the first art museum to have an imaging scientist on staff. In his role as senior imaging scientist, Delaney was tasked with developing and adapting remote sensing imaging cameras and methods and using those advanced digital imaging methods to obtain new information that could be used for conservation and art historical research. Following the end of the Mellon grant–funded position, Delaney permanently joined the Gallery in 2011. In the years since, Delaney and colleagues in the scientific research department have optimized hyperspectral visible and infrared imaging cameras and image processing to better visualize painted-over compositions and map the distribution of pigments. Additional funding from the Mellon Foundation, Samuel H. Kress Foundation, and National Science Foundation has supported the ongoing training by Delaney of fellows in new advanced imaging methods.

In addition to his work at the Gallery, Delaney also served as a research professor in the Department of Electrical and Computer Engineering of the School of Engineering and Applied Science at the George Washington University from 2011 to 2013. He is also an associate editor of *Heritage Science* and was previously an associate editor of *Studies in Conservation* from 2010 to 2017.

Multi-modal imaging spectroscopy of paintings and illuminated manuscripts
John K. Delaney, and Kathryn A. Dooley

The use of non-invasive point measurements with multiple modalities has been shown to yield more complete identification of artist materials' in situ. For example, the combination of x-ray fluorescence elemental analysis combined with molecular reflectance and fluorescence spectroscopy can provide information on colorants (both inorganic and organic lake pigments) as well as paint binders. Extending this to collect 2-D spatial maps with each modality offers unique opportunities to mathematically combine, or fuse, the information derived from all the image cubes to yield material maps with higher confidence, as well as provide new image products. In this talk we will describe the instrumentation used to acquire such 3-D data sets and show examples of information that can be derived from them. The multi-modal imaging system is able to provide XRF image cubes, diffuse reflectance image cubes (400 to 2500 nm), as well as molecular luminescence image cubes (400 to 1000 nm). Using a novel image registration program developed with George Washington University, these image cubes can be spatially aligned with the color image, x-ray radiograph, etc. The case studies will highlight the potential to mine these data sets for new information. For example, an improved model for 'virtual removal of aged varnish' will be presented, which was derived from visible reflectance image cubes collected before and after the removal of an aged varnish on a painting by Georges Seurat titled Haymakers at Montfermeil, c. 1882. Other examples will include improved visualization of prior paintings which have been painted over, such as the Spanish woman beneath Pablo Picasso's Blue Period painting Le Gourmet, c. 1901, and earlier drawn and painted features found in Andrea del Sarto's painting Charity, c. 1528 or 1529 that relate to another of his paintings. Improved material maps obtained by fusion of XRF elemental maps with pigment maps obtained from reflectance imaging spectroscopy of Pacino di Bonaguida's illuminated manuscript titled Christ in Majesty with Twelve Apostles, c. 1320 will also be shown. Finally, all three imaging modalities were used to understand the artist's materials and methods used to create an encaustic Greco-Roman portrait painting from the 2nd century AD Egypt. The increasing availability of such image cubes collected in several modalities offers new opportunities for the development of new processing algorithms which will be expected to yield new information about important cultural objects.

Patrice Abry is currently «Senior Scientist» (Directeur de Recherche) for the French National Center of Scientific Research (CNRS) at Ecole Normale Supérieure de Lyon, France, where he is in charge of the « Signal, System and Physics » statistical signal processing research group, within the Physics department. He received the degree of Professeur-Agrégé de Sciences Physiques, in 1989, at Ecole Normale Supérieure de Cachan and the Ph. D. degree in physics and signal processing from the Claude-Bernard University, Lyon, France, in 1994.

Patrice Abry has developed a long standing research program dedicated to the multiscale statistical analysis and modelling of scale-free phenomena. His researches are motivated by strong interests in integrating theoretical and applied developments, with real-world applications including, hydrodynamic turbulence, Internet traffic, heart rate variability, neurosciences.

He is the author of a book on wavelet, scale invariance and hydrodynamic turbulence and is also the co-editor of a book entitled Scaling, Fractals and Wavelets.

Dr. Abry received the AFCET-MESR-CNRS prize for best Ph.D. in signal processing for years 1993–1994. He serves in the IEEE SPS Signal Processing Theory and Methods Committee since 2014 as well as in the French National Committee for Scientific Research (CoNRS) since 2017. He was also elected IEEE fellow in 2011.

Anisotropic multiscale representations for an automated and reproducible analysis and classification of photographic paper

Patrice Abry (CNRS DR, Laboratoire de Physique, ENS de Lyon), Stéphane G. Roux, Nicolas Tremblay, Pierre Borgnat (Laboratoire de Physique, ENS de Lyon, France), Stéphane Jaffard (LAMA, Université Paris-Est Créteil, France), Herwig Wendt (IRIT, Université de Toulouse, France), Béatrice Vedel (LMBA, Université de Bretagne Sud, France), Andrew G. Klein (Western Washington University, USA), C. Richard Johnson (Cornell University, USA), Paul Messier (Yale University, USA), Jim Coddington and Lee Ann Daffner (MoMA, USA)

Surface texture is a critical feature in the manufacture, marketing, and use of photographic paper. Hence, texture characterization of photographic prints can provide scholars with valuable information regarding photographers' aesthetic intentions and working practices. Currently, texture assessment is strictly based on the visual acuity of a range of scholars associated with collecting institutions, such as conservators. Natural interindividual discrepancies, intraindividual variability, and the large size of collections present a pressing need for computerized and automated solutions for the texture characterization and classification of photographic prints. Recently, an automated and digital raking light procedure has been designed (by P. Messier et al.) that reveals texture through a stark rendering of highlights and shadows.

The present work aims to provide evidence that the combination of this automatic, computer-based raking light based measure of texture with advanced anisotropic multiscale image processing representations permits to achieve relevant characterization and classification of photographic paper textures. The intuition behind anisotropic multiscale representation, originally developed for measuring rugosity, or irregularities, in physics and biomedical applications, consists in analyzing a texture across a collection of views at different scales, or resolutions and relies on a change of paradigm: The information is not extracted in what is seen at each scale, but rather in how what is seen changes when scales vary.

In this work, this recent statistical image processing tool is customized for and applied to two different photographic paper datasets. For proof of concept, it is first applied to a small-size reference data set of historic (silver gelatin, 120 prints) photographic papers that yet combines in purpose several levels of similarity. Second, it is used on a large data set (2491 prints) of culturally valuable photographic prints held by the Museum of Modern Art in New York. The promising results achieved with this fully automatized and non-supervised procedure for the characterization and clustering of photographic paper are interpreted in collaboration with art scholars with an aim toward developing new modes of art historical research and humanities-based collaboration.

Hélène Dubois is a painting conservator and art historian. She trained at the Université Libre de Bruxelles and at the Hamilton Kerr Institute, University of Cambridge (UK). She worked at the Doerner-Institut in Munich, the J. Paul Getty Museum in Malibu, the Royal Museums of Fine Arts in Brussel and the Limburg Conservation Institute in Maastricht where she taught conservation of Old Masters Paintings. Attached to the Royal Institute for Cultural Heritage in Brussels (KIK-IRPA) and to Ghent University (UGent), she leads the conservation project of the brothers van Eyck's *Adoration of the Mystic Lamb* since October 2016 and is researching the material history of the altarpiece for a PhD dissertation at Ghent University.

The Conservation of the Brothers van Eyck's Ghent Altarpiece
Hélène Dubois

Since 2012, the Royal Institute for Cultural Heritage of Belgium has been in charge of the conservation treatment of the brothers van Eyck's altarpiece of the Adoration of the Mystic Lamb (1432). This monumental masterpiece has been admired for centuries for its stunningly convincing suggestion of space and light, materials and expressions, reflecting extraordinary sensibility and level of technical accomplishment that determined the course of western European painting. The tumultuous material history of the altarpiece, marked by wars, civil unrest, fires, changes in tastes and values has deeply influenced its condition and its appearance. The conservation campaign is carried out in public view by a team of conservators in the Ghent Museum of Fine Arts. The discoveries carried through multidisciplinary research involving the conservators, chemical analysis, technical imagery, archival and art historical investigations reveal completely unknown, yet fundamental facets of the altarpiece.

This paper will illustrate the discovery of early overpainting campaigns, the diagnostic of the condition the original paint layers and the recovery of the original.

Dubravka Đukanović, associated professor at the Academy of Arts in Novi Sad, University of Novi Sad, Master Academic Study Program *Conservation and Restoration of works of fine and applied arts*. Research fellow of the Institute of Architecture and Urban & Spatial Planning of Serbia. Architect with more than twenty years of experience in designing and conservation. Owner and leading designer of Studio D'ART. Senior Expert Architect in the HTSPE-Eurotrends Expert Team for the implementation of the EU project for the rehabilitation of the complex of the Franciscan Monastery in Bac. Author of a number of papers in the national and international professional and academic journals, co-author of several publications and author of two monographs – *Serbian Orthodox Churches of the XVIII and XIX Centuries in Backa* (2009) and *Architecture of Roman Catholic Churches in Vojvodina from 1699 to 1939* (2015). Winner of the annual "Ranko Radović Award" for 2010 in the field of the theoretical texts as well as several prizes at national competitions in the field of architectural and urban planning. Member of the Chamber of Architects of Serbia, DANS, Society of Conservators of Serbia, IIC and Serbian National Committee of ICOMOS.

The renovation and protection of the rich cultural and artistic heritage of the Hilandar Monastery

Dubravka Đukanović (Academy of Arts in Novi Sad, University of Novi Sad, Serbia) and Siniša Zeković (The Provincial Institute for the Protection of Cultural Monuments, Petrovaradin, Serbia)

For contemporary art conservation practice, the renovation and protection of the rich cultural and artistic heritage of the Hilandar Monastery is a great and complex challenge. The access to the material is granted only under specific conditions and the works can only be performed in the intervals between daily rituals. Due to the specific circumstances and inability to take the objects outside the territory of Mount Athos, the application of modern research, and especially conservation techniques, was limited to in situ examination, using mobile equipment of small dimensions, limited analysis of movable samples for laboratory diagnostics and techniques of conservation-restoration work applicable in the conditions of fieldwork.

Apart from conservation works on numerous individual items, the experts from the Provincial Institute have completed three total reconstructions of iconostases since 2004, while a team of experts from the Provincial Institute has been commissioned since 2015 on the protection of decorations and most valuable sacred items, kept in the main facilities of the monastery – the church of the Entrance of the Blessed Virgin Mary into the Temple (1321) and the Grand Dining Room (XII-XIII). The multiannual conservation works on the throne featuring the icons of the Virgin with Three Hands, St. Nicholas and Holy Three Hierarchs, which cover the southwest stone pillar started in 2017. It the first phase, the third zone of engraved, gilded elements was conserved and restored. The degree of damage of the icon on the east throne – the Virgin with Three Hands (XIV) was established with visual inspection, while a more detailed analysis is planned for the forthcoming period.

The analysis of hidden layers may bring new findings about historical facts and events on the territory of medieval Serbia. One such object is King Milutin's Charter, written in 1324. The text of the charter, written on a parchment, was "repainted" by Bulgarian monks from Mount Athos at the beginning of the 19th century. It is still not known whether the new base contains a copy of the original text or whether the content of the text was changed. The original letters and signs are visible on the parts of the parchment from which the newly-applied layer has fallen off. Determining the content of the original text would be of great art-historical and iconographical value. The ultimate goal of our work is to determine adequate methodology that would allow the conservators to implement the methods of contemporary conservation in the specific working conditions in the monastery.

The research of **Massimo Fornasier** embraces a spectrum of problems in mathematical modeling, analysis and numerical analysis. Fornasier is particularly interested in the concept of compression as appearing in different forms in data analysis, image and signal processing, and in the adaptive numerical solutions of partial differential equations or high-dimensional optimization problems.

Fornasier received his doctoral degree in computational mathematics in 2003 from the University of Padua, Italy. There he worked also for the realization of the Mantegna Project, i.e., the complete restoration of the Mantegna's frescoes in the Eremitani Church in Padua, which were destroyed by a bombing in World War II.

After spending from 2003 to 2006 as a postdoctoral research fellow at the University of Vienna and University of Rome "La Sapienza", he joined the Johann Radon Institute for Computational and Applied Mathematics (RICAM) of the Austrian Academy of Sciences where he served as a senior research scientist until March 2011. He was an associate researcher from 2006 to 2007 for the Program in Applied and Computational Mathematics of Princeton University, USA. In 2011 Fornasier was appointed Chair of Applied Numerical Analysis at the Technical University of Munich.

The Mategna frescoes in Padua: computer assisted puzzle solving and recolorization
Massimo Fornasier

In 1944, near the end of World War II, an allied bombing campaign destroyed the Eremitani church in Padua, Italy. The church was famous among art lovers for its magnificent frescoes, which included a series by the early renaissance painter Andrea Mantegna (1431-1506). Over 88.000 small pieces of painted plaster, of an average area of only 4-5 square centimeters, had been lovingly collected and conserved after the bombing; together, they accounted for less than 80 square meters – only a very small fraction of the area covered by the frescoes originally. From 1992 onwards, art conservation experts attacked the task of cleaning and photographing every piece, sorting them and hoping to reconstruct at least some fragments. The herculean task seemed hopeless – until mathematics came to the rescue. We developed an approach that made it possible, for each small piece of plaster that still showed an element of the design of the fresco, to find where it belonged exactly. The resulting very fragmented and mosaic-like reconstruction of the color scheme of each fresco was then used, via another algorithm, to fill in the color information for the whole fresco.

Ella Hendriks is full Professor of Conservation and Restoration of Moveable Cultural Heritage at the University of Amsterdam, The Netherlands).

From 1999 to 2016 she was Senior Paintings Conservator at the Van Gogh Museum in Amsterdam and from 1987 to 1999 Head of Conservation at the Frans Hals Museum in Haarlem. In 2007 she collaborated with Prof. C. Richard Johnson, Jr. (Cornell University) to organize the first and second IP4AI conferences held at the Van Gogh Museum in Amsterdam and Museum of Modern Art, New York. Since then she has been fortunate to collaborate with a broad range of experts in the field of image processing for projects ranging from the Thread Count Automation Project (TCAP) established in 2007, to *ReViGo* (Reassessing Vincent van Gogh) within the Science4arts programme supported by NWO (2012-2017).

Transforming conservation
Ella Hendriks

This presentation broadly considers the ways in which image processing is transforming conservation and restoration of cultural heritage in terms of both methods and approach. At the University of Amsterdam conservation and restoration training programme, past years have witnessed an increased demand for image processing expertise involved in research and practice across the nine different tracks of specialisation. This together with the speaker's own projects in the field of paintings conservation, provides a representative picture of the types of problem that can benefit from collaboration with image processing specialists (in the broadest sense of the word). While by no means an exhaustive survey, it provides an opportunity to stand still and evaluate this development from a user's perspective in the context of this sixth IP4AI conference.

A primary task of the modern-day conservator is to manage undesirable processes of change in cultural heritage objects. A limitation is that we can only analyse the present, while to do this we also need to understand the past and to predict the future. Increasingly, image processing plays an important role in all three related areas of activity, ranging from diagnostics to preventive conservation and treatment. These concepts will be used as a convenient way to structure this talk, in which case study examples of image processing application will be grouped according to these three, overlapping areas of conservation practice.

Diagnostics. Image processing-based methods have greatly expanded the range of diagnostic tools available to the conservator. One example is automated mapping of thread count and thread angle in woven fabrics, such as canvases used in paintings. This combined information is useful to answer art historical questions relating to dating, provenance or attribution, but has also proved informative to the conservator as it helps to distinguish features belonging to the original object from later damages and alterations, as examples will demonstrate.

Combined optical and chemical mapping of art works is especially helpful to support process-based analysis in conservation. On the one hand it can help to determine whether certain degradation processes have stabilized or are still ongoing, information that is crucial to the conservator's assessment of condition (an example of metal soap aggregation will be given). On the other hand it enables the process of conservation treatment to be monitored, in order to evaluate its effects as a basis for adaptive decision-making. For example, while still at the developmental stage, a combination of optical coherence tomography (OCT) and reflection mid-Fourier Transfer Infrared (mid-FTIR) proves highly promising to visualize the progressive effects of varnish removal on a painting, mapping both the area and thickness of the layer(s) removed.

Preventive conservation. A combination of non-invasive, micro-analytical scanning techniques can also provide a 'risk' mapping of e.g. light sensitive colour areas in paintings (see paper Koen Janssens et al. in this conference). Potentially this opens the way for selective lighting of different parts of a painting, or of paintings that belong to different risk category groups in a collection, according to their level of vulnerability. A recent development has been to model *future* states of discoloration in digital visualizations that make the problem tangible. In effect, this has taken research on pigment deterioration out of the laboratory and into the stakeholder's office, where the visualizations have been used as a basis for discussions on what constitutes acceptable damage and deciding an appropriate lighting policy (example Van Gogh Museum).

Restorations. In the past, a common goal of 'restoration' was to return objects to a past, or perceived original state. Nowadays we may question the validity of this approach, while the irreversible removal or addition of material required to achieve this goal may no longer be considered an ethically viable option. Image processing methods can provide an alternative by creating visualizations of former states of an object, without changing the object itself (therefore offering a solution for objects that are too fragile to withstand handling and manipulation). These visualizations can take on various forms- from computer screen images to physical or virtual 3D replicas of the object- and allow different options for restoration to be explored. Also, light can be projected onto the object itself to perform a 'virtual restoration', reversing effects of colour change on the original. Case examples considered in this talk will include 3-D reconstructions of (maritime-) archaeological objects, virtual varnish removal from paintings, reversal of colour change in painted surfaces using digital visualisations, and light retouching of furniture.

In those instances when it is decided to physically restore objects, incorporating image processing methods can help to improve the efficiency, cost and ethics of the restoration procedure, as well as benefiting the result. Examples for this talk will include the use of computer generated recipes for non-metameric inpainting of losses, and the use of 3-D scanning and casting/printing technologies for loss compensation. It should be noted that while new computer technologies may play an important role in treatment, the traditional skills and judgement of the trained conservator remain of paramount importance for the quality of the result. The new technologies may be seen to complement existing methods, extending the range of possibilities at the conservator's disposal.

Koen Janssens is full professor of general and analytical chemistry at the University of Antwerp in Belgium. He obtained his PhD in 1989 on a thesis dealing with the use of Artificial Intelligence techniques for automated treatment of X-ray analysis data. Since then, he has been actively making use of strongly focused X-ray micro- and nano-beams, produced in large accelerator complexes called Synchrotron Storage Rings, for non-destructive materials analysis. Such beams are useful to gain information on the distribution and speciation state of (heavy) metals in polluted natural materials such as soils, sediments and airborne particulates and in industrial materials such as heterogeneous catalysts. A combination of X-ray fluorescence spectrometry, X-ray absorption spectroscopy and X-ray diffraction usually is employed to characterize these materials in 2D or 3D imaging mode. He applies the same suite of techniques for better understanding naturally occurring alteration and degradation processes in cultural heritage materials such as historic glass, inks and painters' pigments. In recent years, these investigations on the micrometre scale were augmented with macroscopic imaging, performed using mobile scanning equipment by means of millimetre-sized X-ray beams. Such imagery has proven to be useful for art historians and art conservators in order to understand better both the past and future of works of art.K. Janssens is (co)author of ca 240 scientific papers and has served as (co)editor of four scientific books, dealing with non-destructive analysis in the cultural heritage area. He has organised several conferences on analytical chemistry, X-ray micro beam analysis and its applications. He is currently vice-dean of the Faculty of Science of the University of Antwerp. Since 2016, he was appointed 'Senior Scientist' (hon.) at the Rijksmuseum, Amsterdam, The Netherlands.

Paintings' alternations in the past and in the future: non-invasive x-ray based imaging of subsurface information and how to improve upon it

K. Janssens, G. Van der Snickt, S. Legrand, F. Vanmeert, S. De Meyer

In the last decade we (and other groups) have developed and applied the method of Macroscopic X-ray fluorescence (MA-XRF) imaging to ca 100-150 oil paintings in diverse musea of fine arts in Europe and the US. This series includes works by famous 15-19th century artists Flemish and Dutch artists such as Van Eyck, Memling, Rubens, Van Dyck, Rembrandt, Magritte and Van Gogh. Also stained glass windows, illuminated manuscripts and artistic drawings can be examined by means of MA-XRF to yield information that is relevant for different purposes, be it authentication, art-historical study or conservation.

MA-XRF shares several useful characteristics with commonly employed methods for paintings' inspection such as X-ray radiography (XRR), Infra-red reflectograhy (IRR) and the more recently developed hyperspectral imaging methods that make use of camera's sensitive to parts of the ultraviolet, infrared and visual wavelength spectrum:

(i) While the X-rays employed energize the material in the irradiated spot (< 1 mm) on the artwork during a brief (< 1 s) period, no (discernable) damage results;

(ii) The penetrative x-rays allow to visualize 'hidden' layers in altered paintings that, for various reasons, were covered up by paint strata during/after the artwork creation

MA-XRF differs from XRR and IRR in three important ways:

(a) it is a scanning method, involving stepwise irradiation and spectral data recording, thus taking (substantially) longer than XRR, IRR and more modern full-field imaging methods, to record relevant imaging data, typically taking 1-10 h per m2

(b) it provides multiple elemental images of the investigated painting areas, allowing (sometimes only approximate) identification of the inorganic pigments that were employed and their distribution at the brushstroke level; the largely orthogonal nature of the data usually permits a straightforward interpretation of the data cube.

(c) its information depth depends on the energy of the X-ray fluorescence radiation employed, varying between a few to hundreds of micrometers, which is in a suitable range to allow superficial and subsurface visualization on the projected distribution of various elemental constituents while the less relevant structure of the ground layer/substrate panels usually is not or only vaguely observed. The strong and weak points of MA-XRF (and its more recent extensions) for visualization of different types of paintings alterations will be illustrated by examples involving artworks by the above-mentioned painters; in each case, the benefits that might be gained by (automated) post-processing of the data, e.g. towards noise reduction, feature extraction and/or visualization of anomalous patterns will be addressed.

Henri Maître is an emeritus professor at Telecom ParisTech. He has taught there digital picture processing, was Head of Department and Deputy Director for Research.

His research included works on digital holography, image analysis, image understanding and computer vision, with applications in the domains of medical, satellite and fine arts image processing. Present work concerns the relation between Aesthetics and Photography.

Automatic Appreciation of Aesthetics in Photography: Where are we going?
Henri Maître

Under the impulse of machine learning techniques, digital aesthetics assessment received a renewed interest in recent years. In the last 3 years, deep neural networks outclassed hand-crafted feature methods based on image processing and classification. Since then, a handful of studies claim their capacity to separate nice images from the run of the mill production and exhibit scores of almost 80 % agreement with human experts. But what do these methods measure? Implicitly they are based on the "objectivist" tradition of aesthetics dating back to the Greek philosophers, and highly influential on the artistic field up to the 18th century. However, the "subjectivist" point of view, as pioneered by Locke and Burke gained in popularity in the 19th and 20th centuries. Rested on the psychoanalytic school, then by experimental psychology and social studies, and at last in recent days by neuro-biology (and the so called "neuro-aesthetic" trend), the "subjectivist" school gained in support in the scientific community, ... but not in the image-processing and artificial intelligence body! We will show how the history of "scientific beauty evaluation" since the early works of C. Henry (1885) and G. Birkhoff (1933) until DNN is indeed following an identical slope where only few attention is paid to the viewer when most of the literature on aesthetics tells us that other tracks may be more valuable.

Eric Postma is a professor in Artificial Intelligence at the Cognitive Science & AI department at Tilburg University and at the Jheronimus Academy of Data Science in ‘s-Hertogenbosch.

He received his M.Sc. in 1989 at the University of Nijmegen; his thesis about a connectionist model of implicit and explicit memory was largely based on an internship at Leiden University. His Ph.D. in 1994 at Maastricht University concerned a biologically inspired model of covert attention and it served as an inspiration for his current research, which focusses on the use of data science (machine learning) in image recognition and cognitive modelling.

In 2008 Postma and his team ranked second in the annual “Academische Jaarprijs” with a presentation on the breakthroughs of digital painting analysis. Together with Laurens van der Maaten, he received the AAAI-08 Most Innovative Video Award for a scientific video of the digital painting analysis. Currently, Postma is coordinating the REVIGO project. Over the years, his main interest remained human perception and cognition, and the modelling thereof with AI techniques. Professor Postma (co-)supervised some 70 M.Sc. and 20 Ph.D. students. He has published in numerous cognitive science and AI journals.

Improving the reliability of CNNs for digital artwork analysis
Eric Postma

The use of convolutional neural networks (CNNs) for the digital analysis of artworks is now commonplace. Although CNNs achieve superior performance in a wide variety of visual tasks, their internal models are fallible which can give rise to erroneous predictions. CNNs misclassify adversarial instances (i.e., dataset instances that are intentionally distorted by small perturbations), which indicates that their successful performances on visual tasks are based on visual cues that are inferior to those of humans. Clearly, this limitation of CNNs hampers their applicability on the domain of digital artwork analysis. The presentation reports on our attempts to encourage CNNs to employ human-like visual cues. We show that more human-like cues give rise to improved predictions and strengthen the ability to deal with adversarial instances. With the improvement, CNNs become more suitable to support in the analysis of artworks.

Carola-Bibiane Schönlieb is a Reader in Applied and Computational Mathematics at the Department of Applied Mathematics and Theoretical Physics at the University of Cambridge (UK). There she heads the Cambridge Image Analysis Group, is Director of the Cantab Capital Institute for the Mathematics of Information and of the EPSRC Centre for Mathematical Imaging in Healthcare, Fellow of Jesus College, Cambridge and a Faculty Fellow of the Alan Turing Institute.

Carola graduated from the Institute for Mathematics, University of Salzburg (Austria) in 2004. From 2004 to 2005 she held a teaching position in Salzburg. She received her PhD degree from the University of Cambridge in 2009. After one year of postdoctoral activity at the University of Göttingen (Germany), she became a Lecturer in at DAMTP in 2010, promoted to Reader in 2015.

In her research, she is interested in the interaction of mathematical sciences and imaging. She studies non-smooth and possibly non-convex variational methods and nonlinear partial differential equations for image analysis and inverse imaging problems, among them image reconstruction and restoration, object segmentation, and dynamic image reconstruction and analysis such as fast flow imaging, object tracking and motion analysis in videos. Moreover, she works on computational methods for large-scale and high-dimensional problems appearing in, e.g. image classification and 3D and 4D imaging.

Unveiling the invisible - mathematical approaches for virtual image restoration
Carola-Bibiane Schönlieb

In this talk I will discuss mathematical approaches based on partial differential equations and variational models for the virtual restoration of paintings and illuminated manuscripts. The latter in particular provide an interesting opportunity for digital manipulation because they traditionally remain physically untouched. Showcasing restoration examples we have derived in collaboration with the Fitzwilliam Museum in Cambridge, I will also explain the main mechanisms behind the mathematical methods used.

My presentation will include joint works with Spike Bucklow (Hamilton Kerr Institute, Cambridge, UK), Luca Calatroni (Ecole Polytechnique, Paris, Franca), Marie D’Autume (ENS Cachan, Paris, Franca), Rob Hocking (Faculty of Mathematics, Cambridge, UK), Stella Panayotova (Fitzwilliam Museum, Cambridge, UK), Paola Ricciardi (Fitzwilliam Museum, Cambridge, UK) and Simone Parisotto (Faculty of Mathematics, Cambridge, UK).

Daniele Zavagno has a master degree in Art Conservation from the University of Udine and a Ph.D. in Experimental Psychology from the University of Padoa (Italy). He worked as a postdoc at the University of Padoa, as research associate at NEC Research Institute (now NEC Laboratories America) in Princeton and at the University of Maryland, College Park (USA). He was visiting scientist at the University of Nagoya and visiting faculty at Tohoku Gakuin University (Japan). He is currently associate professor at the University of Milano-Bicocca.

The Visual Science of Art Conference: History and aims

Daniele Zavagno and Rossana Actis-Grosso

The Visual Science of Art Conference was born in 2012 when Baingio Pinna organized a satellite meeting to the 36 th European Conference on Visual Perception (ECVP) held in Alghero, Italy. Since then the conference has continued to be organized as a satellite conference to ECVP, though the two sister events are sometimes organized by different people, though in the same city. Hence, like ECVP, VSAC does not have a stable location nor a stable organization. Its flavour changes and is strongly influenced by the locations and the organizers. In 2014 it was held in Belgrade, in 2015 in Liverpool, in 2016 in Barcelona, and in 2017 in Berlin. In 2018 ECVP comes back to Italy, and will take place in Trieste.

The aims of the conference are conveniently summarized by each year organizers: “There is a growing interest in studying interactions between perception and art [...] VSAC welcomes all kinds of work and approaches, from phenomenological to biological and computational, exploring the link between the science of perception and the arts.” (http://ecvp.org/2015/sac.html). “The study and production of Art has always fascinated both artists and scholars in equal measure throughout history. Nevertheless, in the 20th century there has been little encouragement for artists and scientists to meet and collaborate. Fortunately this state of affairs is currently changing” (http://www.ub.edu/ecvp/about-vsac). “Its main focus is to better connect the communities of visual scientists and artists in order to deepen our understanding of aesthetic phenomena. The VSAC is an ideal venue to debate and collaborate on all topics associated with the perception and evaluation of artworks (https://vsac2017.org). Summarizing, the aim of VSAC is that “of connecting the communities of visual scientists and artists to promote cross-fertilization between the two domains”: open as it is to all those who are interested in art, “VSAC invites all people that connect visual perception and the arts (e.g., empirical, experimental, philosophical, phenomenological, computational approaches)” (http://www.vsac2018.eu).

Information separation in art investigation; a survey

Jan Blažek, Barbara Zitová.

Dept. of Image Processing, Institute of Theory of Information and Automation, The Czech Academy of Sciences, Czech Republic

***Abstract*— The goal of artwork analyzes is often to detect of pentimenti, retouches, overpaintings, or varnishes in order to understand a painting structure. A common model of a painting used for interpretation of an artwork multimodal dataset is based on its multilayer characteristics. Another possibility how to address an artwork structure is to study an *information gain* of a particular modality. We have developed a new approach [2] for the *information gain* extraction and demonstrated its applicability. We present a comparison of four methods for the information separation [4, 1, 3, 2] applied on a multimodal dataset. Their ability to uncover concealed features of paintings will be presented together with their requirements and limitations. The separation limits will be shown using a concept of the intensity correspondence matrix (ICM), which can well describe the correlation and the mutual information. ICM also gives evidence of possibility to achieve an effective signal separation.**

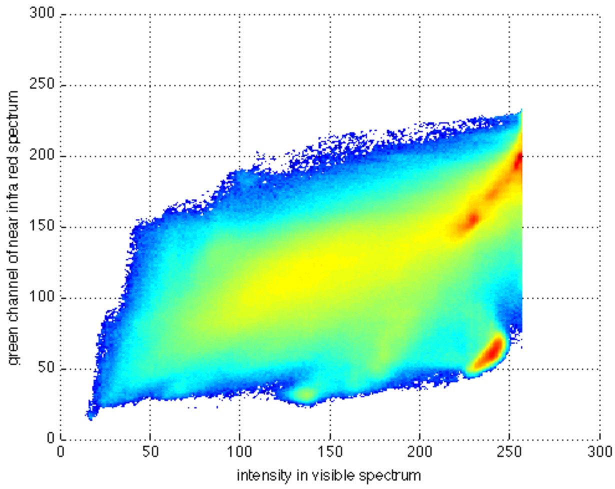

Figure 1: A 2D histogram of the intensities correspondence. In the area of VIS intensity level 230 there are two peaks. The low intensity in NIR corresponds with underdrawings while the high intensity come from areas without underdrawings. Such dataset is separable.

1 Introduction

The general problem of multimodal datasets (if we neglect the major problem of their registration) is their very high cross-modal correlation caused by the principle of reflection or transmission measurements. The information content of a painted surface affects all radiation passing through; reflectance in the visible part of the spectra (VIS) as well as penetrating modalities (terahertz (THz), X-Ray (RTG), near-infrared (NIR)). The contrast of deeper layers is significantly lowered and often falls down to the level of noise. Thanks to this, the modal images are often hardly readable. Moreover, they can go to be completely useless for art investigation (there are exceptions from this concept e.g.[5]).

The identification of features of the covered painted layers is relevant task for image processing. For this purpose the methods for a signal separation come to the scene. Making an assumption, that the VIS modality is less penetrating than the THz, the X-ray or the NIR, respectively, we can assimilate the idea of painted layers. While the VIS modal image is affected just by the surface "layers visible in VIS modality" the more penetrating modalities can be affected also by "some deeper layers" too. For simplicity, we split the painting into two parts: *the surface* or *top layer* which represents the layers affecting VIS reflectogram and *the layer underneath* which contains everything else.

In our survey we would like to compare four methods [4, 1, 3, 2] used for the multimodal dataset separation in order to visualize concealed features hidden in the images obtained in penetrating modalities.

2 ICM and its patterns

We offer a new perspective to signal separation by an *intensity correspondence matrix* (ICM), which can be used as the common denominator for all mentioned studies. Moreover, the patterns for ICM describe the problems and the limits of the information separation of a multimodal dataset.

The ICM is a matrix, which contains the frequency of correspondent pixel(s) intensities of two different modalities. As a matrix *hyper-column* we name a VIS vector while a *hyper-row* denotes a vector in a target modality. The number of hyper-rows and hyper-columns corresponds to the intensity levels recognized in each modality, while the dimensionality of the hyper-row and hyper-column is given by the pixel vector length and the pixel neighborhood size taken into account. E.g. for two modalities with intensity levels $l \in L = \{0, 1, 2, ..., 255\}$ the ICM is 2D histogram with $\|L\| = 256$ bins in both dimensions (see Figure (1)) while for approach in [3] we have 6D histogram with $(n \times d \times \|L\|)^2$ bins, where n is the number of pixels in the patch and d the length of per pixel intensity vector.

The visualization (if possible) of the (low) dimensional ICM can give us a notion of the potential separability effectiveness.

In the ICM, it is easier to recognize more probable correspondences of modal vectors from the less probable ones in the context of modality. In general, the dataset defines a **mapping** between VIS and the second modality. But an algorithm for separation is just a **function** (see Figure (2)). The discrepancy between the mapping and the function causes that all the corresponding intensities in the target modality (for one *hyper-column*) are reduced by the separation function to just one output vector. If we ask for the effectiveness of mapping to function reduction we can recognize peripheral but relevant patterns of the ICM:

- The good case - the highest peaks in ICM hyper-columns correspond to the *top layer* effect in both modalities

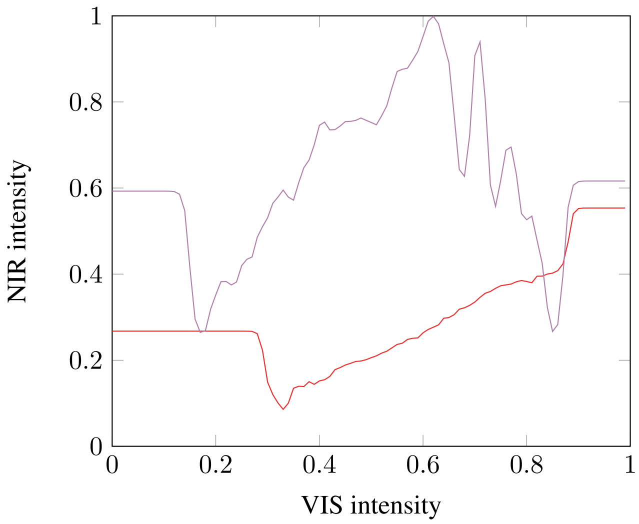

Figure 2: Two successfully trained transfer functions from the intensity in the VIS to intensity in the NIR. The purple line demonstrates how non-trivial such function can be. On the contrary, the red one is an example of the identity like function. Both lines were trained on the dataset presented in [2], the purple one on the dataset in the Figure (10), the red line on the dataset in Figure (9). The VIS input were reduced to the one dimensional intensity gray values.

- The bad case - multiple peaks - several relevant peaks per hyper-column with similar frequencies
- The bad case - smooth distribution - uniform distribution of values per hyper-column

Our results [2] highlight the problem of separability based on the number of materials and their mixtures in the dataset as well as the importance of correlation of both input modalities.

3 Generalization

We connect the presented approaches with the model of ICM, because the ICM and our approach [2] can distinguish cases when the information obtained from the NIR or the X-ray can be effectively separated and when this is impossible.

Firstly we reduce of the problem of the estimation of the separation function to the problem of searching corresponding vector in NIR or X-ray for a vector in VIS. This means an assignment of the hyper-row to a hyper-column in ICM. An assignment of these vectors to all hyper-columns defines the approximation function converting the *top layer* information in the VIS modality to the *top layer* information in the second modality.

In the Gooch's and Tumblin's paper [4] authors estimate the target function for each mean-shift based segment. For the whole segment just one X-ray vector is defined, which is computed as the mean value of the segment in the X-ray. This corresponds to the assignment of an hyper-column to the mean value of relevant hyper-rows. For the application of the method there must be just one significant peak per each hyper-column, otherwise, the mean value which is probably irrelevant will be taken as the X-ray representation.

The method in our paper [2] for a hyper-column extracts the hyper-row which minimizes the square error of approximated and real vectors. In an ideal case this is the most probable hyper-row. Our experiments demonstrated that this peak must be significantly higher than any other hyper-row frequency in the same hyper-column.

In the paper of Anitha et al. [1] the mutual entropy of an approximated *top layer* is minimized and the entropy of XRF *information gain* is maximized. This approach, in most cases, takes the highest peak in the hyper-column as the XRF representation. Other peaks in the hyper-column and low frequency values around them are put into the *under-paintings* category. Other selections of the hyper-column representative are also possible, mostly due to the applied regularization, but they need custom explanation and their stability is lower.

The unsolvable problem for our [2] approach is an existence of more than two peaks per hyper-column. Because in such case, false positives will be produced in the separated signal. A robustness of Anitha's et al. [1] approach is here improved by the wavelet decomposition which causes that a pixel neighborhood also affects output C_u and C_s intensities. But the pixel neighborhood itself does not influence the estimation of an approximation function.

The last approach of Deligiannis et al. [3] includes the influence of a pixel neighborhood into the approximation estimation. The authors use the ICM not only per one pixel intensity vector but they include into one hyper-column/row pixel neighborhood. With the raw pixel intensities this will cause a curse of dimensionality which in the extreme case can produce an unstable algorithm. The authors solve this using the coupled dictionaries with limited number of words and limited linear combination of those words. A notion what is happening here is more difficult, because more dimensions with different meaning are included into our ICM concept. In the optimal case the more dimensions and the different pixel context moves the peaks in hyper-space, defined by hyper-column dimensions, close to each other and a more relevant representative vector could be selected. However, the evaluation of this hypothesis is out of the scope of our paper. In general, stability of this approach can be negatively affected more than all previous methods by decreasing common information in analyzed modalities due to the minimization in higher dimensional space. As we presented in our study [2] the maximum of the second modality *information gain* is around 10%. This condition is for X-ray and XRF modalities hard to meet. In our case, these 10% limit also includes the noise from both modalities.

4 Conclusion

The separation of the modality information to the *top layer* and the *information gain* is in all referred studies done, in principle, by an approximation function. This function estimates the general mapping between modalities which we analyze by the ICM. The computed function in the ICM is realized by hyper-column $\rightarrow$ hyper-row pairs. We have pointed out that there exist patterns in the ICM which cannot be separated by any function, but the ICM in such cases can be constructed in different way. In [3] the ICM is constructed with respect to the pixel neighborhood while the sparsity of such ICM is reduced by dictionary code-words and patch representation rules.

For further research an identification of the bad patterns of ICM as well as an identification of methods limits is crucial. We recommend to test separation methods on phantoms, where the statistical ground truth would be known and non-trivial.

References

[1] Anila Anitha et al. "Restoration of X-ray fluorescence images of hidden paintings". In: *Signal Processing* 93.3 (2013). Image Processing for Digital Art Work, pp. 592–604. ISSN: 0165-1684. DOI: 10.1016/j.sigpro.2012.09.027. URL: http://www.sciencedirect.com/science/article/pii/S0165168412003659.

[2] Jan Blažek et al. "Improvement of the visibility of concealed features in artwork NIR reflectograms by information separation". In: *Digital Signal Processing* 60 (Jan. 2017), pp. 140–151. DOI: 10.1016/j.dsp.2016.09.007.

[3] N. Deligiannis et al. "Multi-Modal Dictionary Learning for Image Separation With Application In Art Investigation". In: *IEEE Transactions on Image Processing* PP.99 (2016), pp. 1–12. ISSN: 1057-7149. DOI: 10.1109/TIP.2016.2623484.

[4] Jack Gooch, Amy and Tumblin. "Visualizing pentimenti: revealing the hidden history of paintings". In: *Journal of Mathematics and the Arts* 00.00 (2007), pp. 1–11. ISSN: 17513472. DOI: 10.1080/1751347YYxxxxxxxx.

[5] David Thurrowgood et al. "A Hidden Portrait by Edgar Degas." In: *Scientific Reports* 6 (2016), p. 29594. DOI: 10.1038/srep29594. eprint: arXiv:1011.1669v3. URL: http://www.nature.com/srep/2016/160804/srep29594/full/srep29594.html.

Disrobing Adam and Eve with the linear osmosis model

Marie d'Autume[1], Enric Meinhardt-Llopis[1].

[1] CMLA, ENS Cachan, CNRS, Université Paris-Saclay, 94235 Cachan, France.

***Abstract*— The linear osmosis model, an alternative to Poisson editing, reconstructs a composite image from an input generally given by the drift fields, invariant to contrast changes, extracted from one or several images. It is well adapted to tasks where the input images contrast vary wildly, as is often the case for multi-spectral images. We show that its stationary local elliptic formulation with mixed boundary conditions is particularly appropriate for the task of digitally removing over-paint in illuminations for which we dispose of underlying information provided by infrared imaging.**

1 Introduction

The osmosis framework was first introduced in [1], [2] as a new model for compact image representation, shadow removal and seamless image cloning. A drift-diffusion PDE problem, the linear osmosis filter is particularly well adapted to tasks where the input images contrast vary wildly, as is the case for the application to image fusion. Recently it was successfully used to solve the light balance problem in Thermal-Quasi Reflectography imaging [5]. Yet in [3], we showed that the osmosis parabolic equation can be advantageously replaced by a the stationary elliptic equation derived from its stable state

$$\Delta u = \operatorname{div}(\mathbf{d}u),$$

where $\mathbf{d}$ is the input vector field and u is the image to be recovered under appropriate boundary conditions. For $\mathbf{d}_v = \nabla v/v$, v is a solution of the elliptic equation. For this reason the authors in [1] defined it as the canonical drift vector field of the image v. It is because $\mathbf{d}_v$ is invariant to multiplicative changes of v that the linear osmosis model is so successful with input images with different contrast.

Multispectral imaging now allows us to look under the layers and recover under-drawings or covered up areas of a painting. These multispectral images usually have very different dynamic range which makes he osmosis equation particularly well suited to deal with them. Therefore we propose to apply this model to the task of digitally removing over-paint in illuminations from the primer of Claude de France, a manuscript from the Fitzwilliam museum [1]. The illuminations illustrating Adam and Eve's creation, temptation and fall from grace were censored by a later owner who had some veils and leaves added to hide their nakedness. But the infrared reflectogram available on the museum website gives some information on the details hidden by the added pigments. It is tempting to try and combine the informations of the colour image and infrared reflectogram to obtain an image closer to the illumination in its original state.

In the following we first recall some theoretical results for the elliptic linear osmosis problem. Then we describe how this model is typically used for the tasks of seamless cloning and shadow removal. And finally we explain how we applied it to this problem of digital restoration.

[1] The images are freely available on the museum website: http://www.fitzmuseum.cam.ac.uk/illuminated/manuscript/discover/the-primer-of-claude-of-france/section/panel-intro/folio/page-4/section/panel-intro

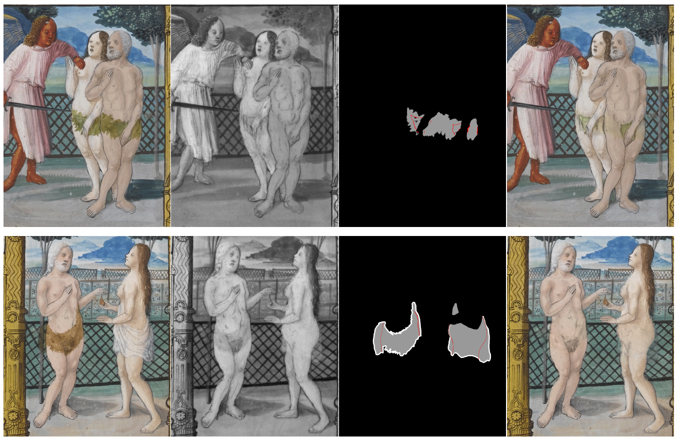

Figure 1: From left to right: colour image, infrared reflectogram, mask and final result. In the mask, the red lines correspond to Neumann boundary conditions and the white line correspond to pure diffusion.

2 Theoretical results

Here we briefly recall some theoretical results that we first enunciated in [3].

Proposition 1. *Let* $0 < \alpha \leq 1$. *Let* D *be a* C^3 *domain. Let* $\mathbf{d} \in C^3(\overline{D})$, $g : \overline{D} \to \mathbb{R}_{+*}$, $g \in C^{2,\alpha}(\overline{D})$. *Then the elliptic Dirichlet problem*

$$\begin{cases} \Delta u = \operatorname{div}(\mathbf{d}u) & \text{on} \quad D \\ u = f(x) & \text{on} \quad \partial D \end{cases} \tag{1}$$

has a unique solution $u \in H^3(D)$.

Remark 1. *For some applications it can be useful to have mixed boundary conditions. In this case we replace the condition* $u = f(x)$ *by* $\langle \nabla u - \mathbf{d}u, \mathbf{n} \rangle = 0$ *on part of the boundary* ∂D.

We adopt the discretization proposed by Weickert *et al.* [1]. We consider a grid size $h = 1$ in x and y direction. To compute the divergence of $\mathbf{d}$ we discretize d_1 on the grid translated by half a pixel in the vertical direction and d_2 on the grid translated by half a pixel in the horizontal direction.

Definition 1. *Let* v *be a positive discrete image, we define the* discrete canonical drift vector field $\mathbf{d}_u$ *by*

$$\begin{cases} d_{1,i+\frac{1}{2},j} = \frac{2(v_{i+1,j} - v_{i,j})}{v_{i+1,j} + v_{i,j}} \\ d_{2,i,j+\frac{1}{2}} = \frac{2(v_{i,j+1} - v_{i,j})}{v_{i,j+1} + v_{i,j}} \end{cases} \tag{2}$$

This yields the osmosis equation

$$
\begin{aligned}
Lu_{i,j} = {} & u_{i+1,j}\left(1-\frac{d_{1,i+\frac{1}{2},j}}{2}\right)+u_{i-1,j}\left(1+\frac{d_{1,i-\frac{1}{2},j}}{2}\right)\\
& + u_{i,j+1}\left(1-\frac{d_{2,i,j+\frac{1}{2}}}{2}\right)+u_{i,j-1}\left(1+\frac{d_{2,i,j-\frac{1}{2}}}{2}\right)\\
& - u_{i,j}\left(4+\frac{d_{1,i+\frac{1}{2},j}-d_{1,i-\frac{1}{2},j}}{2}+\frac{d_{2,i,j+\frac{1}{2}}-d_{2,i,j-\frac{1}{2}}}{2}\right) \qquad (3)
\end{aligned}
$$

We will from now on assume that $\mathbf{d}$ is reasonably small, namely

$$|d_1(\mathbf{x})| < 2,\ |d_2(\mathbf{x})| < 2,\ \forall \mathbf{x} \in \Omega, \qquad (4)$$

thus keeping the weights of all four neighbours of $u_{i,j}$ positive in (3). This condition is always satisfied when $\mathbf{d}$ is a canonical drift vector field of a positive image.

To write this discretization in a more compact notation, we replace the double indexing in each pixel $(i, j) \in \{0, \ldots, h-1\} \times \{0, \ldots, w-1\}$ by a single index $k = i + jh$ and assemble all unknown grey values in a single vector $\mathbf{u} \in \mathbb{R}^{wh}$ and end up with the linear system

$$\mathbf{A}\mathbf{u} = \mathbf{b} \qquad (5)$$

where $\mathbf{A}$ is a pentadiagonal sparse matrix whose nonzero elements are detailed in [3] and $\mathbf{b}$ is a vector of the same size as $\mathbf{u}$ that encodes the Dirichlet boundary conditions.

Proposition 2. *Let* $\mathbf{d} \in \mathbb{R}^{(w+1)h} \times \mathbb{R}^{w(h+1)}$ *such that* $\|\mathbf{d}\|_\infty < 2$. *Then the matrix* $\mathbf{A}$ *corresponding to the mixed boundary problem* (5) *derived from the discretisation* (3) *is invertible.*

3 Seamless cloning and shadow removal

The osmosis model is typically used to create a new image by modifying and combining the canonical drift-vector fields of one or more images.

For the seamless cloning case where one wants to paste part of an image f in the subdomain D of an image g, the problem is solved on D with $\mathbf{d} = \mathbf{d}_f$ in D, $\mathbf{d} = (\mathbf{d}_f + \mathbf{d}_g)/2$ on ∂D.

For the shadow removal problem, the problem is solved on the subdomain of the image g where there is a shadow to remove with $\mathbf{d} = \mathbf{d}_g$ except on the edge of the shadow where it is put to zero. Indeed a shadow can be reduced as a multiplicative change in the domain of the shadowed region of the image while the canonical drift vector field is invariant to multiplicative change. The presence of the shadow is therefore only encoded in the drift vector field on the edge of the shadow.

But as we noted in [3], because of the diffusion occurring on the edge of the shadow from having put the drift-field to zero, this solution only works for cast shadows. For attached shadows, we enforce Neumann boundary conditions on the edge that separates the lit part of an object from its unlit one, Neumann boundary conditions ensuring that no exchange takes place across the boundary.

4 Application to the primer

The problem of digitally removing the over-paint can be viewed as a combination of the applications previously described when one has access to a registered infrared reflectogram of the illumination. The quality of the result and the complexity of the problem are of course directly linked to the pigments used and the wavelength chosen for the infrared reflectogram.

In the ideal case, the pigments added do not appear on the infrared reflectogram while the colours to be restored are perfectly encoded in it. In this case the problem is reduced to a simple seamless cloning application with Dirichlet boundary conditions.

Such an ideal case is however uncommon. The illuminations of the primer provide us with several typical issues. First, in the nice case where the added pigments do not appear on the infrared reflectogram, some original pigments can also be almost absent in the infrared reflectogram. This is the case in the first illumination displayed in figure 1. The colour distinction between the original fig leaves and the skin is slight on the infrared reflectogram. Simply following the seamless cloning method lead to an output image where colours are similar on both sides of the boundary. Thus to ensure that the skin colour and the green of the original fig leaves do not mix, we enforce Neumann boundary conditions along the edges involved (the red lines in the mask).

Another issue is illustrated by the second illumination of Figure 1. Some added pigments appear on the infrared reflectogram. But if the added layer has little to no texture discernible on the infrared and the original image appears by transparency it can be considered as a shadow on the infrared reflectogram. Then we add a shadow removal component to our earlier method: we replace the drift-field of the RGB image by the drift-field of the infrared reflectogram on the subdomain to be restored but we negate the drift-field on the edge of the added layer (white lines in the mask). Still any texture from the added cloth present in the infrared image appears in the final result, as can be seen by the stain on Adam's hip. But it is small enough that the result is still visually acceptable.

Finally in the case where the over-paint contains too much texture appearing in the infrared reflectogram, our method cannot give any good result. A possible solution may be to identify the pigments of the over-paint and select a more adapted wavelength for the infrared reflectogram.

5 Conclusion

We proposed a method to digitally remove over-paint from an illumination. This method requires some patient work from the user for the mask creation and is heavily dependent on the infrared wavelength. For the illuminations from the primer that we worked on, the results were rather satisfying but the method should be tested on a larger dataset.

References

[1] J. Weickert, K. Hagenburg, M. Breuß and O. Vogel, "Linear osmosis models for visual computing", EMMCVPR, 2013.

[2] O. Vogel, K. Hagenburg, J. Weickert and S. Setzer "A fully discrete theory for linear osmosis filtering", ICSSVMCV, 2013.

[3] M. d'Autume, J.-M. Morel and E. Meinhardt-Llopis, "A flexible solution to the osmosis equation for seamless cloning and shadow removal", ICIP, 2018 (to appear).

[4] P. Pérez, M. Gangnet and A. Blake "Poisson image editing", ACM Transactions on Graphics (TOG), 2003.

[5] S. Parisotto, L. Calatroni and C. Daffara "Efficient Osmosis Filtering of Thermal-Quasi Reflectography Images for Cultural Heritage", arXiv preprint arXiv:1704.04052, 2017

Paint Loss Detection via Kernel Sparse Representation

Shaoguang Huang[1], Laurens Meeus[1], Bruno Cornelis[2], Bart Devolder[3], Maximiliaan Martens[4] and Aleksandra Pižurica[1].

[1]Department of Telecommunications and Information Processing, IPI, imec research group at Ghent University, Belgium

[2]Department of Electronics and Informatics (ETRO), Vrije Universiteit Brussel, Belgium

[3]Department of Painting Conservation, Royal Institute for Cultural Heritage (KIK/IRPA), Belgium

[4]Department of Art, music and theatre sciences, Ghent University, Belgium

Abstract— **Automatic paint loss detection is desired for supporting conservation/restoration treatments of paintings. Firstly, producing condition reports with appropriate damage surveys requires now a lot of manual work from the restorers. Secondly, paint losses have to be accurately detected prior to running virtual restoration. Large variation of paint loss in size, shape, intensity as well as varying and complex background make this problem a challenging task. We develop a multimodal paint loss detection method based on sparse representation, which incorporates the information from multiple imaging modalities in a high-dimensional kernel feature space and makes use of the spatial context. To cope with unreliable labelled data, we introduce a majority voting approach. Experimental results with the data set of the *Ghent Altarpiece* demonstrate the effectiveness of the proposed approach.**

1 Introduction

Digital painting analysis has been a rapidly growing field, attracting a lot of interest recently in the signal processing community [1]. The tasks such as characterization of painting style and forgery detection [2, 3], crack detection [4], authorship identification [5], classification of ancient coins [6], canvases [7] and portraits [8], removal of canvas patterns [9] and inpainting [10, 11] have demonstrated the great potential of digital image processing techniques.

Loss of paint is typically caused by abrasion and mechanical fracture. In old oil paintings, paint losses were often overpainted during various restoration campaigns. Modern conservation treatments typically require not only removal of old varnish, but also removal of old retouches and overpaint, which may reveal paint losses underneath [13]. Detection of such paint loss areas is of great importance to painting conservators for estimating the extent of the damaged area, which needs to be maintained for documenting purposes, but also as a crucial step for virtual inpainting to provide simulations for the actual restoration. Despite the importance of automatic paint loss detection, this problem has received little attention in the literature so far. Nowadays, paintings are typically scanned with a multitude of imaging modalities. During restoration campaigns, additional scans are typically made at various stages of the rest oration treatment. Examples are shown in Fig. 1. (a) - (e). We want to exploit such multi-modal information to detect paint losses more reliably. Our approach will be based on constructing (training) a dictionary of prototypes that can be used to effectively, i.e. sparsely, represent paint loss samples.

Sparse Representation Classification (SRC) [14] proved to be effective in various image classification tasks, especially in computer vision and remote sensing. It assumes that each test sample can be sparsely represented as a linear combination of atoms from a dictionary which is constructed by the selected training samples. Directly applying SRC to our task results in poor performance due to the large variability of paint loss, and complex background. To cope with these challenges, it is necessary to incorporate appropriately both spatial context and inter-modal dependencies. Our previous work employed several spatial features within local patches and achieved a good detection performance [13]. However, hand-crafting such features leaves much choice and would involve ad-hoc choices and a lot of manual tuning. Therefore, in this paper we propose a multimodal paint loss detection method based on sparse representation that directly exploits the information from multiple imaging modalities in the kernel feature space and integrates the spatial information of context into the model.

2 The proposed method

The multiple imaging acquisitions are typically captured via different imaging devices and often have different resolutions. Thus image alignment for all the modalities, which is also called image registration, should be first completed. Here we use a joint photometric and geometric image registration technique [15] to register these images. We concatenate the pixels within a square window in the registered data cube into a vector. By using a kernel function, the vector is projected to a high-dimensional kernel feature space. Next to the two classes: 'paint loss' and 'background', we specify a third class 'crack', which is by art restorers treated differently than larger portions of missing paint called paint loss.

The modified SRC model with respect to sparse coefficients of $\mathbf{x} \in \mathbb{R}^m$ in the projected kernel feature space is

$$\hat{\alpha} = \arg\min_{\alpha} \|\phi(\mathbf{x}) - \phi(\mathbf{D})\alpha\|^2 \;\; s.t. \; \|\alpha\|_0 < K_0, \qquad (1)$$

where $\phi : \mathbb{R}^m \to \mathscr{F} \subset \mathbb{R}^{\hat{m}}$ is an implicit mapping function that projects $\mathbf{x}$ to a higher dimensional space; $\phi(\mathbf{D}) = [\phi(\mathbf{d_1}), \phi(\mathbf{d_2}), ..., \phi(\mathbf{d_N})]$ is the dictionary in the projected space and $\mathbf{d_i} \in \mathbb{R}^m$ $(i = 1, 2, ..., N)$ are the training samples. Once the sparse coefficients are calculated, the class-specific residuals can be computed by

$$\begin{aligned} r_i(\phi(\mathbf{x})) &= \|\phi(\mathbf{x}) - \phi(\mathbf{D_i})\alpha_i\|_2 \\ &= \langle \phi(\mathbf{x}) - \phi(\mathbf{D_i})\alpha_i,\; \phi(\mathbf{x}) - \phi(\mathbf{D_i})\alpha_i \rangle^{1/2} \\ &= (\kappa(\mathbf{x},\mathbf{x}) - 2\alpha_i^T \mathbf{K_{D_i}} + \alpha_i^T \mathbf{K_{D_iD_i}} \alpha_i)^{1/2}, \qquad (2) \end{aligned}$$

where $\kappa : \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}$ is a kernel function defined by $\kappa(\mathbf{x_i}, \mathbf{x_j}) = \langle \phi(\mathbf{x_i}), \phi(\mathbf{x_j}) \rangle$; $\mathbb{K}_{\mathbf{D_i}} \in \mathbb{R}^{N_i}$ is a vector associated with

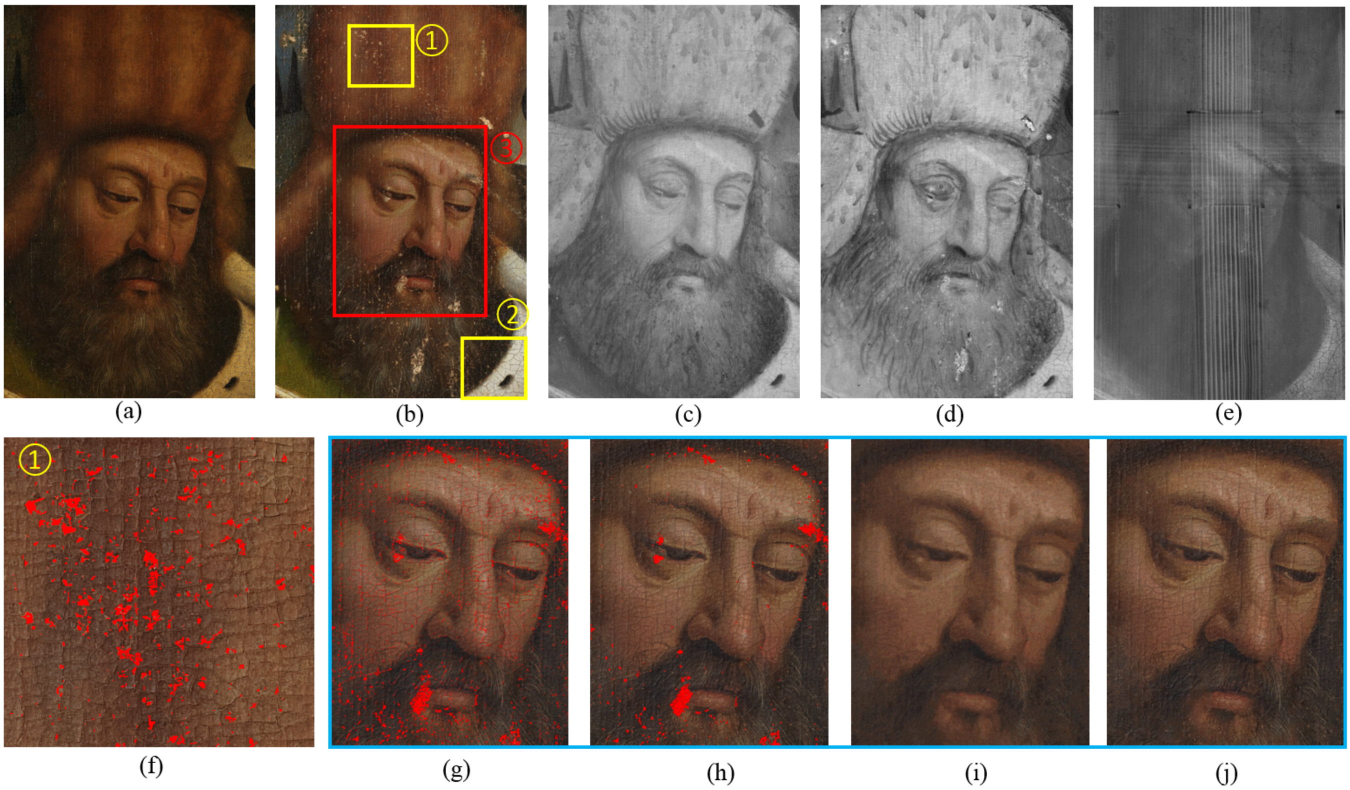

Figure 1: Top row: multiple imaging scans, which include (a) macrophotography before cleaning, (b) macrophotography after cleaning, (c) infrared macrophotography before cleaning, (d) infrared reflectography after cleaning and (e) X-radiography before cleaning. Bottom row: (f) Annotated patch 1 used for training, (g) detection map obtained by applying SRC, (h) detection map obtained by the proposed method, (i) inpainting results using the method of [12] with the SRC map from (g) and (j) inpainting result with the map obtained by our method.

class i in $\mathbb{K}_{\mathbf{D}} \in \mathbb{R}^N = [\kappa(\mathbf{d_i},\mathbf{x}),\cdots,\kappa(\mathbf{d_N},\mathbf{x})]^T$; $\mathbb{K}_{\mathbf{D_iD_i}} \in \mathbb{R}^{N_i\times N_i}$ is a matrix corresponding to class i in $\mathbb{K}_{\mathbf{DD}} \in \mathbb{R}^{N\times N}$ with entries $\mathbb{K}_{\mathbf{DD}}(i,j) = \kappa(\mathbf{d_i},\mathbf{d_j})$ and α_i is a vector associated with class i in α. Then we label the class of a test sample by

$$class(\mathbf{x}) = \arg\min_{i=1,2,3} \ r_i(\phi(\mathbf{x})). \tag{3}$$

We denote by $\mathbf{Map_{crack}}$ the obtained binary crack map. By collecting all the residuals $r_i(\phi(\mathbf{x}_i))$, we form the residual cube. Here we denote by $\mathbf{R} \in \mathbb{R}^{M\times N\times 3}$ the reshaped residual cube, where each layer corresponds to one class.

Typically, paint losses will occupy an area larger than a single pixel. Hence, pixels within a relatively small neighbourhood are likely to belong to the same class and share similar sparse representation coefficients. Therefore we apply a smoothing filter to each layer of the residual cube to make the coefficients of neighbouring pixels similar to each other. In particular, we use for this purpose a weighted least square (WLS) [16] filter. The binary paint loss map, $\mathbf{Map}^{'}$, can be calculated by selecting the smallest smoothed residual. This smoothing has an adverse effect on thin cracks, which tend to be assigned to paint loss (or to background). To solve this, we use the crack map $\mathbf{Map_{crack}}$ generated prior to smoothing, as follows

$$\mathbf{Map} = \mathbf{Map}^{'} \odot \overline{\mathbf{Map_{crack}}}. \tag{4}$$

The training samples in $\mathbf{D}$ of (1) play an important role as they are used to supervise the model to generate the corresponding characteristics of paint loss and background. However, for most cases, compared with the samples of background, the number of paint loss samples is rather small. In addition, accurate annotation on a pixel level is a highly challenging task, which may lead to mislabelled samples. Errors can be caused by blurring in low-resolution images, large transitions and low contrast between target and background, noise, artefacts and so on. To cope with this problem, we suggest a majority voting strategy:

$$identity(\mathbf{x}_j) = \arg\max_{c} \ p_j^c \tag{5}$$

where the fraction $p_j^c = N_j^c/K$ is an empirical probability for the pixel j to belong to the class c. K is the number of simulations and N_j^c the number of times that pixel j was assigned to class $c \in \{Paint\ loss, Other\}$.

3 Results and discussion

We illustrate the detection result on a part of the panel *prophet Zachary*, image patch 3 in Fig. 1 (b). The training samples are from other two image patches in Fig. 1 (b), which were annotated by a painting conservator. Fig. 1 (f) shows one of the annotated image patchs. We set the number of training samples in each class to 80 and K to 10. The imaging modalities in Fig. 1 (a), (b) and (c) are used. Fig. 1 (h) and (j) illustrate paint loss detection results of the proposed approach and virtual inpainting using the detected mask and the inpainting method from [12]. For comparison, we also show the paint loss map in Fig. 1 (g) that is produced by applying the original SRC with multimodal images and majority voting. The corresponding inpainting result is reported in Fig. 1 (i). Obviously the proposed method reduces significantly false detections. Consequently, we avoid previous excessive oversmoothing and undesired removal of cracks during virtual restoration.

References

[1] P. Abry, A. G. Klein, W. A. Sethares, and C. R. Johnson, "Signal processing for art investigation," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 14–16, July 2015.

[2] C. R. Johnson, E. Hendriks, I. J. Berezhnoy, E. Brevdo, S. M. Hughes, I. Daubechies, J. Li, E. Postma, and J. Z. Wang, "Image processing for artist identification," *IEEE Signal Processing Magazine*, vol. 25, no. 4, pp. 37–48, July 2008.

[3] L. Platiša, B. Cornelis, T. Ružic, A. Pižurica, A. Dooms, M. Martens, M. D. Mey, and I. Daubechies, "Spatiogram features to characterize pearls in paintings," in *18th IEEE International Conference on Image Processing (ICIP)*, 2011, pp. 801–804.

[4] B. Cornelis, T. Ružic, E. Gezels, A. Dooms, A. Pižurica, L. Platiša, J. Cornelis, M. Martens, M. D. Mey, and I. Daubechies, "Crack detection and inpainting for virtual restoration of paintings: The case of the ghent altarpiece," *Signal Processing*, vol. 93, no. 3, pp. 605–619, 2013.

[5] P. Abry, S. G. Roux, H. Wendt, P. Messier, A. G. Klein, N. Tremblay, P. Borgnat, S. Jaffard, B. Vedel, J. Coddington, and L. A. Daffner, "Multiscale anisotropic texture analysis and classification of photographic prints: Art scholarship meets image processing algorithms," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 18–27, 2015.

[6] H. Anwar, S. Zambanini, M. Kampel, and K. Vondrovec, "Ancient coin classification using reverse motif recognition: image-based classification of roman republican coins," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 64–74, 2015.

[7] L. van der Maaten and R. G. Erdmann, "Automatic thread-level canvas analysis: A machine-learning approach to analyzing the canvas of paintings," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 38–45, 2015.

[8] R. Srinivasan, C. Rudolph, and A. K. Roy-Chowdhury, "Computerized face recognition in renaissance portrait art: A quantitative measure for identifying uncertain subjects in ancient portraits," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 85–94, 2015.

[9] B. Cornelis, H. Yang, A. Goodfriend, N. Ocon, J. Lu, and I. Daubechies, "Removal of canvas patterns in digital acquisitions of paintings," *IEEE Transactions on Image Processing*, vol. 26, no. 1, pp. 160–171, 2017.

[10] T. Ružic, B. Cornelis, L. Platiša, A. Pižurica, A. Dooms, W. Philips, M. Martens, M. D. Mey, and I. Daubechies, "Virtual restoration of the ghent altarpiece using crack detection and inpainting," in *Advanced Concepts for Intelligent Vision Systems*, 2011, pp. 417–428.

[11] A. Pižurica, L. Platiša, T. Ružic, B. Cornelis, A. Dooms, M. Martens, H. Dubois, B. Devolder, M. D. Mey, and I. Daubechies, "Digital image processing of the ghent altarpiece: Supporting the painting's study and conservation treatment," *IEEE Signal Processing Magazine*, vol. 32, no. 4, pp. 112–122, 2015.

[12] T. Ružic and A. Pižurica, "Context-aware patch-based image inpainting using markov random field modeling," *IEEE Trans. Image Process.*, vol. 24, no. 1, pp. 444–456, 2015.

[13] S. Huang, W. Liao, H. Zhang, and A. Pižurica, "Paint loss detection in old paintings by sparse representation classification," in *Proc. ITWIST*, 2016, pp. 62–64.

[14] J. Wright, A. Yang, A. Ganesh, S. Sastry, and Y. Ma, "Robust face recognition via sparse representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 2, pp. 210–227, 2009.

[15] H. Q. Luong, B. Goossens, A. Pižurica, and W. Philips, "Joint photometric and geometric image registration in the total least square sense," *Pattern Recognition Letters*, vol. 32, no. 15, pp. 2061–2067, 2011.

[16] Z. Farbman, R. Fattal, D. Lischinski, and R. Szeliski, "Edge-preserving decompositions for multi-scale tone and detail manipulation," in *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3, 2008, p. 67.

“NO CHAOS, DAMN IT!” Extracting paint maps from Macro-X-Ray Fluorescence scanning data to deconstruct Jackson Pollock’s “action painting” in *Number 1A, 1948*

Ana Martins

The David Booth Conservation Center, MoMA, 42 W54rd Street, New York, NY 10019, USA

Macro- X-Ray Fluorescence scanning (MA-XRF) of a painting can provide meaningful information on the artist’s materials and process. Results are generally presented in the form of distribution maps for the key chemical elements that can be related directly to the pigments present in the paints based on their known chemical composition. The interpretation of such maps however may prove difficult or confusing if an elemental marker is common to several pigments and thus present in multiple paints. This study proposes extracting the paint distribution maps instead of the elemental maps using a multivariate image analysis approach. The Multivariate Curve Resolution - Alternating Least Squares (MCR-ALS) method can easily retrieve these maps from the hyperspectral data cube, especially when the signature spectra of the pure paints can be inputted in the model. In the case of Pollock’s *Number 1A, 1948* (1948) areas of pure paint were easily identified in the painting and the average spectra over those areas were used as initial estimates for the twelve pure paints and canvas. The MCR-ALS processing method provided distribution maps for paints that could not be visualized individually or visualized at all using elemental maps, thus providing a much better understanding of Pollock’s process.

1 Introduction

In the late 1940s, Jackson Pollock stunned the art world with his unique and radical style of abstract painting which consisted of flinging and dripping enamel paint onto a canvas laid on the floor of his studio. Critics and other personalities were quick to dismiss his paintings as “chaotic” to which he famously replied by telegraphing the *Time* magazine: “NO CHAOS, DAMN IT!” There is indeed a startling level of intentionality and deliberation in Pollock’s painting technique as demystified and immortalized by the photographer Hans Namuth images and footage of Pollock at work in 1950. Can chemical imaging techniques such as Macro-X-Ray Fluorescence scanning (MA-XRF) further deepen our understanding of Pollock’s process and help deconstruct his “action painting”?

A previous study [1] reported on the successful use of MA-XRF combined with Multivariate Curve Resolution solved by Alternating Least Squares optimization (MCR-ALS, a bilinear factor decomposition method to recover the concentration and the pure response profiles for species in unresolved mixtures [2]) to identify, characterize and map the paints used by Pollock in four small sections of *Number 1A, 1948*. This approach is now being revisited and applied to the MA-XRF data for the entire painting in order to get a better understanding of the artist process.

2 Methodology

Number 1A, 1948 (J. Pollock, 1948, oil and enamel paint on canvas, MoMA access.# 77.1950, 173x264cm) was scanned with a Bruker M6-JETSTREAM (28 sections of 40x50cm, 50keV, 0.6mA, 0.3mm step size, 0.4mm spot size resolution and 10ms/pixel dwell time). The SOLO+MIA 8.5.2 software (Eigenvector Research Inc.) was used to pre-process the MA-XRF data (compression to the 1.8-11.4keV range and 100eV spectral resolution followed by Poisson scaling) and to carry out the MCR-ALS analysis (non-negative constraints for both concentration profiles and signature spectra and contrast enhancement for the concentration). MCR-ALS requires both the input of the number of pure components – in this case the number of paints, and an initial estimate of their response profile – the XRF signature spectra of the paints. A total of eleven paints had been identified during the exploratory study [1] and an additional black glossy paint was confirmed by visual examination. The study had also established that the paints had mostly been applied straight out of the tube or paint can and rarely mixed; therefore the signature spectra of the paints can be obtained by selecting areas of pure paint and extracting the corresponding average spectra from the MA-XRF data. The elemental maps were obtained with the Bruker instrument software using the Bayes deconvolution method for comparison with the paint maps. Gimp was used to stitch the final overall maps.

3 Results

The formulation of the paints used by modern artists, whether they were manufactured for artistic or commercial use, is generally complex, containing multiple pigments, fillers and additives. Consequently, it is often difficult or even impossible, in the case of the MA-XRF scanning of modern paintings, to identify a specific element marker for every single paint, and thus it would be more advantageous to map the paints instead, as done for example in multispectral imaging [3].

In the case of *Number 1A, 1948* for example, Ba (in the form $BaSO_4$ / barite – all pigments were confirmed by Raman and/or FTIR spectroscopy) is a major element in the white paint, together with Ti (TiO_2 / titanium white) and Zn (ZnO / zinc white). Ti however is also a major element in the cream paint, together with Zn. Moreover, Ba and Ti contributions are usually difficult to separate due to the overlap of the characteristic Ba L-lines and Ti K-lines, even when using deconvolution methods that are either empirical, like the Bayes deconvolution provided by the instrument software, or using a fundamental parameters approach with for example PyMca [4]. Zn, on the other hand, is present in the yellow paint, together with Ba, Cd and S ($CdS.(Zn).yBaSO_4$ / cadmium barium yellow), and in the red with Ba, Cd, Se and S ($Cd(S/Se).yBaSO_4$ / cadmium barium red). It is therefore impossible to visualize separately the distribution of the white and the cream paints using elemental maps, or to visualize cadmium yellow without the cadmium red. It is also sometimes difficult to separate and visualize smaller elemental contributions. For example, Ba is also present in the turquoise paint together with a small contribution of Mn ($BaMnO_4.BaSO_4$ / manganese blue) but the concentration of this second element is much stronger in the black glossy paint (Fe, Co, Zn and Pb also present). The turquoise paint appears therefore in the Ba map, together with all the other paints containing Ba, and does not appear in the Mn map unless a very strong contrast is used to reveal its contribution and mixed with the much stronger contribution of the glossy black paint. It is thus impossible to visualize the turquoise paint separately from other paints.

The advantage of using the MCR-ALS approach is clearly illustrated in figure 1 where the Ba, Ti, Mn and Cd elemental maps for a section of the painting are compared with the paint maps extracted with MCR-ALS for the cream, white, turquoise and yellow paints. In the paint maps, the cream and white paint are successfully separated tough both contain Ti. The yellow paint can be visualized separately from the red paint though both contain Cd, S and Ba, and a map for the turquoise paint was successfully extracted despite the fact that it is composed mainly of $BaSO_4$. The only paint that could not be extracted is a glossy black paint that contains mostly carbon black which is not detected by XRF. It also contains small amounts of K, Fe and Ca, but these elements are also present in the canvas in similar concentrations and it thus impossible to map them separately. In figure 2, the overall distribution map for the turquoise paint is compared with the image of the painting to demonstrate how revealing and fascinating the paint maps can be! Pollock's process is clearly revealed, the direction of his gestures, how he must have squeezed the paint out of the tube to create the long lines and how the paint formed tadpole like shapes when it landed on the canvas at the beginning of the line.

Figure 1: comparision of selected elemental and paint distribution maps for a section of the painting.

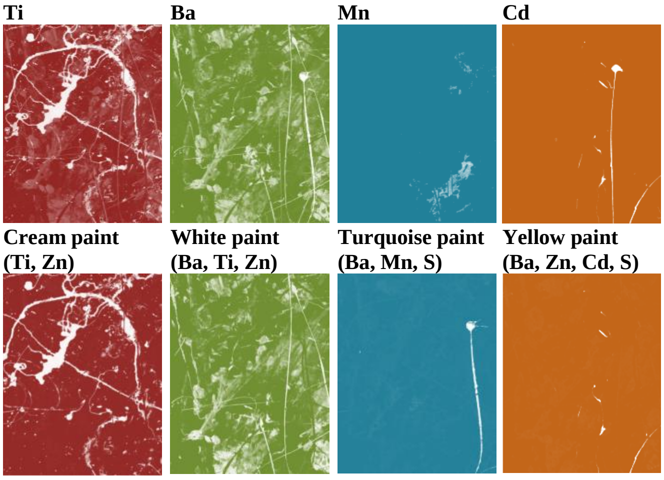

Figure 2: Comparison of the overall distribution map of the turquoise paint with the image of the painting.

Overall turquoise paint map extracted by MCR-ALS

***Number 1A, 1948* (1948) and highlighted section of the painting corresponding to the smaller maps above.**

4 Conclusions

The overall paint maps reveal Pollock's process like never before: the way he approached the canvas, the idiosyncrasy of each gesture and the rhythm that transpires in every paint layer. Positively "NO CHAOS, DAMN IT!"

Acknowledgements: The author wishes to acknowledge Bruker for lending the M6-Jetstream instrument.

References

[1] A. Martins, J. Coddington, G. Van der Snickt, B. van Driel, C. McGlinchey, D. Dahlberg, K. Janssens and J. Dik, “Jackson Pollock’s *Number 1A, 1948*: a non‑invasive study using macro‑X‑Ray fluorescence mapping (MA‑XRF) and multivariate curve resolution‑alternating least squares (MCR‑ALS) analysis”, Herit. Sci **4**:33, 2016.

[2] A. de Juan, J. Jaumot and R. Tauler, “Multivariate curve resolution (MCR). Solving the mixture analysis problem”, Anal. Methods **6**:4964, 2014.

[3] K.A. Dooley, J. Coddington, C.J. Krueger, D.M. Conover, M. Loew and J.K. Delaney, “Standoff chemical imaging finds evidence for Jackson Pollock's selective use of alkyd and oil binding media in a famous ‘drip’ painting”, Anal. Methods **9**:28-37, 2017.

[4] V.A. Solé, E. Papillon, M. Cotte, P. Walter, J. Susini, “A multiplatform code for the analysis of energy-dispersive X-ray fluorescence spectra”, Spectrochim, Acta B, **62**:63–8, 2007.

Deep Learning for Paint Loss Detection: A Case Study on the Ghent Altarpiece

Laurens Meeus[1], Shaoguang Huang[1], Bart Devolder[2], Maximiliaan Martens[3] and Aleksandra Pižurica[1]

[1]TELIN-IPI, Ghent University - imec, Belgium

[2]Princeton University Art Museum, United States of America.

[3]Department of Art, Music and Theatre Sciences, Ghent University, Belgium.

***Abstract*—Producing damage surveys as part of condition reports prior to and during restoration treatments is often a tedious and time-consuming work for the art restorer. We explore the potential of deep learning for automatic paint loss detection in paintings to facilitate condition reporting and to support restoration treatments. To the best of our knowledge, this is the first reported attempt of employing deep learning in this application. We develop a multiscale deep learning method, based on the recent U-Net architecture which we extend with dilated convolutions, such as to improve the detection stability. Our model is applicable to multimodal acquisitions such as visible, infrared, x-ray, and ultraviolet fluorescence. As a case study we use multimodal data of the *Ghent Altarpiece*. Our results indicate huge potential of the proposed approach in terms of accuracy and also its exceptional speed which allows interactivity and continuous learning.**

1 Introduction

One of the documentation tasks during the conservation/restoration of paintings consists of mapping lacunas as well as larger paint losses. Lacunas are mostly a result of drying and flaking of paint, although rough handling can also introduce losses. Currently, the mapping involves a lot of manual work since available software can only give a coarse estimation of the paint loss. This makes the process rather slow and tedious. In order to improve the automated mapping, smarter image processing techniques are sought.

Paintings are nowadays typically scanned in different modalities prior to restoration treatments and during their various stages. Hence, our approach will be designed to make use of the multimodal data. As the size of losses can range from a few to hundreds of pixels, the algorithm should not only take into account spectral information, but also have a large enough spatial support.

Technical literature on paint loss detection is limited. Huang et al [1] reported promising results with sparse representation classification (SRC), surpassing common machine learning approaches like linear regression classification and support vector machines in this task. We propose an alternative method based on deep learning, motivated by the huge success of convolutional neural networks in many other image classification and segmentation problems. We will validate our method on the panels of the *Ghent Altarpiece* [2], a monumental triptych made by the brothers van Eyck in the 15^{th} century. To the best of our knowledge we are the first to report a deep learning method for paint loss detection.

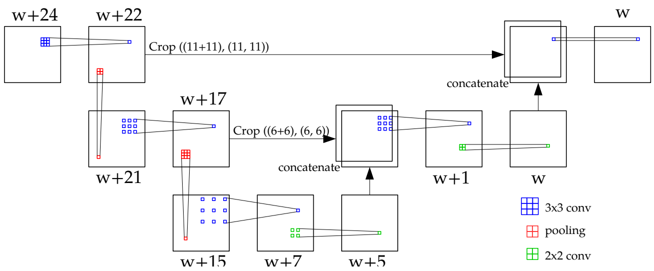

Figure 1: Proposed network architecture: a multiscale, multipath network with dilated convolutions.

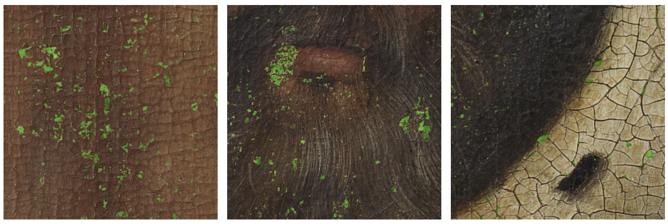

(a) Details from panel *Prophet Zachary*.

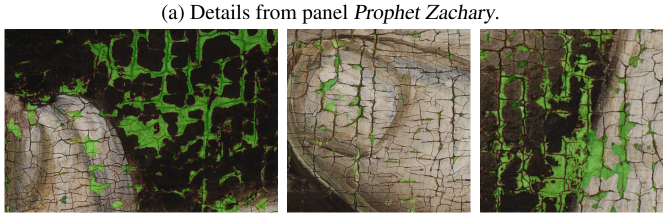

(b) Details from panel *John the Evangelist*.

Figure 2: Annotations made by the art restorer. Craquelure, formed by ageing, is not considered paint loss and is treated differently from it.

2 Methods

The proposed neural network architecture is visualised in figure 1. Similar to the U-Net [3] it consists of an encoder (left), a decoder (right) and skip-connect layers between the encoder and decoder (top) [4]. The difference between the U-Net and the proposed architecture is the removal of the decimation in the pooling layers. This way we maintain the same resolution in all layers and enforce true translation invariance. While this makes the bottom layers more dense than in the original U-Net, the outputs become more averaged out and this improves the stability of the output values. We observe that this leads to an increase in accuracy and learning capability of the model. The encoder consists of 3×3 convolutional layers and for the activation function the Rectified Linear Unit (ReLU, $\sigma(z) = \max(0, z)$) is used. Between these layers, pooling is introduced by taking the maximum in a 2×2 window with overlap to maintain the resolution. To maintain the same receptive field, the subsequent layers are replaced with dilated convolu-

(a) Original image during treatment. (b) Detected paint loss. (c) Original image during treatment. (d) Detected paint loss.

Figure 3: Paint loss detection on parts of the grisaille panel: *John the Evangelist*. The following modalities are provided to the model: visible before restoration, visible after varnish and over-paint removal, and infrared.

tions [5] and the amount of kernel weights remains identical with respect to the original U-Net. The decoder mirrors the encoder and pooling layers are replaced with upsampling 2×2 convolutional layers with linear activation. The skip-connect copies the layer of the encoder and concatenates it with the output of the upsampling layer to combine information of layers working at different resolutions. This gives the network the possibility to learn features on multiple scales simultaneously along the different paths. The last layer is a per-pixel, fully connected layer producing 2 feature maps: the probabilities of a pixel being paint loss or not. These probabilities are a result of the non-linear activation function Softmax:

$$\sigma(\mathbf{z})_j = \frac{e^{z_j}}{\sum_{k=1}^{2} e^{z_k}}. \quad (1)$$

It converts each pixel values $\mathbf{z}$ to a normalised probability vector $\hat{\mathbf{y}} = [p_0, p_1], p_0 + p_1 = 1)$.

To train the filters of the CNN, annotated data is requested. In our case, these were provided by art restorers of the *Ghent Altarpiece*. For each pixel, the annotation is converted to a vector $\mathbf{y}_i = [1; 0]$ for not paint loss and $\mathbf{y}_i = [0; 1]$ for paint loss. The CNN is trained to minimise the cross entropy:

$$C = -y_0 \cdot \log \hat{y}_0 - y_1 \cdot \log \hat{y}_1 \quad (2)$$

using Adaptive Moment Estimation [6]. The final prediction map is obtained by thresholding the probability p_1 of the output. We obtained the highest Intersection over Union score by thresholding at 0.5.

For the input of the network, the different modalities are first registered, concatenated and then cropped to a fixed size. Since each convolution and pooling operation reduces the output area, the input patch of the network is larger than the output patch to account for the receptive field. Because the input shape is a fixed amount bigger than the output shape and all layers operate at the same resolution, there is freedom in selecting the size of the patch to be segmented. Instead of classifying each pixel individually by setting the output shape to 1×1, it is more efficient to classify a big patch of pixels at once. When classifying nearby pixels, the overlap of the receptive field allows the convolutional layers to share computations. This speeds up the inference significantly and this means for the end user a big difference for practical usage.

3 Results and discussion

Figure 3 visualises the detection results on a larger part of the panel *John the Evangelist*. The 6 regions of the *Ghent Altarpiece* annotated by the art restorer, illustrated in figure 2, are from the panels *Prophet Zachary* and *John the Evangelist*. In total there are $807,740$ annotated pixels available of which 8.3% is paint loss. This amount is increased by a factor 8 after data augmentation by rotations of $90°$ and flips. These annotated regions are divided into smaller patches after which the network is trained on 80% of these patches. The remainder is used for picking the optimal hyperparameters and testing the accuracy. The following modalities were given to the model: optical images before and during treatment, infrared, infrared reflectography, X-ray, and ultraviolet fluorescence.

By segmenting patches of 10×10 or 100×100 instead of per pixel, we observe a speed increase of a factor 40 and 300 respectively for the inference. The results in figure 3 illustrate the binary prediction of a relatively large image (size 5954×7545), processed in less than a minute on a GeForce GTX 1070. Our experiments indicate a stable performance even with relatively few annotations. While our technique achieves similar results as the SRC-based method of [1], it is orders of magnitude faster. The art restorers appreciate the achieved results and the speed shows a huge potential for practical use of the proposed approach.

References

[1] S. Huang, W. Liao, H. Zhang, and A. Pižurica, "Paint loss detection in old paintings by sparse representation classification," in *International Traveling Workshop On Interactions Between Sparse Models And Technology*, (Aalborg, Denmark), pp. 62–64, August 2016.

[2] KIK/IRPA, "Belgian art links and tools," 2018.

[3] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," *ArXiv e-prints*, May 2015.

[4] M. Drozdzal, E. Vorontsov, G. Chartrand, S. Kadoury, and C. Pal, "The importance of skip connections in biomedical

image segmentation," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2016.

[5] F. Yu and V. Koltun, "Multi-Scale Context Aggregation by Dilated Convolutions," *ArXiv e-prints*, November 2015.

[6] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," pp. 1–15, 2014.

Mathematical osmosis imaging for multi-modal and multi-spectral applications in Cultural Heritage conservation

Simone Parisotto[1], Luca Calatroni[2] and Claudia Daffara[3].

[1]CCA, University of Cambridge (UK). [2]CMAP, École Polytechnique (France) . [3]University of Verona (Italy).

***Abstract*— In this work we present a dual-mode mid-infrared workflow [6], for detecting sub-superficial mural damages in frescoes artworks. Due to the large nature of frescoes, multiple thermal images are recorded. Thus, the experimental setup may introduce measurements errors, seen as inter-frame changes in the image contrast, after mosaicking. An approach to lowering errors is to post-process the mosaic [10] via *osmosis* partial differential equation (PDE) [12, 13], which preserves details, mass and balance the lights: efficient numerical study for osmosis on large images is proposed [2, 11], based on operator splitting [8]. Our range of Cultural Heritage applications include the detection of sub-superficial voids in *Monocromo* (L. Da Vinci, Castello Sforzesco, Milan) [5], the light-balance for multi-spectral imaging and the data integration on the Archimedes Palimpsest [10].**

1 Introduction

Nowadays, image processing is an active research field in Cultural Heritage (CH) conservation. In particular, multi-spectral image analysis can reveal hidden features of artworks and help specialists during the restoration, e.g. [1, 3, 7, 9]. These non-destructive methods offer high-resolution images and large datasets to be processed via efficient algorithms, eventually.

In this work, we describe a new multi-modal mid-infrared (MWIR, 3–$5\,\mu\mathrm{m}$) imaging technique [6, 5] to detect sub-superficial defects in fresco walls, by combining the emitted and reflected information recorded by a thermal camera with suitable lighting sources. Here, multiple images are acquired and mosaicked together so as to form a large MWIR mosaic. However, the experimental setup may introduce measurement errors, seen in our applications as inter-frame changes in contrast (i.e. constant *shadows*). Thus, errors are lowered via a recently introduced drift-diffusion equation [12, 13], called *osmosis*, already able to remove constant shadows in images. We implemented osmosis efficiently for large images (up to 30 MP), with standard numerical splitting approaches [8, 2].

Applications. We discuss the dual-mode workflow for the detection of sub-superficial voids and cement patches in the restoration of *Monocromo*, a fresco wall by Leonardo Da Vinci (Castello Sforzesco, Milan). Results are validated by specialist restorers from *Opificio delle Pietre Dure* (Florence).

We applied the efficient osmosis model for different CH images acquired in different electromagnetic bands, where a light balance post-processing enhances the information: MWIR thermal quasi-reflectography, UV fluorescence and IR falsecolor imaging. We also discuss the use of osmosis for a multi-spectral data-fusion on the Archimedes Palimpsest.

Organisation. In Section 2, we describe the core idea of the multi-modal MWIR workflow for the detection of mural defects. In Section 3, we recall the osmosis equation and the splitting approach for an efficient numerical solution of the light balancing problem. In Section 4, we present other applications of the osmosis model in different CH imaging problems.

2 Dual-mode mid-infrared imaging

In [4], a new MWIR methodology called *Thermal Quasi-Reflectography* (TQR) is proposed for revealing more sub-superficial mural features than traditional approaches. At the core of TQR there is the observation that an object, at room temperature, emits in the MWIR spectrum only 1% of its thermal energy: thus, its (quasi-)reflected radiation ρ can be recorded by a MWIR thermal camera and suitable light sources.

In [6], this idea is evolved into a dual-mode (reflected plus emissivity) MWIR infrared approach. Here, TQR is the first step of the workflow since it provides the emissivity $\varepsilon = 1 - \rho$ via the simplified Kirchhoff's law. Such ε is then plugged in the recording of a thermal cooling down sequence in the emissivity mode. Both steps are performed by the same thermal device in a fixed geometry lighting setup, tuned safely for artworks.

Thus, the dual-mode complementary information can be superimposed to unveil hidden mural features, e.g. different cooling down profiles. Also, TQR details allow to register the two aligned datasets onto a visible orthophoto, an impossible task for the emissivity-only mode due to the heat diffusion blurring.

In Fig. 1 we report (top row) the visible fresco, the reflected TQR, the emitted thermal images and the fusion of the complementary data; we also show (bottom row) the discovering of sub-superficial cement patches via the dual-modal approach.

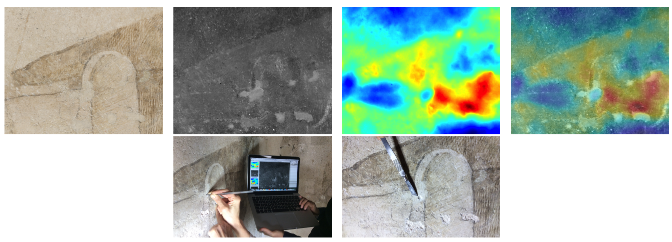

Figure 1: Dual-mode MWIR imaging for discovering hidden cement patches.

In Fig. 2, sub-superficial voids are effectively detected: different void shapes at sub-millimetric precision are revealed and measured with pins by restorers of *Opificio delle Pietre Dure*.

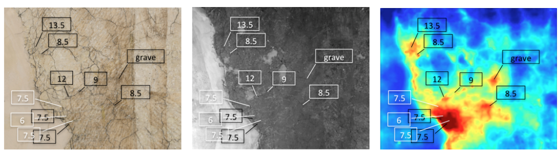

Figure 2: Defect measures (in mm.) of identified damages via dual-mode setup.

3 Efficient osmosis for light balancing

Let $\Omega \subset \mathbb{R}^2$, $u, f, v : \Omega \to \mathbb{R}$, $\mathbf{d} : \Omega \to \mathbb{R}^2$ with $\mathbf{d} = \nabla \log v$. The following model is called *osmosis* [12, 13]:

$$\begin{cases} u_t = \operatorname{div}(\nabla u - \mathbf{d}u) & \text{on } \Omega \times (0, T]; \\ u(0) = f(x) & \text{on } \Omega \text{ at } t = 0; \\ \langle \nabla u - \mathbf{d}u, \mathbf{n} \rangle = 0 & \text{on } \partial\Omega. \end{cases} \tag{1}$$

Interestingly, the non-symmetric drift-diffusion PDE (1) has nice properties, e.g. conservation of the mass, positivity and convergence to non-constant steady states, which makes its use appealing for imaging applications, when $\mathbf{d}$ is locally modified. This is the case for removing constant shadows in images, where $\mathbf{d} = 0$ on the shadow edges, i.e. the jumps produced by sharp shadows. Numerical solutions of (1) are given by solving the linear system $u_t = \mathbf{A}u$ where the 5-diagonal spatial discretisation $\mathbf{A}$ of $\operatorname{div}(\nabla u - \mathbf{d}u)$ is given by $\mathbf{A} = \mathbf{A_1} + \mathbf{A_2}$ with

$$\begin{aligned} \mathbf{A_1}(u) :=& \frac{u_{i+1,j} - 2u_{i,j} + u_{i-1,j}}{h^2} \\ &- \left(d_{1,i+\frac{1}{2},j} \frac{u_{i+1,j} + u_{i,j}}{2h} - d_{1,i-\frac{1}{2},j} \frac{u_{i,j} + u_{i-1,j}}{2h} \right) \\ \mathbf{A_2}(u) :=& \frac{u_{i,j+1} - 2u_{i,j} + u_{i,j-1}}{h^2} \\ &- \left(d_{2,i,j+\frac{1}{2}} \frac{u_{i,j+1} + u_{i,j}}{2h} - d_{2,i,j-\frac{1}{2}} \frac{u_{i,j} + u_{i,j-1}}{2h} \right). \end{aligned}$$

From [12, Prop. 2], $\mathbf{A}$ is irreducible, with non-negative off-diagonal entries and column sums equal to 0: thus, scale-space properties are guaranteed for Implicit and Explicit (with time-step restriction) Euler [12, Prop. 1].

For large images, a way to avoid the computational bottleneck due to the large bandwidth of $\mathbf{A}$ is to consider $\mathbf{A_1}$ and $\mathbf{A_2}$ separately, by ADI splitting [8]. Among the others, the Multiplicative Operator Splitting (MOS) scheme preserves the scale-space properties of (1) and provides effective results for large time-steps $\tau > 0$ [11]:

$$u^{k+1} = \prod_{n=1}^{2} \left(\mathbf{I} - \tau \mathbf{A}_n \right)^{-1} u^k, \tag{MOS}$$

In Fig. 3 we applied the osmosis filter to a TQR mosaic obtained from the workflow described in Section 2. Thus, we aimed to balance the inter-frame contrasts while preserving intra-frame details: this helped restorers in the visual inspection and comparison of hidden features of the mural painting.

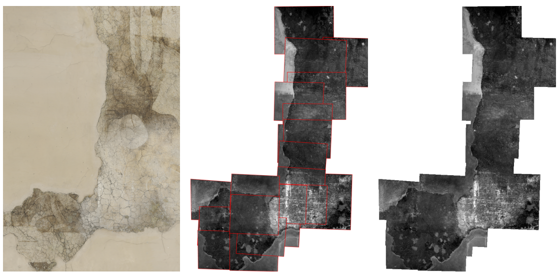

Figure 3: *Monocromo*, L. Da Vinci. Visible (left), TQR mosaicked (middle), osmosis result (right, 28 MP processed in 629 s. via MOS splitting)

4 Other CH applications

In CH imaging there exists other situations in which multiple images are sampled under different light conditions, which may introduce errors and affect the quality of the final image to inspect. Again, the use of an efficient numerical scheme of the osmosis model (1) makes the task of homogenizing the light differences possible in a reasonable time. In Figs. 4 and 5 we report a case study for the UV fluorescence inspection and the falsecolor imaging (IR-Red-Green) of a Russian icon.

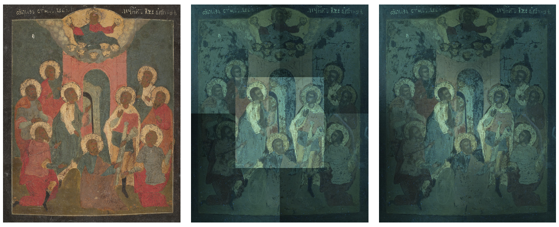

Figure 4: Light balancing in UV Fluorescence. Visible (left), 5 UV Fluorescence tiles, result (right, 18 MP and 3 color channels, 1693 s. [10])

Figure 5: Light balancing in IR falsecolor imaging. IR before and after osmosis (left and middle-left), IR-R-G before and after osmosis (middle-right and right, detail of 18 MP, 721 s. [10])

In Fig. 6 we report a detail of the *Archimedes Palimpsest*, a X century copy of the works by Archimedes overwritten in XIII century. Here we used the property of (1) to converge to a rescaled version of v, encoded in $\mathbf{d}$, to fuse the original written text with the current colours.

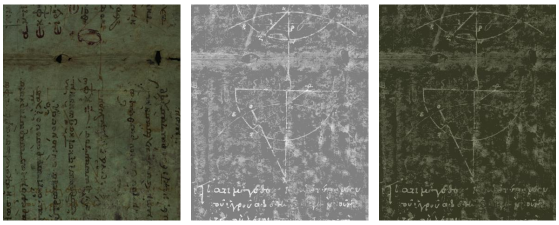

Figure 6: Data integration: *Archimedes Palimpsest* (detail, 2 MP, 137 s. [10]). Visible parchment (left), hidden text (center), data fusion (right).

5 Conclusion

In this work, we highlighted the dual-mode MWIR analysis for the restorations of frescoes [6]. Also, we tested an efficient numerical approach of the osmosis equation [12, 13] for a variety of CH multi-spectral applications, from balancing lights in multi-spectral images to data fusion on Archimedes Palimpsest [10].

Data statement. The *Monocromo* data are sensitive: access subjected to *Soprintendenza Castello Sforzesco*, Milan. The Archimedes Palimpsest[1] is released under CC-BY 3.0 license.

[1] http://openn.library.upenn.edu/Data/0014/ArchimedesPalimpsest

References

[1] J.R.J. van Asperen de Boer. “Infrared reflectography: a method for examination of paintings”. Applied Optics. **7**(9):1711-14, 1968.

[2] L. Calatroni, C. Estatico, N. Garibaldi and S. Parisotto, “Alternating Direction Implicit (ADI) schemes for a PDE-based image osmosis model”, Journal of Physics: Conference Series, **904**(1):012014, 2017.

[3] C. Daffara, R. Fontana. “Multispectral infrared reflectography to differentiate features in paintings”. Microscopy and Microanalysis **17**(5):691-5, 2011.

[4] C. Daffara, D. Ambrosini, L. Pezzati and D. Paoletti. "Thermal quasi-reflectography: a new imaging tool in art conservation". Opt. Express, OSA **20**(13):14746–14753, 2012.

[5] C. Daffara, S. Parisotto, P. I. Mariotti. “Mid-infrared thermal imaging for an effective mapping of surface materials and sub-surface detachments in mural paintings: integration of thermography and thermal quasi-reflectography”. Optics for Arts, Architecture, and Archaeology V, Proc. SPIE **9527**, 2015.

[6] C. Daffara, S. Parisotto, D. Ambrosini. “Multipurpose, dual-mode imaging in the 3–5 µm range (MWIR) for artwork diagnostics: A systematic approach”. Optics and Lasers in Engineering, **104**:266-273, 2018.

[7] D. Gavrilov, R.G. Maev, D.P. Almond. “A review of imaging methods in analysis of works of art: thermographic imaging method in art analysis”. Journal of Physiscs **92**(4):341-64, 2014.

[8] W. Hundsdorfer and J.G. Verwer. "Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations". ISBN 9783540034407, Springer Series in Computational Mathematics, 2007.

[9] X.P. Maldague. “Theory and practice of infrared technology for nondestructive testing”. ISBN 978-0-471-18190-3, New York: Wiley-Interscience, 2001.

[10] S. Parisotto, L. Calatroni and C. Daffara. “Digital Cultural Heritage imaging via osmosis filtering”. ArXiv e-prints, 1802.05468, 2018 (accepted).

[11] S. Parisotto, L. Calatroni, M. Caliari and C.B. Schönlieb. “Dimensional splitting and anisotropic modelling of a drift-diffusion osmosis equation for imaging applications”. (in preparation).

[12] O. Vogel, K. Hagenburg, J. Weickert, S. Setzer. A Fully Discrete Theory for Linear Osmosis Filtering. In LNCS Space and Variational Methods in Computer Vision, (Springer), **7893**:368-379, 2013.

[13] J. Weickert, K. Hagenburg, M. Breuß, O. Vogel. “Linear Osmosis Models for Visual Computing”. In LNCS EMMCVPR, Springer, **8081**:26-39, 2013.

Un-mixing X-Ray Images with the Application in Art Investigation

Zahra Sabetsarvestani*, Francesco Renna†, Franz Kiraly* , Miguel R. D. Rodrigues*
*University College London, U.K.
†Instituto de Telecomunicações and Faculdade de Ciências da Universidade do Porto, Portugal

Fig. 1: Panels from the Ghent Altarpiece: (left) open panel, (centre) closed panel, (right) corresponding X-ray images containing a mixture of components.

Historical paintings – very valuable objects of our cultural heritage – have motivated intensive scientific research in order to gain a more profound insight about their creation and restoration history as well as optimize conservation and preservation processes. Currently, scientific research in art investigation is often aided by various non-invasive imaging processes, such as infrared imaging, X-ray imaging, hyper-spectral imaging, and more, to aid the analysis of brushstrokes, canvas thread counting, digital in-painting of cracks, to name a few [1]. In particular, X-ray scans are able to penetrate through the painting, revealing information about the inner structures of the painting including underpaintings and underdrawings, composition of the materials and cracks in the different layers. [1].

In this work, we concentrate on X-ray unmixing problems arising in art investigation applications. In particular, in view of the fact that the X-ray scan of a double-sided painting consists of the super-position of individual X-rays corresponding to each side of the painting, we propose a new multi-modal image separation algorithm to unmix the mixed X-ray onto its constituents by leveraging the availability of other image modalities as side information – such as visual scans –to aid the separation task in hand.

A well-known piece of artwork containing double-sided painted panels is the Ghent Altarpiece[1] shown in Fig. 1. In particular, Fig. 1 depicts the X-ray images from a double sided panel of the Ghent altarpiece exhibiting features associated with the painting in the front and rear panels.

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

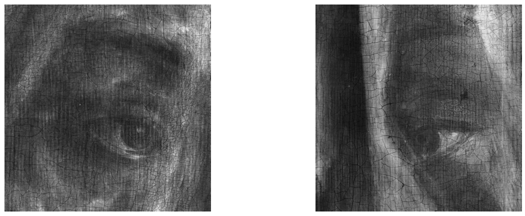

Fig. 2: X-ray images cropped from double-sided panels of the altarpiece.The resolution is 1024×1024 pixels.

Our approach is based on the use of a joint Gaussian mixture model (GMM) to relate an X-ray scan to a visual scan of a painting. This modelling approach – which has been motivated by the fact that the use of GMMs can provide state-of-the-art results in various image processing applications [2]– exhibits various advantages: 1) First, the training process can be done very efficiently using effective algorithms such as expectation maximization, 2) Second, the image separation process can also be done very efficiently using simple closed form expressions, and 3) finally, this method exhibits various performance gains in relation to other state-of-the-art image separation approaches [2].

Finally, we use our algorithm to separate the X-ray images obtained from double-sided panels in the well-known Ghent Altarpiece dataset shown in Fig. 2.

Here, the parameters of the joint GMM – the mean and the covariance matrix of the Gaussian components – are learned from X-ray image patches and their corresponding visual scans taken from single-sided panels of the Ghent Altarpiece, via the Expectation Maximization (EM) algorithm. The learned parameters are then used to separate the mixture of X-rays from double sided panels using a simple closed-form expression.

The results of our algorithm are depicted in Fig. 3 and Fig. 4. We compare our algorithm against other state-of-the-art algorithms in separating X-ray images, coupled dictionary learning [1] and morphological component analysis (MCA) [3]. We use the same training data from single sided panels to learn the coupled dictionaries. We follow the setting in the original work on the same dataset [1] and report the best performing multi-scale approach. It is worth mentioning that, as opposed to the the results reported in [1], we do not remove the crack patterns from the painting images. For MCA

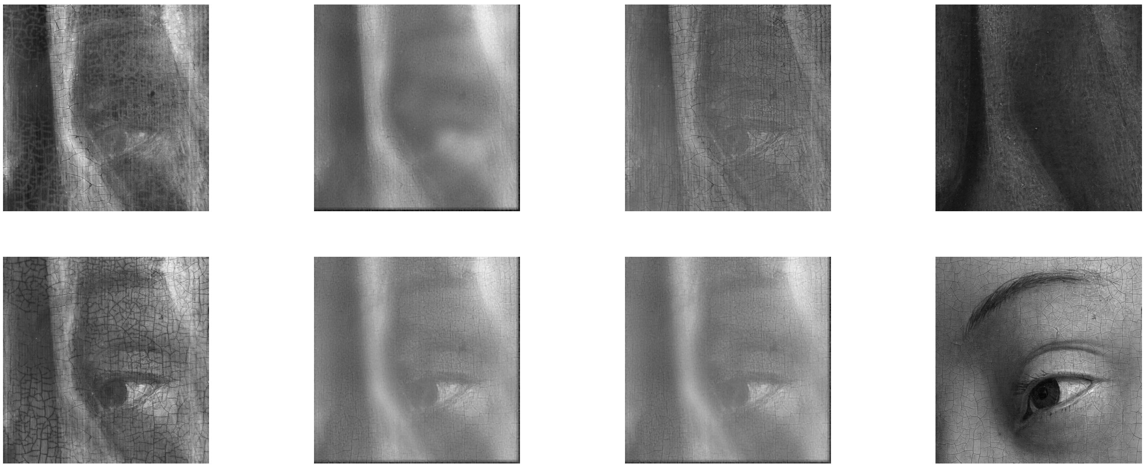

Fig. 3: Visual evaluation of the proposed algorithm in comparison with other algorithms [1], [3] in the separation of X-ray images; first row separated side 1; second row separated side 2; first column, the results of the proposed algorithm; second column, the results of the multi-modal dictionary learning algorithm; third column, the results of the MCA algorithm with fixed dictionaries; fourth column the grey scale image of each side.

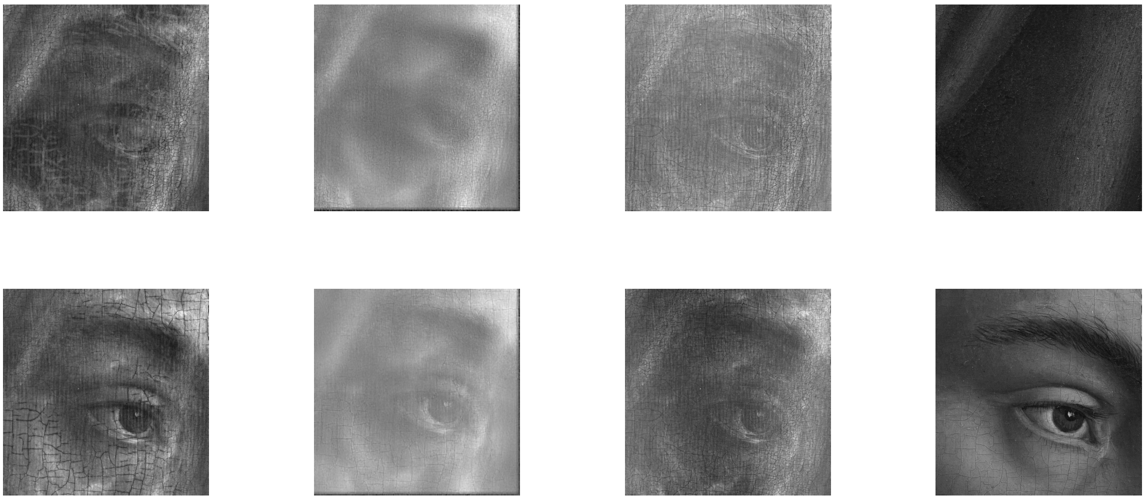

Fig. 4: Visual evaluation of the proposed algorithm in comparison with other algorithms [1], [3] in the separation of X-ray images; first row separated side 1; second row separated side 2; first column, the results of the proposed algorithm; second column, the results of the multi-modal dictionary learning algorithm; third column, the results of the MCA algorithm with fixed dictionaries; fourth column the grey scale image of each side.

algorithms we use fixed curvelet and wavelet dictionaries.

Due to lack of ground truth for this dataset, a quantitative evaluation of the performance of the algorithms is not available and the comparison will be based on the visual evaluation of the results. It is crucial to note that, training the GMM model can be done very efficiently using the effective EM algorithm and the source separation can be done very effectively using simple closed form expression; on the contrary the coupled dictionary based algorithm which uses a modified OMP algorithm for dictionary learning and reconstruction tends to incurred in an increased computationally complexity when the dimensions of the dictionary grow. Our results show our GMM based algorithm is more effective than coupled dictionary learning based algorithms [1] in the separation of X-ray images.

References

[1] N. Deligiannis, J. F. C. Mota, B. Cornelis, M. R. D. Rodrigues, and I. Daubechies, "Multi-modal dictionary learning for image separation with application in art investigation," *IEEE Transactions on Image Processing*, vol. 26, no. 2, pp. 751–764, Feb 2017.

[2] F. Renna, L. Wang, X. Yuan, J. Yang, G. Reeves, R. Calderbank, L. Carin, and M. R. D. Rodrigues, "Classification and reconstruction of high-dimensional signals from low-dimensional features in the presence of side information," *IEEE Transactions on Information Theory*, vol. 62, no. 11, pp. 6459–6492, Nov 2016.

[3] J. Bobin, J. L. Starck, J. Fadili, and Y. Moudden, "Sparsity and morphological diversity in blind source separation," *IEEE Transactions on Image Processing*, vol. 16, no. 11, pp. 2662–2674, Nov 2007.

DOCUMENTATION AND DIGITALIZING OF ROYAL ALBUMEN PHOTOGRAPHIC COLLECTION OF KING FAROUK, DATING FROM 19TH CENTURY

*Rasha Ahmed SHAHEEN [*1], Mona Fouad ALI [2], Medhat EL-DABAA[3]*

[1] Conservation Department, Egyptian Textile Museum, Ministry of Archaeology, Egypt.
[2] Conservation Department, Faculty of Archaeology, Cairo University, Giza, Egypt.
[3] Archaeology Conservation dept., Faculty of Archaeology, South Valley Univ., Egypt.

Abstract

This paper presents the application to treatment and conservation of a set of albumen photographic print-out that exist within King Farouk collections and kept in Royal of Vehicles Museum, dating from 19th century. Museum contains five albums of photographs, and these photos preserved in a good way somewhat distant from the sources of moisture, which helped to save them in good condition. Archaeological and photographic documentation was carried out. The optical microscope was tested for damage to the surface of the album and photographs. The scanning electron microscope (SEM-EDX) was used to examine the surface of the photograph, using an elemental analysis unit to determine the constituents of the black and white photograph. The infrared spectrum analysis (FTIR-ATR) was done to identify the chemical composition of the medium used for depositing light-sensitive silver salts. The test was carried out using UV photography to ensure that there were no watermarks on the paper used as a secondary support for the photographs.

1. Introduction

Historical Vehicles Museum is a one of the rarest historical museums in the world. The museum dates back to the reign of Khedive Ismail Pasha, who ruled Egypt (1879 - 1863). It was the first thought of in the construction of a special building Khedive vehicles and horses. This phase ended when the royal family was leaving the rule that the building was rehabilitated to be a historical museum for vehicles, as it contains vehicles and chariots of King Farouk I. All vehicles documenting by photographs to put it in the fifth albums which sealed the royal stamp. This paper presents the practical study, include the treatment and conservation of a set of albumin photographic print-out, these photos preserved in a good way somewhat distant from the sources of moisture, which helped to save them in good condition. It was selected as one of this royalty albums and make a plan for preservation.

2. Documentation and Visual Investigation

Visual inspection of the album turned out to be in stable condition and good with some dirt in the sides and edges are shown in the following figures:

Fig1: Interface of Ismail Pasha album

Fig2: Royal Logo

Fig3: Flake and damage of the edges

Fig4: Dirty in the edges of the paper

Fig5: Digitalizing of the albumin photos

Fig6: First page of the Royal album

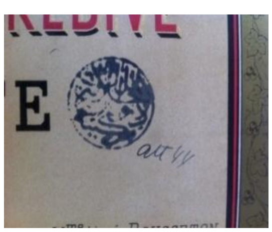

Fig7: Royal Stamp

Fig8: second page which stat the photos for service

Fig9: The opposite age of the first one, and we can see the effect of ink

Fig10: Separation in the connector section

Fig11: Type of photo that takes in one shot

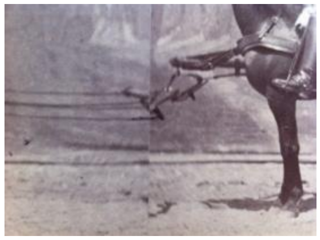

Fig12: Completing in the area the bond between the two images

Fig13: Type of photos that takes in two shots merge they in one photo

Fig14: Another type of photo

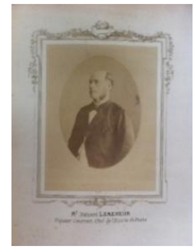

Fig15: Portiere for head of serves

3. Digital Microscope

The album checked by using digital microscope connected by computer (USB Digital Microscope) strongly enlarge 300 ×.

Fig17: First photo, surface of paper, surface of photo, we were note the shape of albumin.

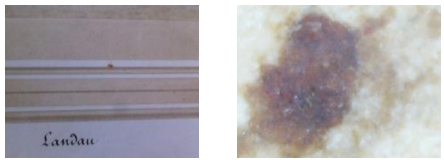

Fig18: Spot of steel the result of the use of iron ink.

Fig19: Peeling and color change on the surface of the album as a result of exposure to light.

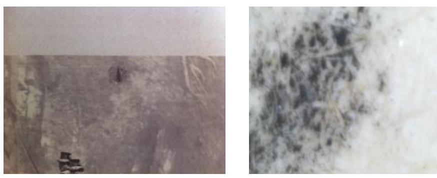

Fig20: Spot of silver halides due to manufacturing.

Fig21: Ink in Second page.

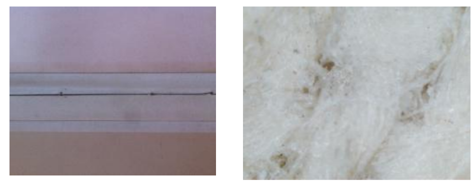

Fig22: Textile structure of the joint between the user pages.

4. Examination and Analysis Methods of Albumin

4. 1. SEM-EDX

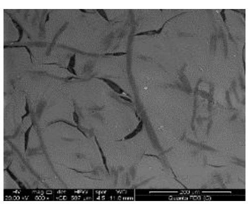

Fig23: SEM (strongly enlarge 200 ×) of an albumin sample of photographs.

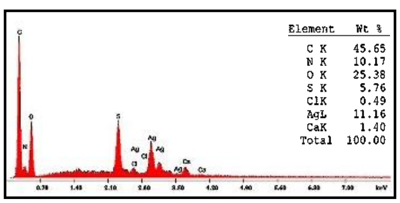

Fig24: Pattern analysis of (SEM-EDX) of a sample of the emulsion albumin Photo.

4.2. FTIR-ATR

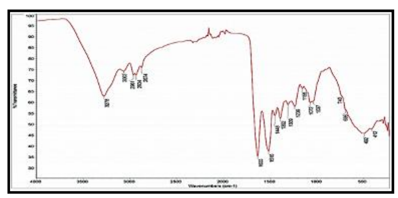

Fig25: The spectrum of Albumin on photographs measured by ATR technique.

Functional Group	Characteristic Absorption(s) (cm^{-1})
N-H stretching	3429 vs
N-H stretching (Amide)	3673 m
N-H stretching (Amide)	3715 m
CH stretching	2918 vs
Amide I (C=O stretching)	1632 m
Amide II (N-H bending)	1537 m
C=O stretching	1428 s
C=N stretching	2356 m
N-H bending	1651 m
CH bending	1055 s
CH bending	1372 w
CH rocking	560 m
C-O-C (ester) stretching	1113 m

Table1: FTIR-ATR characteristic ab-sorptions bands of white egg.

5. Conclusion

Digitalizing of photographic prints by copying and duplicating is one of the principal way to conserving photographic prints. Storage of albumin photographic prints in a museum or place closed to the river makes it susceptible to insecticide as a result of high humidity and salt levels in the surrounding storage area.

The Albumin thin films on photographic materials were investigated by optical microscope, SEM-EDX, and FTIR-ATR spectroscopy. Egypt was among the first countries that used photography to document the vehicles which used in all special services for khedive and keep it in royal album.

Reference

[1] Blackmar. J. 2002. *The Importance & Care of Archeological Records*, Nebraska State Historical Society, Pp. 8-13.

[2] Bugner. 2006. "*A survey of environmental conditions relative to the storage and display of photographs in consumer homes*", Vol.50, No.4.

[3] Fischer. M. 2008. "*Fundamentals of the Conservation of Photographs, Introduction to Collection-level Surveys and Condition Reporting*", Getty Conservation Institute, Pp.1-6.

[4] Lavédrine B. 2016. "*A Guide to the Preventive Conservation of Photograph Collections*", Getty.

A deep learning approach to crack detection in panel paintings

R. Sizyakin[1], B. Cornelis[2], L. Meeus[3], M. Martens[4], V. Voronin[1] and A. Pižurica[3].

[1]Department of Radio-Electronics Systems, Don State Technical University, Russian Federation .

[2]Department of Electronics and Informatics, Vrije Universiteit Brussel, Belgium.

[3]Department of Telecommunications and Information Processing, TELIN-IPI-imec, Ghent University, Belgium.

[4]Department of Art, Music and Theatre Sciences, Ghent University, Belgium.

Abstract— **The accurate detection of cracks in paintings, which generally portray rich and varying content, is a challenging task. Furthermore, traditional crack detection methods are often not well suited for recent acquisitions of paintings as they are not designed for high-resolution images and do not fully exploit the information from the different imaging modalities at hand. In this work, we propose a fast crack detection framework that alleviates the aforementioned challenges. The method consists of a morphological filtering operation followed by a classification step by means of a convolutional neural network architecture. The proposed online method is capable of continuously learning from newly acquired visual data, thus further improving classification results as more data becomes available.**

1 Introduction

Paint cracking (or *craquelure*) is the most common type of deterioration encountered in old master paintings. Generally speaking, cracks appear in paint layers when pressure develops within or on it through the influence of various factors and cause the material to break [1]. The automatic detection of crack patterns is desirable for many reason. Most importantly, crack patterns can offer insights on the structural condition and conservation history of a painting [1]. Crack detection is also used as a preprocessing step for the digital restoration of paintings [2].

Many crack detection methods have been developed over the recent years, see e.g. [2–6] and the reference therein. Still, some important challenges remain especially in terms of feature selection, which often has to be adapted for different paintings, parameter tuning, and complexity, which limit practical applicability.

In this paper, we propose a new deep learning based approach for crack detection in paintings. The method consists of two processing stages: (i) a morphological filtering stage, and (ii) a classification stage. The morphological filtering essentially ensures that the amount of pixels to be classified by the CNN in the second stage is strongly reduced as only those pixels that are similar in structure to cracks are selected. In the classification stage, we employ a convolutional neural network (CNN). CNNs demonstrated recent success in many application where they have outperformed, often by a substantial margin, traditional machine learning algorithms [7–10]. We are not aware of any reported works that apply CNNs to crack detection in paintings. Some recent works applied CNN to detect cracks in roads [11, 12]. Our problem is, however, much more challenging not only because of the huge variability of cracks in paintings but also due to complex background and the fact that some painted details can closely resemble cracks. Therefore, our approach needs to incorporate multimodal data, which together with huge spatial resolution of digitized paintings poses additional challenges for the classifier.

2 Proposed approach

The first step in the proposed method is morphological filtering of the available modalities. This operation creates a preliminary crack map in a similar way as in [4, 5]. Each filtered result is followed by a thresholding step, producing binary images. The threshold is set based on the method of Otsu [13]. The binary maps are then combined into a single one using the logical OR operation. The morphological filtering step improves greatly the classification speed. The classifier is run only on pixels marked as crack in this first stage.

The input of our classification network consists of tensors of size $m \times m \times N$. These tensors are formed by concatenating $m \times m$ sized patches extracted from the N considered modalities (the three color channels of the visual macro-photographs, the single-channel infrared macro-photographs and X-ray images, along with their grayscale morphologically filtered results add up to $N = 9$ modalities.[1]). An input sample is represented by the tensor $x(u_1, u_2, \upsilon_0) \in R^{m \times m \times N}$, where u_1, u_2 are spatial coordinates and υ_0 is the index that identifies the chosen modality. For our experiments, we fix $m = 8$, resulting in tensors with dimensions of $8 \times 8 \times 9$. The convolutions over υ_0 are calculated in the first layer of the CNN as follows:

$$x_1(u_1, u_2, \upsilon_1) = \rho\big(x(u_1, u_2, \upsilon_0) * w_{\upsilon_1}(u_1, u_2, \upsilon_0)\big), \qquad (1)$$

where $x_1(u_1, u_2, \upsilon_1)$ is the feature map obtained by the convolution of $x(u_1, u_2, \upsilon_0)$ with $w_{\upsilon_1}(u_1, u_2, \upsilon_0)$, indexed by υ_1 (in our architecture $1 \leq \upsilon_1 \leq 12$ for the first, $1 \leq \upsilon_2 \leq 24$ for second and $1 \leq \upsilon_2 \leq 48$ for the third convolution layers), and where ρ is an activation function. We choose the well-known rectified linear unit (ReLU) [7], defined as $\rho = max(0, x)$. All kernels are initialised randomly in the beginning of the training procedure. The core of our CNN architecture consists of performing a cascade of convolutions at each layer $j \geq 2$ as follows:

$$x_j(u_1, u_2, \upsilon_j) = \rho\big(x_{j-1}(\cdot, \upsilon_{j-1}) * w_{\upsilon_j}(\cdot)\big), \qquad (2)$$

where, as we navigate through the subsequent layers, the resolution of $x_j(u_1, u_2, \upsilon_j)$ progressively reduces [14, 15]. The last layer of the architecture consist of a *softmax* function, which ensures that all class probabilities sum up to 1.

[1] A disk-shaped structuring element with a diameter of 5 pixels was used for all experiments.

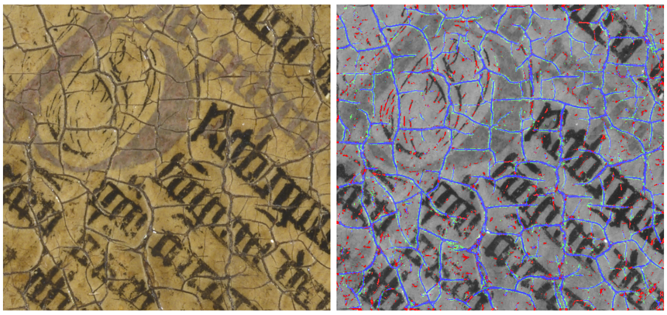

Figure 1: **Left**: A fragment of the panel *Virgin Annunciate*. **Right**: Results of the proposed method in comparison with BCTF. Blue: cracks identified by both methods; red: cracks detected by BCTF only; green: cracks detected by the proposed method only.

3 Experimental Results

All experiments are performed on a high-resolution multi-modal dataset of the *Ghent Altarpiece*, publicly available on the *Closer to Van Eyck* website[2]. In particular, we focus on three panels of the polyptych, named *Virgin Annunciate*, *Singing Angels* and *John Evangelist*. For comparsion we use the Bayesian Conditional Tensor Factorisation (BCTF) method from [3], which was also evaluated and compared to other methods in [6].

Figure 1(Left) shows crack detection results on a part of the panel *Virgin Annunciate* which is particularly challenging because some painted features resemble cracks. The results displayed in Fig. 1(Right) show that BCTF falsely labels a significant amount of pixels (such as the decorative elements around the big letter "P"), while our method successfully differentiates true cracks from those painted features. It should be noted that both methods use the same image modalities.

Similar conclusions follow from the results on the panel *Singing Angels* displayed in Fig. 2(Top). In general, the proposed method detects more cracks while reducing false detections. Furthermore, the proposed method can be trained in an online fashion, i.e. without re-training the whole network, continuously improving detection results (Fig. 2(Bottom)). An important asset is also rapid processing, especially once the network is trained. This makes our framework integrable in a fast and interactive tool that can be used by art professionals.

Figure 3 depicts results on a small area from the panel of *John the Evangelist*. For this example, only one modality was considered, namely the color photograph, for both methods. The BCTF method was trained specifically for this image, while the proposed method was pre-trained on other panels and only a small fraction of labels (approximately 3,000 patches of two types) from the present panel were used to re-train the network. This experiment indicates that the proposed method can be deployed for a variety of paintings, with relative little effort.

In general, the proposed method demonstrates potential to improve upon the current state-of-the-art methods by detecting more cracks while also reducing false detections. Furthermore, the proposed method can be trained in an online fashion, i.e. without re-training the whole network, continuously improving detection results. An important asset is also rapid processing, especially once the network is trained. This makes our framework integrable in a fast and interactive tool that can be used by art professionals.

[2]Link: http://closertovaneyck.kikirpa.be/

Figure 2: **Top**: A fragment of the panel *Singing Angels*. **Bottom**: Results of the proposed method in comparison with BCTF. Blue: cracks identified by both methods; red: cracks detected by BCTF only; green: cracks detected by the proposed method only.

4 Conclusion

In this paper, we propose a novel crack detection framework capable of handling acquisitions from different modalities. In a first step, we apply morphological filtering for a coarse initial identification of crack pixels. This step substantially reduces the amount of data to be processed later on. We then train a CNN architecture with user labelled data to further refine the results obtained in the first step. We show that our method improves upon the current state-of-the-art in this application. An additional advantage is the possibility of re-training the network using newly available data. This feature allows to improve the already obtained result without significant time costs.

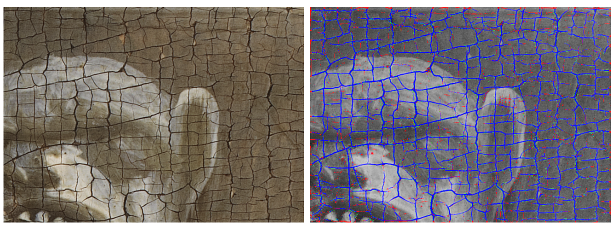

Figure 3: **Left**: A fragment of the panel *John Evangelist*. **Right**: Results of the proposed method (pre-trained on other panels and re-trained with few labels from the present panel) in comparison with BCTF (trained for this particular image). Blue: cracks identified by both methods, red: cracks detected by BCTF only, and green: cracks detected by the proposed method only.

References

[1] J. Mohen, M. Menu, and B. Mottin, "Mona lisa: Inside the painting," *Harry N. Abrams, New York, NY, USA,*, 2006.

[2] B. Cornelis, T. Ružić, E. Gezels, A. Dooms, 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 process*, vol. 93, no. 3, pp. 605–619, 2013.

[3] B. Cornelis, Y. Yang, Joshua T. Vogelstein, A. Dooms, I. Daubechies, and D. B. Dunson, "Bayesian crack detection in ultra high resolution multimodal images of paintings," *18th International Conference on Digital Signal Processing, IEEE*, 2013.

[4] A. Gupta, V. Khandelwal, A. Gupta, and M. C. S. Thammasat, "Image processing methods for the restoration of digitized paintings," *Thammasat Int. J. Sc. Tech*, vol. 13, no. 3, 2008.

[5] I. Giakoumis, N. Nikolaidis, and I Pitas, "Digital image processing techniques for the detection and removal of cracks in digitized paintings," *IEEE Trans. Image Processing*, vol. 15, no. 1, pp. 178–188, 2006.

[6] A. Pižurica, L. Platiša, T. Ružlić, B. Cornelis, A. Dooms, M. Martens, H. Dubois, B. Devolder, M. De Mey, and I. Daubechies, "Digital image processing of the Ghent altarpiece: supporting the painting's study and conservation treatment," *Signal Processing Magazine, IEEE* , vol. 32, 2015.

[7] A. Krizhevsky and I. Sutskever and G. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in Neural Information Processing Systems*, pp. 1097–1105, 2012.

[8] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich, "Going deeper with convolutions," *CoRR*, vol. abs/1409.4842, 2014.

[9] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014.

[10] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015.

[11] Z. Lei, Y. Fan, D.Z. Yimin, and J. Z. Ying, "Road crack detection using deep convolutional neural," *Image Processing (ICIP), 2016 IEEE International Conference*, 2016.

[12] Y.-Jin Cha, W. Choi, and O. Büyüköztürk, "Deep learning-based crack damage detection using convolutional neural networks," *Comput.-Aided Civ. Infrastruct. Eng.*, vol. 32, pp. 361–378, 2017.

[13] N. Otsu, "A threshold selection method from gray-level histogram," *IEEE Transaction on Systems, Man, and Cybernatics*, vol. SMC-9, no. 1, pp. 62–66, 1979.

[14] S. Mallat, "Understanding deep convolutional networks," *Phil. Trans. R. Soc. A*, vol. 374(2065):20150203, 2016.

[15] J. H. Jacobsen, E. Oyallon, S. Mallat, and A. W. M. Smeulders, "Multiscale hierarchical convolutional networks," *CoRR*, vol. abs/1703.04140, 2017.

OmniArt: A Large Scale Artistic Benchmark

Gjorgji Strezoski[1] and Marcel Worring[1].

[1]University of Amsterdam, Amsterdam, Netherlands

***Abstract*— Baselines are the starting point of any quantitative multimedia research, and benchmarks are essential for pushing those baselines further. In this paper, we present baselines for the artistic domain with a new benchmark dataset featuring over 2 million images with rich structured metadata dubbed OmniArt. It contains annotations for dozens of attribute types and features semantic context information through concepts, IconClass labels, color information, and (limited) object level bounding boxes. We establish and present baseline scores on multiple tasks like artist attribution, creation period estimation, type, style, and school prediction. To accomplish this we develop multi-task deep learning pipelines that are able to train and evaluate on multiple attributes at the same time. As an example of additional types of analyses, we explore the color spaces of art through different types and evaluate a transfer learning object recognition pipeline.**

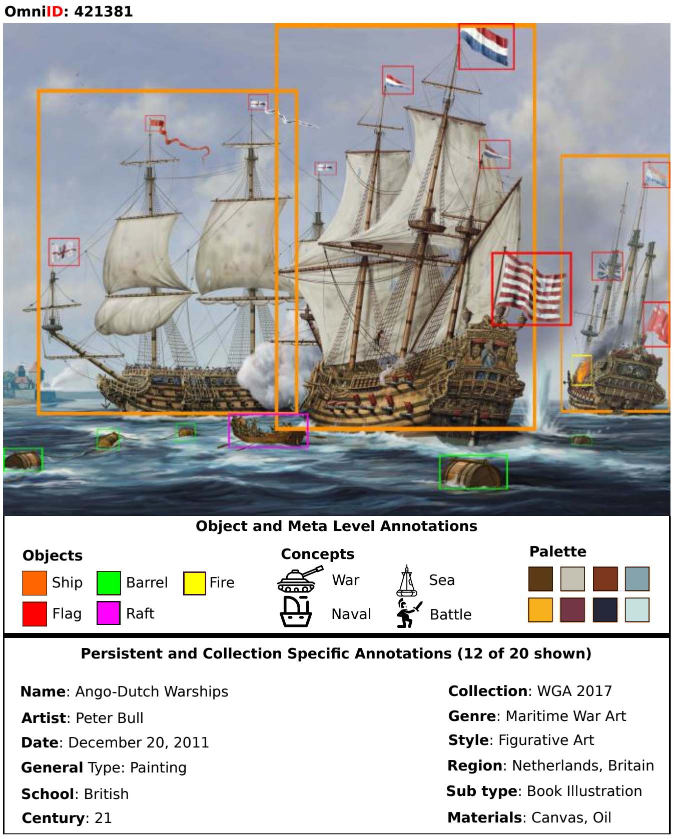

Figure 1: Sampled image from the OmniArt dataset with bounding boxes for object level annotations, key content concepts, the extracted color palette and textual metadata

1 Introduction

OmniArt features 1,348,017 indexed images with full annotations and 702,000 more unlabeled images with incomplete metadata which extends upon our preliminary 500K dataset presented in [3]. These data samples are described by four independent metadata types. Figure 1 illustrates the separate metadata types and their instantiations. The persistent and collection specific annotations are obtained from the collection of origin, while the object and meta level annotations are inserted with the help of intelligent models or manual annotation. Additionally, a modular metadata structure improves redistribution efficiency by minimizing the data volume overhead. For example, different tasks such as object recognition or creation period estimation, might require a different annotation strategy, so having distributed information can increase efficiency and reduce data processing complexity. A good example where a different annotation strategy is required, is in images annotated with the *IconClass* taxonomy. *IconClass* [1] annotations offer a hierarchical semantic description of the visual content of the artwork. They are constructed of an *IconClass* code and a suitable description. In the OmniArt dataset we have assembled a collection of more than 450.000 photographic reproductions of artworks annotated both with the persistent metadata and *IconClasses*. As one artwork can have multiple *IconClasses* (see Figure 1), having it in the same structure as the persistent metadata would increase query durations, which is a problem a modular structure addresses. This modularity of information, even at this scale fits nicely with current software engineering paradigms, thus querying and application development are simplified.

Being able to query and explore a dataset efficiently is crucial, not only for developing software solutions, but for understanding it as well. Especially with a semi-automatically generated dataset of these proportions containing metadata from various sources, an efficient tool for exploration and control is important both for quality assurance and easy access. For that purpose, an integrated exploration and annotation tool on an image and object level comes with the OmniArt dataset. Imagewise, the exploration tool provides insight into attention and saliency maps based on the task of interest [4], while the annotation tool allows for editing the persistent and meta level metadata. On the object level, with the annotation tool a user has the ability to draw, store and alter inner image bounding boxes. Each of the bounding boxes represents an area where the labeled object is present. One image can have multiple objects and multiple annotations from multiple users at the same time. We also take into account color information and enable exploration by colors. The exploration tool designed for OmniArt provides a way to navigate, browse and modify the different annotation types in a seamless process.

2 Benchmark Tasks

2.1 Experimental setup

For the purpose of the experiments and establishing the mentioned baseline scores we deploy a VGG-like architecture [2] pretrained on ImageNet and then fine-tuned for the task at hand. On top of the feature extracting neural network we deploy a multilayer perceptron for classification/prediction purposes. It

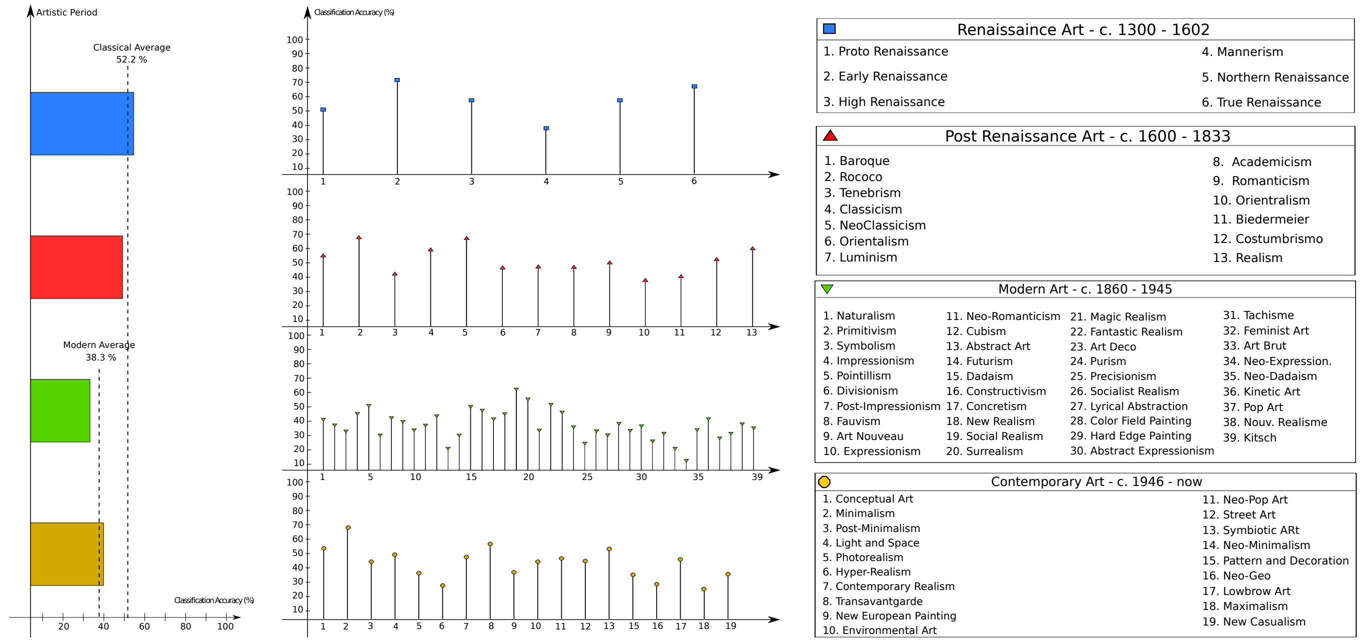

Figure 2: Style classification performance comparison between classical artistic movements and more recent ones. This evaluation shows that modern artistic styles are visually less distinguishable than classic ones.

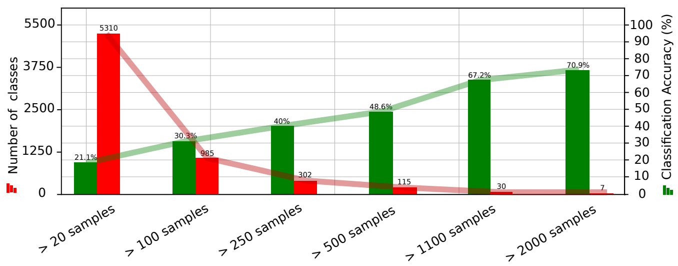

Figure 3: Artist attribution performance relative to the number of classes and number of samples in the used dataset.

consists of two fully connected layers with 1536 ReLu units each and a dropout factor of 0.2. During training we also fine-tune the final convolutional block of the network to obtain features better adjusted to the art domain. The optimization is guided by the Adam optimizer with a dynamic learning rate starting at 0.001 which decreases during the epochs. With a batch size of 96 and a square input size of 224px (random crop) in RGB we perform training for 30 epochs for each of the tasks.

2.2 Artist Attribution

In artist attribution we attempt to determine the artwork's creator based on the visual information contained in the image. For the purposes of this experiment we devised several cue points in terms of the class restrictions specific for the task. The baseline experiments were performed on selected subsets of data, as well as the whole dataset. Each subset has been limited by the minimum number of examples per class, namely 20, 100, 250, 500, 1100 and 2000. Classification performance is displayed in Figure 3.

2.3 Creation Period Estimation

Since this task can be considered both as a prediction or classification, we report scores in both settings. For the prediction task we report scores on creation year estimation, while we consider the creation century estimation both a classification and prediction task. Creation years are estimated with the mean absolute error as the function for which we optimize. In a categorical setting of the creation century estimation we softmax 84 classes, while in the prediction setting we continue to use mean absolute error. Classification performance in all 2 million images is at 18.3%accuracy, and in artworks dating after the 1500s rises to 42.8% accuracy.

2.4 Style Classification

One of the most notorious attributes when it comes to categorization is what makes an artwork beautiful - style itself. An artistic style is a collective title given to artworks which share the same artistic ideals, technical approach, context or timeline. Experimental results in Figure 2 show that our simple model is better suited in distinguishing artistic styles from the previous centuries then more recent ones. We further continue the analysis to conclude that artworks containing larger distinguishable shapes and colors in the Contemporary and Modern Art quadrants (Minimalism, Post-Minimalism, Light and Space, Symbolism) are easier for the model to distinguish from other recent styles.

2.5 Type Prediction

For the artwork type prediction task we performed predictions on both the general type and the subtype of the artworks. Since the general artwork type is a superset of the subtype it can be also considered a hierarchical classification task. There is a total of thirteen general type categories and 2449 subtypes to be taken into account, excluding the artworks for which a general type has yet to be determined. Overall classification accuracy of the type classification task is 39.9% accuracy and paintings are the visually most distinguishable type with 75.4% accuracy.

2.6 School Prediction

In the OmniArt dataset there are more than 200 registered schools on a national basis from the A.D. period. For the artworks created in the B.C. period the school attribute is unknown. The overall predictive performance of the models is 19.7% while the Italian school of painting is most distinguishable with 62.1% accuracy.

References

[1] L. D. Couprie. Iconclass: an iconographic classification system. *Art Libraries Journal*, 8(02):32–49, 1983.

[2] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[3] G. Strezoski and M. Worring. Omniart: Multi-task deep learning for artistic data analysis. *CoRR*, abs/1708.00684, 2017.

[4] G. Strezoski and M. Worring. Plug-and-play interactive deep network visualization. In *VADL: Visual Analytics for Deep Learning*, pages 0100–0106, 2017.

The Monitoring of Cracks in Historical Silver with Image Processing Techniques

Ellen van Bork[1,2], Bruno Cornelis[3].

[1] Rijksmuseum Amsterdam.

[2] University of Amsterdam.

[3] ETRO, Department of Electronics and Informatics at Vrije Universiteit Brussel - imec.

***Abstract*— Cracks have been observed in historical silver objects. The formation of these cracks is most likely the result of the alloy compositions as well as the deformation of the material. To this day, there is no standard or trusted way to identify cracks in silver and determine whether they progress over time. In this work, a number of image processing techniques, originally developed for the detection of cracks in paint surfaces, were explored and evaluated. Initial promising results show that cracks in silver can be detected using image processing techniques, opening up the possibility of monitoring their progress over time.**

1 Introduction

The 17th century was not only a *Golden Age* for Dutch painting, other areas of Dutch art such as gold- and silversmithing equally flourished. Silversmiths from the Low Countries were internationally renowned for their wrought silverwares, which show extraordinary skill in raising, embossing and chasing. Silver and gold objects by artists like Paulus van Vianen (1570-1613) and Johannes Lutma (1584-1669) demonstrate a masterly skill in the extreme manipulation of silver, unique to the Low Countries and sometimes referred to as Dutch raising. It is often assumed that silver, other than being highly sensitive to sulphur compounds, is a stable material. On several accounts however, the formation of (micro-)cracks was observed on silver objects (see Figure 1). Examples displaying this problem are the Matthias Melin (1589-1653) and van Paulus van Vianen (1570-1613) plaques in the Rijksmuseum collection.

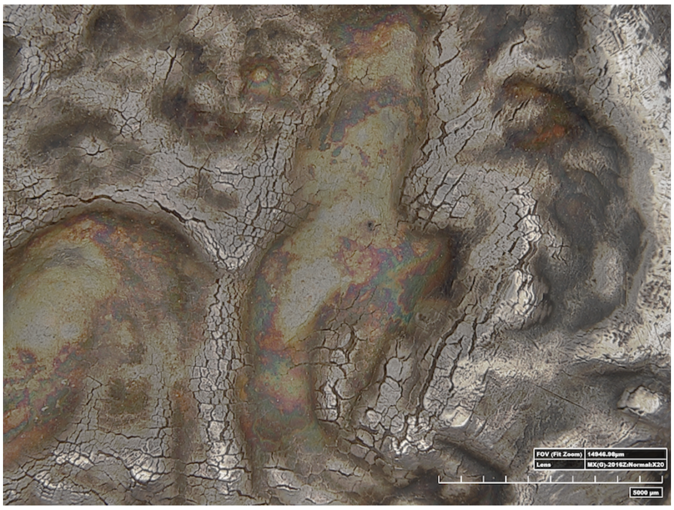

Figure 1: Detail of the back of a silver plaque displaying cracks, anonymous, 1595, Rijksmuseum collection, inv.nr. BK-1982-4-A, magnification 20×.

2 The formation of cracks

It would seem logical to attribute the cracking to the extreme deformation of the silver, however cracks were also observed on areas that have been chased in low relief. This is most likely due to the fact that historical silver alloys contain low amounts of lead (> 0,8%) as trace elements from the ore, but also from the cupellation of the silver. The lead can precipitate in the matrix and on the grain boundaries, often causing embrittlement. It is unknown if these cracks in historical silver are stable or if they are progressing over time and will start to embrittle in a couple of centuries or even decades in ambient environments, as is seen in archaeological silver.

The following research questions were formulated, with the third question being the most relevant:

1. What causes the cracking phenomena and/or embrittlement in historical artefacts?
   (a) What is the influence of trace elements?
   (b) What is the correlation between the presence of trace elements and the amount of cracks observed?
   (c) What is the influence of the manufacturing techniques used?
2. Did cracking form during manufacturing or at a later stage?
3. Are these crack patterns progressing?

Most research into silver embrittlement has been done on ancient archaeological artefact. This has shown that the burial time, the temperature, moisture content, pH and chemical composition of the burial environment, especially the salt, nitrate and nitrite content, can have a huge impact on the deterioration process and can accelerate the corrosion hereof [1]. To gain a deeper insight into this phenomenon on historical silver it is essential to closely study surviving objects from this period as well as mock-ups of historical silver alloys. For reconstruction purposes we are fortunate to have a unique and very informative Dutch manual surviving from the early 18th century; Van Laer's *Weg-Wyzer voor Aankoomende Goud en Zilversmeeden* (1721). Van Laer describes in detail, for example, the use of certain chemicals and the heat source being used during manufacturing or the alloy composition. These descriptions could provide us with essential information on why the cracking occurs. Reconstructions could be used to confirm findings gathered from objects and historical source research, making them an essential part of this research. Surface composition analyses like area scanning by XRF, could provide us with information on whether more lead is present around the grain boundaries of the cracks. X-ray computed tomography and ultrasound could provide information on the morphological and physical properties of an object.

3 Monitoring of cracks by image processing techniques

Up to this day there is no standard or trusted way to monitor the (in)stability of the cracks. Therefore, the use of image processing techniques for the detection of cracks in silver objects is being explored. Already vast amounts of literature exists on detecting crack-like patterns in digital images, often referred to as ridge-valley structure extraction. Examples include the detection of veins or vessels in medical images, fingerprint analysis, or even the segmentation of roads and rivers from satellite imagery. Common approaches include different types of thresholding, the use of multi-oriented filters and various morphological transforms. Often, the results obtained with image processing techniques are further refined by using either supervised or unsupervised machine learning algorithms. In an unsupervised setting, the algorithm attempts at further clustering pixels by means of their properties (also called features), such as colour, whether the pixel is part of an edge, or its value after a filtering operation. In a supervised setting however, the process starts with an expert user manually annotating crack pixels or areas of high crack density. These labelled pixels can then be employed to train a chosen machine learning algorithm, often yielding better results compared to the unsupervised approach.

For this application, we first rely on prior art developed for the detection of cracks in the surface of paintings. An initial result can be observed in Figure 2, where the detected cracks are marked in red. The cracks were obtained by filtering the image with oriented and elongated filters that are designed to emphasize elongated structures in images [2]. The filtered images are then thresholded and combined into a single binary image, marking the locations of the cracks in the original image.

4 Conclusion

First results show that image processing techniques can be used for the detection of cracks in silver objects. Future research will consist of further exploring the image processing and machine learning tools at our disposal for the detection of cracks in different images of silver objects at various zoom levels. Also, it will be important to develop a standardized way to compare the results over time, to determine whether the cracks are in fact progressing.

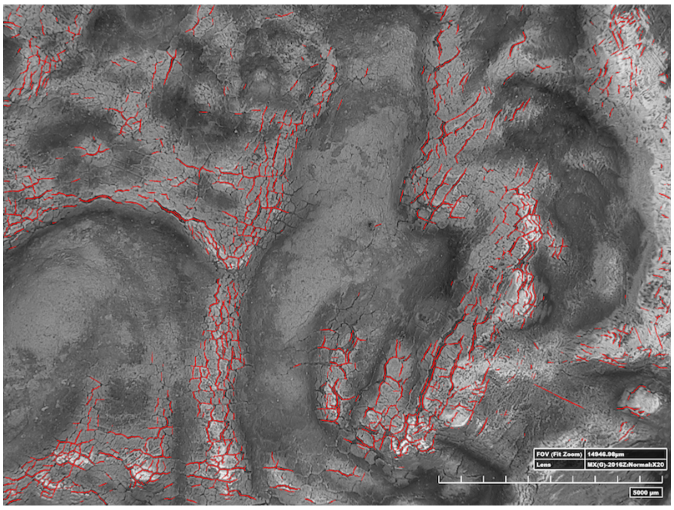

Figure 2: Initial results of crack detection using the oriented filters of [2].

References

[1] Wanhill, R.J.H., *"Ancient silver embrittlement: significances of copper, lead and cold-deformation"*, National Aerospace Laboratory (NLR), Amsterdam-NL: 7, 2003.

[2] Cornelis, B., Ruzic, T., Gezels, E., Dooms, A., Pizurica, A., Platisa, L., Cornelis, J., Martens, M., De Mey, M. and Daubechies, I., *Crack detection and inpainting for virtual restoration of paintings: The case of the Ghent Altarpiece*, Signal Processing, 93(3), pp. 605–619, 2013.

Imaging Ancient Chinese Ivory Puzzle Balls: Deducing the make process

R. van Liere[1], C-L. Wang[2] , A. Kostenko[1], I. Garachon[2], and K.J. Batenburg[1].

[1]Centrum Wiskunde en Informatica, Amsterdam. [2]Rijksmuseum, Amsterdam .

***Abstract*— In this paper we address questions related to the make process of ancient Chinese ivory puzzle balls. The approach we have taken is to image the puzzle balls using x-ray scanning and image processing to measure morphological properties of the balls. From the measurements, we can deduce the size and shapes of the tools used to carve the ivory balls.**

1 Introduction

Ancient Chinese ivory puzzle balls are known for their beauty, finesse and their ability to arouse the curiosity of the viewer. Puzzle balls consist of several concentric sphere shaped "layers". Each layer can rotate freely and has a surface carved with an ornate decorative pattern. In the 18th century, ivory balls were crafted starting from a single block of ivory using only a lathe and a collection of sharp knives and L-shaped scalpels.

The Rijksmuseum have two ivory balls in the collection. AK-NM-7020 (ca 1780) contains 9 concentric layers and has a radius of 4.3 cm. Ball AK-NM-7019 (dated ca 1750) has a radius of 8 cm.

Figure 1: Rijksmuseum archive photographs of AK-NM-7020 (left) and AK-NM-7019 (right). The chains, used to hang balls from ceilings, are also shown. Both balls have 14 'peepholes' through which the enclosed layers can be seen. In the case of AK-NM-7020, the craftsman has covered 11 peepholes on the outer layer with ivory carved caps. On AK-NM-7019, a cap has placed on the top peephole.

The goal of this work is to develop multi-scale acquisition, image processing, pattern recognition, and visualization techniques that measure morphological properties of ivory puzzle balls. These properties are used to deduce the make process of the balls.

2 Methods

Computer tomography scans of the ivory balls were acquired using the custom build and highly fexible Flex-ray CT scanner, developed by XRE NV and located at CWI. In order to capture the fine details in the ivory carvings, balls were acquired at a resolution of approximately 60 micron, resulting in reconstructed three dimensional data volumes in the order of $4K^3$ voxels.

2.1 Segmentation

Segmentation of the layers is not straight forward. Due to the decorations on each layer, the contact regions between layers are irregular and ill defined. As a result, traditional edge based, region based, or feature based segmentation methods do not suffice.

A sphere-fitting based segmentation method has been developed to overcome these difficulties. Informally, the method can be described in the following steps (see diagram 2) :

1. create binary volume
2. compute the center of mass of the volume
3. trace rays from center of mass to boundary of volume
4. for each ray, determine intersection point of ray with the inner boundary of the largest sphere
5. iterate until least-square of fitting residuals is small:
   - fit sphere with set of intersection points
   - remove points with large residuals
6. label white voxels with distance to center greater than radius
7. go to step 2.

There are some additional details to be dealt with, but the general idea is that the steps 2-5 segment the outer layer from the rest of the data. In practice, a few thousand rays are required in order to obtain adequate fitting values.

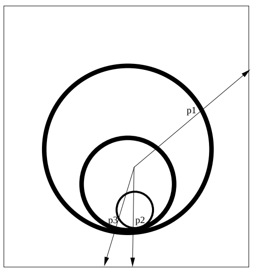

Figure 2: Tracing three rays through three enclosing spheres. Rays 1 and 2 both have 4 intersection points. Ray 3 has 6 intersection points. The intersection points of the largest sphere are labeled p_1, p_2, p_3. The sphere fit procedure returns the center, radius of the sphere and residuals of p_i.

Finally, isosurface meshes of each layer are constructed from the labeled volume (see figure 3). The meshes are then aligned to a common center and orientation.

2.2 Measurements

Various morphological properties are computed from the aligned surface meshes. Computed properties include

- thickness of each layer

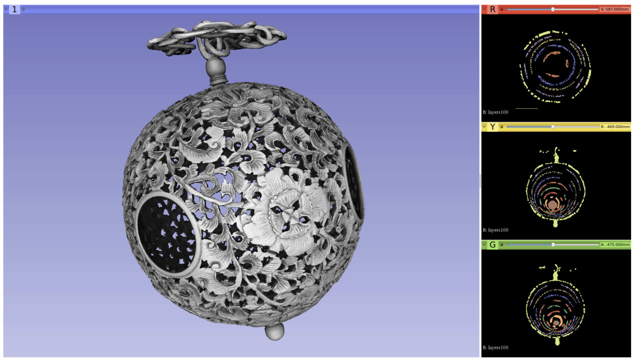

Figure 3: Photo-realistic rendering of layer 1 of AK-NM-7020 (left panel), XY, YZ, XZ cross-sections of labeled volume (right panels).

- annulus (the distance between layers)
- visibility of each triangle in a surface mesh
- size and shape the geometric patterns on each layer

The layer thickness and annulus between spheres for AK-NM-7020 are tabulated. Note that the sum of the layer thickness is smaller then sum of the space between the layers.

Layer	Thickness	Layers	Annulus
1	3.5 mm	1-2	2.8 mm
2	1.6 mm	2-2	2.8 mm
3	1.6 mm	3-4	2.7 mm
4	1.6 mm	4-5	2.8 mm
5	1.7 mm	5-6	2.8 mm
6	1.7 mm	6-7	2.7 mm
7	1.8 mm	7-8	2.7 mm
8	1.8 mm	8-9	2.7 mm
9	5.0 mm		
total	2.0 cm		2.2 cm

Table 1: Layer thickness (left), annulus between layers (right)

Tool marks left by L-shaped scalpel reside on the inner surface of each layer. These can be seen in the data as concentric rings around the peepholes on the inner surface of the layer. Figure 4 shows these patterns on the inner surface of the first layer. The radius of each circle are measured. For ball AK-NM-7020 these radii are 27mm , 25mm and 23mm.

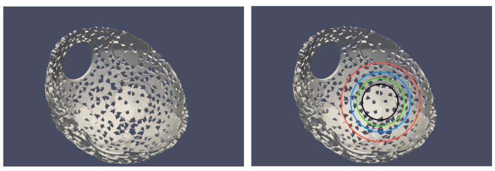

Figure 4: Tool marks on the inner surface of the first layer. Three concentric circles around each peephole are faintly visible (top), annotated view (bottom).

3 Interpretation

The answers to many questions related to the make process of ivory balls can be deduced by analyzing the measurements. We pose two example questions:

1. What was the size and shape of L-shaped scalpels used by the craftsman when separating layers ?

   The height of the L-shaped scalpels is bounded by the annulus. For AK-NM-7020, this is approximately 2.8 millimeter. The length of the scalpel is bounded by the patterns found on the inner surface of each layer. On the first layer there are three such patterns. Hence, the lengths of the three L-scalpels are 27mm , 25mm and 23mm.

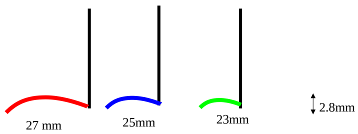

Figure 5: Sizes of three L-shaped scalpels used to separate the first layer.

2. What can the craftsman see when carving geometric patterns on each layer? Is there sufficient light?

   This question has been answered by performing a simplistic light simulation on the aligned meshes. A computer graphics lighting model was used to simulate various lighting conditions. Simulation parameters include the position of a light source, the brightness of a light source, and the amount of ambient lighting. Figure 6 illustrates the output of one simulation. It can be seen that the first few layers receive sufficient direct light, while the deeper layers receive much less light. For AK-NM-7020, the first 5 layers are clearly visible, while layers 6,7,8 and 9 are very much darker.

Figure 6: Light simulation on aligned meshes.

4 Conclusion

We have addressed a number questions related to the make process of ivory balls. In particular, we have shown that three L-shaped scalpels were used to separate the first layer form the rest of the sphere. Also, the shape and size of the scalpels were computed.

The general approach taken in this paper is to first scan an art artifact in a laboratory x-ray setup to acquire a 3D tomogram, segement the tomogram, and measure various morphological proprieties of the artifact. These measurements are used to deduce aspects of the make process. We believe that this approach can be applied to study the make process of other art artifacts.

Acknowledgements

The authors acknowledge financial support from the Netherlands Organisation for Scientific Research (NWO), project number 341-60-001. We acknowledge XRE NV for their role in the FleX-ray collaboration.

Pixel context encoders for painting region inpainting

Nanne van Noord[1] and Eric Postma[2]

[1]University of Amsterdam [2] Tilburg University

Abstract— **In this work we use Pixel Content Encoders (PCE), a light-weight image inpainting model, to inpaint missing regions of paintings. Based on a Convolutional Neural Network (CNN) the PCE leverages dilated convolutions such that it is able to preserve fine grained spatial information and input large missing regions of paintings. Besides image inpainting, we show that without changing the architecture, the PCE can be used for image extrapolation, expanding the painting beyond its existing boundaries.**

1 Introduction

Reconstructing missing or damaged regions of paintings has long required a skilled conservator or artist. Retouching or inpainting is typically only done for small regions, for instance to hide small defects. Inpainting a larger region requires connoisseurship and imagination: the context provides clues as to how the missing region might have looked, but generally there is no definitive evidence. Therefore, sometimes the choice is made to inpaint in a conservative manner. However, with the emergence of powerful computer vision methods specialising in inpainting [1, 2], it has become possible to explore what a potential inpainting result might look like, without physically changing the painting.

Although image inpainting algorithms are not a novel development, previous works has typically only explored inpainting of small regions (i.e., cracks) in paintings [1]. Whereas recent work applied to natural images has shown that approaches based on Convolutional Neural Networks (CNN) are capable of inpainting large missing image regions in a manner which is consistent with the context [2]. In this work we explore (digital) inpainting of large regions in paintings using Pixel Context Encoder (PCE). PCE are a light-weight alternative to previously proposed CNN-based inpainting models which are based on complex network architectures consisting of many trainable parameters, resulting in a necessity of large amounts of data, and often long training times. Our results show that PCE outperform previous work on an inpainting task, and that PCE can be used - without modification - for painting extrapolation.

2 Pixel Context Encoders

The architecture of Pixel Context Encoders (PCE) follows that of encoder-decoder Convolutional Neural Networks used for image generation[3]. In such an architecture the encoder compresses the input, and the decoder uses the compressed representation (i.e., the bottleneck) to generate the output. By incorporating dilated convolutions we reduce the spatial compression in the encoding stage which allows the model to preserve fine grained spatial information. Moreover, as compared to *regular* convolutions, dilated convolutions can cover the same receptive field with significantly fewer parameters.

PCEs are trained through self-supervision; an image is artificially corrupted, and the model is trained to regress back the uncorrupted ground-truth content. The PCE F takes an image x and a binary mask M (the binary mask M is one for masked pixels, and zero for the pixels which are provided) and aims to generate plausible content for the masked content $F(x, M)$. During training we rely on two loss functions to optimise the network: a L1 loss and a GAN loss. For the GAN loss we specifically use the PatchGAN discriminator introduced by Isola et al. [3].

3 Experimental results

To demonstrate the potential of PCE as a tool for painting reconstruction we demonstrate it on two tasks, inpainting and painting extrapolation.

3.1 Datasets

The main dataset used in this work is the "*Painters by Numbers*" dataset (PaintersN) as published on Kaggle[1], and consists of $103,250$ photographic reproductions of artworks by well over a thousand different artists.

Additionally, for the quantitative evaluation we report the performance of the inpainting models on the subset of $100,000$ images that Pathak et al. [2] selected from the ImageNet dataset [4]. The performance is reported on the complete ImageNet validation set consisting of $50,000$ images.

3.2 Inpainting

To compare PCE to previous work we perform centre region inpainting. In centre region inpainting the central 64×64 region is removed from a 128×128 image and subsequently inpainted by the inpainting models. In Table 1 the results of this comparison between Context Encoders [2] (CE) and PCE are shown. Both models are trained and evaluated on the ImageNet dataset and the PaintersN dataset, to explore the generalisability of the models. The performance of the PCE model exceeds that of the model by Pathak et al. for both datasets. Nonetheless, both models perform better on the PaintersN dataset, implying that this might be an easier dataset to inpaint on. Overall, the PCE model trained on the $100,000$ image subset of ImageNet performs best, achieving the lowest RMSE and highest PSNR on both datasets.

In addition to quantitative results we show in Figure 1 the results of a qualitative comparison between CE [2] and the PCE model.

[1]https://www.kaggle.com/c/painter-by-numbers

Table 1: Center region inpainting results on 128×128 images with a 64×64 masked region. RMSE and PSNR for models trained on the ImageNet and PaintersN datasets (horizontally), and evaluated on both datasets (vertically).

		ImageNet		PaintersN	
Trained on	Model	RMSE	PSNR	RMSE	PSNR
Imagenet	CE [2]	43.12	15.44	40.69	15.94
	PCE	**22.88**	**20.94**	**22.53**	**21.08**
PaintersN	CE [2]	43.69	15.32	40.58	15.96
	PCE	24.35	20.40	23.33	20.77

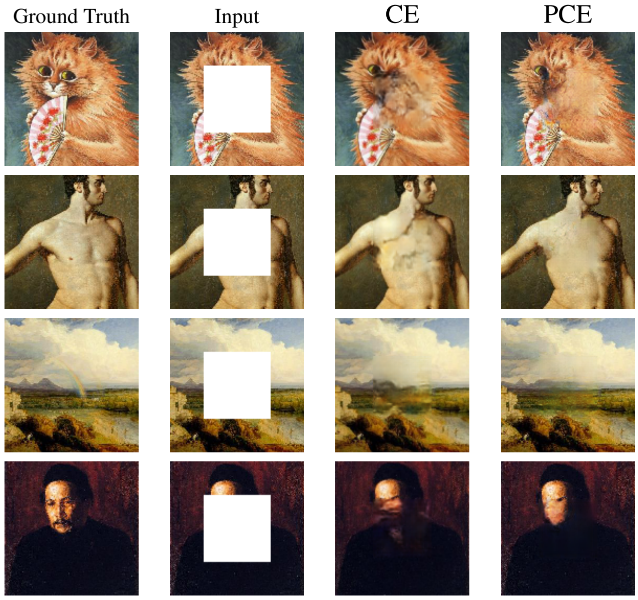

Figure 1: Comparison between CE [2] and PCE, on inpainting a 64×64 region in 128×128 images taken from the PaintersN dataset.

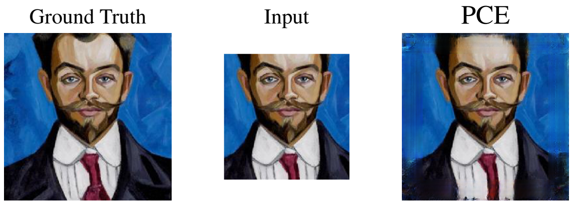

Figure 2: Examples produced by the PCE model, on extrapolating 192×192 regions taken from 256×256 images.

3.3 Painting Extrapolation

Besides inpainting we also explore painting extrapolation; generating novel content beyond the existing painting boundaries. By training a PCE to reconstruct the content on the boundary of an image (effectively inverting the centre region mask), we are able to teach the model to extrapolate paintings.

In Figure 2 we show two examples obtained through painting extrapolation. Based on only the provided input the PCE is able to generate novel content for the 64 pixel band surrounding the input. Although the output does not exactly match the input, the generated output does appear plausible.

Additionally, in Figure 3 we show images obtained by applying the PCE trained for painting extrapolation to uncorrupted images, resized to 192×192 pixels. By resizing the images to the resolution of the region the model was train on, the model will generate a band of 64 pixels of novel content, for which there is no ground truth.

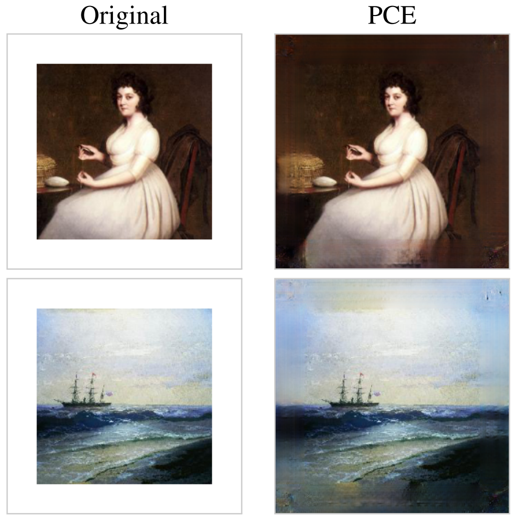

Figure 3: Examples produced by the PCE model, on extrapolating 192×192 images beyond their current boundaries.

4 Conclusion

For this work we trained and evaluated the inpainting performance of PCE on a dataset of paintings and a dataset of natural images (ImageNet). The results show that regardless of the dataset PCE were trained on they outperform previous work on either dataset, even when considering cross-dataset performance (i.e., training on natural images and evaluating on paintings, and vice versa). Based on the cross-dataset performance we pose that PCE solve the inpainting problem in a largely data-agnostic manner. By encoding the context surrounding the missing region PCE are able to generate plausible content for the missing region in a manner that is coherent with the context.

We conclude that PCE offer a promising avenue for image inpainting and the digital restoration of paintings. With an order of magnitude fewer model parameters than comparable inpainting models, PCE obtain state-of-the-art performance on benchmark datasets of paintings and natural images. Moreover, due to the flexibility of the PCE architecture it can be used for other image generation tasks, such as inpainting larger regions and image extrapolation.

References

[1] Cornelis, B., Ružić, T., Gezels, E., Dooms, A., Pižurica, A., Platiša, L., Cornelis, J., Martens, M., De Mey, M., Daubechies, I.: Crack detection and inpainting for virtual restoration of paintings: The case of the Ghent Altarpiece. Signal Processing **93**(3) (mar 2013) 605–619

[2] Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR. (2016)

[3] Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. arXiv (2016) 16

[4] Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet large scale visual recognition challenge. International Journal of Computer Vision **115**(3) (2015) 211–252

Title:

Iron Age Hebrew Epigraphy in the Silicon Age - An Algorithmic Approach To Study Paleo-Hebrew Inscriptions

Author:

Barak Sober

Abstract:

Handwriting comparison and identification, e.g. in the setting of forensics, has been widely addressed over the years. However, even in the case of modern documents, the proposed computerized solutions are quite unsatisfactory. For historical documents, such problems are worsened, due to the inscriptions' preservation conditions. In the following lecture, we will present an attempt at addressing such a problem in the setting of First Temple Period inscriptions, stemming from the isolated military outpost of Arad (ca. 600 BCE). The solution we propose comprises: (A) Acquiring better imagery of the inscriptions using multispectral techniques; (B) Restoring characters via approximation of their composing strokes, represented as a spline-based structure, and estimated by an optimization procedure; (C) Feature extraction and distance calculation - creation of feature vectors describing various aspects of a specific character based upon its deviation from all other characters; (D) Conducting an experiment to test a null hypothesis that two given inscriptions were written by the same author. Applying this approach to the Arad corpus of inscriptions resulted in: (i) The discovery of a brand new inscription on the back side of a well known inscription (half a century after being unearthed); (ii) Empirical evidence regarding the literacy rates in the Kingdom of Judah on the eve of its destruction by Nebuchadnezzar the Babylonian king.