

Master Erasmus Mundus in
Color in Informatics and Media Technology (CIMET)



Computer-Aided Reclamation of Lost Art
Master Thesis Report

Presented by

Maria-Lena Demetriou

and defended at

Gjøvik University College

Academic Supervisor(s): Prof. Jon Yngve Hardeberg
Gabriel Adelman, MD.

Jury Committee: Prof. Alain Trémeau
Dr. Ville Heikkinen

Computer-Aided Reclamation of Lost Art

Maria-Lena Demetriou

2012/07/15

Abstract

Based on previous research on super-resolution and colour correction by example, the present project employs such techniques towards a high-quality reclamation of lost art. The aim is to produce high-quality images of nowadays destroyed or missing paintings using the correspondence between similar artworks available in both low and high quality images. Several approaches for both super-resolution and colour correction techniques have been studied, implemented and tested to result to the most efficient and appropriate to be used. This project is an attempt that has never been done before and the successful reclamation of lost art initialises a new area of colour imaging applications in fine art. It reveals unlimited possibilities in the domain and establishes the potential of further attempts in this direction.

Summary

The present document covers the entirety of a proposed approach and an implemented solution towards *Reclamation of Lost Art*. Employing a combination of example-based super-resolution and colour correction techniques, the project attempts a unique step in colour imaging applications in fine art. It explores the applicability and performance of state-of-the-art techniques in super-resolution, inverse halftoning, image registration, colour transfer and image blending methods to the demanding image set of fine art specimens.

Firstly, in Chapter 1, the problem at hand is introduced and stated along with the motivation behind this aim. It expands on the significance of contribution of this endeavour, the implicated considerations and relevant previous research.

Next, a background overview in Chapter 2 is completed, providing all the necessary information and background knowledge to the reader, for an effortless comprehension of the rest of the document. The notions of halftoning, inverse halftoning, image registration, super-resolution, colour correction and image quality metrics are defined, described and explained with regards to the stated problem.

Proceeding to Chapter 3, related literature work is presented. Relevant theory is outlined, including an in-depth explanation of the selected methods to be utilised in the project. This literature survey is divided into the three major implementation steps of the system: inverse halftoning, super-resolution and colour correction. State-of-the-art techniques are introduced for each of the domains, proceeding to more detailed explanation of the implemented approaches.

With related work covered, the document proceeds via Chapter 4 into a description of the implemented system. This chapter discusses all considerations that arose during development, the selected solutions, the criteria of this selection, the necessary adjustments towards the goal and finally the carried out evaluation.

Chapter 5 then demonstrates the outcomes of the system, restricting to only a few examples due to space limitations. It also includes a qualitative and quantitative analysis of the results. Appendix C includes a detailed quantitative evaluation, while supplementary material is available for visual observation and examination of the complete set of the reclaimed images.

Finally, Chapter 6 concludes the document presenting the significance of the study's outcomes and the contribution to research. It further includes a discussion of possibilities towards extension of the present work through further research, suggesting methods of both improving and expanding it.

In the appendices, two scientific papers are available, along with program documentation and supplementary material from the evaluation procedures.

Preface

Most importantly, I thank my family who have supported and encouraged me into seizing all opportunities and making the best out of every little thing. For all the calls I ended in a rush because I had to be somewhere or do something, you can be sure I substitute them with most of my everyday thoughts.

I would also like to thank my academic supervisor, Prof. Jon Yngve Hardeberg for introducing me to this research area and providing me with the opportunity to focus on the intensely interesting domain of fine art colour applications. Moreover, for your highly valuable support and guidance that made the path clearer. In addition, I thank my second supervisor Gabriel Adelman, who initially introduced the idea and its possibilities. I want to express my gratitude for both the financial support for the project and the provision of required data. Most importantly though, I thank you for your transmitting enthusiasm and fascination which were the ultimate motivation to sustain the same initial excitement throughout the study.

To my two incredibly understanding guys, Michalis and Yiannis, thank you for all your patience to my mumbling about things I am sure you did not really care about, comprehended or had anything to comment on. Even more, thank you for your 'brutal' responses that could in a split second bring me back to reality and rationalise me once again.

To every member of my CIMET family, Rémy, Mike, Lynn, Alex, Pamela, Tatiana, Vignesh, Kiks, Victor and Victor, Natalia, Annick, Piotr, Kica, Oscar, Melkamu, Rahul, Alex, Alina, Owais and Janet, thank you for being there and for being how you are. I would not change a single thing from the last two years with you. This also reaches you, Rahat and Rizwan, for your big-brother support for whatever I asked for.

Lastly but definitely not least, for the two people that have been with me for more time in the last two years than they should be able to handle, Laksmi and Vamsi, I owe to you the person I am today. Thank you for handling the good, the bad and the ugly me with so much success. You have changed me radically, bringing out - I would like to believe - a better me. You should know that no matter the distance, you will always have me.

Marielena Demetriou, 2012/07/15

Contents

Abstract	iii
Summary	v
Preface	vii
Contents	ix
List of Figures	xi
List of Tables	xiii
1 Introduction	1
1.1 Aim	1
1.2 Objectives	1
1.3 Research	2
2 Background	3
2.1 Halftoning	3
2.2 Inverse Halftoning	5
2.3 Image Registration	5
2.4 Super-Resolution	7
2.5 Colour Correction	8
2.6 Image Quality Metrics	9
3 Literature Survey	11
3.1 Inverse Halftoning	12
3.1.1 Using a Look-Up Table	13
3.1.2 By Decision Tree Learning	16
3.1.3 Via Super-Resolution	18
3.2 Super-Resolution	19
3.2.1 Subjective Evaluations of EB, TV, and Combined Regularization	20
3.2.2 Single Image Scale-Up Using Sparse-Representations	24
3.3 Colour Correction	27
3.3.1 Utilisation of Colour Space	28
3.3.2 Naturalness Constraints	29
3.3.3 Image stitching	30
3.3.4 Underwater Images	31
3.3.5 Using Content-Based Image Retrieval	32
4 Implementation	35
4.1 Training and Test Image Sets	35
4.2 Pre-processing Steps	36
4.3 Example-Based Super-Resolution	37
4.4 Colour Correction	39
4.5 Evaluation	40
5 Results and Discussion	43
5.1 Misalignment Simulation	43
5.2 Super-Resolution	43

5.3 Colour Correction	44
6 Conclusions and Further Work	49
6.1 Contributions	49
6.2 Conclusions	49
6.3 Future Work	50
Bibliography	53
A Scientific Papers	59
B Program Documentation	77
C Evaluation Scores	79

List of Figures

1	Error Diffusion	4
2	Halftoning and Inverse Halftoning	5
3	Image Registration	6
4	Resolution	7
5	Application of super-resolution	7
6	Colour Correction based on van Gogh's painting	8
7	Example of colour correction application	9
8	Halftoning pattern in the scanned images	12
9	Reconstruction Results from TV and EB-TV approach	23
10	Colour correction in different colour spaces	28
11	Colour adjustment to a specified CCT and luminance	29
12	Panorama registration	30
13	Panorama rendered with multi-band blending	30
14	Colour correction of underwater images	32
15	CBIR CC Approach Results	33
16	Outline of the implemented system	35
17	Super-Resolution's Boundary Effect	39
18	Results from SR misalignment simulation	44
19	Effect of misalignment in training for test image 'Madonna'	45
20	Difference Map for test image 'The Judgement of Paris'	47
21	Example result of test image 'Satyrs and Bacchants'	48

List of Tables

1	REC vs. C/R	1
2	Performance of Floyd-Steinberg error diffusion restoration methods	13
3	Neighbourhood Used in IH (Rectangular Support, ‘Rect’ template)	14
4	Neighbourhood Used in IH (‘16pels’ and ‘19pels’ templates)	14
5	Performance of IH by Minami <i>et al.</i> approach	18
6	Selected colour correction approaches	40
7	PSNR Measurements for simulated misalignment in training of SR	43
8	Super-Resolution evaluation for ground-truth test set	44
9	Evaluation of final results	46
10	S-CIELAB evaluation of CC results	47
11	PSNR Scores for ground-truth test image ‘The rape of the daughters of Leukippos’	80
12	PSNR Scores for ground-truth test image ‘Liebesgarten’	81
13	PSNR Scores for ground-truth test image ‘Madonna’	82
14	PSNR Scores for ground-truth test image ‘The Judgement of Paris’	83
15	PSNR Scores for ground-truth test image ‘Putti’	84
16	SSIM Scores for ground-truth test image ‘The rape of the daughters of Leukippos’	85
17	SSIM Scores for ground-truth test image ‘Liebesgarten’	86
18	SSIM Scores for ground-truth test image ‘Madonna’	87
19	SSIM Scores for ground-truth test image ‘The Judgement of Paris’	88
20	SSIM Scores for ground-truth test image ‘Putti’	89
21	S-CIELAB Scores for ground-truth test image ‘The rape of the daughters of Leukippos’	90
22	S-CIELAB Scores for ground-truth test image ‘Liebesgarten’	91
23	S-CIELAB Scores for ground-truth test image ‘Madonna’	92
24	S-CIELAB Scores for ground-truth test image ‘The Judgement of Paris’	93
25	S-CIELAB Scores for ground-truth test image ‘Putti’	94

1 Introduction

Preserving the artistic heritage for present and future audiences has been traditionally dealt with by the sciences of *conservation* (C) and *restoration* (R), aimed, respectively, at arresting and reversing the action of such factors as aging of the materials, dirt accumulation, or mechanical, biological, or chemical injuries to the work of art. These time and skill intensive technologies have been revolutionized by computer-based methodologies, but remain essentially operator-dependent. Art object imaging plays an ancillary, albeit very important role, for purposes such as documentation, teaching, and mass reproduction.

A self-evident prerequisite of any C/R endeavour is the survival of that work of art, either in its entirety, or in a sufficient ‘critical mass’ to allow appreciation of the creator’s artistic intention and skill.

1.1 Aim

The present project aimed at exploring an apparent ‘lost cause’ - that of heritage preservation as applied to paintings documented to have been destroyed at times of major turmoil, such as WW II. The question at hand is not whether these objects can ever be recuperated (clearly, they cannot), but rather, whether it is possible to recuperate, by computer manipulation of existing photographs, a reasonably accurate image of the lost original.

For this novel question, the term art *reclamation* (REC) was proposed. In this context, art object imaging plays a central role, since it is, literally, an image that is reclaimed, the original being irretrievably lost.

The underlying reasoning was that if the answer turns out to be affirmative, a high-resolution electronic sample and a corresponding hard-copy reproduction can be obtained and publicly displayed, to the interest and delight of a potentially very large audience, encompassing all age and interest groups, whether technically, artistically, or historically minded. This would represent the sole possibility of contact with lost creations belonging to the core heritage of humanity.

Table 1 circumscribes two main points of difference between REC and C/R.

	C/R	REC
Object still existent	Yes	No
Methodology	Acting on the object itself (imaging ancillary)	Acting on existing images of the object and/or of similar objects (imaging essential)

Table 1: REC vs. C/R

1.2 Objectives

The basic concept is that of using pre-WW II low-quality colour and B/W photographs of lost paintings on one hand, and contemporary high-quality photographs of surviving

stylistically similar paintings by the same master (in our attempt, P.P. Rubens) on the other, to endeavour an enhancement of the former to an acceptable quality standard.

This unique idea evolved from the interaction of people with different backgrounds, sharing an equal interest in the two core manifestations of the human spirit - art and technology. Traditionally viewed as opposites, they appeared to this team as inextricably related.

1.3 Research

As usual with a pioneering attempt, it was important to underscore not only what our project attempted to be, but also how it departed from similar endeavours - or, in other words, what it was not. Emphatically, then, this work was not set as a colourisation technology, since the target was not to *recreate*, but to *redeem (reclaim)* colour-related information, in a manner as objective, operator-independent, and art-historically accurate as possible.

Investigating a pioneering concept, the research team was in uncharted territory, and had no previous body of research to draw, expand, or improve upon. Most important expectation was therefore an objective assessment as to whether the aim of this work was achievable or not, at an adequate level of quality.

2 Background

Before proceeding into the technical sectors of the work, it is necessary to firstly introduce all terms used in the rest of the document, so that the reader is able to follow and comprehend the material without any difficulties. The present chapter introduces all notions that are essential for accurately and fully apprehend the involved topics and thus the scope of this project.

2.1 Halftoning

The paintings aimed to be reclaimed in the scope of the study, are nowadays available only in low quality prints in art books [1, 2]. As prints, the involved low quality images are not of continuous-tone but instead, they are halftoned.

Halftoning is the physical process of converting a continuous-tone image, such as a photograph, to black and white dots for its reproduction in a printing device [3]. Halftoning is an old technique which started out with the use of *halftone screens*, which were mechanical screens of regular grids in front of high-contrast photographic film [4]. When exposed to light, stronger light would make a larger area developable and thus produce a larger dot on the developed film. The film would contain many dots of various sizes, where each would be proportional to the local radiance of the original image. This would then be used to etch the printing plate. Later on, the *contact screens* came up, which were halftone screens made of photographic film. Halftone colour prints were achieved by the use of four *colour separations*, usually CMYK which stands for cyan, magenta, yellow and black inks respectively. The colour separations were individual halftone images made by placing the opposite filter or a combination of filters in the light path of the original image and the contact screen. A different screen angle would be used for each colour separation (Y at 0°, C at 15°, K at 45° and M at 75°), based on the sensitivity of the human visual system to horizontal and vertical gratings. The four colour separations halftone images would then be overprinted on the output medium to form the final colour image.

Nowadays, halftoning is applied using what is called *digital halftoning*. In this technique, the original image is digital, often with 8-bit or more grey levels. This is thresholded into a binary, halftone, digital image that is then sent to a printer. Thus, digital halftoning can be defined as the conversion from a many-bit digital image to a 1-bit binary digital image [4]. Digital printers can perform halftoning by printing dots of equal size and uniformly spaced so that different shades of grey can be produced using $N \times N$ dots as a *halftone cell*. This technique is called *ordered dithering*. Another technique commonly applied by digital printers is *error diffusion*, in which the size of the dots and the ink thickness are variable.

Some of the most popular modern halftoning techniques [4] are

- **Clustered-dot and dispersed-dot halftoning** Both techniques involve a coverage of the image area with halftone cells of equal area, with a black dot printed in the centre of each cell. In the case of clustered-dot pattern, the printer cells are turned on as a cluster of the pixel neighbourhood, while in dispersed-dot pattern the black dot is

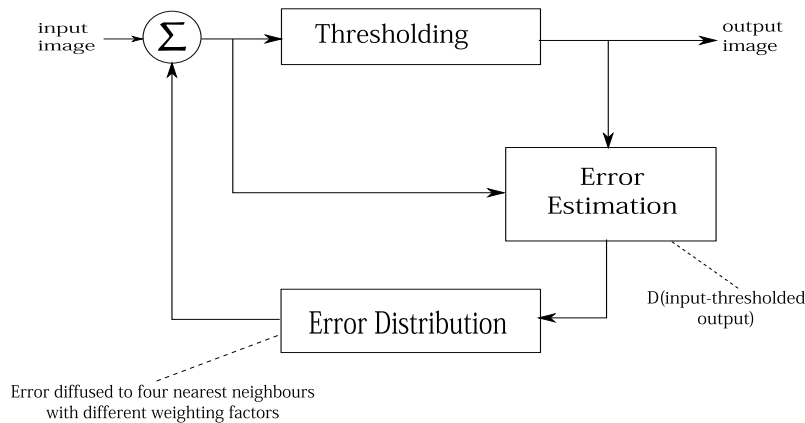


Figure 1: Error Diffusion Procedure Diagram

broken into many smaller dots which are then distributed evenly. The results of the dispersed-dot pattern are much smoother and pleasing, while in the clustered-dot halftoning, if the halftone cell is big enough to be visually resolvable, the patterns are visible as halftoning noise. In the digital halftoning application of these techniques, the halftone cell is divided into smaller printer cells and this is how the size of the dots is controlled.

- **Random Dither** In this method, random noise is added to the input image signal and then a fixed threshold is applied to convert it to binary output. This produces sharp contours and does not provide the illusion of a continuous-tone image. This technique is not considered of high quality, as it contains noise in all spatial frequencies and is not pleasant to the human visual system for low frequency mottles. Nevertheless, it is useful in applications for which contours caused by quantisation need to be made less visible in grey level images.
- **Ordered Dither** Similar to random dither halftoning, ordered dither halftoning is based on thresholding, using a deterministic, two-dimensional threshold array to generate a binary image from a grey-level image. The input image is divided to square non-overlapping consecutive blocks. Point operations follow next, where each pixel is compared to the threshold array and then assigned to a binary value of 0 or 1, independent of the neighbouring pixels' values. Naturally, the output binary image contains a fundamental frequency of a period equal to the size of the threshold array. Ordered dither halftoning is visible and thus it is not considered pleasant.
- **Error Diffusion** This method differs in the fact of taking into consideration the local information and relation between neighbouring pixels in an image. Instead of point operations, neighbourhood operations are performed, leading to much better looking halftone images. The notion behind error diffusion is that when quantising a grey-level pixel into a binary black or white dot, an error in brightness at that point in the image is created. However, the neighbouring pixels can compensate by being turned in the direction that will cancel the present error once all the reflectances in the neighbourhood are integrated by the visual system. Figure 1 demonstrates the method for error diffusion halftoning.

The numerous halftoning techniques have been developed taking into consideration the functionality of the human visual system, devices/media models, models of computational architecture and system calibration as to be able to define how the dot patterns would be physically rendered. If halftoning is not done properly, the halftone's texture pattern can interfere with the observer's perception of an image and it would therefore not be truly invisible. This is why it is necessary to know the spatial resolution for which a halftoning technique is good. The several implicated models are designed in parallel, but still, the final result needs to be calibrated and optimised. If parameters cannot be measured or controlled, the final image quality cannot be ensured.

2.2 Inverse Halftoning

The above study on halftoning techniques aimed mainly to the understanding of the process which led to the resulting input low quality images, as well as their nature. Realisation of the implications in the production of these images allowed a better overall understanding of the problem itself. Moreover, it was a necessary prerequisite for the study of inverse halftoning (IH).

Inverse halftoning (or descreening) was considered an essential preprocessing step onto the input low quality halftones, as it is commonly accepted that image processing cannot be applied on halftone images directly but instead, these must be converted to continuous-tone first. This is by definition the task of inverse halftoning, i.e. the reconstruction of the 8-bit image from its halftoned version, so that image processing operations can be performed [5]. Figure 2 demonstrates the application of halftoning and that of inverse halftoning.

As more than one continuous-tone images can give rise to a particular halftoned image, IH is a difficult task, approached by many different methods. A detailed description of the study of inverse halftoning algorithms and the selection of the most appropriate can be found in section 3.1.



Figure 2: An illustration of halftoning and inverse halftoning. Detail of the Lena image: original (left), Jarvis error diffusion halftone (centre), and inverse halftone result (right). From Foi *et al.* [6]

2.3 Image Registration

Registration is a fundamental task in image processing, used to match and align two or more images. Such images could be for example pictures taken at different times, from different sensors or from different viewpoints [8]. In other words, image registration is the alignment of two or more images which might be different in terms of rotation, scaling, translation and further more. Figure 3 displays an example approach of image registration (IR) from Zitova and Flusser [7].




Figure 3: Four steps of image registration. Top row: feature detection (corners were used as the features in this case). Middle row: feature matching by invariant descriptors (the corresponding pairs are marked by numbers). Bottom left: transform model estimation exploiting the established correspondence. Bottom right: image resampling and transformation using appropriate interpolation technique. From Zitova and Flusser [7]



Figure 4: The same image at 72x100 pixels (right) and 288x400 pixels. From Elad [9]

The step of image registration was necessary in the training phase of super-resolution, as the corresponding low and high quality pairs of images were not perfectly aligned. Factors such as the resolution difference and misalignments in the scanning of the low quality images, made image registration a necessary step before the training phase.

2.4 Super-Resolution

Resolution of an image is described in the total number of pixels of an image, or regarding the size of the smallest detail that can be seen in the image. As one can observe in Figure 4, these two definitions are closely linked. Super-resolution (SR) is therefore the enhancement of an image to higher resolution. Figure 5 displays two examples of application of three state-of-the-art SR approaches.



Figure 5: Flower and girl images magnified by a factor of 3. Left to right: input, bicubic interpolation, neighbour embedding [10], Yang *et al.* method [11], and the original image. From Yang *et al.* [11]

While the term of super-resolution is generally defined as the enhancement of resolution, it is sometimes referred to as the process of fusing a number of low resolution images of the same scene into a high resolution image [12]. The use of ‘super-resolution’ in this document refers to the resolution enhancement definition.

2.5 Colour Correction

The process of colour correction (CC) involves the adjustment and refinement of the colours in an image. In the case that this approach is based on examples, the colour correction procedure is defined according to given ‘correct’ images, which demonstrate the desirable characteristics. Naturally, in order for the source image to be adjusted with appealing results, the target image has to show some similarity to the source image. Otherwise, the results will possibly be disturbing.

Applications of colour correction involve artistic experimentation with images that give the mood of an image to another and enhancement of an image that has distortions of some kind, using examples of high quality and free of distortion. An example of the first approach is the rendering of a winter scene so as to look like a spring outdoor scene [13], or adjustments as the one in Figure 6, in which the image of a real scene is rendered using the colour characteristics of a painting. The second application, targeting image enhancement, includes removal of artifacts such as cast or effects from taking pictures under different illumination conditions or viewpoints. Moreover, image enhancement through colour correction is performed for restoration, as in the case of restoring old photographs of which the colours have faded.

Example-based colour correction is closely related to the term of *colour transfer*, the extraction of colour properties from a reference image and their transmission to the input image. This is essentially what is performed in colour correction by example, as the characteristics of the target image are desired onto the source image.

Figure 7 demonstrates an example of applying colour correction by example. Figure 7(a) presents the original source image while Figure 7(b) displays the target image. Figure 7(c) shows the resulting colour corrected image using Reinhard *et al.* approach [14].



Figure 6: Applying the mood of Vincent van Gogh’s painting titled Café Terrace on the Place du Forum, Arles, at Night (Arles, September 1888, oil on canvas; image from the Vincent van Gogh Gallery, <http://www.vangoghgallery.com>) to a photograph of Lednice Castle near Brno in the Czech Republic. The blues of the sky in both images are matched, as the yellows of the cafe and the castle, and the browns of the tables at the café and the people at the castle separately. From: Reinhard *et al.* [14]




Figure 7: Applying Colour Correction. From: Reinhard *et al.* [14]

2.6 Image Quality Metrics

Image quality metrics were introduced to support quantitative analysis and evaluation regarding the fidelity of images. Typically, the quality of a rendered image is evaluated against a reference image. In these cases, the metric is a full-reference metric, quantifying the fidelity using image to image comparisons [15].

Three full-reference image quality metrics were utilised in this study: Peak signal-to-noise ratio (PSNR), structural similarity index metric (SSIM) [16] and S-CIELAB [17].

PSNR is defined as

$$\text{PSNR} = 10 \log_{10} \left(255^2 \cdot N / \sum_i (\hat{y}_i - y_i)^2 \right), \quad (2.1)$$

with $y, \hat{y} \in [0, 255]^N$. Higher PSNR scores correspond to better image quality.

SSIM improves on PSNR, providing better correlation with human judgement of image quality [16]. This metric is defined as a combination of luminance, contrast and structure components:

$$\text{SSIM}(a, b) = [l(a, b)]^\alpha \cdot [c(a, b)]^\beta \cdot [s(a, b)]^\gamma, \quad (2.2)$$

where $l(a, b) = \frac{2\mu_a\mu_b + A_1}{\mu_a^2 + \mu_b^2 + A_1}$, $c(a, b) = \frac{2\sigma_a\sigma_b + A_2}{\sigma_a^2 + \sigma_b^2 + A_2}$, $s(a, b) = \frac{\sigma_{ab} + A_3}{\sigma_a + \sigma_b + A_3}$. Components μ_a and μ_b are the mean luminance values of windows a and b respectively; σ_a and σ_b are the standard variance of the windows a and b respectively; σ_{ab} is the auto-covariance between a and b . Constants $A_1 = (0.01 * L)^2$, $A_2 = (0.03 * L)^2$ and $A_3 = A_2/2$ where $L = 255$ are set to avoid dividing by zero, while $\alpha = \beta = \gamma = 1$ control the weighting between the three components. The higher value of the SSIM, the higher the structural similarity.

Finally, S-CIELAB metric [17, 18] measures the accuracy of reproduction of a colour against its original when this is viewed by a human observer. Lower S-CIELAB values indicate better image quality.

3 Literature Survey

The present project attempts to achieve a pioneering goal, a reclamation of lost artworks in a way that has never been done before. While there is no previous study on the exact same concept and extent, the components consisting the work: *Inverse Halftoning* (IH), *Super-Resolution* (SR) and *Colour Correction* (CC), are each one on their own a section of research that has been extended greatly, especially during the last two decades. Combining all three of these enhancement techniques allowed to accomplish the set target and brought up opportunities in applications of colour imaging in fine art, with extends onto the possibility of altering art knowledge as we know it.

It is key to note here that in every aspect of the project, an *example-based* (EB) approach was employed. This allowed to make the best out of the available high quality images, using them as reference. In the IH process, the approach by Kim *et al.* [19] was considered as the most appropriate to the goals of the necessary pre-processing. The approach illustrates a look-up table based solution for producing a continuous-tone image from the halftone image, building upon the method by Meşe and Vaidyanathan [20]. Despite the fact that the approach is not considered to provide the highest performance [6], it was selected due to its applicability to all kinds of halftoning. The higher in performance algorithms are by definition applicable to only a specific type of halftoning techniques, such as error diffusion, which in the time of production of the project's art material were definitely not available. Thus, a procedure that is more generic seemed as the most promising and was therefore selected and applied.

Regarding the first step of the main process, that of SR, an extensive study of various previous works was carried out. Although there are many approaches built for applying example-based super-resolution (EB-SR), the applicability of each one on the a specific dataset was questionable due to the morphology of the images. Thus, the three most promising state-of-the-art techniques [21, 22, 11] in literature were studied further and discussed with experts in the domain of SR, leading to the selection of two approaches [21, 11] to be implemented, applied and evaluated.

Throughout the study of related literature in the domain of CC, familiarity with the diversity of the applications in this domain was gained. Numerous approaches have been proposed, each dealing with different types of colour correction for either artistic purposes [23], or towards a specific correction [24]. While most of these approaches are considered to be applicable to any kind of images, the success of their performance is arguable, as how pleasing the results appear can be subjectively evaluated.

The images involved in this project have several differences from those employed in previous studies. First of all, the input images are not bi-level halftones but instead, they are scanned halftones, which introduces additional considerations due to the possible presence of misalignment. Secondly, the images are dissimilar to the standard image datasets used in the studies of state-of-the-art techniques in terms of size and content. Most of the previous works include experiments on natural images rather than artistic and of much smaller size. Finally, CC is on its own a subjective domain, as the effect on colour is sometimes user-defined [14] or automated [23]. So far, CC has limited extend to

applications in fine art, making it impossible to define expectations regarding the output of application of CC on the involved dataset.

Validating experiments were carried out in order to assess the efficiency of the results. These included quantitative evaluation using image quality metrics PSNR, SSIM and S-CIELAB against ground-truth data which were available in both low and high quality. In addition, qualitative evaluation was performed on the same ground-truth image set and a real-test image set of actual lost paintings.

The rest of this chapter covers the state-of-the-art techniques with reference to their relation to the purposes of the present study.

3.1 Inverse Halftoning

Despite the fact that the IH procedure was merely a pre-processing step, its importance to the project was vital. The input images of the project are scans of images in art books, as these were printed over than 50 years ago. Naturally, as these images were initially printed for the art books publication, the resulting digital images from the scanning demonstrate the halftoning effect in a very strong fashion. Figure 8 demonstrates an example of how the images were obtained from the scanning procedure.

It was necessary to resolve this effect so that the input images could be utilised as continuous-tone into the next steps of the project. The presence of halftoning in the images was unacceptable and thus IH process was deployed. In general, to ensure that an IH technique is the most appropriate one, it is best to be aware of the followed printing procedure. However, the age of the art books did not allow access to this information, while the available methods at that time were limited and very different from modern techniques.




Figure 8: Detail of scanned print of low quality image of P.P. Rubens' painting titled 'Madonna'. The halftoning pattern is evident

This fact introduced difficulties in applying IH to the input images, as most of the recently developed IH approaches aim to invert modern halftoning methods which either provide more information about the actual halftoning, or are digitally performed. Thus, the efficiency of the best known IH algorithms was unpredictable and this step included detailed study and extended experimentation with various methods. Different works focused on different input e.g. error diffusion [6, 25, 26] and/or required a priori

Inverse halftoning technique	Lena Image	Peppers Image
Anisotropic LPA-ICI [6]	32.4	31.6
WinHD (Neelamani et al.) [25]	32.1	31.2
Wavelet-Vaguelette (Neelamani et al.) [28]	31.9	31.0
Wavelet (Xiong et al.) [29]	31.7	30.7
Gradient (Kite et al.) [26]	31.3	31.4
Kernel (Wong) [30]	32.0	30.3
LUT (Meşe and Vaidyanathan) [20]	31.0	-
LMS-MMSE (Chang et al.) [31]	31.4	31.2
POCS-SVD (Hein and Zakhor) [32]	30.4	-
POCS-Wavelet (Bozkurt and Çetin) [33]	32.2	30.9

Table 2: PSNR (dB) performance of restoration from Floyd-Steinberg error diffusion methods. From: Foi *et al.* [6]

knowledge about the halftoning method, as for example [5] and/or demonstrated results only on greyscale images [27].

The majority of past studies focused on error diffused halftoning, which could not have been applied for the project’s input images simply because of the technology present at the time. Instead of blindly researching through IH approaches, the search for the most appropriate and efficient one used the work of Foi *et al.* [6] as a basis. This study provided a complete PSNR comparison amongst state-of-the-art algorithms. Thus, it was the best starting point in initiating research into IH algorithms and from there examining the pre-requisites and structure of each one before concluding to the most promising ones. Table 2 demonstrates the final PSNR comparison onto Floyd-Steinberg error diffusion images as it was presented in the original work of Foi *et al.* [6].

The comparison does not include the study by Kim *et al.* [19], despite it demonstrating the highest performance, as it was tested on images of bigger dimensions (1050×1050) than the rest of the algorithms, which used the standard size of images (512×512). Assuming that the higher performance from Kim *et al.* method is due to bigger input images, this did not exclude it from our considerations, as the images that were involved are of much larger dimensions.

From Table 2, the approaches that were selected as most promising and were considered for implementation were the LUT-based ones. This is due to the other algorithms’ special arrangements towards a specific halftoning technique and their pre-requisites. For example, even if the algorithm by Foi *et al.* [6] demonstrates the best as shown in Table 2, it assumes that the error diffusion kernel is known, which makes it inapplicable in the present project.

The rest of this section covers in detail the two most appropriate LUT-based approaches, Meşe and Vaidyanathan’s [20] and that of Kim *et al.* [19]. Lastly, Section 3.1.3 covers the findings towards merging the IH procedure with SR, since IH is only a pre-requisite to the application of SR. Merging them together makes the system less complex and much more efficient.

3.1.1 Using a Look-Up Table

As mentioned earlier, the most appropriate method for applying IH in the concept of the present project, was to follow an example-based approach that could be merged with the other sections of the system and also put in use the available high quality images. The

approach by Meşe and Vaidyanathan [20] follows an example-based strategy, based on a look-up-table (LUT) which is built from a histogram obtained from a few sample halftone images and their corresponding original images, a perfect fit to the present problem.

Another factor that makes this approach ideal is the fact that there is no necessity for a priori knowledge, as it does not depend on the halftoning method. No linear filtering is applied and no image model is assumed. The LUT is trained from sample images by identifying a pixel's neighbourhood and then assigning a contone value from the LUT depending on the distribution of pixels in the neighbourhood. The approach was applied onto binary images using a neighbourhood of pixels of size 16, while scanned halftones were yet to be explored.

a	a	a	a
a	a	a	a
a	a	O	a
a	a	a	a

Table 3: Neighbourhood Used in IH (Rectangular Support, 'Rect' template)

b	a	a	a	b
a	a	a	a	b
a	a	O	a	a
	a	a	a	
		a		

Table 4: Neighbourhood Used in IH ('16pels' and '19pels' templates)

The templates (a collection of pixels) for the LUT could be obtained through two different approaches: they were either experimentally found, or designed during training. In the first case, a non-linear function is applied on the surrounding halftone pixels. Three different forms of templates were chosen, a rectangular template 'Rect', one with the use of 16 pixels '16pels' in the pattern of a symmetric part and some additional pixels and another with 19 pixels '19pels' which is an extension upon 16pels. Table 3 demonstrates the 'Rect' template, while table 4 demonstrates both 16pels template (formed by letters a) and template 19pels (formed by both letters a and b). In all tables, letter O denotes the estimated pixel. During the IH process, the pixels in the region of support are changed in a specific order, and the contone value is obtained from the LUT. Thus, the templates are obtained in a two-step approach:

1. Obtain expected contone value for each template pattern.
2. Assign contone value to corresponding LUT position for that pattern.

In the second method of obtaining the templates, these are designed during training using a recursive template selection algorithm. Firstly, the expected contone value for each template pattern is obtained. This contone value is then assigned to the corresponding LUT position for the specific pattern. If the number of occurrences of a pattern

$$(p_0 p_1, \dots, p_{N-1}) \quad (3.1)$$

in the sample halftone images is

$$K(p_0, p_1, \dots, p_{N-1}) \quad (3.2)$$

then the histogram value of the pattern in Equation 3.1 is denoted by Equation 3.2. Pattern 3.1 can occur in any of the sample images, any number of times. If the contone values in the original images' corresponding patterns 3.1 are

$$C(p_0, p_1, \dots, p_{N-1}, i) \text{ for } i = 0, 1, \dots, K(p_0, p_1, \dots, p_{N-1}) - 1. \quad (3.3)$$

If a pattern occurs in a sample halftone image, the LUT value for that pattern will be the mean of the corresponding contone values, i.e. the contone values in the original images which correspond to the patterns. If however, the histogram value (Equation 3.2) is equal to 0, the pattern 3.1 does not exist in the sample halftone images. Thus, for the non-existent patterns, the contone values should be estimated from other existing patterns. Naturally, the pattern that is chosen is the one that most closely mimics the missing one. This pattern can be obtained through low-pass filtering, Hamming distance or through a best linear estimator which is based on a least-squares solution. From experimentation, the best predictor method proved to be subjectively performing better than both low-pass filtering and the hamming distance approach.

The quality of the IH outcome is mainly affected from the handling of the non-existent patterns, on the content of the training set and the template selection procedure. These are interrelated as well, as the number of non-existent patterns depends on the defined template and the halftoning method that it is dealt with. The training set is naturally of great importance, with the LUT method performing better for images with high frequency content, especially if these are included in the training. While the algorithm does not perform as good onto smooth images, it still performs better than known algorithms, such as Kite *et al.* [26].

The template selection changes drastically the nature and performance of the algorithm, as the actual LUT changes. It was found that 'Rect' template always led to better results than '16pels' and even better with '19pels'. Despite the fact that adding more pixels to the template provides a better quality IH, it is not ideal to do so as this makes the LUT bigger and leads to a trade-off between performance and efficiency.

The template selection procedure is the prediction of the contone value of the current pixel from a template. The template is of fixed size M and with a fixed neighbourhood. With a recursive algorithm, the mean-square error is calculated between the contone image set and the estimate image set, searching for the minimum error value $\min \text{Err}(D, D_e)$.

The algorithm is a two-step greedy method, providing a non-globally optimum template:

1. Obtain estimated images using the pixels in the previous template and one more candidate pixel.
2. Choose one pixel from the candidate pixels such that when that pixel is added to the previous template in the LUT, the estimated images are closer to their corresponding contone images.

Using the optimisation approach enables the avoidance of trial-and-error which is applied in normal template selection. It therefore allows for obtaining the optimum template without any requirement for human interaction and extensive experimentation.

Coloured halftones are dealt with in the same way, processing each colour plane separately. Template selection and LUT table design is performed individually for each chan-

nel. Better results are possible by using the high correlation between the halftone values of any colour plane and the contone value of a specific colour plane. Pixels from different colour planes can be included in the prediction of the contone value of a particular colour plane. In this case, it needs to be decided exactly which pixels from which colour planes will be added to the template. The template selection algorithm can be adapted to colour halftones by changing its search space. The pixel locations in R, G, B colour planes can be denoted as (R, p, r) , (G, p, r) , and (B, p, r) for $(p, r) \in \mathbb{N}_L$. A fixed neighbourhood around the current pixel to include pixels from all colour planes is defined as

$$N_{\text{colour},L} = \{(c, p, r) | c \in \{R, G, B\}, p \in \{-L/2, \dots, L/2\} \text{ and } r \in \{-L/2, \dots, L/2\}\} \quad (3.4)$$

This algorithm demonstrated good results on stochastic halftones which involve only blue noise, while for dithered halftones, a median filter was applied after the LUT IH to smooth the image. This improves the image quality, as periodic structures are suppressed. An issue of low screen frequency is present, as small templates cannot cover a full cycle and periodic frequency content appears in the image. If the templates are made big enough though, the storage requirements are infeasible.

The matter that could have arose from the use of this approach, is the uncertainty of its performance on scanned halftones, as it is defined to assume bi-level halftones as input. One strong factor for this would be the alignment of the images, as these are certainly not perfectly aligned as in the case of digital halftones.

3.1.2 By Decision Tree Learning

The approach applied by Kim *et al.* [19] is centred around orthographic greyscale halftone images, as was the case in Mese and Vaidyanathan [20]. Following a decision-tree (DT) learning, it makes it possible to reuse a number of algorithms that had already been developed. The central idea behind this work is the maximisation of entropy gain, which allows the learning algorithm to automatically select the ideal window as the decision-tree is constructed. Structured to deal with greyscale images, it was applied on images of size 1050x1050, providing PSNR values of several dB above most well-known IH algorithms. Not only it is of high efficiency, but it also has a very fast implementation.

A theoretic and algorithmic framework based on DT learning is used. This framework, titled ID3 demonstrates two very important advantages over the tree structure LUT approach by Meşe and Vaidyanathan [31]. Firstly, the template selection is automatic, without the user needed to select explicitly a window. Instead, the attributes most necessary to decide the output contone colour are chosen automatically. Secondly, the initial window for DT-learning is defined as much bigger, with the appropriate peepholes being defined in appropriate order, based on the entropy gain maximisation. Thus, its authors consider this work as an improvement upon Meşe and Vaidyanathan's work [31].

Let $Q^x : \mathbb{Z}^2 \rightarrow \{0, 1\}$ be the set of binary images and $Q^y : \mathbb{Z}^2 \rightarrow [0 \dots 255]$ be the greyscale images' set. Binary image Q^x is mapped to a greyscale image Q^y using function Ψ . Function Ψ is defined through a window $W = \{W_1, \dots, W_w\}$, $W_i \in \mathbb{Z}^2$ of w points, and a characteristic mapping function $\psi : \{0, 1\}^w \rightarrow [0 \dots 255]$ as

$$Q^2(p) = \Psi(Q^x)(p) = \psi(Q^x(W_1 + p), \dots, Q^x(W_w + p)) \text{ where } p \in \mathbb{Z}^2. \quad (3.5)$$

Each element W_i of window is called a peephole or a feature. Training input image

A^x and training output image A^y set the training set, while Q^x denotes the image to be inverse halftoned and Q^y the (supposedly unknown) ideal output image. Then a_p^x is a training instance denoted as the content in A^x of the window W shifted to $p \in \mathbb{Z}^2$ at pixel p

$$a_p^x = (A^x(W_1 + p), A^x(W_2 + p), \dots, A^x(W_w + p)) \in \{0, 1\}^w. \quad (3.6)$$

Each pattern a_p^x is associated with an output-colour $A^y(p) \in [0 \dots 255]$. When all pixels of A^x and A^y are scanned, the data obtained will be the training set

$$a = \{(a_{p_1}^x, A^y(p_1)), \dots, (a_{p_m}^x, A^y(p_m))\}, \quad (3.7)$$

where m is the amount of pixels of images A^x and A^y . The test set is obtained in the same way, using images Q^x and Q^y as

$$q = \{(q_{p_1}^x, Q^y(p_1)), \dots, (q_{p_n}^x, Q^y(p_n))\}, \quad (3.8)$$

where n is the amount of pixels in images Q^x and Q^y . Each $q_{p_i}^x$ is a query-pattern, while the ideal output colour is output $Q^y(p_i) \in [0 \dots 255]$.

Applying IH, operator $\hat{\Psi}$ based on A^x and A^y is constructed so that when applied to Q^x , the resulting image $\hat{Q}^y = \hat{\Psi}(Q^x)$ is similar to the ideal output image Q^y . The difference between the ideal and processed outputs is defined using a loss (error) function l . Learning algorithm A constructs a characteristic function $\hat{\psi}$ based on the training set a so that when $\hat{\psi}$ is applied to query pattern $q_{p_i}^x$ yielding the output-colour $\hat{Q}^y(p_i) = \hat{\psi}(q_{p_i}^x)$, the loss l is low with high probability.

Implemented as DT-learning, this method generalises beyond the training set, has a fast implementation, does not need excessive memory space and the training phase does not take long to complete. As mentioned above, the version of DT-learning that was used is the IDE algorithm [34, 35].

The input-pattern space $\{0, 1\}^w$ is split into two and all sample input patterns with black colour in the splitting attribute W_s will belong to one half-space and those with white colour to another. For each of the two half-spaces, the splitting process continues recursively, reducing the dimensions of the space by one $\{0, 1\}^{w-1}$ at each iteration. At each split, an internal node is created and the splitting attribute s is stored in it. The recursion stops when each space contains either only samples with the same output-colour, or only samples with the same input pattern but with two or more different output colours. If the first case takes place, a terminal node is created and the output colour is stored in it. In the second case, a ‘conflict’ takes place, where a terminal node is again created, but this time the average of output-values is evaluated and stored.

In the case of no occurrences of conflicts, the training set is classified perfectly. This could however lead to overfitting, with a hypothesis that treats the training examples poorly, but performs better over the entire distribution of instances. In order not to complicate the algorithm, overfitting was avoided by computing the average (or median) whenever there were k or fewer samples in a pattern subspace. If $k = 1$ the algorithm is performed as ID3, thus this approach can be denoted as k -ID3.

The most crucial point in the algorithm is the choice of the attribute that will be used to perform the splitting of the input pattern space at each node of the tree, as it defines the generation of the DT. If one of two DTs has to be chosen, it is widely accepted that

the shorter and thus simpler one is preferred, a technique known as ‘Occam’s Razor’. A possible way of doing so is to create all DTs and then choose the simpler one, but this is impractical as it would take too long to compute. Therefore, a criterion similar to ‘Occam’s Razor’ is used instead, in which high information gain attributes are placed close to the root. The entropy of a sample set \mathbf{a} where each sample can take one of c different output values is defined as

$$\text{Entropy}(\mathbf{a}) \equiv \sum_{i=1}^c -p_i \log_2 p_i, \quad (3.9)$$

where p_i is the proportion of examples in \mathbf{a} that belong to the class i . To compute entropy, 255 possible output values are quantised into $c = 16$ categories: $[0 \dots 15], [16 \dots 31]$ etc. The precise value is not lost however, but instead it is stored in the DT leaves. The information gain, as defined in Equation 3.10 is the expected reduction in entropy caused by partitioning the examples according to attribute s .

$$\text{Gain}(\mathbf{a}, s) \equiv \text{Entropy}(\mathbf{a}) - \sum_{v \in \{0,1\}} \frac{|a_v|}{|\mathbf{a}|} \text{Entropy}(a_v), \quad (3.10)$$

where a_v is the subset of \mathbf{a} for which attribute s has value v . In each internal node, the learning algorithm chooses the attribute that maximises the information gain. Once the DT is constructed, when given a query pattern q_p^x , DT is traversed from top to bottom, until a terminal node is reached. The information contained in that leaf is then used as the output-value $\hat{Q}^y(p_i)$.

3.1.3 Via Super-Resolution

The survey in search for the best available approach for solving IH did not cover only well-known IH approaches, but also included literature that combined IH and SR. This is a suited match for this project, as IH can be considered as a necessary pre-processing step before continuing onto SR. If the two steps can be merged, this improves the structure of the system and suggests that IH pre-processing step can be ignored all together and employ SR directly for both IH and the actual SR solution.

The unification of the two approaches is more or less an experimental concept, but examples of studies in literature support the possibility of applying IH via SR. Such an example is the work of Minami *et al.* [36] which demonstrated good results onto halftones obtained by error diffusion (Jarvis, Floyd&Steinberg). The results of the approach were compared against the application of a Gaussian filter and demonstrated higher PSNR. Upon comparison with other more popular IH approaches [32, 29, 6, 30, 26], they produce similar PSNR values, with the advantage of no a priori knowledge necessary. Table 5

IH Method	Lena Image	Peppers Image	Prior Knowledge
Super-Resolution [36]	30.3	29.9	none
POCS [32]	30.4	-	method & kernel
Wavelet [29]	31.7	30.7	method & kernel
Anisotropic LPA-ICI [6]	32.4	31.6	method
Kernel [30]	32.0	30.3	method
Gradient [26]	31.3	31.4	method

Table 5: PSNR performances of Minami *et al.* proposed method and previous methods

displays these results.

The results are very encouraging towards merging the two approaches. As there is no a priori knowledge available in this project's concept, the trade-off to performance is affordable and it is reasonable to expect resolution of IH when SR is applied directly on the halftones.

3.2 Super-Resolution

Conventional super resolution (SR) is performed using a number of measured low-quality images to produce a single higher optical resolution sample. The set of low-quality images can include several images of the same scene [37] or even be frames extracted from videos [38, 39, 40, 41]. SR recovers the high frequencies in the resulting image by exploiting the given images, where each has a different aliasing effect [42]. Using the measurements in aliased form, the high frequencies and thus the details in the final image can be recovered.

In mathematical terms, one can consider the SR problem as the inversion of the relation in Equation 3.11.

$$z_l = SHy_h + v \quad (3.11)$$

where z_l represents the low resolution noisy image, S performs a decimation by an integer factor s , H describes low-pass filtering, y_h is the high resolution image and v stands for additive noise. If z_l is given, the problem is to find $\hat{y} \in \mathbb{R}^{N_h}$ such that $\hat{y} \approx y_h$. A maximum-likelihood estimation is done by minimising

$$\|SH\hat{y} - z_l\|_2 \quad (3.12)$$

Since SH is rectangular, the inversion cannot be stable and thus, there are many possible solutions to the system. The difference amongst the SR approaches is the use of various priors on the image to stabilise this inversion.

There are several circumstances under which SR can be applied. In the more trivial cases, many instances of the same image will be available. This set of images might be of low quality, but because of their large number, the retrievable information is large. This happens as the low resolution images include different details and information from the same original image and therefore, when used all together, the final image will constitute of a sum of all of them. Unfortunately though, this project does not have involve numerous available input images, as there are only a few records of these images which are also of old age. Additionally, as the images under examination are those of paintings, there is no availability of images from different perspective neither.

Conclusively, there was the need of an SR solution which required as input only a single image, greyscale or coloured. This is where example-based SR (EB-SR) comes into play, as it removes the necessity for several input images. EB-SR describes the use of examples from both template and corresponding low quality input images for training the SR system and then use the extracted relation to apply it onto a further low quality images. Based on the condition that well-chosen training images that would share the same nature as the test images, the method's performance seemed most promising.

As there was no confirmation that EB-SR would indeed perform well or at all onto the project's task, the selection of the most appropriate EB-SR method did not depend on simply the literature, but also on experimentation and visual evaluation of the resulting output from various EB-SR approaches. Naturally though, the list of algorithms to be

experimented with could not be extensively long. Thus, three algorithms [21, 11, 43] were selected and then reduced to two through discussion and guidance from experts in the domain. These two algorithms involve the works of Zeyde *et al.* [21] and Yang *et al.* [11].

It is worth noting here that the algorithms were not applied using real low quality images with their corresponding high quality pairs originally, but instead the low quality images were created by scaling down the high quality images. This introduced a doubt on the effectiveness of the algorithms' performance, as a realistic low quality image set introduces the further issue of image registration and the possibility of misalignment between the low and high quality images pair in the training phase. Moreover, the images in this project were of larger resolution than the images the algorithms were tested on. Thus, the results from experimentation with these algorithms were unpredictable and included various adaptations according to the project's needs.

During research for the most appropriate method, the approach of dictionary learning appeared the most promising and appropriate to the group, based on the recent good results it has demonstrated.

The next sections cover the three considered techniques in the selected algorithms applied for EB-SR. As the approach by Zeyde *et al.* is an optimisation of the work by Yang *et al.* [11], a detailed analysis is included only for Zeyde *et al.* [21].

3.2.1 Subjective Evaluations of EB, TV, and Combined Regularization

Familiarity with the work of Anderson *et al.* [22] made the study a good starting point to approach the solution of EB-SR. In this work, three different techniques are compared to result in the most efficient one regarding noisy deblurring (image SR) and denoising. The evaluations include a *Total-Variation* (TV) approach, an *Example-Based* (EB) approach and a *Combined Regularisation* EB-TV approach. The study demonstrated that a joint approach performs better when the examples are of similar structure as the test images. If that is not the case, it is better to use overcomplete dictionaries which cover certain common fixed patterns through images. The combined regularisation proved to perform the best in the case of images with strong structure and when the dictionary is based on similar example images. The evaluations of the study were based on SSIM evaluation, which correlates best with human judgement of image quality [22].

Given a low-resolution image y , it can be defined as $y = D\tilde{H}x$ where \tilde{H} describes an ideal anti-aliasing filter which is applied onto an unknown high resolution image x . D describes a decimation operator, i.e. reduction of the image in the number of samples. One can consider the non-decimated image as $z = \tilde{H}x$, where x can be decomposed into $x = x_{LR} + x_{HR}$. The low-resolution part x_{LR} lies in the row space of \tilde{H} , while the high resolution part x_{HR} lies in the null-space of \tilde{H} such that

$$\tilde{H}x = \tilde{H}x_{LR} + \tilde{H}x_{HR} = \tilde{H}x_{LR}. \quad (3.13)$$

Single-frame SR is aiming to estimate x_{HR} given $y = Dz$, D , \tilde{H} , and a set of high resolution example-images. As the solution to example-based deblurring is the estimation of x_{HR} given $z = Hx + w$, H , and a set of high resolution images. Therefore, SR can be considered as a special case of deblurring, via up-sampling y to obtain z :

$$z = Hx + w, w = 0, H = \tilde{H} \quad (3.14)$$

Based on work by Freeman *et al.* [44], the approach towards SR evolves around nearest-neighbour patches. Image x_{HR} can be estimated by comparing blocks of y (or z), to blocks of example images after anti-aliasing and down-sampling. The nearest neighbour match is found and y 's missing high-frequencies are estimated to be the same as the nearest neighbour example image block. High-frequencies can be estimated from image blocks with similar mid-frequencies, because edges induce correlation of adjacent frequency bins [45, 46]. As the approach deals with the image in blocks, the continuity of the image was ensured via a boundary condition between adjacent blocks, e.g. finding the nearest-neighbour selection in raster-scan order and overlap each new patch with previously reconstructed image pixels [44]. However, the approach by Freeman *et al.* created objectionable artifacts in the reconstructed image, overfitting to noise in the test image. Later on, Anderson *et al.* applied this approach to deblurring and observed similar reconstruction artifacts [43]. Conclusively, Anderson *et al.* proposed a cost function for deblurring which uses a joint total variation and example-based regulariser while remaining faithful to the observed image, similar to Datsenko and Elad's approach [47, 48].

Via experimentation, the researchers found that it is better to learn the EB regularisation reconstruction jointly with the estimated restored image x . Following Chambolle's work [49], an auxiliary image was also introduced in terms of efficiency. The cost function thus solves for a restored image $x \in \mathbb{R}^{MN}$, an auxiliary image $v \in \mathbb{R}^{MN}$, and weights $\{\alpha_i \in \mathbb{R}^k\}_{i=1}^B$ for which

$$\min_{x,v,\{\alpha_i\}_{i=1}^B} \|Hv - z\|_2^2 + \lambda_1 \|x - v\|_2^2 + \lambda_2 \left\| x - \sum_i T_i \alpha_i \right\|_2^2 + \lambda_3 \|x\|_{\text{TV}}, \quad (3.15)$$

where λ_1, λ_2 and λ_3 are regularisation parameters, the columns of each $T_i \in \mathbb{R}^{MN \times k}$ contain the k most similar $m \times m$ training blocks to the i th block of the input image z . The i th training blocks are weighted by the corresponding vector α_i , and the α_i are jointly solved to minimise the error between x and the nearest neighbour reconstruction $\sum_i T_i \alpha_i$. The term $\|\cdot\|_{\text{TV}}$ denotes total variation as

$$\|x\|_{\text{TV}} \triangleq \sum_{j,k} \sqrt{|\nabla_v x|_{j,k}|^2 + |\nabla_h x|_{j,k}|^2}, \quad (3.16)$$

where the subscripts v and h describe vertical and horizontal gradients respectively, and the subscript (j, k) indexes the $((k-1)N + j)$ th entry of the vector.

Equation 3.15 demonstrates the objective function. Term $\|Hv - z\|_2^2$ ensures that the auxiliary image v matches the measured signal if blurred by H , making the operation faithful to the observations through a standard least-squares approach to deblurring. Detail enhancing is handled by term $\|x - \sum_i T_i \alpha_i\|_2^2$ by inferring high resolution information from middle-resolution information using training examples. Smoothness of the reconstruction and sharp edges are treated by term $\|x\|_{\text{TV}}$. Finally, term $\|x - v\|_2^2$ forces a close relation between the auxiliary image v and the estimated image x .

While based on Datsenko and Elad [48, 47], the regulariser differs on three points: the regularisation is carried out towards a linear combination of neighbourhood blocks instead of summed regularisations to each block, a more sophisticated method is used to adaptively regularise depending on the quality of the examples and the example-based regularised is coupled with the total variation regularisation.

The rest of the section shall now cover the most important aspects and an overall description of the approach. For further technical procedures and mathematical background, the reader is directed to the documentation of the original study [22].

As previously mentioned, this approach is block-based and so to ensure the continuity of the image, overlapping is used. Based on Freeman *et al.* [50], subsequent blocks are overlapped by one row (or column) of pixels. The distance calculation between the blocks used is the same as in the original work, involving Euclidean distance and a scalar factor $\alpha = 0.1$.

The parameters λ_1 , λ_2 and λ_3 , as defined in Equation 3.15, can lead to a variety of reconstructions when adjusted to different values. For example, if $\lambda_2 = \lambda_3 = 0$, Equation 3.15 leads to a least-squares reconstruction $x = v$, whereas if only $\lambda_2 = 0$, then x is the TV reconstruction as this was implemented by Bresson and Chan [51]. In the case though that $\lambda_1 = \lambda_3 = 0$, it gives a reconstruction based on the blockwise k -nearest neighbour reconstruction s , analogous to Freeman *et al.* approach [44].

In general, it is hypothesised that the setup of $\lambda_1 > 0$, $\lambda_2 > 0$, $\lambda_3 > 0$ will give the best reconstruction. Nevertheless, a manual setting of the optimal values for the parameters appropriate in each case of image is impractical for many applications. Instead, the choice of parameters is based on the following hypothesis: that a good set of parameter choices for an image is possible by finding good parameters for similar images and then use these similar images to train the parameters. The ideal case would be if both test and training images have similar edge content since human perceptual quality is generally scene dependent. The training images used for setting the deblurring parameters should be similar to the test images in a psycho-visual sense. The definition of *similarity* of two images is debatable, but in this case this is done according to coarse image content categories, such as ‘face’ and ‘houses’. Aside from the content of the training images, the number of the images included in the training phase also affects, as using a fewer number of training images will result in greater diversity in the nearest neighbours. This leads to overfitting and artifacts in the resulting image.

Conclusively, the severity of the artifacts depends on how well the training blocks match image blocks in the image z and thus the selected training set is of great importance for the efficiency of the system. To result in the most appropriate parameters setting, each instance is evaluated using human judges and five-full reference image quality metrics (IQMs). Evaluation of the performance with a choice of the parameters based on a full-reference IQM would allow the automation of the parameters selection. Parameters would get cross-validated by applying reconstruction on one of the training images. With ground-truth data available, this way the selection of parameters would be evaluated correctly.

Leave-one-out cross-validation was employed for the optimisation of λ_1 , λ_2 and λ_3 in the training set. This ensures that the nearest neighbour blocks T_i for a given training image are extracted from the other images.

In the experiments of this study, each sharp image is blurred via convolution with a Gaussian blur kernel with a 2.25 pixel bandwidth and adding Gaussian white noise so that the PSNR between the image and the original sharp image is 34 dB PSNR. In the λ_2 reconstructions and the case of a test image reconstruction, example image blocks from all the images in the training set were used. If the reconstruction was applied on a training image (e.g. for parameters cross-validation), then example image blocks were

extracted only from those images in the training set that were not being reconstructed.

For simplification, instead of a colour-plane-dependent TV deblurring [52], the study involved the conversion of RGB colour images to the YCbCr colour space and the deblurring of only the Y plane. The Y plane corresponds to human perception of luminance and luminance edges are most visible, therefore the approximation was considered reasonable for a full deblurring. Colour space YCbCr was chosen over more perceptually accurate colour spaces such as CIELAB, as channel Y is a linear function of the original RGB channels and thus the blur affecting Y is still H. Once the deblurring of Y channel was completed, the YCbCr resulting image was converted back to a RGB image, proceeding with visual comparisons between RGB images. In the IQM's evaluation however, only the Y colour plane was used.

In the evaluation of the results, three expert observers provided scores on 36 images in 11 experiments of random order. The evaluation of the reconstructions was categorised as very good (score=5), good (score=4), fair (score=3), unsatisfactory (score=2) and objectionable (score=1). The scoring was averaged to obtain the 'human score' for each reconstruction. While the scores were highly correlated, they were not a perfect match, as each one of the experts evaluated differently in terms of blur, overall noise and objectionable artifacts.

Displayed in Figure 9 are example results from reconstructions with no TV application (middle-left), TV reconstruction with the highest quality as judged by the three observers (middle-right) and the EB-TV reconstruction with the highest quality as judged by the observers (far-right). The reconstruction with no TV creates objectionable artifacts in the reconstruction, while the EB-TV reconstruction is sharper but noisier.

The study concluded that overall, example-based regularisation can sometimes provide better performance than TV alone and that it is possibly more useful when the test image is better matched to the training images. Moreover, the hypothesis of use of similar images was supported by the results, leading to the conclusion that if a test image is similar to the training set, parameters preferred for image reconstruction of training images may be applied to the test image to obtain a reconstruction close to that produced by the optimal parameter set for that test image. Automation of the parameter set selection was found to be best supported by SSIM, which provided the strongest correlation to hu-

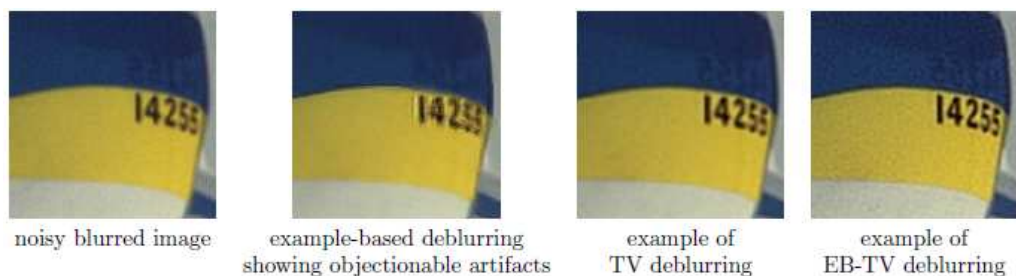


Figure 9: A noisy blurred image (far left), an example-based reconstruction that exhibits crisp edges at the expense of objectionable artifacts (second from left), total-variation (TV) reconstruction (second from right), and example-based TV (EB-TV) reconstruction (far left). Regularization parameters for TV and EB-TV were optimized over a finite set of parameter choices by maximizing the average score of three independent observers

mans' perspective and thus performed better than the rest of the examined IQMs (PSNR, ΔE of CIELAB, ΔE of s-CIELAB [17], structural similarity index (SSIM) [16] and visual information (VIF) [53]). Thus, the comparison between the reconstruction for parameters cross-validation involved SSIM, selecting the parameter set $(\lambda_1, \lambda_2, \lambda_3)$ for which the SSIM is maximised on the training set of images. However, the parameters selected from SSIM cross-validation led to worse results, sometimes dramatically worse, than the results obtained from human cross-validated parameters.

Conclusively, these experiments demonstrated that good deblurring parameters are learnable by cross-validation with human scores, especially if the training and test data are well-matched and in the case of EB-TV, if the images have structure. Compared to TV deblurring, EB-TV deblurring was found to be significantly better when the test and training images were matched in structural content and given parameters cross-validated by humans. Most importantly, the effect of the selection of the training set appeared to be crucial, as the edge content and the appropriate set of parameters for different images may differ and thus lead to different results if not appropriately chosen.

The work concludes that the use of a joint approach performs better when the training examples have similar structural elements as the test image. If that is not the case, the authors suggested using an overcomplete dictionary which includes certain common fixed patterns in actual images.

3.2.2 Single Image Scale-Up Using Sparse-Representations

The major EB-SR technique which was applied was that of Zeyde *et al.* [21]. This approach builds upon the work of Yang [11], with modifications and simplifications regarding its computational complexity and the algorithm's architecture. The present section covers in detail the approach in [21].

In comparison to Yang *et al.* method, Zeyde *et al.* approach follows a different training approach for the dictionary pair and also approaches the ability to operate without a training set by bootstrapping the scale-up task from the given low-resolution image.

This approach is based on dictionary learning, relying on two toolboxes, 'OMP' [54] and 'K-SVD' [55] to achieve a faster algorithm through sparse representations. A *Sparse-Land* local model is the main tool in the process, assuming that each patch from the considered images can be well represented using a linear combination of few atoms from a dictionary. In other words, assuming that each patch from the images can be represented with the multiplication of a dictionary by a sparse vector of coefficients [11, 56, 57, 58].

As common in SR studies, in the original works of these algorithms, the input low quality images were obtained by scaling down the high quality images. Since the scale-down of an image might result to a blurry result, the process of scale-down includes a pre-filtering that averages local pixels, so that aliasing effects are reduced.

The proposed enhancement, operates on patches extracted from scaled-up image $y_l \in \mathbb{R}^{N_h}$, aiming to estimate the corresponding patch from y_h .

Let $p_h^k = R_k y_h \in \mathbb{R}^n$ be a high resolution image patch of dimensions $\sqrt{n} \times \sqrt{n}$, extracted using $R_k : \mathbb{R}^{n_h} \rightarrow \mathbb{R}^n$ extraction operator from high resolution image y_h in location k . The locations considered, k , are only these centred around true pixels in the low-resolution image y_l (as opposed to filled-in pixels due to interpolation). The resulting set of samples will be from now on referred to as set Ω .

Using the sparse-land model, it is assumed that $p_h^k \in \mathbb{R}^n$ can be sparsely represented by $q^k \in \mathbb{R}^m$ over dictionary $A_h \in \mathbb{R}^{n \times m}$ as

$$p_h^k = A_h q^k, \quad (3.17)$$

where $\|q^k\|_0 \ll n$, where the l_0 -pseudo-norm counts the number of non-zeros in the vector q^k . Matrix A_h thus describes the high resolution patches (dictionary). A given low-resolution patch p_l^k can be represented by the same sparse vector q^k over the effective dictionary $A_l = LA_h$, with a controlled error ϵ . This implies that for a given low-resolution patch p_l^k , its sparse representation is found and the p_h^k can be recovered by multiplying this representation by the dictionary A_h .

The algorithm applies the above concept using a simple structure, split into the training phase and the reconstruction (scale-up) phase. The training phase is done off-line and thus it is unnecessary to repeated it on every application of the algorithm. The reconstruction phase uses the trained model which is saved in a file after the completion of the training phase.

In the training phase, a training set is firstly constructed from a set of high resolution training images $\{y_h^j\}$ and the low resolution images training set y_l^j is constructed by scaling-down the high resolution images by a pre-defined scale-down operator. Through this process, pairs of matching patches of high and low resolution training images $P = \{p_h^k, p_l^k\}$ are extracted. Each of these patch-pairs undergoes a pre-processing stage which extracts the luminance channel Y by transforming the image to YCbCr colour space [59, 60]. After that, high-pass filtering is applied onto p_h^k to focus the training on characterising the relation between the low-resolution patches and the edges and texture content within the corresponding high resolution images. Then, features are extracted from p_l^k by applying a Laplacian and a Gradient high-pass filters. Local patches are then extracted based on $\{k\}$, forming $P = \{p_h^k, p_l^k\}_k$. For every position $k \in \Omega$, p_h^k patch of size $[\sqrt{n}, \sqrt{n}]$ is extracted from the high resolution images and from the filtered images, p_l^k is extracted correspondingly. Every R corresponding low resolution patches are concatenated into one vector \tilde{p}_l^k of length nR . Next, dimensionality reduction is applied using *Principal Component Analysis* (PCA) on the features of the low resolution patches p_l^k , making the dictionary training step much faster and saving from computations in the subsequent algorithms.

A dictionary A_l is then trained for the low resolution patches, so that they can be represented sparsely. The $A_l \in \mathbb{R}^{n_l \times m}$ dictionary is obtained by applying K-SVD dictionary training [55]. Dictionary A_h is also constructed at this stage, formed to match the low resolution dictionary. As patch p_h^k is recovered by $p_h^k \approx A_h q^k$ approximation, the found sparse representation vector for the low resolution patch is multiplied by the high resolution dictionary to recover p_h^k . The approximation must be as exact as possible and this is why the A_h dictionary is defined as the one that minimizes the mean approximation error

$$A_h = \operatorname{argmin}_{A_h} \sum_k \|p_h^k - A_h q^k\|_2^2 = \operatorname{argmin}_{A_h} \|P_h - A_h Q\|_F^2, \quad (3.18)$$

where matrix P_h contains the high resolution training patches $\{p_h^k\}_k$ as its columns and Q contains $\{q^k\}_k$ as its columns. This is the same approach that appears in the study of Wang *et al.* [57] as well. Given that Q has full row rank, the solution to the problem is

obtained by applying the Pseudo-Inverse expression in

$$A_h = P_h Q^+ = P_h Q^T (Q Q^T)^{-1}. \quad (3.19)$$

Note that overlap is not taken into consideration. A_h should be optimised so that the resulting image is as close as possible to the original, as the reconstructed high resolution image is constructed from the positioning and averaging of high resolution patches.

Each patch of size $n \times n$ from the high resolution image is extracted by a pre-defined operator R_k at location k . Reconstructed high resolution image \hat{y}_h is constructed using Equation 3.21 and Equation 3.22

$$\hat{y}_h = y_l + \left[\sum_{k \in \Omega} R_k^T R_k \right]^{-1} \left[\sum_{k \in \Omega} R_k^T A_h q^k \right]. \quad (3.20)$$

The high resolution patch is built by term $R_k^T A_h q^k$ and then it positions it in the k -th location of the high resolution image. The term $W = \sum_k R_k^T R_k \in \mathbb{R}^{N_h \times N_h}$ describes a weight diagonal matrix, which weights every pixel in the high resolution outcome, based on the number of contributions it gets from the overlapped patches. The error computation includes y_l , as the patches in P_h are constructed from the difference image $e_h = y_h - y_l$. For the construction of image \hat{y}_h though, the algorithm should return these low-frequencies. Thus, the best dictionary A_h is defined as the solution to the optimisation task

$$A_h = \underset{A_h}{\operatorname{argmin}} \|y_h - \hat{y}_h\|_2^2 = \underset{A_h}{\operatorname{argmin}} \left\| y_h - y_l - \left[\sum_{k \in \Omega} R_k^T R_k \right]^{-1} \left[\sum_{k \in \Omega} R_k^T A_h q^k \right] \right\|_2^2. \quad (3.21)$$

Defining $W_k = R_k W^{-1} \in \mathbb{R}^{n \times N_h}$ and writing $\hat{y}_h = y_l + \sum_k W_k^T A_h q^k$, minimisation of $\|y_h - \hat{y}_h\|_2^2$ is performed with respect to A_h . If $e_h = y_h - y_l$, dictionary A_h is obtained by the minimisation of

$$A_h = \underset{A_h}{\operatorname{argmin}} \left\| y_h - y_l - \sum_k W_k^T A_h q^k \right\|_2^2 = \underset{A_h}{\operatorname{argmin}} \left\| e_h - \sum_k W_k^T A_h q^k \right\|_2^2. \quad (3.22)$$

Given $X \in \mathbb{R}^{n \times m}$, define $x \equiv \operatorname{cs}(X)$ as the column-stack version of X (where $x_{i+nj} = X_{ij}$). Using the Kronecker product property

$$\operatorname{cs}(BAC) = (C^T \otimes B) \operatorname{cs}(A) = (C \otimes B^T)^T \operatorname{cs}(A), \quad (3.23)$$

we reach

$$\operatorname{cs}(e_h) = e_h = \sum_k W_k^T A_h q^k = \left(\sum_k q^k \otimes W_k \right)^T \operatorname{cs}(A_h) = M \cdot \operatorname{cs}(A_h), \quad (3.24)$$

where $M \in \mathbb{R}^{N_h \times mn}$ is defined as $M^T = \sum_k q^k \otimes W_k$. Thus, the direct formula $M^\dagger e_h$ can be used to obtain the optimal A_h . In the study of this algorithm [21], the authors also propose an alternative approach to avoid too large matrices, based on an iterative scheme (Conjugate Gradient method).

Once the dictionaries have been constructed, the reconstruction phase can be applied. Given a test low resolution image z_l , the reconstruction phase will scale it up. In the original implementations of the algorithm, z_l is assumed to have been generated from a high resolution image y_h by the same blur and scale-down operations as used in the training phase. The trained dictionaries A_h and A_l are used to migrate from the low resolution domain to high resolution, applying K-SVD for the low resolution images and Pseudo-Inverse for the high resolution ones.

The first step in the reconstruction phase is z_l 's scale-up to the destination size y_l by scale factor s and using bicubic interpolation Q . The image is then filtered by the same R high-pass (Laplacian and Gradient) filters that were used in feature extraction step in training. Then, p_l^k patches of size $\sqrt{n} \times \sqrt{n}$ are extracted from the R images at each location $k \in \Omega$. The corresponding to the same location patches are concatenated to form a patch vector \tilde{p}_l^k leading to a collection of patches, the set $\{\tilde{p}_l^k\}_k$. Dimensionality reduction is then applied, multiplying $\{\tilde{p}_l^k\}_k$ by projection operator B to obtain set $\{p_l^k\}_k$, where each patch is of length n_l . Orthogonal Matching Pursuit (OMP) algorithm is next applied to obtain sparse-coded $\{p_l^k\}_k$ [54]. L atoms are allocated to their representation and the sparse representation vectors $\{q^k\}_k$ are calculated. These representation vectors $\{q^k\}_k$ are multiplied by the high resolution dictionary A_h to obtain the approximated high resolution patches $\{A_h q^k\}_k = \{\hat{p}_h^k\}_k$. Finally, the resulting super-resolved image \hat{y}_h is constructed from \hat{p}_h^k through the solution of the minimisation problem with respect to \hat{y}_h in

$$\hat{y}_h = \underset{\hat{y}_h}{\operatorname{argmin}} \sum_k \|R_k(\hat{y}_h - y_l) - \hat{p}_h^k\|_2^2. \quad (3.25)$$

The extracted patches from the resulting difference image $\hat{y}_h - y_l$ should be as close as possible to the approximated patches \hat{p}_h^k . The solution to this problem is a closed-form Least-Squares solution of

$$\hat{y}_h = y_l + \left[\sum_k R_k^T R_k \right]^{-1} \sum_k R_k^T \hat{p}_h^k, \quad (3.26)$$

which describes the positioning of patches \hat{p}_h^k to their appropriate locations, averaging in overlap regions and the addition of y_l to obtain the final image \hat{y}_h .

Finally, this approach also proposes a technique where the training set is extracted from the test image itself. If cases where there no access to an external set of images for training, a 'bootstrapping' technique as proposed by Glasner *et al.* [61] is possible. For the training process, z_l is used as the 'high resolution' image and its scaled-down version x_{ll} as the low resolution image. The trained dictionaries enable the reconstruction phase which will scale z_l up to y_h . This means that if the training data from the test image itself are enough to build a valid sparse-land model, scale-up of a single image using a sparse-land model is possible using one single image.

3.3 Colour Correction

This section covers an overview of different approaches towards colour correction. As done in most cases, the methods have been developed with a specific aim, therefore the section is divided according to the theme of each approach. The approaches included here spread over the categories of *colour transfer* and *image blending*, both of which consist the term of colour correction [62].

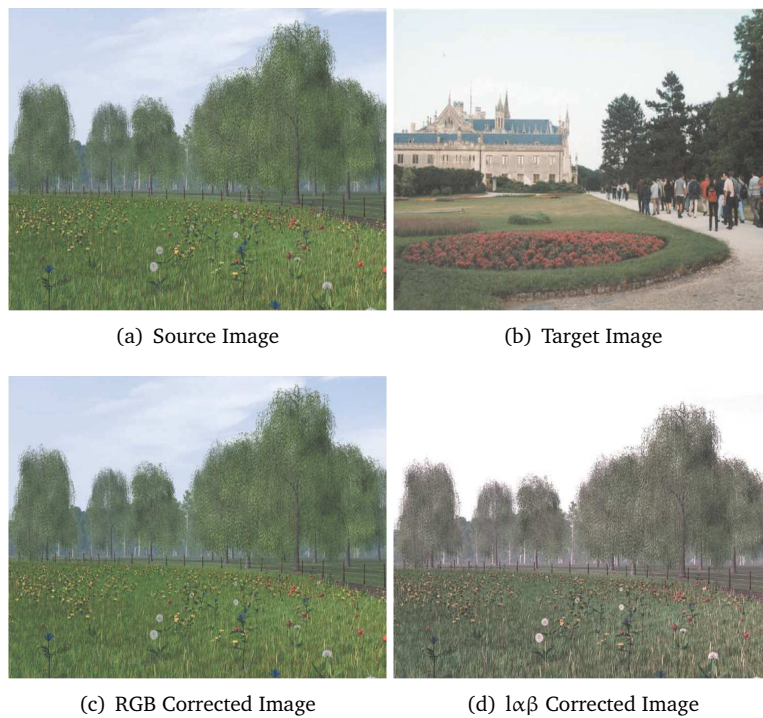


Figure 10: Colour correction in different colour spaces. From: Reinhard *et al.* [14]

3.3.1 Utilisation of Colour Space

One of the baseline approaches of colour transfer is that of Reinhard *et al.* [14]. In this approach, the authors use perception-based colour space $l\alpha\beta$ as this was developed by Ruderman *et al.* [63]. The motivation behind using colour space $l\alpha\beta$ is its minimisation of correlation between the colour channels for many natural scenes. It was developed based on human visual system understanding and provides little correlation between its axes. Reinhard *et al.* used the specific colour space to be able to apply different operations to different colour channels without dealing with the correlations between them.

The three different channels correspond to

l: Achromatic Channel

α : Chromatic Yellow-Blue Opponent Channel

β : Chromatic Red-Green Opponent Channel

This approach is based on the *Gray World Assumption*, which assumes that the average of all colours in an image is grey and thus the amount the image average differs from grey can be used to determine the illuminant's RGB [64]. Initially, both the source and the target image are transformed from RGB colour Space to $l\alpha\beta$ colour Space. From there, the mean and standard deviation along each of the colour space axes are calculated for both images and in the case of source image, the mean is subtracted from the original image. The colour correction is then performed by adding the averages of the target image onto the mean-subtracted source image values.

The algorithm can be performed either directly globally on the source image as in Figure 10, or using only parts of the target image for the desired effect, as in Figure 6. The latter removes the requirement of composition similarity between the source and

target images. One can notice that the results in $\alpha\beta$ colour space (Figure 10(c)) are significantly better than those in RGB (Figure 10(d)).

An extension to the above approach was performed by Xiao and Ma [65], to result to an approach that can directly deal with the colour transfer in any 3D space. The approach follows the same translation and scaling procedures as Reinhard *et al.*, while rotation follows a matrix obtained by decomposing the covariance matrix using SVD algorithm. Based on the fact that straightforward colour statistics can capture subjective notions of style and appearance in images [66], Xiao and Ma developed a statistics-based method that can perform colour transfer in RGB colour space by the simple use of mean and covariance matrix. The transformations between the source and target image simulate morphing an ellipsoid to fit another one. The mean value is the centre of an ellipsoid, while the eigenvalues and eigenvectors of the covariance matrix indicate the length and orientation of the three axes of the ellipsoid.

3.3.2 Naturalness Constraints

Zhang *et al.* [67] approached colour correction using the colour adaptation model *RLAB* by Fairchild and Berns [68], which simulates incomplete adaptation to a coloured target illuminant. The authors chose *RLAB* against other colour adaptation models due to its calculations in the form of matrix multiplication sequence, which allows real-time implementation. Moreover, through its strong luminance-dependent colour adjustment element, *RLAB* enables colour saturation adjustment of an image in a way that is consistent with variations people might see under different illuminant changes.

This approach is divided in the following steps

1. Characterisation of source image colours with a limited number of parameters. In this step, the source image is characterised using illumination-related properties, namely the scene CCT and the scene luminance level.
2. The source image is adjusted by applying the colour adaptation model *RLAB*.
3. ‘Tone transfer’ is performed on top of the colour transfer using the mean and standard deviation of the L^* values of the image pixels.



Figure 11: Example of images with colour adjustment to a specified CCT and luminance. Starting from the left column: original images, colour adjusted images to $cct=2000K$ and $luminance=3000cd/m^2$, colour adjusted images to $cct=11000K$ and $luminance=0.1cd/m^2$. From: Zhang *et al.* [67]

The motivation behind this work is to enable several images in the same composition to share common colour properties. In the approach, a ‘focal’ image is selected as the target and the colours of all the other images are adjusted towards this target. This way, compositions like collages could be transformed to have a more uniform and appealing presentation, while still enforcing naturalness. Figure 11 displays some results of applying colour transformation using this method, based on the CCT and luminance levels.

3.3.3 Image stitching

A common example where colour correction is required is panorama construction [69, 64]. As many images are stitched together to create a panorama, the differences between these images give undesired artifacts to the resulting panorama. The ordering, orientation, scale and illumination in these images can be different, which often leads to a panorama such as the one presented in Figure 12. These artifacts are mainly due to different viewpoints and scales, as well as the use of different capturing devices. Distortions are evidently noticeable, especially at points where the illumination changes rapidly. Figure 13 displays an example of a colour corrected panorama.

Tian *et al.* [64] approached such colour discontinuities in panoramas using *colour histograms matching*. For each pair of images that would be stitched for creating the panorama, the maximising overlapping area in the two stitching images is firstly identified. Then, a histogram mapping is applied over the overlapping areas, resulting to an estimated transformation matrix which is then applied to the two stitching images. In simple words, the first image is altered by using the colour histogram of the second image.




Figure 12: Registered images for a panorama. From: Brown *et al.* [69]



Figure 13: Registered images rendered with multi-band blending. From: Brown *et al.* [69]

Brown *et al.* [69] on the other hand, targeted a *multi-band blending* strategy, as this was proposed by Burt and Adelson [70]. Smooth transitions between images with dif-

ference overall intensities was enabled, preserving sharp details. A weight function is assigned to each image $w(x, y) = w(x)w(y)$, where $w(x)$ varies linearly from 1 at the centre of the image to 0 at the edge. Low frequencies are blend over a large spatial range using a linear weighted sum, while high frequencies are blend over a short range from the image with the maximum weight. In this manner, blurring high frequency details in the case of small registration errors between the images of the panorama is avoided.

Nevertheless, illuminant variation does not only occur on outdoors images. A study by Bascle *et al.* focused specifically on indoors illuminants [71]. Since the choices of indoors illuminants can somewhat be defined, the study involved fluorescent, incandescent light sources and natural light which describe what is usually used in offices and houses. Goal of the study was to deal with shadows, sunlight variation and changes due to the artificial illumination, i.e. switching lights on and off. The most important aspect of this approach is its automation. Using a modified *multi-layer perceptron* (MLP), the authors implemented a learning automatic colour image correction method. Since it works pixel-wise, the approach has real-time performance which makes it a complete automatic system. However, as it is the case with MLPs, for a given specific input there are multiple possible outputs. Thus, the authors defined a fourth input, other than the three RGB inputs, which defines the desired output luminance and is inserted into the middle layer of the network. Therefore, the input image can be transformed as it would be seen under the desired illumination.

An application that demonstrates similar considerations to image stitching is that of multiview video, a technique where an object or a scene is recorded using a setup of several synchronous cameras from different positions, making it possible to record moving objects or scenes. Due to the calibration differences, heterogeneous cameras and the difference in appearance from different angles, the recorded multiview video data demonstrate significant variations regarding the chrominance and luminance components [72]. These differences between the recorded camera views leads to incorrect illumination and colour reproduction after view interpolation. Fecker *et al.* [73] implemented the idea of Hekstra *et al.* [74] for applying histogram matching to adapt all camera views to a reference view in the centre of the camera setup. The cumulative histogram of the distorted image is basically adapted to match the cumulative histogram of the reference image. This simple approach allows for no assumptions on the type of distortion, the consideration of non-linear operations and correction of global discrepancies in the luminance and chrominance components, while preserving the local illumination changes that originate from surface orientation differences. In the approach by Fecker *et al.* [73], the reference image is chosen from the camera views that are closer to the centre of the camera setup and all the other camera views are corrected so that their histograms fit the histogram of the chosen reference view. Since the reference and the target camera views have a displacement, the calculation of the histograms is purely based on the overlapping area between them. This process is applied individually on each channel of the colour space that describes the camera view recordings.

3.3.4 Underwater Images

The work of Torres-Mendez and Dudek [24] focused on underwater images with the purpose of implementing image restoration for use in aquatic robots. Underwater images appear bluish, or as officially defined ‘colour depleted’, blurry and out of focus.




Figure 14: The training pairs are indicated by the squares in the original and input images. From: Torres-Mendez and Dudek [24]

These distortions take place because of the ambient light underwater and also because of the frequency-dependent scattering and absorption between the camera itself and the environment, as well as between the light source and the local environment [24].

Building on the work of Freeman [50], Torres-Mendez and Dudek utilised a Markov Random Fields' model to train learning relationships between each of the colour training images (reference images) with a corresponding colour depleted image. The system is trained on the relationship of the corresponding images, instead of only the reference images.

The authors describe two different scenarios on how the model can be used. Firstly, ground truth data are used to create simulated effects of the water using an attenuation filter. The training set is then defined by small image regions which are extracted from the ground truth data and the simulated depleted image. The training images correspond to regions of interest in terms of the variations in pixel colour values, so as to capture intrinsic statistical dependencies between the colour depleted and the ground truth pixel values. The resulting corrected image, an example of which can be seen in Figure 14, appears good, with no over-smoothing and no sudden changes in colour. The second scenario of use of the model was defined for aquatic robots real-time processing. While an aquatic robot is in the ocean, it captures video images. Being equipped with its own source of white light, the robot will stop in certain time intervals, turn its light on and take an image. As these images will be better in both clarity and colour, they can be used in the algorithm for colour correction of neighbouring frames.

3.3.5 Using Content-Based Image Retrieval

Qiu et al. approached the problem of selecting appropriate training images by using content-based image retrieval [13]. This way, one can get clusters of homogeneous image statistics and then use those as examples for example-based image processing, or specifically in this case, colour correction. The cluster of images to be used as examples can be selected according to the processing effects.

This work is based on Reinhard *et al.* [14], working in $L\alpha\beta$ colour space. As mentioned in Section 3.3.1, the colour statistics from one image can be transferred to the colour appearance of another image. This approach achieves optimisation of the process of finding correct reference images, as it identifies appropriate examples which guaran-

tee non-failure of the system and also helps to avoid manual selection of the example images.



(a) Thumbnail images in a cluster



(b) Colour distorted images

(c) Colour corrected image

Figure 15: CBIR Approach Results. From: Qiu *et al.* [13]

4 Implementation

~The present section covers all the implementation steps along with the necessary adaptations to available software, the testing procedure and the difficulties that were encountered. All processes were carried out on Intel Core 2 Duo at 2.66GHz with 3.25GB of RAM, using the environment of Matlab R2010a.

The diagram in Figure 16 outlines the implemented system. The initial steps involved the preparation and pre-processing of the datasets. Once the training and test images were defined, SR was applied starting with the training phase and continuing with application of it onto the test set. The super-resolved images would then be input to the CC phase, obtaining thus the final super-resolved colour corrected images. These were then evaluated, concluding on the efficiency of the implementation.

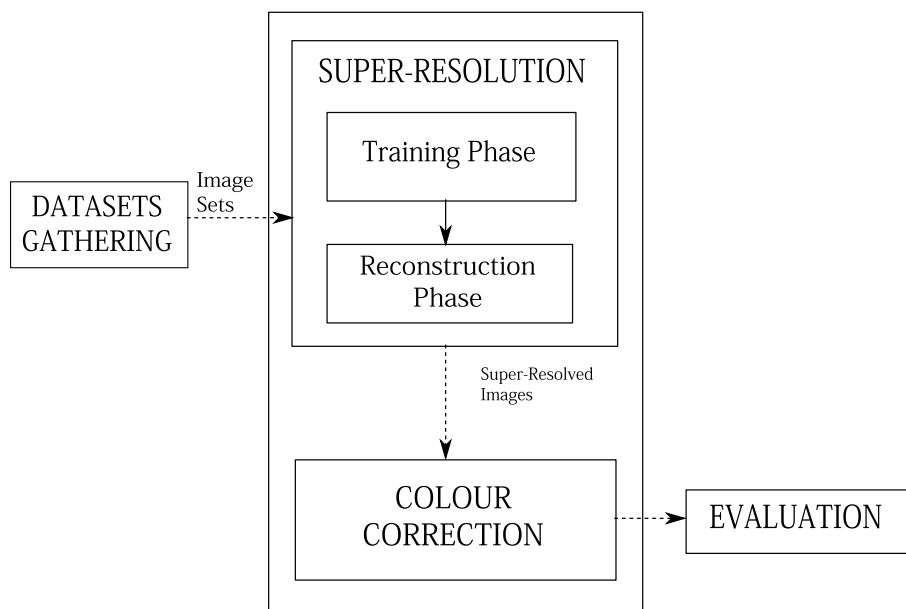


Figure 16: Outline of the implemented system

4.1 Training and Test Image Sets

Initial step in the implementation was naturally the gathering of the training and test sets. The training set included two corresponding subsets: one of low quality images which resembles the problem at hand and a 'correct' reference high quality image set. The test image set involved two subsets as well, one of which included ground-truth data, in order to later be able to quantitatively evaluate the performance of the system, and a real-test set which constituted by images of three paintings which are now considered lost.

Five ground-truth images were gathered in low and high quality by scanning pre-WWII artbooks [1, 2] and from Bridgeman Art Library Limited respectively. The high

quality images were provided at a resolution of 300dpi, while the low quality images were scanned using a Microtek ScanMaker 9800XL at a resolution of 1600dpi with no automatic adjustments, as to preserve the originality of the image, especially the halftoning pattern. These paintings were - as titled - ‘The Rape of the Daughters of Leukippos’, ‘Putti playing with garlands of fruit’, ‘The Judgement of Paris’, ‘Madonna’ and ‘Liebesgarten’, all by P.P. Rubens.

The training set was defined by 59 patches of dimensions 540×540 pixels, all extracted from one of the ground-truth images, more specifically, the ‘Rape of the Daughters of Leukippos’. These patches were manually chosen, as to cover all the variations present in the image. Being a ground-truth image, the corresponding subsets of the training set were extracted from the low and the high quality version respectively. This image, along with the rest of the ground-truth images also consisted the ground-truth test set. The real-test images were scanned following the same process as for the ground-truth data and included three paintings, also by P.P. Rubens: the ‘Resurrection of Lazarus’, ‘Diana At Her Bath, Surprised by Satyrs’ and ‘Satyrs and Bacchants’.

The same artbook [1], was used for obtaining the low quality sets both for training and testing. Only one real-test image, that of ‘Satyrs and Bacchants’ was not available from the same source, and was thus scanned from a different artbook [2]. Therefore, considerations regarding the printing technology, such as the form of the ink and the used paper, were minimised, ensuring the same conditions on almost all the gathered low quality images.

The scanning process of the low quality images was performed at high resolution, leading to large sized images. Due to the memory demanding handling that this introduces and the inability of Matlab environment to handle such large data, the scans were scaled down to ease the processing. The five ground-truth low quality images were scaled down to match the dimensions of their corresponding high quality pairs in a range of [3898..10197, 3408..10917, 3]. On the other hand, the real-test images were scaled down to dimensions in the range of [1000..1500, 1000..1500, 3].

4.2 Pre-processing Steps

Apart from the major processes of the implementation, those of super-resolution and colour correction, there were two necessary pre-processing tasks to be carried out: inverse halftoning and image registration. These were applied to prepare the data for suitable image processing, avoiding artifacts from external conditions in the system setup.

As mentioned before, the low quality training images demonstrated a halftoning pattern evidently. Therefore, before being able to proceed to image processing, it was essential to perform *inverse halftoning* (IH) to convert the halftone images to continuous-tone.

As described in Section 3.1, the most appropriate IH baseline techniques to this concept were the example-based approaches. In specific, the proposals of Kim and de Queiroz [19] and Meşe and Vaidyanathan [20] were concluded as the best suiting approaches. As the approach by Kim and de Queiroz [19] extends the work of Meşe and Vaidyanathan [20], it was solely selected to be applied, disregarding the original proposal by Meşe and Vaidyanathan. The software of the approach is freely available and it was used as on the author’s website [19].

In parallel to the application of IH by Kim and de Queiroz [19], a joint solution to halftoning was attempted via SR directly. As discussed in Section 3.1.3, EB-SR has

been successfully used in a small extent to resolve inverse halftoning. Based on the encouraging results of Minami *et al.* [36], IH was also approached directly via the EB-SR procedure of the implementation. As the approach by Kim and de Queiroz [19] provided incomplete continuous-tone images, the inverse halftoning was in the end resolved completely through the EB-SR step.

A different pre-processing was performed in the gathering of the training set. As mentioned, the training set consists of two corresponding subsets, of images of the same scene, but different quality. In the implementation, the sources of the two subsets were different, unlike most of previous studies. Most of past studies obtain the low quality image set by manipulating the high quality images. However, as this would not be possible in a real-world occurrence of the problem at hand, the low quality images were obtained from a source that would be available in any instance, the artbooks. Due to this difference between the sources of the training subsets, the corresponding patches in the training subsets were misaligned. It was thus necessary to perform *image registration* (IR).

The 59 patches of the low quality training subset underwent image registration, to be aligned with their high quality corresponding patches. Due to the uncertainty of the extent this issue could have in the final results of the implementation, the research for image registration techniques was short and concise. From research on the available techniques, image registration was in the end performed in a supervised fashion. An approach by Sheikh and Bovik [53] was applied under supervision. The method performed affine, translation, rigid and projective transforms to register the input images. This method was carried out for each of the 59 training patches, with the results from each of the transforms being subjectively evaluated and then selected.

4.3 Example-Based Super-Resolution

Through consultation and discussion with experts in the domain of EB-SR [43], it was concluded that the best approach for its implementation was dictionary-based. As referred to in Section 3.2, the selected dictionary-based method was that of Zeyde *et al.* [21] known for its performance and efficiency.

The implementation by Zeyde *et al.* [21] optimised the algorithm of Yang *et al.* [11]. It improved the execution times using the toolboxes K-SVD [55] and OMP [54], while maintaining the high performance of the original algorithm. This approach was embodied in the present implementation, adjusted from the original to the needs of this project, as this was provided from the authors [21]. Both the implementations of Zeyde *et al.* and Yang *et al.* were experimented with. However, as the implementation of Yang *et al.* is time and memory demanding, with similar performance to the optimised implementation of Zeyde *et al.*, it was finally discarded.

All steps of the original implementation, except for the core of SR, were adjusted to fit the purposes of the project. From the training to the evaluation phase, the implementation was different due to the different conditions of the present problem.

One major extension to the original algorithm that affected the structure of the implementation was the application of EB-SR onto coloured instead of greyscale images that were targeted in the original studies of Zeyde *et al.* [21] and Yang *et al.* [11]. Following the baselines of previous works, EB-SR was performed only on the luminance channel Y of the colour space YCbCr. Before any processing in the training or testing phase, the in-

volved images were transformed from RGB colour space to YCbCr. Any processing in the terms of EB-SR was performed only on the luminance channel, while the channels Cb, Cr were interpolated. The channels were then merged together into one single image, which was converted back to RGB colour space.

Next, the training set was gathered differently, which on its own introduced additional considerations. In the original works, the low quality training subset was obtained by scaling down and degrading the high quality training subset, simulating the problem under examination. This removed considerations such as image registration and size relation between the corresponding high and low quality images. Since the low quality training subset in the present terms is obtained from a different source and not the high quality image, it varies in alignment and resolution. It has no apparent relation to its high quality corresponding images, except for the scenes of the images.

By scaling down the low quality images as described in Section 4.1, a set of corresponding ground-truth data with matching dimensions between low and high quality data was obtained. However, this does not necessarily mean the alignment of the corresponding pairs. This is the reason why the pre-processing as described in Section 4.2 was necessary at this point. By completing image registration on the low quality training subset, both of the two training subsets included 59 corresponding patches, each of which were aligned to their corresponding pair. The size of all training patches was set at 540×540 , for both the low and the high quality subsets. Once the training subsets preparation was completed, the training was carried out as proposed by Yang *et al.* [11] and described in Section 3.2.2.

The parameter setting was decided via experimentation and observation of the results. These parameters were common in both the training and the test phase of the implementation, therefore the efficiency of the parameter setting was evaluated from the final super-resolved image. The experimentation led to the definition of the overlap window, the window for partial overlap for speeding the system, to $[1 \ 1]$, the number of scale-ups to be performed to 1 and the window in the low resolution images to $[3 \ 3]$. The scale-up factor s was set to $s = 3$. The parameters regarding the pre-processing in the training phase were selected as in the base study, applying high-pass filtering on the high quality images and local feature extraction from the low quality images. The filters used were the same as defined in the original works, employing a gradient $[1 \ 0 \ -1]$ and a laplacian $[1 \ 0 \ -2 \ 0 \ 1]/2$ filter.

In the training phase, around 470,000 training patch-pairs were collected, with patches of size $n = 81$. Principal component analysis was applied to reduce the feature dimensions to $n_1 = 26$. The dictionary included 1,000 atoms of $n_1 = 26$. With the above settings, the low-resolution dictionary learning takes approximately 28 minutes for 40 iterations of the K-SVD algorithm.

After the training, the core process of the SR was used as originally implemented by Zeyde *et al.* [21]. Nevertheless, the testing phase was different, mainly because of the nature of the test image set. As described in Section 4.1, the dimensions of the test images were notably large. This large size of the files, uncommon when compared to standard image sets, introduced memory handling issues in the implementation. Being memory demanding, the images could not be processed as a whole in the EB-SR procedure. Therefore, each test image was split into row extracts, with full width but a length of 600 pixels. SR was applied individually in a sequential manner on these row extracts.



Figure 17: Boundary effect as a result of SR application

The resulting super-resolved extracts were then merged back together following the original image's format to obtain the complete super-resolved image. This reveals the reason behind the setting of another parameter, that of the border of the image to be ignored, which was set to $[0\ 0]$, to preserve the border information in the image.

Despite the above arrangements, there was one artefact of the original implementation that had to be dealt with. In the implementation of Zeyde *et al.* [21], the application of the overlap-and-add method produces a boundary effect in the resulting image. Figure 17 demonstrates this artefact on a detail of an example result. In the previous study, this is resolved by cropping the resulting images' boundary. However, this solution is not applicable to the problem at hand, as that information cannot be discarded but instead, it is essential for the correct reconstruction of the complete super-resolved image. Instead, the problem was resolved using the overlapping between the row extracts of the image, discarding the erroneous pixels of a specific row extract and replacing them with erroneous-free pixels from the overlapping area of another row extract.

Reconstructing the original image as a whole after the EB-SR had been applied to the images, lead to a set of super-resolved images of the test set images. These were of the same dimensions as the original test images, as the EB-SR was applied targeting the details in the image and not a scale-up. Thus, before the application of EB-SR, these images were scaled down by the inverse of the set scale-up factor s of the EB-SR, as to target details enhancement. Therefore, the resulting super-resolved images had dimensions of range $[3898..10197, 3408..10917, 3]$.

4.4 Colour Correction

Application of colour correction required the loading of both the target and reference images at the same time. Since the involved images were of large dimensions, loading two images of over 3000×3000 pixels was not possible in the environment of Matlab. To go around this problem, the super-resolved images from the previous stage were scaled down to dimensions of range $[1000..1500, 1000..1500, 3]$. As mentioned in Section 4.1, the above range was also the selected one for the real-test image set.

Once all test images and reference images were scaled down, CC was carried out. Six different approaches were applied onto the test set, covering a selection of colour

correction techniques from both colour transfer and image blending. For the selection of the approaches, the evaluation study of Xu and Mulligan [62] was employed as a baseline for the top performing colour correction algorithms. The selected approaches included standard baseline techniques [14] and latest techniques of colour correction. Both model-based parametric approaches [69, 14, 65] and model-less non-parametric approaches [73] were applied, involving both global [14, 65] and local [69, 73] approaches, as well as operations in different colour spaces (RGB, $l\alpha\beta$, CIECAM). The selected approaches are listed in Table 6.

Index#	Name of approach	Colour Space
1	Gain Compensation [69]	intensity
2	Global Colour Transfer [14]	$l\alpha\beta$
3	Global Colour Transfer [14]	RGB
4	Global Colour Transfer [14]	CIECAM
5	Global Colour Transfer in correlated colour space [65]	RGB
6	Cumulative Histogram Mapping [73]	YCbCr

Table 6: Selected colour correction approaches [62]

The implementation for the CC approaches was adapted from the implementation of Xu and Mulligan [62], as this was available on the author’s website. Every super-resolved test image was colour corrected by all six methods and using all ground-truth high quality images as references individually. Therefore, for each super-resolved image, the application of CC resulted in 30 super-resolved colour corrected versions for each test image. It should be noted here that CC was performed using the corresponding high quality image of the test image as well, providing the highest barrier to the CC performance.

Experimentation with collages of images as reference was also carried out, but the results were not as good as when using a single image as reference and the approach was thus discarded.

4.5 Evaluation

Completion of the CC step concluded the reclamation system, allowing for evaluation. The selection of the evaluation criteria followed the guidelines of Wang *et al.* [16]. According to this study, image quality evaluation of colour altered images should not only include colour coherence, but also structural coherence, as colour correction may also affect the structure of an image. Thus, measuring the fidelity of the resulting images against the ground truth images was of double importance, both for pure evaluation of the SR application, as well as for CC.

Evaluation was carried out using three full-reference image quality metrics. Peak signal-to-noise ratio (PSNR) and structural similarity index metric (SSIM) [16] were used for computing the overall structural enhancement, while S-CIELAB metric [17] targeted the reproduction accuracy of the colours in the resulting images. Higher scores for PSNR and SSIM correspond to better image quality, while S-CIELAB denotes the colour difference between two images. SSIM improves on PSNR, combining the components of luminance, contrast and structure, correlating in this way better with human judgement of image quality [43]. S-CIELAB metric measures the accuracy of reproduction of a colour against its original when this is viewed by a human observer. Lower S-CIELAB

values indicate lower difference between two images and thus higher colour coherence.

The five ground-truth test images were evaluated using the image quality metrics, as a reference was available. However, for the real-test image set, for which the actual paintings are lost, a reference image was naturally not available. These were therefore evaluated subjectively.

An additional section of the evaluation session included an assessment of the performance of the original super-resolution algorithm against the consideration of misalignment, as this was described in Section 4.2. The effect of misalignment in training was evaluated by simulating it in the training phase of the original implementation by Zeyde *et al.* [21]. The image sets included the standard dataset of natural images as in the original implementation. Since in the original implementation, the low quality training subset is obtained directly from the high quality training subset, there is no misalignment present between the subsets. For the purposes of misalignment simulation, the patches of the low quality training subset were misaligned before the training phase. Two experimental settings were included in this simulation: (a) anti-clockwise rotation of 5° and (b) anti-clockwise rotation of 10° . The low quality training subset would still consist of patches directly obtained from the high quality training subset, but before proceeding to training, the low quality training patches would be transformed according to the simulation's settings. The rest of the procedure was performed as in the original implementation.

5 Results and Discussion

5.1 Misalignment Simulation

The carried out experimentation towards evaluation of the effect of misalignment in the training phase was performed on both a test image from the original standard set of the original implementation by Zeyde *et al.* [21] and a ground-truth test image from the present project.

Examination of the robustness of the original implementation against a simulated misalignment showed that the algorithm is strongly affected by misalignment in training. In Table 7, the performance of the algorithm seems to strongly decrease as the level of the (rotational) misalignment increases. Figure 18 demonstrates examples of super-resolved images as these were obtained from the original application of the algorithm and the two simulated misalignment applications.

Image	Original	Rotation +5°	Rotation +10°
Baboon	23.5	20.8	20.4
Barbara	26.8	22.3	21.2
Bridge	25.0	19.8	19.0
Coastguard	27.1	23.0	22.1
Comic	24.0	18.1	17.3
Face	33.5	27.5	26.6
Flowers	28.4	20.8	19.8
Foreman	33.2	23.4	21.8
Lena	33.0	24.3	22.9
Man	27.9	21.7	20.6
Monarch	31.1	21.4	20.3
Pepper	34.0	24.1	22.4
ppt3	25.2	17.3	16.3
Zebra	28.5	18.2	17.2
Average	28.66	21.62	20.56

Table 7: PSNR Measurements for simulated misalignment in training of SR

Including misalignment in the EB-SR of a ground-truth test image from the problem at hand demonstrated similar behaviour, with strong artifacts present in the resulting image. For comparison, the same image was tested with alignment corrected by image registration. Figure 19 demonstrates an extract from an example result of the application of EB-SR when the training set involves misalignment (middle) and when image registration is employed to resolve the misalignment issue (right). The left-most image in Figure 19 displays the ground-truth high quality image. In order to focus on the effects on SR, only the luminance channel (which was super-resolved in implementation) is presented.

5.2 Super-Resolution

The performance of SR was evaluated separately, before proceeding to the completion of the system through CC. In other words, the resulting super-resolved images were




Figure 18: A subset of the resulting SR-ed images, displaying starting from the left column: the original test image, the original application of the algorithm on the image, application with training misalignment by rotation of 5° and application with misalignment with rotation of 10° .

evaluated before being sent to the CC procedure. This set of measurements focused on structural coherence and thus PSNR and SSIM image quality metrics were computed for each of the super-resolved results for the ground-truth test images. Table 8 displays the evaluation scores of the super-resolved images in comparison to the scores obtained by the original input, the low quality test images, in terms of both PSNR and SSIM, testing against the original high quality ground-truth image.

Image	PSNR		SSIM	
	Original	Super-Resolved	Original	Super-Resolved
Leukippos	16.059	16.427	0.449	0.347
Liebesgarten	15.542	15.599	0.375	0.396
Madonna	15.462	16.133	0.342	0.393
Paris	13.475	13.616	0.424	0.447
Putti	13.953	14.100	0.326	0.346

Table 8: Super-Resolution evaluation for ground-truth test set

5.3 Colour Correction

Upon completion of the CC step, the final images of the implementation were gathered. Thus, the evaluation of the output of the CC step did not only evaluate the CC as is, but also the application as a whole.

As mentioned in Section 4.5, the application of CC was expected to affect not only the colour coherence of the test images, but also their structural coherence [16]. Therefore, measurements in terms of the structural metrics of PSNR and SSIM were repeated for the colour correction results as well.



Figure 19: Effect of misalignment in training of SR. The image is a detail from painting 'Madonna' by P.P. Rubens. The left-most image is the original high-quality image in greyscale, the middle image is the resulting image from SR when there is rotational misalignment of 10° counter-clockwise in the training low-quality set and the right-most image displays the resulting image from SR when image registration is applied on the low-quality training set to correct the misalignment

For each of the eight test images, both real-test and ground-truth, 30 colour corrected super-resolved resulting images were gathered. Quantitative evaluation using the three full-reference image quality metrics of PSNR, SIIM and S-CIELAB was executed only on the ground-truth data for which an actual reference image existed. Qualitative and subjective evaluation was carried out for the total of the resulting images, both to evaluate the real-test images and to compare the resulting conclusions to the quantitative evaluation's results.

Due to space limitations, it was impossible to include here the whole of the results. The reader is directed to Appendix C, which includes all the detailed quantitative measurements. Table 9 presents the mean and standard deviation statistics of the results from the CC procedure, categorised by CC algorithm. As previously described, these data include only the ground-truth test set. The results have been categorised according to the CC method applied, due to their common input, the super-resolved images. Each CC method lead to a sum of 25 resulting images: five ground-truth test images, each colour corrected using each of the five ground-truth high quality images as reference. To ease the reader, Table 9 is sorted from higher overall performance to lower.

Method	PSNR		SSIM		S-CIELAB	
	μ_{PSNR}	σ_{PSNR}	μ_{SSIM}	σ_{SSIM}	$\mu_{\text{S-CIELAB}}$	$\sigma_{\text{S-CIELAB}}$
Alg. 4 [14]	14.974	1.147	0.347	0.045	22.816	3.774
Alg. 3 [14]	14.942	1.180	0.347	0.043	22.442	3.849
Alg. 6 [73]	14.814	1.191	0.333	0.054	22.937	3.714
Alg. 5 [65]	14.705	1.165	0.341	0.043	24.378	4.381
Alg. 2 [14]	13.800	1.696	0.319	0.056	29.984	4.623
Alg. 1 [69]	10.563	3.089	0.274	0.066	27.470	6.355

Table 9: PSNR, SSIM and S-CIELAB mean (μ) and standard deviation (σ) statistics per CC algorithm

The statistics show a large variation among the test set, especially regarding the colour reproduction. The results for each of the algorithm vary vastly, probably because of the large variation between the reference images. The two algorithms that performed the best (Alg. 3 and Alg. 4) [14] demonstrated similar scores both regarding structural and colour coherence. Alg. 3 showed the best performance in colour coherence, but Alg. 4 performed slightly better in terms of structural coherence. However, apart from the two best performing algorithms, an absolute ranking of the rest of the algorithms is not possible, as they vary in performance for structural and colour coherence.

The metrics of PSNR and SSIM correlate with each other, except for the case of Alg. 5 and Alg. 6, in which have an interchangeable performance. The ranking of colour coherence followed the same ranking as for the structural coherence in overall, except for Alg. 1 - Alg. 2 and Alg. 3 - Alg. 4 which were switched in order.

One might observe that the average values for the obtained S-CIELAB scores in Table 9 are reportedly high. Therefore, the results were also studied in terms of enhancement instead of solely a colour coherence evaluation. Table 10 demonstrates the change in colour difference against the high quality ground-truth image for (a) the initial input test set and (b) the super-resolved colour corrected output. These values coincide to the best cases of the super-resolved colour corrected results for each of the ground-truth test images, excluding the cases where the corresponding high quality image was used



Figure 20: An example colour difference map, demonstrating ground-truth test image ‘The Judgement of Paris’ at a difference threshold $T = 10$. The white parts of the image correspond to difference $< T$

for colour correction of a specific test image. The colour difference across the images demonstrates a drop in colour difference of the range [1.549..6.901].

Test Image	S-CIELAB		
	Original image	SR+CC result	Improvement
Leukippos	22.330	18.291	4.039
Liebesgarten	18.708	15.438	3.270
Madonna	22.826	21.277	1.549
Paris	24.152	17.251	6.901
Putti	30.347	23.501	6.846

Table 10: S-CIELAB scores between (a) original low quality test images and (b) super-resolved (best) colour corrected images against ground truth high quality images

The evaluation of the results from the application onto the ground-truth test set was not limited to quantitative scores. In order to verify that the numerical results corresponded to a stronger change in the significant parts of the test image (e.g. faces and skin tones), difference maps were obtained for each of the resulting images. The difference threshold was set to 10 and 15 in two individual experiments. Figure 20 shows an example difference map, for the best result regarding ground-truth test image ‘The Judgement of Paris’.

The difference maps supported the quantitative measurements, with the quantitatively best cases demonstrating the least difference in the area of faces, difficult shading textures and skin surfaces. This evaluation also revealed results that while quantitatively they were considered of low performance, they produced very low difference specifically on areas of skin tones. Such an observation contributed to further consideration towards better utilisation of difference maps.

A full-reference evaluation was inapplicable for the results from the set of real-test



Figure 21: Detail of the image from the P.P. Rubens' lost painting 'Satyrs and Bacchants', as originally in low quality (left) and after SR+CC (right)

images, as a reference for these was unavailable. Instead, they were qualitatively, subjectively evaluated via observation. Figure 21 displays an example result from the process applied on the image of the lost painting by P.P. Rubens, titled 'Satyrs and Bacchants'.

The evaluation of these images was performed subjectively. Due to space limitations, the images cannot be included in this document, but the total of the image sets is provided at the author's website: <http://www.stud.hig.no/~101530/>. The low quality images, the high quality images, the super-resolved results and the colour-corrected images are all included for supportive observation of the results on both the real-test images as well as the ground-truth test images.

6 Conclusions and Further Work

6.1 Contributions

Through this study, the endeavour towards reclamation of lost art has been successfully completed. Most importantly, the present work proves that such an attempt is possible with numerous possibilities in this direction. Initialising a novel branch of colour imaging applications in fine art, the results of this work and their evaluation are strong motivation for further studies. Not only regarding the possibilities in reclamation of lost art, but also into utilisation of the outcomes of the work for the domains of super-resolution, inverse halftoning and colour correction.

By applying this unique image dataset to state-of-the-art techniques of different domains, it extends their applicability and performance evaluation to a brand new, highly demanding dataset. It brings on board the necessity of image registration in the case the training sets are misaligned, as well as the application of super-resolution to a real-world problem. It confirms the efficiency of the colour correction technique by Reinhard *et al.* [14], despite its simplicity. It demonstrates however, that in the given context, the colour space of $\text{l}\alpha\beta$ does not provide the best results as in the case of different datasets, but instead, it is less efficient than the colour spaces of RGB and CIECAM.

The method moreover questions the ranking of colour correction well-known approaches and certifies the efficiency of dictionary-based super-resolution. It furthermore verifies the applicability and efficiency of super-resolution based approaches towards inverse halftoning, extending the work of Minami *et al.* [36].

6.2 Conclusions

Firstly, the study extends current knowledge of super-resolution efficiency, by examining the robustness of EB-SR technique against misalignments in the training set. As discussed in Section 5.1, the degree of misalignment in the training set causes severe artifacts in the resulting image. As the level of misalignment increases, the performance of the algorithm decreases drastically. This effect is evident both visually as one can observe from the resulting images (see Figure 19) and also quantitatively (see Table 7). With this issue brought up, further research works onto real-world problems in which the training subsets of low and high quality are from different sources can attend this matter. Proving that image registration is more than sufficient to resolve this artefact, also gives a solution to the raised matter. Thus, the contribution to the domain of super-resolution technique not only covers evaluation of its robustness against misalignment, but also proposes and verifies a solution via image registration.

Secondly, a more demanding dataset is employed, extending the applicability of EB-SR by demonstrating the efficiency of the technique on a new type of images. This study involved the first application of EB-SR on images of fine art and verified the successful completion of the technique onto these exclusive images.

Furthermore, the experimental application of EB-SR directly on halftone images supported and extended the attempt by Minami *et al.* [36] towards an IH method via SR.

The halftoning pattern was absent from the resulting super-resolved images, successfully producing continuous-tone images with IH being embedded in SR instead of being a separate pre-processing step.

A set of colour correction techniques was additionally utilised towards a solution to the problem at hand, evaluating methods from both colour transfer and image blending. The colour corrected images demonstrated a significant enhancement regarding colour coherence in the originally distorted low quality images. The scores of S-CIELAB, evaluating the fidelity of the results in colour, showed a vast decrease in colour difference, which was supported by visually pleasing results.

The state-of-the-art techniques that were selected for the completion of this task were evaluated and then ranked according to their performance. It was concluded that while the top performing algorithms, Reinhard *et al.* in RGB and CIECAM colour space, were evidently the most efficient ones, an absolute and definite general ranking of the techniques in terms of both structural and colour coherence was not possible to be stated.

Observation of the results and experimentation in the implementation lead to the belief that the training set plays a critical role regarding the performance of the technique. This was naturally expected, due to the example-based approach of the implementation. It became even more evident in the application of CC, where the effect of the selected reference was strongly affecting the resulting image. Therefore, the study identified the high importance of correct ‘examples’ and their significant effect on the final output.

The above point is closely related to a consideration that arose from the beginning of the project, that of ‘similarity’ between the context of two images. Since a reference image is not available for real-test images, a good reference would have to represent a different scene and naturally, a different painting. Initially, this similarity was targeted via the selection of the same theme (mythology) and the same artist (P.P. Rubens). However, these two factors are not enough to absolutely secure the similarity of two images of paintings. The best result for a real-test image could be obtained only if the reference image and the training image were similar enough. The question of how to judge this similarity is not easy to answer. It is not possible to quantify, but is more subjective and only art experts would be able to identify a perfect match to a given test image. Application of the implementation onto ground-truth data though, revealed similarities in the images of ‘Leukippos’-‘Madonna’, ‘Paris’-‘Liebesgarten’ and others. Use of specific images as reference performed better for some test images, revealing an underlying relation. Use of this relation to support the selection of the perfect reference image is a strong extend to the present work.

6.3 Future Work

The current work is a first step in an entirely new area of computer-assisted heritage preservation, and suggests extensive possibilities for further study.

As a first step, future work would attempt a more extensive evaluation of image quality, through a combination of both structure and colour coherence metrics [62]. It would also explore the notion of similarity between two artworks, in order to automate and optimise the selection of a reference image for a specific test image. Additional approaches could be applied and tested towards both SR and CC, based on the revelation of the efficiency of state-of-the-art algorithms in the present concept.

Identification of the importance of training motivates an approach of a more complete

training, employing a cross-validation training to better generalise the solution. Such an extension of the training phase would cover a much bigger image set, with better applicability.

Upon gathering of the results from EB-SR, it was observed that the colours were altered, changed towards the colours included in the training set. This suggests a possibility of applying colour correction via super-resolution completely. Further work towards this direction could involve the examination of efficiency of colour correction via SR.

As mentioned in Section 5.3, difference maps were estimated to display the areas at which colour difference was handled in a better fashion. While the best performing algorithms covered most of the images' area, there were some less efficient ones which performed exceptionally well on the areas of faces or skin tones. Utilisation of this knowledge could lead to a combined approach of colour correction, which deals with the input image locally and adapts accordingly.

Another important limitation in the project was the minimal evaluation of the results for the real-test images. Qualitative evaluation was not possible and there is not a non full-reference image quality metric that is perfectly applicable in the context of this problem. A more complete evaluation of the outcome could be performed via psychophysical experiments involving a large group of observers, both experts and non-experts.

A concluding step of this work would be the high quality printing of the resulting images and their extensive evaluation via psychophysical experiments. Such experimentation would support the resulting quality of the artworks and define whether their quality is high enough for such a print to be displayed alongside surviving, 'real' paintings. Demonstrating that it is possible to do so would be the ultimate milestone for this work.

Bibliography

- [1] Dulberg, F. 1932. *Rubens*. E.A. Seeman.
- [2] Koetschau, K. 1911. *Album des Kaiser-Friedrich - Museums Zu Berlin*. Seemann & Co.
- [3] Lau, D. L. & Arce, G. R. 2001. *Modern Digital Halftoning*. CRC Press.
- [4] Lee, H.-C. & Sharma, G. 2005. *Introduction to Color Imaging Science*. Cambridge University Press.
- [5] Karni, Z., Freedman, D., & Shaked, D. 2010. Fast inverse halftoning. *31st International Congress on Imaging Science*.
- [6] Foi, A., Katkovnik, V., Egiazarian, K., & Astola, J. September 2004. Inverse halftoning based on the anisotropic lpa-ici deconvolution. In *Proceedings of the International TICSP Workshop on Spectral Methods and Multirate Signal Processing, SMMSP 2004*, 49–56, Vienna.
- [7] Zitova, B. & Flusser, J. 2003. Image registration methods: A survey. *Image and Vision Computing*, 21(11), 977–1000.
- [8] Brown, L. G. December 1992. A survey of image registration techniques. *ACM Computing Surveys*, 24(4), 325–376.
- [9] Elad, M. November 2011. From sd to hd improving video sequences through super-resolution. final Workshop Presentation, Computer-Science Department, The Technion - Israel.
- [10] Chang, H., Yeung, D.-Y., & Xiong, Y. 2004. Super-resolution through neighbor embedding. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [11] Yang, J., Wright, J., Ma, Y., & Huang, T. 2008. Image super-resolution as sparse-representation of raw image patches. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [12] Park, A. S., Park, M., & Kang, M. G. May 2003. Super-resolution image reconstruction, a technical overview. *IEEE Signal Processing Magazine*, 20(3), 21–36.
- [13] Qiu, G. 2004. From content-based image retrieval to example-based image processing. *School of Computer Science, The University of Nottingham*.
- [14] Reinhard, E., Ashikhmin, M., Gooch, B., & Shirley, P. 2001. Color transfer between images. *IEEE Computer Graphics and Applications*.
- [15] Chalmers, A., Daly, S., McNamara, A., Myszkowski, K., & Troscianko, T. 2000. Image quality metrics. In *IN SIGGRAPH TUTORIALS 2000*.

- [16] Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. April 2004. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Processing*, 13(4), 600–612.
- [17] Zhang, X. & Wandell, B. A. 1996. A spatial extension of cielab for digital color image reproduction. *SID Symposium Technical Digest*, 27, 731–734.
- [18] Johnson, G. M. & Fairchild, M. D. 2003. top down description of s-cielab and ciede2000. *Color Research and Application*, 28, 425–435.
- [19] Kim, H. Y. & de Queiroz, R. L. 2003. Inverse halftoning by decision tree learning. In *Proc. IEEE International Conference on Image Processing (ICIP)*, Barcelona, Spain. <http://www.lps.usp.br/~hae/software/invhalf/index.html>.
- [20] Meşe, M. & Vaidyanathan, P. October 2001. Look-up table (lut) method for inverse halftoning. *IEEE Transactions on Image Processing*, 10(10).
- [21] Zeyde, R., Elad, M., & Protter, M. 2010. On single image scale-up using sparse-representations. In *Curves and Surfaces*, Boissonnat, J.-D., Chenin, P., Cohen, A., Gout, C., Lyche, T., Mazure, M.-L., & Schumaker, L. L., eds, volume 6920 of *Lecture Notes in Computer Science*, 711–730. Springer. <http://www.cs.technion.ac.il/~elad/software/>.
- [22] Anderson, H. S., Gupta, M. R., & Hardeberg, J. Y. 2012. Subjective evaluations of example-based, total variation, and combined regularization for image processing. *Proceedings SPIE 8296*, 82960S.
- [23] Huang, Y.-H., Wang, P.-C., Chou, C.-C., & Lin, K.-P. 2011. An automatic selective color transfer algorithm for images. In *Proceedings of the 2011 ACM Symposium on Applied Computing, SAC '11*, 66–71, New York, NY, USA. ACM.
- [24] Torres-Mendez, L. & Dudek, G. 2005. A statistical learning-based method for color correction of underwater images. *Research on Computer Science, Centre for Intelligent Machines, McGill University*.
- [25] Neelamani, R., Nowak, R. D., & Baraniuk, R. G. Oct. 2002. Winhd: Wavelet-based inverse halftoning via deconvolution. *IEEE Transactions on Image Processing*.
- [26] Kite, T. D., Damera-Venkata, N., Evans, B. L., & Bovik, A. C. Sept. 2000. A fast, high-quality inverse halftoning algorithm for error diffused halftones. In *IEEE Trans. Image Processing*, volume 9, 1583–1592.
- [27] Wang, L., Hua, B.-S., & Li, X. 2010. Adaptive energy diffusion for blind inverse halftoning. In *PCM (1)*, 470–480.
- [28] Neelamani, R., Nowak, R., & Baraniuk, R. September 2000. Model-based inverse halftoning with wavelet vaguelette deconvolution. In *Proc. IEEE Int. Conf. Image Processing - ICIP*, volume 3, 973–976, Vancouver, Canada.
- [29] Xiong, Z., Orchard, M. T., & Ramchandran, K. 1999. Inverse halftoning using wavelets. In *IEEE Transactions on Image Processing*, volume 8(10), 1478–1482.

- [30] Wong, P. W. 1995. Inverse halftoning and kernel estimation for error diffusion. In *IEEE Transactions on Image Processing*, volume 4(4), 486–498.
- [31] Meşe, M. & Vaidyanathan, P. June 2002. Tree-structured method for lut inverse halftoning and for image halftoning. In *IEEE Trans. on Image Processing*, volume 11 (6), 644–655.
- [32] Hein, S. & Zakhor, A. 1995. Halftone to continuous-tone conversion of error-diffusion coded images. In *IEEE Transactions on Image Processing*, volume 4(2), 208–216.
- [33] Bozkurt Unal, G. & Çetin, A. December 2001. Restoration of error-diffused images using projection onto convex sets. In *IEEE Trans. Image Processing*.
- [34] Mitchell, T. M. 1997. *Machine Learning*. WCB/McGraw-Hill.
- [35] Quinlan, J. R. 1986. *Induction of Decision Trees*, volume 1. Machine Learning.
- [36] Minami, Y., Azuma, S.-i., & Sugie, T. Aug. 2010. An inverse halftoning algorithm based on super-resolution image reconstruction. In *SICE Annual Conference 2010, Proceedings of*, 1110–1113.
- [37] Elad, M. & Feuer, A. December 1997. Restoration of single super-resolution image from several blurred, noisy and down-sampled measured images. In *IEEE Trans. Image Process.*, volume 6 of 12, 1646–1658.
- [38] Farsiu, S., Robinson, D., Elad, M., & Milanfar, P. 2004. Advances and challenges in super-resolution. *International Journal of Imaging Systems and Technology, Special Issue on High Resolution Image Reconstruction*, 14(2), 47–57.
- [39] Farsiu, S., Elad, M., & Milanfar, P. January 2006. Multi-frame demosaicing and super-resolution of color images. *IEEE Trans. on Image Processing*, 15(1), 141–159.
- [40] Farsiu, S., Robinson, D., Elad, M., & Milanfar, P. October 2004. Fast and robust multi-frame super-resolution. *IEEE Transactions on Image Processing*, 13(10), 1327–1344.
- [41] Farsiu, S., Elad, M., & Milanfar, P. 2006. Video-to-video dynamic superresolution for grayscale and color sequences. *EURASIP Journal of Applied Signal Processing*, 2006(Article ID 61859), 1–15. Special Issue on Superresolution Imaging.
- [42] Elad, M. & Datsenko, D. 2009. Example-based regularization deployed to super-resolution reconstruction of a single image. *The Computer Journal*, 52(1), 15–30.
- [43] Anderson, H. S. & Gupta, M. R. 2009. Joint deconvolution and imaging. In *Proc. SPIE Conf. on Computational Imaging*.
- [44] Freeman, W. T., Jones, T. R., & Pasztor, E. C. Mar/Apr 2002. Example-based super-resolution. *IEEE Comp. Graph. and Appl.*, 22(2), 56–65.
- [45] Simoncelli, E. 1999. Modeling the joint statistics of images in the wavelet domain. *Proc. SPIE*, 188–195.

- [46] Crouse, M. S., Nowak, R. D., & Baraniuk, R. G. 1998. Wavelet-based statistical signal processing using hidden markov models. In *IEEE Trans. Signal Processing*, volume 46, 886–902.
- [47] Datsenko, D. & Elad, M. 2007. Example-based single document image super-resolution: a global map approach with outlier rejection. In *Multidim. Syst. Sig. Process*, volume 18, 103–121.
- [48] Elad, M. & Datsenko, D. 2007. Example-based regularization deployed to super-resolution reconstruction of a single image. *The Computer Journal*, 50(4), 1–16.
- [49] Chambolle, A. 2004. An algorithm for total variation minimization and applications. *J. Math. Imaging and Vision*, 20, 89–97.
- [50] Freeman, W., Pasztor, E., & Carmichael, O. 2000. Learning low-level vision. *International Journal of Computer Vision*, 20(1), 25–47.
- [51] Bresson, X. & Chan, T. F. 2008. Fast dual minimization of the vectorial total variation norm and applications to color image processing. *Inverse Problems and Imaging*, 2(4), 455–484.
- [52] Wen, Y., Ng, M., & Huang, Y. November 2008. Efficient total variation minimization methods for color image restoration. *IEEE Trans. Image Processing*, 17(11).
- [53] Sheikh, H. R. & Bovik, A. C. Feb 2006. Image information and visual quality. *IEEE Trans. Image Processing*, 15(2), 430–444.
- [54] Elad, M., Rubinstein, R., & Zibulevsky, M. Efficient implementation of the k-svd algorithm using batch orthogonal matching pursuit. Technical report, CS, Technion, April 2008.
- [55] Aharon, M., Elad, M., & Bruckstein, A. November 2006. The k-svd: An algorithm for designing of overcomplete dictionaries for sparse representation. In *IEEE Trans. on Signal Processing*, volume 54 of 11, 4311–4322.
- [56] Yang, J., Wright, J., Huang, T. S., & Ma, Y. November 2010. Image super-resolution via sparse representation. *Trans. Img. Proc.*, 19(11), 2861–2873.
- [57] Wang, J., Zhu, S., & Gong, Y. January 2010. Resolution enhancement based on learning the sparse association of image patches. *Pattern Recognition Letters*, 31(1).
- [58] Lou, Y., Bertozzi, A., & Soatto, S. August 2010. Direct sparse deblurring. *J. Math. Imag. Vis.*
- [59] Poynton, C. 1996. *A Technical Introduction to Digital Video*. John Wiley & Sons, Inc.
- [60] (1982-1986-1990-1992-1994-1995). Studio encoding parameters of digital television for standard 4:3 and wide-screen 16:9 aspect ratios. Rec. ITU-R BT.601-5. Section 3.5.
- [61] Glasner, D., Bagon, S., & Irani, M. October 2009. Super-resolution from a single image. *International Conference on Computer Vision ICCV*.

- [62] Xu, W. & Mulligan, J. 2010. Performance evaluation of color correction approaches for automatic multi-view image and video stitching. In *IEEE Int. Conference on Computer Vision and Pattern Recognition*, San Francisco, USA. http://ia.cs.colorado.edu/~wxu/color_correction.htm.
- [63] Ruderman, D., Cronin, T., & Chiao, C. 1998. Statistics of cone responses to natural images: Implications for visual coding. *J. Optical Soc. of America*, 15(8), 2036–2045.
- [64] Tian, G., Gledhill, D., & Taylor, D. 2002. Colour correction for panoramic imaging. *Proceedings of the Sixth International Conference on Information Visualisation, IV'02*, 483–488.
- [65] Xiao, X. & Ma, L. 2006. Color transfer in correlated color space. *Proc. 2006 ACM International conference on virtual reality continuum and its applications*, 305–309.
- [66] Tumblin, J. & Ferwerda, J. A. September 2001. Guest editors' introduction: Applied perception. *IEEE Comput. Graph. Appl.*, 21(5), 20–21.
- [67] Zhang, X., Chao, H., & Tretter, D. 2011. Image color transfer with naturalness constraints. *Society for Imaging Science and Technology, 19th Color and Imaging Conference Final Program and Proceedings*.
- [68] Fairchild, M. & Berns, R. 1993. Image color-appearance specification through extension of cielab. *Color Research & Application*, 18, 178–190.
- [69] Brown, M. & Lowe, D. G. 2003. Recognising panoramas. *Proc. ICCV*, 2, 1218–1225.
- [70] Burt, P. J. & Adelson, E. H. October 1983. A multiresolution spline with application to image mosaics. *ACM Trans. Graph.*, 2(4), 217–236.
- [71] Bascle, B. & Bernier, ŃO. andăLemaire, V. 2006. Illumination-invariant color image correction. *Advances in Machine Vision, Image Processing And Pattern Analysis, Lecture Notes in Computer Science*, 4153, 359–368.
- [72] López, J., Kim, J. H. nad Ortega, A., & Chen, G. December 2004. Block-based illumination compensation and search techniques dfor multiview video coding. *Proc. Picture Coding Cymp. (PCS 2004), San Fransisco, CA*, 15–17.
- [73] Fecker, U., Barkowsky, M., & Kaup, A. 2008. Histogram-based prefiltering for luminance and chrominance compensation of multiview video. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(9), 1258–1267.
- [74] Hekstra, A., Beerends, J., Ledermann, D., de Caluwe, F., Kohler, S., Koenen, R., Rihs, S., Ehram, M., & Schlauss, D. 2002. PVQM - a perceptual video quality measure. *Signal Processing: Image Communication*, 17(10), 781 – 798.

A Scientific Papers

The present chapter includes two scientific papers which resulted from the research carried out in the project. The first paper ‘Computer-Aided Reclamation of Lost Art’ was submitted to ECCV Visart Workshop in 2012, (<http://printart.isr.ist.utl.pt/visart/>). It includes a description of the system as a whole, focusing on the contribution to cultural heritage. It is therefore the official paper submission for the purposes of CIMET evaluation.

The second paper that is included here was submitted to CVMP 2012 (<http://www.cvmp-conference.org/>), as ‘Learning Super-Resolution from misaligned examples’. This paper concentrates specifically on extended evaluation of SR approaches when misalignment is present between the low and high quality training sets.

Computer-Aided Reclamation of Lost Art

Anonymous ECCV submission

Paper ID ***

Abstract. There are numerous approaches towards restoration of art, including computer applications as aid to manual performance. However, to our knowledge, it has not been attempted to recuperate high quality images of missing or presumably destroyed works of art. While, by definition, these works will never again be available in their original form, it may be feasible to considerably enhance the quality of preserved photographic reproductions. In this paper a pioneering approach to reclaim high quality images of lost works of art is presented, employing a combination of super-resolution and colour correction. The techniques are performed by example, utilising correspondence between artworks of similar nature, which are currently available both in low and high quality. With extensive prior knowledge in the domains of super-resolution and colour correction, specific applicable approaches were studied, implemented and tested, concluding to the most efficient. Experimental results are highly promising and encouraging, revealing a new research path in colour imaging for fine art.

1 Introduction and Motivation

There has not been a previous attempt as the present endeavour towards reclamation of lost art. We propose a novel strategy for making the images of selected pieces of our lost artistic heritage accessible to the public, an endeavour we term *‘lost art reclamation’*. To have the ability to regain paintings considered to have been lost forever was an unrivalled motivation for this study.

With no background studies to support this effort, the possibilities and expectations of the work were unpredictable. The efficiency of a chosen technique was uncertain, not only due to the question at hand, but also because of the nature of the involved images. Being a fine art application, the images of paintings that were included demonstrated major characteristic differences with standard datasets of natural images, against which most image processing approaches were tested. Thus, the reported performance of state-of-the-art techniques was under examination, in parallel to the search for a solution.

The most important records of these lost works are colour photos published in period art books, especially between the two World Wars. At publication, these were at the cutting age of printing technology, but as compared to contemporary images, these samples are of low quality, with obvious halftoning patterns, erroneous colours, and low resolution. As the only records of significant artworks, these samples are of inadequate quality regarding their importance. They are

045 however the only available starting point to work towards a high quality image 045
046 of a lost artwork, aiming to obtain good enough quality for the reclaimed artwork 046
047 to be exhibited alongside surviving works. 047

048 With halftoned, erroneously coloured prints of low resolution as a starting 048
049 point, the study targeted a combined solution which resolves these matters. 049
050 Super-resolution (SR) was concluded as the most appropriate technique towards 050
051 resolution enhancement. More specifically, *example-based super-resolution* (EB- 051
052 SR) [1, 2] was the selected approach, incorporating art images available in both 052
053 low and high quality. 053

054 Super-resolution involves three different categories: interpolation approaches, 054
055 reconstruction-based approaches and learning (example-based) approaches. In 055
056 the concept of the stated problem however, only a learning approach is appli- 056
057 cable. Interpolation and reconstruction-based families of techniques require as 057
058 input a series of low resolution images of the same scene. This is impossible 058
059 in the examined case, as the scene of a painting’s image cannot change. No- 059
060 tions as obtaining several low resolution images of the same scene from different 060
061 perspectives, or under different viewing conditions are unsuitable in this case. 061

062 A learning based method involves the training of two corresponding dictio- 062
063 naries, one of low quality A_l and one of high quality A_h [3, 4]. The included 063
064 patches correspond to the exact same scene in an image. Given a low resolution 064
065 image, the learning is employed to correctly replace patches of the test image 065
066 with the most appropriate ones from the dictionary. Identifying a training patch 066
067 in the dictionary A_l as the closest match to the given test patch, its correspond- 067
068 ing high quality patch from dictionary A_h will be returned as the SR result 068
069 for that specific patch. Repeating this method for all patches in the test image 069
070 results in a complete, super-resolved high quality image. 070

071 Naturally, the low quality training dictionary includes halftones, erroneously 071
072 coloured low resolution patches which correspond to continuous-tone high qual- 072
073 ity patches in the high quality dictionary, mimicking the problem at hand. The 073
074 application of EB-SR on the low quality test image, enhances the resolution 074
075 and also, due to the example-based approach, resolves the matters of halfton- 075
076 ing and incorrect colours automatically. The process of *inverse halftoning* [5] is 076
077 thus merged with the SR approach [6]. Upon completion of SR, the resulting 077
078 images are of continuous-tone, high resolution images with attempted corrected 078
079 colours. Nevertheless, to ensure the colour correctness of the test image, further 079
080 processing is necessary. A crucial reason to that is the patch-oriented applica- 080
081 tion of EB-SR, which leads into inconsistencies of colour in the general image. 081
082 Colour correction is necessarily applied next to ensure the global correction of 082
083 the image. 083

084 *Colour correction* (CC) [7] has been developed in various studies, either as 084
085 *colour transfer* [8, 9], which relates to the artistic transformation of an image to- 085
086 wards the appearance of another, or *image blending* [10, 11], which is a common 086
087 notion amongst image stitching techniques. The extensive set of implemented ap- 087
088 proaches and their large variety of applications account for the lack of extensive 088
089 evaluation regarding the performance of CC approaches. 089

Combining the above techniques, the present work attempted to resolve the problems identified in the low quality images: the halftoning in the images, their low resolution and erroneous colours. The joint solution followed an example-based approach, utilising ground-truth data in low and high quality and applying this learning to given input.

2 Methodology

The proposed method for a complete system of lost art reclamation involves two major steps: example-based super-resolution and colour correction. Before the processing however, the training and test sets are gathered. The training set involves two corresponding subsets: one of low quality images which resembles the problem at hand and a ‘correct’ reference high quality image set.

The test image is transformed to colour space $YCbCr$ and EB-SR is applied on the luminance channel Y , while channels $CbCr$ are interpolated. The resulting super-resolved images are transferred to the CC subsystem.

2.1 Super-Resolution

Due to the large amount of possible solutions to a high quality image that corresponds to a given test low quality image, regularisation is necessary in SR. The proposed regularisation involves the use of a local sparse-land model on image patches as performed by Zeyde *et al.* [4]. The main concept is that each patch of the test images can be well represented using a linear combination of few atoms from a dictionary. In other words, it can be represented by the multiplication of the dictionary by a sparse vector of coefficients. The algorithm approaches the problem as the minimisation of

$$\|SH\hat{y} - z_l\|, \quad (1)$$

where z_l is the low quality image, \hat{y} is the reconstructed high quality image, S is the blurring factor and H is the decimation operator.

Firstly, image z_l is scaled up, using bicubic interpolation Q , to bring it to the dimensions of the high quality image y_h . This results in $y_l = Qz_l$, from which patches are extracted, estimating the corresponding patch of y_h . A set of k locations of true pixels (which have not been obtained by interpolation) define the position of patches in the image y_l to be extracted by operator R_k . The same operation is carried out onto y_h , resulting to set Ω of patches p_h^k . According to the sparse-land model [12, 13], patches p_h^k can be represented sparsely by q^k vectors of dictionary A_h , as $p_h^k = A_h q^k$. Obtaining the same sparse representation against y_l , the set p_h^k corresponds to set $p_l^k = R_k y_l$. Operator $L^{all} = QSH$ transforms the complete high resolution image y_h to the low resolution one y_l . It can therefore be assumed that $p_l^k = Lp_h^k + \tilde{v}_k$, where L is a local operator being a portion of L^{all} and \tilde{v}_k is the additive noise in this patch. As $p_h^k = A_h q^k$,

$$Lp_h^k = LA_h q^k. \quad (2)$$

135 Utilising the relation between the low and high resolution patches $p_l^k = Lp_h^k + \tilde{v}_k$, 135
 136 it can be stated that 136

$$137 LA_h q^k = Lp_h^k = p_l^k - \tilde{v}_k, \quad (3) \quad 137$$

139 implying that 139

$$140 \|p_l^k - LA_h q^k\|_2 \leq \epsilon, \quad (4) \quad 140$$

142 where ϵ is related to the noise power σ of v . Thus, patch p_l^k can be represented 142
 143 by q^k over the dictionary $A_l = LA_h$, and patch p_h^k can be recovered from q^k 143
 144 multiplied by A_h . 144

145 In the training phase, sets $\{y_h^j\}$ and $\{y_l^j\}$ are constructed from the high and 145
 146 low quality images respectively. Before proceeding to the processing of these sets, 146
 147 it has to be noted, that the training set consists of extracted parts of the training 147
 148 images. The training low and high quality images are firstly resized to match in 148
 149 physical resolution. *Image registration* is applied onto the low quality training 149
 150 images so that they match their corresponding high quality images, in order to 150
 151 avoid artifacts due to misalignment in training. Image registration is applied in 151
 152 a supervised fashion to ensure that the most efficient transformation is applied. 152

153 Once sets $\{y_h^j\}$ and $\{y_l^j\}$ are gathered, patches are extracted from them, 153
 154 leading to $P = \{p_h^k, p_l^k\}$. Pre-processing is performed on p_h^k to remove low fre- 154
 155 quencies and feature extraction is performed on p_l^k as to extract local features 155
 156 that correspond to their high-frequency content. Two filters are involved in this 156
 157 feature extraction, a gradient and a Laplacian filter. Next, principal component 157
 158 analysis is performed onto the p_l^k features to reduce the feature vector, resulting 158
 159 to $\{\tilde{p}_l^k\}$. Dictionary A_l is then trained so that p_l^k can be represented sparsely. 159
 160 Finally, dictionary A_l is trained so that it matches A_l . 160

161 Given a test image z_l , this is interpolated to y_l and sharpened by spatial non- 161
 162 linear filtering. Then, for every location $k \in \Omega$, patches p_l^k are sparse-coded using 162
 163 A_l resulting to $\{q^k\}$. These q^k are then used to recover p_h^k , by multiplying with 163
 164 A_h . Merging $\{p_h^k\}$ to obtain \hat{y}_h , the overlapping areas are averaged to get the 164
 165 final resulting image. This last step is performed as the following minimisation 165
 166 with respect to \hat{y}_h : 166

$$167 \operatorname{argmin}_{y_h} \sum_k \|R_k(\hat{y}_h - y_l) - \tilde{p}_h^k\|_2^2. \quad (5) \quad 167$$

170 As the involved painting images are highly memory demanding, they cannot 170
 171 be processed as a whole. Therefore, the image to be tested is firstly split into 171
 172 row extracts of full width but length of 600 pixels. EB-SR is performed on these 172
 173 extracts individually and the resulting super-resolved extracts are merged to 173
 174 obtain the complete image. 174

175 An additional consideration is the boundary effect when the patches are 175
 176 overlapped and added. In the proposed method, cropping the boundary to ignore 176
 177 the effect is not possible [4], as the boundary of an extract is necessary for the 177
 178 final composition of the image as a whole. Instead, the overlap between the 178
 179 extracts is arranged so that this effect is dealt with. 179

2.2 Colour Correction

Upon gathering of the super-resolved results, CC process is carried out. The ground-truth high quality images are used one by one as reference for each test image. Depending on the similarity of the test image with the reference image, the performance of CC varies for each application.

Upon completion of the image processing procedures, evaluation of the results follows, using image quality metrics which correspond to both the structure and the colours of the resulting images.

3 Evaluation Setup

3.1 Selection of approaches

In our experimental setup, EB-SR followed the implementation of Zeyde *et al.* [4], who optimised the algorithm by Yang *et al.* [3] using toolboxes K-SVD [14] and OMP [15]. The implementation of Zeyde *et al.* enhanced the algorithm both in terms of computational complexity and algorithm architecture. The final results were obtained from this optimised implementation due to the long execution time and demanding memory handling of the original by Yang *et al.*

Regarding the application of CC, the selected algorithms covered a variety of approaches, incorporating six methods from both colour transfer and image balancing. These included standard baseline approaches [9] and latest techniques of colour correction. Both model-based parametric approaches [10, 9, 8] and model-less non-parametric approaches [11] were applied, involving both global [9, 8] and local [10, 11] approaches, as well as operations in different colour spaces (RGB, $l\alpha\beta$, CIECAM).

The first applied approach was an exposure/gain compensation technique by Brown and Lowe [10] (Alg. 1). In this colour balancing approach the intensity gain level of the component images was adjusted to compensate for appearance differences caused by different exposure levels. The second approach was the standard colour correction method by Reinard *et al.* [9], which deals with global colour transfer (Alg. 2). Despite the simplicity of this approach, its reported efficiency is significant, making it a baseline work which was extended in later works [8]. Its structure is centred around a linear transformation based on the mean and standard deviation of the global colour distributions of the source and target images in the uncorrelated $l\alpha\beta$ colour space. Developed by Ruderman *et al.* [16], $l\alpha\beta$ colour space minimises correlation between channels. The colour value of a target pixel was defined as:

$$g(G_t) = \mu_s + \frac{\sigma_s}{\sigma_t}(C_t - \mu_t). \quad (6)$$

This approach was also applied in the colour spaces RGB (Alg. 3) and CIECAM (Alg. 4). Based on the work by Reinhard *et al.* [9], Xiao *et al.* [8] (Alg. 5) extended to a correlated RGB global colour approach which makes use of an ellipsoid mapping scheme. Lastly, the cumulative colour histogram mapping approach by

Fecker *et al.* [11] (Alg. 6) was applied. This involved a closest neighbour mapping scheme to select the corresponding colour level of the source image to each level of the target image.

3.2 Training and Test Image Sets

Five ground-truth data, both in low and high quality were obtained by scanning pre-WWII artbooks [17] and from Bridgeman Art Library Limited respectively. The artworks involved were all by Peter Paul Rubens, titled ‘Madonna’, ‘The rape of the daughters of Leukippos’, ‘Liebesgarten’, ‘The Judgement of Paris’ and ‘Putti’. These ground-truth data were divided in a training and a test set. Four of the images formed the test set in order to quantitatively evaluate the performance of the processes, while the fifth image was used for training.

Additionally, a real-test image set was defined from images of paintings by P.P. Rubens which are nowadays lost, namely the ‘Resurrection of Lazarus’, ‘Diana At Her Bath, Surprised by Satyrs’ and ‘Satyrs and Bacchant’s’.

The total of the scans except for image ‘Satyrs and Bacchant’s’ was gathered using the same artbook [17] to minimise effects from different printing technologies on the input data. The scans were performed using a Microtek ScanMaker 9800XL at a resolution of 1600dpi with no automatic adjustments, as to preserve the originality of the image, especially the halftoning pattern. The ground-truth data were scaled, using bicubic interpolation, to match their corresponding high quality pairs’ dimensions in a range of [3898..10197, 3408..10917, 3], while the images of the lost paintings were scaled down to dimensions in the range of [1000..1500, 1000..1500, 3].

3.3 Evaluation Criteria

As stated by Wang *et al.* [18], image quality evaluation of colour altered images should not only include colour coherence, but also structural coherence, as colour correction may also affect the structure of an image. Thus, measuring the fidelity of the resulting images against the ground truth images was of double importance, both for pure evaluation of the SR application, as well as for CC.

Three full-reference image quality metrics were used in the evaluation of the ground-truth test set. The metrics of peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) [18] were used to evaluate the overall structural enhancement evaluation, while the metric S-CIELAB [19] targeted the reproduction accuracy of the colours in the resulting images. Higher scores for PSNR and SSIM correspond to better image quality, while S-CIELAB denotes the colour difference between two images. SSIM improves on PSNR, combining the components of luminance, contrast and structure, correlating in this way better with human judgement of image quality [20]. S-CIELAB metric measures the accuracy of reproduction of a colour against its original when this is viewed by a human observer. Lower S-CIELAB values indicate lower difference between two images and thus higher colour coherence.

3.4 Implementation details and parameter settings

The code for SR was adapted from the original implementation of Zeyde *et al.* [4], as this is available on the author’s website. The CC approaches were also used directly from the implementation by Xu and Mulligan [7]. The system was implemented as a complete process, merging the steps of EB-SR and CC along with the evaluation.

The approaches and the evaluation criteria in the implementation use the same parameters as stated in the original papers. The only exception regards the dimensions of the involved images, which are larger than the standard input.

4 Results and Discussion

The ground-truth test and real-test sets were inserted in the application, resulting with 30 super-resolved, colour corrected images per input test image. The single output of SR was colour corrected by six different approaches, using all five ground-truth high quality images as reference. It can be noted here that CC was performed using the corresponding high quality image of the test image as well, providing the highest barrier to the CC performance. Table 1 presents the resulting performance measurements for the ground-truth test data. The results have been categorised according to the CC method applied, due to their common input, the super-resolved images. Each CC method included evaluation of 25 resulting images of five ground-truth data, each colour corrected using each of the five ground-truth high quality images as reference.

Table 1. PSNR, SSIM and S-CIELAB mean (μ) and standard deviation (σ) statistics per CC algorithm

Method	μ_{PSNR}	σ_{PSNR}	μ_{SSIM}	σ_{SSIM}	$\mu_{S-CIELAB}$	$\sigma_{S-CIELAB}$
Alg. 1 [10]	10.563	3.089	0.274	0.066	27.470	6.355
Alg. 2 [9]	13.800	1.696	0.3194	0.056	29.984	4.623
Alg. 3 [9]	14.942	1.180	0.347	0.043	22.442	3.849
Alg. 4 [9]	14.974	1.147	0.347	0.045	22.816	3.774
Alg. 5 [8]	14.705	1.165	0.341	0.043	24.378	4.381
Alg. 6 [11]	14.814	1.191	0.333	0.054	22.937	3.714

First of all, a large variation among the test set is observed, especially concerning the colour reproduction of the images. Such variation is natural, as the results is strongly dependent on the reference image. Moreover, the overall top performing algorithms (Alg. 3 and Alg. 4) have very similar scores, with Alg. 3 having obtained the best performance in colour coherence, but with Alg. 4 proved the best in structural coherence. One can observe that in PSNR terms, the lowest performance was given by Alg. 1, but the same algorithm performed better in

S-CIELAB terms. Therefore, while the top performing algorithms are evident, a definite general ranking amongst the rest of the algorithms cannot be stated.

Moreover, the obtained S-CIELAB scores are reportedly high. However, when compared to the initial difference between the original low quality test image and the high quality image as demonstrated in Table 2, the enhancement on colour reproduction is significant. There is a colour difference drop of a range of [1.549..6.901]. The S-CIELAB values for the super-resolved colour corrected images correspond to the best performing case in terms of S-CIELAB, when the reference image is of a different scene than the test image.

Table 2. S-CIELAB scores between (a) original low quality test images and (b) super-resolved (best) colour corrected images against ground truth high quality images

Test Image	Original S-CIELAB	SR+CC S-CIELAB
Leukippos	22.330	18.291
Liebesgarten	18.708	15.438
Madonna	22.826	21.277
Paris	24.152	17.251
Putti	30.347	23.501

The real test images could not be evaluated using the full-reference metrics as a reference does not exist. Instead, they were subjectively evaluated. Figure 1 displays an example result from the process applied on the image of the lost painting by P.P. Rubens, titled ‘Satyrs and Bacchants’.

5 Conclusions and Further Work

The present work is the first study that attempts a reclamation of lost art. Not only it proves that such an endeavour can indeed be successful, but it also reveals the possibilities of further exploration of the domain. The results of the work, along with their evaluation, suggest the feasibility of an entirely novel branch of colour imaging applications in fine art, as well as observations and considerations for the domains of super-resolution, inverse halftoning and colour correction.

Applying this unique image dataset to state-of-the-art techniques, it extends their applicability and performance evaluation to a brand new, more demanding dataset. It brings on board the necessity of image registration in the case the training sets are misaligned, as well as the application of super-resolution to a real-world problem. It proves that in this case the colour correction technique by Reinhard *et al.* [9], is the most effective, despite its simplicity. It questions the ranking of colour correction well-known approaches and verifies the efficiency of dictionary-based super-resolution. It furthermore certifies a solution of inverse halftoning via super-resolution, extending the work of Minami *et al.* [6].




Fig. 1. Detail of the image from the P.P. Rubens' lost painting 'Satyrs and Bacchants', as originally in low quality (left) and after SR+CC (right)

The current work is a first step in an entirely new area of computer-assisted heritage preservation, and suggests extensive possibilities for further study. Firstly, future studies could attempt a better evaluation of image quality, through a combination of both structure and colour coherence metrics [7]. It could further explore the notion of similarity between two artworks, in order to automate and optimise the selection of a reference image for a specific test image. Further approaches could be applied and tested both towards SR and CC, based on the revelation of the efficiency of state-of-the-art algorithms in the problem's concept.

A concluding step of this work would be the printing of the resulting images and their extensive evaluation via psychophysical experiments. Such experimentation would support the resulting quality of the artworks and define whether their quality is high enough for such a print to be displayed alongside surviving, 'real' paintings. Demonstrating it is possible to do so would be the ultimate milestone for our work.

References

1. Freeman, W., Pasztor, E., Carmichael, O.: Learning low-level vision. *International Journal of Computer Vision* **20** (2000) 2547
2. Freeman, W.T., Jones, T.R., Pasztor, E.C.: Example-based super-resolution. *IEEE Comp. Graph. and Appl.* **22** (2002) 56–65

- 405 3. Yang, J., Wright, J., Ma, Y., Huang, T.: Image super-resolution as sparse- 405
406 representation of raw image patches. IEEE Conference on Computer Vision and 406
407 Pattern Recognition (CVPR) (2008) 407
- 408 4. Zeyde, R., Elad, M., Protter, M.: On single image scale-up using sparse- 408
409 representations. In Boissonnat, J.D., Chenin, P., Cohen, A., Gout, C., Lyche, 409
410 T., Mazure, M.L., Schumaker, L.L., eds.: Curves and Surfaces. Volume 6920 of 410
411 Lecture Notes in Computer Science., Springer (2010) 711–730 411
- 412 5. Wong, P.W.: Inverse halftoning and kernel estimation for error diffusion. In: IEEE 412
413 Transactions on Image Processing. Volume 4(4). (1995) 486498 413
- 414 6. Minami, Y., Azuma, S.i., Sugie, T.: An inverse halftoning algorithm based on super- 414
415 resolution image reconstruction. In: SICE Annual Conference 2010, Proceedings 415
416 of. (2010) 1110–1113 416
- 417 7. Xu, W., Mulligan, J.: Performance evaluation of color correction approaches for 417
418 automatic multi-view image and video stitching. In: IEEE Int. Conference on 418
419 Computer Vision and Pattern Recognition, San Francisco, USA (2010) 419
- 420 8. Xiao, X., Ma, L.: Color transfer in correlated color space. Proc. 2006 ACM 420
421 International conference on virtual reality continuum and its applications (2006) 421
422 305–309 422
- 423 9. Reinhard, E., Ashikhmin, M., Gooch, B., Shirley, P.: Color transfer between im- 423
424 ages. IEEE Computer Graphics and Applications (2001) 424
- 425 10. Brown, M., Lowe, D.G.: Recognising panoramas. Proc. ICCV **2** (2003) 1218–1225 425
- 426 11. Fecker, U., Barkowsky, M., Kaup, A.: Histogram-based prefiltering for luminance 426
427 and chrominance compensation of multiview video. IEEE Transactions on Circuits 427
428 and Systems for Video Technology **18** (2008) 1258–1267 428
- 429 12. Elad, M., Aharon, M.: Image denoising via learned dictionaries and sparse repre- 429
430 sentation. In: In CVPR. (2006) 17–22 430
- 431 13. Elad, M., Aharon, M.: Image denoising via sparse and redundant representations 431
432 over learned dictionaries. Image Processing, IEEE Transactions on **15** (2006) 3736 432
433 –3745 433
- 434 14. Aharon, M., Elad, M., Bruckstein, A.: The k-svd: An algorithm for designing of 434
435 overcomplete dictionaries for sparse representation. In: IEEE Trans. on Signal 435
436 Processing. Volume 54 of 11. (2006) 4311–4322 436
- 437 15. Elad, M., Rubinstein, R., Zibulevsky, M.: Efficient implementation of the k-svd 437
438 algorithm using batch orthogonal matching pursuit. Technical report, CS, Technion 438
439 (2008) 439
- 440 16. Ruderman, D., Cronin, T., Chiao, C.: Statistics of cone responses to natural images: 440
441 Implications for visual coding. J. Optical Soc. of America **15** (1998) 2036–2045 441
- 442 17. Dulberg, F.: Rubens. E.A. Seeman (1932) 442
- 443 18. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: 443
444 From error visibility to structural similarity. IEEE Trans. Image Processing **13** 444
445 (2004) 600–612 445
- 446 19. Zhang, X., Wandell, B.A.: A spatial extension of cielab for digital color image 446
447 reproduction. SID Symposium Technical Digest **27** (1996) 731–734 447
- 448 20. Anderson, H.S., Gupta, M.R.: Joint deconvolution and imaging. In: Proc. SPIE 448
449 Conf. on Computational Imaging. (2009) 449

Learning Super-Resolution from misaligned examples

Maria Lena Demetriou
Gjøvik University College
Gjøvik, Norway

marialena.dem@gmail.com

Jon Y. Hardeberg
Gjøvik University College
Gjøvik, Norway

jon.hardeberg@hig.no

Gabriel Adelmann
West Wing Doctors
United Kingdom
adel_g@yahoo.com

ABSTRACT

This paper examines the applicability and performance of Example-Based Super-Resolution (EB-SR) techniques to problems in which the low-quality training set is not constructed from the high-quality training set, but instead, it is obtained from a different source. This is the case when EB-SR is applied on images that, by their specific nature, cannot be perfectly simulated by processing high-quality images towards the gathering of the low-quality training images. Applying EB-SR to printed images can represent such a problem. The low-quality training images will in that case be gathered from scanning actual halftones, leading to a somewhat incoherent relation between the low- and high-quality training sets, despite their image scene correspondence. Assembling the corresponding training sets from different sources, gives rise to one major consideration, that of **misalignment**. The present work examines the effect of misalignment in the corresponding low- and high-quality pairs of the training set in an EB-SR approach.

Two sets of experiments were carried out. In the first phase of experimentation, EB-SR was performed including simulated misalignment on standard datasets on which a selected EB-SR method has previously reported high performance. Then, experiments were extended to a real-world problem in which the high-quality images are digitally available, but low-quality images are available as scanned prints. The conceptual framework for this particular choice is that of **Reclamation of Lost Art (RLA)**. In the second experiment, image registration is enforced on the RLA to identify at what degree it reduces the effects of misalignment.

Being a pioneering research step, the present work explores a new dimension in the robustness of EB-SR approaches. It examines this in terms of a simulated problem and a real-world problem. **Image registration** is also employed in the real-world problem to examine whether it can eliminate the effects caused by misalignment.

Keywords

Super-Resolution, Fine Art, Image Registration, Misalignment

1. INTRODUCTION

The investigations reported in this paper were initially suggested as part of the feasibility assessment of an original idea, that of virtual recovery of lost paintings [2]. It was considered that surviving low-resolution colour prints of paintings subsequently lost in conflagrations such as World War II could theoretically be much enhanced, in an objective, fact-based manner, to yield better-quality images. Such images would make available, for the first time in decades, at least a fair approximation of the original aspect of physically destroyed core pieces of our cultural heritage. Additionally, the paradigm could be extended to other types of low-quality, culturally or historically important images.

While this concept, termed as **reclamation of lost art**, theoretically encompasses large number of (presumably) destroyed artworks, for the purposes of the project, three paintings by Peter Paul Rubens were selected, titled ‘The Resurrection of Lazarus’, ‘Satyrs and Bacchants’, and ‘Diana At Her Bath, Surprised by Satyrs’. These three paintings, from the inventory of the Kaiser-Friedrich Art Museum in Berlin, are thought to have been lost during the last days of World-War II. All that is available regarding their original aspect are images from art books published in the 1930s [3]. Naturally, the quality of these images (at the time of publication, at the cutting edge of printing technology) appears, to the contemporary eye, poor indeed, with low resolution and distorted colours. The physical resolution at which these are available is also very small.

It soon became obvious that one very important (though not the only) task when tackling the problem at hand is that of super-resolution (SR). Conversely, it appeared that **Image Reclamation** could be an especially interesting context for discussion of the concept of SR. For completeness and perspective, the team endeavoured to discuss SR both on a simulated and a real-world concept. It is desirable to make the point of the dual relationship between an exciting new idea and a well-established theoretical concept.

The rest of the paper covers an overview of SR approaches in Section 2 including a discussion regarding the approach selection most suitable for the purposes of the present project. Next, the implementation of the simulation and the real-world problems is covered in Section 3. Section 4 then presents the results from the implementation, while Section 5 completes the paper with final observations and conclusions.

2. METHODOLOGY

The notion of enhancing the resolution of an image is closely related to numerous applications. Satellite imaging, medical imaging

and surveillance videos analysis have made it necessary to be able to increase the resolution, without producing any visual artifacts. Mechanical approaches cover the adjustment of the image acquisition technologies by reducing the size of the pixels on the sensor or increasing the size of the chip. These techniques however, produce undesirable artifacts or in the case of the chip's size increase, implicate very high cost. Therefore, instead of a physical solution, an image processing one was preferable, leading to the rapid development of the super-resolution methods over the last three decades.

The term of **Super-Resolution** is often used to describe the fusion of several low-resolution images of the same scene to obtain a high-resolution image [6] instead of the more general concept of the increase of resolution in an image. Overall, the SR methods can be divided into two categories according to their general approach, either as **frequency domain based**, or as **spatial domain based**. Frequency domain techniques are not used in practical applications though, due to their inefficiency in embodying prior knowledge. There are three different subcategories of SR in the spatial domain, based on their baseline as this falls in the concept of **interpolation**, **reconstruction** or **learning**.

2.1 Selecting the SR method

The criterion behind the selection of the SR approach to be applied in this work, was its applicability to the real-world problem of reclamation of lost art, as this was mentioned previously. Dealing with images of paintings brings on board its own several difficulties and limitations into play. The nature of the scenes is completely different than that of a standard dataset's, raising the doubt of applicability of known methods simply because of the structure of the images themselves. Additionally, the availability of several images of the same scene, which are often employed in methods such as SR is not present in this context. A painting's scene describes one and only one scene and does not provide the possibility of manipulations the same way natural images, for example, do. Concepts such as the use of images from different view points or different capture times are absent when dealing with archival images of paintings.

Therefore, in the processing of a fine art image, there is only one specific scene that incorporates all available information. This is a major constraint especially in terms of SR, as the methods of SR of both interpolation and reconstruction require as input a series of low resolution images. This fact solely makes these two SR methods inapplicable to the problem at hand. Luckily, the last possible method of SR, that of learning (or example-based) SR is indeed applicable. With the original painting being lost, the 'gold standard images' in an EB-SR approach were high-quality contemporary images of stylistically similar, but still-existent, paintings by the same artist. An EB-SR approach is the most appropriate solution to this problem, as it can utilise the correspondence between low- and high-resolution images of paintings that still exist. Learning from a database of low- and high-resolution images (or from a single image), EB-SR can be applied to reclaim a high-quality image of a painting even if there is only a low-quality sample available.

There are numerous studies in the field of EB-SR, but it was impossible to examine each one against the identified implications of the project. Completing a survey on the state-of-the-art in EB-SR and through discussion with experts in the field, a single EB-SR approach was chosen and applied. This approach was chosen not only based on the evaluation purposes of the present work, but also as the most promising one for a complete system for the RLA prob-

lem. The selected approach is that of Zeyde, Protter and Elad [8]. The specific study optimises the implementation of the algorithm by Yang et. al. [7] through the use of toolboxes K-SVD [1] and OMP [4]. The dictionary-based technique incorporated by this approach appeared as the most promising and practical one. Its short execution times and high performance ensured it as the optimal choice. Despite the fact that in preliminary results the original algorithm by Yang et. al. performs visually better on art images, its long execution times made it difficult to apply it on the images implicated in the project, as these are very memory demanding due to their high physical resolution.

3. IMPLEMENTING SUPER-RESOLUTION

The main consideration in the referred algorithm is the employment of **Sparse Representation** towards EB-SR. The positioned hypothesis is that each patch from a given image can be well represented using a linear combination of few atoms from a dictionary, i.e. the multiplication of the dictionary by a sparse vector of coefficients [8]. Therefore, obtaining two dictionaries, one of high-quality A_h and one of low-quality A_l , with related sparse representation, this information can be used to enhance the resolution of a test low-quality image, using the corresponding information in the dictionaries.

An overview of the implementation of the EB-SR approach as adapted from the original implementation by Zeyde et. al. [8] follows in the rest of this section. The description covers the adaptation of the algorithm towards its application to the RLA problem as previously defined. The misalignment simulation implementation follows the same approach, with the only difference involving the assembly of the training images. The low-quality training images were obtained as in the original studies, by scaling down the high-quality images. To then introduce misalignment in the simulation implementation, the low-quality training images were rotated by 5° and 10° in the counter-clockwise direction, while retaining the original dimensions of the image.

3.1 Training Phase

In the original implementation of the algorithm, the learning phase, i.e. obtaining the trained two dictionaries A_h and A_l was done using the relation between the two dictionaries. The images used as input to EB-SR in the original work were simulated low-quality images, obtained by scaling down by the inverse of the desired SR's level, $s = 3$. In the present case however, the low-quality images are different from the high-quality images in many aspects, including physical resolution. All low-quality images, both training and test images, were scanned from art books at a resolution of 1600dpi, without any processing, in order to preserve the original nature of the images as much as possible. Once the training low-quality set was obtained, its images were scaled down to match the dimensions of the high-quality training images. Despite the fact that scaling down low-quality images to match the high-quality ones might seem unusual, it should be noted that the low-quality images were obtained by scanning at a resolution of 1600dpi, while the high-quality images were originally scanned at 300dpi, which explains the above difference in the images' physical resolution.

Having gathered the corresponding training images' pairs of $\{y_h^j\}$ and $\{y_l^j\}$, 59 square patches of dimensions 540×540 , to which we will referred to as *extracts* from now on, were selected. The selection of these extracts was done manually, so that they cover a large spectrum of the images' information.

The training phase was applied as in the original implementation, using the selected extracts as the final training set. Patches were extracted from these images, leading to set $P = \{p_h^k, p_l^k\}$, where k describes the defined locations of the patches to be extracted from the training extracts by operator R_k . Pre-processing is applied to the patches, with low frequencies being removed from p_h^k and features extracted from p_l^k . The extracted features were local and correspond to their high-frequency content. There were two filters used in the feature extraction, a gradient filter $G = [1 \ 0 \ -1]$ and a Laplacian filter $L = [1 \ 0 \ -2 \ 0 \ 1]/2$. Once the patches' sets had been established, Principal Component Analysis (PCA), reducing the patch size from $n = 81$ to a feature dimension of $n_l = 30$ was performed. The dictionary training procedure was applied in 40 iterations of the K-SVD algorithm, with $m = 1,000$ atoms in the dictionary, and allocating $L = 3$ atoms for each representation vector. Dictionary A_h was trained as to match A_l and it was defined so that it minimises mean approximation error

$$A_h = \underset{A_h}{\operatorname{argmin}} \sum_k \left\| p_h^k - A_h q^k \right\|_2^2 = \underset{A_h}{\operatorname{argmin}} \|P_h - A_h Q\|_F^2 \quad (1)$$

where q^k is the sparse representations vectors of p^k .

3.2 Reconstruction Phase

Due to the high physical resolution of the test images, it was impossible to process a complete image as a whole in the EB-SR reconstruction process. Therefore, the input test images were firstly split in rows of full width and length of 600 pixels. These row parts of the image were processed individually in EB-SR and upon completion, they were merged back together to the original image.

Given a test image z_l , this was firstly scaled down by the inverse of the desired scale $s = 3$, so that the application of SR would target the details of the image and not a simple scale-up. Before proceeding to SR application, the image was converted to YCbCr colour space, and EB-SR was applied on the luminance Y channel only. The channels CB and Cr were interpolated and then added to the super-resolved image. Regarding the EB-SR application onto Y channel, the image was interpolated to y_l to match the desired resolution and was then sharpen by spatial non-linear filtering. For every defined location k , patches p_l^k were sparsely coded using dictionary A_l and OMP was applied to obtain the sparse representation vectors $\{q^k\}_k$. Multiplying q^k with A_h , the patches' set p_h^k was recovered. Merge of patches $\{p_h^k\}$ provided the reconstructed high-quality image \hat{y}_h .

In the approach by Zeyde et. al., the resulting boundary effect in the image is simply ignored, by cropping the boundary and using the rest of the image. In the present real-world problem this was not applicable, as if the boundary was ignored at every reconstructed image, which represents a row of the actual test image, this would result in a strong line artefact in the final image. Therefore, an overlapping area of 20 pixels was included in the rows consisting the final image, which was later used to overwrite and remove the boundary artefact in the image.

It is worth noting here another consideration present in the solution towards the real-world problem. Firstly, the input low-quality images, both in terms of training and testing, were scanned images of prints from the early 1930's. Scanned at the high resolution of 1600dpi, their halftone pattern was evident and strong in the images. Encouraged by research performed by Minami [5], the necessary **Inverse Halftoning** was left to be dealt with through the

Table 1: PSNR Measurements

Image	Original	Rotation +5	Rotation +10
Baboon	23.5	20.8	20.4
Barbara	26.8	22.3	21.2
Bridge	25.0	19.8	19.0
Coastguard	27.1	23.0	22.1
Comic	24.0	18.1	17.3
Face	33.5	27.5	26.6
Flowers	28.4	20.8	19.8
Foreman	33.2	23.4	21.8
Lena	33.0	24.3	22.9
Man	27.9	21.7	20.6
Monarch	31.1	21.4	20.3
Pepper	34.0	24.1	22.4
ppt3	25.2	17.3	16.3
Zebra	28.5	18.2	17.2
Average	28.66	21.62	20.56

EB-SR solution itself.

4. RESULTS

Examination of the robustness of the original implementation against a simulated misalignment showed that the algorithm is strongly affected by the misalignment. As one can observe in the PSNR measurements in Table 1, the performance of the algorithm decreases as the level of misalignment, in this case rotational misalignment, increases. Figure 1 demonstrates examples of super-resolved images as these were obtained from the original application of the algorithm and the two simulated misalignment applications.

Applying EB-SR to the real-world problem demonstrated a similar behaviour, with strong artifacts present in the resulting images. Towards a solution to reduce the effect of misalignment in the training sets, **Image Registration** was applied. Unfortunately, PSNR measurements cannot be applied to the results, since PSNR is a full-reference metric. Nevertheless, as it can be observed in Figure 2, the algorithm's performance increases if image registration is applied. The figure displays the extracts from the original high-quality image on the left, the super-resolved image with misalignment in the middle and the super-resolved image after image registration has been applied on the right. It is worth mentioning here, that colour is actually an important component of reclamation of lost art, since it is an essential element in the art of painting. It was, however abstracted away from at this early stage, in order to better focus on the crucial issue of misalignment and allow for better comparison between the misalignment simulation experimentation and the results from the original works.

5. CONCLUSIONS

Examining the results of the evaluation of the EB-SR approach, the importance of misalignment in the training set was considered as crucial. Since the method is an example-based one, erroneous matches of low-quality extracts to mismatched high-quality ones leads to artifacts in the image and lower performance of the algorithm. This is supported by the results, both qualitatively and quantitatively.

These conclusions also apply to the practical application of lost art reclamation. Thus, encouragingly, this novel idea appears amenable



Figure 1: A subset of the resulting SR-ed images, displaying starting from the left column: the original test image, the original application of the algorithm on the image, application with training misalignment by rotation of 5° and application with misalignment with rotation of 10° .



Figure 2: Effect of misalignment in training of SR, for the RLA problem. The image is a detail from painting ‘Madonna’ by P.P. Rubens. The left-most image is the original high-quality image in greyscale, the middle image is the resulting image from SR when there is rotational misalignment of 10° counter-clockwise in the training low-quality set and the right-most image displays the resulting image from SR when image registration is applied on the low-quality training set to correct the misalignment.

to a methodological approach involving existing concepts, innovatively applied. In this context, image registration appears necessary, in order to ensure the elimination of misalignment effects on the algorithm's performance. Future work in this context will address the present and other challenges.

6. REFERENCES

- [1] M. Aharon, M. Elad, and A. Bruckstein. The k-svd: An algorithm for designing of overcomplete dictionaries for sparse representation. In *IEEE Trans. on Signal Processing*, volume 54 of 11, pages 4311–4322, November 2006.
- [2] M. Demetriou, J. Hardeberg, and G. Adelman. Computer-aided reclamation of lost art - a new concept. *Proceedings of Gjøvik Color Imaging Symposium*, pages 67–72, 2011.
- [3] F. Dulberg. *Rubens*. E.A. Seeman.
- [4] M. Elad, R. Rubinstein, and M. Zibulevsky. Efficient implementation of the k-svd algorithm using batch orthogonal matching pursuit. Technical report, CS, Technion, April 2008.
- [5] Y. Minami, S.-i. Azuma, and T. Sugie. An inverse halftoning algorithm based on super-resolution image reconstruction. In *SICE Annual Conference 2010, Proceedings of*, pages 1110–1113, Aug. 2010.
- [6] A. S. Park, M. Park, and M. G. Kang. Super-resolution image reconstruction, a technical overview. *IEEE Signal Processing Magazine*, 20(3):21–36, May 2003.
- [7] J. Yang, J. Wright, Y. Ma, and T. Huang. Image super-resolution as sparse-representation of raw image patches. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [8] R. Zeyde, M. Elad, and M. Protter. On single image scale-up using sparse-representations. In *Curves and Surfaces*, volume 6920 of *Lecture Notes in Computer Science*, pages 711–730. Springer, 2010.

B Program Documentation

This chapter includes information on how to utilise the system's implementation to reproduce the results and/or experiment further. This code will soon be freely available on the author's website, and data will be available upon request.

The system is centralised in the `main` function. All procedures are accessed from there. The individual steps are included in different folders in the same directory as `main`, namely 'SR', 'CC', 'IR', 'IH', 'Tests' for super-resolution, colour correction, image registration, inverse halftoning and testing procedures respectively. Two additional folders, 'Data' and 'Results' are also included, containing the training and test data in the first and the resulting images from each step of the system in the latter. The access to the training and test data is defined in function `main` and can be adjusted.

Implementation of SR as provided by Zeyde *et al.* [21] on the authors' website should be included in the folder 'SR'. Folder 'CC' must include the implementations for the six selected CC methods, as these are provided by Xu and Mulligan [62] on their webpage. Each approach shall be included in a different folder named after the first author of each CC approach. Finally, calculations of image quality metrics as defined by their authors [16, 17] should be included in folder 'Tests', under a sub-folder named 'PSNR', 'SSIM', 'SCIELAB' respectively.

The following functions are necessary for the execution of the program:

`pathsSetup` establishes the environmental setup for the execution.

`imgcut` performs the split of input image in square patches or row patches. The dimensions of the square and the length of the row patches are defined as an argument to the function. This function is applied by default to every test image before carrying out the procedure.

`imgstitch` is used after completion of SR to merge the super-resolved extracts from `imgcut` together and obtain the super-resolved image as a whole.

`training` is the function responsible for the training of SR method.

`load_YCbCr` transforms the input image to the YCbCr colour space.

`scaleup` performs SR on the input image using the approach by Zeyde *et al.* [21].

`CC_Reinhard` executes CC using the colour transfer method by Reinhard *et al.* [14]. The colour space of the transfer can be specified as 'rgb', ' $\alpha\beta$ ' or 'ciacam'.

`CC_Brown` executes CC using the image blending technique by Brown *et al.*.

`CC_Xiao` performs CC using the colour transfer technique proposed by Xiao *et al.* [65].

`CC_Fecker` applies CC via the cumulative histogram mapping approach by Fecker *et al.* [73].

`calc_PeakSNR` calculates the PSNR scores between the input target and source images.

`SSIMcolor` computes the SSIM values for a given target and source colour images' pair.

`batchSCIELAB` computes the S-CIELAB colour difference between two given images.

The next functions are necessary as pre-processing and can also be executed via the `main` by setting to 1 the appropriate flags. These functions are not necessary for every

execution of the system.

`img_scaledown` is responsible for the scale down of the input images by the input scale.

`alignPair` performs image registration using the method by Sheikh *et al.* [53]. The resulting registered images are presented to the user for supervision.

`getDiffMap` calculates the difference map between two images of S-CIELAB difference, with white pixels corresponding to a difference lower than the defined threshold.

C Evaluation Scores

This chapter includes all quantitative evaluation results for purposes of verification and validation of the testing of the ground-truth test data. The measurements are grouped according to the image quality metrics in the following order: PSNR, SSIM, S-CIELAB. Each ground-truth test image is presented in a separate table sorted from best performance to worst, with details regarding the CC method applied and the reference image used.

CC method	Reference Image	PSNR
Brown [69]	The rape of the daughters of Leukippos	17.2355
Xiao [65]	The rape of the daughters of Leukippos	17.1179
Reinhard [14] in RGB	The rape of the daughters of Leukippos	17.0979
Fecker [73]	The rape of the daughters of Leukippos	17.056
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	17.0269
Brown [69]	Liebesgarten	16.9446
Reinhard [14] in CIECAM	Putti	16.8511
Reinhard [14] in RGB	Putti	16.8415
Brown [69]	The Judgement of Paris	16.734
Fecker [73]	Putti	16.6966
Xiao [65]	Putti	16.6092
Reinhard [14] in RGB	The Judgement of Paris	16.0707
Xiao [65]	The Judgement of Paris	16.0359
Reinhard [14] in CIECAM	The Judgement of Paris	16.014
Reinhard [14] in $\alpha\beta$	Putti	15.9515
Fecker [73]	The Judgement of Paris	15.8753
Brown [69]	Putti	15.8406
Reinhard [14] in $\alpha\beta$	The rape of the daughters of Leukippos	15.4097
Reinhard [14] in CIECAM	Liebesgarten	15.2546
Reinhard [14] in RGB	Liebesgarten	15.2369
Reinhard [14] in CIECAM	Madonna	14.9999
Reinhard [14] in RGB	Madonna	14.8852
Fecker [73]	Liebesgarten	14.8029
Brown [69]	Madonna	14.6702
Fecker [73]	Madonna	14.6244
Xiao [65]	Liebesgarten	14.4611
Reinhard [14] in $\alpha\beta$	Madonna	14.2951
Xiao [65]	Madonna	14.0223
Reinhard [14] in $\alpha\beta$	Liebesgarten	13.6397
Reinhard [14] in $\alpha\beta$	The Judgement of Paris	13.337

Table 11: Test image ‘The rape of the daughters of Leukippos’: PSNR scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	PSNR
Reinhard [14] in CIECAM	Liebesgarten	16.4212
Fecker [73]	Liebesgarten	16.3569
Reinhard [14] in RGB	Liebesgarten	16.2992
Reinhard [14] in CIECAM	Madonna	16.134
Reinhard [14] in RGB	Madonna	16.1197
Fecker [73]	Madonna	15.8458
Xiao [65]	Liebesgarten	15.7904
Xiao [65]	Madonna	15.481
Reinhard [14] in $\alpha\beta$	Madonna	15.3082
Reinhard [14] in CIECAM	Putti	15.1852
Reinhard [14] in RGB	Putti	15.1308
Xiao [65]	Putti	15.0608
Fecker [73]	Putti	14.9603
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	14.7607
Xiao [65]	The rape of the daughters of Leukippos	14.4966
Reinhard [14] in RGB	The rape of the daughters of Leukippos	14.4834
Reinhard [14] in $\alpha\beta$	Putti	14.4296
Fecker [73]	The rape of the daughters of Leukippos	14.3909
Reinhard [14] in CIECAM	The Judgement of Paris	14.1941
Xiao [65]	The Judgement of Paris	14.0577
Reinhard [14] in RGB	The Judgement of Paris	14.0142
Reinhard [14] in $\alpha\beta$	Liebesgarten	13.7745
Fecker [73]	The Judgement of Paris	13.702
Reinhard [14] in $\alpha\beta$	The rape of the daughters of Leukippos	13.4487
Reinhard [14] in $\alpha\beta$	The Judgement of Paris	12.2475
Brown [69]	The rape of the daughters of Leukippos	9.6462
Brown [69]	Liebesgarten	8.6708
Brown [69]	The Judgement of Paris	8.4365
Brown [69]	Putti	8.1958
Brown [69]	Madonna	7.3969

Table 12: Test image ‘Liebesgarten’: PSNR scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	PSNR
Reinhard [14] in RGB	Madonna	15.9315
Reinhard [14] in CIECAM	Madonna	15.8105
Fecker [73]	Madonna	15.7642
Reinhard [14] in RGB	Liebesgarten	15.2912
Xiao [65]	Madonna	15.2837
Reinhard [14] in CIECAM	Liebesgarten	15.2551
Fecker [73]	Liebesgarten	15.1515
Reinhard [14] in $\alpha\beta$	Madonna	14.8689
Xiao [65]	Liebesgarten	14.847
Reinhard [14] in CIECAM	Putti	14.6519
Reinhard [14] in RGB	Putti	14.5864
Fecker [73]	Putti	14.5789
Xiao [65]	Putti	14.5213
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	13.854
Reinhard [14] in $\alpha\beta$	Putti	13.6085
Reinhard [14] in RGB	The rape of the daughters of Leukippos	13.5816
Xiao [65]	The rape of the daughters of Leukippos	13.5624
Fecker [73]	The rape of the daughters of Leukippos	13.5372
Reinhard [14] in CIECAM	The Judgement of Paris	12.8825
Reinhard [14] in RGB	The Judgement of Paris	12.6974
Xiao [65]	The Judgement of Paris	12.6971
Reinhard [14] in $\alpha\beta$	The rape of the daughters of Leukippos	12.6634
Fecker [73]	The Judgement of Paris	12.6395
Reinhard [14] in $\alpha\beta$	Liebesgarten	12.0215
Brown [69]	The rape of the daughters of Leukippos	11.6967
Reinhard [14] in $\alpha\beta$	The Judgement of Paris	10.7269
Brown [69]	Liebesgarten	10.471
Brown [69]	Putti	10.3271
Brown [69]	The Judgement of Paris	10.209
Brown [69]	Madonna	9.4265

Table 13: Test image ‘Madonna’: PSNR scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	PSNR
Reinhard [14] in CIECAM	The Judgement of Paris	16.1844
Xiao [65]	The Judgement of Paris	16.1319
Fecker [73]	The Judgement of Paris	16.1245
Reinhard [14] in RGB	The rape of the daughters of Leukippos	16.0885
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	16.0688
Reinhard [14] in RGB	The Judgement of Paris	16.0633
Xiao [65]	The rape of the daughters of Leukippos	16.0484
Fecker [73]	The rape of the daughters of Leukippos	16.0058
Reinhard [14] in RGB	Putti	15.5114
Fecker [73]	Putti	15.3772
Xiao [65]	Putti	15.3439
Reinhard [14] in CIECAM	Putti	15.333
Reinhard [14] in $\alpha\beta$	The rape of the daughters of Leukippos	14.9855
Reinhard [14] in $\alpha\beta$	Putti	14.6466
Reinhard [14] in RGB	Liebesgarten	14.4181
Fecker [73]	Liebesgarten	14.38
Reinhard [14] in CIECAM	Liebesgarten	14.3725
Xiao [65]	Liebesgarten	14.064
Reinhard [14] in $\alpha\beta$	The Judgement of Paris	13.8664
Reinhard [14] in RGB	Madonna	13.7637
Reinhard [14] in $\alpha\beta$	Liebesgarten	13.718
Reinhard [14] in CIECAM	Madonna	13.6467
Fecker [73]	Madonna	13.5716
Xiao [65]	Madonna	13.3464
Reinhard [14] in $\alpha\beta$	Madonna	13.2867
Brown [69]	The rape of the daughters of Leukippos	9.8685
Brown [69]	Liebesgarten	8.99
Brown [69]	Putti	8.9216
Brown [69]	The Judgement of Paris	8.8103
Brown [69]	Madonna	8.1935

Table 14: Test image ‘The Judgement of Paris’: PSNR scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	PSNR
Reinhard [14] in RGB	Putti	14.6861
Reinhard [14] in CIECAM	Putti	14.6676
Fecker [73]	Putti	14.6551
Xiao [65]	Putti	14.6316
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	14.2355
Reinhard [14] in RGB	The rape of the daughters of Leukippos	14.0975
Xiao [65]	The rape of the daughters of Leukippos	14.0515
Fecker [73]	The rape of the daughters of Leukippos	13.9382
Reinhard [14] in RGB	Madonna	13.841
Fecker [73]	Madonna	13.8013
Reinhard [14] in CIECAM	Madonna	13.7328
Reinhard [14] in RGB	Liebesgarten	13.6121
Reinhard [14] in CIECAM	Liebesgarten	13.5459
Reinhard [14] in $\alpha\beta$	Putti	13.507
Fecker [73]	Liebesgarten	13.488
Xiao [65]	Madonna	13.458
Xiao [65]	Liebesgarten	13.3285
Reinhard [14] in CIECAM	The Judgement of Paris	13.2764
Reinhard [14] in RGB	The Judgement of Paris	13.1885
Xiao [65]	The Judgement of Paris	13.1812
Reinhard [14] in $\alpha\beta$	Madonna	13.0874
Fecker [73]	The Judgement of Paris	13.0141
Reinhard [14] in $\alpha\beta$	The rape of the daughters of Leukippos	12.7376
Reinhard [14] in $\alpha\beta$	Liebesgarten	11.119
Reinhard [14] in $\alpha\beta$	The Judgement of Paris	10.8082
Brown [69]	The rape of the daughters of Leukippos	9.5881
Brown [69]	Liebesgarten	8.7208
Brown [69]	Putti	8.647
Brown [69]	The Judgement of Paris	8.5426
Brown [69]	Madonna	7.8858

Table 15: Test image ‘Putti’: PSNR scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	SSIM
Brown [69]	The rape of the daughters of Leukippos	0.3801
Fecker [73]	The rape of the daughters of Leukippos	0.3604
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	0.3532
Xiao [65]	The rape of the daughters of Leukippos	0.3523
Reinhard [14] in RGB	The rape of the daughters of Leukippos	0.3521
Reinhard [14] in $\alpha\beta$	Madonna	0.3481
Reinhard [14] in $\alpha\beta$	The rape of the daughters of Leukippos	0.3473
Reinhard [14] in RGB	Putti	0.347
Brown [69]	Liebesgarten	0.3428
Xiao [65]	Putti	0.3414
Reinhard [14] in RGB	Liebesgarten	0.3369
Brown [69]	Madonna	0.3368
Brown [69]	The Judgement of Paris	0.3354
Fecker [73]	Madonna	0.3341
Reinhard [14] in CIECAM	Liebesgarten	0.3319
Brown [69]	Putti	0.3316
Reinhard [14] in RGB	Madonna	0.3268
Xiao [65]	Liebesgarten	0.3222
Reinhard [14] in $\alpha\beta$	The Judgement of Paris	0.3221
Reinhard [14] in CIECAM	Putti	0.3191
Reinhard [14] in $\alpha\beta$	Liebesgarten	0.3173
Reinhard [14] in RGB	The Judgement of Paris	0.314
Xiao [65]	Madonna	0.3135
Xiao [65]	The Judgement of Paris	0.3121
Fecker [73]	Putti	0.3112
Reinhard [14] in CIECAM	Madonna	0.3101
Reinhard [14] in CIECAM	The Judgement of Paris	0.3096
Reinhard [14] in $\alpha\beta$	Putti	0.3073
Fecker [73]	Liebesgarten	0.3056
Fecker [73]	The Judgement of Paris	0.2874

Table 16: Test image ‘The rape of the daughters of Leukippos’: SSIM scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	SSIM
Xiao [65]	Liebesgarten	0.3797
Reinhard [14] in CIECAM	Liebesgarten	0.3786
Fecker [73]	Liebesgarten	0.3749
Reinhard [14] in RGB	Liebesgarten	0.371
Reinhard [14] in $\lambda\beta$	Liebesgarten	0.3565
Fecker [73]	The rape of the daughters of Leukippos	0.3492
Xiao [65]	Madonna	0.3409
Reinhard [14] in CIECAM	The Judgement of Paris	0.3338
Fecker [73]	The Judgement of Paris	0.333
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	0.3304
Xiao [65]	The Judgement of Paris	0.3293
Reinhard [14] in $\lambda\beta$	The Judgement of Paris	0.3292
Reinhard [14] in CIECAM	Madonna	0.3287
Reinhard [14] in RGB	Madonna	0.3283
Reinhard [14] in RGB	The Judgement of Paris	0.3207
Fecker [73]	Madonna	0.32
Reinhard [14] in RGB	The rape of the daughters of Leukippos	0.3185
Xiao [65]	The rape of the daughters of Leukippos	0.3168
Reinhard [14] in $\lambda\beta$	The rape of the daughters of Leukippos	0.2889
Fecker [73]	Putti	0.2755
Reinhard [14] in $\lambda\beta$	Madonna	0.2728
Reinhard [14] in CIECAM	Putti	0.2724
Reinhard [14] in RGB	Putti	0.2656
Xiao [65]	Putti	0.2563
Brown [69]	The rape of the daughters of Leukippos	0.251
Brown [69]	Madonna	0.2327
Brown [69]	Liebesgarten	0.2256
Brown [69]	The Judgement of Paris	0.2225
Brown [69]	Putti	0.2213
Reinhard [14] in $\lambda\beta$	Putti	0.1767

Table 17: Test image ‘Liebesgarten’: SSIM scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	SSIM
Reinhard [14] in $\lambda\alpha\beta$	Madonna	0.4091
Reinhard [14] in CIECAM	Madonna	0.4069
Reinhard [14] in RGB	Madonna	0.4064
Fecker [73]	Madonna	0.4058
Reinhard [14] in CIECAM	Putti	0.3858
Reinhard [14] in RGB	Putti	0.3798
Xiao [65]	Madonna	0.3745
Xiao [65]	Putti	0.3681
Fecker [73]	Putti	0.3666
Reinhard [14] in CIECAM	Liebesgarten	0.3657
Reinhard [14] in $\lambda\alpha\beta$	The rape of the daughters of Leukippos	0.3652
Reinhard [14] in RGB	Liebesgarten	0.3613
Xiao [65]	Liebesgarten	0.3592
Reinhard [14] in RGB	The rape of the daughters of Leukippos	0.3591
Xiao [65]	The rape of the daughters of Leukippos	0.3529
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	0.35
Reinhard [14] in $\lambda\alpha\beta$	Liebesgarten	0.3423
Reinhard [14] in $\lambda\alpha\beta$	Putti	0.3404
Fecker [73]	The rape of the daughters of Leukippos	0.3343
Reinhard [14] in $\lambda\alpha\beta$	The Judgement of Paris	0.3236
Reinhard [14] in RGB	The Judgement of Paris	0.3192
Xiao [65]	The Judgement of Paris	0.3185
Brown [69]	Putti	0.3106
Reinhard [14] in CIECAM	The Judgement of Paris	0.309
Fecker [73]	Liebesgarten	0.3032
Brown [69]	The rape of the daughters of Leukippos	0.2934
Brown [69]	Madonna	0.2869
Fecker [73]	The Judgement of Paris	0.2615
Brown [69]	Liebesgarten	0.238
Brown [69]	The Judgement of Paris	0.2319

Table 18: Test image ‘Madonna’: SSIM scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	SSIM
Xiao [65]	Liebesgarten	0.4413
Reinhard [14] in CIECAM	Liebesgarten	0.4389
Reinhard [14] in RGB	Liebesgarten	0.4365
Fecker [73]	The rape of the daughters of Leukippos	0.4334
Fecker [73]	The Judgement of Paris	0.4232
Fecker [73]	Liebesgarten	0.421
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	0.4163
Reinhard [14] in CIECAM	The Judgement of Paris	0.4157
Reinhard [14] in $l\alpha\beta$	The Judgement of Paris	0.4155
Xiao [65]	The Judgement of Paris	0.412
Reinhard [14] in RGB	The Judgement of Paris	0.4098
Reinhard [14] in $l\alpha\beta$	Liebesgarten	0.4051
Reinhard [14] in RGB	The rape of the daughters of Leukippos	0.4042
Xiao [65]	The rape of the daughters of Leukippos	0.3957
Reinhard [14] in RGB	Madonna	0.3737
Xiao [65]	Madonna	0.3715
Reinhard [14] in CIECAM	Madonna	0.3664
Reinhard [14] in $l\alpha\beta$	The rape of the daughters of Leukippos	0.3589
Fecker [73]	Madonna	0.3539
Brown [69]	The rape of the daughters of Leukippos	0.3467
Brown [69]	Putti	0.3402
Reinhard [14] in CIECAM	Putti	0.3375
Brown [69]	Madonna	0.3371
Fecker [73]	Putti	0.3329
Reinhard [14] in RGB	Putti	0.3216
Brown [69]	Liebesgarten	0.3174
Xiao [65]	Putti	0.3162
Brown [69]	The Judgement of Paris	0.3139
Reinhard [14] in $l\alpha\beta$	Madonna	0.2914
Reinhard [14] in $l\alpha\beta$	Putti	0.2145

Table 19: Test image ‘The Judgement of Paris’: SSIM scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	SSIM
Reinhard [14] in RGB	Putti	0.3784
Reinhard [14] in CIECAM	Putti	0.3765
Xiao [65]	Putti	0.3689
Reinhard [14] in CIECAM	Madonna	0.3563
Reinhard [14] in RGB	Madonna	0.3557
Fecker [73]	Madonna	0.3453
Fecker [73]	Putti	0.3439
Xiao [65]	Madonna	0.3284
Reinhard [14] in $\alpha\beta$	Madonna	0.3246
Reinhard [14] in RGB	The rape of the daughters of Leukippos	0.32
Xiao [65]	The rape of the daughters of Leukippos	0.3181
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	0.3156
Reinhard [14] in $\alpha\beta$	Putti	0.3115
Reinhard [14] in $\alpha\beta$	The rape of the daughters of Leukippos	0.3054
Reinhard [14] in CIECAM	Liebesgarten	0.3003
Reinhard [14] in RGB	Liebesgarten	0.2939
Xiao [65]	Liebesgarten	0.2829
Fecker [73]	The rape of the daughters of Leukippos	0.2787
Xiao [65]	The Judgement of Paris	0.2623
Reinhard [14] in RGB	The Judgement of Paris	0.2615
Brown [69]	Putti	0.2581
Reinhard [14] in $\alpha\beta$	Liebesgarten	0.2568
Reinhard [14] in CIECAM	The Judgement of Paris	0.2544
Reinhard [14] in $\alpha\beta$	The Judgement of Paris	0.254
Fecker [73]	Liebesgarten	0.2527
Brown [69]	Madonna	0.2264
Fecker [73]	The Judgement of Paris	0.2183
Brown [69]	The rape of the daughters of Leukippos	0.1889
Brown [69]	Liebesgarten	0.1405
Brown [69]	The Judgement of Paris	0.1375

Table 20: Test image ‘Putti’: SSIM scores for the structural coherence between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	S-CIELAB
Reinhard [14] in RGB	The rape of the daughters of Leukippos	16.94633
Fecker [73]	The rape of the daughters of Leukippos	17.01028
Xiao [65]	The rape of the daughters of Leukippos	17.039
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	17.16043
Fecker [73]	The Judgement of Paris	18.29091
Reinhard [14] in CIECAM	The Judgement of Paris	18.53061
Reinhard [14] in RGB	The Judgement of Paris	18.98055
Xiao [65]	The Judgement of Paris	19.3154
Fecker [73]	Putti	19.52252
Reinhard [14] in CIECAM	Putti	19.53027
Brown [69]	The Judgement of Paris	19.53167
Reinhard [14] in RGB	Putti	19.90329
Xiao [65]	Putti	20.85254
Reinhard [14] in RGB	Liebesgarten	21.81857
Reinhard [14] in CIECAM	Liebesgarten	22.07612
Fecker [73]	Liebesgarten	22.10175
Reinhard [14] in CIECAM	Madonna	23.17844
Reinhard [14] in $\lambda\alpha\beta$	The rape of the daughters of Leukippos	23.67672
Fecker [73]	Madonna	23.8792
Reinhard [14] in RGB	Madonna	23.89826
Reinhard [14] in $\lambda\alpha\beta$	The Judgement of Paris	26.16369
Reinhard [14] in $\lambda\alpha\beta$	Liebesgarten	26.4703
Xiao [65]	Liebesgarten	27.31326
Reinhard [14] in $\lambda\alpha\beta$	Putti	28.42991
Reinhard [14] in $\lambda\alpha\beta$	Madonna	28.48176
Xiao [65]	Madonna	29.55819
Brown [69]	Madonna	31.02939
Brown [69]	Putti	32.88975
Brown [69]	Liebesgarten	35.7924
Brown [69]	The rape of the daughters of Leukippos	37.34785

Table 21: Test image ‘The rape of the daughters of Leukippos’: S-CIELAB scores for the colour difference between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	S-CIELAB
Brown [69]	The Judgement of Paris	15.43763
Reinhard [14] in RGB	Liebesgarten	15.95823
Fecker [73]	Liebesgarten	15.99388
Reinhard [14] in CIECAM	Liebesgarten	16.28374
Reinhard [14] in RGB	Madonna	19.02853
Reinhard [14] in CIECAM	Madonna	19.4754
Xiao [65]	Liebesgarten	19.73947
Reinhard [14] in CIECAM	The Judgement of Paris	19.91066
Xiao [65]	The Judgement of Paris	20.36514
Fecker [73]	Madonna	20.62259
Reinhard [14] in RGB	The Judgement of Paris	20.67187
Fecker [73]	The Judgement of Paris	21.03328
Reinhard [14] in RGB	The rape of the daughters of Leukippos	21.25687
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	21.30221
Brown [69]	Madonna	21.86493
Xiao [65]	The rape of the daughters of Leukippos	22.00864
Reinhard [14] in $\alpha\beta$	Liebesgarten	22.32485
Fecker [73]	The rape of the daughters of Leukippos	22.61512
Brown [69]	Putti	22.75221
Xiao [65]	Madonna	23.20874
Reinhard [14] in RGB	Putti	23.54283
Xiao [65]	Putti	24.25135
Brown [69]	Liebesgarten	24.64488
Reinhard [14] in $\alpha\beta$	Madonna	25.34362
Reinhard [14] in CIECAM	Putti	25.48326
Brown [69]	The rape of the daughters of Leukippos	25.72361
Reinhard [14] in $\alpha\beta$	The Judgement of Paris	26.35416
Fecker [73]	Putti	26.82046
Reinhard [14] in $\alpha\beta$	The rape of the daughters of Leukippos	32.28415
Reinhard [14] in $\alpha\beta$	Putti	36.64407

Table 22: Test image ‘Liebesgarten’: S-CIELAB scores for the colour difference between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	S-CIELAB
Reinhard [14] in RGB	Madonna	20.78423
Brown [69]	Madonna	21.27462
Brown [69]	Putti	21.2771
Fecker [73]	Madonna	21.87578
Reinhard [14] in CIECAM	Madonna	22.216
Fecker [73]	Liebesgarten	23.13826
Reinhard [14] in RGB	Liebesgarten	23.3081
Reinhard [14] in CIECAM	Liebesgarten	23.32694
Brown [69]	The rape of the daughters of Leukippos	23.33348
Brown [69]	Liebesgarten	23.68315
Xiao [65]	Madonna	24.60183
Reinhard [14] in RGB	Putti	25.77445
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	25.78278
Xiao [65]	Liebesgarten	25.82488
Reinhard [14] in RGB	The rape of the daughters of Leukippos	26.30583
Fecker [73]	The rape of the daughters of Leukippos	26.31856
Reinhard [14] in $\lambda\alpha\beta$	Madonna	26.56547
Xiao [65]	Putti	27.13883
Xiao [65]	The rape of the daughters of Leukippos	27.23304
Reinhard [14] in CIECAM	Putti	27.26577
Fecker [73]	Putti	27.73048
Reinhard [14] in CIECAM	The Judgement of Paris	27.85441
Brown [69]	The Judgement of Paris	28.20637
Fecker [73]	The Judgement of Paris	28.31664
Reinhard [14] in RGB	The Judgement of Paris	28.63891
Xiao [65]	The Judgement of Paris	29.40506
Reinhard [14] in $\lambda\alpha\beta$	Liebesgarten	29.96567
Reinhard [14] in $\lambda\alpha\beta$	The rape of the daughters of Leukippos	32.13543
Reinhard [14] in $\lambda\alpha\beta$	The Judgement of Paris	34.00861
Reinhard [14] in $\lambda\alpha\beta$	Putti	37.66694

Table 23: Test image ‘Madonna’: S-CIELAB scores for the colour difference between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	S-CIELAB
Reinhard [14] in RGB	The Judgement of Paris	16.86876
Xiao [65]	The Judgement of Paris	17.00878
Brown [69]	The Judgement of Paris	17.21829
Reinhard [14] in RGB	The rape of the daughters of Leukippos	17.25156
Fecker [73]	The Judgement of Paris	17.29789
Reinhard [14] in CIECAM	The Judgement of Paris	17.54533
Xiao [65]	The rape of the daughters of Leukippos	18.00943
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	18.69222
Fecker [73]	The rape of the daughters of Leukippos	18.83801
Reinhard [14] in RGB	Putti	20.6368
Xiao [65]	Putti	21.2709
Reinhard [14] in RGB	Liebesgarten	21.46166
Fecker [73]	Liebesgarten	21.96408
Reinhard [14] in CIECAM	Liebesgarten	22.18017
Reinhard [14] in CIECAM	Putti	22.74743
Fecker [73]	Putti	23.32192
Reinhard [14] in $l\alpha\beta$	The Judgement of Paris	23.623
Reinhard [14] in $l\alpha\beta$	Liebesgarten	23.74451
Reinhard [14] in RGB	Madonna	24.46477
Xiao [65]	Liebesgarten	24.48991
Reinhard [14] in CIECAM	Madonna	25.34758
Fecker [73]	Madonna	25.8368
Brown [69]	Madonna	27.60301
Xiao [65]	Madonna	28.34111
Brown [69]	Putti	28.69085
Reinhard [14] in $l\alpha\beta$	The rape of the daughters of Leukippos	29.59106
Brown [69]	Liebesgarten	30.35176
Reinhard [14] in $l\alpha\beta$	Madonna	30.67331
Brown [69]	The rape of the daughters of Leukippos	32.40809
Reinhard [14] in $l\alpha\beta$	Putti	34.27688

Table 24: Test image ‘The Judgement of Paris’: S-CIELAB scores for the colour difference between the resulting super-resolved colour corrected images against the ground-truth high quality image

CC method	Reference Image	S-CIELAB
Reinhard [14] in RGB	Putti	23.50144
Fecker [73]	Putti	23.71506
Reinhard [14] in CIECAM	Putti	24.45673
Xiao [65]	Putti	24.51938
Reinhard [14] in CIECAM	The rape of the daughters of Leukippos	24.72682
Fecker [73]	The rape of the daughters of Leukippos	24.74931
Reinhard [14] in RGB	The rape of the daughters of Leukippos	25.42769
Xiao [65]	The rape of the daughters of Leukippos	26.17275
Brown [69]	The Judgement of Paris	26.46324
Fecker [73]	The Judgement of Paris	27.19918
Fecker [73]	Madonna	27.36245
Reinhard [14] in RGB	Madonna	27.41521
Reinhard [14] in CIECAM	The Judgement of Paris	27.72311
Fecker [73]	Liebesgarten	27.87417
Reinhard [14] in RGB	The Judgement of Paris	28.26109
Reinhard [14] in CIECAM	Madonna	28.50174
Xiao [65]	The Judgement of Paris	28.62188
Reinhard [14] in RGB	Liebesgarten	28.93468
Reinhard [14] in CIECAM	Liebesgarten	29.09675
Xiao [65]	Madonna	31.2578
Brown [69]	Madonna	31.73253
Reinhard [14] in $\lambda\alpha\beta$	The rape of the daughters of Leukippos	31.73528
Xiao [65]	Liebesgarten	31.9035
Reinhard [14] in $\lambda\alpha\beta$	Madonna	32.26575
Brown [69]	Putti	33.12741
Reinhard [14] in $\lambda\alpha\beta$	Putti	34.70046
Reinhard [14] in $\lambda\alpha\beta$	The Judgement of Paris	36.2189
Reinhard [14] in $\lambda\alpha\beta$	Liebesgarten	36.2467
Brown [69]	Liebesgarten	36.87095
Brown [69]	The rape of the daughters of Leukippos	37.4875

Table 25: Test image ‘Putti’: S-CIELAB scores for the colour difference between the resulting super-resolved colour corrected images against the ground-truth high quality image