

Rochester Institute of Technology

RIT Digital Institutional Repository

Theses

7-27-2017

Scientific Imaging of Cultural Heritage for Computer Graphics Rendering

Brittany D. Cox
bdh2631@rit.edu

Follow this and additional works at: <https://repository.rit.edu/theses>

Recommended Citation

Cox, Brittany D., "Scientific Imaging of Cultural Heritage for Computer Graphics Rendering" (2017). Thesis. Rochester Institute of Technology. Accessed from

This Dissertation is brought to you for free and open access by the RIT Libraries. For more information, please contact repository@rit.edu.

Scientific Imaging of Cultural Heritage for Computer Graphics Rendering

Brittany D. Cox

Rochester Institute of Technology



College of Science
Program of Color Science
Munsell Color Science Laboratory

Ph. D. Dissertation

A Dissertation Submitted in
Partial Fulfillment of the Requirements for the
Degree of Doctorate of Philosophy in Color Science

Scientific Imaging of Cultural Heritage for Computer Graphics Rendering

Brittany D. Cox

July 27th, 2017

Signature of Author _____

Accepted by _____
Graduate Program Director Date

Brittany D. Cox

Scientific Imaging of Cultural Heritage for Computer Graphics Rendering

Ph. D. Dissertation, July 27th, 2017

Advisor: Roy S. Berns

Committee Chair: Tina Lent

Committee Members: Mark D. Fairchild and Joe Geigel

Rochester Institute of Technology

College of Science

1 Lomb Memorial Dr

Rochester New York, 14623

COLLEGE OF SCIENCE
ROCHESTER INSTITUTE OF TECHNOLOGY
ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

Ph.D. DEGREE DISSERTATION

The Ph.D. Degree Dissertation of Brittany D. Cox
has been examined and approved by the
committee as satisfactory for the
dissertation required for the
Ph.D. degree in Color Science

Tina Lent, Chair

Roy S. Berns, Advisor

Mark D. Fairchild

Joe Geigel

Date

Declaration

This dissertation is a presentation of my original research work. Wherever contributions of others are involved, every effort is made to indicate this clearly, with due reference to the literature, and acknowledgment of collaborative research and discussions.

I hereby grant permission to Wallace Center Library of R.I.T. to reproduce my thesis in whole or in part. Any reproduction will not be for commercial use or profit.

Rochester, New York

Brittany D. Cox

Abstract

Viewing paintings and drawings is a dynamic experience where observers can move about the object discerning surface structure such as brushwork and gloss in addition to color and composition. Conventional studio photography reduces the experience to a single viewpoint where the photographer uses experiential learning and aesthetic judgments. This dissertation had several research goals: improving the quality of estimated surface normal maps using only four studio strobe lights, evaluating the viability of a virtual viewing experience using computer graphics rendering software, understanding how gloss and texture are perceived through psychophysical experiments, and using this understanding to inform image rendering using computer graphics or conventional photography. This research improved the four strobe imaging technique by replacing the previously used polarizers on the camera and strobes with a mathematical thresholding technique. This reduced the number of images required to estimate surface normal and diffuse color by half while creating surface normal and color maps that could be used in a rendering software to generate virtual fine art images. With the need to improve fine art rendering practices in mind, a series of psychophysical experiments were performed. Observers were asked to assess the amount of perceived gloss and texture in sets of images that used different lighting configurations of the same samples. It was determined that observers have difficulty separating gloss and texture while performing discrimination tasks. From the results of the observations, it is assumed that texture influences gloss perception largely at raking lighting angles and gloss influences texture perception primarily at near normal lighting angles. Further exploring the perceived amount of gloss and texture in fine art images with different lighting geometries, observers were asked to rate the aesthetic appeal of these images. On average, observers preferred images with small amounts of both perceived gloss and texture. This research resulted in a successful imaging and rendering protocol that used a limited amount of studio equipment without sacrificing the capture of surface

texture information as well as additional knowledge regarding the desirable amount of perceived gloss and texture that can be used in making lighting and rendering choices that are pleasing to observers while still representing the physical features of a work of art.

Acknowledgements

I would like to acknowledge the Andrew W. Mellon Foundation for providing financial support for this research.

This research would not have been possible without the support and guidance of my advisor, Dr. Roy Berns, and committee, Dr. Mark Fairchild, Dr. Joe Geigel, and Dr. Tina Lent.

I would like to thank all of the students, faculty, and staff that participated in each of my experiments. I owe much of my success to the MCSL faculty, Dr. Susan Franand, Dr. Mike Murdoch, and Dr. Dave Wyble for providing feedback and engaging in many brainstorming sessions. And especially Dr. Dave Wyble for not only providing academic support, but also helping me navigate and adjust to such a strenuous environment that graduate school can be.

I would also like to thank Val Hemink for being a wonderful lab mom and helping with all the necessary paperwork and for being an excellent cheerleader. Additionally, I would like to thank Angela Owen for all of the proof reading and whimsy she provided throughout my academic career in both under grad and grad school.

Finally, I would like to thank my husband, Nathanael, for providing endless support, understanding, and distractions while we were working on our PhDs together. I could not have done this without you and I am so thankful that we were able to go through this adventure together.

Contents

List of figures	xii
List of tables	xxii
1 Introduction	1
1.1 Motivation	1
1.2 Approach	1
1.3 Dissertation Structure	2
2 Introduction to Imaging and Rendering Fine Art	4
2.1 Imaging Fine Art: Color	4
2.1.1 RGB dSLR	5
2.1.2 Dual-RGB	6
2.2 Imaging Fine Art: Surface Macrostruture	10
2.2.1 Photometric Stereo	10
2.2.2 Reflectance Transformation Imaging	15
2.3 Imaging Fine Art: Surface Microstruture	16
2.3.1 Gloss and Material Appearance	17
2.3.2 Measuring BRDF	18
2.3.3 Modeling BRDF	19
2.3.4 Artist Material Database	20
2.4 Rendering Fine Art	21
2.4.1 Rendering Workflow	22
2.4.2 RTIviewer, ArtViewer, and Commercial Graphics Rendering Software	22
2.5 Related Work	24
3 Deriving Surface Normal Maps From Traditional Photographs	28

3.1	Building on an Existing Foundation	28
3.2	Threshold Method Enhanced Four Light Imaging	29
3.3	Rendering Images using Threshold Data	30
4	Lighting for Scientific Imaging	33
4.1	Background	33
4.2	The Addition of Honeycomb Grids	35
4.3	Honeycombs and Normal Maps	38
4.3.1	Normal Map Difference Images	39
4.4	Further Investigations of Honeycomb Use	42
4.4.1	Rendered Images and Color Differences	43
4.5	Conclusion	48
5	Combining Color and Surface Normal Maps in a Rendering Software	50
5.1	Introduction to Maya and mental ray	50
5.2	Rendering with Maya	51
5.3	Rendering Spatially Varying Gloss	57
5.3.1	Approach	57
5.3.2	Results	58
5.3.3	Conclusion	59
6	Observer Preferences and Digital Fine Art Reproduction	61
6.1	Experiment 3: Creating Stimuli	61
6.2	Experiment 3: Procedure	64
6.2.1	Texture Perception	64
6.2.2	Gloss Perception	65
6.2.3	Observer Viewing Preference	66
6.3	Experiment 3: Results	70
6.3.1	Gloss Perception	70
6.3.2	Texture Perception	79
6.3.3	Observer Preferences	91
6.4	Conclusion	99
7	Observer Preferences and Modeling	100
7.0.1	Gloss Perception Modeling	104

7.0.2	Texture Perception Modeling	112
7.0.3	Modeling Preference	119
7.1	Conclusion	126
8	Conclusions	127
8.0.1	Future Work	128
A	Appendix A: 4LIS and Maya: A Case Study	130
A.1	MoMA Imaging Setup	131
A.2	4LIS Imaging Setup	132
A.3	Post Processing	133
A.4	Rendered Images and Results	133
B	Appendix B: Relating Ward Model to Maya's Anisotropic Shader	139
B.1	Experiment 1: Creating Stimuli	140
B.2	Experiment 1: Procedure	142
B.3	Experiment 2: Procedure	143
B.4	Experiment 2: Results	145
B.5	Conclusions	149
C	Appendix C: Supplementary Materials	151
C.1	Frequency Matrices	151
C.2	Probability Matrices	156
C.3	Texture, Gloss, and Preference Surveys	161
	Bibliography	162

List of Figures

2.1	Camera sensitivities when using a blue-green filter (solid lines) and a yellow filter (dashed lines) [Chen and Berns, 2012c].	7
2.2	Five sampling wavelength regions from using the Dual-RGB method of imaging [Chen and Berns, 2012c].	8
2.3	Five light setup for the technique described by Rushmeier, et al. [Rushmeier et al., 1997].	12
2.4	Flow chart showing the steps in the hierarchal sorting technique to remove highlights and shadows [Sun et al., 2007].	13
2.5	Comparison of four different methods for calculated albedo a) three lights, b) Rushmeier's five light technique, c) all six lights, and d) six lights without shadows and specularities and a comparison of the same four methods for rendered images using e) three lights, f) Rushmeier's five light technique, g) all six lights, and h) six lights without shadows and specularities [Sun et al., 2007].	14
2.6	A diagram of the four light with polarizers technique shows the relative locations of a painting, camera and four different lights [Chen and Berns, 2012a]. . . .	15
2.7	A comparison of a traditional photograph (left), specular enhanced image (middle), and a normal map (right) of a stone carving [Duffy, 2013].	16
2.8	A diagram showing (from right to left) specular gloss, sheen, contrast gloss, absence of bloom, and distinctness-of-image gloss in terms of incoming and outgoing light [Hunter and Harold, 1987].	17
2.9	A diagram of a gonireflectometer showing a reflectance detector and a moving light source [Ward, 1992].	18
2.10	A diagram of the linear light source set up [Gardner et al., 2003].	19
2.11	An example grid from the Artist Material Database (AMD) showing unvarnished, matte acrylic varnish, gloss acrylic varnished, gamvar applied with two different techniques [Ashbaugh et al., 2009].	21

2.12	Generalized workflow for rendering fine art requires surface normal information, diffuse color, BRDF of microstructure, and known lighting and viewing conditions.	22
2.13	A screenshot of the ArtViewer software with a Van Gogh style painting loaded.	23
2.14	Diagram of the setup used by Y. Chen to record information about fine art under varying illumination angles [Chen, 2008].	25
2.15	Examples of the Phong and Torrance-Spaarow models' ability to render samples [Chen, 2008].	26
2.16	Example images using the eight different systems to image the same painting [Chen, 2013].	27
3.1	Diffuse image of the unvarnished (top) and varnished (bottom) paintings from cross-polarized (left) and traditional images (right). Diffuse color image of varnished painting from cross-polarized (bottom left) and traditional images (bottom right) [Cox and Berns, 2015].	30
3.2	Normal maps of the unvarnished painting. Left image used cross-polarized data to generate a normal map using four lights. The center image used thresholded data from three or four lights when appropriate. The right image is a difference map [Cox and Berns, 2015].	31
3.3	Normal maps of the varnished painting. Left image used cross-polarized data to generate a normal map using four lights. The center image used thresholded data from either three or four lights when appropriate. The right image is a difference map [Cox and Berns, 2015].	31
3.4	Rendered images of the unvarnished painting from cross-polarized normal and diffuse maps (far left), from traditional/thresholded normal and diffuse color maps (left-middle), the varnished painting from cross-polarized normal and diffuse maps (right-middle), and from traditional/thresholded normal and diffuse color maps (right) [Cox and Berns, 2015].	32
4.1	Left: Lighting setup used by the Department of Imaging and Visual Resources at MoMA using two strobes (highlighted with red circles) directed toward the ceiling from above the paintings. Right: An example of fine art lighting used by the Museum of Modern Art in Queens, New York on Pablo Picasso's Vase of Flowers.	34

4.2	A section of an oil painting imaged with two different types of lighting normal(left) and raking (right) [Webber, 2008].	34
4.3	a) Aiming strobes to achieve illumination centered in the field of view of the camera requires an off-axis alignment shown by the solid black lines for the left and right strobes. Dotted-gray lines show what would be considered on-axis lighting. b) An example of a Broncolor xenon strobe with P70 reflector attached [Broncolor, Bron Elektronik AG, 2015].	36
4.4	a) Normal map with holes, or dark areas where not enough information was available to accurately calculate the surface normal.	36
4.5	Honeycomb grids that control the beam spread of the source to ten, twenty, and thirty degrees [B & H Foto & Electronics Corp, 2015].	37
4.6	Falloff images of a white background using (from top left to bottom right) no grid, a 30-degree, 20-degree, 10-degree, and 5-degree grid.	38
4.7	Scene used to compare 4LIS imaging with different honeycomb grids. The scene includes all necessary elements (cue ball, color chart, and painting) to facilitate efficient imaging for all honeycombs.	40
4.8	Normal maps resulting from using different grids. From top left to bottom left (clockwise) 20 degrees, 10 degrees, 5 degrees, 4LIS without a grid.	40
4.9	Close-up section from the bottom left of each normal map using no grid and the 5 to 20 degree grids.	41
4.10	Difference images, comparing normal maps to normal maps generated from the use of other grids.	41
4.11	Van Gogh style painting imaged using both the twenty- and thirty- degree grids.	42
4.12	Normal map of a Van Gogh style painting created from images taken with a thirty-degree honeycomb grid.	43
4.13	Normal map of a Van Gogh style painting created from images taken with a twenty-degree honeycomb grid.	44
4.14	Normal map of a Van Gogh style painting cropped to a small section of a dark glossy region (pine tree) using a 20 degree honeycomb (left) and a 30 degree honeycomb (right).	45
4.15	20 degree (left) and thirty degree (right) rendered sections of the dark pine tree in a Van Gogh Style painting, with areas containing "holes" circled in red.	46

4.16	Color difference plots showing the difference between measured data (circles) and calculated data (tip of each arrow) for the twenty-degree grid.	46
4.17	Color difference plots showing the difference between measured data (circles) and calculated data (tip of each arrow) for the thirty-degree grid.	47
4.18	Color differences shown in an sRGB rendered image showing the measured data in the center surround by data calculated from images taken (from top left and clockwise around) at short distance with a 20-degree grid, 30-degree grid, then a long distance with a 20-degree grid and 30-degree grid.	48
5.1	Diffuse color (left) and normal (right) images of a van Gogh style painting resulting from the simplified four-light method.	52
5.2	Maya workspace showing a painting illuminated with two different lights, a directional light from above and a diffuse fill light from below.	52
5.3	Rendered images of the painting using Maya to simulate directional lighting from above and fill lighting from below with its matte, existing varnish (left) and simulated picture varnish (right).	54
5.4	Rendered image of a close up section of a painting with directional lighting coming from top left (top left), top center (top right), and top right (bottom) generated using Maya and mental ray.	55
5.5	Rendered image of a painting with simulated raking light (left) and specular enhancement (right).	56
5.6	A section of furniture attributed to Joseph Baumhauer belonging to the Sculpture and Decorative Arts Department of the J. Paul Getty Museum Collection. (left) The diffuse color map (middle) and normal map (right) of a flat section of the piece indicated by the red box in the left image.	57
5.7	Masks to separate the color map into images of three different material types: black lacquer (left), high gloss gold (middle), and matte gold and red accents (right).	58
5.8	Images of the rendered panel with lighting from four different directions showing how the gold areas change differently from the black lacquered areas. Three different shaders were used to render the different materials in the top half of the renderings. The renderings in the bottom half used only one.	59

6.1	Schematic of lighting angle used to image each sample. Strobes outlined in red (angles 82 and 88) fired at the same time as a diffuse fill light from the opposite side of the painting.	62
6.2	Images of each reference stimuli of (from top left to bottom right) an unvarnished gray patch of acrylic paint, an unvarnished acrylic painting, a varnished acrylic painting, a matte varnished gray patch of oil paint, matte print, woven tapestry, and two varnished oil paintings.	63
6.3	Flowers sample showing each image used in the experiment starting with the largest lighting angle in the top left and the smallest angle in the bottom right.	64
6.4	Example image of the experiment interface showing three images of an acrylic painting, two anchor images and one test image.	66
6.5	Example image of the experiment interface showing three images of a varnished acrylic painting, two anchor images and one test image.	67
6.6	A section of the gallery shown to observers using modeling lights located above the paintings and aimed towards the vertical center of the paintings on display.	68
6.7	Example screenshot of the third task where a reference and test image were shown side by side and observers were asked to rate the right image's appearance in comparison to the left image.	69
6.8	Average gloss perception rating for each of the five artwork samples: Tulip, OilTrees, Gray Patch 1, Flowers, and DarkOil.	71
6.9	Average gloss perception rating for Flowers with standard deviation from the mean as error bars.	72
6.10	Average gloss perception rating for Gray Patch 1 with standard deviation from the mean as error bars.	72
6.11	Average gloss perception rating for DarkOil with standard deviation from the mean as error bars.	73
6.12	Average gloss perception rating for OilTrees with standard deviation from the mean as error bars.	73
6.13	Average gloss perception rating for Tulip with standard deviation from the mean as error bars.	74
6.14	Average gloss perception rating for Flowers with standard deviation from the mean as error bars.	74

6.15	Average gloss perception rating for Gray Patch 1 with standard deviation from the mean as error bars.	75
6.16	Average gloss perception rating for DarkOil with standard deviation from the mean as error bars.	75
6.17	Average gloss perception rating for OilTrees with standard deviation from the mean as error bars.	76
6.18	Average gloss perception rating for Tulip with standard deviation from the mean as error bars.	76
6.19	Images of three of the five samples that have roughly the same gloss rating.	78
6.20	Images of two of the five gloss samples that have roughly the same gloss rating.	79
6.21	Average texture perception rating for all 8 artwork samples: OilTrees, Gray Patch 1, Tulip, DarkOil, Print, Textile, Gray Patch 2, and Flowers.	80
6.22	Average texture perception rating for 5 of the 8 artwork samples: OilTrees, DarkOil, Tulip, Gray Patch 1, and Flowers.	80
6.23	Average texture perception rating for 3 of the 8 artwork samples: Gray Patch 2, Textile, and Print.	81
6.24	Average texture perception rating for Gray Patch 1 sample.	82
6.25	Average texture perception rating for DarkOil sample.	83
6.26	Average texture perception rating for OilTrees sample.	83
6.27	Average texture perception rating for Tulip sample.	84
6.28	Average texture perception rating for Flowers sample.	84
6.29	Close up sections of the images of the DarkOil painting and the Tulip painting taken with the light source at 67 degrees that emphasize the difference between gloss depicted in the Dark Oil and Tulip paintings.	85
6.30	Images of 3 of the 5 samples that have roughly the same texture rating. . . .	86
6.31	Images of 2 of the 5 samples that have roughly the same texture rating. . . .	87
6.32	Average texture perception rating for Gray Patch 1 sample using data from all samples.	88
6.33	Average texture perception rating for DarkOil sample using data from all samples.	88
6.34	Average texture perception rating for OilTrees sample using data from all samples.	89
6.35	Average texture perception rating for Tulip sample using data from all samples.	89
6.36	Average texture perception rating for Flowers sample using data from all samples.	90

6.37	Image preferred (left) next to the image that best represents the original in a gallery setting (right).	92
6.38	Image preferred (left) next to the image that best represent the original in a gallery setting (right). A close up section of both images is shown underneath.	92
6.39	Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.	93
6.40	Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.	94
6.41	Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.	95
6.42	Gray Patch 1 sample showing each image used in the experiment starting with the largest lighting angle in the top left and the smallest angle in the bottom right.	96
6.43	Image preferred (left) next to the image that best represent the original in a gallery setting (right). A close up section of both images is shown underneath.	96
6.44	Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.	97
6.45	Image preferred (left) next to the image that best represents the original in a gallery setting (right).	97
6.46	Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.	98
6.47	Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.	98
7.1	Two images with the same mean luminance but different skewness values, -1.34 (left) and 2.40 (right) and different amounts of perceived gloss [Motoyoshi et al., 2007].	101
7.2	A diagram showing the pixel of interest's nearest neighbors as defined by [0 1], [1 1], and [1 0].	102
7.3	A Gabor filtered image of the Tulip sample using a wavelength of 8 and 0 degrees.	103
7.4	Stepwise fit modeling the gloss of the Tulip sample (blue) and the observer rated data (black).	105
7.5	Perceived gloss rating (blue) and predicted rating (black) for the OilTrees sample.	106

7.6	Perceived gloss rating (blue) and predicted rating (black) for the Gray Patch 1 sample.	106
7.7	Perceived gloss rating (blue) and predicted rating (black) for the DarkOil sample.	107
7.8	Perceived gloss rating (blue) and predicted rating (black) for the Flowers sample.	107
7.9	Stepwise fit modeling the gloss of the Tulip sample (blue) and the observer rated data (black) using all samples to produce the model.	109
7.10	Stepwise fit modeling the gloss of the OilTrees sample (blue) and the observer rated data (black) using all samples to produce the model.	109
7.11	Stepwise fit modeling the gloss of the Gray Patch 1 sample (blue) and the observer rated data (black) using all samples to produce the model.	110
7.12	Stepwise fit modeling the gloss of the DarkOil sample (blue) and the observer rated data (black) using all samples to produce the model.	110
7.13	Stepwise fit modeling the gloss of the Flowers sample (blue) and the observer rated data (black) using all samples to produce the model.	111
7.14	Perceived texture rating (blue) and predicted rating (black) for the Tulip sample.	113
7.15	Perceived texture rating (blue) and predicted rating (black) for the Gray Patch 1 sample.	113
7.16	Perceived texture rating (blue) and predicted rating (black) for the DarkOil sample.	114
7.17	Perceived texture rating (blue) and predicted rating (black) for the Flowers sample.	114
7.18	Perceived texture rating (blue) and predicted rating (black) for the Tulip sample using all samples as input.	116
7.19	Perceived texture rating (blue) and predicted rating (black) for the Gray Patch 1 sample using all samples as input.	116
7.20	Perceived texture rating (blue) and predicted rating (black) for the DarkOil sample using all samples as input.	117
7.21	Perceived texture rating (blue) and predicted rating (black) for the OilTrees sample using all samples as input.	117
7.22	Perceived texture rating (blue) and predicted rating (black) for the Flowers sample using all samples as input.	118
7.23	Preference rating (blue) and predicted rating (black) for the Tulip sample. . .	120
7.24	Preference rating (blue) and predicted rating (black) for the OilTrees sample.	120

7.25	Preference rating (blue) and predicted rating (black) for the Gray Patch 1 sample.	121
7.26	Preference rating (blue) and predicted rating (black) for the DarkOil sample. .	121
7.27	Preference rating (blue) and predicted rating (black) for the Flowers sample.	122
7.28	Preference rating (blue) and predicted rating (black) for the Tulip sample modeled using all samples as input.	123
7.29	Preference rating (blue) and predicted rating (black) for the OilTrees sample-modeled using all samples as input.	124
7.30	Preference rating (blue) and predicted rating (black) for the Gray Patch 1 sample modeled using all samples as input.	124
7.31	Preference rating (blue) and predicted rating (black) for the DarkOil sample-modeled using all samples as input.	125
7.32	Preference rating (blue) and predicted rating (black) for the Flowers sample-modeled using all samples as input.	125
A.1	Painting 1: <i>Girl on a Divan</i> , Kirchner and Painting 2: <i>Vase of Flowers</i> , Picasso	130
A.2	Lighting setup used by the Department of Imaging and Visual Resources at MoMA using two strobes (highlighted with red circles) directed toward the ceiling from above the paintings.	131
A.3	Lighting setup used for Four-Light Imaging technique. Four lights are placed to the left, above, to the right, and below the painting at approximately 45 (highlighted with red circles).	132
A.4	Diffuse color image (top left), normal map (top right), and a close up of the same region of both the diffuse color and normal map (bottom left and right) of Painting 1.	134
A.5	Diffuse color image (top left), normal map (top right), and a close up of the same region of both the diffuse color and normal map (bottom left and right) of Painting 2.	135
A.6	Photograph and rendering (outlined in red) of Painting 1 with close ups of three different sections.	136
A.7	Photograph and rendering (outlined in red) of Painting 2 with close ups of two different sections.	137

A.8	Left: Rendering of Painting 1 (Kirchner) with high levels of gloss and texture produced by using directional lighting in Maya. Right: Rendering of the same area of Painting 1 with limited gloss and texture produced by using ambient lighting in Maya.	138
B.1	A screenshot of the Attribute Editor menu in Maya showing some of the attributes associated with the anisotropic shader.	140
B.2	Images of each reference stimuli taken with two strobes firing simultaneously.	143
B.3	A diagram of the lighting used to image each AMD panel. Light coming from the left, illuminates from the bottom of the rotated panel with a diffuser, while the Light from the right is stronger.	144
B.4	A screenshot of the Experiment 1 interface where observers are presented an reference stimuli (top) and an adjustable test stimuli (bottom).	145
B.5	Experimental setup with a physical sample presented under an LED source on the left to be compared to renderings displayed on the screen on the right.	147
B.6	Scale values calculated from z-scores for sample 1.	148
B.7	Scale values calculated from z-scores for sample 4.	148
B.8	Scale values calculated from z-scores for sample 5.	149
B.9	Scale values calculated from z-scores for sample 6.	149
B.10	Scale values calculated from z-scores for sample 7.	150
B.11	Scale values calculated from z-scores for sample 9.	150
B.12	Scale values calculated from z-scores for sample 10.	150

List of Tables

4.1	Percent falloff produced by each honeycomb grid and the reference.	38
4.2	Color difference (CIEDE2000) of images taken with the 20 and 30 degree grids compared to the spectrophotometer data for each patch for a close distance. Larger color differences are denoted by bold text.	45
4.3	Minimum, maximum, average, and ninetieth percentile color difference values for the twenty- and thirty-degree images.	46
4.4	Color difference (CIEDE2000) of images taken with the 20 and 30-degree grids compared to the spectrophotometer data for each patch for a long distance. Higher color differences are denoted by bold text.	47
4.5	Minimum, maximum, average, and ninetieth percentile color difference values for the twenty- and thirty-degree images for the long distance.	47
5.1	Blinn Shader parameter settings for each rendering.	53
6.1	Recorded angles for the light location from the surface normal of the painting for each image taken.	62
6.2	A list of questions asked to each observer after completing the experimental tasks. The intention of these questions is to provide insight into each observers ratings or selections.	70
6.3	Observer preferences compared to observer selected representative images for each of the eight samples denoted by the lighting angle used to capture each image.	91
7.1	RMSE values for samples used in the gloss perception modeled using the stepwise function method.	105
7.2	RMSE values for samples used in the gloss perception modeled using the stepwise function method with all samples used as input simultaneously. . .	108

7.3	RMSE values for samples used in the texture perception modeled using the stepwise function method	112
7.4	RMSE values for samples used in the texture perception modeled using the stepwise function method with all samples used as input simultaneously. . .	115
7.5	RMSE values for each predicted versus observer selected preference rating using the stepwise function method.	119
7.6	RMSE preference values for each sample using the stepwise function method with all samples used as input simultaneously.	123
B.1	Samples from the Artist Material Database and details about the medium, substrate, application technique, overcoat, and the Ward BRDF model parameters rho and alpha.	141
B.2	The spread values (1 to 100) and the roughness values (0.01 to 1.00) used to create 144 different renderings for each of the 12 samples.	142
B.3	A frequency matrix for sample 4 that shows how many times sample k was selected over sample j	146
B.4	A probability matrix for sample 4 that shows how many times sample k was selected over sample j	146
B.5	A matrix for sample 4 that shows z-scores and an interval scale for how often k was selected over sample j	146
B.6	My caption	148
C.1	A frequency matrix for sample 1 that shows how many times sample k was selected over sample j	151
C.2	A frequency matrix for sample 2 that shows how many times sample k was selected over sample j	152
C.3	A frequency matrix for sample 3 that shows how many times sample k was selected over sample j	152
C.4	A frequency matrix for sample 4 that shows how many times sample k was selected over sample j	152
C.5	A frequency matrix for sample 5 that shows how many times sample k was selected over sample j	153
C.6	A frequency matrix for sample 6 that shows how many times sample k was selected over sample j	153

C.7	A frequency matrix for sample 7 that shows how many times sample k was selected over sample j	153
C.8	A frequency matrix for sample 8 that shows how many times sample k was selected over sample j	154
C.9	A frequency matrix for sample 9 that shows how many times sample k was selected over sample j	154
C.10	A frequency matrix for sample 10 that shows how many times sample k was selected over sample j	154
C.11	A frequency matrix for sample 11 that shows how many times sample k was selected over sample j	155
C.12	A frequency matrix for sample 12 that shows how many times sample k was selected over sample j	155
C.13	A probability matrix for sample 1 that shows how many times sample k was selected over sample j	156
C.14	A probability matrix for sample 2 that shows how many times sample k was selected over sample j	156
C.15	A probability matrix for sample 3 that shows how many times sample k was selected over sample j	157
C.16	A probability matrix for sample 4 that shows how many times sample k was selected over sample j	157
C.17	A probability matrix for sample 5 that shows how many times sample k was selected over sample j	157
C.18	A probability matrix for sample 6 that shows how many times sample k was selected over sample j	158
C.19	A probability matrix for sample 7 that shows how many times sample k was selected over sample j	158
C.20	A probability matrix for sample 8 that shows how many times sample k was selected over sample j	158
C.21	A probability matrix for sample 9 that shows how many times sample k was selected over sample j	159
C.22	A probability matrix for sample 10 that shows how many times sample k was selected over sample j	159

C.23	A probability matrix for sample 11 that shows how many times sample k was selected over sample j	159
C.24	A probability matrix for sample 12 that shows how many times sample k was selected over sample j	160
C.25	Summary of responses to Gloss and Texture survey.	161
C.26	Summary of responses to preference versus representative survey.	161

Introduction

1.1 Motivation

In the museum world, physical proximity to a work of art is not always possible. Often museums and institutions lend artwork to other venues for long periods of time or remove paintings from exhibition for conservation treatment. Cases also occur where all pieces in a series of work are not owned by the same museum or simply not enough wall space exists and therefore they are rarely, if ever, exhibited together. Visitors to a museum would benefit from a digital rendering that was capable of simulating the appearance of work that had either been removed for conservation tasks or lent to a different institution. Although no substitute exists for the real work of art, being able to watch a video that allows each viewer to experience a variety of viewing or lighting directions that accurately simulates changes occurring on the surface of the work due to the interaction of lighting and surface material could be an adequate alternative to static images. Providing stills that convey more information than traditional images could be as beneficial when a video is not appropriate. These renderings could also be used to experience an entire series that cannot be physically brought together for exhibition or as a tool to enrich an art history class by replacing still photographs with a more engaging alternative. In addition to using these renderings as part of a museum or classroom setting, images that convey texture information would also be valuable to institutions interested in setting up a virtual gallery, where individuals could experience artwork from their homes or work places using their personal computers.

1.2 Approach

A method of capturing and rendering surface information about textured two dimensional surfaces, i.e. paintings, that uses traditional equipment with only a few inexpensive additions is presented. Texture capturing and rendering is the primary focus of this research, as a copious amount of literature dedicated to capturing and reproducing color information of

artwork already exists [Martinez et al., 2002; Tominaga and Tanaka, 2002; Berns, 2005; Chen and Berns, 2012b]. The texture capture system is used to generate normal maps of a set of painting samples. These normal maps are then used to generate digital renderings of each sample using Maya® and mental ray, here after referred to as Maya, a commonly used graphics rendering software and plugin.

To assess the viability of using Maya, a set of psychophysical experiments was conducted that compared the rendered data in the form of a video as well as in single frames to the original samples. Observers were asked to assess the software's ability to convey realistic texture information as well as changes in viewing direction that are often exemplified by changes in the perception of highlights across the painting's surface determined by the apparent gloss of the material.

The two main goals of this research are to evaluate the plausibility of capturing and rendering fine artworks using typical studio equipment and commercially available software and to explore the appearance of texture and gloss in static images. To achieve these goals, this research will be comprised of four main parts outlined below:

1. Implementing a method that captures surface normal from a traditional photography studio set-up with little-to-no additional equipment using photometric stereo,
2. Investigating the viability of using Maya as a rendering software to display virtual paintings based on real samples,
3. Using images of fine art samples taken with different lighting angles to relate image statistics to perceived gloss and texture in images as determined by observers, and
4. Identifying preferences for virtual fine art reproduction that focuses on the reproduction of gloss and texture related to lighting.

1.3 Dissertation Structure

This dissertation is broken up into chapters. Chapter 2 discusses previous work that was influential to the current research including color imaging, photometric stereo, and rendering fine art. Chapter 3 discusses deriving surface normal maps using traditional photographs

and a typical studio environment building off of the pre-existing Four Light Imaging technique. Chapter 4 presents nuances of lighting for the purpose of scientific imaging rather than fine art photography. Chapter 5 is an in-depth depiction of using Maya together with normal and color maps to render fine art. Chapters 6 and 7 delves into texture and gloss perception in images and how observer ratings relate to image statistics as well as viewing preferences associated with viewing fine art online. These preferences include how much gloss is rendered for a variety of artist medium types. Conclusions and future research is presented in Chapter 8.

Introduction to Imaging and Rendering Fine Art

The imaging of fine artworks most commonly focuses on the accurate capture and reproduction of color. A variety of methods have been explored in this area including scanning technologies, traditional RGB camera Capture, Dual-RGB capture, and multi-spectral and hyper-spectral imaging [Besser, 1991; Hardeberg et al., 2002; Berns et al., 2005; Zhao et al., 2005; Fischer and Kakoulli, 2006; Chen and Berns, 2012c]. However, an increase in the amount of research related to capturing surface information is also saturating the available literature, including photometric stereo, reflectance transformation imaging (RTI), laser scanning, and photogrammetry [Woodham, 1980; Debevec et al., 1996; Levoy et al., 2000; Malzbender et al., 2001]. These two foci, color and surface information, do not quite complete the necessary data sets for rendering fine art. The third necessary area of interest involves capturing gloss, or surface microstructure, often in the form of bidirectional reflectance distribution function, or BRDF [Nicodemus et al., 1977]. BRDF relates the amount of light incident on an object for a given direction to the amount of light reflected from the object in a different direction.

In this chapter, previous and influential approaches to imaging fine artwork are discussed briefly. The areas of these approaches include the capture of color, surface macrostructure, and surface microstructure.

2.1 Imaging Fine Art: Color

Since nearly the beginning of capturing color images of artwork or of other colorful scenes, the desire and often need to capture or render color more accurately has been recognized. With the introduction of digital media, a variety of techniques have been explored to address this issue. These techniques include using a monochrome charged couple device (or CCD) and multiple filters to acquire spectral data. One technique includes the use of a

liquid crystal tunable filter and monochrome sensor [Berns et al., 2005]. Another technique known as the Dual-RGB approach, replaces the liquid crystal tunable filter with a filter wheel containing two specifically selected filters and the monochrome camera with camera containing a color-filter-array that has had the the built-in infrared cut-off filter replaced with a visible bandpass filter to create a six-channel, multi-spectral imaging system from a commercial three-color filter array [Hardeberg et al., 2002; Zhao et al., 2005; Fischer and Kakoulli, 2006; Berns and Chen, 2012; Berns, 2016].

2.1.1 RGB dSLR

The appeal of a traditional color filter array digital single-lens reflex, or dSLR, camera is its commercial availability, well-documented user guidelines, availability of high resolution arrays, and low cost. Although the use of a traditional camera is not ideal for color capture on its own, the incorporation of flat-fielding and colorimetric characterization greatly improve color accuracy. Flat-fielding is performed to remove non-uniformities in lighting and the camera sensor. Images taken of a uniform white background using the same camera and lighting setup can be flat-fielded using Equation 2.1, where the *sampleImage* is an image of the object or painting of interest, *whiteImage* is an image of the uniform white background, and *w* is a scaling factor determined by the white patch of a color calibration chart such as the X-rite ColorChecker® Digital SG. Often, the image of the white background is filtered with a low-pass Gaussian filter to remove any non-uniformities in the white board or background [Berns and Chen, 2012]. Details on the benefits of flat-fielding can be found in [Witwer and Berns, 2015].

$$ffImage = w * \frac{sampleImage}{whiteImage} \quad (2.1)$$

Colorimetric characterization is performed by relating average, flat-fielded RGB digital counts from the camera to XYZ tristimulus values of a color calibration chart, by a linear transformation matrix, *M*, shown by Equation 2.2. Equation 2.2 is written for a single color calibration patch where *XYZ_m* are tristimulus values calculated from spectrophotometer measurements and *RGB* are average photometric data. To account for any stray light or error associated with lighting geometry mismatch between measured and imaged data, Equation 2.2 can be written as Equation 2.3 where *R_o*, *G_o*, and *B_o* are offsets. The matrix coefficients and the offset terms are optimized simultaneously to minimize the average CIEDE2000

($S_L=1$) color difference. In the case of n number of color chart patches, Equation 2.4 can be used.

$$\begin{bmatrix} X_m \\ Y_m \\ Z_m \end{bmatrix} = \begin{bmatrix} m_1 & m_2 & m_3 \\ m_4 & m_5 & m_6 \\ m_7 & m_8 & m_9 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad (2.2)$$

$$\begin{bmatrix} X_m \\ Y_m \\ Z_m \end{bmatrix} = \begin{bmatrix} m_1 & m_2 & m_3 \\ m_4 & m_5 & m_6 \\ m_7 & m_8 & m_9 \end{bmatrix} \begin{bmatrix} R - R_o \\ G - G_o \\ B - B_o \end{bmatrix} \quad (2.3)$$

$$\begin{bmatrix} X_{m,1} & \dots & X_{m,n} \\ Y_{m,1} & \dots & Y_{m,n} \\ Z_{m,1} & \dots & Z_{m,n} \end{bmatrix} = \begin{bmatrix} m_1 & m_2 & m_3 \\ m_4 & m_5 & m_6 \\ m_7 & m_8 & m_9 \end{bmatrix} \begin{bmatrix} R_1 - R_{1,o} & \dots & R_n - R_{n,o} \\ G_1 - G_{1,o} & \dots & G_n - G_{n,o} \\ B_1 - B_{1,o} & \dots & B_n - B_{n,o} \end{bmatrix} \quad (2.4)$$

The CIEDE2000 metric is calculated from CIELAB values computed from tristimulus values and their predicted counterparts. This approach can be implemented regardless of commercial camera choice [Fairchild et al., 2008]. The use of the commercially available, professional-quality camera makes this research more relevant to a wider range of museum and studio photographers.

2.1.2 Dual-RGB

It has been shown by Berns, et al. (2005) that spectral imaging is preferred for color capture over traditional three-channel capture systems for color accuracy; however, the steps are tedious and the equipment involved is sometimes expensive. An alternative to using these cumbersome systems is to introduce a set of optimized filters to a three-channel camera. This system still incorporates flat-fielding and color characterization with the use of a color chart. In some cases, the system has been shown to outperform spectral imaging acquired by a monochrome sensor and multiple filters or a tunable liquid crystal filter [Berns et al., 2005].

Although the use of dual filters with a commercial dSLR is possible, for the purpose of this research, the higher-resolution Sinar 86H 48 megapixel camera will be considered for use with the dual-RGB technique. The dual-RGB technique replaces the infrared filter with clear

glass and alternates the use of two filters, one yellow and one bluish green, in front of the lens. The filters have VIS-bandpass and anti-reflection coatings that limit transmission to the visible spectrum. The use of alternating filters transforms this three-channel camera into a five-channel camera due to the different spectral sensitivities produced by the red and blue channels when accompanied by each filter. Camera sensitivities when using both filters is shown in Figure 2.1 and the five sampling wavelength regions produced by this method is shown in Figure 2.2. (The green channel is unaffected by filtering, reducing dual-RGB to five channels.) Similar flat-fielding and colorimetric estimation procedures as those mentioned above are performed, producing an acceptable compromise between colorimetric and spectral accuracy without dramatically increasing imaging and processing time [Chen and Berns, 2012c].




Fig. 2.1. – Camera sensitivities when using a blue-green filter (solid lines) and a yellow filter (dashed lines) [Chen and Berns, 2012c].

The Dual-RGB approach uses two calibration routines, a colorimetric calibration and a spectral calibration. The colorimetric calibration relates camera signals to estimated CIE tristimulus values, \hat{X} , \hat{Y} , and \hat{Z} , according to Equation 2.5 where R_b , G_b , and B_b are image data associated with the blue-green filter and R_y , G_y , and B_y are image data associated




Fig. 2.2. – Five sampling wavelength regions from using the Dual-RGB method of imaging [Chen and Berns, 2012c].

with the yellow filter. Coefficients for Matrix \mathbf{M} are optimized values and $R_{b,k}$ through $B_{y,k}$ are offset terms used to account for stray light and differences in the lighting geometry between the imaged data and the spectrophotometer used to estimate the reference tristimulus values [Berns, 2016]. The coefficients of the $\mathbf{M}_{3 \times 6}$ matrix are optimized by minimizing the average CIEDE2000 between the average image data for each patch of a color calibration chart and spectrophotometer colorimetric data by using the unconstrained minimization function in Matlab, `fminunc`. The optimization is performed twice. The first optimization uses only data associated with the blue filter to determine the values of the first three columns of matrix $\mathbf{M}_{3 \times 6}$. The second optimization uses these values as starting values that are then used to optimize the entire 3 by 6 matrix. Optimizing for the blue filter first improves image quality while using only the blue filter data to estimate \hat{Y} and only using the green camera

signals from the blue filter image improves spatial image quality and reduces noise [Berns, 2016].

$$\begin{bmatrix} \hat{X} \\ \hat{Y} \\ \hat{Z} \end{bmatrix} = \begin{bmatrix} m_{1,1} & m_{1,2} & m_{1,3} & m_{1,4} & 0 & m_{1,6} \\ m_{2,1} & m_{2,2} & m_{2,3} & 0 & 0 & 0 \\ m_{3,1} & m_{3,2} & m_{3,3} & m_{3,4} & 0 & m_{3,6} \end{bmatrix} \begin{bmatrix} R_b - R_{b,k} \\ G_b - G_{b,k} \\ B_b - B_{b,k} \\ R_y - R_{y,k} \\ G_y \\ B_y - B_{y,k} \end{bmatrix} \quad (2.5)$$

The spectral calibration relates camera signals, R_b , G_b , B_b , R_y , and B_y , to spectral reflectance factor \hat{R} according to Equation 2.6. This calibration occurs in four steps. The first step is to estimate the terms in matrix \mathbf{M} using a pseudo-inverse as shown in Equation 2.7 where PINV represents the Moore-Penrose pseudo-inverse. Equation 2.7 is in terms of n , number of samples. The second step is to calculate the spectral reflectance factor RMSE. The third step is to calculate the maximum and minimum reflectance factor for each patch of the color calibration chart at each wavelength and then calculate the SS error for the maximum and minimum differences. Step four is to optimize Equation 2.8, where e refers to error, using `fminunc` and the pseudo-inverse matrix as starting values. Equation 2.8 is used to improve the accuracy of neutrals spectrally [Berns, 2016].

$$\begin{bmatrix} \hat{R}_{\lambda=380} \\ \vdots \\ \hat{R}_{\lambda=7300} \end{bmatrix} = \mathbf{M} \begin{bmatrix} R_b \\ G_b \\ B_b \\ R_y \\ B_y \end{bmatrix} \quad (2.6)$$

$$\mathbf{M} = \begin{bmatrix} \hat{R}_{\lambda,1=380} & \dots & \hat{R}_{\lambda,n=380} \\ \vdots & \ddots & \vdots \\ \hat{R}_{\lambda,1=730} & \dots & \hat{R}_{\lambda,n=730} \end{bmatrix} PINV \begin{bmatrix} R_{b,1} & \dots & R_{b,n} \\ G_{b,1} & \dots & G_{b,n} \\ B_{b,1} & \dots & B_{b,n} \\ R_{y,1} & \dots & R_{y,n} \\ B_{y,1} & \dots & B_{y,n} \end{bmatrix} \quad (2.7)$$

$$Z = e_{RMS} + 10e_{maxSS} + 50e_{minSS} \quad (2.8)$$

2.2 Imaging Fine Art: Surface Macrostruture

The surface macrostructure is commonly referred to as depth, topography, or surface structure. In the case of fine art, surface macrostructure can include brush strokes, canvas weave, cracks in paint, impasto, and much more that affects the surface in visual and structural ways.

Along with color, surface macrostructure has been an area of interest in the museum and archiving fields for many years. Several different techniques have been developed or enhanced for the purpose of determining surface normal, or texture information, of fine art. These techniques include shape from shading [Horn, 1970], photogrammetry [Waldhäusl, 1992], structured light [Akca et al., 2006], photometric stereo [Woodham, 1980; Rushmeier et al., 1997; Tominaga and Tanaka, 2002; Barsky and Petrou, 2003; Sun et al., 2007; Berns and Chen, 2012], laser scanning [Levoy et al., 2000], reflectance transformation imaging (RTI) with polynomial texture mapping (PTM) [Malzbender et al., 2001; Padfield et al., 2005; Cultural Heritage Imaging, 2013], raking light imaging [Tisato and Parraman, 2014], and linear source imaging [Gardner et al., 2003]. These techniques range from the use of a small number of images taken with an equal number of lighting directions to thousands of images, depending on the size of the sample, taken from different vantage points.

The main focus of this research utilizes photometric stereo; however, the technique is also compared to RTI in some applications.

2.2.1 Photometric Stereo

A commonly used technique for capturing surface information is photometric stereo, first introduced in 1980 [Woodham, 1980]. Woodham observed that a reflectance map was capable of incorporating reflectance from a surface, the geometry in which an object is imaged, and the illumination of a scene into a function of surface orientation in terms of image intensity. This depends on Lambert's cosine law that states that the intensity reflected from a diffuse, lamberian surface depends on the orientation of the the light source, but is independent of the viewing direction. The first step in calculating a surface-normal map was to determine the lighting direction for each scene. The lighting direction, $L = [l_x, l_y, l_z]$, for a

point with a given surface normal N and diffuse albedo ρ were used to define the intensity, I , of each pixel value according to Equation 2.9.

$$I = \rho L * N \quad (2.9)$$

This can be written out for n number of light sources as

$$\begin{bmatrix} I^1 \\ I^2 \\ \vdots \\ I^n \end{bmatrix} = \begin{bmatrix} l_x^1 & l_y^1 & l_z^1 \\ l_x^2 & l_y^2 & l_z^2 \\ \vdots & \vdots & \vdots \\ l_x^n & l_y^n & l_z^n \end{bmatrix} \times \begin{bmatrix} \rho N_x \\ \rho N_y \\ \rho N_z \end{bmatrix} \quad (2.10)$$

This equation can be solved for ρN using linear regression, resulting in Equation 2.11, where \dagger denotes a pseudo-inverse, with the constraint of the unit vector shown in Equation 2.12. The equation is constrained to the unit vector because of the desired knowledge of vector direction rather than magnitude.

$$\begin{bmatrix} \rho N_x \\ \rho N_y \\ \rho N_z \end{bmatrix} = \begin{bmatrix} l_x^1 & l_y^1 & l_z^1 \\ l_x^2 & l_y^2 & l_z^2 \\ \vdots & \vdots & \vdots \\ l_x^n & l_y^n & l_z^n \end{bmatrix}^{\dagger} \times \begin{bmatrix} I^1 \\ I^2 \\ \vdots \\ I^n \end{bmatrix} \quad (2.11)$$

subject to

$$\sqrt{N_x^2 + N_y^2 + N_z^2} = 1 \quad (2.12)$$

Since Woodham's introduction of using two (theoretically) or more images to determine surface orientation, many articles have been written in hopes of improving the accuracy of the photometric stereo technique. A common theme in previous research has been to increase the number of lights to improve the accuracy of the resulting normal map. Three different approaches are of particular interest. The first was described by Rushmeier, et al. in 1997. Rushmeier, et al. used a photometric stereo technique with five lights, a fixed video camera, and a video capture board [Rushmeier et al., 1997] shown in Figure 2.3.




Fig. 2.3. – Five light setup for the technique described by Rushmeier, et al. [Rushmeier et al., 1997].

The method presented in their work eliminated maximum and minimum values according to the assumption that highlights and shadows were represented by these values, respectively. This method allowed resulting maps to be used to render objects without requiring the reconstruction of the original geometry of the object. This five-light photometric stereo technique was capable of representing fine scale and self-shadowing surfaces but fell short in its ability to represent objects with large changes in geometry, (i.e. three dimensional objects) because of the possibility of multiple intra-object inter-reflections. Surfaces with widely varying BRDF values across the objects were also difficult to represent because they violated the lambertian surface assumption required by Woodham's method.

In the second approach, Sun, et al. expanded the traditional three-light photometric stereo technique to include the use of six lights and a hierarchal sorting technique [Sun et al., 2007]. This selection method was used to exclude highlights and shadows. Six lights were incorporated into their technique to guarantee that each location on a sample was illuminated by at least three lights once highlights and shadows were excluded. A flow chart of the hierarchal sorting is shown in Figure 2.4.

The results of the six-light method with hierarchal selection were compared to: using six lights without hierarchical selection; Rushmeier's five-light technique; and the traditional

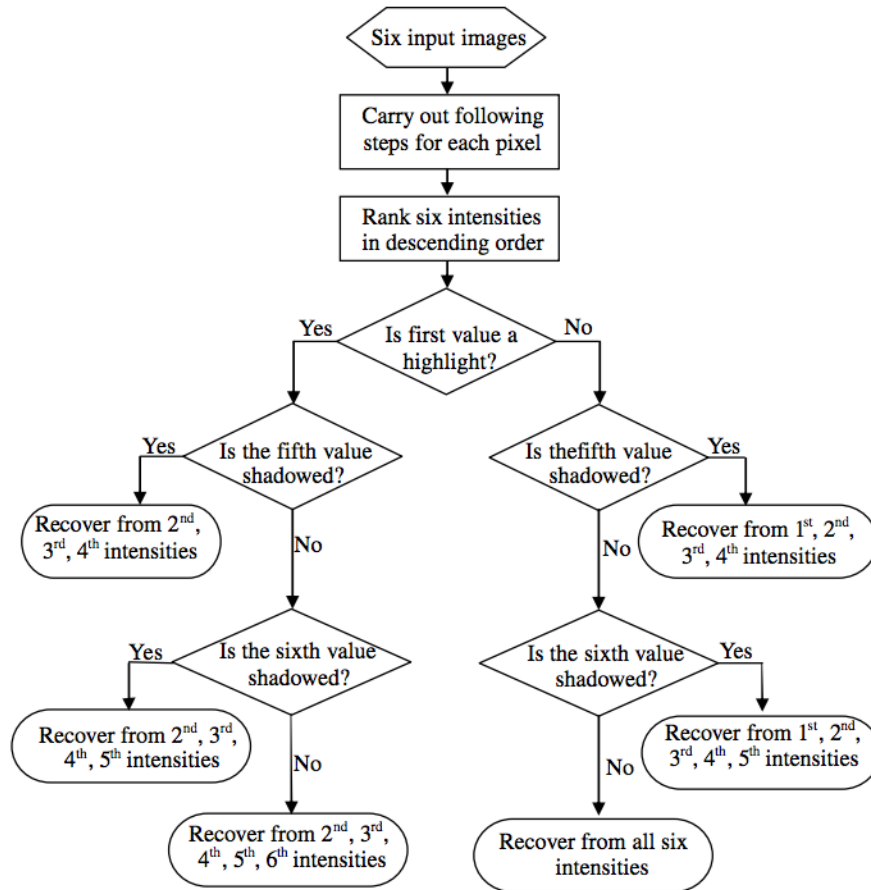


Fig. 2.4. – Flow chart showing the steps in the hierarchical sorting technique to remove highlights and shadows [Sun et al., 2007].

three-light technique shown in Figure 2.5. After a comparison of speed of implementation of each algorithm, reconstructed albedo images, height maps, and rendered surface orientation images, the authors concluded that the six-light method with selection out-performed the other methods. However, the five-light technique (maximum and minimum exclusion) was a very close second [Sun et al., 2007].

Both of the previously mentioned techniques increased the number of lights used in their respective techniques with the addition of removing pixel values that corresponded to highlights and shadows to improve the accuracy of the surface normal maps. The third technique used an alternative to using a selection or elimination technique. Berns, et al. employed the use of polarizers with a four-light setup to remove specular highlights [Berns and Chen, 2012]. The use of polarizers was implemented to satisfy the Lambertian surface




Fig. 2.5. – Comparison of four different methods for calculated albedo a) three lights, b) Rushmeier's five light technique, c) all six lights, and d) six lights without shadows and specularities and a comparison of the same four methods for rendered images using e) three lights, f) Rushmeier's five light technique, g) all six lights, and h) six lights without shadows and specularities [Sun et al., 2007].

assumption required by Woodham [Woodham, 1980]. Cross-polarization eliminated the specular component from the images.

The Berns, et al. technique placed a linear polarizer in front of each light (placed at forty five degrees from the surface of the painting and ninety degrees from each neighboring light along an annulus) and camera. A diagram of this set up is shown in Figure 2.6. Once cross-polarization was achieved by rotating each polarizer in front of each light source until highlights on a reflective surface were minimized, images of the object, calibration target, cue ball, and white background were captured using one light source at a time. The problems associated with highlights were addressed with the use of polarizers, but problems created by the presence of shadows were not directly addressed. However, similar to Rushmeier, a shadow was assumed to occur in only one light direction, therefore a shadow would reduce the number of light directions to three since the signal was null. A null signal would have the same effect as excluding the shadowed pixel from surface normal calculations.




Fig. 2.6. – A diagram of the four light with polarizers technique shows the relative locations of a painting, camera and four different lights [Chen and Berns, 2012a].

2.2.2 Reflectance Transformation Imaging

Reflectance transformation imaging, or RTI, stems from the work of Malzbender, et al. on polynomial texture maps [Malzbender et al., 2001]. RTI is an imaging technique that uses multiple lights and a fixed camera position, similar to photometric stereo, but uses a polynomial texture mapping method to fit reflectance distribution data [Duffy, 2013]. The model proposed by Malzbender, et al. took luminance information for each pixel and mathematically described the information in terms of the direction of the incoming light in the form of a biquadratic polynomial stored with color pixel data. This model provided per pixel information about surface inter-reflection, self-shadowing, and subsurface scattering [Mudge et al., 2005]. For diffuse objects, the biquadratic for each texture element was represented by Equation 2.13, where (l_u, l_v) are normalized light vector projections into the (u, v) coordinate system and L is the surface luminance at the (u, v) coordinate. The a_0 - a_5 coefficients were fit to each texture element using data taken from all images and the singular value decomposition technique [Malzbender et al., 2001].

$$L(u, v; l_u l_v) = a_0(u, v)l_u^2 + a_1(u, v)l_v^2 + a_2(u, v)l_u l_v + a_3(u, v)l_u + a_4(u, v)l_v + a_5(u, v) \quad (2.13)$$

This technique allows control of lighting conditions as well as specular enhancement, diffuse gain, and light-direction extrapolation in rendering. Specular enhancement uses the normal vector to add specular properties to enhance surface shape [Malzbender et al., 2001]. A comparison of a traditional photograph (left), specular enhanced image (middle), and a normal map (right) of a stone carving is shown in Figure 2.7 [Duffy, 2013].



Fig. 2.7. – A comparison of a traditional photograph (left), specular enhanced image (middle), and a normal map (right) of a stone carving [Duffy, 2013].

The use of RTI has increased in popularity with classes and starter kits being made available by Cultural Heritage Imaging, a nonprofit corporation focused on digital capture and documentation of artwork, examples of cultural heritage, and historic artifacts [Cultural Heritage Imaging, 2014]. The RTI system is also portable and has been used for imaging in the field. For example, the RTI system was used at a remote site in southern Armenia to image rock art [Duffy, 2013].

The available training, inexpensive equipment requirements, camera-independent performance, and portability make the RTI technique an appealing option for archiving artwork and provides the motivation for evaluating its performance in conjunction with the provided RTIviewer along side the photometric stereo technique.

2.3 Imaging Fine Art: Surface Microstructure

The bidirectional reflectance distribution function, or BRDF, is a function that relates how much the irradiance from a known direction contributes to the radiance in a given direction [Nicodemus et al., 1977]. The amount of light reflected from a surface in a given direction is

related to the surface orientation and roughness on a micro-scale. The measurement and fitting of the BRDF for a given surface is an important part of rendering due to its contribution of information about how light interacts with a surface and affects the perceptual attribute of gloss.

2.3.1 Gloss and Material Appearance

Gloss is an important part of the visual experience that helps in the identification of objects. Gloss can inform an observer about the type of material an object is made out of, if paint is wet or dry, and if roads are icy. Hunter and Harold categorized gloss into six different visual attributes: specular gloss—shininess or brightness of highlights; contrast gloss—contrast between specular and diffuse areas; distinctness-of-image gloss—sharpness of reflected image; absence of bloom gloss—lack of cloudiness near highlights; sheen—grazing angle shine; and absence of texture gloss (surface uniformity gloss)—uniform or smooth surface appearance shown in Figure 2.8 [Hunter and Harold, 1987]. Absence of texture gloss is not included in the figure because it is not a function of reflectance.

A previous study was performed to evaluate cross-media comparison of gloss in 2013 [Fores et al., 2013]. In this research, a light booth was constructed that could be simulated by a typical LCD display. The physical samples were black cylinders of varying gloss levels placed inside the light booth on a black and white checkerboard bottom. Observers were asked to adjust the gloss level of a black cylinder on a checkered bottom on the display until it matched the physical sample. The results of this experiment showed that observers were able to accurately match low gloss samples but slightly underestimated the level of gloss for materials with mid- to high-gloss levels [Fores et al., 2013]. The results of this cross-media study has set a precedent for the cross-media experiments described in later chapters.




Fig. 2.8. – A diagram showing (from right to left) specular gloss, sheen, contrast gloss, absence of bloom, and distinctness-of-image gloss in terms of incoming and outgoing light [Hunter and Harold, 1987].

2.3.2 Measuring BRDF

Measuring BRDF comprehensively can be a laborious and time-consuming task because of the numerous angles that need to be used for measuring incoming and outgoing light. Ward proposed a means for measurement using a gonireflectometer, shown in Figure 2.9. The general design of a gonireflectometer includes a photometer that moves in relationship to a sample and a sample holder that moves in relationship to light sources. These moving parts allow for the measurement associated with two incident angles and two excitant angles as is required for BRDF [Ward, 1992].

In addition to using a gonireflectometer, BRDF can be measured using a linear light source. In the research done by Gardner, et al. a single linear light source was used to take a series of images with a fixed camera while the light source scanned over the sample. A second pass was made with a different geometry to provide lighting from all possible directions. Using images taken with this setup, Gardner, et al., were able to estimate diffuse color, specular intensity, and roughness properties that could be fit to the Ward BRDF model [Gardner et al., 2003]. A diagram of the linear light source set up is shown in Figure 2.10.

Inverse rendering is another method for estimating BRDF, the difference being that the technique uses images to estimate the BRDF of an object inside a real scene instead of directly measuring the objects' reflectance. Boivin and Gagalowicz proposed a method for inverse rendering from a single image. This method estimated different types of param-




Fig. 2.9. – A diagram of a gonireflectometer showing a reflectance detector and a moving light source [Ward, 1992].




Fig. 2.10. – A diagram of the linear light source set up [Gardner et al., 2003].

ters of reflectance using a single photograph of a scene taken under known illumination conditions, accompanied with a three-dimensional model of the scene which allowed the production of photorealistic rendered images [Boivin and Gagalowicz, 2002].

2.3.3 Modeling BRDF

BRDF models can be placed into two main categories: physically based models and empirical models. Empirical models include Phong and Ward models, and physical models include those like Torrance-Sparrow [Chen et al., 2007].

The Torrance-Sparrow model is considered a predecessor to other models and one of the most physically complete [Montes and Ureña, 2012]. The Torrance-Sparrow model represents rough surfaces as being composed of small mirror-like facets. The Phong model is a simplification of the Torrance-Sparrow model that does not obey reciprocity or energy conservation but whose simplicity makes it a frequent choice in computer graphics rendering [Montes and Ureña, 2012]. The Ward model is also not physically plausible but it is computationally inexpensive and capable of fitting measured reflectance data [Montes and Ureña, 2012].

The isotropic Ward reflection model is shown in Equation 2.14 where $\rho(\theta_i, \phi_i, \theta_o, \phi_o)$, represents the bidirectional reflectance distribution function, θ_i , ϕ_i , θ_o , and ϕ_o are coordinates for incoming, i , and outgoing, o , directions, δ is the half angle between the incoming and

outgoing vectors, ρ_d is the diffuse reflection, ρ_s is the specular component, and α is the spread of the specular lobe [Ward, 1992].

$$\rho(\theta_i, \phi_i, \theta_o, \phi_o) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{\exp[-(\tan^2 \delta / \alpha^2)]}{4\pi\alpha^2 \sqrt{\cos \theta_i \cos \theta_o}} \quad (2.14)$$

More recently, Löw, et al. introduced two BRDF models for glossy surfaces [Löw et al., 2012]. The first model, referred to as the Smooth Surface BRDF model, was based on the Rayleigh-Rice theory for optically smooth surfaces. The second model was the Microfacet BRDF model, which was a modified version of the Cook-Torrance model. Löw, et al. claimed that the new models possessed more flexibility in the types of surfaces that they could characterize, however they are geared towards modeling glossy surfaces and do not outperform other models when used to characterize material like fabric, due to the difference in the scattering properties of these materials.

2.3.4 Artist Material Database

The BRDF of an object is an important contributor to how the object is rendered. Because this research is primarily focused on the rendering of fine art, knowing the BRDF of different mediums, such as acrylic paint, oil paint, or watercolor, is an important first step.

The Artist Material Database is a public-domain image and optical database of artist materials consisting of a variety of mediums, substrates, colors, and overcoats, resulting in almost 600 39 millimeter-square samples on 26 panels. The database includes twelve different mediums and a total of six different substrates (medium specific) [Ashbaugh et al., 2009]. An example of one of the panels is shown in Figure 2.11.

Each patch in the database was carefully imaged and measured using the gonio-imaging system developed at the Munsell Color Science Laboratory (MCSL) at Rochester Institute of Technology. The system used a sample holder that moves about three axes and three different detectors, a Canon 5D camera, the multispectral camera developed by MCSL, and a Konica Minolta CS 2000 spectroradiometer. Data taken with the Canon 5D was fit to the Ward BRDF model, by first parameterizing the measurements with spherical coordinates, then using photometric stereo to compute diffuse surface normals, computing specular



Fig. 2.11. – An example grid from the Artist Material Database (AMD) showing unvarnished, matte acrylic varnish, gloss acrylic varnished, gamvar applied with two different techniques [Ashbaugh et al., 2009].

surface normals for some number, typically 100, of the brightest pixels; and finally fitting to a BRDF model using the Levenber-Marquardt method and all brightest pixels used to calculate specular surface normal and taking their average [Ashbaugh et al., 2009; Chen et al., 2011].

2.4 Rendering Fine Art

Rendering a realistic image of fine artwork requires knowledge of the diffuse color, surface structure, and how light interacts with the surface of the work, usually in the form of the material's BRDF. Although there are many ways to capture data associated with each of these three properties, the general process for rendering an image with these data are similar, regardless of which BRDF model is used to describe the data [Blinn and Newell, 1976; Gardner et al., 2003; Mudge et al., 2005; Chen et al., 2007; Chen, 2008].

2.4.1 Rendering Workflow

A generalized rendering workflow is shown in Figure 2.12. This workflow uses examples of measurement techniques and BRDF models; the actual methods used are at the discretion of the renderer. By knowing color and surface structure information, along with how light interacts with the object, a defined lighting direction can be modeled for a defined viewing direction using an appropriate model assuming that the surface is of uniform gloss.




Fig. 2.12. – Generalized workflow for rendering fine art requires surface normal information, diffuse color, BRDF of microstructure, and known lighting and viewing conditions.

Although knowing the BRDF of a material that is to be rendered is beneficial, it is possible to render a surface using a rendering software to approximate the way light would react when it comes in contact with the object without having measured the BRDF. The use of rendering software is improved by knowing the BRDF of a reference material.

2.4.2 RTIviewer, ArtViewer, and Commercial Graphics Rendering Software

Measuring the BRDF of every sample or artifact for rendering quickly becomes tedious, and, with a large number of objects, impractical. By measuring and modeling a database like the MCSL Artist Material Database (AMD), a reference material can be used to model an artifact using software that takes into account these material attributes. The current research relies on the use of two noncommercial softwares– RTIviewer from Cultural Heritage Imaging

and ArtViewer from MCSL –and one commercial software, Maya[®] from Autodesk (with a mental ray plugin) [Berns and Chen, 2012; Cultural Heritage Imaging, 2013; Autodesk, Inc., 2015b].

ArtViewer was written based on the RTIviewer and, therefore, works similarly. Users select the incoming lighting direction, but the viewing direction is restricted to 0° from normal. A screenshot of ArtViewer is shown in Figure 2.13. ArtViewer allows the user to select a similar material and finish from the Artist Material Database. By making this selection, the BRDF properties of this reference sample are applied to the artifact allowing small changes in specularities and roughness [Chen, 2013]. The RTIviewer allows adjustment of diffuse color, specularities, and highlight size, as well as allowing different viewing modes. These include normal viewing and specular enhancement [Cultural Heritage Imaging, 2013]. The appeal of these rendering softwares comes from their simplicity and physically based adjustments that occur in realtime. However, all of these platforms are limited to rendering that does not include spatially-varying BRDF. Using the commercial software Maya increases the number




Fig. 2.13. – A screenshot of the ArtViewer software with a Van Gogh style painting loaded.

of user-defined input options as well as increasing the number of available viewing angles. Maya was chosen because of its prevalence in the graphics community, availability to the student population, well-documented user guides and tutorials, and its compatibility with the mental ray plug in.

Maya is an animation, modeling, simulation, and rendering software [Autodesk, Inc., 2015b]. The software includes a general-purpose renderer called mental ray. Mental ray uses rendering techniques such as a scanline rendering algorithm for primary visible surface determination and binary space partitioning algorithms to render ray tracing effects. Mental ray supports physically correct simulation of general global illumination. Supported simulation includes any combination of diffuse, glossy, and specular reflection and transmission [Autodesk, Inc., 2015a].

The mental ray software uses multiple shaders, or algorithms that specifies how a surface reacts to light, to create geometric features like bump and displacements maps, (grayscale textures mapped to a surface), to create an illusion of surface relief, and light sources. These shaders provide parameters that can be manipulated to adjust material properties, ambient lighting, diffuse and specular color, and reflectivity to name a few [Autodesk, Inc., 2015a].

All three of these softwares use the diffuse color and normal maps generated from measuring the artifacts as part of the rendering process. The introduction of gloss or micro-structure information come from either selecting a reference material or a model associated with a predetermined shader, such as a Phong or Blinn shader in Maya [Autodesk, Inc., 2015c]. More information about rendering with Maya can be found in Chapter 5 with brief descriptions of various rendering methods discussed throughout multiple chapters.

2.5 Related Work

This research was influenced by the MCSL Master's theses of Y. Chen (2008) and L. Chen (2013). The goal of Y. Chen's research was to develop a system to record information about fine art paintings when measured under different illumination from different angles. This data was then used to render images of the paint surface using two different models [Chen, 2008]. Y. Chen's research employed a digital camera with one degree of freedom for measurement that was used to document the relative radiance of a painted surface as a function of illumination angle. A diagram of this set up is shown in Figure 2.14 [Chen, 2008].




Fig. 2.14. – Diagram of the setup used by Y. Chen to record information about fine art under varying illumination angles [Chen, 2008].

The painted surfaces included thirteen paint samples of varying gloss levels, materials, and colors. The Phong model and Torrance-Sparrow model were used to represent the radiance data collected. The ability of each model to render accurate images of each of the thirteen samples was compared computationally and visually using paired comparison experiments. An example of a uniform sample and painted sample are shown in Figure 2.15. Computationally, both the Phong and Torrance-Sparrow models performed well for matte samples. However, both models performed poorly for glossy samples with errors increasing with the increase of gloss. Although, the Torrance-Sparrow model performed more accurately computationally and visually for glossy samples. The two models also performed similarly for different materials and samples with complex surface structures [Chen, 2008].

L. Chen's research began with a survey used to identify the most important features of paintings that need to be considered when photographing and reproducing each work of art. Her survey asked nine questions of twenty-three experts in the fine art reproduction field. This survey identified color, gloss, and texture as these three features. The information gathered in this survey was used to inform the experiments designed as part of her research.

Chen selected eight different systems to compare how well each system rendered the total appearance of paintings. These eight systems included conventional digital reproductions,




Fig. 2.15. – Examples of the Phong and Torrance-Sparrow models' ability to render samples [Chen, 2008].

a linear system using a digital camera and a linear light source, the Three-Light system, the Four-Light system, the Reflectance Transformation Imaging (RTI) system, a soft-box, a laser scanning system using a commercial laser-scanner software, and the diffuse system that calculates color information from the average of all four diffuse albedo images captured by the four-light system [Berns et al., 2012; Chen and Berns, 2012a; Chen and Berns, 2012c; Chen, 2013]. Images created using six of these methods are shown in Figure 2.16. (Images from laser scanning and the diffuse system color calculation are not shown.)

Paired comparison experiments performed as part of Chen's research identified the four-light system as one of the top performing system for the reproduction of total appearance and texture reproduction of samples and thus has influenced the choice to improve this system [Chen, 2013].




Fig. 2.16. – Example images using the eight different systems to image the same painting [Chen, 2013].

Deriving Surface Normal Maps From Traditional Photographs

3.1 Building on an Existing Foundation

The expectation of the aforementioned Berns, et al. method of Four Lights Imaging, 4LI, was that the four images could be used to render images equivalent to conventional imaging without polarizers so that a single set of images could be used for computer graphics and reprographics. Unfortunately, cross-polarization resulted in images that appeared blurry, due in part to the use of polarizers placed in front of the camera and a decrease in high frequency information [Berns et al., 2012]. Therefore, reproductions would require additional images with either linear or uncontrolled polarization. The need for images that are taken both with and without cross polarization increases the number of required images, the time required for imaging, and computational time by doubling the number of necessary images.

The use of polarization addressed the problems associated with highlights but did not directly eliminate the presence of shadows. Cross polarization was used to eliminate first surface reflectance including specular highlights. Similar to Rushmeier et al. [1997], a shadow was assumed to occur in only one light direction. That reduced the number of light directions to three since the signal was null. A psychophysical experiment was performed to compare techniques that directly eliminated shadows and those that did not. The results of this experiment showed that the presence of shadows did not negatively affect the appearance of a rendered painting [Chen, 2013]. This conclusion was used to inform the decision to continue using four lights instead using a laser scanning system to determine a true height map to be used for renderings.

3.2 Threshold Method Enhanced Four Light Imaging

The following is a modified excerpt from the SPIE conference proceeding "Imaging artwork in a studio environment for computer graphics rendering," published February 2015 [Cox and Berns, 2015].

In an effort to continue using only four lights but without polarizers, a method of thresholding was developed. This thresholding method was inspired by both the work of Rushmeier et al. and Sun et al. [Rushmeier et al., 1997; Sun et al., 2007]. The thresholding method, written mathematically in Equation 3.1, assumes that a highlight is only present in a given pixel location for one direction at a time.

$$\begin{array}{ll} \text{if } imgA(x, y) > avgImg(x, y) + avgImg(x, y) * k, & \\ \text{then } imgA(x, y) \text{ is a highlight} & \end{array} \quad (3.1)$$

In the thresholding equation, (x, y) denotes a pixel location, $imgA$ is a single image taken from one of the lighting directions, $avgImg$ is the calculated average without the maximum value of all four light directions, and k is a constant, accounting for any skewing that may have occurred if a pixel was in shadow. The value 0.75 was used for constant k , selected by visual evaluation. Pixels above the threshold were assumed to be highlights and removed, resulting in three light directions for calculating the surface normal. Images taken using the new method, referred to as 4LI Simplified or 4LIS, were used to calculate the diffuse color of a painting.

This new method was compared to the 4LI method with polarization proposed by Berns, et al. using the same samples [Berns and Chen, 2012]. The 4LI Simplified technique allowed for the generation of normal maps with more high-frequency detail than the normal maps generated from cross-polarized images, with the trade-off of less-smooth transitions between surface changes, especially where a highlight was captured. These tradeoffs were credited to the abrupt changes in pixel intensity that are only present in the highlight regions and not to the cross-polarized images. The 4LI Simplified method, with its unique averaging method, also generated diffuse normal maps that were sharper than their cross-polarized

counterparts, which is believed to be due to the increased level of contrast in the non-polarized images as compared to the polarized images shown in Figure 3.1. Normal maps of the unvarnished and varnished paintings derived from the 4LI and 4LIS methods, along with a difference map highlighting the difference, between the two techniques, are shown in Figure 3.2 and 3.3.



Fig. 3.1. – Diffuse image of the unvarnished (top) and varnished (bottom) paintings from cross-polarized (left) and traditional images (right). Diffuse color image of varnished painting from cross-polarized (bottom left) and traditional images (bottom right) [Cox and Berns, 2015].

3.3 Rendering Images using Threshold Data

Surface normal maps from both acrylic paintings along with diffuse color maps captured using the 4LI and 4LI Simplified method were imported into Maya, and mental ray's ray




Fig. 3.2. – Normal maps of the unvarnished painting. Left image used cross-polarized data to generate a normal map using four lights. The center image used thresholded data from three or four lights when appropriate. The right image is a difference map [Cox and Berns, 2015].




Fig. 3.3. – Normal maps of the varnished painting. Left image used cross-polarized data to generate a normal map using four lights. The center image used thresholded data from either three or four lights when appropriate. The right image is a difference map [Cox and Berns, 2015].

tracing option was used to render images of the paintings simulating a D65 spotlight forty five degrees to the right of the center of the painting and a camera, or observer, forty five degrees to the left of the painting. Each painting was assigned a *miamaterialX* shader. This shader is equipped with many preset material attributes for easy rendering. The virtual paintings were assigned properties of the preset matte plastic, and approximate index of refraction values of 1.40 (unvarnished) and 1.52 (varnished). Matte plastic was chosen because it was the closest preset material with characteristics similar to dry acrylic paint. The reflective glossiness property for the varnished painting was also changed to 0.600, compared to 0.500 for the unvarnished painting to account for the differences in gloss level between the two paintings.

A visual comparison of renderings (Figure 3.4) generated using Maya, composed of images from both techniques, showed a similar performance in rendering surface texture between the new approach and previous one. However, the images from the new approach generated overall sharper renderings of both varnished and unvarnished acrylic paintings without the addition of extraneous equipment or training.




Fig. 3.4. – Rendered images of the unvarnished painting from cross-polarized normal and diffuse maps (far left), from traditional/thresholded normal and diffuse color maps (left-middle), the varnished painting from cross-polarized normal and diffuse maps (right-middle), and from traditional/thresholded normal and diffuse color maps (right) [Cox and Berns, 2015].

These examples and comparisons indicate that the thresholding method is a successful simplification to the 4LI technique. As such the 4LIS technique will be used for all artworks used in this research moving forward.

Lighting for Scientific Imaging

Lighting fine art for image capture is a complex task that requires a great deal of skill and attention to detail. Different combinations of lighting produce varying effects of gloss and shadowing. For example, raking lighting produces long dark shadows to accentuate the surface structure of a painting. Different institutions choose to light works differently based on the outcome they prefer for representing archiving data and images. In the case of the the Museum of Modern Art in New York, New York, a more diffuse lighting configuration is created to image artwork. An example of their set up is shown in Figure 4.1 (left) where images were taken with a calibrated Hasselblad camera and two Broncolor strobes above the painting. Two strobes pointed vertically at the ceiling and a diffuse white board positioned below the painting were used to provide lighting for each shot. Images were taken with both strobes firing simultaneously. Lighting fine art in this fashion produces an image like that shown in Figure 4.1 (right).

Lighting artwork to create an aesthetically pleasing image is often very different from lighting artwork for scientific imaging. Scientific Imaging can include many different techniques such as using ultraviolet or infrared illumination, radiography (X-ray), and the previously mentioned raking light [Webber, 2008]. A section of an oil painting imaged with raking light (right) and normal lighting (left) is shown in Figure 4.2 [Webber, 2008]. The image taken with raking light exhibits details in surface distortions such as cracks and impressions while the image taken with normal lighting shows minimal cracking. All of these techniques provide different information. For example, ultraviolet, or UV, radiation can provide information about the type of varnish used on a painting based on what color the surface fluoresces [Webber, 2008].

4.1 Background

The 4LIS system focuses on providing information about the texture of a painting. The system is based on Woodham's photometric stereo technique and as such, the positioning of




Fig. 4.1. – Left: Lighting setup used by the Department of Imaging and Visual Resources at MoMA using two strobes (highlighted with red circles) directed toward the ceiling from above the paintings. Right: An example of fine art lighting used by the Museum of Modern Art in Queens, New York on Pablo Picasso's Vase of Flowers.



Fig. 4.2. – A section of an oil painting imaged with two different types of lighting normal(left) and raking (right) [Webber, 2008].

the light sources is a vital part of achieving successful normal maps. When using Broncolor strobes with a P70 reflector, as is the case of imaging done using the 4LIS technique, the lights are "centered" by aiming the strobe significantly off axis of center. An example of aiming the left and right lights is shown in Figure 4.3a. By aiming the lights off-center, a circular area radiating from the center of the image plane outward is created with the area of highest illumination in the center of the field. It has been hypothesized that the off-axis alignment of lights creates a hot-spot in the center of the field due to the more collimated light that is being reflected from the outside edges of the P70 reflectors to the center, an example of which is shown in Figure 4.3b. Because the strobes cannot be simply aimed towards the center of a painting or background, a significant amount of time is spent adjusting the lights to achieve the appropriate illumination.

4.2 The Addition of Honeycomb Grids

The imaging and calculation of normal maps of paintings with dark or glossy regions has been improved by using the threshold based averaging method when compared to the use of polarizers [Cox and Berns, 2015]. However, these areas continue to be difficult to accurately represent, leaving holes, or an absence of physical information, in the calculated normal maps. An example of a normal map with holes generated from a data set containing a lack of pixel information in high gloss and dark areas is shown in Figure 4.4.

Honeycomb grids have been added to the 4LIS system to investigate how their addition affects both the aiming process of each strobe as well as the calculation of normal maps for challenging works of art.

Honeycomb grids are used to control the beam spread of the light source to which they are attached. The grids are generally made up of a metal lattice with openings of different sizes and grid thicknesses. The different configurations of the honeycomb pattern control the beam spread to a specific amount. Four different honeycomb grids were used in the 4LIS imaging process. These grids focused the light to five, ten, twenty, and thirty degrees solid angle. An example of the ten, twenty, and thirty degree grids can be seen in Figure 4.5. The largest openings in the grid pattern correspond to the thirty-degree focus area.




Fig. 4.3. – a) Aiming strobes to achieve illumination centered in the field of view of the camera requires an off-axis alignment shown by the solid black lines for the left and right strobes. Dotted-gray lines show what would be considered on-axis lighting. b) An example of a Broncolor xenon strobe with P70 reflector attached [Broncolor, Bron Elektronik AG, 2015].




Fig. 4.4. – a) Normal map with holes, or dark areas where not enough information was available to accurately calculate the surface normal.




Fig. 4.5. – Honeycomb grids that control the beam spread of the source to ten, twenty, and thirty degrees [B & H Foto & Electronics Corp, 2015].

Images were first taken to visualize the falloff of a light source when using different grids of varying degrees of focus. Each image was taken at the same distance from a uniform white background using the same light source set at a fixed distance from the background and camera.

Due to the nature of each grid and the strobe, the strobes were redirected to center the light in the field of view of the camera. An image of the white background was taken without the use of a honeycomb grid as a reference image. To quantify the falloff from the center of each image to the edge, digital counts were collected. These digital counts are displayed on the falloff visualizations in Figure 4.6. The ideal digital count value in the center of each image is around 15000. This value ensures proper exposure while leaving room for highlights and glossy areas to be captured appropriately.

The image representing the setup that does not use a grid has the least amount of falloff from the center image to the edge with a twenty five percent falloff from 15118 to 11300 digital counts. The image taken with the five degree honeycomb grid, however, has the most falloff of ninety six percent. The maximum digital count, minimum digital count, and percent falloff for each grid is recorded in Table 4.1.




Fig. 4.6. – Falloff images of a white background using (from top left to bottom right) no grid, a 30-degree, 20-degree, 10-degree, and 5-degree grid.

Tab. 4.1. – Percent falloff produced by each honeycomb grid and the reference.

Grid	5°	10°	20°	30°	No Grid
Center DC	14167	14412	14244	14553	15118
Edge DC	559	1769	5852	8498	11300
Percent Falloff	96%	88%	59%	42%	25%

The five and ten degree grids are the easiest to aim to achieve centered lighting because they are the most focused, however the falloff for these grids is very large and poses some concerns for proper exposure of scenes that fill the frame.

4.3 Honeycombs and Normal Maps

Due to the motivation of examining how focusing the light effects the normal maps, the thirty-degree grid was not used in the preliminary investigations because it most closely resembled the original setup that did not use a grid during imaging. The decision to concentrate on the smaller degree grids was based on the hypothesis that the more focused the light, the less chance there is for multiple highlights to be generated in the same pixel local from different lighting directions.

Images were taken of a painting with some gloss and heavy impasto. These images also included a black, glossy cue ball, and ColorChecker® Classic to minimize the number of images necessary to perform the 4LIS technique. The scene that was imaged is shown in Figure 4.7.

Images taken of this scene were used in the calculation of normal maps for lighting configurations using five-, ten-, and twenty-degree grids along with no grid as a reference normal map. These normal maps are shown in Figure 4.8, cropped to show the region containing the painting. A close up section of the bottom left region of each normal map is shown in Figure 4.9.

Normal maps generated from images taken with the five and ten-degree grids appear to be the most different from the normal map created without a grid. These differences were explored further by creating difference images.

4.3.1 Normal Map Difference Images

A difference image was created by layering two normal maps and then subtracting each channel from the corresponding channel of the second normal map. The difference image between the twenty-degree honeycomb and the standard normal map is black indicating that there is no perceptible difference between the two normal maps. The difference images comparing the ten degree and five degree normal maps to the normal map created without the use of any grid are darker on the right than the left indicating that the normal maps were not uniformly affected by the use of these honeycomb grids. Examination of these normal maps in Figure 4.10 confirms this observation due to the left side of each of the ten and five degree maps being darker and less blue compared to the light bluish color comprising most of the reference normal map. After examining the difference maps and the fall off images, the five and ten-degree honeycombs were eliminated from consideration due in part to their non-uniformity. Although these images are flat-fielded, the large amount of fall-off in the five and ten-degree images generated such low signals that flat-fielding was not enough make the images viable.



Fig. 4.7. – Scene used to compare 4LIS imaging with different honeycomb grids. The scene includes all necessary elements (cue ball, color chart, and painting) to facilitate efficient imaging for all honeycombs.




Fig. 4.8. – Normal maps resulting from using different grids. From top left to bottom left (clockwise) 20 degrees, 10 degrees, 5 degrees, 4LIS without a grid.




Fig. 4.9. – Close-up section from the bottom left of each normal map using no grid and the 5 to 20 degree grids.




Fig. 4.10. – Difference images, comparing normal maps to normal maps generated from the use of other grids.

4.4 Further Investigations of Honeycomb Use

An imaging sequence that used the twenty-degree honeycomb and the thirty-degree honeycomb was implemented to further investigate the use of honeycombs in conjunction with the 4LIS system. Each honeycomb grid was used to image a Van Gogh style painting featuring a dark, textured pine tree shown in Figure 4.11. This painting was selected due to the limitations of the 4LIS imaging revealed while imaging dark and glossy paintings previously. It was hypothesized that holes in the normal maps are due to having highlights in more than one image at the same pixel location or due to a lack of information being captured in the dark areas of the artwork because these holes do not appear in other regions of this painting's normal map.



Fig. 4.11. – Van Gogh style painting imaged using both the twenty- and thirty- degree grids.

Although the thirty-degree honeycomb was eliminated in the preliminary experiment, it was reintroduced after the five-, and ten-degree grids were eliminated. Therefore, images of the Van Gogh style painting, were taken using both the twenty- and thirty-degree grids, while

keeping the imaging process and distance between the lights and painting and the distance between the painting and camera as consistent as possible. Images of the normal maps are shown in Figures 4.12 and 4.13 for the twenty- and thirty-degree grids respectively, with a close up of the pine tree from both the twenty- and thirty-degree images shown in Figure 4.14.



Fig. 4.12. – Normal map of a Van Gogh style painting created from images taken with a thirty-degree honeycomb grid.

The thirty-degree images show higher contrast in the detail of each brush stroke and texture; however there appears to be no noticeable difference in the number of holes in each normal map. The difference image of these two normal maps is predominately dark, with no clear pattern of green or red pixels. This small difference indicates that there is no improvement of one normal map over another.

4.4.1 Rendered Images and Color Differences

Two different images were rendered using the software ArtViewer [Berns and Chen, 2012], which was developed at the Munsell Color Science Laboratory at Rochester Institute of



Fig. 4.13. – Normal map of a Van Gogh style painting created from images taken with a twenty-degree honeycomb grid.

Technology based on the RTIviewer [Cultural Heritage Imaging, 2013]. The images were rendered using normal maps from the twenty and thirty degree images separately. A cropped section showing a section of the pine tree of each rendering is shown in Figure 4.15. Areas where the rendered images contain dark spots are circled in red in both the twenty- and thirty-degree rendered sections. These black spots are areas where holes were present in the normal map image. The rendering from the twenty-degree image has more dark areas than the thirty-degree image. Both images are affected by holes in their respective normal maps. However, after careful examination of these renderings, the twenty-degree images appears to be slightly more affected.

To further investigate the advantages of using either the twenty-degree or thirty-degree grid over the other, the color difference between color corrected diffuse images was evaluated between the twenty-and thirty-degree images at the standard distance from the camera that was used to take images associated with the normal maps previously mentioned as well as images taken from a distance roughly twice as long. The diffuse color images were calculated and corrected using the method described by Cox and Berns [Cox and




Fig. 4.14. – Normal map of a Van Gogh style painting cropped to a small section of a dark glossy region (pine tree) using a 20 degree honeycomb (left) and a 30 degree honeycomb (right).

Berns, 2015]. A vector plot showing color differences compared to data measured using a spectrophotometer for both the twenty-and thirty-degree images taken from the closer distance are shown in Figures 4.16 and 4.17 respectively. The color difference for each patch along with the maximum, minimum, mean, and ninetieth percentile are shown in Tables 4.2, 4.3, 4.4, and 4.5 for both the short and long distances.

Tab. 4.2. – Color difference (CIEDE2000) of images taken with the 20 and 30 degree grids compared to the spectrophotometer data for each patch for a close distance. Larger color differences are denoted by bold text.

Grid	20/30	20/30	20/30	20/30	20/30	20/30
Row 1	0.69 /0.67	1.90/1.90	0.55/ 0.71	1.71/ 1.83	1.72/ 1.75	1.42 /1.22
Row 2	1.35/ 1.38	1.30 /1.25	0.90 /0.89	1.12/ 1.25	1.19 /0.91	0.99/ 1.26
Row 3	0.67 /0.58	1.90 /1.89	1.89 /1.87	0.83/ 1.09	1.40/ 1.45	1.26/1.26
Row 4	0.92/ 1.17	1.49/ 1.67	1.26/ 1.46	0.41/ 0.63	0.79/ 0.95	0.10 /0.00



Fig. 4.15. – 20 degree (left) and thirty degree (right) rendered sections of the dark pine tree in a Van Gogh Style painting, with areas containing "holes" circled in red.




Fig. 4.16. – Color difference plots showing the difference between measured data (circles) and calculated data (tip of each arrow) for the twenty-degree grid.

Tab. 4.3. – Minimum, maximum, average, and ninetieth percentile color difference values for the twenty- and thirty-degree images.

	Min DE00	Max DE00	Avg DE00	90th percentile
Grid 20/30	0.10/0.00	1.90/1.90	1.16/ 1.21	1.84/ 1.86




Fig. 4.17. – Color difference plots showing the difference between measured data (circles) and calculated data (tip of each arrow) for the thirty-degree grid.

Tab. 4.4. – Color difference (CIEDE2000) of images taken with the 20 and 30-degree grids compared to the spectrophotometer data for each patch for a long distance. Higher color differences are denoted by bold text.

Grid	20/30	20/30	20/30	20/30	20/30	20/30
Row 1	0.81 /0.71	2.07 /1.88	0.52/ 0.68	1.60 / 1.77	1.72 / 1.74	1.35 /1.15
Row 2	1.09/ 1.42	1.51 /1.49	1.07 /0.94	1.40 / 1.45	1.45 / 1.54	0.83/ 0.95
Row 3	0.6/ 0.79	1.20/1.89	2.07 /1.89	0.60/ 0.61	1.81 /1.48	0.84/ 1.43
Row4	1.22/ 1.44	1.69 /1.68	1.40 /1.35	0.33/0.33	1.01/ 1.09	0.01/ 0.03

Tab. 4.5. – Minimum, maximum, average, and ninetieth percentile color difference values for the twenty- and thirty-degree images for the long distance.

	Min DE00	Max DE00	Avg DE00	90th percentile
Grid 20/30	0.33/0.33	2.07 /1.89	1.21/ 1.24	1.99 /1.85

Although it appears that the twenty-degree images out perform the thirty-degree images at both distances the difference between the average color differences, 1.16 and 1.21 for the close distance and 1.21 and 1.24 for the long distance, are so small that the performance of the twenty-degree grid is not enough to edge out the thirty-degree images. To visualize how small these color differences are an sRGB image showing the measured color values for each patch (in the center) compared to all four calculated colors for each patch, one for each distance and grid is shown in Figure 4.18.




Fig. 4.18. – Color differences shown in an sRGB rendered image showing the measured data in the center surround by data calculated from images taken (from top left and clockwise around) at short distance with a 20-degree grid, 30-degree grid, then a long distance with a 20-degree grid and 30-degree grid.

4.5 Conclusion

The use of honeycomb grids of various degrees of focus were tested in conjunction with the Four Lights-Simplified technique. Using a grid of any size simplified the aiming process compared to aiming the lights to achieve centered illumination without any grid. The smaller the area of focus the easier it became to aim the strobes, however the five and ten degree honeycombs were eliminated as viable choices early on due to the large amount of falloff

they create, ninety-six and eighty-eight percent respectively, as well as affecting the normal maps they generated non-uniformly from side to side as well as top to bottom.

The investigation into honeycomb use continued with a more detailed look into the use of the twenty- and thirty-degree grids. A quick look at the normal maps from both the twenty-degree and thirty degree images did not show a clear advantage to using one grid over the other. A closer look at the differences between the two normal maps indicated that the differences did not indicate a large difference or clear pattern but appeared to be more like noise. Rendered images created from each normal map were used to determine that the thirty-degree rendering was less affected by holes in the normal map than the twenty-degree rendering. When color differences were compared for both grids, the twenty-degree grid did not significantly out perform the thirty-degree honeycomb. With all of these findings considered in addition to the falloff being less for the thirty-degree grid (42 percent versus 59 percent), the thirty-degree honeycomb grid will be used with the 4LIS system. Although both grids would be acceptable choices, the twenty-degree honeycombs are easier to aim and have a small but not significantly better color difference performance, but the thirty-degree honeycomb has a smaller number of holes in the corresponding normal maps and a smaller falloff percentage, which is useful for imaging samples that fill the frame.

Combining Color and Surface Normal Maps in a Rendering Software

Although surface normals and diffuse color information provide important details about the appearance of fine art, combining the two in a rendering software can provide an meaningful experience to observers. One way to combine this information is use a rendering software. Using a rendering software also allows the user to provide information about the surface gloss of the artwork. The rendering software, Maya, was selected due to the large amount of support material available, its name recognition, as well as its availability for educational purposes at low cost.

5.1 Introduction to Maya and mental ray

Maya is an animation, modeling, simulation, and rendering software [Autodesk, Inc., 2015b]. The software includes a general-purpose renderer called mental ray. Mental ray uses rendering techniques such as a scanline rendering algorithm for primary visible surface determination and binary space partitioning algorithms to render ray tracing effects. Mental ray supports physically correct simulation of general global illumination. Supported simulation includes any combination of diffuse, glossy, and specular reflection and transmission [Autodesk, Inc., 2015a].

The mental ray software uses multiple shaders, or algorithms that specify how a surface reacts to light, to create geometric features like bump and displacements maps, (grayscale textures mapped to a surface), to create an illusion of surface relief, and light sources. These shaders provide parameters that can be manipulated to adjust material properties, ambient lighting, diffuse and specular color, and reflectivity [Autodesk, Inc., 2015a].

Maya with the mental ray plug-in, here after referred to as Maya, is one of the most popular commercial products for producing computer graphics (CG) renderings. This software was used in conjunction with the diffuse color and normal maps generated from measuring the artifacts as part of the rendering process. The introduction of gloss, or micro-structure, information comes from either selecting a reference material or a model associated with a predetermined shader, such as a Phong, Anisotropic, or Blinn shader in Maya [Autodesk, Inc., 2015c].

5.2 Rendering with Maya

Rendering realistic images of two-dimensional fine art requires information about diffuse color, macrostructure defined by a surface normal or height map, and microstructure defined by a bidirectional reflection distribution function (BRDF). All three of these characteristics can be calculated or measured in different ways as previously discribed.

Diffuse color and macrostructure can be derived from using the 4LIS technique. Although it is beneficial to know the BRDF of an object, it is not necessary when using CG. The software is generally capable of approximating the way light will react when it comes in contact with the surface without knowing its specific BRDF. The limitation is when the artwork has appreciable spatially varying microstructure.

Using the 4LIS method to capture the Van Gogh style painting results in the diffuse and normal map images shown in Figure 5.1. As a diffuse image, macrostructure is not readily seen. The normal map contains the macrostructure information. The resulting diffuse color and surface-normal maps were combined in Maya and then rendered with a user supplied lighting and viewing combination, as well as surface material attribute settings. The workspace, where a painting is illuminated from above with a directional light and from below with a low intensity and wide angle fill light created from Maya's spotlight option, is shown in Figure 5.2.

The surface representing the painting in the workspace was assigned the Blinn shader. There is an extensive list of shaders that can be used in Maya. Each shader has a unique combination of attributes that can be changed to achieve a variety of effects. This section



Fig. 5.1. – Diffuse color (left) and normal (right) images of a van Gogh style painting resulting from the simplified four-light method.

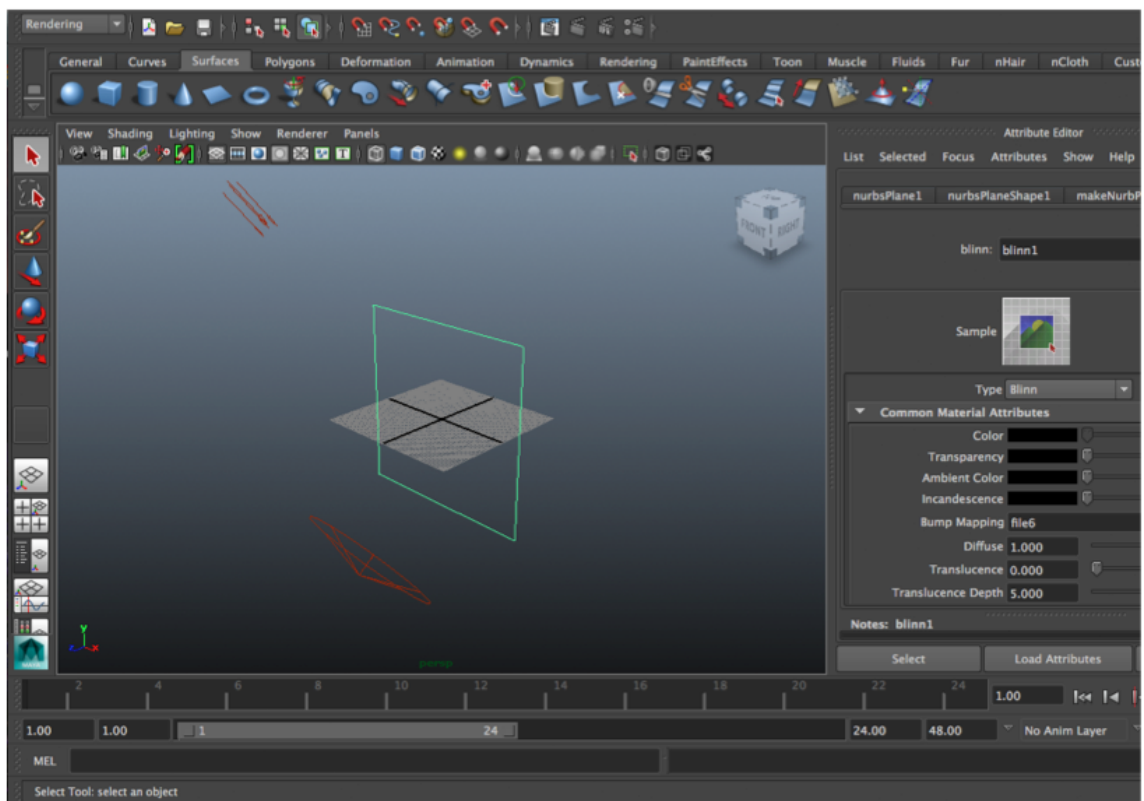


Fig. 5.2. – Maya workspace showing a painting illuminated with two different lights, a directional light from above and a diffuse fill light from below.

used the Blinn shader while other shaders are used in the following chapters. Shader choice is often based on the preference of the user. Many shaders can achieve the same or similar effects. This shader allows control of the size of highlights and how the surface reflects its surroundings. This is done by adjusting the values for *Eccentricity* and *Specular Roll Off*, respectively [Autodesk, Inc., 2015c]. The Blinn shader also has parameters that are common to many of the other surface shaders. These parameters include *Specular Color*, *Diffuse*, and *Reflectivity* [Autodesk, Inc., 2014b; Autodesk, Inc., 2014c].

The *Specular Color* parameter represents the color of highlights on the surface of an object. *Reflectivity* represents an object's ability to reflect its surroundings where a value of 0 represents no reflections and 1 represents clear reflections. For example a value of 1 would be used for chrome or mirror materials. *Diffuse* represents a value that determines the object's ability to reflect light. The higher the *Diffuse* value the closer to the (user defined) surface color the object will appear. A value of 0 will result in no light being reflected. These are only a few of the parameters that can be manipulated as part of the Blinn shader.

The rendered image shown in Figure 5.3 (left) was rendered with user defined parameter settings starting with the default settings for the Blinn shader. These parameters were adjusted until the rendered images appropriately emphasized the glossy nature of the painting. The parameters that were changed were *Diffuse*, *Eccentricity*, *Specular Roll Off*, *Specular Color*, and *Reflectivity*. These parameters were changed to 1, 0, 0, White, and 0 respectively. Parameter settings are given in Table 5.1.

Tab. 5.1. – Blinn Shader parameter settings for each rendering.

Parameter	Matte Painting	Glossy Painting	Specularly Enhanced
Diffuse	1	1	0.3
Eccentricity	0	0.152	0.3
Specular Roll Off	0	0.545	1
Specular Color	White	White	White
Reflectivity	0	0.4	0.3

The image was re-rendered to simulate the addition of a glossy picture varnish across the surface of the painting. The glossier rendering was obtained by adjusting the parameter values for *Eccentricity* from 0.000 to 0.152, *Specular Roll Off* from 0.000 to 0.545, and *Reflectivity* from 0.000 to 0.400. The varnish-simulated image is shown in Figure 5.3 (right).

The "varnished" painting exhibits larger and more pronounced highlights compared to the original rendering, which are most noticeable in the top portion of the painting due to the proximity of the overhead light source.



Fig. 5.3. – Rendered images of the painting using Maya to simulate directional lighting from above and fill lighting from below with its matte, existing varnish (left) and simulated picture varnish (right).

To display what varying the lighting would produce without changing any physical parameters of the painting, a series of close ups are shown in Figure 5.4 using matte parameter settings with the lighting direction changing from top left, to top center, to top far right. The shadows cast from each different lighting configuration change accordingly as evident in these images. Maya is capable of rendering images with an endless combination of lighting configurations from just a diffuse color map, normal map, and appropriate shader or shaders.

Not only is this CG rendering process capable of changing material properties to simulate a matte or glossy surface, it is also capable of rendering images with different lighting configurations to achieve effects similar to raking light imaging.

By changing the configuration of lighting shown in Figure 5.2 to include more grazing angle illumination with high intensity light coming from above the painting and two low intensity fill




Fig. 5.4. – Rendered image of a close up section of a painting with directional lighting coming from top left (top left), top center (top right), and top right (bottom) generated using Maya and mental ray.

lights to provide ambient lighting, a rendering with a similar appearance to raking lighting is achieved. This rendering is shown in Figure 5.5 (left). More pronounced shadows occurs with this type of lighting than in the previous renderings.

Specular enhancement renderings are also possible with this CG rendering process. By using similar lighting to that of raking lighting and by changing a few key attributes, a metal-like appearance is achieved. The benefit of these types of renderings is to emphasize the painting's surface structure [Cultural Heritage Imaging, 2013]. The attributes that were changed to achieve the specular enhancement effect were the same ones that were changed to make the painting appear varnished, however they were changed to different values. A detailed investigation into using Maya to represent real objects is described in Appendix B.

The *Diffuse* value was changed to 0.4 to reduce the effect of the surface color on the rendering. *Eccentricity* was changed to 0.3, to increase the size of the highlights. *Specular Roll Off* was changed to 1, to increase how well the surface reflects the light in the scene; and *Reflectivity* was changed to 0.3 to increase the painting's overall reflectivity. A rendering of the specular enhanced painting is shown in Figure 5.5 (right). A list of parameters and their settings for each rendering is given in Table 5.1.




Fig. 5.5. – Rendered image of a painting with simulated raking light (left) and specular enhancement (right).

Maya was used to generate images of a painting from a static lighting direction, with a lighting direction that moves across the surface of the painting, and with raking light. These rendered images accentuate the texture in the painting in addition to emphasizing the reflective properties by showing how the moving light changes the areas that are highlighted or in shadow without having to reshoot the primary images.

This CG rendering process and software combination was also used to simulate changes in surface properties from a matte surface to a glossier and even metallic-like surface. By selectively changing a few parameters surface material properties were changed without affecting the surface texture defined by the painting's normal map.

5.3 Rendering Spatially Varying Gloss

A significant challenge in visually representing fine art is the ability to fully replicate the appearance of the work. Static images are typically generated using a single camera and lighting geometry. Due to this limiting setup, texture and gloss information is often lost in these images, especially when gloss changes across the surface of the work. Using a rendering software allows the possibility of multiple viewing or lighting geometries. A process to render a mostly two-dimensional cultural heritage object with SVBRDF is described below.

5.3.1 Approach

A section of a commode, attributed to Joseph Baumhauer, was imaged using the Four Lights Simplified (4LIS) technique. The piece is a part of the Sculpture and Decorative Arts Department of the J. Paul Getty Museum Collection. The panel is oak veneered with ebony, set with panels of Japanese lacquer on Japanese arborvitae with vernis Martin (imitation lacquer), with gilt-bronze mounts [*Commode*]. Four different images were taken, each using a different lighting direction. An example of one of these images is shown in Figure 5.6 (left). The panel has an overall curvature, therefore a smaller section of the images were used for rendering. This flat section is outlined in red in Figure 5.6.



Fig. 5.6. – A section of furniture attributed to Joseph Baumhauer belonging to the Sculpture and Decorative Arts Department of the J. Paul Getty Museum Collection. (left) The diffuse color map (middle) and normal map (right) of a flat section of the piece indicated by the red box in the left image.

The determined color map and surface normal map are shown in Figure 5.6 (middle, right). Because the materials used to create this piece have different gloss properties, masks were created to separate the color map into three different images by using Photoshop to select areas of similar material and gloss. These images are shown in Figure 5.7. Each image consists of a color region with the same gloss level, and a transparent region. There is a black lacquer image, matte gold and red image comprised mostly of facial features and accents, and the gold image.



Fig. 5.7. – Masks to separate the color map into images of three different material types: black lacquer (left), high gloss gold (middle), and matte gold and red accents (right).

The color maps and normal maps are loaded into Maya using a layered shader. Using a layered shader allows the user to define different shaders to different elements of the object. In the case of the furniture panel, three different shaders are used, one for each color mask. A single normal map is used for the entire object.

5.3.2 Results

By using the layered shader technique, each material in the piece can be rendered with a different effect. The gold gilded regions will appear the most glossy from using a shader and settings to mimic the appearance of a metallic surface, while the black lacquered and accent areas will appear less glossy with settings used to approximate a lacquered material. Images of the rendered object are shown in Figure 5.8 from four different lighting directions. The lighting starts from the left of the panel and moves in an arc about its center to the

right with the light ending to the far right at a nearly grazing angle. The top portion of the renderings use three different shaders for each material while the bottom portion uses only one shader with the same settings for all materials. Although subtle, differences in the gilded areas of the top and bottom portions are clearly visible with particular focus on the leg and chest of the middle horse (circled in the left image).



Fig. 5.8. – Images of the rendered panel with lighting from four different directions showing how the gold areas change differently from the black lacquered areas. Three different shaders were used to render the different materials in the top half of the renderings. The renderings in the bottom half used only one.

5.3.3 Conclusion

The rendered images and outlined procedure demonstrate how Maya can be used to display paintings of different, uniform gloss levels for a virtual gallery with a variable light source. The location, angle, and type of lighting can be altered to accentuate different surface properties of a painting in accordance with each museum or institution's lighting preferences considered. The images generated are capable of showing impasto, highlights, and shadows that correspond and react to the lighting chosen for each rendering. Multiple images with varying lighting directions can also be rendered and combined to produce a short video of how light interacts with the surface of the painting as an added feature to the virtual gallery. Although, ArtViewer is capable of creating renderings of paintings that do not have a complex surface structure or coating, Maya is able to render images with higher amounts of gloss more completely because it allows multiple bounces of photons off of surfaces whereas ArtViewer only uses a single bounce to calculate renderings.

By creating layers that correspond to different materials used in the piece, renderings can be created that depict the way light would interact with these different materials in the

same rendering. The effects of multiple materials being rendered simultaneously can be further emphasized by rendering scenes with multiple lighting configurations to accent the differences in gloss level, as shown in Figure 5.8. A single still image, like the one shown in Figure 5.6, does not afford as much information about gloss and material composition as a sequence of renderings that separate individual materials and renders them with multiple lighting geometries.

Observer Preferences and Digital Fine Art Reproduction

Fine art is presented in a variety of ways to observers. Not all observers find themselves in a position to view an original work in a museum under carefully thought-out lighting. Often, observers are viewing reproductions. These representations include prints, book figures, photographs, digital renderings, and high resolution images. All of these substitutions are incapable of recreating the viewing experience perfectly; however, some steps can be taken to improve how artwork is reproduced. An important part of creating an appropriate reproduction is selecting the optimal lighting configuration. Different lighting configurations create diverse material appearances based on the type of artwork. The choice of lighting is explored with a focus on texture and gloss appearance.

6.1 Experiment 3: Creating Stimuli

A set of experimental tasks was designed and implemented to explore gloss and texture appearance as perceived by observers. This process began with the creation of representative image stimuli. Images were created of eight different samples spanning different representative materials used in fine art. Each sample was imaged using a high resolution camera and a strobe placed at different angles from the center of the sample. Thirteen different images were taken with the angle of lighting ranging from 6 to 88 degrees from the surface normal of the painting. The angles created with the strobe and the samples for each image are shown in Table 6.1. The thirteenth image, taken at 6 degrees, of each sample was excluded from the experimental data set because it displayed such a large amount of gloss for varnished materials that it completely eliminated any information about appearance of the painting. The change in lighting angle between each image is, on average, seven degrees. This is considered to be a finer sampling than a typical gallery visitor would move about a painting on display. The first two images also employed the use of a fill light. This fill light was angled from the opposite side of the paintings and used a diffuser. It was used

to achieve proper exposure for these two extreme angles. A schematic showing the strobe placement in conjunction with the paintings and camera is shown in Figure 6.1 where the angles that used a fill light are outlined in red.

Tab. 6.1. – Recorded angles for the light location from the surface normal of the painting for each image taken.

Index	Angle	Index	Angle	Index	Angle	Index	Angle
Img_88	88°	Img_67	67°	Img_42	42°	Img_23	23°
Img_82	82°	Img_59	59°	Img_34	34°	Img_17	17°
Img_77	77°	Img_52	52°	Img_28	28°	Img_12	12°




Fig. 6.1. – Schematic of lighting angle used to image each sample. Strobes outlined in red (angles 82 and 88) fired at the same time as a diffuse fill light from the opposite side of the painting.

Each sample was imaged with a strobe at the same twelve angles to keep consistency between sample sets. Images were flat-fielded and minimally sharpened. Three different experiments were performed using images from this dataset; one experiment focused on texture, one on gloss, and one on aesthetics. The eight different samples that were imaged are shown in Figure 6.2. These samples include (from top left to bottom right) an unvarnished gray patch of acrylic paint, an unvarnished acrylic painting, a varnished acrylic

painting, a matte varnished gray patch of oil paint, matte print, woven tapestry, and two varnished oil paintings. An example of a sample, in this case Flowers, imaged using each of the 12 different lighting directions is shown in Figure 6.3.



Fig. 6.2. – Images of each reference stimuli of (from top left to bottom right) an unvarnished gray patch of acrylic paint, an unvarnished acrylic painting, a varnished acrylic painting, a matte varnished gray patch of oil paint, matte print, woven tapestry, and two varnished oil paintings.




Fig. 6.3. – Flowers sample showing each image used in the experiment starting with the largest lighting angle in the top left and the smallest angle in the bottom right.

6.2 Experiment 3: Procedure

6.2.1 Texture Perception

Observers were asked to examine three images presented simultaneously, shown in Figure 6.4. The top two images were considered anchor images (reference images). These images were assigned the values 30 and 75. The left anchor image, assigned the value of 30 had a low amount of perceived texture. The right anchor image, assigned the value of 75, had a large amount of perceived texture. The third image was considered the test image. The values of 30 and 75 were selected because they create a range that is not evenly divided by 2 and also leave room for observers to make ratings above or below the range while still being with a 0 to 100 range [S. Farnand, personal communication, January 2017]. The observers were asked to assign the test image a value that represented how much texture

they perceived in this image compared to the anchor images. The observer could assign any positive integer value above, between, below, and including the anchor image values. For example, if the test image appears to have less texture than the high anchor image but more texture than the low anchor image, the observer might assign the test image a value of 50.

Each observer was asked to repeat this task for every image in the sample set and for multiple samples. This resulted in 96 separate texture assessments per observer. The experiment was presented to observers in a darkened room on an Apple iMac 27-inch display that had been colorimetrically calibrated using the Day model based on work done by Fairchild and Wyble [Day et al., 2004; Fairchild and Wyble, 1998]. Each image was color corrected using the procedure outlined in section 2.1.1. Observers were shown test images in a randomized order without any prior knowledge or viewing of the physical sample each image represented. An example of the experiment interface is shown in Figure 6.4. Observers were read the following script in preparation of performing the experiment.

"You will be shown three images. Two of the images are anchor images. The low anchor image is assigned the value of 30. This image represents the lower texture range of the images. The high anchor image is assigned the value of 75. This image represents the higher texture range of the images. You will be shown the third image and be asked to assign a value to the this image that represents the amount of texture in the image compared to the anchor images, keeping in mind that the value can be less than, greater than, or the same as the anchor image values."

6.2.2 Gloss Perception

The second task of the experiment was very similar to the first task where observers were asked to examine three images presented simultaneously where the top two images were considered anchor images and the third image was the test image. The anchor images were assigned the same values of 30 and 75. The left anchor image in task two, was assigned to an image with a low amount of perceived gloss (instead of texture) and the right anchor image had a large amount of perceived gloss. Similarly, the observers were asked to assign the test image a value that represented how much gloss was perceived in this



Fig. 6.4. – Example image of the experiment interface showing three images of an acrylic painting, two anchor images and one test image.

image compared to the anchor images using any positive integer value above, between, below, and including the anchor image values. The observers were given the same set of instructions as before but with gloss replacing texture.

Each observer was asked to repeat this task for every image in the sample set and for multiple samples. This task included samples from the same set but excluded matte samples. This resulted in seventy two separate gloss assessments per observer. An example of the experiment interface is shown in Figure 6.5. The setup for the first and second tasks were nearly identical and used the same sample images with the exception of the images of the textile, print, and matte gray patch being excluded from the gloss judgements due to their matte appearance for all lighting geometries. Because of the similar nature of these two experiments observers were randomly assigned which task they completed first. The intentions of this randomization was to account for any influence one task would have on the other.

6.2.3 Observer Viewing Preference

The third task began with observers entering a gallery containing all of the artwork used in the first two tasks. This gallery was created by positioning the artwork on easels arranged



Fig. 6.5. – Example image of the experiment interface showing three images of a varnished acrylic painting, two anchor images and one test image.

in a row with overhead lighting generated from using three modeling lights. An image of the middle section of the gallery using this lighting is shown in Figure 6.6. Observers were instructed to view each piece in the gallery spending as much or as little time viewing each piece as they desired. The observers were informed that they could move freely about the space to examine the artwork as long as they did not touch the easels or artwork. The gallery was separated from the experiment area by a thick black curtain. Once the observer elected to leave the gallery area, a simple color vision test was administered by showing the observer a series of Ishihara Plates. After completing the color vision test, observers were shown a series of representative artist materials and were informed about each type of material. For example, observers were shown what effect different kinds of varnish had on the same underlying paint. The following script was read to each observer while pointing out each material as it was mentioned.

“Paintings and other fine art are made with many different materials and techniques. These include acrylic paint, oil paint, watercolor, ink, colored pencils, and markers. Heavy brush strokes, heavy impasto or texture, light/minimal texture, glossy varnish, matte varnish, and no varnish at all are all possible application or finishing methods. Fine art can also include textiles like tapestries with either coarse or fine yarn. All of these different materials influence the overall appearance of the work. Some works may be very glossy while others are very



Fig. 6.6. – A section of the gallery shown to observers using modeling lights located above the paintings and aimed towards the vertical center of the paintings on display.

matte based more on just what type of varnish was used. Not only does the application of materials influence the appearance of the work, so does the substrate they are applied to. Canvas, masonite, watercolor paper: all have their own texture and influence the overall look of the piece.”

After completing the artist material tutorial, observers were directed to the 27-inch Apple display to complete the third task. Observers were shown two images simultaneously. An example of the screen shown to observers as part of this task is shown in Figure 6.7. The image on the left side of the screen was considered the reference image and was given a value of 85. This image did not represent the ideal digital reproduction but an image with a medium amount of texture and gloss. The reference value of 85 was selected to leave room for observers to rate the test image higher without being subconsciously bound by a limit of 100.

The observers’ task was to assign a value to the second image that represented how much better or worse the second image represents the real object compared to the reference image. For example, if the second image did almost the same quality of job representing the original work and the materials that were used to create it, the observer would assign the



Fig. 6.7. – Example screenshot of the third task where a reference and test image were shown side by side and observers were asked to rate the right image’s appearance in comparison to the left image.

image a number close to 85. If the second image did not represent the sample nearly as well as the reference image, the observer would assign a much lower value than 85. Observers could assign any positive integer value below, above, or equal to 85. The observer was then re-shown every image in a different order and asked to repeat the experiment by assigning a value to the second image based on preference rather than which image best represented the original work. Preference was defined by asking the observer to make ratings based on how likely they would be to purchase a postcard or print of either the reference image or the test image.

After completing each task, observers were asked a series of questions regarding their observations. A list of questions asked is shown in Table 6.2. Observers were asked these questions in the hopes of providing insight into the ratings they made as well as to increase time between tasks to allow observers to adjust to the new experiments with minimal influence from previous tasks.

Tab. 6.2. – A list of questions asked to each observer after completing the experimental tasks. The intention of these questions is to provide insight into each observers ratings or selections.

Task	Question
1	What features were you looking at to determine how the gloss in the test image compared to the references?
	What features made it difficult to assess the gloss level?
2	What features were you looking at to determine how the texture in the test image compared to the references?
	What features made it difficult to assess the amount of texture?
3	What influenced your rating?
	What features were most important to you to be well represented?
	How did the images that best represent the materials and the real object compare with the images you preferred? Were you evaluating the images the same or differently for each question?

6.3 Experiment 3: Results

A total of 12 different observers, five female and seven male, participated in both the gloss and texture perception tasks. The observers ranged in age from 18 to 55. Observers were allowed to complete the tasks back to back or with hours and sometimes days between them. The type of task was randomized so that half of the observers rated gloss first and half of the observers rated texture first.

6.3.1 Gloss Perception

The average gloss perception rating was calculated for each of the five images. The averages for the 12 observers are plotted in Figure 6.8. Plots of all observer ratings and the average are shown in Figures 6.9 through 6.13. The average gloss ratings are also plotted separately with error bars in Figures 6.14, 6.15, 6.16, 6.17, and 6.18. These error bars were calculated by taking the standard deviation of observer ratings for each image.

An overall trend showing a decrease in gloss perception from 88 degrees to about 60 degrees followed by an increase in gloss perception as the lighting angle decreased is evident for all five samples. The increase in gloss perception from 60 to roughly 10 degrees follows what is known about measuring specular gloss. As the lighting angle approaches




Fig. 6.8. – Average gloss perception rating for each of the five artwork samples: Tulip, OilTrees, Gray Patch 1, Flowers, and DarkOil.

the equal but opposite measurement angle, an increase in measured gloss is expected. In the experimental setup used to image each sample, the camera was placed perpendicular to the sample surface. The equal but opposite angle to place the light source to maximize specular gloss is 0 degrees, or perpendicular to the surface of the sample.

The higher gloss rating at 88 degrees is hypothesized to be due to the influence of increased texture appearance as the angle of lighting increased moved closer to raking. This could be due to the edge of thick brush strokes or impasto catching the light from the strobe more than surrounding, or shadowed areas, and in-turn mimicking a highlight.

For many of these samples, there tended to be a wider range of ratings given for images taken at raking angles. One observer actually rated the gloss of the Gray Patch 1 sample highest at 88 degrees. It appears that observers tend to agree more closely with each other for mid range angles and disagree more frequently for the extreme angles, with a few




Fig. 6.9. – Average gloss perception rating for Flowers with standard deviation from the mean as error bars.




Fig. 6.10. – Average gloss perception rating for Gray Patch 1 with standard deviation from the mean as error bars.




Fig. 6.11. – Average gloss perception rating for DarkOil with standard deviation from the mean as error bars.




Fig. 6.12. – Average gloss perception rating for OilTrees with standard deviation from the mean as error bars.




Fig. 6.13. – Average gloss perception rating for Tulip with standard deviation from the mean as error bars.




Fig. 6.14. – Average gloss perception rating for Flowers with standard deviation from the mean as error bars.




Fig. 6.15. – Average gloss perception rating for Gray Patch 1 with standard deviation from the mean as error bars.




Fig. 6.16. – Average gloss perception rating for DarkOil with standard deviation from the mean as error bars.




Fig. 6.17. – Average gloss perception rating for OilTrees with standard deviation from the mean as error bars.




Fig. 6.18. – Average gloss perception rating for Tulip with standard deviation from the mean as error bars.

exceptions evident for the Tulip sample around 50 degrees as shown in Figures 6.9 through 6.13.

When observers were asked what made it difficult to assess the appearance of gloss in each image, nearly every observer mentioned texture in their list of factors. The influence of the amount of texture in each sample may be a source of variance in each sample's gloss rating. Figures 6.19 and 6.20 show images with similar gloss ratings but very different lighting configurations. These images emphasize how much texture influences the ratings selected by observers when making judgements about gloss.




Fig. 6.19. – Images of three of the five samples that have roughly the same gloss rating.



Fig. 6.20. – Images of two of the five gloss samples that have roughly the same gloss rating.

6.3.2 Texture Perception

The average texture perception rating was also calculated for each of the eight texture samples. The averages are plotted in Figure 6.21 and then again in Figures 6.22 and 6.23 where the matte and gloss samples have been separated for clarity. These samples include the five samples from the gloss perception experiment.

An overall trend showing a decrease in texture perception from 88 degrees to about 60 degrees followed by an increase in texture perception as the lighting angle decreases is evident for 4 of the 8 samples. Three of the remaining four samples continue to decrease from 60 degrees onward as seen in Figures 6.22 and 6.23. The remaining sample, DarkOil




Fig. 6.21. – Average texture perception rating for all 8 artwork samples: OilTrees, Gray Patch 1, Tulip, DarkOil, Print, Textile, Gray Patch 2, and Flowers.




Fig. 6.22. – Average texture perception rating for 5 of the 8 artwork samples: OilTrees, DarkOil, Tulip, Gray Patch 1, and Flowers.




Fig. 6.23. – Average texture perception rating for 3 of the 8 artwork samples: Gray Patch 2, Textile, and Print.

behaves similarly to both sets, and has only a slight upward turn in ratings between 10 and 30 degrees. It was hypothesized that texture perception ratings should decrease monotonically as lighting angle decrease as indicated by the dashed lines in Figures 6.22 and 6.23. This was hypothesized based on the practice of using raking lighting to accentuate texture. Because the matte samples followed the expected trend and have no gloss components, it is further hypothesized that gloss influences the observers' ratings of texture and therefore the matte samples will be excluded from further investigations.

Higher values of perceived texture are expected at higher angles, like those used for raking lighting. It follows that as the angle is decreased, the perceived texture will decrease for most samples. The four higher gloss samples that have texture ratings that slowly increase from about 60, shown in Figures 6.24, 6.26, 6.27 and 6.28 are believed to be influenced by the presence of gloss in these images. With the exception of the DarkOil painting, shown in Figure 6.25, each of the samples from the gloss task strongly follow this trend. It is believed that the gloss visible in the DarkOil painting has more of an appearance of noise rather than gloss as indicated by some of the observers during the questionnaire portion

of the experiment. Close up sections of the images of the DarkOil painting and the Tulip painting taken with the light source at 67 degrees are shown in Figure 6.29 for comparison to illustrate this claim. In the DarkOil image, the highlights follow more of a speckled pattern whereas the highlights in the Tulip image follow the edges of impasto.




Fig. 6.24. – Average texture perception rating for Gray Patch 1 sample.

To further illustrate the influence of gloss on texture, images of the same sample with similar texture ratings but different lighting angles are shown in Figures 6.30 and 6.31. In these images, the areas of highlights often accentuate the structure of the painting in the same way that shadows accentuate the surface structure. Based on the texture ratings of the three matte samples, it is believed that the texture ratings of these glossy images would be rated differently if they were matte but still possessed the same amount of impasto or surface features.

In these images, highlights produced by the the glossy nature of the samples and the lighting geometry accentuated the presence of brush strokes and layered paint in much the




Fig. 6.25. – Average texture perception rating for DarkOil sample.




Fig. 6.26. – Average texture perception rating for OilTrees sample.




Fig. 6.27. – Average texture perception rating for Tulip sample.




Fig. 6.28. – Average texture perception rating for Flowers sample.



Fig. 6.29. – Close up sections of the images of the DarkOil painting and the Tulip painting taken with the light source at 67 degrees that emphasize the difference between gloss depicted in the Dark Oil and Tulip paintings.

same way that shadows accentuated the thickness of impasto and heaviness of the brush strokes.

Plots of observer ratings for each sample and their averages are shown in Figures 6.32 through 6.36. As seen in these plots, ratings near normal angles have the highest errors. These angles are where the samples are expected to appear most glossy. For the DarkOil, OilTrees, and Flowers samples it appears that a small number of observers were able to discount the influence of gloss and rate images in a manner similar to the predicted trend while others were not. This created the larger errors from 3 to 10 degrees.




Fig. 6.30. – Images of 3 of the 5 samples that have roughly the same texture rating.



Fig. 6.31. – Images of 2 of the 5 samples that have roughly the same texture rating.




Fig. 6.32. – Average texture perception rating for Gray Patch 1 sample using data from all samples.




Fig. 6.33. – Average texture perception rating for DarkOil sample using data from all samples.




Fig. 6.34. – Average texture perception rating for OilTrees sample using data from all samples.




Fig. 6.35. – Average texture perception rating for Tulip sample using data from all samples.




Fig. 6.36. – Average texture perception rating for Flowers sample using data from all samples.

6.3.3 Observer Preferences

Observers were asked to rate each image compared to a reference image based on their memory of viewing each sample in a gallery setting. Following this task, observers were asked to rate which image they preferred compared to the same reference image. The highest rated images for each task and sample are shown in Table 6.3 denoted by the degree at which the lighting was configured for each sample.

For 5 out of the 8 samples, observers preferred images taken with lighting from higher angles than the angles used to capture images they considered to be most representative. Larger lighting angles create images with more texture and less gloss in many cases. From these data, it is inferred that observers typically prefer images that display the texture of the piece without the effects of gloss compared to what may be more representative of the original work when viewed in a gallery setting. The images most often selected as most representative and preferred are shown in Figures 6.45, 6.43, 6.44, 6.46, and 6.47, along with closeup sections of some samples to highlight the differences between the preferred and representative images.

Tab. 6.3. – Observer preferences compared to observer selected representative images for each of the eight samples denoted by the lighting angle used to capture each image.

Sample	Gray Patch 1	Flowers	Oil Trees	Tulip	Gray Patch 2	Dark Oil	Print	Textile
Preferred	67°	67°	52°	67°	52°	52°	12°	67°
Representative	77°	52°	42°	59°	28°	67°	42°	52°

Although some of the differences are subtle, the preferred images of the Tulip and Flowers samples show more shadows and in turn more texture which can be seen in the layered paint in the blue area of the Tulip painting and in the light purple area in the Flowers image. The preferred image of the DarkOil sample has less gloss in the dark areas of the painting than the representative image. The preferred image of the Gray Patch 1 sample was taken with a lower angle of lighting and has less perceived texture than the representative image. This is shown in the darker shadows especially in the canvas area of the sample. The differences in the OilTrees sample are harder to distinguish partially due to the high spatial frequency nature of the image. In the close up section more highlights can be seen in the representative image than the preferred images.




Fig. 6.37. – Image preferred (left) next to the image that best represents the original in a gallery setting (right).



Fig. 6.38. – Image preferred (left) next to the image that best represent the original in a gallery setting (right). A close up section of both images is shown underneath.

For most samples, preferences did not stray far from representative images. In general, observers preferred images with more texture and less gloss than that of the representative



Fig. 6.39. – Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.

images with a few exceptions. For the Gray Patch 1 sample, observers preferred the image with less perceived texture possibly due to the texture coming mostly from the canvas substrate rather than the applied paint. Observers preferred a slightly lower lighting angle for the DarkOil sample. Similar to the Gray Patch 1 sample, the lower angle highlights the structure of the canvas more than the structure of the paint as evident in the close up image taken with lighting at 67 degrees. Both of these samples have some gloss inherent to either the paint or varnished used. The last sample that observers preferred a lower lighting angle for was the matte Print sample. Because this sample is matte, many of the lower angle samples look very similar and are difficult to differentiate between them. However, when comparing the image the observers preferred to the image they felt was most representative, the observers preferred the image that did not emphasize the structure of the paper used. Small bumps, as well as more shadows can be seen in the representative images, that are not noticeable in the preferred image. It can thus be inferred that observers tend to prefer more perceived texture in images of artwork unless the texture is generated primarily



Fig. 6.40. – Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.

from the substrate the artwork is created on instead of the medium used. Plots of both preference and representative image ratings are shown in Figures 6.45 through 6.47.

The Gray Patch 1 sample exhibits a unique appearance relating to the changes in the lighting direction that the other samples do not exhibit. When the light angle changes the emphasis on the texture directly related to the application of paint transfers to the texture created by the canvas. This shift from one main texture appearance to another is shown in Figure 6.42 where each image is shown starting with the largest lighting angle in the top left and the smallest angle in the bottom right. In the images taken with the smaller lighting angles both physical features are easily recognized with a nearly uniform canvas weave pattern visible underneath the layer of paint as well as two ridges and a pock mark towards the bottom of the sample. As the angle increases the visibility of the paint ridges and pock mark lessens while the canvas weave becomes more pronounced in the higher angle images compared to the middle angle values. Other samples do not exhibit this shift in texture appearance. For the textile, the fabric weave provides all of the information about



Fig. 6.41. – Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.

texture to the observer because there is not a gloss component to the the fabric as well as it having no other contributing features. The DarkOil painting has paint layered on top of a canvas but the canvas weave provides little to the texture appearance because of how thickly the paint is applied compared to the Gray Patch 1 sample. It is believed that the switch between which features were responsible for the majority of the texture appearance for the Gray Patch 1 sample caused different observers to categorize the perceived texture differently than others resulting in large errors especially for the higher lighting angles. The overall shape of the ratings could also be influenced by the change in how each physical feature was emphasized by the changes in lighting.




Fig. 6.42. – Gray Patch 1 sample showing each image used in the experiment starting with the largest lighting angle in the top left and the smallest angle in the bottom right.




Fig. 6.43. – Image preferred (left) next to the image that best represent the original in a gallery setting (right). A close up section of both images is shown underneath.




Fig. 6.44. – Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.




Fig. 6.45. – Image preferred (left) next to the image that best represents the original in a gallery setting (right).




Fig. 6.46. – Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.




Fig. 6.47. – Image preferred (left) next to the image that best represents the original in a gallery setting (right). A close up section of both images is shown underneath.

6.4 Conclusion

Observers were asked to rate perceived texture and gloss for eight different samples. When both texture and gloss were present in the same sample, observers were often influenced by the feature they were not asked to rate. In other words, in an observer was asked to rate the gloss level of a sample, their rating was often influenced by the presences of texture in images taken at near raking angles. From these observations it is hypothesized that it is difficult for observers to separate the two features when rating samples.

Observers were also asked about the aesthetics of images by rating images based on preference as well as how representative the images were to artwork shown in a gallery setting. Observers typically preferred images with lighting between 67 and 52 degrees, while glossy samples were typically considered representative with light at 52 or 42 degrees, indicating that observers preferred images with more visible texture. Combining the survey question responses, found in Appendix C, to the experimental data, further strengthens the conclusion that observers prefer images of artwork that contain a higher perceived texture and lower perceived gloss than images that represent artwork viewed in a simulated gallery environment. This indicates that the appropriate lighting for a gallery may be different from the lighting that should be used to create digital reproductions.

Observer Preferences and Modeling

As described in the previous chapter, an overall trend showing a decrease in gloss perception from 88 degrees to about 60 degrees followed by an increase in gloss perception as the lighting angle decreases is evident for all five glossy samples. The higher gloss rating at 88 degrees is hypothesized to be due to the influence of increased texture appearance as the angle of lighting increased from 60 to 88 degrees. The appearance of both texture and gloss is minimized around 60 degrees.

These trends were characterized using statistical data derived from each separate image. Each image was converted from original 16-bit encoded images to grayscale using the Matlab function *rgb2gray*. These calculations, from original pixel data, included histogram skewness, kurtosis, mean, and standard deviation. The equations for skewness and kurtosis are shown below. Kurtosis was included as a measure whether the data contains few or many outliers. Kurtosis is also calculated from the image histogram. It is hypothesized that an image with many specular highlights may have more outliers than images that do not. The kurtosis for a normal distribution is 3, higher values indicate longer histogram tails or a larger number of outliers and a kurtosis less than 3 indicates a lack of outliers [National Institute of Standards and Technology (NIST), 2013]. Kurtosis is calculated according to Equation 7.1 where n is sample size, and \bar{x} is the mean of sample [MathWorks, Inc., 2017]. Skewness was included because it has been correlated with judgements of both gloss and lightness [Motoyoshi et al., 2007]. Skewness is considered highly correlated with gloss judgments and a histogram with positive skew is indicative of increased gloss compared to images with negatively skewed histograms. Histogram skewness can be calculated according to Equation 7.2 where $I(x,y)$ is the luminance at pixel (x,y) , m is the mean, N is the number of pixels, and $s.d$ is the standard deviation [Motoyoshi et al., 2007].

$$kurtosis = \frac{\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^4}{\left(\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \right)^2} \quad (7.1)$$

$$skew = \frac{\sum (I(x,y) - m)^3}{N(s.d.)^3} \quad (7.2)$$

An example of how histogram skewness is related to perceived gloss of a stucco-like surface is shown in Figure 7.1 where the image on the left has a negative skewness of -1.34 and the image on the right has a positive skewness of 2.40. Although these images have the same mean luminance, the image on the right appears to be glossier [Motoyoshi et al., 2007].




Fig. 7.1. – Two images with the same mean luminance but different skewness values, -1.34 (left) and 2.40 (right) and different amounts of perceived gloss [Motoyoshi et al., 2007].

The gray level covariance matrix, or GLCM, was also calculated for each image. The GLCM is a matrix of relative frequencies that denote how often two neighboring pixels separated by a defined distance have specified gray values [Haralick, Shanmugam, et al., 1973]. GLCMs have been used to successfully characterize different textural features for a variety of applications including analyzing sea ice imagery, assessing wear of carpet fibers, and image classification [Soh and Tsatsoulis, 1999], [Siew et al., 1988], [Haralick, Shanmugam, et al., 1973].

GLCMs were calculated using the quantization value of 64, in terms of bit depth, and offset values of [0 1], [1 1], and [1 0] to select nearest neighbors. Offset values are written in

terms of [row offset, column offset]. A diagram showing the nearest neighbors to the pixel of interest, as defined by these three offset values, is shown in Figure 7.2. GLCMs from the three different offsets were averaged together to create a single GLCM to be used for further calculations. These matrices were used to calculate dissimilarity, entropy, and contrast of each image. The equations for dissimilarity (f_d), entropy (f_e), and contrast (f_c) are shown in Equations 7.3, 7.4, and 7.5 where $p(i,j)$ is the (i,j) term in the GLCM and N_g is the quantized gray levels [Soh and Tsatsoulis, 1999]. Entropy values denote randomness or homogeneity. A low value indicates inhomogeneity while a larger value indicates a large degree of homogeneity [Albregtsen et al., 2008]. Contrast characterizes local intensity difference.

$$f_d = \sum_i \sum_j |i - j| \cdot p(i, j) \quad (7.3)$$

$$f_e = - \sum_i \sum_j p(i, j) \log(p(i, j)) \quad (7.4)$$

$$f_c = - \sum_{n=0}^{N_g-1} n^2 \left\{ \sum_{i=1}^N \sum_{j=1}^N p(i, j) \middle| |i - j| = n \right\} \quad (7.5)$$




Fig. 7.2. – A diagram showing the pixel of interest's nearest neighbors as defined by [0 1], [1 1], and [1 0].

In addition to calculating GLCMs, a series of Gabor filters were applied to each image. These filters included combinations of four angles (0, 45, 90, and 135 degrees) and four wavelengths (2, 4, 8, and 16). These values were selected to sample both small scale differences and larger scale differences in texture in multiple directions. It is hypothesized that 2 and 4 wavelengths are on the same order as texture associated with canvas weave

and highlights, while 8 and 16 are more associated with changes in brush strokes. Averages were then calculated for each filtered image. The average of all filtered images, and the average of all filters using a distance of 2 were of particular interest and used in the modeling process.

A Gabor filter is used for texture analysis and is considered to be similar to the human visual system. It has been argued that the visual cortex representation is closely related to Gabor filtering and can be used in pattern recognition [Marčelja, 1980]. An example of a Gabor filtered image is shown in Figure 7.3.



Fig. 7.3. – A Gabor filtered image of the Tulip sample using a wavelength of 8 and 0 degrees.

The calculated data from each image, GLCM, and Gabor filtered image were used to create a model that predicted the gloss and texture perception of each sample as a function. A stepwise regression technique was used to predict gloss and texture ratings using the Matlab function *stepwisefit*. This function performs a stepwise multilinear regression. This method adds and removes terms to determine the best fit using supplied data. The stepwise function was used to model the gloss and texture perception of each sample separately.

7.0.1 Gloss Perception Modeling

Each of the different samples was modeled using a different combination of statistics and constants. The gloss perceptions of the samples were modeled using mean, entropy, mean of the Gabor filtered images using a distance of 2, standard deviation, and contrast. Each predicted and observer rated gloss value is shown in Figures 7.4 through 7.8. The gloss perception for the Tulip sample is modeled using Equation 7.6. This equation uses mean, standard deviation, and contrast to model the observer perceived gloss. This resulted in an RMSE of 1.97. The gloss perception for the OilTrees sample was modeled using Equation 7.7 and only standard deviation. This resulted in an RMSE of 4.36. The Gray Patch 1 sample was modeled using Equation 7.8 with mean and standard deviation resulting in an RMSE of 2.35. The gloss perception of the DarkOil sample was modeled using Equation 7.9 with entropy resulting in an RMSE of 3.27. Gloss of the Flowers sample was modeled using Equation 7.10 with the average of the Gabor filtered images using the distance of 2. This resulted in an RMSE of 5.70. This is the highest RMSE value for the gloss samples. Each of the RMSE values are recorded in Table 7.1

All RMSE values are under 6, with the highest values being associated with the varnished and most glossy samples. The Flowers sample has the highest error value which could be due to the type of varnish that was applied. Varnish was sprayed, rather than brushed onto the Flowers sample which impart some texture from the varnish layer itself, making it more difficult to make gloss judgements. The lowest errors occur for the Gray Patch 1 and Tulip samples. These samples are comprised of large blocks of similar color, the Gray Patch 1 sample being a single color in its entirety. It is believed that this makes it easier to discern changes in gloss because the observer is not distracted by high frequency details in the image.

$$Tulip_{gloss} = (368.588 \times Mean) - (1817.400 \times StdDev) + (0.901 \times Contrast) + 86.242 \quad (7.6)$$

$$OilTrees_{gloss} = (418.919 \times StdDev) - 30.836 \quad (7.7)$$

$$Gray1_{gloss} = (182.846 \times Mean) + (8.249 \times Entropy) - 63.139 \quad (7.8)$$

$$DarkOil_{gloss} = (33.707 \times Entropy) - 177.08 \quad (7.9)$$

$$Flowers_{gloss} = (787.638 \times Gabor_2) + 16.333 \quad (7.10)$$




Fig. 7.4. – Stepwise fit modeling the gloss of the Tulip sample (blue) and the observer rated data (black).

Tab. 7.1. – RMSE values for samples used in the gloss perception modeled using the stepwise function method.

Sample	Gray Patch 1	DarkOil	Flowers	Tulip	OilTrees
RMSE	2.35	3.27	5.70	1.97	4.36

The gloss of each sample was remodeled using all samples simultaneously. This resulted in a single equation for all samples shown in Equation 7.11. This modeled used seven out of the nine image statistics to model the gloss of each sample: the average of all Gabor filtered images (*Gaboravg*), skewness (*Skew*), kurtosis (*Kurtosis*), mean (*Mean*),




Fig. 7.5. – Perceived gloss rating (blue) and predicted rating (black) for the OilTrees sample.




Fig. 7.6. – Perceived gloss rating (blue) and predicted rating (black) for the Gray Patch 1 sample.




Fig. 7.7. – Perceived gloss rating (blue) and predicted rating (black) for the DarkOil sample.




Fig. 7.8. – Perceived gloss rating (blue) and predicted rating (black) for the Flowers sample.

contrast (*Contrast*), entropy (*Entropy*), and the average of Gabor filtered images that used a wavelength of 2 (*Gabor2*). The RMSE values for this new model are shown Table 7.2. Plots showing the new model are shown in Figures 7.9 through 7.13.

$$\begin{aligned}
 Sample_{gloss} = & (19.119 \times Gabor_{avg}) + (13.502 \times Skew) - (2.100 \times Kurtosis) \\
 & + (204.889 \times Mean) - (0.270 \times Contrast) - (8.244 \times Entropy) \quad (7.11) \\
 & + (699.727 \times Gabor_2) - 23.672
 \end{aligned}$$

Tab. 7.2. – RMSE values for samples used in the gloss perception modeled using the stepwise function method with all samples used as input simultaneously.

Sample	Gray Patch 1	DarkOil	Flowers	Tulip	OilTrees
RMSE	4.87	4.94	5.20	5.56	5.84

All of the RMSE values are under 6 but the model uses seven terms when the individual models used 3 at the most with better results. However, it is expected that the individual sample models would perform better because modeling all the samples simultaneously introduces more variance in the data that has to be accounted for in a single equation instead of five separate equations.




Fig. 7.9. – Stepwise fit modeling the gloss of the Tulip sample (blue) and the observer rated data (black) using all samples to produce the model.




Fig. 7.10. – Stepwise fit modeling the gloss of the OilTrees sample (blue) and the observer rated data (black) using all samples to produce the model.




Fig. 7.11. – Stepwise fit modeling the gloss of the Gray Patch 1 sample (blue) and the observer rated data (black) using all samples to produce the model.




Fig. 7.12. – Stepwise fit modeling the gloss of the DarkOil sample (blue) and the observer rated data (black) using all samples to produce the model.




Fig. 7.13. – Stepwise fit modeling the gloss of the Flowers sample (blue) and the observer rated data (black) using all samples to produce the model.

7.0.2 Texture Perception Modeling

The texture perceptions of the samples were modeled using skewness, kurtosis, mean, entropy, and dissimilarity. Plots of the predicted and observer rated perceived texture are shown in Figures 7.14 through 7.17. The texture perception for the Tulip sample was modeled using Equation 7.12. This equation used skewness and kurtosis to model the observer perceived texture. This resulted in an RMSE of 4.63. The Gray Patch 1 sample was modeled using Equation 7.13 with mean and entropy resulting in an RMSE of 3.09. The texture perception of the DarkOil sample was modeled using Equation 7.14 with kurtosis and entropy resulting in an RMSE of 3.22. Texture of the Flowers sample was modeled using Equation 7.15 with mean and dissimilarity. This resulted in an RMSE of 3.84. Each of the RMSE values for the texture perception modeling are recored in Table 7.3. The texture of OilTrees could not be modeled using these image statistics.

The RMSE values for the texture perception models are all under 5. The lowest RMSE value occurs for the Gray Patch 1 sample. This sample, again, is a single color and may be more easily rated and modeled due to its simplicity. The OilTrees sample, which could not be modeled, has a high frequency pattern that could contribute to its poor performance.

$$Tulip_{texture} = (41.284 \times Skew) + (31.065 \times Kurtosis) - 33.152 \quad (7.12)$$

$$Gray1_{texture} = (-67.530 \times Mean) - (19.037 \times Entropy) - 8.271 \quad (7.13)$$

$$DarkOil_{texture} = (19.392 \times Kurtosis) + (24.468 \times Entropy) - 211.268 \quad (7.14)$$

$$Flower_{texture} = (-368.67 \times Mean) + (10.325 \times Dissimilarity) + 138.970 \quad (7.15)$$

Tab. 7.3. – RMSE values for samples used in the texture perception modeled using the stepwise function method

Sample	Gray Patch 1	DarkOil	Flowers	Tulip	OilTrees
RMSE	3.09	3.22	3.84	4.63	—




Fig. 7.14. – Perceived texture rating (blue) and predicted rating (black) for the Tulip sample.




Fig. 7.15. – Perceived texture rating (blue) and predicted rating (black) for the Gray Patch 1 sample.




Fig. 7.16. – Perceived texture rating (blue) and predicted rating (black) for the DarkOil sample.




Fig. 7.17. – Perceived texture rating (blue) and predicted rating (black) for the Flowers sample.

The texture perception of each sample was also remodeled using all samples simultaneously. This resulted in a single equation for all samples shown in Equation 7.16. This modeled used three out of the nine image statistics to model the texture of each sample: contrast (*Contrast*), dissimilarity (*Dissimilarity*), and the average of Gabor filtered images that used a wavelength of 2 (*Gabor2*). The RMSE values for this new model are shown Table 7.4. Plots showing the new model are shown in Figures 7.18 through 7.22.

$$\begin{aligned} Sample_{texture} = & (-0.301 \times Contrast) + (15.823 \times Dissimilarity) \\ & - (512.911 \times Gabor_2) + 14.973 \end{aligned} \quad (7.16)$$

Tab. 7.4. – RMSE values for samples used in the texture perception modeled using the stepwise function method with all samples used as input simultaneously.

Sample	Gray Patch 1	DarkOil	Flowers	Tulip	OilTrees
RMSE	7.02	11.87	13.86	9.83	16.89

All of the RMSE values are are over 7. The model uses three terms when the individual modes used 2 at the most with better results. However, it was not possible to create a model for the texture perception of the OilTrees sample alone from the image statistics used. Modeling texture using all samples simultaneously is significantly less accurate with RMSE values being at least twice as high for most samples, with some of the predicted texture ratings losing the overall shape of the observer ratings as shown in Figure 7.20 where predicted ratings between 90 and 70 degrees are significantly lower than their observer selected counterparts for the DarkOil sample. Another example is evident for the OilTrees sample where predicted ratings form a horizontal line from 70 to 30 degrees instead of being bowl shaped like the observer ratings shown in Figure 7.21.




Fig. 7.18. – Perceived texture rating (blue) and predicted rating (black) for the Tulip sample using all samples as input.




Fig. 7.19. – Perceived texture rating (blue) and predicted rating (black) for the Gray Patch 1 sample using all samples as input.




Fig. 7.20. – Perceived texture rating (blue) and predicted rating (black) for the DarkOil sample using all samples as input.




Fig. 7.21. – Perceived texture rating (blue) and predicted rating (black) for the OilTrees sample using all samples as input.




Fig. 7.22. – Perceived texture rating (blue) and predicted rating (black) for the Flowers sample using all samples as input.

7.0.3 Modeling Preference

After modeling the texture and gloss of each sample, a model was created for the preference of each sample. Entropy was used for 3 out of the 5 samples. Mean was used for 2 out of the 5 samples. Skewness, standard deviation, and the average of Gabor filtered images using a wavelength of 2 were also used. Plots of the modeled preference data compared to the observer rated preferences are plotted in Figures 6.43 through 6.47. Equations 7.17 through 7.21 are the equations optimized using the *stepwisefit* function for each of the 5 samples. The model for the OilTrees samples used standard deviation, Gray Patch 1 used mean and entropy, Flowers used entropy, Tulip used skewness, mean, and mean Gabor filtered images with a wavelength of 2, and DarkOil used entropy. The RMSE values for each sample's preference modeling is shown in Table 7.5.

$$Tulip_{preference} = (-54.36 \times Skew) - (292.69 \times Mean) + (1245.300 \times Gabor_2) + 150.87 \quad (7.17)$$

$$OilTrees_{preference} = (-467.70 \times StdDev) + 151.14 \quad (7.18)$$

$$Gray1_{preference} = (-83.823 \times Mean) - (12.012 \times Entropy) + 153.72 \quad (7.19)$$

$$DarkOil_{preference} = (-34.38 \times Entropy) + 294.62 \quad (7.20)$$

$$Flowers_{preference} = (-34.25 \times Entropy) + 291.84 \quad (7.21)$$

Tab. 7.5. – RMSE values for each predicted versus observer selected preference rating using the stepwise function method.

Sample	Gray Patch 1	DarkOil	Flowers	Tulip	OilTrees
RMSE	4.86	3.66	5.60	3.53	5.52

The RMSE values for preference were all under 6. The highest values occurred for the Flowers and OilTrees samples. The lowest error occurred for the Tulip sample. It is of interest to note these RMSE values follow the same trend as the perceived gloss RMSE




Fig. 7.23. – Preference rating (blue) and predicted rating (black) for the Tulip sample.




Fig. 7.24. – Preference rating (blue) and predicted rating (black) for the OilTrees sample.




Fig. 7.25. – Preference rating (blue) and predicted rating (black) for the Gray Patch 1 sample.




Fig. 7.26. – Preference rating (blue) and predicted rating (black) for the DarkOil sample.




Fig. 7.27. – Preference rating (blue) and predicted rating (black) for the Flowers sample.

values. For example, Tulip has the smallest error for both gloss and preference. Similarly, Flowers and OilTrees having the highest RMSE values for both gloss and preference rating. It is plausible that observers are looking at similar visual cues to assess gloss as well as preference; however more investigation is necessary to determine if these trends are related.

A second modeling approach for each sample was employed that used all samples simultaneously resulting in a single equation for all samples shown in Equation 7.22. This model used two out of the nine image statistics to model the preference of each sample: mean (*Mean*) and dissimilarity (*Dissimilarity*). The RMSE values for this new model are shown Table 7.6. Plots showing the results of the new model are shown in Figures 7.28 through 7.32.

$$Sample_{preference} = (-62.306 \times Mean) - (4.193 \times Dissimilarity) + 110.461 \quad (7.22)$$

Tab. 7.6. – RMSE preference values for each sample using the stepwise function method with all samples used as input simultaneously.

Sample	Gray Patch 1	DarkOil	Flowers	Tulip	OilTrees
RMSE	10.49	8.35	8.82	11.52	10.01

All of the RMSE values are over 8. The previous models had RMSE values all under 6. The new model uses two terms when the individual models used 2 and sometimes 1 term with better results. It is not unreasonable for the new model to behave poorly compared to the original model that modeled each sample individually. The newly predicted ratings do not match the overall shape of the observer ratings with many predictions flattening between 70 to 30 degrees instead of following the curve shape of the the observer rating data.




Fig. 7.28. – Preference rating (blue) and predicted rating (black) for the Tulip sample modeled using all samples as input.




Fig. 7.29. – Preference rating (blue) and predicted rating (black) for the OilTrees sample-modeled using all samples as input.




Fig. 7.30. – Preference rating (blue) and predicted rating (black) for the Gray Patch 1 sample-modeled using all samples as input.




Fig. 7.31. – Preference rating (blue) and predicted rating (black) for the DarkOil sample-modeled using all samples as input.




Fig. 7.32. – Preference rating (blue) and predicted rating (black) for the Flowers sample-modeled using all samples as input.

7.1 Conclusion

A stepwise linear regression method was used to model perceived gloss and texture and observer preference. The modeling routine started with nine image statistic sets of input data, histogram skewness, kurtosis, mean, standard deviation, contrast, dissimilarity, entropy, mean of a set of Gabor filtered images using a range of angles and wavelengths, and the mean of a set of Gabor filtered images that used only a wavelength of two. Each sample and trait of interest was modeled differently. Out of all nine sets of data, only the mean of all Gabor filtered images was excluded. Entropy and mean were used the most often while dissimilarity, and contrast were used the least often. Mean and entropy are likely candidates for modeling gloss and texture because these values are likely to change when the amount of highlights or shadows change. Highlights and shadows were both mentioned as features used to make judgements by observers. For most samples, the image statistics used in their respective models were different from texture to gloss to preference with the exception of the Gray Patch 1 sample which used mean and entropy for all three, OilTrees that used standard deviation for gloss and preference, and DarkOil that used entropy for all three with kurtosis also included to model texture.

This approach to modeling did not uncover a single or set of image statistics that could account for all ratings related to each trait, so a second model that used all samples simultaneously was investigated for each trait. These three models performed poorly compared to the original set of models with corresponding RMSE values often being at least twice as large as the RMSE values from the original models. Although a single set of statistics or models was not uncovered does not mean that one does not exist, it can only be concluded that one does not exist in the set of image statistics tested. Future work would be required to model gloss, texture, and preference more completely. This would include imaging more samples as well as exploring more image statistics and filters.

Conclusions

Representing fine art requires an attention to detail that focuses on more than just color, but the entire material appearance. Texture and gloss are two very important features to consider in reproducing artwork. In representing fine art, it is important to first utilize a capturing technique that will provide all the necessary information without requiring additional and often expensive equipment or training that many smaller institutions cannot access.

This research had four main parts:

- This research improved the four lights imaging technique
- Psychophysical experiments were conducted to better understand the observer perceptions of gloss and texture
- Image statistics were used to model the gloss and texture ratings
- Psychophysical experiments related to observer preference were conducted

This research improved the four lights imaging technique by replacing the previously used polarizers with a mathematical thresholding routine. This improvement reduced the number of images required to produce surface normal maps, reducing the time necessary to capture and process images. This research also showed that images taken with this imaging technique could be used to create a virtual viewing experience using a computer graphics rendering software. This viewing experience could even recreate spatially varying gloss across the surface of a painting. These results and improvements on the four lights imaging technique could be used to create virtual exhibits in museums and institutions that are unable to acquire all of the pieces in a series or to augment a current show.

In the second half of this research psychophysical experiments were conducted to better understand the observer perceptions of gloss and texture. It was determined that observers were influenced by texture when making gloss ratings and were influenced by gloss when making texture ratings. Image statistics were used to model the gloss and texture ratings.

Though no single set of statistics was found to model the gloss or texture of all samples, entropy and mean were most commonly used, indicating that randomness and changes in average luminance are important features when discriminating between gloss and texture levels.

Although the success of these models is limited to samples with material properties that lie within the sample set and no single set of statistics could be used to model all samples, the modeling process provided a valuable starting point for future endeavors in modeling gloss and texture from image statistics. Experiments related to observer preferences were an important part of this section of work. Observers recorded preferences related to accentuating the amount of texture visible in images while limiting the presence of gloss. The preferred images differed from the images that were selected as representative of art shown in a gallery setting, indicating that the best way to display fine art and render fine art may be different. These findings can be used to inform future renderings and reproductions of fine art to create educational or promotional material that is pleasing to consumers while maintaining a physically possible result.

8.0.1 Future Work

A beneficial area of future work is further testing of the 4LIS technique. This testing could include more complex samples with varying regions of gloss and different textures. Samples of interest could include artwork that has a range of textures that represent a fine texture similar to canvas weave or smaller to samples with texture that varying on the order of centimeters rather than millimeters. Future work to determine the limitations of the 4LIS system would provided valuable information about what types of fine art can be accurately imaged. Current limitations of imaging artwork with occluded structures, mirror-like gloss, and dark colors have already been observed. Further investigation could provide more precise limitations.

Further testing of the perception models described in Chapter 7 is also recommended. A limited number of samples was used during the gloss and texture experiments. The intention of including the chosen samples was to have examples of a range of different textures and gloss levels. Because the matte samples proved to behave in an expected

fashion, it would be most beneficial to focus on samples that contained both gloss and texture. Additional imaging should occur with a larger number of samples that include a set of single color patches as well as samples with high frequency details. These samples should first be considered separately to identify factors that influence ratings that do not include the structure of the artwork. Investigation into more appropriate image statistics and image filters is also necessary to identifying a set of data that can be used for all samples. A test of how well the models predict gloss and texture for a sample that was not used to create subsequent models should be selected and imaged at a series of different lighting angles. The image statistics should then be calculated for each of the images taken and used to predict the perceived gloss and texture for each image according to the appropriate equation. The samples could then be used to create a similar set of experiments to generate observer assigned gloss and texture values that could be compared to the predicted values.

The results from the preference based experiments of texture and gloss appearance in Chapter 6 could be used to design and perform an experiment related to lighting artwork in a gallery environment. The lighting in galleries and museums have been carefully considered and adjusted by experts interested in providing adequate lighting for viewing as well as lighting that will limit the amount of damage the displayed artwork may undergo while on exhibit. An experiment could be developed that explored the usefulness of changing lighting direction to enhance material related features while allowing observers to remain at a safe distance.

Appendix A: 4LIS and Maya: A Case Study

Demonstration of the Four-Light Simplified Technique to Capture and Render Fine Art Images at MoMA Queens

Painting 1: *Girl on a Divan*

Artist: Ernst Ludwig Kirchner

Date: 1906

Medium: oil on composition board

Dimensions: 18 5/8 x 26 3/4 in

Painting 2: *Vase of Flowers*

Artist: Pablo Picasso

Date: Paris, Spring 1908

Medium: oil on canvas

Dimensions: 36 1/4 x 28 3/4 inches




Fig. A.1. – Painting 1: *Girl on a Divan*, Kirchner and Painting 2: *Vase of Flowers*, Picasso

The purpose of the visit to the Museum of Modern Art Queens (MoMA) facility was to observe typical imaging practices performed by their Department of Imaging and Visual Resources and to demonstrate the Four-Light Simplified (4LIS) technique developed in the Munsell Color Science Laboratory as part of the Studio of Scientific Imaging and Archiving of Cultural Heritage.

A.1 MoMA Imaging Setup

The Imaging and Visual Resources departmental staff was asked to set up each painting, Kirchner's *Girl on a Divan* (Painting 1) and Picasso's *Vase of Flowers* (Painting 2), according to their typical imaging workflow. The paintings were mounted on a white wall centered at approximately sixty inches from the floor with color calibration targets mounted to the left of the painting. The images were taken with a calibrated Hasselblad camera in single- and multi-shot modes. Two Broncolor strobes above the painting, pointed vertically at the ceiling, and a diffuse white board positioned below the painting were used to provide lighting for each shot. Images were taken with both strobes firing simultaneously. This setup was also used to image paintings and color calibration targets provided by the Studio for Scientific Imaging and Archiving of Cultural Heritage. Figure A.2 is a photograph of the setup used for the typical imaging process employed by MoMA Department of Imaging and Visual Resources photographer.



Fig. A.2. – Lighting setup used by the Department of Imaging and Visual Resources at MoMA using two strobes (highlighted with red circles) directed toward the ceiling from above the paintings.

A.2 4LIS Imaging Setup

After all images were taken using the department workflow, the paintings were left mounted on a white wall with color calibration targets mounted to the left of the painting. New images were taken with the same calibrated Hasselblad camera in the single-shot mode, however, the configuration of the lighting was changed to comply with the 4LIS workflow.¹ Four strobes were arranged around an annulus about the painting approximately ninety degrees from each other and forty-five degrees from the painting front surface. The lighting setup for the 4LIS technique is shown in Figure A.3.



Fig. A.3. – Lighting setup used for Four-Light Imaging technique. Four lights are placed to the left, above, to the right, and below the painting at approximately 45 (highlighted with red circles).

This setup was used to image paintings, color calibration targets, the diffuse white wall, and a glossy black sphere (cue ball). Each element was imaged four different times using only one of the four strobes, resulting in four separate images for each object of interest illuminated from a different direction according to the process described by Cox and Berns and Berns and Chen [Cox and Berns, 2015; Berns and Chen, 2012].

A.3 Post Processing

Images taken using the typical setup used by MoMA were tone corrected to match L^* values for middle gray and black patches on the GoldenThread Object-Level target. (For comparison purposes the images taken with the multi-shot setting of the Hasselblad camera are not used from MoMA workflow.)

The images taken in accordance with the 4LIS technique were used to flatfield, color correct, identify the source direction of each light, calculate a surface normal map using Woodham photometric stereo technique, and generate a diffuse color image [Berns and Chen, 2012; Woodham, 1980]. The diffuse color image and the surface normal map were combined in the computer graphics rendering software, Maya, and rendered using the mental ray plugin to generate an image that would mimic the image taken by the staff photographer. The diffuse color image and normal map of Painting 1 and Painting 2 accompanied by a close up are depicted in Figures A.4 and A.5, respectively.

The normal map associated with Painting 1 contains holes where dark paint had been applied to the composition board. The term holes refers to areas on the surface normal map that do not correspond to physical features of the painting. These holes are attributed to the lack of signal captured by the camera due to the dark and glossy nature of the paint in these areas. There was not enough data captured by the imaging system to generate surface normals that correspond to these dark regions of this painting. However, the normal map of Painting 2 contains no noticeable holes. Painting 2 is less glossy than Painting 1 resulting in fewer highlights being produced and thus less pixel data being computationally removed by the thresholding method described by Cox and Berns [Cox and Berns, 2015].

A.4 Rendered Images and Results

The surface normal map and the diffuse color images are combined in Maya and placed in a scene with directional lighting coming from above the painting and diffuse lighting coming from below the painting. This lighting configuration was chosen to mimic the setup used by the MoMA department of Imaging and Visual Resources. The renderings created were positioned next to the single-shot image taken by the photographer for ease




Fig. A.4. – Diffuse color image (top left), normal map (top right), and a close up of the same region of both the diffuse color and normal map (bottom left and right) of Painting 1.

of comparison. Close ups of sections of each painting are also shown of the photograph and the corresponding rendering in Figures A.6 and A.7. Rendered images are denoted by a red outline.

Sections of the rendered Painting 1 with dark blue paint contain dark pixels that correspond to the location of holes in the normal map. The renderings also contain more pronounced shadows and fewer highlights than the images taken by the photographer. These differences can be addressed by selecting different property settings in Maya such reflected color, specular color, and reflectivity. The selection of these properties is a time-intensive process that will be addressed in future work with the goal of identifying useful starting points for different objects with a diverse set of surface properties such as gloss level.




Fig. A.5. – Diffuse color image (top left), normal map (top right), and a close up of the same region of both the diffuse color and normal map (bottom left and right) of Painting 2.

The low gloss level inherent to Painting 2 simplified the rendering process considerably, by reducing the number of material properties that needed altering, compared to the rendering process of the glossy painting. The rendering produced very similar amounts of shadows and highlights to those present in the photograph. The texture of the canvas weave and thicker areas of paint are more evident in the renderings than the photograph. This is due to the incorporation of the normal map in the rendering software to computationally produce shadows based on surface topography of the painting. More diffuse lighting could be used




Fig. A.6. – Photograph and rendering (outlined in red) of Painting 1 with close ups of three different sections.

to render images with fewer, less pronounced shadows, if the amount of visible detail too pronounced.

Having separate diffuse and normal maps facilitates re-rendering to accentuate or reduce surface properties. This would produce images that are optimal for different applications, for example, publication, web delivery, documentation, and conservation. This is shown in Figure A.8 where the rendering on the left accentuates the number of shadows and highlights while the rendering on the right reduces the amount of visible texture by only changing the lighting.




Fig. A.7. – Photograph and rendering (outlined in red) of Painting 2 with close ups of two different sections.

The Four-Light Simplified technique was presented to the Department of Imaging and Visual Resources at the Museum of Modern Art in Queens, New York, in December 2014. The



Fig. A.8. – Left: Rendering of Painting 1 (Kirchner) with high levels of gloss and texture produced by using directional lighting in Maya. Right: Rendering of the same area of Painting 1 with limited gloss and texture produced by using ambient lighting in Maya.

technique was used to image two paintings provided by MoMA. The first painting (Kirchner's *Girl on Divan*) was a glossy work with significant impasto. The second painting (Picasso's *Vase of Flowers*) was a matte painting with very little impasto. The 4LIS and rendering workflow was compared to images of each painting taken by a skilled photographer using MoMA standard lighting setup. A comparison of the two techniques shows that the 4LIS system has deficiencies in capturing high gloss materials with dark, heavily textured regions resulting in noisy and less accurate renderings. The system does not have the same deficiencies when used with a less glossy and less textured work. Renderings of the second painting were successful representations of the artwork, shown by the comparisons in Figure A.7.

Appendix B: Relating Ward Model to Maya's Anisotropic Shader

Maya is equipped with many different shaders and rendering options. Some of these shaders are named after BRDF models, one of which is the anisotropic shader. The anisotropic shader is based on the Ward model for anisotropic materials, meaning this shader is intended to represent materials that have a grooved texture like that of brushed metal. The attributes of this shader can be set to influence the appearance of highlights, set the roughness and reflectivity of the material, and control the physical characteristics of grooves [Autodesk, Inc., 2014a]. For example, the shader can be used to represent an isotropic object by setting the Spread X and Spread Y values equal to each other. A screenshot showing some of the attributes of the anisotropic shader are shown in Figure B.1.

This shader was selected for the following experiments because it is believed that the attributes of this shader will better relate to Ward model (the model used to fit the BRDF data from the Artist Material database) parameters than other shaders. Because the samples composing the Artist Material Database are considered to be isotropic, Spread X and Spread Y will be set to the same value to produce the effect of an isotropic material. The attributes Spread X/Spread Y and Roughness are the main attributes of interest. It is expected that Spread will relate to the Ward parameter α and Roughness will relate to the Ward parameter ρ . When values for Spread X and Y are increased the specular highlight becomes smaller in size. When the value is decreased, the highlight spreads out. These attributes can be set to any value between 0.1 and 100.0. The roughness parameter can range between 0.01 to 1.0 and determines the overall roughness of the surface. The smaller the value the smoother the surface of the material being rendered. Small values correspond to concentrated highlights like those that can be seen on mirrored objects while larger values correspond to spread out highlights [Autodesk, Inc., 2014a].

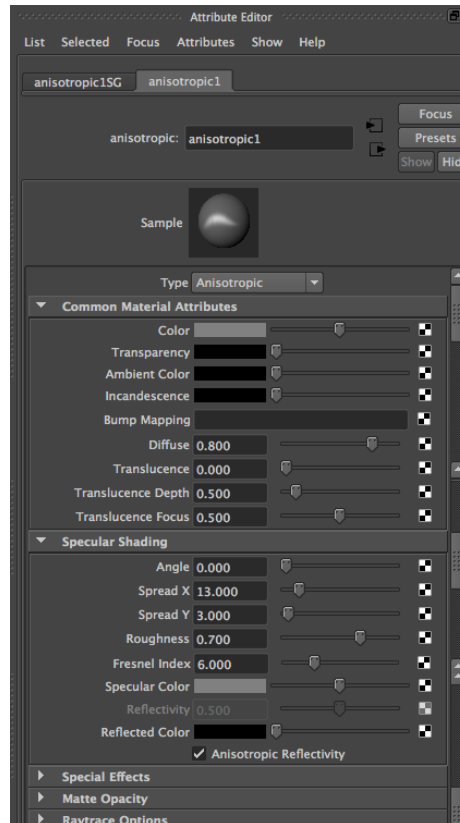


Fig. B.1. – A screenshot of the Attribute Editor menu in Maya showing some of the attributes associated with the anisotropic shader.

B.1 Experiment 1: Creating Stimuli

The Experiment 1 dataset was created by first selecting samples from the Artist Material Database that represented a range of BRDF measurements as well as a variety of mediums, substrates, application techniques and overcoats. A list of the samples selected and the BRDF parameters associated with each sample are recorded in Table B.1. There was a total of twelve different individual patches selected for the data set. The maximum and minimum rho values are 0.0830 and 0.0001 and the maximum and minimum alpha values are 0.2629 and 0.0146.

Each of the twelve different samples were then rendered using the same viewing and lighting geometry in Maya. The parameters SpreadX, SpreadY, and Roughness were varied incrementally while all other parameters were set to their default values. This resulted in

Tab. B.1. – Samples from the Artist Material Database and details about the medium, substrate, application technique, overcoat, and the Ward BRDF model parameters rho and alpha.

Sample	Medium	Substrate	Technique	Varnish	rho	alpha
1	oil	Masonite-Gessoed	Palette knife-"bumpy" impasto	None	0.0539	0.0235
2	oil	Masonite-Gessoed	Palette knife-"bumpy" impasto	Matte	0.0056	0.0146
3	acrylic	Canvas-Gessoed Coarse Linen	Palette knife-slight impasto	None	0.0202	0.0778
4	acrylic	Canvas-Gessoed Coarse Linen	Palette knife-slight impasto	Matte	0.0001	0.1825
5	acrylic	Canvas-Gessoed Coarse Linen	Palette knife-slight impasto	Gloss	0.0621	0.0212
6	acrylic	Masonite-Gessoed	Brushed-(vertically)	None	0.0195	0.0686
7	acrylic	Masonite-Gessoed	Brushed-(vertically)	Matte	0.0001	0.1825
8	acrylic	Masonite-Gessoed	Brushed-(vertically)	Gloss	0.0482	0.0212
9	acrylic	Canvas-Raw Coarse Linen	Palette knife-"bumpy" impasto	None	0.0188	0.06886
10	marker	Cold-Pressed Watercolor Paper	Colored-(all directions)	None	0.0310	0.2207
11	gouache	Hot-Pressed Watercolor Paper	Brushed-(horizontally)	None	0.0001	0.1825
12	ink	Hot-Pressed Watercolor Paper	Brushed-(horizontally)	None	0.0830	0.2629

Tab. B.2. – The spread values (1 to 100) and the roughness values (0.01 to 1.00) used to create 144 different renderings for each of the 12 samples.

Spread X/Y	1	5	10	20	30	40	50	60	70	80	90	100
Roughness	0.01	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0

144 renderings for each of the 12 samples. The spread values range from 1 to 100 and the roughness values range from 0.01 to 1.00. A list of these values are shown in Table B.2.

B.2 Experiment 1: Procedure

In a darkened room, observers were presented with a reference stimuli (top) and a test stimuli (bottom) displayed on an Apple iMac 27-inch display with a middle gray background. The test stimuli was changed by moving a slider bar from left to right, displayed underneath the stimuli on the screen. After an appropriate match had been made between the test and reference stimuli, observers were asked to rank how well the test stimuli matched the reference stimuli by moving a second slider bar ranging from 1 to 7 where 1 represents "Not a Good Match" and 7 represents a "Superior Match" . The reference stimuli was an image taken with strobe to the left and right of the AMD panel that had been rotated 90 degrees counter-clockwise. Images of the 12 different reference stimuli are shown in Figure B.2. A diagram of the lighting used to capture images of each panel is shown in Figure B.3.

Due to the large number of samples, the dataset was divided into two different set. During the first set of the experiment, observers were asked to adjust a slider that represented the Spread value while the Roughness parameter was held constant. The second set asked observers to adjust Roughness while Spread was held constant. A screen shot of the experiment interface is shown in Figure B.4.

Five observers participated in each part with two participating in both parts. Resulting in a total of eight different observers. The data collected was used to select meaningful pairs for comparison as part of Experiment 2. Any sample that received a rating of less than 3.5 was eliminated from the set of samples that were to be used in Experiment 2. This method of elimination created an uneven number of renderings to be considered for each reference stimuli. For example, the ink on watercolor paper sample had fifteen renderings




Fig. B.2. – Images of each reference stimuli taken with two strobes firing simultaneously.

that rated higher than 3.5 while the acrylic on raw linen had nine renderings considered to be an acceptable match.

B.3 Experiment 2: Procedure

Experiment 2 used the method of paired comparison. In this method, an observer is presented a pair to compare to a reference stimuli. The observer is then asked which of the two samples has more of the property of interest [Engeldrum, 2000]. In this experiment, a random sequence of pairs is presented to an observer that then selects which of the pair best represents the gloss level of the reference stimuli. In this experiment, the reference stimuli are the panels from the AMD displayed with an LED light source illuminating the samples from above.

The experimental set up is shown in Figure B.5, where the test samples are displayed on the right on the iMac and the reference stimuli are located on the left on a thirty degree




Fig. B.3. – A diagram of the lighting used to image each AMD panel. Light coming from the left, illuminates from the bottom of the rotated panel with a diffuser, while the Light from the right is stronger.

sample holder. The test stimuli were randomized in which order they appeared and in which location (left versus right). This is done to eliminate as much bias as possible that might be created based on the location of the stimuli. All of the renderings that rated high enough in Experiment 1 were included in Experiment 2. Experiment 2 was also divided into two tasks, due to the large number of necessary comparisons. Each observer was asked to participate in both tasks.




Fig. B.4. – A screenshot of the Experiment 1 interface where observers are presented an reference stimuli (top) and an adjustable test stimuli (bottom).

B.4 Experiment 2: Results

Fourteen observers participated in Experiment 2. A matrix of observer responses was created for each of the twelve different samples. These matrices are referred to as frequency matrices. The frequency matrix for Sample 12 is shown in Table B.3. This matrix reports how many times sample i is selected over sample j by recording the number of selections in the corresponding column i and row j [Engeldrum, 2000].

The frequency matrices were then used to create a second set of matrices of proportions, an example of which is shown in Table B.4. These proportion matrices were derived by dividing each element of the frequency matrix by the number of observers that compared each pair [Engeldrum, 2000]. The proportion matrices were then used to create z-score matrices. These z-score matrices are used to derive the psychophysical scale for the compared renderings. An example matrix and interval scale is shown in Table B.5.

The frequency, proportion, and z-score matrices for each of the twelve samples can be found in the Appendix C.

The z-score matrices are used to calculate the psychophysical scale by averaging the z-score values in each column, according to Case V of Thurstone's Law of Comparative Judgement

Tab. B.3. – A frequency matrix for sample 4 that shows how many times sample k was selected over sample j .

		patch k									
		1	2	3	4	5	6	7	8	9	10
patch j	1	0	6	9	7	10	9	7	4	4	4
	2	8	0	12	4	11	10	9	8	8	7
	3	5	2	0	3	4	9	9	6	4	7
	4	7	10	11	0	12	11	11	8	8	8
	5	4	3	10	2	0	11	9	7	5	5
	6	5	4	5	3	3	0	11	6	5	5
	7	7	5	5	3	5	3	0	4	4	5
	8	10	6	8	6	7	8	10	0	9	9
	9	10	6	10	6	9	9	10	5	0	6
	10	10	7	7	6	9	9	9	5	8	0

Tab. B.4. – A probability matrix for sample 4 that shows how many times sample k was selected over sample j .

		patch k									
		1	2	3	4	5	6	7	8	9	10
patch j	1	0.50	0.43	0.64	0.50	0.71	0.64	0.50	0.29	0.29	0.29
	2	0.57	0.50	0.86	0.29	0.79	0.71	0.64	0.57	0.57	0.50
	3	0.36	0.14	0.50	0.21	0.29	0.64	0.64	0.43	0.29	0.50
	4	0.50	0.71	0.79	0.50	0.86	0.79	0.79	0.57	0.57	0.57
	5	0.29	0.21	0.71	0.14	0.50	0.79	0.64	0.50	0.36	0.36
	6	0.36	0.29	0.36	0.21	0.21	0.50	0.79	0.43	0.36	0.36
	7	0.50	0.36	0.36	0.21	0.36	0.21	0.50	0.29	0.29	0.36
	8	0.71	0.43	0.57	0.43	0.50	0.57	0.71	0.50	0.64	0.64
	9	0.71	0.43	0.71	0.43	0.64	0.64	0.71	0.36	0.50	0.43
	10	0.71	0.50	0.50	0.43	0.64	0.64	0.64	0.36	0.57	0.50

Tab. B.5. – A matrix for sample 4 that shows z-scores and an interval scale for how often k was selected over sample j .

		patch k									
		1	2	3	4	5	6	7	8	9	10
patch j	1	0.00	-0.18	0.37	0.00	0.57	0.37	0.00	-0.57	-0.57	-0.57
	2	0.18	0.00	1.07	-0.57	0.79	0.57	0.37	0.18	0.18	0.00
	3	-0.37	-1.07	0.00	-0.79	-0.57	0.37	0.37	-0.18	-0.57	0.00
	4	0.00	0.57	0.79	0.00	1.07	0.79	0.79	0.18	0.18	0.18
	5	-0.57	-0.79	0.57	-1.07	0.00	0.79	0.37	0.00	-0.37	-0.37
	6	-0.37	-0.57	-0.37	-0.79	-0.79	0.00	0.79	-0.18	-0.37	-0.37
	7	0.00	-0.37	-0.37	-0.79	-0.37	-0.79	0.00	-0.57	-0.57	-0.37
	8	0.57	-0.18	0.18	-0.18	0.00	0.18	0.57	0.00	0.37	0.37
	9	0.57	-0.18	0.57	-0.18	0.37	0.37	0.57	-0.37	0.00	-0.18
	10	0.57	0.00	0.00	-0.18	0.37	0.37	0.37	-0.37	0.18	0.00
interval scale		0.06	-0.28	0.28	-0.45	0.14	0.30	0.42	-0.19	-0.15	-0.13




Fig. B.5. – Experimental setup with a physical sample presented under an LED source on the left to be compared to renderings displayed on the screen on the right.

[Montag, 2006]. A confidence interval was then computed according to Equation B.1. This is a 95% confidence interval around each scale value. σ_{obs} is the standard deviation of observations [Montag, 2006].

$$Cl = 1.96\sigma_{obs} \quad (B.1)$$

Case V is not appropriate for all datasets. The Case V calculations rely on there not being a significant difference between stimuli. To assess if using Case V was appropriate for each sample's set of test stimuli, a chi-square test was performed. If the calculated chi-square value is less than the critical chi-square value, the use of Case V is considered valid [Engeldrum, 2000].

This test is performed by first calculating estimated z-scores from the psychophysical scale. This is achieved by subtracting each scale value from one another. Estimated probabilities are then calculated from the estimated z-score values. The arcsine of these estimated probabilities and the measured probabilities is then computed. This calculation changes the proportions into normal deviates, that have a mean value of zero and a constant variance of 0.1 or 1 divided by the number of observers [Engeldrum, 2000]. The difference of the arcsine of the measured and estimated probabilities is calculated and then added together. This value is then multiplied by the number of observations made of each pair resulting in the chi-squared calculated value.

The critical chi-squared value is calculated for a 95% confidence interval with a given number of degrees of freedom. The number of degrees is calculated by Equation B.2 where n is the number of samples [Engeldrum, 2000]. The chi-squared value and the critical value are then compared. A list of calculated and critical chi-squared values are shown in Table B.6. Samples 2, 3, 8, 11, and 12 all failed the chi-squared test. The individual images used for comparison in this experiment contained images that were too different from the other samples and therefore were never selected by the observers as a better match to the physical sample. If all of these samples were to be removed from the data set, Samples 2, 3, 8, 11, and 12 would pass the chi-squared test.

$$df = (n - 1)(n - 2)/2 \quad (\text{B.2})$$

Tab. B.6. – My caption

sample	1	2	3	4	5	6	7	8	9	10	11	12
calculated χ^2	41.06	54.66	63.68	33.15	127.38	13.34	13.6	57.8	39.85	47.88	109.24	162.36
critical χ^2	41.34	41.34	51	51	137.7	25	41.34	32.67	41.34	73.31	99.62	114.27

Plots of the relative scales created for each of the seven samples that passed the chi-squared test are shown in Figures B.6, B.7, B.8, B.9, B.10, B.11, and B.12.




Fig. B.6. – Scale values calculated from z-scores for sample 1.




Fig. B.7. – Scale values calculated from z-scores for sample 4.

Sample 5




Fig. B.8. – Scale values calculated from z-scores for sample 5.

Sample 6




Fig. B.9. – Scale values calculated from z-scores for sample 6.

B.5 Conclusions

As indicated by the numerous data points with overlapping error bars, there is no clear set of parameters that better represents the perceived gloss level of each sample. There are, however, samples that can be eliminated from the set of representative samples.

This experiment could be improved with the addition of many more observers or by selecting fewer images from the first part of the experiment by limiting the approval rating to a much higher threshold. The results of this experiment in its current state, indicates that the scale for which spread and roughness can be altered as a part of the anisotropic shader in the rendering software is too fine to generate images that vary significantly for the samples selected for rendering. It is unclear that further investigation would be beneficial to future users of Maya trying to render paintings with a known amount of gloss.

Sample 7




Fig. B.10. – Scale values calculated from z-scores for sample 7.

Sample 9




Fig. B.11. – Scale values calculated from z-scores for sample 9.

Sample 10




Fig. B.12. – Scale values calculated from z-scores for sample 10.

Appendix C: Supplementary Materials

C.1 Frequency Matrices

Twelve different samples were compared to a varying number of pairs of rendered images. The following matrices represent how many times one rendering was selected over another for each sample. In these matrices renderings corresponding to column k were selected over renderings corresponding to row j .

Tab. C.1. – A frequency matrix for sample 1 that shows how many times sample k was selected over sample j .

		patch k								
		1	2	3	4	5	6	7	8	9
patch j	1	0	13	13	13	13	13	12	13	13
	2	1	0	2	8	7	6	6	7	6
	1c 3	1	12	0	13	11	13	13	12	13
	4	1	6	1	0	10	4	13	11	10
	5	1	7	3	4	0	5	9	8	11
	6	1	8	1	10	9	0	9	13	11
	7	2	8	1	1	5	5	0	8	7
	8	1	7	2	3	6	1	6	0	3
	9	1	8	1	4	3	3	7	11	0

Tab. C.2. – A frequency matrix for sample 2 that shows how many times sample k was selected over sample j .

		patch k								
		1	2	3	4	5	6	7	8	9
patch j	1	0	4	3	3	7	10	7	7	6
	2	10	0	6	12	11	9	10	10	7
	3	11	8	0	9	10	7	8	8	10
	4	11	2	5	0	8	5	9	7	10
	5	7	3	4	6	0	8	6	4	7
	6	4	5	7	9	6	0	3	3	13
	7	7	4	6	5	8	11	0	12	12
	8	7	4	6	7	10	11	2	0	12
	9	8	7	4	4	7	1	2	2	0

Tab. C.3. – A frequency matrix for sample 3 that shows how many times sample k was selected over sample j .

		patch k									
		1	2	3	4	5	6	7	8	9	10
patch j	1	0	2	6	6	1	9	3	1	1	1
	2	12	0	11	7	2	11	5	1	8	1
	3	8	3	0	5	2	6	4	2	2	0
	4	8	7	9	0	3	11	1	0	8	1
	5	13	12	12	11	0	11	6	5	9	3
	6	5	3	8	3	3	0	1	1	3	1
	7	11	9	10	13	8	13	0	2	4	2
	8	13	13	12	14	9	13	12	0	11	2
	9	13	6	12	6	5	11	10	3	0	0
	10	13	13	14	13	11	13	12	12	14	0

Tab. C.4. – A frequency matrix for sample 4 that shows how many times sample k was selected over sample j .

		patch k									
		1	2	3	4	5	6	7	8	9	10
patch j	1	0	6	9	7	10	9	7	4	4	4
	2	8	0	12	4	11	10	9	8	8	7
	3	5	2	0	3	4	9	9	6	4	7
	4	7	10	11	0	12	11	11	8	8	8
	5	4	3	10	2	0	11	9	7	5	5
	6	5	4	5	3	3	0	11	6	5	5
	7	7	5	5	3	5	3	0	4	4	5
	8	10	6	8	6	7	8	10	0	9	9
	9	10	6	10	6	9	9	10	5	0	6
	10	10	7	7	6	9	9	9	5	8	0

Tab. C.5. – A frequency matrix for sample 5 that shows how many times sample k was selected over sample j .

		patch k															
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
patch j	1	0	4	3	12	11	7	9	11	10	8	12	7	10	9	8	10
	2	10	0	12	10	12	14	14	13	12	13	11	13	10	12	13	12
	3	11	2	0	14	12	13	13	13	13	12	10	11	12	13	13	10
	4	2	4	0	0	4	4	4	6	6	4	6	8	2	4	3	3
	5	3	2	2	10	0	5	5	9	6	3	5	7	4	6	5	4
	6	7	0	1	10	9	0	8	8	5	7	10	9	9	6	7	7
	7	5	0	1	10	9	6	0	2	7	2	1	4	4	4	5	2
	8	3	1	1	8	5	6	12	0	6	6	2	5	2	1	3	2
	9	4	2	1	8	8	9	7	8	0	3	6	8	7	8	6	4
	10	6	1	2	10	11	7	12	8	11	0	8	10	9	8	7	6
	11	2	3	4	8	9	4	13	12	8	6	0	9	5	6	2	3
	12	7	1	3	6	7	5	10	9	6	4	5	0	2	7	6	5
	13	4	4	2	12	10	5	10	12	7	5	9	12	0	7	7	5
	14	5	2	1	10	8	8	10	13	6	6	8	7	7	0	5	6
	15	6	1	1	11	9	7	9	11	8	7	12	8	7	9	0	6
	16	4	2	4	11	10	7	12	12	10	8	11	9	9	8	8	0

Tab. C.6. – A frequency matrix for sample 6 that shows how many times sample k was selected over sample j .

		patch k						
		1	2	3	4	5	6	7
patch j	1	0	0	12	6	11	7	12
	2	14	0	14	14	14	14	14
	3	2	0	0	2	8	1	7
	4	8	0	12	0	14	13	13
	5	3	0	6	0	0	3	8
	6	7	0	13	1	11	0	12
	7	2	0	7	1	6	2	0

Tab. C.7. – A frequency matrix for sample 7 that shows how many times sample k was selected over sample j .

		patch k								
		1	2	3	4	5	6	7	8	9
patch j	1	0	7	8	7	8	5	8	10	6
	2	7	0	6	7	5	7	6	7	6
	3	6	8	0	6	8	7	7	6	9
	4	7	7	8	0	8	5	6	7	7
	5	6	9	6	6	0	7	5	5	9
	6	9	7	7	9	7	0	9	8	7
	7	6	8	7	8	9	5	0	7	9
	8	4	7	8	7	9	6	7	0	6
	9	8	8	5	7	5	7	5	8	0

Tab. C.8. – A frequency matrix for sample 8 that shows how many times sample k was selected over sample j .

		patch k							
		1	2	3	4	5	6	7	8
patch j	1	0	10	10	14	14	13	14	7
	2	4	0	2	10	8	8	5	0
	3	4	12	0	14	14	14	13	3
	4	0	4	0	0	6	3	1	2
	5	0	6	0	8	0	1	1	1
	6	1	6	0	11	13	0	2	2
	7	0	9	1	13	13	12	0	6
	8	7	14	11	12	13	12	8	0

Tab. C.9. – A frequency matrix for sample 9 that shows how many times sample k was selected over sample j .

		patch k								
		1	2	3	4	5	6	7	8	9
patch j	1	0	2	13	11	10	9	10	9	0
	2	12	0	13	13	9	13	13	9	1
	3	1	1	0	10	2	8	7	9	0
	4	3	1	4	0	3	7	11	6	0
	5	4	5	12	11	0	11	11	11	0
	6	5	1	6	7	3	0	6	7	0
	7	4	1	7	3	3	8	0	9	1
	8	5	5	5	8	3	7	5	0	1
	9	14	13	14	14	14	14	13	13	0

Tab. C.10. – A frequency matrix for sample 10 that shows how many times sample k was selected over sample j .

		patch k											
		1	2	3	4	5	6	7	8	9	10	11	12
patch j	1	0	7	5	8	5	5	7	5	5	2	3	5
	2	7	0	7	7	8	7	6	9	8	7	6	7
	3	9	7	0	5	5	4	6	7	5	3	4	5
	4	6	7	9	0	4	4	4	6	5	5	7	6
	5	9	6	9	10	0	3	5	7	8	6	4	4
	6	9	7	10	10	11	0	5	6	2	4	3	7
	7	7	8	8	10	9	9	0	5	5	6	2	5
	8	9	5	7	8	7	8	9	0	4	4	2	4
	9	9	6	9	9	6	12	9	10	0	6	6	7
	10	12	7	11	9	8	10	8	10	8	0	4	7
	11	11	8	10	7	10	11	12	12	8	10	0	8
	12	9	7	9	8	10	7	9	10	7	7	6	0

Tab. C.11. – A frequency matrix for sample 11 that shows how many times sample k was selected over sample j .

		patch k													
		1	2	3	4	5	6	7	8	9	10	11	12	13	14
patch j	1	0	13	11	12	12	13	12	13	13	12	11	10	13	13
	2	1	0	8	7	5	8	4	3	5	1	1	2	1	4
	3	3	6	0	10	9	6	4	3	6	0	2	1	3	1
	4	2	7	4	0	7	5	8	4	4	1	0	2	0	2
	5	2	9	5	7	0	4	6	3	4	1	2	0	1	2
	6	1	6	8	9	10	0	6	4	9	1	0	1	2	0
	7	2	10	10	6	8	8	0	4	7	1	2	3	1	6
	8	1	11	11	10	11	10	10	0	9	6	6	5	9	5
	9	1	9	8	10	10	5	7	5	0	3	3	2	2	1
	10	2	13	14	13	13	13	13	8	11	0	7	6	10	10
	11	3	14	12	14	12	14	12	8	11	7	0	8	11	8
	12	4	12	13	12	14	13	11	9	12	8	6	0	10	10
	13	1	13	11	14	13	12	13	5	12	4	3	4	0	5
	14	1	10	13	12	12	14	8	9	13	4	6	4	9	0

Tab. C.12. – A frequency matrix for sample 12 that shows how many times sample k was selected over sample j .

		patch k														
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
patch j	1	0	9	9	8	9	7	10	8	5	5	3	3	4	2	4
	2	5	0	5	4	8	6	3	6	2	2	0	0	0	1	0
	3	5	9	0	5	9	7	5	6	4	3	1	1	0	0	0
	4	6	10	9	0	11	8	4	4	3	6	0	0	1	1	0
	5	5	6	5	3	0	4	6	6	7	5	1	0	0	0	0
	6	7	8	7	6	10	0	7	8	2	4	0	1	0	1	1
	7	4	11	9	10	8	7	0	5	5	6	1	1	0	0	0
	8	6	8	8	10	8	6	9	0	5	3	0	1	0	0	0
	9	9	12	10	11	7	12	9	9	0	8	2	2	3	4	2
	10	9	12	11	8	9	10	8	11	6	0	3	4	2	1	2
	11	11	14	13	14	13	14	13	14	12	11	0	10	10	11	14
	12	11	14	13	14	14	13	13	13	12	10	4	0	8	11	12
	13	10	14	14	13	14	14	14	14	11	12	4	6	0	7	12
	14	12	13	14	13	14	13	14	14	10	13	3	3	7	0	10
	15	10	14	14	14	14	13	14	14	12	12	0	2	2	4	0

C.2 Probability Matrices

Twelve different samples were compared to a varying number of pairs of rendered images. The following matrices of probability were calculated from the frequency matrices by dividing each frequency by the number of observations made of each corresponding pair. In these matrices renderings corresponding to column k were selected over renderings corresponding to row j .

Tab. C.13. – A probability matrix for sample 1 that shows how many times sample k was selected over sample j .

	patch k								
	1	2	3	4	5	6	7	8	9
1	0.50	0.93	0.93	0.93	0.93	0.93	0.86	0.93	0.93
2	0.07	0.50	0.14	0.57	0.50	0.43	0.43	0.50	0.43
3	0.07	0.86	0.50	0.93	0.79	0.93	0.93	0.86	0.93
4	0.07	0.43	0.07	0.50	0.71	0.29	0.93	0.79	0.71
patch j 5	0.07	0.50	0.21	0.29	0.50	0.36	0.64	0.57	0.79
6	0.07	0.57	0.07	0.71	0.64	0.50	0.64	0.93	0.79
7	0.14	0.57	0.07	0.07	0.36	0.36	0.50	0.57	0.50
8	0.07	0.50	0.14	0.21	0.43	0.07	0.43	0.50	0.21
9	0.07	0.57	0.07	0.29	0.21	0.21	0.50	0.79	0.50

Tab. C.14. – A probability matrix for sample 2 that shows how many times sample k was selected over sample j .

	patch k								
	1	2	3	4	5	6	7	8	9
1	0.50	0.29	0.21	0.21	0.50	0.71	0.50	0.50	0.43
2	0.71	0.50	0.43	0.86	0.79	0.64	0.71	0.71	0.50
3	0.79	0.57	0.50	0.64	0.71	0.50	0.57	0.57	0.71
4	0.79	0.14	0.36	0.50	0.57	0.36	0.64	0.50	0.71
patch j 5	0.50	0.21	0.29	0.43	0.50	0.57	0.43	0.29	0.50
6	0.29	0.36	0.50	0.64	0.43	0.50	0.21	0.21	0.93
7	0.50	0.29	0.43	0.36	0.57	0.79	0.50	0.86	0.86
8	0.50	0.29	0.43	0.50	0.71	0.79	0.14	0.50	0.86
9	0.57	0.50	0.29	0.29	0.50	0.07	0.14	0.14	0.50

Tab. C.15. – A probability matrix for sample 3 that shows how many times sample k was selected over sample j .

		patch k									
		1	2	3	4	5	6	7	8	9	10
patch j	1	0.50	0.14	0.43	0.43	0.07	0.64	0.21	0.07	0.07	0.07
	2	0.86	0.50	0.79	0.50	0.14	0.79	0.36	0.07	0.57	0.07
	3	0.57	0.21	0.50	0.36	0.14	0.43	0.29	0.14	0.14	0.00
	4	0.57	0.50	0.64	0.50	0.21	0.79	0.07	0.00	0.57	0.07
	5	0.93	0.86	0.86	0.79	0.50	0.79	0.43	0.36	0.64	0.21
	6	0.36	0.21	0.57	0.21	0.21	0.50	0.07	0.07	0.21	0.07
	7	0.79	0.64	0.71	0.93	0.57	0.93	0.50	0.14	0.29	0.14
	8	0.93	0.93	0.86	1.00	0.64	0.93	0.86	0.50	0.79	0.14
	9	0.93	0.43	0.86	0.43	0.36	0.79	0.71	0.21	0.50	0.00
	10	0.93	0.93	1.00	0.93	0.79	0.93	0.86	0.86	1.00	0.50

Tab. C.16. – A probability matrix for sample 4 that shows how many times sample k was selected over sample j .

		patch k									
		1	2	3	4	5	6	7	8	9	10
patch j	1	0.50	0.43	0.64	0.50	0.71	0.64	0.50	0.29	0.29	0.29
	2	0.57	0.50	0.86	0.29	0.79	0.71	0.64	0.57	0.57	0.50
	3	0.36	0.14	0.50	0.21	0.29	0.64	0.64	0.43	0.29	0.50
	4	0.50	0.71	0.79	0.50	0.86	0.79	0.79	0.57	0.57	0.57
	5	0.29	0.21	0.71	0.14	0.50	0.79	0.64	0.50	0.36	0.36
	6	0.36	0.29	0.36	0.21	0.21	0.50	0.79	0.43	0.36	0.36
	7	0.50	0.36	0.36	0.21	0.36	0.21	0.50	0.29	0.29	0.36
	8	0.71	0.43	0.57	0.43	0.50	0.57	0.71	0.50	0.64	0.64
	9	0.71	0.43	0.71	0.43	0.64	0.64	0.71	0.36	0.50	0.43
	10	0.71	0.50	0.50	0.43	0.64	0.64	0.64	0.36	0.57	0.50

Tab. C.17. – A probability matrix for sample 5 that shows how many times sample k was selected over sample j .

		patch k															
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
patch j	1	0.50	0.29	0.21	0.86	0.79	0.50	0.64	0.79	0.71	0.57	0.86	0.50	0.71	0.64	0.57	0.71
	2	0.71	0.50	0.86	0.71	0.86	1.00	1.00	0.93	0.86	0.93	0.79	0.93	0.71	0.86	0.93	0.86
	3	0.79	0.14	0.50	1.00	0.86	0.93	0.93	0.93	0.93	0.86	0.71	0.79	0.86	0.93	0.93	0.71
	4	0.14	0.29	0.00	0.50	0.29	0.29	0.29	0.43	0.43	0.29	0.43	0.57	0.14	0.29	0.21	0.21
	5	0.21	0.14	0.14	0.71	0.50	0.36	0.36	0.64	0.43	0.21	0.36	0.50	0.29	0.43	0.36	0.29
	6	0.50	0.00	0.07	0.71	0.64	0.50	0.57	0.57	0.36	0.50	0.71	0.64	0.64	0.43	0.50	0.50
	7	0.36	0.00	0.07	0.71	0.64	0.43	0.50	0.14	0.50	0.14	0.07	0.29	0.29	0.29	0.36	0.14
	8	0.21	0.07	0.07	0.57	0.36	0.43	0.86	0.50	0.43	0.43	0.14	0.36	0.14	0.07	0.21	0.14
	9	0.29	0.14	0.07	0.57	0.57	0.64	0.50	0.57	0.50	0.21	0.43	0.57	0.50	0.57	0.43	0.29
	10	0.43	0.07	0.14	0.71	0.79	0.50	0.86	0.57	0.79	0.50	0.57	0.71	0.64	0.57	0.50	0.43
	11	0.14	0.21	0.29	0.57	0.64	0.29	0.93	0.86	0.57	0.43	0.50	0.64	0.36	0.43	0.14	0.21
	12	0.50	0.07	0.21	0.43	0.50	0.36	0.71	0.64	0.43	0.29	0.36	0.50	0.14	0.50	0.43	0.36
	13	0.29	0.29	0.14	0.86	0.71	0.36	0.71	0.86	0.50	0.36	0.64	0.86	0.50	0.50	0.50	0.36
	14	0.36	0.14	0.07	0.71	0.57	0.57	0.71	0.93	0.43	0.43	0.57	0.50	0.50	0.50	0.36	0.43
	15	0.43	0.07	0.07	0.79	0.64	0.50	0.64	0.79	0.57	0.50	0.86	0.57	0.50	0.64	0.50	0.43
	16	0.29	0.14	0.29	0.79	0.71	0.50	0.86	0.86	0.71	0.57	0.79	0.64	0.64	0.57	0.57	0.50

Tab. C.18. – A probability matrix for sample 6 that shows how many times sample k was selected over sample j .

		patch k						
		1	2	3	4	5	6	7
patch j	1	0.50	0.00	0.86	0.43	0.79	0.50	0.86
	2	1.00	0.50	1.00	1.00	1.00	1.00	1.00
	3	0.14	0.00	0.50	0.14	0.57	0.07	0.50
	4	0.57	0.00	0.86	0.50	1.00	0.93	0.93
	5	0.21	0.00	0.43	0.00	0.50	0.21	0.57
	6	0.50	0.00	0.93	0.07	0.79	0.50	0.86
	7	0.14	0.00	0.50	0.07	0.43	0.14	0.50

Tab. C.19. – A probability matrix for sample 7 that shows how many times sample k was selected over sample j .

		patch k								
		1	2	3	4	5	6	7	8	9
patch j	1	0.50	0.50	0.57	0.50	0.57	0.36	0.57	0.71	0.43
	2	0.50	0.50	0.43	0.50	0.36	0.50	0.43	0.50	0.43
	3	0.43	0.57	0.50	0.43	0.57	0.50	0.50	0.43	0.64
	4	0.50	0.50	0.57	0.50	0.57	0.36	0.43	0.50	0.50
	5	0.43	0.64	0.43	0.43	0.50	0.50	0.36	0.36	0.64
	6	0.64	0.50	0.50	0.64	0.50	0.50	0.64	0.57	0.50
	7	0.43	0.57	0.50	0.57	0.64	0.36	0.50	0.50	0.64
	8	0.29	0.50	0.57	0.50	0.64	0.43	0.50	0.50	0.43
	9	0.57	0.57	0.36	0.50	0.36	0.50	0.36	0.57	0.50

Tab. C.20. – A probability matrix for sample 8 that shows how many times sample k was selected over sample j .

		patch k							
		1	2	3	4	5	6	7	8
patch j	1	0.50	0.71	0.71	1.00	1.00	0.93	1.00	0.50
	2	0.29	0.50	0.14	0.71	0.57	0.57	0.36	0.00
	3	0.29	0.86	0.50	1.00	1.00	1.00	0.93	0.21
	4	0.00	0.29	0.00	0.50	0.43	0.21	0.07	0.14
	5	0.00	0.43	0.00	0.57	0.50	0.07	0.07	0.07
	6	0.07	0.43	0.00	0.79	0.93	0.50	0.14	0.14
	7	0.00	0.64	0.07	0.93	0.93	0.86	0.50	0.43
	8	0.50	1.00	0.79	0.86	0.93	0.86	0.57	0.50

Tab. C.21. – A probability matrix for sample 9 that shows how many times sample k was selected over sample j .

		patch k								
		1	2	3	4	5	6	7	8	9
patch j	1	0.50	0.14	0.93	0.79	0.71	0.64	0.71	0.64	0.00
	2	0.86	0.50	0.93	0.93	0.64	0.93	0.93	0.64	0.07
	3	0.07	0.07	0.50	0.71	0.14	0.57	0.50	0.64	0.00
	4	0.21	0.07	0.29	0.50	0.21	0.50	0.79	0.43	0.00
	5	0.29	0.36	0.86	0.79	0.50	0.79	0.79	0.79	0.00
	6	0.36	0.07	0.43	0.50	0.21	0.50	0.43	0.50	0.00
	7	0.29	0.07	0.50	0.21	0.21	0.57	0.50	0.64	0.07
	8	0.36	0.36	0.36	0.57	0.21	0.50	0.36	0.50	0.07
	9	1.00	0.93	1.00	1.00	1.00	1.00	0.93	0.93	0.50

Tab. C.22. – A probability matrix for sample 10 that shows how many times sample k was selected over sample j .

		patch k											
		1	2	3	4	5	6	7	8	9	10	11	12
patch j	1	0.50	0.50	0.36	0.57	0.36	0.36	0.50	0.36	0.36	0.14	0.21	0.36
	2	0.50	0.50	0.50	0.50	0.57	0.50	0.43	0.64	0.57	0.50	0.43	0.50
	3	0.64	0.50	0.50	0.36	0.36	0.29	0.43	0.50	0.36	0.21	0.29	0.36
	4	0.43	0.50	0.64	0.50	0.29	0.29	0.29	0.43	0.36	0.36	0.50	0.43
	5	0.64	0.43	0.64	0.71	0.50	0.21	0.36	0.50	0.57	0.43	0.29	0.29
	6	0.64	0.50	0.71	0.71	0.79	0.50	0.36	0.43	0.14	0.29	0.21	0.50
	7	0.50	0.57	0.57	0.71	0.64	0.64	0.50	0.36	0.36	0.43	0.14	0.36
	8	0.64	0.36	0.50	0.57	0.50	0.57	0.64	0.50	0.29	0.29	0.14	0.29
	9	0.64	0.43	0.64	0.64	0.43	0.86	0.64	0.71	0.50	0.43	0.43	0.50
	10	0.86	0.50	0.79	0.64	0.57	0.71	0.57	0.71	0.57	0.50	0.29	0.50
	11	0.79	0.57	0.71	0.50	0.71	0.79	0.86	0.86	0.57	0.71	0.50	0.57
	12	0.64	0.50	0.64	0.57	0.71	0.50	0.64	0.71	0.50	0.50	0.43	0.50

Tab. C.23. – A probability matrix for sample 11 that shows how many times sample k was selected over sample j .

		patch k													
		1	2	3	4	5	6	7	8	9	10	11	12	13	14
patch j	1	0.50	0.93	0.79	0.86	0.86	0.93	0.86	0.93	0.93	0.86	0.79	0.71	0.93	0.93
	2	0.07	0.50	0.57	0.50	0.36	0.57	0.29	0.21	0.36	0.07	0.07	0.14	0.07	0.29
	3	0.21	0.43	0.50	0.71	0.64	0.43	0.29	0.21	0.43	0.00	0.14	0.07	0.21	0.07
	4	0.14	0.50	0.29	0.50	0.50	0.36	0.57	0.29	0.29	0.07	0.00	0.14	0.00	0.14
	5	0.14	0.64	0.36	0.50	0.50	0.29	0.43	0.21	0.29	0.07	0.14	0.00	0.07	0.14
	6	0.07	0.43	0.57	0.64	0.71	0.50	0.43	0.29	0.64	0.07	0.00	0.07	0.14	0.00
	7	0.14	0.71	0.71	0.43	0.57	0.57	0.50	0.29	0.50	0.07	0.14	0.21	0.07	0.43
	8	0.07	0.79	0.79	0.71	0.79	0.71	0.71	0.50	0.64	0.43	0.43	0.36	0.64	0.36
	9	0.07	0.64	0.57	0.71	0.71	0.36	0.50	0.36	0.50	0.21	0.21	0.14	0.14	0.07
	10	0.14	0.93	1.00	0.93	0.93	0.93	0.93	0.57	0.79	0.50	0.50	0.43	0.71	0.71
	11	0.21	0.93	0.86	1.00	0.86	1.00	0.86	0.57	0.79	0.50	0.50	0.57	0.79	0.57
	12	0.29	0.86	0.93	0.86	1.00	0.93	0.79	0.64	0.86	0.57	0.43	0.50	0.71	0.71
	13	0.07	0.93	0.79	1.00	0.93	0.86	0.93	0.36	0.86	0.29	0.21	0.29	0.50	0.36
	14	0.07	0.71	0.93	0.86	0.86	1.00	0.57	0.64	0.93	0.29	0.43	0.29	0.64	0.50

Tab. C.24. – A probability matrix for sample 12 that shows how many times sample k was selected over sample j .

		patch k														
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
patch j	1	0.50	0.64	0.64	0.57	0.64	0.50	0.71	0.57	0.36	0.36	0.21	0.21	0.29	0.14	0.29
	2	0.36	0.50	0.36	0.29	0.57	0.43	0.21	0.43	0.14	0.14	0.00	0.00	0.00	0.07	0.00
	3	0.36	0.64	0.50	0.36	0.64	0.50	0.36	0.43	0.29	0.21	0.07	0.07	0.00	0.00	0.00
	4	0.43	0.71	0.64	0.50	0.79	0.57	0.29	0.29	0.21	0.43	0.00	0.00	0.07	0.07	0.00
	5	0.36	0.43	0.36	0.21	0.50	0.29	0.43	0.43	0.50	0.36	0.07	0.00	0.00	0.00	0.00
	6	0.50	0.57	0.50	0.43	0.71	0.50	0.50	0.57	0.14	0.29	0.00	0.07	0.00	0.07	0.07
	7	0.29	0.79	0.64	0.71	0.57	0.50	0.50	0.36	0.36	0.43	0.07	0.07	0.00	0.00	0.00
	8	0.43	0.57	0.57	0.71	0.57	0.43	0.64	0.50	0.36	0.21	0.00	0.07	0.00	0.00	0.00
	9	0.64	0.86	0.71	0.79	0.50	0.86	0.64	0.64	0.50	0.57	0.14	0.14	0.21	0.29	0.14
	10	0.64	0.86	0.79	0.57	0.64	0.71	0.57	0.79	0.43	0.50	0.21	0.29	0.14	0.07	0.14
	11	0.79	1.00	0.93	1.00	0.93	1.00	0.93	1.00	0.86	0.79	0.50	0.71	0.71	0.79	1.00
	12	0.79	1.00	0.93	1.00	1.00	0.93	0.93	0.93	0.86	0.71	0.29	0.50	0.57	0.79	0.86
	13	0.71	1.00	1.00	0.93	1.00	1.00	1.00	1.00	0.79	0.86	0.29	0.43	0.50	0.50	0.86
	14	0.86	0.93	1.00	0.93	1.00	0.93	1.00	1.00	0.71	0.93	0.21	0.21	0.50	0.50	0.71
	15	0.71	1.00	1.00	1.00	1.00	0.93	1.00	1.00	0.86	0.86	0.00	0.14	0.14	0.29	0.50

C.3 Texture, Gloss, and Preference Surveys

Observers who participated in the experiments described in Chapter 6, were asked a series of questions about their ratings and choices to help inform the experimenter of their thinkings during each task. The responses are recorded below. Many of the observers echoed the same responses but they are only recorded once.

Tab. C.25. – Summary of responses to Gloss and Texture survey.

Features used to determine gloss level					
contrast	white pixels	increase/decrease in global brightness	brush strokes	highlights	shimmer
Features that made determining gloss level difficult					
shadows, texture, noise, thick brush strokes, high contrast					
Features used to determine texture level					
sharp edges, shadows, brush strokes, highlights, topography, tactile response (how would it feel)					
Features that made determining texture level difficult					
gloss, highlights, noise, high frequency patterns					

Tab. C.26. – Summary of responses to preference versus representative survey.

Features used to determine preference and representative images	
texture, gloss, artist intent, contrast	
Features preferred to representative images	
too much gloss was bothersome, gloss level lower than in real life, good reproduction of texture, minimize gloss and texture, contrast	

Bibliography

[Akca, Remondino, Novák, Hanusch, Schrotter, et al., 2006]

Akca, Devrim, Fabio Remondino, David Novák, Thomas Hanusch, Gerhard Schrotter, et al. (2006). *Recording and modeling of cultural heritage objects with coded structured light projection systems*. Institute of Geodesy and Photogrammetry, ETH Zurich (Cited on page 10).

[Albregtsen et al., 2008]

Albregtsen, Fritz et al. (2008). “Statistical texture measures computed from gray level cooccurrence matrices”. In: *Image processing laboratory, department of informatics, university of oslo* 5 (Cited on page 102).

[Ashbaugh, Berns, Darling, and Taplin, 2009]

Ashbaugh, Justin C, Roy S Berns, Benjamin A Darling, and Lawrence A Taplin (2009). “Artist material BRDF database for computer graphics rendering”. In: *Color and Imaging Conference*. Vol. 2009. 1. Society for Imaging Science and Technology, pp. 62–68 (Cited on pages 20, 21).

[Autodesk, Inc., 2014a]

Autodesk, Inc. (2014a). *Anisotropic*. Online. URL: http://download.autodesk.com/global/docs/maya2014/en_us/index.html?url=files/Shading_Nodes_Anisotropic.htm,topicNumber=d30e635744 (Cited on page 139).

[Autodesk, Inc., 2014b]

Autodesk, Inc. (2014b). *Common surface material attributes*. Online. URL: http://download.autodesk.com/global/docs/maya2014/en_us/index.html?url=files/Shading_Nodes_Blinn.htm,topicNumber=d30e636070 (Cited on page 53).

[Autodesk, Inc., 2014c]

Autodesk, Inc. (2014c). *Common surface material Spectral Shading attributes*. URL: http://download.autodesk.com/global/docs/maya2014/en_us/index.html?url=files/Shading_Nodes_Blinn.htm,topicNumber=d30e636070 (Cited on page 53).

[Autodesk, Inc., 2015a]

Autodesk, Inc. (2015a). *mental ray Help*. Autodesk Inc. (Cited on pages 24, 50).

[Autodesk, Inc., 2015b]

Autodesk, Inc. (2015b). *Overview Comprehensive 3D animation software*. Autodesk Inc. URL: <http://www.autodesk.com/products/maya/overview> (Cited on pages 23, 24, 50).

[Autodesk, Inc., 2015c]

Autodesk, Inc. (2015c). *Shader Types*. Autodesk Inc. URL: http://download.autodesk.com/global/docs/maya2014/en_us/ (Cited on pages 24, 51, 53).

[B & H Foto & Electronics Corp, 2015]

B & H Foto & Electronics Corp (2015). *Broncolor Honeycomb Grid Set of 3, for P70 Reflector*. Online. URL: http://www.bhphotovideo.com/c/product/6899-REG/Broncolor_33_207_00_Honeycomb_Grid_Set_of.html (Cited on page 37).

[Barsky and Petrou, 2003]

Barsky, Svetlana and Maria Petrou (2003). "The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows". In: *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 25.10, pp. 1239–1252 (Cited on page 10).

[Berns, Taplin, Nezamabadi, Mohammadi, and Zhao, 2005]

Berns, Roy, Lawrence Taplin, Mahdi Nezamabadi, Mahnaz Mohammadi, and Yonghui Zhao (2005). "Spectral imaging using a commercial colour-filter array digital camera". In: *The Conservation Committee of the International Council of Museums (ICOM-CC)* (Cited on pages 4–6).

[Berns, 2005]

Berns, Roy S (2005). "Color-accurate image archives using spectral imaging". In: *Proceedings of the National Academy of Science-Scientific Examination of Art: Modern Techniques in Conservation And Analysis*, pp. 105–119 (Cited on page 2).

[Berns, 2016]

Berns, Roy S. (2016). *Theory and Practice of Dual-RGB Imaging*. Tech. rep. Rochester Institute of Technology (Cited on pages 5, 8, 9).

[Berns and Chen, 2012]

Berns, Roy S and Tongbo Chen (2012). *Update:* Practical Total Appearance Imaging of Paintings*. Society for Imaging Science and Technology (Cited on pages 5, 10, 13, 23, 29, 43, 132, 133).

[Berns, Chen, and Chen, 2012]

Berns, Roy S., Tongbo Chen, and Lin Chen (2012). *Evaluating the Quality of Surface Normal Maps of Paintings Using a Scanning Linear Source or Four-Strobe Photometric Stereo*. Tech. rep. Rochester, New York, United States: Rochester Institute of Technology, College of Science, Center for Imaging Science, Munsell Color Science Laboratory. URL: <http://www.cis.rit.edu/DocumentLibrary/admin/uploads/CIS000175.pdf> (Cited on pages 26, 28).

[Besser, 1991]

Besser, Howard (1991). "Imaging: Fine arts". In: *Journal of the American Society for Information Science* 42.8, pp. 589–596. URL: [http://onlinelibrary.wiley.com/doi/10.1002/\(SICI\)1097-4571\(199109\)42:8%3C589::AID-ASI9%3E3.0.CO;2-K/epdf](http://onlinelibrary.wiley.com/doi/10.1002/(SICI)1097-4571(199109)42:8%3C589::AID-ASI9%3E3.0.CO;2-K/epdf) (Cited on page 4).

[Blinn and Newell, 1976]

Blinn, James F. and Martin E. Newell (1976). "Texture and Reflection in Computer Generated Images". In: *Commun. ACM* 19.10, pp. 542–547. URL: <http://doi.acm.org/10.1145/360349.360353> (Cited on page 21).

[Boivin and Gagalowicz, 2002]

Boivin, Samuel and André Gagalowicz (2002). "Inverse rendering from a single image". In: *In Proceedings of IST CGIV* (Cited on page 19).

[Broncolor, Bron Elektronik AG, 2015]

Broncolor, Bron Elektronik AG (2015). *Standard Reflector P70*. Online. URL: <http://www.bron.ch/broncolor/products/light-shapers/showproduct/standard-reflector-p70/?cHash=bb5eb9c79259fd500c9ed1e9f7b7fd70#.V1zHEmSrTow> (Cited on page 36).

[Chen, 2013]

Chen, Lin (2013). "Systems Evaluation for Computer Graphics Rendering of the Total Appearance of Paintings". MA thesis. 54 Lomb Memorial Drive, Rochester, NY 14623: Rochester Institute of Technology (Cited on pages 23, 26–28).

[Chen and Berns, 2012a]

Chen, Lin and Roy S. Berns (2012a). *Visual Validation of Measuring the Total Appearance of Paintings Using Computer Graphics Techniques*. Tech. rep. Rochester, New York, United States: Rochester Institute of Technology, College of Science, Center for Imaging Science, Munsell Color Science Laboratory. URL: <http://www.cis.rit.edu/DocumentLibrary/admin/uploads/CIS000176.pdf> (Cited on pages 15, 26).

[Chen and Berns, 2012b]

Chen, Tongbo and Roy S. Berns (2012b). *Measuring the Total Appearance of Paintings Using a linear Source, Studio Strobes and a Dual-RGB Camera*. Tech. rep. Rochester Institute of Technology (Cited on page 2).

[Chen and Berns, 2012c]

Chen, Tongbo and Roy S. Berns (2012c). *Measuring the Total Appearance of Paintings Using a Linear Source, Studio Strobes, and a Dual-RGB Camera*. Tech. rep. Rochester, New York, United States: Rochester Institute of Technology, College of Science, Center for Imaging Science, Munsell Color Science Laboratory. URL: <http://www.cis.rit.edu/DocumentLibrary/admin/uploads/CIS000174.pdf> (Cited on pages 4, 7, 8, 26).

[Chen, Talpin, and Berns, 2011]

Chen, Tongbo, Lawrence A. Talpin, and Roy S. Berns (2011). *Munsell Color Science Laboratory Artist Material Database BRDF Fitting*. Tech. rep. Rochester Institute of Technology (Cited on page 21).

[Chen, 2008]

Chen, Ying (2008). “Model Evaluation and Measurement Optimization for the Reproduction of Artist Paint Surfaces through Computer Graphics Renderings”. M.S. Thesis. Rochester, New York, United States: Rochester Institute of Technology, , Center for Imaging Science. URL: <http://art-si.org/PDFs/Acquisition/YChenThesis08.pdf> (Cited on pages 21, 24–26).

[Chen, Berns, and Taplin, 2007]

Chen, Ying, Roy S. Berns, and Lawrence A. Taplin (2007). “Model Evaluation for Computer Graphics Renderings of Artist Paint Surfaces”. In: *Proc. of Fifteenth Color Imaging Conference: Color Science and Engineering, Systems, Technologies and Applications, Color Imaging Conference*. IST/SID. Springfield, VA, United States, pp. 54–59. URL: http://art-si.org/PDFs/Metric/CIC2007_Chen.pdf (Cited on pages 19, 21).

[Cox and Berns, 2015]

Cox, Brittany D. and Roy S. Berns (2015). “Imaging artwork in a studio environment for computer graphics rendering”. In: *Conference on Measuring, Modeling, and Reproducing Material Appearance*. Ed. by Maria Ortiz Segovia, Philipp Urban, and Fransico Imai. Vol. 9398. SPIE, pp. 1–7 (Cited on pages 29–32, 35, 44, 132, 133).

[Cultural Heritage Imaging, 2013]

Cultural Heritage Imaging (2013). *Reflectance Transformation Imaging: Guide to RTIViewer*. Cultural Heritage Imaging. URL: http://culturalheritageimaging.org/What_We_Offer/Downloads/rtiviewer/RTIViewer_Guide_v1_1.pdf (Cited on pages 10, 23, 44, 55).

[Cultural Heritage Imaging, 2014]

Cultural Heritage Imaging (2014). *Cultural Heritage Imaging: About Us*. Online. Cultural Heritage Imaging. URL: http://culturalheritageimaging.org/About_Us/ (Cited on page 16).

[Day, Taplin, and Berns, 2004]

Day, Ellen A, Lawrence Taplin, and Roy S Berns (2004). “Colorimetric characterization of a computer-controlled liquid crystal display”. In: *Color Research & Application* 29.5, pp. 365–373 (Cited on page 65).

[Debevec, Taylor, and Malik, 1996]

Debevec, Paul E., Camillo J. Taylor, and Jitendra Malik (1996). "Modeling and Rendering Architecture from Photographs: A Hybrid Geometry- and Image-based Approach". In: *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '96. New York, NY, USA: ACM, pp. 11–20. URL: <http://doi.acm.org/10.1145/237170.237191> (Cited on page 4).

[Duffy, 2013]

Duffy, S Bryan (2013). *Multi-Light Imaging Techniques for Heritage Application*. English Heritage (Cited on pages 15, 16).

[Engeldrum, 2000]

Engeldrum, Peter G (2000). *Psychometric scaling: a toolkit for imaging systems development*. Imcotek press (Cited on pages 143, 145, 147, 148).

[Fairchild and Wyble, 1998]

Fairchild, Mark and David Wyble (1998). *Colorimetric characterization of the apple studio display (flat panel LCD)*. Tech. rep. Rochester Institute of Technology (Cited on page 65).

[Fairchild, Wyble, and Johnson, 2008]

Fairchild, Mark D, David R Wyble, and Garrett M Johnson (2008). "Matching image color from different cameras". In: *Electronic Imaging 2008*. International Society for Optics and Photonics, 68080E–68080E (Cited on page 6).

[Fischer and Kakoulli, 2006]

Fischer, Christian and Ioanna Kakoulli (2006). "Multispectral and hyperspectral imaging technologies in conservation: current research and potential applications". In: *Studies in Conservation* 51.Supplement-1, pp. 3–16 (Cited on pages 4, 5).

[Fores, Ferwerda, Tastl, and Recker, 2013]

Fores, Adria, James Ferwerda, Ingeborg Tastl, and John Recker (2013). "Perceiving Gloss in Surfaces and Images". In: *Color and Imaging Conference*. Vol. 2013. 1. Society for Imaging Science and Technology, pp. 44–51 (Cited on page 17).

[Gardner, Tchou, Hawkins, and Debevec, 2003]

Gardner, Andrew, Chris Tchou, Tim Hawkins, and Paul Debevec (2003). "Linear light source reflectometry". In: *ACM Transactions on Graphics (TOG)*. Vol. 22. 3. ACM, pp. 749–758 (Cited on pages 10, 18, 19, 21).

[Haralick, Shanmugam, et al., 1973]

Haralick, Robert M, Karthikeyan Shanmugam, et al. (1973). "Textural features for image classification". In: *IEEE Transactions on systems, man, and cybernetics* 3.6, pp. 610–621 (Cited on page 101).

[Hardeberg, Schmitt, and Brettel, 2002]

Hardeberg, Jon Y, Francis Schmitt, and Hans Brettel (2002). "Multispectral color image capture using a liquid crystal tunable filter". In: *Optical engineering* 41.10, pp. 2532–2548 (Cited on pages 4, 5).

[Horn, 1970]

Horn, Berthold KP (1970). *Shape from shading: A method for obtaining the shape of a smooth opaque object from one view*. Tech. rep. Massachusetts Institute of Technology (Cited on page 10).

[Hunter and Harold, 1987]

Hunter, Richard Sewall and Richard W Harold (1987). *The measurement of appearance*. 2nd. John Wiley & Sons (Cited on page 17).

[Levoy, Pulli, Curless, Rusinkiewicz, Koller, et al., 2000]

Levoy, Marc, Kari Pulli, Brian Curless, Szymon Rusinkiewicz, David Koller, et al. (2000). "The digital Michelangelo project: 3D scanning of large statues". In: *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., pp. 131–144 (Cited on pages 4, 10).

[Löw, Kronander, Ynnerman, and Unger, 2012]

Löw, Joakim, Joel Kronander, Anders Ynnerman, and Jonas Unger (2012). "BRDF models for accurate and efficient rendering of glossy surfaces". In: *ACM Transactions on Graphics (TOG)* 31.1, p. 9 (Cited on page 20).

[Malzbender, Gelb, and Wolters, 2001]

Malzbender, Tom, Dan Gelb, and Hans Wolters (2001). "Polynomial texture maps". In: *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, pp. 519–528 (Cited on pages 4, 10, 15, 16).

[Marçelja, 1980]

Marçelja, S (1980). "Mathematical description of the responses of simple cortical cells". In: *JOSA* 70.11, pp. 1297–1300 (Cited on page 103).

[Martinez, Cupitt, Saunders, and Pillay, 2002]

Martinez, K., J. Cupitt, D. Saunders, and R. Pillay (2002). "Ten years of art imaging research". In: *Proceedings of the IEEE* 90.1, pp. 28–41 (Cited on page 2).

[MathWorks, Inc., 2017]

MathWorks, Inc. (2017). *kurtosis*. <https://www.mathworks.com/help/stats/kurtosis.html>: MathWorks, Inc. URL: <https://www.mathworks.com/help/stats/kurtosis.html> (Cited on page 100).

[Montag, 2006]

Montag, Ethan D (2006). "Empirical formula for creating error bars for the method of paired comparison". In: *Journal of Electronic Imaging* 15.1, pp. 010502–010502 (Cited on page 147).

[Montes and Ureña, 2012]

Montes, Rosana and Carlos Ureña (2012). *An overview of BRDF models*. Tech. rep. Granada, Spain: University of Granada. URL: http://digibug.ugr.es/bitstream/10481/19751/1/rmontes_LSI-2012-001TR.pdf (Cited on page 19).

[Motoyoshi, Nishida, Sharan, and Adelson, 2007]

Motoyoshi, Isamu, Shin'ya Nishida, Lavanya Sharan, and Edward H Adelson (2007). "Image statistics and the perception of surface qualities". In: *Nature* 447.7141, pp. 206–209 (Cited on pages 100, 101).

[Mudge, Voutaz, Schroer, and Lum, 2005]

Mudge, Mark, Jean-Pierre Voutaz, Carla Schroer, and Marlin Lum (2005). "Reflection Transformation Imaging and Virtual Representations of Coins from the Hospice of the Grand St. Bernard". In: *Proceedings of the 6th International Conference on Virtual Reality, Archaeology and Intelligent Cultural Heritage*. VAST'05. Pisa, Italy: Eurographics Association, pp. 29–39. URL: <http://dx.doi.org/10.2312/VAST/VAST05/029-039> (Cited on pages 15, 21).

[National Institute of Standards and Technology (NIST), 2013]

National Institute of Standards and Technology (NIST) (2013). *Engineering Statistics Handbook*. NIST. URL: <http://www.itl.nist.gov/div898/handbook/eda/section3/eda35b.htm> (Cited on page 100).

[Nicodemus, Richmond, Hsia, Ginsberg, and Limperis, 1977]

Nicodemus, Fred E, Joseph C Richmond, Jack J Hsia, Irving W Ginsberg, and Thomas Limperis (1977). *Geometrical considerations and nomenclature for reflectance*. Vol. 160. US Department of Commerce, National Bureau of Standards Washington, DC, USA (Cited on pages 4, 16).

[Padfield, Saunders, and Malzbender, 2005]

Padfield, Joseph, David Saunders, and Tom Malzbender (2005). "Polynomial texture mapping: a new tool for examining the surface of paintings". In: *ICOM Committee for Conservation* 1, pp. 504–510 (Cited on page 10).

[Rushmeier, Taubin, and Guéziec, 1997]

Rushmeier, Holly, Gabriel Taubin, and André Guéziec (1997). "Applying shape from lighting variation to bump map capture". In: *Rendering Techniques' 97*. Springer, pp. 35–44 (Cited on pages 10–12, 29).

[Siew, Hodgson, and Wood, 1988]

Siew, Lee Hok, Robert M. Hodgson, and Errol J. Wood (1988). "Texture measures for carpet wear assessment". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10.1, pp. 92–105 (Cited on page 101).

[Soh and Tsatsoulis, 1999]

Soh, L-K and Costas Tsatsoulis (1999). "Texture analysis of SAR sea ice imagery using gray level co-occurrence matrices". In: *IEEE Transactions on geoscience and remote sensing* 37.2, pp. 780–795 (Cited on pages 101, 102).

[Sun, Smith, Smith, Midha, and Bamber, 2007]

Sun, Jiuai, Melvyn Smith, Lyndon Smith, Sagar Midha, and Jeff Bamber (2007). "Object surface recovery using a multi-light photometric stereo technique for non-Lambertian surfaces subject to shadows and specularities". In: *Image and Vision Computing* 25.7, pp. 1050–1057 (Cited on pages 10, 12–14, 29).

[The J. Paul Getty Museum]

, The J. Paul Getty Museum. *Commode*. URL: <http://www.getty.edu/art/collection/objects/5359/attributed-to-joseph-baumhauer-commode-french-about-1750/> (Cited on page 57).

[Tisato and Parraman, 2014]

Tisato, Flavia and Carinna Parraman (2014). "An investigation into the micro surface of artworks using alternative lighting techniques". In: *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, pp. 901808–901808 (Cited on page 10).

[Tominaga and Tanaka, 2002]

Tominaga, Shoji and Norihiro Tanaka (2002). "Measuring and rendering art paintings using an RGB camera". In: *Proc. of EUROGRAPHICS*, pp. 299–306 (Cited on pages 2, 10).

[Waldhäusl, 1992]

Waldhäusl, P (1992). "Defining the future of architectural photogrammetry". In: *International Archives of Photogrammetry and Remote Sensing* 29.B5, pp. 767–770 (Cited on page 10).

[Ward, 1992]

Ward, Gregory J (1992). "Measuring and modeling anisotropic reflection". In: *ACM SIGGRAPH Computer Graphics* 26.2, pp. 265–272 (Cited on pages 18, 20).

[Webber, 2008]

Webber, S (2008). "Technical imaging of paintings". In: *Williamstown Art Conservation Center Technical Bulletin* (Cited on pages 33, 34).

[Witwer and Berns, 2015]

Witwer, Joel and Roy S Berns (2015). "Increasing the Versatility of Digitizations through Post-Camera Flat-Fielding". In: *Archiving Conference*. Vol. 2015. 1. Society for Imaging Science and Technology, pp. 110–113 (Cited on page 5).

[Woodham, 1980]

Woodham, Robert J (1980). "Photometric method for determining surface orientation from multiple images". In: *Optical engineering* 19.1, pp. 191139–191139 (Cited on pages 4, 10, 14, 133).

[Zhao, Taplin, and Nezamabadi, 2005]

Zhao, Yonghui, Lawrence Taplin, and Mahdi Nezamabadi (2005). "Using the matrix R method for spectral image archives". In: *Association Internationale de la Couleur (AIC)-International Color Association* (Cited on pages 4, 5).

