Exploring the Cognitive Correlates of Artistic Practice using a
Parameterized Non-Photorealistic Toolkit

by

Stephen DiPaola

B.S., State University of New York, Stony Brook, 1981

M.A., New York Institute of Technology, 1991

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE STUDIES

(Interdisciplinary Studies)

[Computer Science / Psychology / Philosophy]

The University Of British Columbia
(Vancouver)

July 2013

© Stephen DiPaola, 2013

Abstract

Artists and scientists have different approaches to knowledge acquisition, usage and dissemination. This research work is one attempt to bridge these different fields, through a multi-stage process involving the creation of a software toolkit for non-photorealistic rendering (NPR). Our domain of inquiry is the creation and viewing of fine art painting -- we are interested in elucidating cognitive and perceptual mechanisms or ‘cognitive correlates’ which correspond and relate to artists’ techniques and conceptions regarding fine art painting in general and portraiture in particular.

As our starting point, we analyse an extensive corpus of art-theory literature to identify broadly accepted understandings and techniques, which might be relevant to human perception and cognition. We further condense this artistic knowledge into a concise set of heuristics, which are suitable for parameterization and algorithmic implementation, and examine findings from psychology and neuroscience, which correlate to each heuristic. We present our system design for a painterly NPR toolkit, which is informed by these heuristics within a suitable object-oriented, cognitively inspired architecture.

By interpreting artistic and cognitive science knowledge into a well-defined computational framework, we gain opportunities to formalize and test new hypotheses. We demonstrate the productive power of such an approach by examining in depth two particular techniques (lost-and-found edges and varying fine detail level) used by a particular artist (Rembrandt). We formulate four experiments based on eye tracking of human viewers, using our painterly NPR toolkit to generate example artworks with manipulated generation parameters. We obtain significant findings suggesting that artists such as Rembrandt use techniques, which leverage perceptual and cognitive function to exert control over viewer’s gaze patterns, which in turn influences the experienced artistic merit of a painting.

Preface

Papers from this thesis work include:

- 1. DiPaola S, (2007). Painterly Rendered Portraits from Photographs using a Knowledge-Based Approach, In Proc: SPIE Human Vision and Imaging, Int. Society for Optical Engineering, Keynote paper. pp 33-43.**
- 2. DiPaola S, (2008). The Trace and the Gaze: Textural Agency in Rembrandt's Late Portraiture from a Vision Science Perspective, Proceedings of Electronic Visualisation and the Arts, 8 pages, London.**
- 3. DiPaola S, (2009). Exploring a Parameterized Portrait Painting Space, International Journal of Art and Technology, Vol 2, No 1-2, pp 82-93.**
- 4. DiPaola S, Riebe C, Enns J, (2010). Rembrandt's Textural Agency: A Shared Perspective in Visual Art and Science, Leonardo, Vol 43, No 3, pp 145-151.**
- 5. DiPaola S, (2010). Face, Portrait, Mask – Using a Parameterized System to Explore Synthetic Face Space, Proceedings of Electronic Visualisation and the Arts, pp. 296-302. British Computer Society, London.**
- 6. DiPaola S, Riebe C, Enns J, (2013). Portrait Viewing and Appreciation Is Guided by Selective Detail, Perception, in press.**

The Rembrandt Textural Agency theory (Paper 2) was my original idea and research; I took this hypothesis, with questions on how to test it empirically, to my committee member James Enns and on his advice, we began working together to design and run the four studies done in his UBC Vision Lab. I was a member of his lab. Much of the study design expertise, study techniques and results discussion in our papers on these studies (Paper 4 and 6) came from his significant experience. These four studies with their results and discussion are combined and the major part of the Cognitive Science Studies chapter (Chapter 7). James Enns and I also supervised a UBC undergraduate student Caitlin Riebe (who won a Department of Psychology Quinn Summer research award to work on the project). She was subsequently included as a co-author on Papers 4 and 6 to acknowledge her contribution in data collection and discussion of the ideas. For these four studies, the creation of the source

Rembrandt sitter photos, and the many final manipulated Rembrandt facsimile output from them using my NPR toolkit was solely my responsibility based on advice from Enns.

Approval for the experiments described in this dissertation (Paper 4 and 6) was obtained from the UBC Office of Research Services under certificate numbers B98-0398, H10-01707 and H11-02768. (UBC Vision Lab – James Enns)

The painterly NPR system detailed in this thesis (Paper 1 and 3) was originally designed and coded solely by me. Over the years, I have paid a part time undergraduate Java software programmer to help implement parts of the system based on my designs and management. The undergraduate student was first Thomas Johnson, then Dustin Dunsmuir and then Brendan Vance, all from Simon Fraser University's School of Interactive Arts and Technology (SFU SIAT). The latter, Brendan Vance, was SFU SIAT's valedictorian, a very bright individual who made strong suggestions about implementation details based on my designs and management.

Table of Contents

Abstract	ii
Preface	iii
Table of Contents.....	v
List of Figures	viii
Acknowledgements	xii
1: Introduction.....	1
1.1 Motivations and Approach	1
1.2 Background: Fine Art Practice, Cognitive Artist	3
1.2.1 Source Photograph – Final Painting Pairing As Domain of Inquiry	5
1.2.2 Salience of What? Artists Directed Salience: Portrait Knowledge	6
1.3 NPR Painterly Rendering	7
1.4 Goals and Contribution of Thesis.....	9
1.4.1 Goals	9
1.4.2 Contributions	10
1.5 Overview of Thesis	12
2: Related Work.....	14
2.1 Non Photorealistic Rendering	14
2.1.1 NPR Introduction	14
2.1.2 NPR Related Work.....	14
2.2 Perceptual and Cognitive Sciences: Visual Art.....	19
3: Identifying Cognitively-Based Techniques within Artistic Painting Practice	22
3.1 Introduction	22
3.1.1 Motivation and Goals	24
3.1.2 Methods and Scope of Artist Process Categorization	24
3.1.3 Poetics Methodology	25
3.2 Soft Rules of Fine Art Painting	26
3.2.1 Tonal	26
3.2.2 Colour.....	29
3.2.3 Textural Highlights	32
3.2.4 Shadows and Halftone ‘Core’	33
3.2.5 Shapes, Edges and Centre of Interest	34
3.2.6 Brush Strokes	37
3.2.7 Backgrounds.....	38
3.2.8 Working from Photograph as Source Image	39

3.2.9	Process Plan	39
3.2.10	Summary	41
3.3	Cognitive Correlates of Artistic Painter Practice	42
3.3.1	Introduction and Soft Rules.....	42
3.3.2	Cognitive Correlates Discussion	44
3.3.3	Natural Scenes and Visual Cognition.....	45
3.3.4	Foveal Vision and Authored Pattern of Eye Gaze	48
3.3.5	Tonal Value and Colour	49
3.3.6	Global and Relative Analysis: Tonal Masses, Volume and Content.....	53
3.3.7	Summary	55
3.4	NPR Toolkit Parameters.....	56
3.4.1	Introduction	56
3.4.2	NPR Toolkit Summary.....	62
3.5	Conclusion.....	63
4:	Rembrandt and Textural Agency – Art Critical Theory.....	64
4.1	Introduction	64
4.2	Rembrandt and Textural Agency	65
4.2.1	Early Modern Painting: Monolithic Mechanistic View	65
4.2.2	Systems of Early Modern Painting - Painting Modes	66
4.3	The Science of Human Vision and Eye Gaze	71
4.3.1	Dynamic Vision.....	72
4.3.2	Eye Movement Research in Art	73
4.4	Rembrandt, Eye Movement and Intent.....	76
4.5	Conclusion.....	82
4.5.1	Future Evidence Gathering.....	82
5:	Painterly System – Parametric Approach	83
5.1	Overview	83
5.2	Initial Phase One NPR Toolkit.....	85
5.3	ThinkerPainter (Phase Two) Framework	91
5.3.1	Thinker Painter Concepts: BlobThinker, ActionPainter, RelativeJChPalette	94
5.3.2	Blobs, Stroke and Concerns	96
5.3.3	Limitations	98
5.4	Conclusion.....	100
6:	Painterly System – Scripting, Output and Implementation.....	101
6.1	Scripting, Concerns and Examples.....	101
6.1.1	Use in Industry Research: 3D Pipeline Automatic Control.....	106
6.1.2	Component Details and Implementation.....	108
6.1.3	The Top-Level Control Structure	109
6.1.4	BlobThinker, ActionPainter, RelativeJChPalette in Practice	110
6.1.5	Algorithm – The Thinker/Painter Loop in Detail.....	114
7:	Cognitive Science Studies – Rembrandt’s Textural Agency and Eye Fixations.....	118
7.1	Introduction	118
7.2	Background	118
7.3	Testing the Detail-Gaze Hypothesis.....	122

7.4	Experiment 1	124
7.4.1	Method (Experiment 1)	127
7.4.2	Results (Experiment 1)	128
7.5	Experiment 2	133
7.5.1	Method (Experiment 2)	134
7.5.2	Results (Experiment 2)	135
7.6	Experiment 3	139
7.6.1	Method (Experiment 3)	141
7.6.2	Results (Experiment 3)	142
7.7	Experiment 4	144
7.7.1	Method (Experiment 4)	144
7.7.2	Results (Experiment 4)	144
7.8	General Discussion of Experiments	145
7.8.1	Summary of Our Findings	145
7.9	Implications of These Findings	147
7.9.1	From a Vision Science Perspective	147
7.9.2	From a Relationship for Art, Art History and Vision Science	147
8:	Conclusion.....	149
8.1	Discussion of Studies	149
8.1.1	Implications for Vision Science	154
8.2	A Toolkit and Process for Visual Art and Science.....	155
8.3	Thesis Summary and Future Work.....	158
8.3.1	Future Work	159
	References	161
	Appendix A: Painterly System Documentation and XML language	188
	Appendix B: Table of Artistic Practice Source Material	196

List of Figures

Figure 1	Main Process Chart: From input qualitative data through art and cognitive filtering to NPR toolkit and its output results and benefits.	3
Figure 2	Using a source sitter photograph (lower right), our NPR system uses learned semantic and artistic constructs to output a Rembrandt facsimile (left) with the ability for repeatable/customizable manipulation (here to test effect of textural highlights all pointing to a dominant eye) for both science and art studies.	9
Figure 3	(Zhao & Zhu, 2011, 2013) work with brush templates for portraits.	17
Figure 4	DeCarlo and Santella [2002] results using their abstraction using eye-tracking data (left) and later Bottom Row: Our automatic abstraction, and (Winnemöller et al., 2006) version also using eye tracking.	18
Figure 5	Source photograph and final painting by fine artist James Horton from (Horton, 1995). How does a human artist ‘remap’ the source photo (left) to a finished painting and what cognitive processes does the remapping exploit?	22
Figure 6	‘Tone first’, then map using some colour method – lie warm/cool colours. From (Sanden & Sanden, 1999; Saper, 2001)	27
Figure 7	Unequal Dominate Value: Left: Strong: big dominate value and unequal sub values. Right: Weak: little dominate and in addition, equal balance. From (Brown & Lehrman, 2004).	29
Figure 8	Tone (Value, Luminance) First (A), then balancing (separating) warm and cool colour temperature across colour wheel (B). Images: Bruce MacEvoy.	30
Figure 9	Unequal Colour (Brown & Lehrman, 2004)– Warm/Cool colour wheel (bottom left) shows dominant colour (black dot), with its compliment directly across the wheel, then the 2 discords equidistant offshoots and how they are all unequally weighted in the two upper abstract images (green dominate and blue dominate respectively). ...	32
Figure 10	Exaggerating the core between light and dark areas. (Saper, 2001)	34
Figure 11	Centre of Interest, Lost and Found Edges. From (Sanden, 2004) and (Brown & Lehrman, 2004).	35
Figure 12	An example of both 1) edges, then lost for gaze pathing toward a 2) centre of focus detailed eye, which came from our later research studies in Chapter 7 with Rembrandt’s <i>Beret</i> and other work.	36
Figure 13	Photo versus painting – Value Clumping in light and shadow of face. (Saper, 2001)	37
Figure 14	Brush stroke direction techniques (Stewart, 1997).	38
Figure 15	From a black and white source (inset), our NPR toolkit created this painterly output using a hierarchical blob tree structure and semantic regions. The colour	

palette was remapped from a fine art master's portrait where tonal values map into the semantic regions of the source colour painting.	56
Figure 16 Left: Blob Tree with false colour semantic top regions (optional) and next level luminance regions; Middle: the pre-specified semantic regions and labels; Right: the third level of tree luminance regions. All exist in a tree structure.....	59
Figure 17 J.H. Sanden, focus and eye fixation paths to sharp eye. (Sanden, 2004)	67
Figure 18 H. Brown, 'Lost & Found Edges' – contrast based edges intact, then lost at elbows to drive you back to face. (Brown & Lehrman, 2004).	69
Figure 19 (a) The fovea's size and placement in the eye. (b) The middle spike and fast fall off show the fovea's highly packed density of cones compared to the rest of the eye. From (Lindsay & Norman, 1977).	72
Figure 20 Yarbis (1967) Eye movement: Images of a: face of young girl and b: bust of Queen Nefertiti along with the eye scans, both movements (lines) and fixations (points) (Noton & Stark, 1971).	74
Figure 21 a) Yarbis (1967), Repin study, 7 targeted scans. b) Molnar (1997) Rembrandt, 2 scans.	75
Figure 22 Extending on Durand's bottom up cognitive flow, low level constructs (on the textural plane) can affect the viewer's reception of high level semantics in a painting.....	77
Figure 23 Early self-portrait, 1629 versus a late self-portrait (right), 1661 (close-up). Note increased textural level details (e.g. graduation, edges, and focus differences).....	78
Figure 24 <i>Self Portrait with Beret</i> , 1659. Focused left eye, lost left lower cheek edge & dark background send viewers eye toward cent of interest.	80
Figure 25 Top) Portrait of an Old Lady, 1662. Focused left eye, lines, and contours wrapping the face. Bottom) <i>Elderly Man</i> , 1667. Disjointed highlight lines pull your gaze in all directions, emphasizing a confused almost drunk depiction.	81
Figure 26 Main Process Chart: From input qualitative data through art and cognitive filtering to NPR toolkit and its output results and benefits.	83
Figure 27 Parameterization Process: From soft knowledge to a space of millions of individual paintings, a parameter space (middle) defines a set of factors (params) whose values determine the characteristics of a system (dimensions of the space). With one individual (a painting) being a point in that space.....	85
Figure 28 For the output image created (right) of the Jane Series, our smartpalette system used content maps (left) to better inform the colour choices as well as started our cognitive based hierarchical blob tree as level 0, with luminance calculated blobs at lower and lower levels of the tree (levels 1,2,3 from right to left – top) that allow local and global relative communication within the toolkit.	88
Figure 29 Using semantic regions, FullRangeValuePalette(), the sitter images of PhD students (lower right), were transformed in 1 of 4 customizable Rembrandts for our eye-tracking studies.	89
Figure 30 Using semantic regions, FullRangeValuePalette(), the sitter images of PhD students (lower right), were transformed in 2 of 4 customizable Rembrandts for our eye-tracking studies.....	90

Figure 31	Controlling the parameters, allows the toolkit to generate several or hundreds of different painterly portrait scenarios in the specified multi-dimensions as shown by these four images.....	90
Figure 32	The example set taken from a run of hundreds, shows moving through a one parameter (above - brush size), then other examples (cropped) of many paths through the portrait space from the same input. Note these early results were prior to our implementation of the JCh CIELAB colour model.	91
Figure 33	Process flow and main modules of ThinkerPainter (phase 2) NPR system.	93
Figure 34	A ‘Concern’ to use either a 3D depth buffer (left) or regular luminance (right) or some intelligent mix for stroke direction decisions.....	98
Figure 35	Test Jane Series Image which showing the Maps (inset) that affect its brush parameters across the canvas (see script) using a SmartPalette from a historically painting by Hopper for its tone mapped regional colour choices.....	101
Figure 36	Black and white source image for the Jane Series. Since no colour exists in the source, all colour is generated through using tone to remap into a rescaled and regional art historical colour palette.....	103
Figure 37	Jane Series One: with details eye and Concern script to value-blend smart pallets....	104
Figure 38	Jane Series Two	105
Figure 39	Jane Series Three	105
Figure 40	Jane Series Four	106
Figure 41	Images stills from industry / research partnership (NSERC) to control our NPR toolkit automatically from a 3D animation pipeline.	107
Figure 42	Animation stills from industry / research partnership (NSERC) to control our NPR toolkit automatically from a 3D animation pipeline.	108
Figure 43	Thinker painter main data flow.	109
Figure 44	Left, the reference image; right, a visual representation of hierarchical leaf blobs in a structure derived from this reference.....	111
Figure 45	Left, a gradient map of the reference image; right, an error map derived from the reference.....	112
Figure 46	A colour map generated by PSDPortraitPalette and RelativeJChPalette using the ‘Hopperself’ palette.....	113
Figure 47	ThinkerPainter Flow in detail.	115
Figure 48	One pass, rendered to the canvas. This pass uses the same blob, gradient and error maps pictured above.....	117
Figure 49	J.H. Sanden’s <i>Reverend Cole close-up</i> . From (Sanden & Sanden, 1999).....	120
Figure 50	Source image (lower right), original Rembrandt <i>Hendrickje</i> close up (upper right) and NPR system output using style and palette from <i>Hendrickje</i> , in this case with manipulated ‘textural highlights’ supporting main eye for Exp 3.....	122
Figure 51	Four original Rembrandts and our photographs of the human model analogues used in the eye tracking study.	124

Figure 52	Top: Close-ups of 2 of 4 images for Beret (+Mirrors) where we vary the right or left detail of the eye (centre of focus) as well as the right and left lost and found ‘sharp’ edge. For all 4 Rembrandts. The original Rembrandts and the original sitter photographers were used for our studies. Bottom: <i>Hendrickje</i>	126
Figure 53	All fixations made for one of the Rembrandt-renderings during the 5 sec viewing period for all 32 participants in Experiment 1. Each participant’s fixations are coded in a different colour.	129
Figure 54	Mean time to first fixation (in milliseconds) in one of the two eye regions for photos and Rembrandt-renderings in Experiment 1. Error bars are +/- one standard error.	130
Figure 55	Mean number of fixations made to high and low detailed eye regions, separately as a function of whether the more detailed collar is on the same or opposite side of the more detailed eye region. Error bars are +/- one standard error.	131
Figure 56	Mean artistic ratings for the 40 portraits in Experiment 1, ordered from highest to lowest. Selected sample portraits are shown below in miniature form.	132
Figure 57	Scatterplot of the correlation between the difference in fixation frequency for the two eye regions and the mean artistic rating in Exp. 1. Each data point represents the mean rating of 8 participants on 1 of the 40 photos and critical portrait.	133
Figure 58	Mean number of fixations made during the viewing of each type of portrait in Experiment 2. Error bars are +/- one standard error.	136
Figure 59	Mean time (in milliseconds) taken to make a first fixation to an eye region in Experiment 2. Error bars are +/- one standard error.	137
Figure 60	Mean artistic ratings for the 10 filler portraits and 4 Rembrandt original portraits in Experiment 1, rank from highest to lowest. Selected sample portraits are shown below in thumbnail form.	138
Figure 61	Image showing <i>Beret</i> close ups of two different highlight variant, towards bad eye (left) and towards good eye (right), also no highlights and scattered highlights were used.	140
Figure 62	Image showing <i>Hendrickje</i> close ups of two different highlight variant, towards bad eye (left) and towards good eye (right), also no highlights and scattered highlights were used.	141
Figure 63	Mean number of fixations made to each eye region in Experiment 3, plotted as a function of five different portrait types. Error bars are +/- one standard error.	142
Figure 64	Scatterplot showing each critical portrait’s mean artistic rating as a function of the difference in fixation frequency to the sharp versus coarse eye. Separate regression lines show baseline difference in artistic ratings for two portraits.	143
Figure 65	Scatterplot showing each filler portrait’s mean artistic rating as a function of the difference in fixation frequency to the sharp versus coarse eye. Several examples of the portraits are shown below the graph.	145
Figure 66	Close up of original, showing significant form of textural highlights on <i>Hendrickje</i> ’s cheek from the Metropolitan Museum of Art, NYC.	151
Figure 67	Rembrandt’s <i>Elderly Man</i> with untypical scattered textural highlights.	152

Acknowledgements

I would thank my wife Ligia and son Forrest for their help, encouragement and patience through this process. I would also like to thank my committee: Wolfgang Heidrich, Dominic Lopes, Rhodri Windsor-Liscombe and especially James Enns for their advice, stewardship and encouragement.

1: Introduction

This thesis offers a new approach to understanding painterly artistic practice and perception. Our work transfers and synthesizes knowledge among several fields including cognitive and perception sciences, artistic practice and computer science; here we clarify the approach and connections and give background to the key topics of painterly practice, perception science and non-photorealistic rendering (NPR) and vision science experiments. We explain our choice and scope of specific problem within this domain - namely, fine art painting in general and specifically portrait painting based on a source photograph.

1.1 Motivations and Approach

Cognitive science seeks to understand how human experience is linked to behaviour, whether that behaviour is from motor actions or from neurons. Those links are what we define in this thesis as ‘cognitive correlates’. The correlates point to properties of brain, mind and cognition.

Here we are interested in the process of fine art painting, how it is created and perceived as well as the cognitive phenomena, which underlay it. Therefore the process of fine art painting is the behaviour we attempt parse from a ‘cognitive correlate’ perspective. We are motivated in part by the growing recognition that ‘Artists are Neuroscientists’: i.e. they have discovered valuable ways of understanding and working with human cognitive and perceptual mechanisms to achieve authorship techniques to convey desired experiences and narrative in their created art works (Cavanagh, 2005; Zeki, 2001). Whether artists are or act like a type of neuroscientist as many cognitive scientists like to state or more simply that artists have a passed down methodology and talent space where they use their eyes, perception and mind to produce a painting in a way that when analysed through a certain cognitive lens can benefit both the arts and cognitive sciences, is less the point. Hence, we explore art topics through a cognitive science perspective with the aim of enriching our understanding of both art practice (i.e. the act of fine art painting) and the underlying

perceptual mechanisms. Cognitive science analyses mind on several levels: conceptual, algorithmic, and neurological. Here we construct an overarching methodology, which extracts artistic knowledge and reifies and tests hypotheses about it through computational modelling. Part of this methodology is the thoughtfully considered design of a painterly NPR computational toolkit.

As a key contribution of this work, we identify artistic cognitive based techniques, in general and some more deeply ('lost and found edges' and 'centre of focus' detailing to guide the viewer's pattern of gaze in a painting) and analyse it in art literature, then from a historical, cognitive and psychological perspective. We use our computational toolkit to implement this technique algorithmically. We conduct user studies to verify and quantify the impact of this technique.

We use Non Photorealistic Rendering (NPR) as a computational modelling technique for our work. NPR has several important uses and is an area of active research. Current research involves computer graphics and imaging techniques with some researchers looking at limited aspects of human cognition to inform their systems. While NPR systems are based on the human artistic process, researchers have typically not analysed or computer modelled, in a deep way, the passed down artistic methodology artists use. Moreover, many cognitive scientists have begun to look into how artists work from a cognitive standpoint. The work we present here attempts to use strong analysis of the artistic painterly process through a lens of new scientific understanding of human cognition to create a cognitive knowledge based painterly NPR software toolkit that can have both wider range and improved results compared to many current NPR techniques. By limiting the investigation to fine art portraits (as opposed to all painterly art forms) and how they vary from their photography analogues, the system can use strong knowledge (e.g. salience) of portraits and faces to inform the system. Therefore, the other main benefit of the system is as a scriptable NPR toolkit that can be useful to art history, design and scientific scholars to test hypotheses from both artistic and scientific fronts. Finally, we use our NPR system to conduct four eye-tracking, parameterized modelling experiments both to give strong evidence to an original hypothesis about Rembrandt's intuitive understanding of aspects of vision science as well as generate strong empirical results that show how vision science understanding is used to create fine art. These

studies are conducted to both further the field of cognition and the arts as well as to validate our painterly NPR toolkit as a research tool. The process is outlined in Figure 1.

Figure 1 Main Process Chart: From input qualitative data through art and cognitive filtering to NPR toolkit and its output results and benefits.

1.2 Background: Fine Art Practice, Cognitive Artist

The process of representing a scene, whether from life or a photograph as the source, must by necessity involve simplification, a loss of irrelevant detail in order to better highlight what's important to the artist (DiPaola, 2007). For instance, when using a photograph as source, artists must select the focus that will be rendered in the painting (e.g., will everything be in focus when directly fixated or only selected regions?), the range of value clumping (e.g., will the shadows be darker and the brighter areas lighter, or vice versa?), the range of colour values, and the degree of perspective distortion (e.g., will the depth in the painting be reduced or exaggerated relative to the photo). Artists who have written about this selection process claim to rely on heuristic techniques such as the following. For instance, for colour selection they may rely on source tone over wavelength to map into an artistic colour temperature model. For value clumping, they may use brush and tonal 'sharpness' variability

to create a centre of interest. To influence viewer gaze they may use lost and found edges (sharp edges that at some point fade away) to direct the viewer's gaze towards the centre of interest in the image as well as other techniques to filter and emphasize (DiPaola, 2009). Our collected, systematic review, analysis and alignment of these and other passed down artistic techniques are presented in Chapter 3.

Cognitive scientist Collomosse (2004), paraphrasing art historian E.H. Gombrich (2000), has noted, artists “commonly paint to capture the structure and elements of the scene which they consider important; the remaining detail is abstracted away in some differential style ...”. Cavanagh in his ground breaking article in the journal *Nature* entitled “The Artist as Neuroscientist” (Cavanagh, 2005) wrote that artists in discovering these semantic shortcuts, act as research neuroscientists, and that “a great deal can be learned from tracking down their discoveries”. Cavanagh showed how the difference between the real world and the world artists create could reveal “as much about the brain within us as the artist reveals about the world around us”.

Neuroscientist Stephen Grossberg, well known for establishing a number of cognitive science and neural network theories, has documented how a variety of artists, including Renaissance painters, Fauvists and Impressionists “have developed an artistic understanding of rather deep properties of how the brain sees. These include what we now know about how the brain perceives properties of 2D and 3D perceptual boundaries and surfaces, and organization of opponent colors (Grossberg, 2008).

Zeki (2000) who originally coined the term ‘neuro-aesthetics’, believed that artists study the brain like neurologists albeit with more qualitative techniques that are unique to them. Artists therefore, over 40,000 years, have created a record of experiments in visual science, driven to be “both impactful but economical”. Therefore, many cognitive scientists have noted in essence, all art is a type of cognitive caricature based on the psychological effects aimed to enhance, transcend, and distort reality for a given narrative goal. (Cavanagh, 2005; Ramachandran & Hirstein, 1999; Redies, 2007). Much of this recent thinking by cognitive scientists can be traced back to (Gombrich, 2000, 2006).

This thesis will adopt the claim that artists behave as implicit cognitive scientists as a central premise. This claim will be the foundation that we will rely on to pursue the

possibility of designing and testing an NPR system that can render paintings from source photographs, by exploiting the techniques that master artists use to paint fine art from a source image or scene. For this thesis, we define ‘cognitive correlates’ as making links between the experiences of fine art painting with behaviour. The correlates point to behavioural properties of brain, mind, and cognition. Stated more simply we define a ‘cognitive correlate’ as a cognitive theory that correlates to the passed down methodologies that artist communicate about painterly art making. We limited the scope of cognitive theories to mainly those from vision and perception theory, given that our correlates tie to this visual art.

Some cognitive and behavioural correlates of painterly art making are already known, or have at least been speculated on by researchers, many of whom we have introduced in these introductory chapters and will be referred to in the NPR related work sections. Other cognitive theories that correlate to painterly art making are still being debated or speculated but not tested; still others will take scores of research years to fully parse into testable hypotheses. The research described in this thesis will thus add to the work of these previous researchers in specifying the cognitive and behavioural correlates of painterly art.

Besides our specific process, we use the two following related main techniques to limit the problem space.

1.2.1 Source Photograph – Final Painting Pairing As Domain of Inquiry

By framing the problem as the delta or translational remapping between reality (represented by a photograph) and how an artist depicts that reality (represented by the finished painting from the source photograph), we can better parse the difference between these two forms: the source photograph versus the delta change from it to the painting. The delta of that change is where evidence of these exploited cognitive correlates lies. Painterly NPR and our toolkit in particular allow researchers to create this pairing of the source photo and the final painting allowing us to better investigate this translational difference equation. In our four eye tracking studies (Chapter 7), it can be seen how this technique is used to its fullest. We, in a sense, reverse this pairing process by recreating the Rembrandt source ‘sitter images’, with the painterly NPR script to output believable but variably controlled

Rembrandt facsimiles from the original historical Rembrandt masterpiece. This process allows us to test subjects on how they view the sitter photo compared to the final painting and its cognitive based variants.

1.2.2 Salience of What? Artists Directed Salience: Portrait Knowledge

Fine art painting involves making subjective decisions about what aspects of the source are particularly noticeable, important, prominent or salient. Painterly salience involves emphasising details using painterly techniques. For example, given the goal of creating a regal but approachable portrait of the Bishop sitting in front of me, an artist considers what elements or regions of the portrait should stand out, and modifies these regions using specific artistic manipulations that exert a psychological impact on viewers (details in Chapter 3). Cognitive researchers believe that some of the techniques used by artists include the following:

- “simplify, compose and leave out what’s irrelevant, emphasizing what’s important” (DiPaola, 2007).
- “direct the viewer’s attention to the relevant content and to influence their perception of it” (Santella & Decarlo, 2002).
- create “abstractions of photorealistic scenes in which the salient elements are emphasized” (Vanderhaeghe & Collomosse, 2013).

There has been much debate amongst psychologists, vision scientists, image processing experts and even NPR researchers about the true meaning of salience and how to measure it. Santella and Decarlo (2002) tracked a causal viewer’s gaze through a scene in a NPR painting (with an eye tracker) to determine what part of the scene needed to be emphasized. Our view is that painterly salience (emphasising details via painterly techniques) is authored not by the viewer but by the artist -- to emphasise what she believes she wants to convey. It is a goal-oriented endeavour. In a landscape scene, for example, with mountains and a stream in the background and two horses and a tree in the grassy foreground - what should be emphasised? The tree or perhaps one or both of the horses? Zeng and Zhao have attempted to

deal with this problem in NPR by creating hierarchical salient parse trees of a scene. In the above example, the mountain, stream, background, grass, tree and horse objects would be mapped by attention priority in a salient priority tree structure form. Then emphasis and filtering painterly techniques may be used to complete the painting, say with one house at the top of the priority tree (Zeng et al., 2009). Since we take a knowledge-based approach, we again limit the problem space to one specific genre with considerable knowledge data: fine art portraits. A portrait contains the beginning of a hierarchical tree structure of salience.

The approach to salience taken here is to start with portraiture and then expand to other genres as techniques are developed and refined. Portraiture makes more extensive use of cognitive science research than other art forms such as landscapes (research that extends considerably beyond face-specific knowledge rules).

In this work, we attempt to collect and use both general artistic methodology and (when appropriate or needed) specific portrait knowledge (what we call face semantics, i.e., the fact that eyes are more salient than hair). Hence, our painterly NPR system has no specific software code for portraits and can be used for any painterly work. The system includes processes to deal with face semantics in the scripting systems that allow known portrait rules to be exploited. Once other semantic rules from other kinds of art genres (e.g. landscapes, abstract) are known we will move on to generalize our NPR language-based methods (in data passing and script form) to these other genres.

1.3 NPR Painterly Rendering

Non Photorealistic Rendering (NPR) is a computer graphics technique, which creates imagery with a wide variety of expressive styles inspired by painting, drawing, technical illustration, cartoons and mapping (Gooch & Gooch, 2001). This is in contrast to typical computer graphics, which focus on photorealism. NPR already has applications in video games, animation, movies, architectural and technical illustration, and rising fields such as computational photography. More recently, NPR also has applications in visualization, learning and medicine (e.g. communication systems for autistic children) where filtering out unnecessary detail is important.

NPR is sometimes technically categorized into two broad groups, based on the input and control space a system uses: 1) 3D or object space systems; and 2) 2D or image space systems (Collomosse, 2004). Our NPR system, which is concerned with modelling the perceptual and creative process of human artists and painters, falls into the latter camp of image-based approaches. However, while its main processes use image-based techniques, it is capable of using information such as shape and depth data from 3D source materials. Another more computer-modelling grouping in NPR is to categorize NPR systems by the traditional media they simulate. Our system, given its goal, is typically classified as ‘artistic rendering’ (AR) or ‘painterly rendering’ (PR).

Simulating artistic depiction rather than photorealism builds a computational process that can take advantage of how illustrations differ from real scenes. Drawings, illustrations and artistic painting allow for stylized representation, such as emphasising of shape changes or boundaries and de-emphasising the less important features of the depicted scene. Therefore, artistic depiction is seen as a positive in certain fields.

Many NPR software programmers and researchers rely mainly on 3D or image based computer algorithms to quickly and efficiently create NPR output. A strong subset try to understand the 40,000 year old craft of fine art painting, incorporating at least some known artistic techniques into their algorithms (Barla, 2006; Chen et al., 2004; Haeberli, 1990; Hertzmann, 2001; Winkenbach & Salesin, 1994). A smaller set attempt to exploit current knowledge of vision and cognitive science in their NPR algorithms, sometimes tying them to art craft (Collomosse & Hall, 2002; du Buf et al., 2006; Gooch et al., 2002; Santella & Decarlo, 2002). Some cognitive scientists are interested in understanding artistic techniques from a cognitive science perspective (Campbell, 2004; Cavanagh, 2005; Pinna, 2008; Ramachandran & Hirstein, 1999; Redies, 2007; Solso, 1996). Our interdisciplinary work attempts to combine all these interests, to create a NPR toolkit that supports these interrelated fields and can be used as a research toolkit for further investigation.

Figure 2 Using a source sitter photograph (lower right), our NPR system uses learned semantic and artistic constructs to output a Rembrandt facsimile (left) with the ability for repeatable/customizable manipulation (here to test effect of textural highlights all pointing to a dominant eye) for both science and art studies.

1.4 Goals and Contribution of Thesis

1.4.1 Goals

This thesis pursues the five interrelated and interdisciplinary goals of:

1. making certain artistic painterly practices explicit, by looking for common formalist denominators in what has been written about art process and by looking at artworks themselves.
2. relating these artistic painterly practices summarized in #1 to current knowledge in the perception and cognitive sciences.
3. building a parametric painterly Non Photo-realistic Rendering (NPR) toolkit which attempts to align with these artistic painterly practices in #1 and #2 to create a toolkit for research.
4. testing the NPR toolkit's influence on art and human perception understanding as well as adding to specific science knowledge in #2 by achieving study results for our

‘detail-gaze hypothesis’ (i.e. does textural detail of a painting guide the viewer’s gaze to selected regions). This is performed by tracking the eye movements of viewers examining Rembrandt’s (and others) portraits in which the relative level of detail is systemically varied (decoupling it from other factors such as content, lighting, and spatial layout) using our toolkit. We were further interested in bringing scientific evidence to the possible correlation of art viewer’s looking and liking as well as art scholars theories that Rembrandt was one of the first to explore these detail-gaze techniques.

5. exploring the implications for the tests in #4 for future scholarly research in both the arts and the cognitive sciences.

1.4.2 Contributions

With these goals in mind, using interdisciplinary cognitive science techniques, we report on three related contribution areas in the fields of cognitive understanding of artistic practice and computational NPR modelling:

- 1) Using a cognitive science, poetics-based methodology, we produce a ‘soft rulebook’ of cognitive and perceptually based fine art passed down painting techniques. Analysing a comprehensive corpus of historical and modern art practice documents (~ 200 sources), we extracted a rudimentary treatise of the passed-down artistic painterly processes that functions independently of specific crafts or media and relate clearly to aspects of human cognition and perception, structuring this information for computer modelling in a parameterized form. We believe this can be a useful knowledge base for NPR and cognitive scientists. The material is for general fine art painting as well as for portrait domain specific.

- 2) Informed by the findings from 1), we design and implement a novel painterly NPR system using a parameterized computer modelling approach that takes advantage of cognitive and painterly knowledge soft rules. The scriptable and extensible Java/OpenGL based open source toolkit is intended to function as both an instrument for scholarly research as well as a working painterly NPR system, which appears to give wider and more organic results than typical systems. Specifically it implements the following novel algorithms including:

- Global/regional awareness: hierarchal perception blob trees facilitate regional and global communication from child to parent blob relieving a stated problem in NPR (too local).
- Semantic region parsing language: universal but currently set for portrait/face knowledge.
- Allow decision, iteration and smart palette decisions, based on light, volume or semantic content regions.
- ‘Tone first’ colour model: precise tone (JCh CIELAB) first, then semantic mapping into smart palettes: either regional and/or semantic palettes. More art centric colour model.
- Strokes: cognitive blob based tools: density, detail, complexity, value-blending control. Can utilize volume (Z buffer) and/or tonal gradient with regional intelligence for stroke direction and decisions.

3) We perform and report implications of four vision science, eye-tracking experiments to first test our painterly NPR system as a scholarly tool as well as give original evidence for the mechanisms artists (specifically Rembrandt) have agency over viewer’s pattern of gaze. (Figure 2)

Results showed that the viewers’ gaze was attracted and held longer by regions of relatively finer painted detail (Experiments 1, 2), by textural highlighting (Experiment 3), and that artistic ratings increased when portraits strongly biased gaze (Experiments 1, 2, 4). These findings imply that successful portrait artists, possibly dating back to techniques of Rembrandt, rely on an implicit understanding of how gaze is directed by relative detail. The six major findings are:

1. Rembrandt-renderings lead to a ‘calmer eye’ (versus sitter photos) in the viewer.
2. Eye regions of portraits attract most frequent fixations. 52% - collar: 4.3%

3. An eye with greater detail attracts a first fixation in less time, for a longer duration, and attracts more repeat fixations. - Detailed eye attraction magnifies when it is across from a lost edge collar.
4. Forced-Choice preferences show bias for pro-Rembrandt renderings - when asked, none of participants were able to articulate a reason for their choices
5. Textural highlights influence looking patterns.
6. Artistic merit ratings predicted by looking frequency to the biased eye. The more a viewer's gaze is guided, the more artistic merit they give the work.

1.5 Overview of Thesis

In Chapter 2, we review the Related Work in NPR, and the related areas to artistic practice within cognitive and perception based science.

In Chapter 3, we review our cognitive science based data gathering methodology to extract correlated knowledge of the human artistic painterly process and present a concise overview of our soft rules, from an: artistic, cognitive science and computer software implementation perspective.

In Chapter 4, we present an abridged published paper on original research, which goes in depth to a specific aspect of artistic process, using a mixed art history / cognitive science methodology, on Rembrandt's use of texture agency in his late portraits. This early paper's hypothesis is later tested using our toolkit and eye tracking methodology.

In Chapters 5 and 6, we describe our painterly NPR system, its motivation, design, components and describe its processes and results.

In Chapter 7, we report on our four eye tracking studies using our NPR system, both to validate the toolkit as a tool for scholarly work and to further the field. We report on results specific to vision science and art practice - in general and specific to our early Rembrandt texture agency hypothesis from Chapter 4.

In Chapter 8, we conclude the thesis with a summary of the work presented, its ramifications and future directions.

Chapter 1: Introduction

Appendices include A) system documentation of all parameters of our NPR toolkit scripting language and B) list of sources (books and articles) used for our art practice poetic methodological research.

2: Related Work

We overview related work in 1) Non Photorealistic Rendering and 2) Perception and Cognitive Science as related to visual art practice.

2.1 Non Photorealistic Rendering

2.1.1 NPR Introduction

Non Photorealistic Rendering (NPR) is a computer graphics technique which creates imagery with a wide variety of expressive styles inspired by painting, drawing, technical illustration, cartoons and mapping (Gooch & Gooch, 2001). This is in contrast to typical computer graphics, which focuses on photorealism. The descriptor Non Photorealistic Rendering coined in 1994 by Winkenbach and Salesin (1994) has been criticized as a weak representative title for the field since it describes what the field is not. Therefore, while NPR is still the dominant title, many papers use other terminology including expressive rendering, artistic rendering, painterly rendering, art-based rendering, etc. Within 3D and 2D (image based) NPR techniques, stroke based or painterly rendering typically uses a 2D source such as a photograph and creates a list of strokes to be rendered on a new canvas. NPR already has applications in video games, movies, architectural and technical illustration, animation and rising fields such as computational photography. NPR also has applications in learning and medicine (e.g. communication systems for autistic children) where filtering out un-needed detail is important, as is true in technical illustration.

2.1.2 NPR Related Work

Within NPR research, a number of painterly rendering techniques were developed in the early 1990s, starting with Haeberli's pioneering work (1990) which introduced painting with an ordered collection of strokes described by size, shape, colour and orientation. Litwinowicz (1997) created a fully automated algorithm based on Haeberli's earlier work producing

paintings by using short linear paint strokes tangent to Sobel edge gradients. Relying on Haeberli's work as a base, many approaches appeared that used local image processing techniques in an attempt to achieve relevant high level content semantics to better control stroke/placement/colour choices. Hertzmann (1998, 2001) advanced the field by using a multi-pass system of coarse-to-fine curved b-spline strokes aligned to a layered course-to-fine image difference grid with multiple styles. This was inspired by his observation that artists begin a painting with a first pass of large broad strokes and then refined the process with smaller strokes to create detail. Our system is inspired by his work and the general procedure of using artist knowledge to dictate parameters. For instance, we have moved away from a multi-pass course-to-fine grid for stroke placement choices and toward tonal masses based on lighting and drawing types. We also are developing a colour system that moves away from colour sampling with perturbations to a system, that uses source tone (i.e. luminance) to remap into one of many known semantic colour temperature mapping models. This soft rule that instructs the system to 'sample the tone (greyscale value) first, then with that tone remap into a colour model of choice', came out of our artistic practice investigations and differs from typical NPR colour choice algorithms.

More global oriented computer vision techniques to model scene semantics began to appear in an attempt to both automate the higher-level process and improve aesthetics. Gooch and others (Gooch & Gooch, 2001; Gooch et al., 2002) first proposed segmenting the image into homogeneous greyscale regions leading to a significant reduction in the number of brush strokes. Segmentation was also used by Santella and Decarlo (2004; 2002) who presented a salience extension of guiding the painting process by eye movement, emphasizing which area the user thinks is important (Figure 4). Given our hypotheses that human artists will emphasize the salient parts of a scene and de-emphasize the other parts, we feel their salience-based eye tracking work was a major innovation to NPR. However, they derived the salience of a scene by eye tracking the way in which a person viewed a photograph of the scene and then used that data of where the casual viewer gazed to emphasize that area in the final NPR output. We will discuss the many different painterly 'soft rules' artists used to emphasize or deemphasize an object or area in a painting in the next chapter. Salience in general, is a major issue in cognitive driven NPR research. Our approach takes the salience and eye tracking work begun here, but looks at it as an authored narrative of the artist.

An automatic system for salience adaptive painting, driven by machine learning, was presented by Collomosse and Hall (2002, 2006). Much of Collomosse's work has affected our views on salience techniques and trying to build upon their work. Around this time, Collomosse, Gooch and others (Collomosse & Hall, 2002; Gooch & Gooch, 2001) described a major problem of most NPR approaches – they implemented local not global techniques and therefore restricted “attention to local image analysis, by independently examining small pixel neighbourhoods” and these local techniques “can give no real indication of salience in an image” (Collomosse, 2004). Given that blind local processing/analysis especially in image based NPR was a major impediment to provide artistic salience in NPR, he and several other researchers looked for efficient global solutions (Collomosse & Hall, 2002; Gooch et al., 2002; Gooch & Gooch, 2001; Huang, Fu, & Li, 2011; Kagaya et al., 2011). Inspired by this work, and given our artist practice / cognitive centric approach, we have built a hierarchical perception blob tree structure into our system. Our ‘blob thinker’ algorithm starts with an optional semantic map for portraits (background, face, hair, clothes, ...) as the parent blobs and then iterates through these hierarchically to create progressively smaller child blobs inside their parents via luminance regioning. This approach simulates how artists work through a painting (region by region) and perceive a small cognitive appropriate region to work on, not specific brush strokes (i.e. I am working on the top of the nose now). This technique both gives a more artist centric approach to move through an image (than say Hertzmann's original progressive x, y grid passes) and allows for regional and global communication by allowing child blobs to communicate up the tree to obtain local and global comparative information they can use to make colour/stroke decisions such as “have we already put too many texture highlights on this side of the face, if so I will be sparing in my local area”. While we have built efficient and effective perception blob tree communication systems, to date we have not fully exploited all the local to global communication possibilities. Understanding cognitively how human artists make these creative perceptual relativistic decisions so we can exploit in our system, is a still a source of research (Miall & Tchalenko, 2001).

Hertzmann (2003) summarized and categorized many stroke-based rendering (SBR) techniques. Hertzmann and later Hays and Essa (2004) used the variation of low level parameters such as stroke length to create different painterly styles. Our system relies heavily

on hierarchical unitized parameters, which build up into higher constructs from low-level XML scripted language components based on painter methodology. Collomoose and others (Shugrina et al., 2006) formed a mapping onto parameters using a high level emotional parameterization derived from a facial tracking system. And Setlur and Gooch (2004) used a technique to create rendered faces by exploring facial emotions using low and high spatial frequency details based on human vision knowledge. Gooch as well as Collomosse and Hall relied on several human vision techniques (Collomosse & Hall, 2006b; Gooch et al., 2010, 2002; Shugrina et al., 2006). Rodrigues, du Buf and others (du Buf et al., 2006; Roberto et al., 2006) used human vision techniques to better tie visual perception theories to NPR and painterly rendering. Kang and others (Kang et al., 2006) introduced several techniques to adaptively generate stroke and other artistic parameters or attributes to tailor to a wide variety of SBR styles.

Figure 3 (Zhao & Zhu, 2011, 2013) work with brush templates for portraits.

Zeng and Zhao have attempted to deal with open-ended salience issues by implementing user specified hierarchical salient parse trees of a scene (Zeng et al., 2009) which use a knowledge based philosophy similar to our methods. Our tree structures are automatically generated based on semantic content maps (typically of portrait regions) and progressively refined luminance calculations (i.e. hierarchical blob trees). Zhao has also made strides in

portrait knowledge based NPR (Zhao & Zhu, 2011, 2013) using premade brush templates to describe stroke size and placement (i.e. a brush dictionary) with very strong results. They use a similar knowledge based philosophy, although our work tends to use artistic and cognitive knowledge (see Chapter 3) where their methods use constructed content tree and brush dictionaries (Figure 3). Therefore, much in our systems are complimentary -- therefore we have had discussions with Zhao including the possibility of implementing their brush dictionary into our system in the future, as well as extending it based on additional knowledge semantics. Our language-based toolkit approach, which has adaptive mechanisms that perform scripted logical calculations using image and data 2D image buffers, allows for inclusion of other researchers techniques. To date, we have not implemented their techniques but would like to in the future.

Figure 4 DeCarlo and Santella [2002] results using their abstraction using eye-tracking data (left) and later Bottom Row: Our automatic abstraction, and (Winnemöller et al., 2006) version also using eye tracking.

As NPR moves past its early stage, many scholars have begun to systemically evaluate NPR's worth, current state, best techniques and future direction. Salesin (2002) put out seven grand challenges to the NPR community in terms of NPR moving into a new level of maturity which was recently revisited in a state of the art survey paper (Gooch et al., 2010). Others researchers have recently surveyed the state of the art techniques of NPR or painterly rendering with an eye to the future maturing direction (Collomosse et al., 2012; Hegde et al.,

2013; Isenberg, 2013; Vanderhaeghe & Collomosse, 2013). There has been strong debate over using science-based evaluation methods for NPR (Isenberg, 2013) versus calls that NPR is best left as art form to be appreciated and not measured. (Hall, & Lehmann, 2013). We hope that by making a language based painterly NPR toolkit we can help support best practises, experimentation and evaluation of the field.

2.2 Perceptual and Cognitive Sciences: Visual Art

Recent advances over the last 20 years in our understanding of the cognitive and neuroscience of perception have encouraged cognitive scientists to turn their attention towards understanding the visual art process. While there is a danger with being too reductionist about the science of art making (Hyman, 2010; Ione, 2000a; Seeley, 2006) many see a value in exploring this field from a cognitive science perspective. Painted artworks are visual artefacts intentionally designed by the artist's use of light, colour, shape and contexts to bring about a range of emotional, perceptual and cognitive responses in art viewers. Both the act of painting by artists as well as appreciating artwork show how humans acquire, process and use visual and textural information embedded in a painting in order to understand and evaluate its content (Gombrich, 2000; Gombrich et al., 1972; Hyman, 2010; Seeley, 2006, 2011; Zeki, 2000). Cognitive science, generally, is an investigation of how humans acquire, process and use perceptual information from their environment. Therefore, cognitive science can be used to model the psychological processes by which humans both create and view artworks.

Looking scientifically at art is still controversial in some circles, as well; cognitive scientists have come to the practice from different fields and expertise, albeit all with a theory of mind perspective. Many were influenced by earlier work by Gombrich (2000, 2006). Most scientists in this area use a perception sciences lens (Cavanagh, 2005; Conway & Livingstone, 2007; Grossberg, 2008; Mamassian, 2008; Pinna, 2008; Ramachandran & Hirstein, 1999; Solso, 1996), while others are interested in how humans process or create a thing of beauty, or as Zeki (1999) termed it, a neuro-aesthetic process (Chatterjee, 2002; Hyman, 2010; Ione, 2000a; Ramachandran & Hirstein, 1999; Zeki, 2000), others see the act of creation from human consciousness perspective and some come from pattern recognition

(Graham et al., 2010; Redies, 2007). Many also see the research as a two way street, understanding how the human mind can create artwork for its own sake (i.e. the study of the art process from a cognitive perspective) as well as mining that artistic process and knowledge artists seem to have intuited (intentionally or not) about perception over modern human history as a way to give additional support to the cognitive sciences in general. Our work is interested as well in both of these goals.

Many cognitive scientists have asserted and given varying degrees of evidence that visual artists exploit the properties of the human cognitive and visual systems (Cavanagh, 2005; Grossberg, 2008; Mamassian, 2008; Seeley, 2011; Zeki, 2000). Rather than adhering to physical properties of the world, artistic paintings reflect perceptual shortcuts used by the brain. (Graham et al., 2010; Livingstone, 2002; Mamassian, 2008; Pinna, 2008; Redies, 2007). Researchers note that ambiguities in typical visual perception were resolved thanks to prior constraints that are often derived from a human's knowledge of the makeup of natural scenes (Enns, 2009; Henderson, 2003; Hochberg, 1968). Ambiguities in visual arts, however, while also somewhat resolved by conventions of perceptual priors, more typically are resolved from other sources such as artistic stylistic or arbitrary choices (Mamassian, 2008; Myin, 2000; Zeki, 2000). For all humans, researchers have shown that vision is an active, or a constructive process (Cavanagh, 2011; Enns, 2004; Findlay & Gilchrist, 2003; Zeki, 2000). We 'construct' the internal model based on the representation of these experiences (Cavanagh, 2011; Enns, 2009; Grossberg, 2008; Mamassian, 2008). This emphasizing the salient is what Zeki (2001) refers to as the Law of Abstraction, in which the particular is subordinated to the general, so that it can represent many particulars. Ramachandran and Hirstein (1999) proposed eight principles in their cognitive model of art most notably the peak shift effect, which is similar to the theory of exaggerating the salient features that distinguish a given object of interest from other objects. Zeki and Ramachandran have been criticized for being too reductionist with their theories on art creation and aesthetics (Hyman, 2010; Ione, 2000a, 2000b; Myin, 2000). This reductionism is a constant issue for scientists trying to study the arts.

Another fundamental feature of human vision is that our experience of a scene or an artwork is not uniformly detailed (Henderson, 2003; Hochberg, 1968; Land, 1999; Melcher & Colby, 2008; Molnar, 1981; Yarbus, 1967). Thus our viewing experience is actually one

that extends over time, including periods of fixation, in which the eye position is almost stationary and visual information is taken in, interrupted by saccades, rapid movements of the eye from one image region to another during which we are also effectively blind (Enns, 2006; Henderson, 2003; Land, 1999; Melcher & Colby, 2008). Given then the constructed nature of vision, only the salient stimuli are fully updated (remapped), on average about 3-4 attended items (Prime et al., 2007), the rest can change, disappear or reappear without us noticing the change (Melcher & Colby, 2008; Simons & Rensink, 2005).

Artists employ many different forms of nonlinear luminance compression schemes in order to depict the world (Conway & Livingstone, 2007; Graham et al., 2010; Hochberg, 1980; Livingstone, 2002; Redies, 2007). Artists in scientific terms ‘adjust the image statistics’ of their subjects during the process of creation so that, in their works of art, statistics are closer to those encountered in complex natural scenes (Graham et al., 2010; Redies, 2007). Cognitive Scientists believe that art can be tuned more precisely to nervous system properties than natural scenes are (Pinna, 2008; Solso, 1996, 2000).

3: Identifying Cognitively-Based Techniques within Artistic Painting Practice

3.1 Introduction

Figure 5 Source photograph and final painting by fine artist James Horton from (Horton, 1995). How does a human artist ‘remap’ the source photo (left) to a finished painting and what cognitive processes does the remapping exploit?

Examining it from a scientific perspective, this chapter explores the problem of mapping a source scene, which we can represent as a source photographic image into an often revered fine art painting. For ease of exposition, I will refer to this as the *remapping problem* hereafter. That is, working with mainly concerns of light, but also volume, colour and content, an artist remaps or distorts the reality based source scene or photograph (Figure 5

left) to tell a narrative in the final created output painting (Figure 5 right). They perform this extraordinary and creative feat via passed down artistic knowledge and process techniques that many researchers believe exploit human perception and cognition (Cavanagh, 2005; Gombrich, 2006; Livingstone, 2002; Zeki, 2000). We refer to these passed down painterly heuristics within the thesis as ‘Soft Rules’.

Since we are concerned with making a parameterized computer model of this process in the form of a painterly NPR toolkit, we are interested in systematically:

1. categorizing using a ‘poetic’ methodology, this passed down artistic knowledge which we reviewed from ~ 200 books and source materials (see Appendix B for references) on artistic practice, what we call a ‘soft rulebook’ in Section 3.2,
2. aligning our bracketed subset of these artist ‘soft rules’ with known science-based ‘cognitive correlates’ in Section 3.3 and
3. using knowledge from both, describing how this bracketed subset of ‘soft rules’ informed our NPR toolkit design in Section 3.4 and Chapter 5 and 6.

We take on each of these three perspectives individually in the next three sections but use a convention established in the first section of defining a more specific subset of knowledge ‘soft rules’ that will be discussed within all three sections. The convention uses square brackets rules such as ‘**[R1.1 Tone first then colour]**’ to define and re-discuss a subset soft rule, where the R1.1 refers to the artistic subsection it came from in Section 3.2 for better cross-referencing. This convention gives us the ability to separate sectionally the three methodological discussion styles: 1) artistic, 2) cognitive science and 3) computer software implementation, while still discussing shared major themes presented in the first section of passed down artistic knowledge. The third section is abridged since the computer software implementation of our NPR toolkit is detailed in its own Chapters 5 and 6. In later chapters, we use our NPR toolkit to perform deeper scientific studies on some of our knowledge rules involving gaze pattern authorship.

3.1.1 Motivation and Goals

We begin with our artist passed down painterly heuristics or soft rules. This knowledge is a body of observations, techniques and strategies used or prescribed by artists and art theorists. We have extracted and summarized these painterly heuristics from literature/interviews using a poetics methodology (defined below). Artists do not use scientific methods to collect, refine and pass down knowledge. Therefore, deriving art methodology knowledge from interviews and reference materials with working artists on their passed down artistic process (as well as art experts, art instructors) is obviously fraught with ambiguities, seemingly ad hoc methods, differing opinions and techniques that have no current software based analogy. Moreover, many in the art world are somewhat uncomfortable about overly dissecting an aesthetic pursuit such as fine art painting. With these challenges in mind, the goal was to create an NPR toolkit research system that could be iteratively updated and refined based on information and user studies from the NPR, art criticism and cognitive science sectors. We hope a system that attempts to quantify the process in a parameterized way can serve as a tool in both science and art discussions including respecting and informing the human art making endeavour by trying to model the process.

3.1.2 Methods and Scope of Artist Process Categorization

In general, we tried to collect and systemically categorize universal painterly knowledge practices; however, painterly practices are based on a mixture of prioritizing the process based on 1) light (2D), 2) volume (3D) and 3) content semantics (the meaning of objects), so you will note when content semantics are needed, we mostly default to one specific area, the content area of portraiture. The data presented then in Section 3.2 is a mixture of more universal painting soft rules and content based sub-rules specific to portraits. While portrait knowledge is used as the default semantic placeholder, we believe it is possible, with further research, and using the techniques here as a start, to add more semantic areas to the knowledge base. Specifically, much of the portrait work translates to figurative and still life (i.e. strong binary foreground/background, with known semantic knowledge about the foreground object) and then to other genres such as landscape, abstract and other areas.

3.1.3 Poetics Methodology

Our research approach for this section, Section 3.2 Artistic Soft Rules, used a poetics methodology to uniformly collect and correlate artistic knowledge. Poetics started with Aristotle's Poetics (Aristotle, 1961) and has become a modern methodology for studying literature, drama and art from a formalist perceptive. The methodology uses **“the systematic review, analysis, and alignment (or perhaps ‘synthesis’) of the expressed poetics of expert practitioners”**. For our work we used more modern techniques as described in David Bordwell's *Poetics of Cinema* (Bordwell, 2007). In a strong statement on the value of poetics as scholarship Bordwell states:

The poetics of any medium studies the finished work as a result of a process of construction—a process, which includes a craft component (e.g., rules of thumb), the more general principles according to which the work is composed, and its functions, effects, and uses. Any inquiry into the fundamental principles by which a work in any representational medium is constructed can fall within the domain of poetics.

Using this poetics methodology, we did a systematic review, analysis and categorised alignment of the expressed poetics of fine art painters/instructors/theorists explaining their craft of fine art painting. We first reduced the material of over 200 books/source material to 30 pages of correlated references and documented notes and eventually to ~ thirteen pages of passed down painterly soft rules in the following section. See Appendix B for the list of our approximately two hundred source books and articles for this section. Our goal was to filter or translate craft talk and obfuscation into cognitive based action, process and embodiment based soft rules. Eventually we catalogued, correlated and validated the coherence of known methods that artists have passed down, as the general techniques they use, void of craft jargon. We built up these correlated notes, categorized them into areas from original sources but with several iterations, reducing them to essential processes that were corroborated from many sources – the passed down soft rules of painting. The inclusion criteria for Soft Rules were:

- cognitive or at least human process based
- action, process and embodiment based

- coherence; corroborated from many sources
- known methods that artists have passed down
- categorizable into areas; essential processes
- not specific to particular medium-based technique
- translation of craft talk and obfuscation; void of craft jargon

The following is our data of artist reference materials and artist interviews; a more condensed version appears in the authors paper on the subject (DiPaola, 2007). Two caveats before we begin this section: One, these are notes on how artists currently and historically discuss their passed down artistic practice soft rules, albeit refined, systematized and categorized using our poetics methodology – the notes are still in their non-scientific voice and void of 100s of more academic references or 3rd person studies. Two, note the bracketed subset of soft rules, (e.g. **[R1.3 Unequal value]**), interspersed through this section, as our main methodological convention for discussing a more specific and concise subset of these intertwining and discursive artistic soft rules in our 3.2 cognitive correlates and 3.3 NPR computer software sections.

3.2 Soft Rules of Fine Art Painting

In the following nine soft rules sections plus the summary, if we cited references normally, the approximately two hundred foundational references would overwhelm and obfuscate the text in this section. Therefore, we only cite specific quotes within Section 3.2, and have the bulk of 200 fully cited foundational references, tagged by subsection, in Appendix B.

3.2.1 Tonal

The artist Max Meldrum in 1919 argued that painting was “the science of optical analysis by means of which the artist, in carefully perceiving and analysing tone and tonal relationships, could produce an exact appearance of the thing seen” (Colahan, 1919). Tone

(also called value, and specifically referring to measurements of *luminance* in the language of vision science) was the most important component of the art of painting. The next most important consideration was proportion, referring to “the superficial area occupied by one tone”. The third consideration was colour, which many naïve viewers of art are surprised to learn is the least important component. Pablo Picasso famously stated on this subject, “Colors are only symbols. Reality is to be found in luminance alone ... When I run out of blue, I use red.” (Wenrich, 2013) **[R1.1 Tone first then colour]**.

The tonal value (tone) is the general lightness or darkness of an object. Artists squint in an attempt to focus on tones while ignoring differences in colour. They also adaptively squint to see more in general tonal masses, allowing artists to decide the size of mass they are interested in tone sampling at the moment. **[R1.2 Regions as likeminded tonal masses]**. By masses, they mean a perceptual grouping of similar tonal values that make a loose shape. At the fine or stroke level, final colour is generally derived from a ‘sampled’ tonal value. This ‘tone first approach’ (Figure 6) [R1.1] allows an artist to be more creative in their colour choice without sacrificing the painting’s integrity.

Figure 6 ‘Tone first’, then map using some colour method – lie warm/cool colours. From (Sanden & Sanden, 1999; Saper, 2001)

An artist tries to limit the overall number of values in painting. In portraiture, still-lives and figurative work, values can be categorized into particular ‘types’ where the first three are the most important:

- Body tone (or light) – in the light source; in direct light (usually warm colours).
- Halftone – where the light begins to turn; in between light and shadow.
- Body shadow – away from the light source; darkest area (usually cool colours).
- Other Types: Cast shadows, Reflections, and Highlights.

One commonly used approach in painting is to first paint the gross masses in three divisions of value – body tone, halftone and shadow [R1.2]. Within these gross areas fit degrees of finer tonal gradation (light to dark) which benefit from the semantic information of what gross tone area they are in (i.e. in the body shadow region). Again this demonstrates an artist’s ability to discriminate arbitrary grouping or sampling of likeminded group of tonal value as a shape or region [R1.2] and therefore grouping large masses first, then progressively smaller ones as they procedure through creating a painting.

A painting must have a domain value, either light, or medium, or dark. For example, a small sketch might be dominated by a light value, an impressionist landscape is dominated by a medium value and a Rembrandt portrait is dominated by a dark value. With three values: dark, midtone and light, one should be dominant, and the other two together typically make up less than half of the first, none being equal amount to the others. This ‘rule’ is often referred to as Unequal Dominate Value [**R1.3 Unequal value**] (Figure 7).

Figure 7 Unequal Dominate Value: Left: Strong: big dominate value and unequal sub values. Right: Weak: little dominate and in addition, equal balance. From (Brown & Lehrman, 2004).

3.2.2 Colour

The tonal value points to what colour to pick within the other rules of colour formulation [R1.1], the most important of which is colour temperature. Colour temperature refers to the relative warmth or coolness of any given colour. Warm colours and cool colours work harmoniously against each other both globally and within regions of a painting. Warm light refers to light that appears more yellow, orange or red, while cooler light has a blue, green or violet in it. [R2.1 Warm/Cool Colour] See Figure 8. The foremost concern when picking a colour is getting the value correct [R1.1]. Many artists who exploit colour temperature rules do not simply darken a colour but add its complement to grey it down. Note, while globally in a painting the idea is to separate warm and cool in the major value regions, artists also talk about locally (and minimally – as in small strokes) placing warm and cools side by side. By placing a warm colour next to its cool version, it is said that a painter sets up a beautiful visual vibration. A typical process would be to determine the colour of the light on the area of your subject, say the body tone area. In painting, everything left in shadow takes the opposite temperature. When the light areas are painted in warmed tones, generally speaking,

the artist will paint the shadows in cool tones. Most art portraits, for instance, have warm light (skin) and cool shadows but there are many exceptions. So warm lights produce cool shadows and cool lights produce warm shadows.

Figure 8 Tone (Value, Luminance) First (A), then balancing (separating) warm and cool colour temperature across colour wheel (B). Images: Bruce MacEvoy.

The rule of unequal balance dictates that both light and cool temperatures cannot be shown equally. [R2.2 Unequal colour]. Usually within a painting, the warm light is dominant over the cool in shadows. Alongside a dominant value [R1.3], artists assert that successful paintings have a dominant colour. See Figure 9.

Most painters use the classic three primary colour system around a temperature based colour wheel. More recently many artists have adopted the Munsell Colour Notation system's five 'principal' colours (red, yellow, green, blue and purple) which are spaced equally around a colour wheel. Artists typically make a decision about palette or colour harmony for a given painting before they start (a colour plan). With the standard three primary systems there are many colour harmonies, here are two well-used examples:

- **Analogous Colour Harmony** includes adjacent wedges of colour on the colour wheel, including greyed-down neutrals and light and darker version of the colours themselves. A yellow-green example wedge is: orange - yellow - green. With this

system, an artist would pick the dominant colour and then the wedge. Discords can be added at five and seven o'clock from the dominant (at 12 o'clock). Used very sparingly, discords give energy and capture the eye's attention.

- **Complementary Colour** Harmony is the most versatile for portrait use. Complementary refers to the colour directly across the colour wheel. Since all skin has some aspects of red and yellow in it, it can easily adapt to a red-green, yellow-violet or blue orange schemes, keeping in mind the principle of unequal balance.

Munsell's five 'principal' colours shift the complements, so red's complement is blue-green instead of green, yellow is purple-blue rather than blue. Greying down a complement will allow you to use more of it. On the colour wheel, the closer a colour moves towards the centre the more the complement is introduced to it, greying and reducing its intensity. A colour and its complement cancel each other's identity and form a version of grey. A colour finds its identity and actually is intensified if it is put in juxtaposition to its complement.

Figure 9 Unequal Colour (Brown & Lehrman, 2004)– Warm/Cool colour wheel (bottom left) shows dominant colour (black dot), with its complement directly across the wheel, then the 2 discords equidistant offshoots and how they are all unequally weighted in the two upper abstract images (green dominate and blue dominate respectively).

Here is a summary of how an artist can use colour temperature: 1) a painting should have a dominant colour that is readily evident [R2.2], 2) complementary colour (from across the colour wheel) may be introduced, occupying much smaller space at full intensity, or greater area if greyed [R2.1], 3) discord colours (equidistant on the colour wheel from the dominant hue and from each other) should be added sparingly in approximately equal amounts. Discords bridge together and add visual excitement (see Figure 9).

3.2.3 Textural Highlights

The very last strokes painted on a surface are the highlights representing reflections from the light source itself. Textural highlights are those edge or blob based light-valued brushstrokes that are interpreted by the viewer as specular highlights (i.e., shiny regions

indicating maximum light being reflected from the viewed surfaces). Highlights need colour too. The most effective device to convey the intensity of this is to change the temperature of the light source, i.e. warm light sources yielding cool highlights and vice versa. Highlights can be overused - they should have a location shape and colour and should not be painted if they are not there. Source highlights are the one thing that will move with the viewer's gaze, so they give away the cameras or viewers eye point. **[R3.1 Highlights - eye fixations]**. One of our four cognitive science eye tracking studies using our NPR toolkit presented in Chapter 7 was conducted to give evidence on how and if textural highlights techniques are used as an authorship tool to influence a viewer's eye gaze – in general and specific to Rembrandt's late portraits.

3.2.4 Shadows and Halftone 'Core'

In paintings, many artists assert that there should be no textures/details or strong colours in the shadows. The more gradually an object turns away from the light, the softer and wider its shadow edge will be. Therefore, in a portrait example, a forehead or cheek turns gradually, making a soft shadow edge while a nose would be sharper. Hard edges are found in cast shadows, closest to the object casting the shadow. Soft edges are fuzzy such as in form shadow.

The colour in shadows (and lesser so in light) is influenced by the colour of the background, in the turning down or away planes (e.g. underside of chin and nose). Since the shadow side is darker than the light side, the flow of colour and value is broken up by an object's features (e.g. a nose in the face). Artists generally: avoid strong colour in the shadow (greyed down hues and lower intensity), avoid hard edges in the shadows (softer than those in the light) and avoid strong contrast in shadows (narrower range of values that in the light areas). Areas where two objects (e.g. palm and cheek) touch in shadows have reflected light (light bouncing back and forth). These make an effect that causes a slight change in colour and value, making the area of reflected light slightly lighter than the rest of the shadow. It can often be overstated (because it looks so strong to use). One artist stated that reflected light is always darker than you think.

The soft transition edge that is created where the form turns and light meets the shadow is called 'the core' of the shadow. The core is where the colour and value are the truest - neither obscured by shadow nor bleached by light. Therefore, the core is where the artist has hit their colour most strongly. Artists also widen the core area for aesthetic effect (Figure 10). **[R4.1 Relative widen/reduce/harden/soften].**

Figure 10 Exaggerating the core between light and dark areas. (Saper, 2001)

3.2.5 Shapes, Edges and Centre of Interest

Shapes and composition can be more important than content (e.g. subject). Artists link shapes to create pattern and thereby composition. Shapes are related to edges.

Edges occur wherever shapes meet. By softening or hardening edges or making them disappear entirely, the artist strengthens the illusion of form and gives a painting dramatic

flow - this is known as 'lost and found edges'. **[R5.1 Lost & found edges – eye fixations]**. By losing an edge, we allow it to merge with an adjacent shadow, creating a link between objects, which is a powerful tool for design (Figure 11 right). For instance, a subject can emerge from a strongly textured background, yet remain one with it, with the help of edges softened or entirely lost and shadow cores. Edges can control the viewer's eye movement over the canvas, since eyes always move along sharp edges and coast softly over soft edges. On a lost edge, the viewer finds comfort in seeking out the place where it is found again. Sharp edges can be at or in the centre of interest. An edge plan can create a story for the painting [R4.1].

Lines and shapes moving vertically or horizontally convey formality and solidity. Diagonals convey movement and excitement. Every subject has some diagonals. A painter strengthens those elements and perhaps downplays the more static lines.

Figure 11 Centre of Interest, Lost and Found Edges. From (Sanden, 2004) and (Brown & Lehrman, 2004).

An artist creates and draws the viewer towards the painting's centre of interest through the edge quality or detail in the portrait [R5.2 Centre of focus - eye fixations]. The sharpest edges/areas are at the centre of interest, less sharp edges move the viewer's eyes across the canvas toward the centre of interest and the softest edges are placed where the artist wants the viewer's eye to glide over (Figure 11 left). Note that several of our four cognitive science eye tracking studies using our NPR toolkit presented in Chapter 7 were conducted to give evidence of how and if centre of interest techniques along with lost and found edges are used as an authorship tool to influence a viewer's eye gaze – in general and specific to Rembrandt's late portraits (Figure 12).

Figure 12 An example of both 1) edges, then lost for gaze pathing toward a 2) centre of focus detailed eye, which came from our later research studies in Chapter 7 with Rembrandt's *Beret* and other work.

Figure 13 Photo versus painting – Value Clumping in light and shadow of face. (Saper, 2001)

3.2.6 Brush Strokes

The conditions that dictate the placement, size, shape and direction of a particular brush stroke are less verbalized by artists. They can be based more on muscle memory, lived experience and intuition, similar to a musician's finger technique, and therefore not much data exists on brush stroke techniques. There are however many possible known (but conflicting) techniques artists can pick from. Artists often think about the way the surface or plane is orientated and paint in colour and strokes to suggest the surface itself. Often they stroke the paint on in the same direction as the plane. One technique is to pull or push the stroke in the direction of the plane. Strokes put down in the same direction become monotonous, so they are broken with strokes in opposing directions which it is said brings energy to a painting.

It is possible to use a choppy aggressive and thicker technique for less salient parts of a picture, say the clothing and background in a portrait, which can be contrasted with the painterly look of the salient parts, say of the face. For the lost edge, pull the brush away from the canvas. Start your stroke with pressure where you want a found line. Artists state you

need found 'sharp' and lost 'fuzzy' lines to make the planes of the entire periphery project and recede.

Figure 14 Brush stroke direction techniques (Stewart, 1997).

Therefore, in summary, balance and variety seem to be the key. Stroke to match the shape or light gradient, and then complement it with a stroke in the opposite direction (Figure 14). Never completely finish edges. Always overlap rather than putting objects in a row. Use stroke direction with edges to move views to the centre of interest. Again stroking techniques seemed to be least understood, discussed or planned and are often at a muscle memory or intuition level at execution time [**R6.1 Light, shape, content concerns**].

3.2.7 Backgrounds

The background affects the overall relationships and should be painted first or at the same time. The main rule of backgrounds is that they recede compared to foregrounds. There are four main types of portrait backgrounds: 1) the plain solid coloured background, 2) the imaginary background made-up of shapes and colours, 3) A representational background of what is there and 4) an abstraction of what is there.

For a plain solid coloured background, usually an artist will use a significant colour in the foreground to maintain colour harmony. [**R7-8.1 Overall relativity and rescaling**]. The inverted background is one of the most common. Typically, the same kinds of brushstrokes are used in the background as in the portrait but the background the edges are softer than on

the model. Busy backgrounds confuse the image unless they use an unrelated colour to the subject. Warm backgrounds bring out warm flesh tones while cool backgrounds emphasize cooler tones. A light background shows up the shape of a dark figure, while a darker one blends with dark tones on the subject.

3.2.8 Working from Photograph as Source Image

A painting is not the same as the scene it depicts. How it is different is complex, and multidimensional but one way of starting that dialogue is to look at a painting from a photograph of a scene. There are several problems with painting based on photographs including depth of field, value clumping, and colour distortion. Depth of field from photographs does not provide a centre of interest focus, since everything is in focus or the camera focuses in unnatural vertical planes (although some creative selective focus is possible). This means an artist must create the centre of interest. Photos also give false edge readings, making edges too sharp. Cameras also darken the shadows and lighten the bright areas causing value clumping. Colour and perspective distortions are also common in photographs. **[R7-8.1 Overall relativity and rescaling]**. From the photograph to the painting, an artist attempts to simplify, by leaving out the irrelevant and emphasizing the important. Photographs distort the tonal range, known as value clumping making values on either end of the scale compressed (Figure 13). Artists extend the middle value by extending out areas that look dark in value, as well as extending lighter areas to range to darker in value.

3.2.9 Process Plan

From a process standpoint, artists investigate something of interest about a scene or sitter they have chosen to paint, and come up with a ‘visual narrative’ of this scene that includes how it will be expressed using the language and techniques of painting, and how the painting will deviate from the real scene. Factors that may be taken into account during this process include (1) the content of the scene, (2) their career trajectory (e.g. techniques and innovations they have experimented with, advice they have received from mentors, and so forth) mixed with (3) personal feelings and intuitions as they arise in the moment about such factors as what will invite exploration or provide closure. Much of an artist’s time is spent re-

assessing what is now there in front of them, and re-evaluating what to do next based on that data, a process that Gabora (2002) refers to as honing through re-iterated context-driven actualization of potential. With a mixture of convention and innovation they develop:

- A value plan (their main values as well as their balance of values).
- An edge plan and centre of interest plan (what they will emphasize or manipulate via textural detail).
- A colour plan (their palette choices, their relative colour choice method, and their balance of colours).
- A brush plan (brush stroking styles, regions changes)

Some artists stick tightly to such plans, while others deviate widely based on the source and the narrative. As they begin, they exploit a limited tonal and colour range where they can re-scale and re-centre the relative range of all aspects of the source image and source masses (tonally, tonal balance, colour, colour balance, detail, edging, shapes,...) in ways they perceive fit their style and narrative goal. They work through the painting, taking anywhere from hours to days to weeks. They frequently make mental comparisons of how different elements or regions of the painting appear relative to one another. They usually start with large tonally limited masses (i.e., two or three values) that they tone sample mentally. Some artists squint to help both blur out details and remove colour importance. Then they shift to another region of the painting. As they continue, they increasingly focus in on smaller details *i.e.*, smaller brush strokes, and more colour choices. These details are not only rescaled and re-centred from the source but typically are picked using different aesthetic rules from those used initially, such as a warm/cool remapping. These decisions fall into place depending on what commitments they have already made or how the initial potentiality of the yet-to-be painted has been actualized thus far. At some point in this progressively refined process, they cycle into a ‘see, think, paint’ loop, deciding what area to work on next, sampling a cognitive region (like minded tone sampled) from the source, thinking about it -- remapping it conceptually, and through the craft of stroking on canvas, commit this new conception to paper. With each cycle, some of the potentiality of the previous cycle gets actualized, and

simultaneously, potentialities for future iterations arise. They cycle through this process repeatedly until the semantic area they are working on (*e.g.*, the lower face) is complete. They may then make a more detailed pass involving progressive refinement of the area. Light is the primary consideration in deciding how the painting unfolds (in the form of tone sampled shapes), although artists also consider volume to a lesser extent (*i.e.*, stroke up the volume of the nose) and content areas (working in the background or a cast shadow or for an eye, I do something different and specific).

3.2.10 Summary

Using our poetics methodology, we have attempted to systematically categorize passed down artistic practice in fine art painting in general and specific to a defined set of soft rules (a bracketed subset) which we hope will be useful on their own or paired with the next section to researchers looking critically at artistic practice. We will now use these soft rules to discuss known and conjectured ‘cognitive correlates’ of these rules. We will first review the rules collected in the next section.

This section is by no means a complete or systematic description of the painterly artistic process. Nor do we categorize specific craft based knowledge, for example rules of cross-hatching or watercolour techniques. Nor are we implying these ‘soft rules’ are complete or not broken in certain situations. We have noted that many artists cluster a subset of these soft rules into an integrated whole which works for their style and intention within a career direction or specific to one work. Art practice is a complex and intertwined human endeavour and this is just a start at parsing it systematically. Using our poetics methodology pre-collection rules, we biased for more shared and passed down soft rules associated with a cognitive semantic processes, dealing with sensory input, reflective stepwise processes and authored output. Working with tone, shape, colour, line, composition and process bring salient authorship processes into the mix, with centre of focus, lost and found edges, and other tools for narrative emphasis (or filtering out / de-emphasis). When content semantics make the process clearer, we used portrait-based content knowledge. While trying to maintain generality, the centre of these soft rules might be closest to impasto oil painting, where an artist, using brushes, a pre-set out palette of oil colours, directly strokes on a canvas

those painted colours with simple mixing techniques, as opposed to elaborate pre layering methods and other non-direct approaches.

It should also be noted that we have skillfully avoided the fascinating field of cognitive creativity – how and why humans and fine art painters act creatively. While an area of active research by the author, it is beyond the scope of this thesis.

3.3 Cognitive Correlates of Artistic Painter Practice

3.3.1 Introduction and Soft Rules

The artistic passed down painterly knowledge or soft rules contain some degree of rigour and empirical feedback, since it has been moulded by success or failure of artists over generations, of passed down and evolved techniques. As humans, everything artists do has a cognitive correlate. For this thesis, we define ‘cognitive correlates’ as making links between the experiences (of fine art painters and viewers) with behaviour. The correlates point to behavioural properties of brain, mind and cognition. Having presented a framework of process artist knowledge in the previous section, we now present possible cognitive correlates from the artist practice side (last section) and this cognitive side (this section) to inform our design choices in building an NPR software toolkit. However, it is beyond the scope of this thesis to make exact perception and cognitive science assertions about all activities that fine art painters perform. This is still an area of active research and debate. Instead we are interested in creating a ‘testable framework’ that begins to map some of these cognitive correlates. Those that are obvious we build into the parameters (i.e. scriptable building blocks) for our NPR toolkit, those that we are interested in understanding deeply, we begin to more rigorously test with our system (see Chapter 7) and those that need more research, we loosely detail here and leave for future research or for other researchers with the possibility that our NPR toolkit can be used as a tool in their research.

With that caveat, we would like to summarize the bracketed subset of the presented artistic soft rules from the last section that we will as a bridging convention to discuss specific artistic soft rules use though the next two sections. They are labelled by the

subsection they came from (3.2.1 = 1 and so on) for cross referencing. Following is an overview:

[R1.1 Tone first then colour]

- Sample the tone (luminance) first, with that remapped into a colour model.

[R1.2 Regions as likeminded tonal masses]

- A region, shape or blob is an arbitrary area of similar values, large to small.

[R1.3 Unequal value]

- A painting should have a dominant value and unequal subdominant values.

[R2.1 Warm/Cool colour]

- Artists choose colours based on a warm/cool colour wheel. Separating warm and cool globally in large value regions but pairing them in small local areas.

[R2.2 Unequal colour]

- A painting should have a dominant colour and unequal subdominant colours.

[R3.1 Highlights - eye fixations]

- Textural (specular) highlights can influence the viewer's eye fixations.

[R4.1 Relative widen/reduce/harden/soften]

- Compared to the source: edges, regions, ranges can be manipulated by remapping (rescaling) the range, sharpness or detail.

[R5.1 Lost & found edges - eye fixations]

- Edges can be hard, soft or lost, viewers eyes follow hard edges and move into the picture through lost edges.

[R5.2 Centre of focus - eye fixations]

- Extra detail in an area can influence a viewer's eye fixations and interest to it.

[R6.1 Light, volume, content concerns]

- Artists are concerned foremost with light (tone, gradation, shapes) then less so with the object's physical 3D volume, then the content of an object (i.e. a nose).

[R7, 8.1 Overall relativity and rescaling]

- Almost all decisions (tone, colour, hardness, detail) are relative to what is already painted and are scaled off this relative comparison.

3.3.2 Cognitive Correlates Discussion

[R4.1 Relative widen/reduce/harden/soften]

In the *Nature* article entitled “The Artist as Neuroscientist” Cavanagh (2005) wrote that artists in discovering semantic shortcuts, act as research neuroscientists, and that “a great deal can be learned from tracking down their discoveries”. Cavanagh showed how the difference between the real world and the world artists create could reveal, “as much about the brain within us as the artist reveals about the world around us”. Cognitive scientists have observed that paintings often manipulate, distort, exaggerate or violate the physics of shadows, reflections, colours and contours (Cavanagh, 2005; Grossberg, 2008; Mamassian, 2008; Seeley, 2011). Rather than adhering to physical properties of the world, artistic paintings reflect perceptual shortcuts used by the brain. Artists, in experimenting with forms of depiction, discovered what psychologists and neuroscientists are now identifying as principles of perception (Zeki, 2000a). These assertions fit well with our soft rule [R4.1 Relative widen/reduce/harden/soften] that allows painters to greatly change or manipulate the source data of their paintings by hardening or softening selected edges, widening or reducing dynamic range of tonal, colour and other graduations in selected areas as well as reworking or completely leaving out reflections, shadows and other physical properties in the source material (Graham et al., 2010; Pinna, 2008; Redies, 2007).

Mamassian (2008) surveyed work related to these ambiguities and artistic conventions in composition, spatial scale, illumination and colour, spatial layout, shape and movement. His review emphasized the nature of the conventions and ambiguities common in everyday scenes and art perception, and compared the conventions used in visual arts with the prior knowledge used in everyday perception. Visual perception is ambiguous and in their work, artists play with these ambiguities. Researchers note that ambiguities in typical visual

perception were resolved thanks to prior constraints that are often derived from a human's knowledge of the makeup of natural scenes (Enns, 2009; Henderson, 2003; Hochberg, 1968). Ambiguities in visual arts however, while also somewhat resolved by conventions of perceptual priors, more typically are resolved from other sources such as 'artistic stylistic or arbitrary choices' (Mamassian, 2008; Myin, 2000; Zeki, 2000a). That viewers of art work deal with ambiguities differently than they do with natural scenes and default more to artistic stylistic and prior art as well as what cognitive scientists still do not understand (what Mamassian terms 'arbitrary choices') gives cognitive support to some of our passed down soft rules: [R1.1 Tone first then colour] [R2.1 Warm/Cool colour] [R4.1 Relative widen/reduce/harden/soften] [R7, 8.1 Overall relativity and rescaling]. Of course, not all ambiguities in visual perception are resolved. Similarly, not all ambiguities in art paintings are resolved, and artists probably plan or intuit to leave a certain amount of ambiguities in their final work, to let the viewer contribute to the experience in a personal way (Pinna, 2008; Solso, 1996).

Inconsistencies do not seem to take away from the quality of the painting, but allow the artist to emphasize the really important or salient aspects of the painting, which again support the authorship possible from almost all of our soft rules. In other words, "the artist can take shortcuts, presenting cues more economically, and arranging surfaces and lights to suit the message of the piece rather than the requirements of the physical world." (Cavanagh, 2005).

3.3.3 Natural Scenes and Visual Cognition

It is good to step back and again realize that an artist's visual cognition is like that of all humans. For all of us, vision is an active, or a constructive process (Cavanagh, 2011; Enns, 2004; Findlay & Gilchrist, 2003; Zeki, 2000). Studies have now shown us that vision is not just a simpler but still accurate internal representation of the world a human sees in front of them. The task is severely constrained by the fact that the retinal inputs to the visual system are fragmented and underdetermined. Retinal images are two-dimensional records of the relative luminance of light reflected from discrete points in the visual field over time. Under our normal daily situations, this two dimensional array is in constant change due to alterations in lighting conditions, movements of objects and movements of the perceivers.

Yet what we see is not a pattern of constantly changing points of light, but stable, three dimensional objects and scenes. Perception scientists have given evidence to the active nature of vision through many optical illusions studies (Enns, 2009; Hubel, 1988; Purves et al., 2001; Seeley, 2006). Optical illusions of size, distance and motion show us that vision is an active process where salient features and objects in the natural scene are extracted and manipulated for further mental use. Enns explains, “Vision must be far less photographic than subjective experience suggests it to be Vision involves the construction of a model of the external world, not a detailed registration of the images that arise from such objects” (Enns, 2004). Our prior visual experiences play a large role in forming the basis of full mental model. We ‘construct’ the internal model based on the representation of these experiences (Cavanagh, 2011; Enns, 2009; Grossberg, 2008; Mamassian, 2008). This ‘emphasising the salient’ is what Zeki (2001) refers to as the Law of Abstraction, in which the particular is subordinated to the general, so that it can represent many particulars. Researchers posit the capacity to abstract is probably imposed on the brain by the limitations of its memory system, therefore the need to filter out many non-salient details. Art, too, Zeki believes, abstracts and thus externalizes the inner workings of the brain. The artist also forms abstractions, through a process that may share similarities with the same natural scene visual processes but certainly goes beyond them, in that the constructed ‘abstract idea’ itself mutates with the artist’s development. The central principles of the eight principles proposed by Ramachandran and Hirstein (1999) in their cognitive model they called the peak shift effect. The peak shift effect is the exaggeration of salient features that distinguish a given object of interest from other objects. In their model, art is thought to be composed according to the peak shift principle and is reinforced by other neural mechanisms such as perceptual grouping, in order to optimally stimulate particular visual areas of the brain. It should be noted that Zeki and Ramachandran have been critiqued for being too reductionist with their theories on art creation and aesthetics (Hyman, 2010; Ione, 2000a, 2000b; Myin, 2000). However, it is well established that vision for all humans is a highly interactive and constructive process.

A source of active debate is whether artists are able to convey their inner active construction in a painting for their viewers and not just a stylized version of the source scene. However, what inner construct is it? Necessarily, it must be a personal one for the artist but is

it also narratively authored? As we have discussed, artists typically mix three concerns: 1) a somewhat representational record of the real scene or sitter -- in portraiture: resemblance, 2) the inner feeling of the scene or sitter -- in portraiture: personality of the sitter, after interviewing and conversation and 3) lastly, the experiences and trajectory of artistic direction over their careers – all three combine in this authored construct. They are both using the soft rules we have categorized to manipulate the constructed scenes in a planned way as well as working through it live and intuitively as art process -- stepwise through starting, working out issues in process and ultimately completing the piece. Therefore, they are actively manipulating the scene in front of them to tell a mix of authored personal and professional narrative with these tools. They have some extra assistance as well – they ultimately have a viewer doing their own constructive visioning, and a viewer who is more forgiving with processing conventions and ambiguities (Cavanagh, 2005; Grossberg, 2008; Mamassian, 2008). Therefore, artists use these re-authoring tools and abstractions in a substantial way where their new interpretations from the original scene or photograph, in the best cases, add value to the point they become pieces worthy of museums and galleries.

In emphasising this constructed nature that artists can manipulate and build from, we use one of many examples of exaggeration that already occurs in the minds of perceivers of a natural scene. Enns (2009) documents one of the lowest level functions performed by the eye, that of registering an ‘edge’ at a particular location, as an example of this typical exaggerated emphasis that can be scaled or manipulated even further by an artist using these soft rules including [R4.1 Relative widen/reduce/harden/soften] and [R7,8.1 Overall relativity and rescaling]. “The daylight neural receptors in the eye, known as ‘cones,’ do not merely record edges as faithfully as they can. Rather, they are organized to deliberately exaggerate the contrast of edges they encounter, through a system of circuitry known as ‘lateral inhibition’. The net effect of this circuitry is to highlight the existence of edges that are detected, so as to help bridge and join edges in regions of the image in which the lighting conditions make their detection difficult” (Enns, 2009). Artists are able to exploit this constructive system, in general and specifically to this example with soft rule [R5.1 Lost & found edges - eye fixations].

3.3.4 Foveal Vision and Authored Pattern of Eye Gaze

[R3.1 Highlights - eye fixations]

[R5.1 Lost & found edges - eye fixations]

[R5.2 Centre of focus - eye fixations]

Another fundamental feature of human vision is that our experience of a scene or an artwork is not uniformly detailed (Henderson, 2003; Hochberg, 1968; Land, 1999; Melcher & Colby, 2008; Molnar, 1981; Yarbus, 1967). Each eye contains only a small area in which the cone receptors are packed densely enough to provide us with detailed perception -- a region in our visual field about the size of a thumbnail when viewed at arm's length. Thus our viewing experience is actually one that extends over time, including periods of fixation, in which the eye position is almost stationary and visual information is taken in, interrupted by saccades, rapid movements of the eye from one image region to another during which we are also effectively blind (Land, 2006). This makes seeing a highly interactive process, one in which information acquired in a fixation is influencing the content of our mental experience, while at the same time the content of our mind (including its goals and cognitive strategy) is guiding our eyes to new image regions in order to acquire further high resolution information (Enns, 2006; Henderson, 2003; Melcher & Colby, 2008). While bottom-up information from each fixation is influencing our mental experience, our current mental states, including tasks and goals, are guiding saccades in a top-down fashion to new image locations for more information. Those parts of a natural scene or painted image that are more salient generate more neuronal activity, not only in low-level visual areas but also in higher level regions of our visual system, regions concerned with object representation and memory (Collomosse & Hall, 2002; Conway & Livingstone, 2007; DiPaola et al., 2010). Given the constructed nature of vision (previous section), only the most salient stimuli are fully updated or remapped, on average about 3-4 attended items (Prime et al., 2007), the rest can change, disappear or reappear without us noticing the change (Melcher & Colby, 2008; Simons & Rensink, 2005).

As we have gathered in our soft rules, we documented that Sanden (2004) and other modern artists claim that lost and found edges and increased textural detail in paintings guides the gaze of the viewer. The scientific understanding just detailed of foveal vision

gives a foundation for these passed down rules. However, is there direct evidence to support this detail-gaze hypothesis? When we first began to examine the detail-gaze hypothesis from the perspective of vision science, we were surprised to learn that it had not yet been put to a direct test. There have been previous studies examining the gaze patterns of viewers while they were inspecting works of original art (Livingstone, 2002; Molnar, 1981; Yarbus, 1967) but in each case it was difficult to attribute the viewer's gaze patterns directly to the selective emphasis in the painting involving the degree of textural detail. The reasons for this lack of direct evidence are quite straightforward when considered from the perspective of visual fine art creation. Strong correlations in an artwork between the semantic level, the compositional level and the level of textural detail all likely conspire in a synergistic way to guide the gaze of the viewer to selected regions of the painting. However, for the purposes of putting the current claim, that textural variations in themselves guide the viewer's eye, to a proper scientific test, these correlations in original portraits make it impossible to test by monitoring viewers' gaze while they view portraits.

The research on scene viewing suggest that although eye fixations may initially be attracted to regions of high image salience (Itti & Koch, 2001), defined as local regions that are markedly different from their surroundings in colour, orientation, motion and depth, once a scene has been viewed long enough so that its semantic content can be determined, the eye tends to explore those image regions deemed by the participant to be most informative with regard to the task they are performing (Birmingham et al., 2008). Artists have claimed that these choices guide viewer's gaze, but it has been difficult to test because portrait gaze is typically also associated with greater meaning, stronger lighting and a more central location. We therefore specifically investigate the soft rules discussed in this section [R3.1 Highlights, R5.1 Lost & found edges, R5.2 Centre of focus - eye fixations] and use our NPR toolkit to perform four eye tracking-based experiments in which we monitor viewer gaze using paintings where these factors are decoupled from relative detail in Chapter 7.

3.3.5 Tonal Value and Colour

We now look for cognitive correlates for our soft rule subsections on 3.2.1 Tonal and 3.2.2 Colour where we stated, 1) that artists find the tone (luminance) first, then use that tonal

value to remap into one of many 2) colour approaches like opposite temperatures (i.e. warm/cool) and that asymmetry (unequal dominate value/colour) both of 3) tone and 4) colour are important constructs of this process.

[R1.3 Unequal value]

[R2.2 Unequal colour]

Leyton (1992) observed that the mind assigns to any shape a causal history explaining how the shape was formed. By reducing the study of shape to the study of symmetry, Leyton shows that symmetry is crucial to our everyday cognitive processing. Symmetry is the means by which shape is converted into memory. Santella, an NPR based cognitive scientist paraphrasing the work of Leyton, states “Good art should contain as much information in the form of asymmetry as possible to stimulate viewers, but not too much, which will disturb them” (Santella, 2005).

[R1.1 Tone first then colour]

[R1.2 Regions as likeminded tonal masses]

Colour and luminance (often called tone or value by artists) do different things for us in vision. We have three cone types in our visual system and luminance is acquired by summing the three. The processing of colour is split between luminance and chrominance pathways (Livingstone & Hubel, 1987). While natural objects are determined both by luminance and by chrominance, some objects can be defined exclusively from their luminance properties (i.e. different grey levels) and others from their chromatic properties (i.e. hue and saturation). Depth perception is aided by luminance contrast. Depth perception is generally blind to colour; it does not matter for depth perception what colour shadows are for instance, as long as they are an appropriate luminance.

Neurobiologist Margaret Livingstone’s work in visual perception is based mostly on electrode recordings from the primate visual cortex. Livingston demonstrated the building blocks of colour perception are tuned to fire action potentials in response to certain parts of

the spectrum, and are organized in a system of colour opposites, such as orange and blue. These colour perception tuned individual cells fire from the excitation and inhibition of their neighbouring colour-opponent cells.

Livingstone showed via her contributions to the cellular basis of luminosity, that since shape can be derived from luminance differences, artists use contrast to produce shapes, leaving colour for expressive rather than descriptive purposes, although this is a debated and highly complex issue.

Livingstone tried to demonstrate that a large range of artistically salient dynamic effects could originate from how artists rescale spatial frequency information in their paintings. Leonardo's techniques in painting the *Mona Lisa* are among the most familiar examples of this. Livingstone filtered parts of the *Mona Lisa* to produce low, middle and high spatial frequency cropped image samples of the face area. The luminance contours that define *Mona Lisa's* expression were more in the low and middle, but not the high spatial frequencies. Foveal neurons in the retina are generally blind to low and middle spatial frequencies, but peripheral neurons are fine-tuned to this kind of information. Therefore, Livingstone argues that Leonardo used an authored ambiguity of luminance in *Mona Lisa's* smile. The evocative smile is visible only at low spatial frequencies but conveys a more neutral emotion at all spatial frequencies. Livingstone thus proposed that the ambiguity of *Mona Lisa's* smile will be most striking when an observer makes an eye movement looking first at the mouth and then at the eyes. When the mouth is seen in visual periphery, only the low spatial frequency information is available and the interpretation of the emotion is a smile (Livingstone, 2002; Mamassian, 2008; Seeley, 2006). Other researchers have added random luminance noise in the lower face region of the *Mona Lisa* painting to verify that the change of emotion came mostly from the mouth region, and used the classification image technique to localize the relevant information (Kontsevich & Tyler, 2004). They showed that the perception of smiling in the eyes was solely attributable to a configurational effect projecting from the mouth region. While still a complex area of study, this work begins to unravel why artists use luminance (or tone as we have mainly referred to it in our soft rules) as a primary authoring approach with proper luminance either in a specific location or tone mapped as a working

area or shape, then use that luminance value to remap into different authored colour and area soft rules.

[R2.1 Warm/Cool colour]

We have outlined in our artistic process soft rules [R1.1 Tone first then colour, R2.1 Warm/Cool colour] a very well-developed, systematic and solid passed down artistic knowledge system for using warm/cool colour temperature remapping rules in painterly techniques. This system historically is documented extensively including the regular use of the warm/cool colour wheel in art instruction over an extended period of modern art history (as seen in our Figure 8 and 9 in Section 3.2). This passed down notion of excitement and vibrancy that artists talk about when putting warm adjacent to cool in local areas, is very intriguing to us, even though they warn not to overdo it. Of course, the main process they speak of is to separate remapped (from the source) warm and cool deviations of large areas along the light and dark tonal mass divisions. It seems there must be a neural basis for this process that elicits an emotional affect in viewers and appears to be a major pillar technique for artistic painterly practice. We have spent much time researching this area over the years, talking with many experts and while we thought this might be a major candidate for our scientific experiments, we had to conclude there is still much work to be done in this area. We decided it would be better to spend our time on our patterns of eye gaze tracking work and leave this for future research. Colour is a complex issue and not one where artists and cognitive scientists are on the same page. C. L. Hardin might have said it best, “Painters are the experts in colour phenomenology. Their business is to use colour to affect our feelings. Psychophysicists are expert in making experimental inferences from behavioural responses to the functional mechanisms of perception. The varying aims of these two groups of people mean that much that is of interest to the one is of little concern to the other” (Hardin, 1988). However, in recent years, there is some early work by scientists to correlate warm-cool painterly colour practice.

Katra and Wooten of Brown University traced the connection between the warm-cool of temperature and the warm-cool of colour to the corresponding activation levels of their respective neural systems rather than to stereotypical environmental associations such as red

with fire and blue with water. Colour space has an irregular structure. The structure of colour space is arrived at entirely by comparing the colours with each other, so the irregularity of structure is intrinsic to the domain of the colours (Goguen & Myin, 2000; Goguen, 1999). Katra and Wooten concluded, “The remarkable correspondence between the obtained ratings of warmth and coolness and the activation levels in the opponent channels . . . suggests that the attribution of thermal properties to colours may be linked to the low-level physiological processes involved in colour perception.” Experiments with nonhuman primates strongly suggest that this irregular, intrinsic structure is of biological rather than cultural origins (Byrne & Hilbert, 1997).

Some believe there is DNA evidence for warm-cool colour bias where the blue–yellow system in human beings, and probably other trichromatic mammals, is more ancient than the red–green system (Mollon, 1989). C. L. Hardin (1988, 2000) conjectures “that in general animal lines with colour vision had ancestors with just one chromatic pair in an opponent configuration. Suppose that in the most primitive ancestors the two poles of this chromatic visual system were hard-wired to behavioural mechanisms, one activating the animal, one inhibiting it. As animals became more complicated, the connection between wavelength and behaviour became less immediate and direct, and lingers in us as diffuse affect. When a painting speaks to us, we perhaps discern the faint voices of our forebears” (Hardin, 2000).

Of course, he goes on to say, “Right now, the study of the affective character of colour is rich in folklore and unsupported assertion, but very poor in solid scientific knowledge. . . . These problems are hard and, very likely, ill-posed. Making progress on them will require unusual people with an understanding of visual science, cognitive science and art — perhaps someone reading these words.” It is of course our hope that by building up our NPR toolkit with simple correlated colour space parameters as scripted building blocks for experiments that we can support this effort.

3.3.6 Global and Relative Analysis: Tonal Masses, Volume and Content

[R1.2 Regions as likeminded tonal masses]

[R6.1 Light, volume, content concerns]

[R7, 8.1 Overall relativity and rescaling]

While much is known about the cognitive processing of contours, edges and boundaries, there is still active research in the area of shape recognition from a tonal or luminance grouping point of view. Can a human consciously or unconsciously tone sample their world and clump together similar tonal values into a perceived shape and can this be done at any progressively refined level starting with the classic 3 tones – light, halftone and shadow? What is known is the maximum dynamic range of luminances for paintings is only a fraction of what is possible for natural scenes. Therefore artists employ many different forms of a nonlinear luminance compression schemes in order to depict the world on canvas (Conway & Livingstone, 2007; Graham et al., 2010; Hochberg, 1980; Livingstone, 2002; Redies, 2007).

Graham and others (2010) using statistical techniques showed it is possible to find the tone mapping or artist's look-up table (ALUT), i.e., the histogram resulting transform, to describe an artist's use of luminance scaling. Nonlinear scaling is required to compress the significant range of tonal values present in natural scenes into the far smaller range available for paintings. Several researchers including Graham and Redies have shown that works of art and natural scenes display a similar degree of scale invariance and fractal-like properties. In contrast, some studies have shown other sets of images had little or no aesthetic appeal (i.e. photographs of household objects, plants, scientific illustrations even faces) and were found to display significantly different Fourier power spectra, which indicated that most of these images were not scale-invariant. These results begin to show that artists create art with properties that are not necessarily the same as those of the subject depicted. Rather, artists in scientific terms 'adjust the image statistics' of their subjects during the process of constructed creation so that, in their works of art, statistics are closer to those encountered in complex natural scenes (Graham et al., 2010; Redies, 2007). Tonal mapping seems part of this process.

Redies (2007) believes a work of art is created through constant feedback with the artist's visual system so that a specific state of neural activity is induced. As a result, art objects reflect not only the higher-order statistical properties of natural scenes, to which the visual system has adapted from an evolutionary standpoint but also intrinsic properties of our early visual system. Redies has hypothesized that this additional reflection of intrinsic properties distinguishes art objects from natural scenes. Implicit to this hypothesis, the visual system

can resonate more profoundly to art objects than to natural scenes. In other words, art can be tuned more precisely to nervous system properties than natural scenes are (Pinna, 2008; Solso, 1996, 2000). The process of tonal mapping [R1.2 Regions as likeminded tonal masses] as well as other tonal, colour, range rescaling [R7, 8.1 Overall relativity and rescaling] are part of the process.

3.3.7 Summary

Using a correlated subset of our categorized artist knowledge soft rules, we have attempted to show how artistic passed down knowledge can align with current research in cognitive and perception science and how both can support each other. This correlation is still in its infancy with many researchers from different related fields converging on emerging topics. We hope our systematic artistic painterly practice categorization techniques can support current and future research in this area.

3.4 NPR Toolkit Parameters

3.4.1 Introduction

Figure 15 From a black and white source (inset), our NPR toolkit created this painterly output using a hierarchical blob tree structure and semantic regions. The colour palette was remapped from a fine art master's portrait where tonal values map into the semantic regions of the source colour painting.

From the findings from the previous sections, we design and implement a novel painterly NPR software system using a parameterized computer modelling approach that was informed

by cognitive and painterly knowledge (Figure 15). The scriptable and extensible Java/OpenGL based open source NPR toolkit is intended to function as both an instrument for scholarly research as well as a working painterly NPR system

One overall prerequisite to better implement in software many of our soft rules including [R1.3 Unequal value, R2.2 Unequal colour, R4.1 Relative widen/reduce/harden/soften and R7,8.1 Overall relativity and rescaling], is the need for a system that can evaluate the source material at any state of the output canvas both on a global or any local level. Many NPR researchers (Collomosse & Hall, 2002; Gooch & Gooch, 2001; Hays & Essa, 2004; Santella & DeCarlo, 2004) have warned about the limitations of many NPR systems that only work with local comparisons and calculations. They call for systems that can perform global comparison and analysis, and have used different methods to try to deal with this problem in a speedy and efficient manner. Our global awareness approach attempts to best match its process with artistic practice aligned with cognitive science knowledge.

We achieve a scriptable global/regional awareness using hierarchal perception blob trees where a current child region can communicate up the tree to the parent or higher to receive information about the current working state and use it in its local stroke decisions. We look at a blob as a perceptual region of the source that a human artist can take in, think about and act on – and that artists do this at different progressive levels starting with large gross tones and eventually smaller details regions. ‘Concerns’ like unequal colour or not overdoing highlights are examples of this communication. When our BlobThinker component (detailed in our implementation chapter) of the toolkit is instantiated, it generates a hierarchical blob structure from the reference image. To do this, the ‘root’ blob is first created, which encompasses the entire canvas; if the blobs-from-regions parameter is true, this root blob is then divided into optional created semantic regions (i.e. any labelled region, such as hair, background, clothes), see Figure 15, middle image. Essentially, the blob algorithm subdivides parent blobs into child blobs by locating contiguous groups of pixels whose lightness values (i.e. the ‘J’ in our JCh colour space system) fall above the parent’s mean lightness as well as groups whose mean lightness falls below it.

Thus, in principle, blobs represent different value regions within the reference image that may be perceived grossly (towards the root of the blob structure) or finely (towards the leaf

blobs). The algorithm proceeds iteratively, dividing each leaf blob into yet more leaf blobs until no leaf blob can be further subdivided. The system has many parameters in the scripted language set up for customization of different painter systems. Generally it implements a hierarchical progressive system for our soft rule [R1.2 Likeminded tonal mapping], meaning the system does not work on a stroke level but on a blob level at different levels of detail as supported by the current needs of the artistic process. The process allows the system to be concerned with both content (i.e. you are in the body shadow area of the skin) or light (tonal area) as specified soft rule [R6.1 Light, volume, content concerns] and gives us a mechanism to performed informed calculations for:

- [R1.1 Tone first then colour]
- [R1.3 Unequal value]
- [R2.1 Warm/Cool colour]
- [R2.2 Unequal colour]
- [R4.1 Relative widen/reduce/harden/soften]
- [R5.1 Lost & found edges - eye fixations]
- [R5.2 Centre of focus - eye fixations]

In Figure 16, we see a visual representation of the hierarchal blob system, where the lower middle images showing how the semantic mattes (2D regions) are specified and labelled as Photoshop layers. There can be any number of layers and labels with regions. Our scripting system reads in the info from a PSD Photoshop file and allows calls to any label for any reason (see scripting examples in Chapter 6). These labels and regions can be generated automatically from a computer graphics 3D source (see image examples in Chapter 6). In Figure 16 - left, the first two levels of blob tree are shown, the coloured areas show the top level with the optional semantic regions: 1. Background, 2. Clothes, 3. Hair, 4. SkinLight, 5. SkinShadow. The smaller regions within these colour areas illustrate the next level in the tree of calculated luminance regions, those within the Background region would be lower down the tree as 1.1, 1.2 and so on. Figure 16 - right illustrates the next child level of regions

calculated within the tree structure with the same region as 1.1.1, 1.1.2 and so on. Subdivided regions or blobs continue fully down the tree, allowing for communication up and down the tree structure. Blobs are then used within the logic of our system for different analysis (detailed in Chapter 5 and 6).

Figure 16 Left: Blob Tree with false colour semantic top regions (optional) and next level luminance regions; Middle: the pre-specified semantic regions and labels; Right: the third level of tree luminance regions. All exist in a tree structure

We now provide a brief overview our NPR design implementation as informed by the soft rules presented in this chapter. Chapter 6 and 7 have additional details about the NPR toolkit and its design and function.

[R1.1 Tone first then colour]

Based on this rule, we implemented a precise tone JCh CIELAB colour algorithm in our phase two NPR toolkit, giving us the ability to perform sophisticated remapping and rescaling in both tonal conversion and colour space. While we have many high level parameter techniques (called ‘Concerns’ modules) for the colour palette a user can script in our toolkit, our main palette concern uses the following automated procedure: at each stroke

time, sample the source tone then based on optional semantic regionaling, either based on content (i.e. in clothes or background) or region (in a shadow or halftone area) the system remaps that tonal value (several rescaling options) via smart palettes to a constrained and specific colour choices.

[R1.2 Likeminded tonal mapping]

We use a hierarchal perceptual blob tree to facilitate where a blob is a similar tonal mapping (i.e. perceived shape at some cognitive level). Perceptual blobs use luminance sampling but optionally can use semantics blobs as artists do to both start and work within. For region and global communication this tree structure of more progressively refined (i.e. smaller and smaller) blob regions can communicate from the current child to parent blob relieving a stated problem in NPR (too local).

[R1.3 Unequal value]

At stroke time, a ‘Concern’ can be written to communicate with parent blobs up the tree to evaluate if a tonal value has been overused or in any tonal decision (must be scripted as an equation with our language tools). There are no current high level constructs for global value balance as this is a source of on-going research but scripting tools are available to setup up a structure.

[R2.1 Warm/Cool colour]

The semantic region facility can be set up to redirect the tonal sample into a constrained colour space – say a warm colour palette from a light body tone region at the top of the blob tree. This is done with smart palettes and semantic region labelling. With this facility, it is also possible to constrain regions on the output image to a colour palette in the same regions (the hair) in one of many historical painting templates. This is how we implemented colour resemblance for our Rembrandt NPR paintings in our studies in Chapter 7. In this way, the system is still tone sample from the source image but remap that tonal/colour conversion to

the same semantic area in a Rembrandt or Hopper painting for example. (e.g. Figures 29, 30 and 37).

[R2.2 Unequal colour]

At stroke time, a ‘Concern’ can be written to communicate with parent blobs up the tree if a colour has been overused (have I used too much strong red?) or in any colour balance decision (must be scripted as an equation with our language tools). This is still an area of active research, eventually requiring an Artificial Intelligence approach as colour balance is complex. In Chapter 6, we document how this might be done with Concerns.

[R3.1 Highlights - eye fixations]

Using our semantic region labelling and pass facility, it is possible to affect different qualities of textural highlights, by using specific palettes and stroke line parameters and density of a specific constrained textural highlight pass.

[R4.1 Relative widen/reduce/harden/soften]

Our scripting language and Concerns have a full array of rescaling and remapping options including scriptable splining facility, including a splining system for rescaling two ranges (say tone into colour palette). This is most currently set up for tonal range, tonal to colour remapping and colour range rescaling as well as rescaling blob region shaping. Creating soft edges is a more sophisticated process and must be combined with labels and special passes.

[R5.1 Lost & found edges - eye fixations]

Creating soft edges or hard edges, or rules for losing or softening edges with some global schemes is still a matter for research. However as stated above there are scripting tools to soften or lose edges in a specific region by creating a semantic region in which, on a special

pass, only that region can have specific stroke/edge parameters performed to create different edge qualities.

[R5.2 Centre of focus - eye fixations]

Using the described labelled regions of maps, we can perform any level of extra detail via special passes. This can be automated if there is a 3D CGI source. We briefly documented this automated process in our research with our 3D industry partner in Chapter 6.

[R6.1 Light, volume, content concerns]

Our concerns allow any combination of analysis of light via tonal blobs and/or content via region label and calls, as described above. Artists also use volume; say to brush in the direction of the slope of the nose. Our stroke direction via gradation system, can be set to any source raster buffer of information, typically on the luminance gradient (Hertzmann), we allow for stroke direction to be calculated by a depth buffer (from a 3D source) and have scripted tools (Concerns) to use both depth or luminance in sophisticated combinations as an artist would. Therefore, the painterly calculations can be informed by light, volume and content in logical art-like processes.

[R7, 8.1 Overall relativity and rescaling]

We have documented above that at any point in the process, current global or any level of local regions can communicate with the main ‘Concerns’ system, so a relative decision can be made. See Chapter 5 for details.

3.4.2 NPR Toolkit Summary

This section gave a brief overview of how the painterly soft rules that we have systematically categorized and aligned with cognitive and perception science have informed our NPR toolkit design and implementation. Full details of our NPR toolkit techniques and

process are given in Chapters 5 and 6 including visual examples of these processes in practice.

3.5 Conclusion

We have presented three sections describing passed down artist painterly practice from a 1) systematized poetics of art, 2) cognitive correlates of art practice and 3) informed NPR toolkit design perspective. We used a bracketed soft rule convention to tie art soft rule concepts or painterly heuristics through all three fields. It is hoped that our work in Sections 1 and 2 (3.2 and 3.3), that is, beginning to systematically quantify artistic painterly practice and align it with science-based cognitive correlates, is of value on its own right to a wide group of art, philosophical and science scholars. While this area, as presented in this chapter, is still emerging and speculative, the thesis specifically will detail some of these soft rules involved in eye gaze pattern agency in more rigorous science-based detail with four eye-tracking studies using our NPR toolkit (Chapter 7). Before proceeding to those chapters, however, we will first present the art historical perspective both showing how this work can inform art history scholars as well as forming the hypotheses that drove our studies in later chapters.

4: Rembrandt and Textural Agency – Art Critical Theory

4.1 Introduction

Given the wide systematic approach of the last chapter to artist practice, this chapter documents our work at a more depth specific view as well as bringing in the art historical – art critical perspective (DiPaola, 2008). It looks specifically at evidence we encountered when tracing the source of ‘lost and found edges’ and ‘centre of interest’ techniques from the last chapter. While only fully detailed in modern art books, where and when did they originate historically? Was there a time in art history, when these techniques to guide the eye of the viewer were matured or mastered by one or a group of artists and at that point subsequently passed down or communicated through the more typical art salon mentoring evolution from artists to artist to present day where they are better documented in art technique books? In our poetics research, there seems to be evidence showing their more professionalized use by late Rembrandt. Our research into the art historical perspective specific to late Rembrandt, overviews the basic vision science research from an historical perspective and makes a hypothesis using both vision science and art history critical theory.

The work of this chapter, mainly from our paper on the subject (DiPaola, 2008), was written early in the thesis process; the hypothesis made was more from an art criticism level. We were then interested in proving the hypothesis and with guidance from a supervisory committee member, were convinced to use our NPR toolkit as the basis to perform four vision science / eye tracking studies. Which both gave evidence to the community in this field and as tested our NPR toolkit as a substantial tool for scholars in both art history and cognitive science. From an historical, introductory perspective, it seems correct to present this work in its original earlier published paper, when the hypothesis was first discussed. However, it should be noted that some of the vision science discussion is quite naive at this point. A second more updated work was published in the MIT journal *Leonardo* (DiPaola et al., 2010), with a more sophisticated view of vision science and art – content from which is

details in the later chapters of this work. Issues aside, it is felt that this first published work, better and honestly introduces the texture agency Rembrandt theory, as a non-biased question to be answered.

4.2 Rembrandt and Textural Agency

This interdisciplinary work hypothesizes that Rembrandt, reacting to his Italian contemporaries, developed specific painterly techniques, typically not associated with the early modern period that engaged the viewer and directed their gaze. Though these methods were not based on scientific evidence at the time, it can be argued that they are based on a correct understanding of visual perception. Using scientific and critical sources, this work attempts to support and extend art history theories that artists in the late ‘early modern’ (i.e. Renaissance) period developed painterly techniques associated with optics and texture, alongside the more established perspective construction, to guide and influence the observer’s perception of their work. While discussed in general terms, the work will focus on faces and portraiture and will analyse Rembrandt’s late portraits to support the general thesis.

It has been speculated that fine art painters have evolved an open methodology, which exploits specific human vision and perception techniques that has only recently been validated through modern cognitive and biological scientific methods. This work attempts to explore these issues both from a modern science perspective as well as an art history perspective and will shed early light on how these techniques developed and how they may support critical theory in terms of intuitive versus mechanized views of the early modern period and Rembrandt’s place within it. It is our contention that artists such as Rembrandt were also using a more organic vision approach to communicate with the observer beside perspective and optical lighting effects.

4.2.1 Early Modern Painting: Monolithic Mechanistic View

A commonly held belief about the Renaissance is that science and art were very much intermingled. This was the result of the introduction of the scientific method, which provided a new process for discovery and placed an emphasis on empirical evidence and the importance of mathematics. In Renaissance art, the crowning achievement of this new

intermingling was the development of highly realistic linear perspective. The development of perspective can be seen as one part of a much broader trend in both the sciences and the arts towards realism. Harry Berger Jr., in his 1998 essay (Berger Jr, 1998), describes what he and others see as a monolithic view in academic discourse on the early modern period. Martin Jay termed this view as ‘Cartesian Perspectivalism’ (Foster, 1988) in which the eye is treated as an abstract Cartesian point, or a pinhole camera, which forms the visual image and presents it to the brain for interpretation, rather than as a complex organ that interacts dynamically with brain. Alberti described in his classic *On Painting*, that this trend “hailed mathematics over physics or physiology” (Alberti, 1966) in its recasting of vision. This prompts us to wonder whether there might be other, more organic, forms of painterly intention that have been buried by the static understanding of perspective in which this period has been viewed. To us it also raises the paradoxical possibility that although scientific thinking may have contributed to this mechanized view of vision during the early modern period; science may now be called upon once again to relocate the eye back into a more properly dynamic understanding of vision.

4.2.2 Systems of Early Modern Painting - Painting Modes

To help provide a framework for our historical and scientific discussion, we turn to Berger’s *Systems of Early Modern Painting*, in which he describes four painting modes: decorative, graphic, optical and textural. The decorative mode uses pigment, colour, light and techniques to give a sense of beauty and, even more importantly, to honour the painting and the subject. In the graphic mode, subjects are painted as they are known or thought to be, i.e., as people imagine they really are and appear. Lifelike, naturalistic imitation of 3D forms in space and spatial relations are significant in the graphic mode as well as the visualization of knowledge, for instance incorporating the knowledge from anatomical studies. Berger posits that the transition from decorative mode to graphic came as patronage changed from religious clientele to that of the merchant class.

In the optical mode, things are painted as they are seen with an emphasis on the conditions of visibility that affect, alter or in some cases interfere with the graphic mode. Looking deeper into the optical mode, we see that this mode offers the observer a somewhat more active interpretive role than the graphic mode. The optical mode brings about a shift from

‘objective’ to a ‘subjective’ set of cues. Finally, the textural mode is about ‘the trace,’ the work of the brush in ‘real-time’ and as an extension of the painter’s own body. The textural mode can therefore be read as the painter’s interpretive act, which calls for an interpretive response from the viewer. Texture generates conflicting modes of observership and can be seen as a window to the graphic mode. Berger claims that the textural mode, “obscures where the graphic clarifies [and] softens where it hardens.”

In this work, we wish to extend Berger’s view of the textural mode to include a level of agency. Our primary proposal is that the textural mode can be used to enrich, invite and move the viewer’s gaze via the artist’s intention, so that the oft-cited direct connection between the artist’s trace and the observer’s perception can be understood as more than just a metaphor. Stated in its strongest form, our hypothesis is that the creative act of painting begins with the artist’s hands, and that these hands leave a trace that can be used to hint at, to guide and to sometimes even coerce the viewer’s gaze through the act of completing the painting as a mental experience. As a case study for our exploration of this hypothesis, we will restrict ourselves to Rembrandt portraits, in particular those done late in his life.

Figure 17 J.H. Sanden, focus and eye fixation paths to sharp eye. (Sanden, 2004)

In the case of portraiture, it could be said that this artistic methodology attempts to simplify, compose and leave out what is irrelevant, emphasizing what's important in their subject. While seemingly a qualitative task, artists have used known techniques such as relying on source tone over colour to remap into a colour temperature model, using 'sharpness' to create a centre of interest, using edges to move the viewer's gaze, and other techniques to filter and emphasize their varied goals. One specific technique, associated with 20th century painting, is to use painterly brushwork on the textural plane to direct and coerce the viewer's eye gaze through a painting, thereby influencing the observer's eye gaze movement paths and fixation points within the work. Well known portraitist John Howard Sanden [5] describes his technique of using centre of focus sharpening techniques to move the viewer's eye gaze to the sitter's left eye and eyebrow to emphasize an intrinsic element of the sitter (Figure 17). Sanden also creates a textural spiralling gaze path into this accentuated eye area by using loose directional brush strokes under the eye which follow the asymmetrically curved mouth edge. Exploiting the tendency for eye gaze to follow a line or edge, artist Harley Brown demonstrates in Figure 18, the 'lost and found edges' technique where an edge once followed by the observer's eye can be lost to move the eye into the image (Sanden & Sanden, 1999). While there are several examples in Figure 18, the most blatant is the lost elbows, which lead you into the downward-looking face (Brown & Lehrman, 2004). How these textural agency techniques work in a cognitive sense is still an area of research, since they work at the very edge of consciousness or at a strongly intuitive level – both from an artist and observer point of view.

Figure 18 H. Brown, ‘Lost & Found Edges’ – contrast based edges intact, then lost at elbows to drive you back to face. (Brown & Lehrman, 2004).

Few artists embraced this use of texture during the early modern period like Rembrandt. In his book, *Fictions of the Pose: Rembrandt against the Italian Renaissance*, Berger (1994) uses Rembrandt as a vehicle to demythologize a more standard view of the intellectual and cultural history of the early modern period. One of Berger’s main points concerns what he thinks ‘the Rembrandt look “sees”’. What Rembrandt sees, challenges, critiques and at times parodies, is “the embarrassment of the Renaissance riches” of the Italian model. Berger asserts, expanding on Kenneth Clark (1978), that “Rembrandt transformed his style by the study of Italian Renaissance art.” This dialogue, one-directional as it may seem, between Rembrandt and Italy gave way to both a classical and anti-classical style in Rembrandt’s art and, more importantly, shifted the discourse in the other direction, reading Italy through Rembrandt’s eyes. Berger shows us the process by which painters can “imitate, emulate, appropriate, and sublate” prior art and, more specifically, Rembrandt’s performative mastery of this process. Berger refers to this as “revisionary allusion” which allows Rembrandt’s

implied reconstructions to “creatively distort the past” in order to make it reflect the present and “more immediate[ly] focus on critique.” Rembrandt also used a similar approach to critique his interpretation of patronage as well as “the Italianism residing in the Dutch scene of patron / painter relations”.

The adoption of oil painting as a medium was controversial when it first appeared. Titian quickly recognized its merits, and added several innovations to early oil painting techniques. While the control of oil with stiff brushes allowed for quite hard edges, Titian apparently found the softer edge techniques more to his liking, and used them extensively, as they gave the effect of being slightly out of focus. The fact that the mature Rembrandt was deeply influenced by 16th-century Venetian painting - especially by Titian and Giorgione – wasn't news when Kenneth Clark (1966) published his classic study on Rembrandt and the Italian Renaissance. Rembrandt never went to Italy. But Italy came to the prosperous art houses of Amsterdam in the form of what Clark described as ‘almost unbelievable quantities’ of paintings, drawings, etchings and prints. Historians note that Rembrandt's own voracious art collecting clearly contributed to his bankruptcy in the 1650s. When an inventory was taken of Rembrandt's possessions, an album was devoted almost entirely to Titian's work. As Vasari noted, Titian's late works were “carried out in bold strokes, broadly applied in great patches in such a manner that they cannot be looked at closely but from a distance appear perfect”.

Noted art historian Rosand (1990) described the appeal of these late paintings as “tactile as well as visual, inviting us to touch as well as to look.” A 17th-century Venetian critic Marco Boschini, who claims to report the words of a late-16th-century disciple of Titian, writes that toward the end of Titian painting process, he painted more with his fingers than with the brush, comparing himself to God, who formed the human body, created it out of earth with his hands (Rosand, 1981). Titian never taught his assistants, but, as Vasari reports, each disciple took whatever he could from the master's example. The same can be said for Titian's most significant followers, like Rembrandt. Rosand describes how “Titian's facture, for all the revelation of his textured surfaces, continues to remain just beyond the reach of comprehension. In comparison, Rembrandt's technique seems quite straightforward, much more accessible to direct visual analysis.” We know that Rembrandt revered and emulated late Titian. In typical Rembrandt style, he appropriated and extended Titian's process of

textural and bodily touch, with the result that his touch, while still gestural, embodied a deep communication link with the eye gaze of his viewers.

In further support of this interpretation, Virgil Elliott (2007), an artist and writer on Rembrandt states of Rembrandt's technical innovations, states that "the highly refined imagery of his younger days gradually gave way to a rougher, more painterly finish in his middle and later years, perhaps due to changes in his eyesight". But rather than attributing these changes to a visual impairment — a diagnosis consistent with a mechanical understanding of vision but not of the dynamic interplay between eye and brain — we propose that a mature Rembrandt was simply continuing his lifelong critique and discourse with the Italian Renaissance. Specifically, he was performatively inventing another, more painterly, technique, one 'more cognizant of his discourse with his viewers', more personal and direct than what Berger has called graphic and optical modes.

4.3 The Science of Human Vision and Eye Gaze

What does modern cognitive theory correlate to Berger's discussion of the four modes of Early Modern Painting? We see the world as light enters our eye, passing information through to our visual cortex before being assessed in our associative cortex. The visual cortex processes low level visual constructs such as lines, edges, contours and shapes. These low level constructs are built up in our associative cortex into high level constructs such as patterns, objects, scenes and high level knowledge interpretation. In the human eye, light is sensed through two light receptors – rods and cones. Humans see the world using rods and cones using a central vision strategy.

Figure 19 (a) The fovea's size and placement in the eye. (b) The middle spike and fast fall off show the fovea's highly packed density of cones compared to the rest of the eye. From (Lindsay & Norman, 1977).

Opposite the pupil is the fovea (Figure 19a), which is a small area more densely packed with cones than the rest of the eye (Figure 19b). The fovea is the size of a pinhead yet it is our window to the outside world. Humans have one fovea paired with very fast eye muscles, to concentrate most of their visual acuity at one central point, giving us maximum focus when we need it but also mandating that we sample our world through eye movement. The fovea accounts for only 2 degrees of a human's 180 degree vision, roughly equivalent to the viewing area of the nail of a thumb on an outstretched arm. The fovea is surrounded by the next area of vision called the Parafovea, which is 10 times less acute than the central fovea. The near peripheral, at 60° of vision and peripheral at 180° of vision, are each 10 times less acute than their inner neighbour.

4.3.1 Dynamic Vision

A fundamental feature of human vision is that our experience of a scene or an artwork is not uniformly detailed (Hochberg, 1968). Each eye contains only a small area in which the cone receptors are packed densely enough to provide us with detailed colour and shape perception (a region in our visual field about the size of a thumbnail when viewed at arm's length). Thus our viewing experience is actually one that extends over time, including periods of fixation, in which the eye position is almost stationary and visual information is taken in, interrupted by saccades, rapid movements of the eye from one image region to

another during which we are also effectively blind (Land, 2006). This makes seeing a highly interactive process, one in which information acquired in a fixation is influencing the content of our mental experience, while at the same time the content of our mind (including its goals and cognitive strategy) is guiding our eyes to new image regions in order to acquire further high resolution information (Enns, 2006; Henderson, 2003; Melcher & Colby, 2008).

At differing levels, many modern artists, as we presented with Sanden and Brown, are aware of this understanding of human vision and deliberately seek to incorporate and exploit it in their work. Directing the viewer's gaze to selected regions in a portrait is one of the tools a modern artist has for emphasizing certain character traits of the sitter and for giving viewers a glimpse into the collaboration between sitter and painter in the development of a portrait (Berger Jr, 1998; Miall & Tchalenko, 2001; Nicholls et al., 1999). In addition to being consistent with the modern understanding of mind-eye dynamics, artists' selective application of detail is also consistent with the specialized neural pathways of human vision: coarse brushwork corresponds to low spatial frequency information that is transmitted very rapidly to many regions of the visual system to help orient the eyes to points of possible interest, whereas fine brushwork corresponds to higher spatial frequency channels that transmit more slowly to the centres involved in detailed and prolonged inspection (Livingstone, 2002; Vidyasagar, 1999).

4.3.2 Eye Movement Research in Art

Scientific research towards analysis of art, what we define in this thesis as the 'cognitive correlates' of fine art painting, has only begun in the last 40 years. Based on the current understanding of visual perception, we do not see a painting all at once but form an impression from a large number of eye fixations looking at small details.

Figure 20 Yarbis (1967) Eye movement: Images of a: face of young girl and b: bust of Queen Nefertiti along with the eye scans, both movements (lines) and fixations (points) (Noton & Stark, 1971).

Eye movement research is still in its early stages since the processes at work are very complex and measuring devices for eye movement are still fraught with issues of precision, intrusiveness and complicated data analysis. It was even more so in Russia in 1967 when Yarbis (1967) using primitive equipment performed some of the first eye movement research in the visual arts. He assembled a study where a test group of subjects viewed, for three minutes, images of a bust of Queen Nefertiti and the face of a young girl. With Queen Nefertiti (Figure 20 right), the eye scan gaze paths show that a viewer's eye gaze moves into the image, usually coming from the left, then stops and is diverted by the strong edge contour of the object, following the edge around the bust. Eye gaze moves into the bust at specific points such as through the eyes or following the horizontal edge behind the ear. Most of the gaze follows the contours at the face side of the bust. In the face of the young girl (Figure 20 left), eye movements are so universal that we can recognize the face just by looking at the scan of the eye traces. Eye movements appear to be a window into the viewer's thoughts and tend to follow edges and contours but also centre on high level constructs like eyes. This was made even clearer by Yarbus' later studies using Repin's oil painting titled 'Unexpected Visitor.' Subjects were asked to view the painting for several minutes, but this time were given seven different goals while they viewed the image. In general, what was apparent was that the viewer's gaze flowed to the faces and was affected by edges and contours. This

pioneering work confirmed through scientific methods what artists have appeared to intuit and exploit in their artwork through generations.

Figure 21 a) Yarbhis (1967), Repin study, 7 targeted scans. b) Molnar (1997) Rembrandt, 2 scans.

Over a decade later, Molnar (1997) separated art students into two groups as they viewed ‘The Anatomy Lesson’ by Rembrandt. One group was told to prepare for questions on (Figure 21b top) aesthetic qualities of the painting, the other was told to prepare for (Figure 21b bottom) semantic meaning. Molnar was surprised by the similarity of the scans of these different groups. Molnar reasoned that the similarity might be explained by the fact that both groups were experienced with artistic conventions like background and style. However, the general hypothesis here is to question if Molnar might have underestimated Rembrandt’s own techniques within the painting to affect eye gaze. These cues by a master painter might have influenced the two groups over the simple initial questions asked by Molnar.

Two points need to be addressed before we bring this section to a close. The first is that we are implying no direct link between the science of central vision, eye movement and art reception with late Rembrandt and other artists in the early modern period. We are simply presenting an overview of how science can bring some clarity and food for thought to this uniquely intuitive and human process of how we view our world and art. The second point worth noting is that these very early studies in eye movement and art, together with a handful of others, are the precious few we have on this illuminating area. Even with these few examples, a model unfolds of how humans put a scene together and the influence of low-level constructs like edges and contours together with high-level constructs such as faces, people, objects and light, which can deeply affect how we move through a scene and build up an impression of an artwork. It is our contention, speculative as it may be that even without the knowledge of the science; artists over thousands of years have exploited many vision, perception and cognitive traits of humans, of which science has only recently begun to verify.

4.4 Rembrandt, Eye Movement and Intent

With Berger's vocabulary of painting modes, a review of the recent scientific understanding of human central focused vision and eye movement theory as well as noting early research into eye movement and observing art, it is now possible to examine how artists like Rembrandt might have exploited an intuitive understanding of visual perception to focus the viewer's gaze within their paintings. Virgil Elliott (2007) an artist and writer on the technical innovations of Rembrandt states of Rembrandt's work that "the highly refined imagery of his younger days gradually gave way to a rougher, more painterly finish in his middle and later years, perhaps due to changes in his eyesight (Figure 23). Could it be instead that a mature Rembrandt continued his lifelong critique and discourse with the Italian Renaissance by performatively inventing another, more painterly, technique, one 'more cognizant of his discourse with his viewers', more personal and direct than what Berger has called graphic and optical modes?

The eye-brain loop is significant, intuitive and not very well understood, but it reveals how in real-time, humans can build up a view, i.e. a personal model, of the world in front of us. Our eye-brain loop is constantly re-evaluating our next move in the scene before us –

building from nothing, to an impression, to an understandable scene, to our reaction or intention. This is as true for a painter as they are in the act of creating, and evaluating their next painterly move, as it is for an observer of a painting trying to take in all they see. Besides more mechanical, planned out artistic methods, both individually and historically, like the painting's narrative, its composition, the use of perspective, as well as his strong use of planned lighting effects, did Rembrandt also use his gesture, his hand and his intuition to guide and persuade the viewer through the image via a more painterly textural level. The textural trace can lie on its own plane but deeply affects the observer's recognition, his personal experience of the graphic and optimal plane of an image – it can set up deep conversation between the artist and the observer. Did Rembrandt's reaction to the Italian Model and more specifically Titian, herald a new textural agency between artist and observer which would be exploited by artists in the following modern era, returning a more human and organic eye back to the gifts of the early modern era? There is much circumstantial evidence to support this claim. Titian innovated several oil painting techniques. Rembrandt was significantly influenced by Titian and even owned some of his work.

Figure 22 Extending on Durand's bottom up cognitive flow, low level constructs (on the textural plane) can affect the viewer's reception of high level semantics in a painting.

Figure 23 Early self-portrait, 1629 versus a late self-portrait (right), 1661 (close-up). Note increased textural level details (e.g. graduation, edges, and focus differences).

Scanning Rembrandt's late portraits, modern textural intent and agency techniques that can affect eye movements such as lost and found edges, centre of interest focusing using either tonal, brushwork or contrasts sharpening and eye fixation pathing techniques via brush direction, highlight edges and blobs they can clearly be noted (Figures 24 and 25). This appears to be less true in Titian's late work based on inspection and critical commentary; however, more research that is rigorous is needed. It should also be noted, as seen in the results of Yarbiss and Molnar, that high level cognitive constructs play an equal or greater role in viewer interest and eye movement, although this area is well documented. This work specifically attempts to extend the scholarly discussion to low-level textural constructs like edges, contours and sharpness and how they are mediated in real-time by the painter's trace. We have also limited our analysis to portraiture. Analysis of this kind is fraught with uncertainties as both high and low cognitive areas work in concert with human volition and individuality deeply intertwined in the mix. Extending the work of Durand (2002) on his general science of depiction, we show how the textural plane gives artists another direction to

affects the high level cortex via low level constructs with the uses of more painterly and gestural processes (Figure 23).

Using Berger's foundational discourse of optical, and textural painting modes from this essay, along with once again relying on scientific thought, albeit in a new modern form, it seems possible to expand on Rembrandt's appropriation and redefinition of Italian Renaissance ideals to relocate the eye back to its organic origins and, in so doing, demonstrate that other, more mannered forms of artistic agency besides perspective had their origins in this period. Human centralized vision creates our view of the world and with it, our view of art, based on thousands of rapid eye fixations of small visual areas that build up via an elaborate eye-brain recursive loop, which then builds an internal model of a cohesive image or scene. Is it possible that artists such as Rembrandt, using optical and textural techniques outlined by Berger and others, can direct and influence the viewer through a painting by exploiting this more organic and intent-driven model of vision? Berger's detailing of his four painting modes in early modern painting system, specifically the optical and textural modes, appears to be a way to parse through this possibility. Berger describes the "overlapping debate about the relative merits of the 'smooth' and 'rough' manners of painting which arise as a response to the drift of the optical toward the textural mode" (Berger Jr, 1998, p. 499). Once again performatively reacting to Italy and specifically Titian's late work that introduced significant brushwork and bodily trace, we hypothesize that Rembrandt appropriated and extended Titian's technique from just 'texture as material and bodily trace' to 'texture (with optical, as they are interrelated) as a tool for artistic agency and intent' – directing and influencing Crary's (1992) observer's gaze through a painting, not only through high level cognitive constructs like scenes, people and lighting but through more textural and painterly means (Figure 23). Rembrandt's late portraits seem to display hints of this intended agency through textural tracings, including brushstroke 'centre of interest', focusing techniques and 'lost and found edges.' These brushwork-based, textural plane techniques that coerce and lead the observer's eye through a painting are well noted as standard techniques in late 20th century modern art and portraiture, but are seldom attributed to having their birth in the early modern period or with Rembrandt.

Figure 24 *Self Portrait with Beret*, 1659. Focused left eye, lost left lower cheek edge & dark background send viewers eye toward cent of interest.

Figure 25 Top) *Portrait of an Old Lady*, 1662. Focused left eye, lines, and contours wrapping the face. Bottom) *Elderly Man*, 1667. Disjointed highlight lines pull your gaze in all directions, emphasizing a confused almost drunk depiction.

4.5 Conclusion

On a modest level, the work in this chapter, has attempted to extend Berger's textural mode definition to include textural agency and support historians like Berger to return the 'organic eye' through textural agency constructs back to the early modern period via the same scientific thought that obscured in the first place. We are not implying a direct link between the science of human eye movement with late Rembrandt and other artists in the early modern period. We are putting forth a speculative claim that artists, even without a deep knowledge of the science, have intuitively learned to exploit many human vision, perception and cognitive traits in their craft. Using this knowledge and science intertwined with art scholarship, it is possible to bring some clarity to both art and science questions. We then attempt to use these techniques to hypothesis that Rembrandt, reacting to his Italian contemporaries, developed specific painterly techniques, typically not associated with the early modern period, that influence a observers eye-brain modelling of a work, heralding in a deeply intuitive and communicative new technique that artists have embraced thought-out the modern era.

4.5.1 Future Evidence Gathering

The work from this chapter was conceived and published early in the thesis research (DiPaola, 2008). Once established, we than sought to gather strong quantitative evidence to support its main hypothesis – Rembrandt's texture agency to influence eye gaze patterns of his viewers. We did this by conducting novel eye tracking studies using controllable output paintings from our painterly NPR system documented in Chapter 7. First, however we built up the functionality of painterly NPR toolkit, which is documented in the next two chapters.

5: Painterly System – Parametric Approach

5.1 Overview

We now describe our painterly NPR toolkit, which has been informed by the work from Chapter 3. Our research methodology, as shown in Figure 26, has been to gather and systematically convert qualitative painterly knowledge into a quantitative parameterized software NPR toolkit.

Figure 26 Main Process Chart: From input qualitative data through art and cognitive filtering to NPR toolkit and its output results and benefits.

Our first task has been to collect and quantify painterly soft rules. With this knowledge, we can begin to create both low-level NPR components as well as build initial higher level, more adaptive techniques, which are built up from the lower level parameters. This language based, parameterized approach is the heart of our painterly NPR toolkit, giving it the control, and customization that is needed for its stated goals. The systems parametric goals are that

the language is rigorous, repeatable, aligned to science and customizable, allow it to be used as a toolkit for scholarly investigation. Most NPR systems are concerned mainly with quality output, our systems is also concerned with quality output -- our goal is high quality and wide range of stylized results; however our main concern is that it is built from the ground up as an art and science NPR inquiry tool. To help achieve this in our system, we use knowledge domain based multi-dimensional and hierarchical parameter spaces (Parke, 1982; Valentine, 2001) as depicted in Figure 27, which have been used successfully in several computer science fields areas including generative systems and computer facial animation.

Our portrait painter research comes out of work we have explored with multidimensional parameter spaces for evolutionary generative systems (DiPaola & Gabora, 2009; DiPaola et al., 2013) and computer facial systems (DiPaola & Arya, 2006, 2007a; DiPaola, 1991, 2002). The basis of this approach is to create a language of low level parameters that are object-oriented, encapsulated and mathematically rigorous. These can be thought of as letters in a specialized alphabet, which form the basis for words and phrases (high-level components). The key to our low-level parameters, as with real letters, is their rigour and universality - any high level component should be derived from them. In implementation, we use XML parameters as our low-level dimensions (e.g. axes) in a scriptable large knowledge space, which can be accessed individually, via equations and/or through higher-level constructs, including 2D dimension maps, which are solely comprised of the lower level parameters often with logical, spatial and temporal attributes. (DiPaola & Arya, 2007b).

For example, in facial animation, low level muscle parameters can be built up into a more semantic ‘smile’ parameter then our ‘smile’ with other parameters and temporal considerations can be built up into high level construct ‘joyousness’. In practice, the low level parameters of our NPR toolkit are grouped into functional types:

1. *parameter constants* like brush size or colour weighting which typically have floats as descriptors but can also remap into many different knowledge buffers such as depth maps (see Section 5.3.2).
2. *method parameters* such as `ClosestColourPalette` which uses a method where all brush colour choices for a given pass or area use an remapping into a second palette based on the closest colour

3. *process method parameters* which guide the process flow of the other parameters types (i.e., do this before that).

Figure 27 Parameterization Process: From soft knowledge to a space of millions of individual paintings, a parameter space (middle) defines a set of factors (params) whose values determine the characteristics of a system (dimensions of the space). With one individual (a painting) being a point in that space.

5.2 Initial Phase One NPR Toolkit

We now overview our phase one painterly NPR toolkit, which was informed by and built to exploit artistic knowledge, gathered from the previous chapters. The system is a Java based toolkit that accepts as input our XML configuration language, an input source image and any number of knowledge map files (i.e. palettes, content region maps, salient region maps, brush textures, depth maps, calculation maps) then works automatically to render blob or brush stroke based output images using an OpenGL 3D library. Parameters to our system can be simple floats (i.e. brush curvature = .5) or full remapping through a raster calculation map (i.e. brush curvature at this x,y point in the output painting is based on same point in the calculation map - e.g. Figure 35). There is also a full spline and mathematical expression mechanism within our language for complex remapping and control (see Section 6.1 XML code). The system can work at different scale resolutions for calculation maps both saving space and speed, as well as simulating different levels of needed cognitive detail, less for some things more for others. The system can also output a rendered XML based brush stroke

list file with full details of stroke styles, positions and curvature nodes including temporal and comment data. We have used this stroke script output to render images using Corel Painter; a natural media brush stroke renderer with sophisticated natural brush stroke capabilities via its automatic macro facility.

To exploit tone and colour portrait knowledge, (e.g. soft rules: [R1.1 Tone first then colour], [R1.3 Unequal value], [R2.1 Warm/Cool colour], [R2.2 Unequal colour]) the system uses a multi-layer stroke analyser/renderer, which perceives and lays strokes down in large masses first, progressively using smaller strokes and more detailed analysis. This multi-pass tonal shape based approach approximates how painters squint at first to read large tonal masses and progressively add greater levels of detail over exposed paint from the layer before which fit with our soft rules: [R1.2 Regions as likeminded tonal masses], [R4.1 Relative widen/reduce/harden/soften], [R7, 8.1 Overall relativity and rescaling]. Rather than use progressive difference Hertzmann based grid techniques (Hertzmann, 2001; Huang et al., 2011; Kagaya et al., 2011) to move through a source image, the system progressively iterates over semantic perception blob masses, optionally beginning with major semantic content areas (Figure 28 - below left). We describe our hierarchal perception blob masses in detail in the (phase two) ThinkerPainter section. This technique needs no semantic starting point, but can work like artists do, and iterates through known content-based semantic regions, progressively refining to more detail via luminance-calculated regions. Santella and Decarlo (2004; 2002) also use hierarchical segmentation and a corresponding tree representation however they use a segmentation pyramid mainly for extract meaning edges while we are interested in cognitive blobs -- regions of progressive cognitive attention of a painter. Our 'Jane Series' set of images, demonstrates the hierarchical blob tree where we have false coloured level 1, level 2, and level 3 (from left to right) of the luminance blob level in the top of Figure 28. For level 0, we used the semantic map represented by the eyes, hair, skin and clothes maps (Figure 28 right). Any semantic blob map can be used as they are created, labelled and called on by the scripting language via layers in a Photoshop PSD file. They are currently semi-automatically made in the case of source photographs or automatically made (via depth buffers) in the case of 3D CGI source (see Chapter 6). With our hierarchical sub-blob system, these regions are sub divided into perception regions down to a detailed tree structure for eventual brush stroke tiling. This technique is one attempt to solve what many

researches including Collomosse refer to as one of the biggest issues in most NPR systems. That is they work locally and not globally.

From our cognitive science analysis of artistic practice, much of the decisions are relative to the global and level regions (e.g. soft rules: [R1.3 Unequal value], [R2.1 Warm/Cool colour],[R2.2 Unequal colour], [R4.1 Relative widen/reduce/harden/soften], [R7, 8.1 Overall relativity and rescaling]). Our tree structure of hierarchal perceptual blob regions, allows communicating back up the tree to a parent blob for non-local information, making value, colour, highlight and other plans and decisions at local stroke time.

Figure 28 For the output image created (right) of the Jane Series, our smartpalette system used content maps (left) to better inform the colour choices as well as started our cognitive based hierarchical blob tree as level 0, with luminance calculated blobs at lower and lower levels of the tree (levels 1,2,3 from right to left – top) that allow local and global relative communication within the toolkit.

A value plan is used which can rescale and re-centre the tonal space of source masses. Unequal dominant and sub-dominant value parameters also can rescale how the system analyses/uses the source image. This is an example where painterly knowledge rules can supersede the information from the source sitter image by filtering, emphasizing/de-emphasizing and scaling input information as well as by other rule based means. These rescale and remapping techniques supports soft rules: [R1.2 Regions as likeminded tonal masses], [R1.3 Unequal value], [R2.1 Warm/Cool colour], [R2.2 Unequal colour], [R4.1 Relative widen/reduce/harden/soften], [R7, 8.1 Overall relativity and rescaling].

The system has several low level colour method parameter routines that remap into a colour temperature system via tonal value information at stroke time. Since many artists work in a known palette space, our current high level colour components create a pre-constructed palette by sampling the source photograph(s) as a pre-process along with other image analysis or by choosing historical portrait palettes from any given artwork with optional smart regions. Currently the colour palette remapping can be achieved via many different method parameters at a per image, per pass, per region or per stroke level. With this method we have been able to use any number of historical colour palettes separated into body tone, shadow, clothes, background and hair sub-palettes which we than can map to any input photograph, region by region based on weighed constant parameters such as value or hue. This technique supports our R1.1, R1.3, R2.1 and R2.2 colour soft rules and was used to create our four Rembrandt adaptable eye tracking study images, where we can precisely calculate tone to colour temperature mappings only in the semantic regions that Rembrandt used (i.e. the body tone or warm lighter part of the skin) as a direct mapping to our source sitter images (Figure 29 and 30).

Figure 29 Using semantic regions, `FullRangeValuePalette()`, the sitter images of PhD students (lower right), were transformed in 1 of 4 customizable Rembrandts for our eye-tracking studies.

Figure 30 Using semantic regions, `FullRangeValuePalette()`, the sitter images of PhD students (lower right), were transformed in 2 of 4 customizable Rembrandts for our eye-tracking studies.

Figure 31 Controlling the parameters, allows the toolkit to generate several or hundreds of different painterly portrait scenarios in the specified multi-dimensions as shown by these four images.

Figure 32 The example set taken from a run of hundreds, shows moving through a one parameter (above - brush size), then other examples (cropped) of many paths through the portrait space from the same input. Note these early results were prior to our implementation of the JCh CIELAB colour model.

An artist or scholar has the ability to control different axes of the parameter space, which can generate one, several or thousands of correlated portrait images smoothly ranged in the given multi-dimensions. Figures 31 and 32 show example images, cropped from the high resolution originals, from larger image sets (hundreds or thousands) iterated in different dimensions. We have developed an interactive batch creation and viewing tool that allows users to view paintings from a multidimensional space of premade output using sliders through the axis of the painting space.

5.3 ThinkerPainter (Phase Two) Framework

The phase one system as described above was used to create the parameterized output for our four texture agency – eye tracking studies (i.e. Rembrandt paintings) described in detail in later chapters, as well as to output imagery that has been shown in major gallery and museum shows. However, from lessons learned, our phase two system, while using the same basic art knowledge /cognitive process approach, was rewritten to better emulate (using

parallel object oriented classes) human painter cognitive processes, as well as deal with a level of complexity we had reached with this emulation in our phase one system.

The ‘ThinkerPainter’ framework is our updated system. Rather than thinking in terms of individual blobs, regions or brushstrokes, this new framework thinks in terms of ‘PaintActions’, which represent a hypothetical painter’s intention to act on the current state of the canvas by specifying a general area to paint, a colour to use and some attributes to consider while painting the region (for example, whether it should be sparsely or densely painted). While the program runs, object oriented components called ‘Thinkers’ analyse the current canvas as well as the source photograph, generating PaintActions corresponding to the artist’s high-level intentions. Thinkers pass the PaintActions they create to components called ‘Painters’, which are responsible for rendering these PaintActions as paint on the digital canvas, typically first as perception blobs then strokes within the blobs. Most importantly, lightweight components called ‘Concerns’ may intercept PaintActions between the Thinker and the Painter and modify them to make small changes to the painting based on specific painterly considerations, for example, making regions near a specified centre of interest more detailed. With ‘Concerns’, PaintActions can abstract out specific variations in painterly technique from the larger two component types and maintaining an equitable separation of concerns.

The interoperable and better-distributed nature of components in the ThinkerPainter framework is far more suitable for recombining components, investigating cognitive passed art techniques in studies, and supporting a wider variety of painterly output than was possible using the phase one system. Additionally, because each component is now smaller and functionality is more loosely-coupled, new features can be developed much faster and reused much more easily. Thus, the system has improved both for users and developers, greatly facilitating its continued development and use.

The Thinker is responsible for identifying areas to paint and how to paint them. The Painter is responsible for rendering paint to the canvas. The two objects communicate through PaintActions, which contain sets of high-level parameters. The process flow and major components are shown in Figure 33.

Thinker - Analyses the current canvas as well as the source photograph and creates PaintAction objects representing its high-level intent to paint on a specific region of the canvas in a certain way. References any number of Concern objects that can suggest modifications to the PaintActions it produces.

Painter - Receives a PaintAction object from a Thinker, and decides how to render it to the canvas. It may, for instance, decide to make many small strokes for high-detail regions, use a palette knife to fill gross regions, etc...

Concern - A modular, lightweight object that can be plugged into any Thinker to suggest modifications to the PaintActions it produces (Palettes would be a type of concern).

PaintAction - An interface containing high-level parameters specifying a region to be painted as well as how it should be painted:

Area: Alpha channel representing the canvas region to paint.

Detail: How fine brushwork should be.

Density: How densely the paint should be laid.

Colour: The colour to use.

Colour Variation: How uniform the colour should be.

Flow

Figure 33 Process flow and main modules of ThinkerPainter (phase 2) NPR system.

Thinker Painter re-examines the canvas after individual strokes rather than entire passes (as our phase one system did), more accurately modelling the painterly process, reducing redundant strokes and permitting the use of many different end conditions. It splits our many techniques into two modular segments (Thinker and Painter) for better reusability, permitting the painting of image regions by means other than individual brushstrokes.

5.3.1 Thinker Painter Concepts: BlobThinker, ActionPainter, RelativeJChPalette

BlobThinker

As a fine art painting process encourages artists to see and draw in regions of light and dark, BlobThinker seeks to regard the reference (or sitter) image in the same way, creating a knowledge structure about the reference image based on regions of similar value. This structure is hierarchical, dividing the image into progressively fine regions of ‘lightness’ and ‘darkness’ joined in a tree structure. Its purpose is to provide some means by which our system can think about the reference image that is higher-level than pixels or strokes, and which approaches a painterly cognitive process of thinking about what area of the source/canvas to act on next, viewing that area, processing it internally and acting to paint that semantic cognitive region of the source image on the canvas – what we call a semantic blob or sub blob. See Figure 28 - top for a false colour imaging of the blob tree (Levels 1,2,3), where child blobs fall within parent blobs.

As a Thinker component, BlobThinker’s task is to repeatedly identify areas of the canvas it would like to act on. By default, and if there are no ‘Concerns’ that would specify a specific process are active, it does this by selecting every ‘leaf’ blob with no regard for the larger tree in which it exists. However, we have implemented a number of high-level constructs including an error metric that is used to decide when a cluster of leaf blobs should instead be painted as a single larger parent region. (See Figures 38 and 39).

‘Meta’ Palettes

Although belonging to the same class, two Palettes are used not for selecting colours but rather for recombining the colour selections of other Palette objects. PSDPortraitPalette is a ‘meta’ Palette used for applying various different Palettes to different regions of the canvas (the skin, the background, the clothing, etc...), while BlendPalette is used to perform weighted colour blending between two different Palettes (an especially helpful process for animation work). PSDPortraitPalette was used for our Rembrandt work (albeit in our phase

one system), as well as most high level output work, where the regions of the source portrait (or scene) can best be semantically matched to a corresponding palette from an existing masterwork regardless of the specific tone to colour matching techniques used. So while a user can script to use our best RelativeJChPalette, colour systems (the colour that is picked for a given stroke), it will still only choose colours from a palette in that same semantic area (e.g. the body shadow, the hair or lips) from a Rembrandt original painting. We have a database of scores of smart semantic regional palettes from noted paintings. It takes approximately 20 minutes using Adobe Photoshop to make any smart palette from a fine art painting source. As is true with our general methodology with our toolkit, once we gain confidence that an experimental high-level construct has long standing merit, we then can write image processing software or Artificial intelligence software to automate the process.

RelativeJChPalette

RelativeJChPalette is the end result of an experimentation process begun with FullRangeValuePalette and carried through SmartHSVPalette and RelativeHSVPalette. This toolkit's high-level constructs all use some form of rescaling and remapping tonal to colour conversion techniques – attempting to better model artist practice. The underlying idea is that paintings contain 'value systems' -- that is, distributions of light and dark values across the canvas -- that are of great importance to their composition. RelativeJChPalette accepts a 'colour-value system' (that is, a lightness distribution with specific hues/saturations mapped to specific lightness values) as input, and attempts to apply this system to the reference image in the same way a painter might apply her preferred (remapped and rescaled) value system and colour palette to her subject. It should be noted that it took months of extra development work to re-implement a full CIELAB JCh colour space within our system, which slowed the system down by a significant amount. However, we felt that the new results were worth of the effort. This is an example of learning from artistic practice through a cognitive science lens and re-implementing that knowledge into our system. Our soft rule research in Chapter 3 has demonstrated that luminance is much more important to sample off the source image. With accurate luminance, more sophisticated semantic methods of converting tone to a correlated colour temperature can be used. This 'tone first' innovation from art practice,

demonstrated to us that a regular software colour conversion space system was simply not sophisticated enough for this artist driven method.

ActionPainter

ActionPainter, being a Painter component, is not tasked with analysing or understanding the reference image, but instead with rendering faithfully the PaintActions it receives (PaintActions, again, represent intent to act on the canvas in a particular way). ActionPainter directs strokes by laying splines along the reference image's gradients. The vast majority of ActionPainter's logic is concerned with respecting the shapes defined in PaintActions as well as responding to their 'detail' and 'density' parameters (which, respectively, signal intent to paint grossly/finely and densely/sparsely).

ActionPainter, more than BlobThinker, is constrained greatly by the ThinkerPainter interface (specifically, the PaintAction object). Because the parameters received via PaintActions are supposed to determine the way in which it paints, it is critical to design a set of parameters that is expressive and allows for versatility.

5.3.2 Blobs, Stroke and Concerns

Our system allows for relative scaling, comparisons, iteration decisions and smart palette decisions, based on the main factors that artists use, which as we have reported in Chapter 3 are an authored combination mainly working with light, but also volume or content. We have already described our hierarchically blob tree facility and our semantic content maps/labels allowing the user to work with light (tone mapping) and content (e.g. if I am working in the 'background', I should use large brushes of constrained colour). We would like to present a facility to work with volume knowledge and use this example to describe our Concerns mechanism.

A specific scripted Concern intercepts PaintActions between the Thinker and the Painter and modifies them to make small changes to the painting based on specific painterly considerations, for example, making regions near a specified centre of interest more detailed with specialized stroke density and detail parameters. We have created a number of new

Concerns as a lightweight development means to test and implement our more open ended and less explored artistic and cognitive soft rules knowledge. Concerns allows us to try out algorithms that emulate new soft rules, allowing for experimentation. One such area is performing specific adaptable brush stroke setup in a given pass, content area, region or situation based on a Concern. Concerns can communicate up the blob tree, so density, detail, complexity, value blend, stroke direction, stroke type can all be programmed based on a Concern via relative and global comparisons.

With volume, we documented from our stroke artistic soft rules in Section 3.2.6: “One technique is to pull or push the stroke in the direction of the plane.” To do this, the plane or 3D volume of the scene or object must be known. Our system is open to communicate with any knowledge equation, reference data or 2D buffer (2D map). Using a CGI 3D face model as source, a Concern can specify to, rather than stroke by luminance gradient, the system should stroke by source depth buffer gradient fulfilling the above ‘stroke in the direction of the plane’ soft rule. Figure 34 shows a Concern that can stroke by volume gradient (right), that is the direction of the plane, (e.g. the side nose slope), or by more typical luminance gradient (left). All stroke parameters, like size, length, density, curvature, transparency are still active but work on the depth rather than the luminance gradient.

Since artists use both techniques (switching between depth or luminance) at one time, we are experimenting with a sophisticated artist process Concern that moves between the two methods based on content and other knowledge, (for instances, a Concern could be: if I am currently in [Nose, Cheek or Eyes] content regions and the depth planes is over a certain error angle: use depth gradient 80% of the time, else use luminance gradient). Eventually we see Concerns tied to artificial intelligence modules which evaluate more cognitively rich processes.

Figure 34 A ‘Concern’ to use either a 3D depth buffer (left) or regular luminance (right) or some intelligent mix for stroke direction decisions.

5.3.3 Limitations

Our painterly NPR system is built with rigour, extensibility and scripting in mind, so researchers can attempt to experiment with new NPR techniques and use the system for a variety of research uses in both art and vision science. . It is less built for speed and we have found that waiting for a result does hinder the research process. Rendering, as of this writing, currently takes from 1 to 10 minutes per output image on a PC or laptop with a mid-level graphics card. Implemented software code better aligned to the computer’s accelerated graphics card (i.e. GPU) is of interest to us. We do this to a great extent already, since our main graphics techniques use the 3D language OpenGL, which is hardware accelerated. However, the JCh colour space conversions are our current bottleneck and we have yet to see a GPU (Graphics Processing Unit) implementation of these colour conversions.

With so many language based scripting choices, the toolkit is a bit daunting to use. We hope to implement a high-level authoring system that automatically writes out scripts that can then be executed by the system. We have begun looking at using Artificial Intelligence techniques (currently Genetic Algorithms) for this high-level front end to ease this issue.

As a toolkit, the system is not automatic. It is built as a scripted language that can communicate with knowledge buffers of many kinds including depth buffers, colour palettes, layers of labels maps and mattes, which give it its knowledge based, customizable abilities. However, in practice many of these image or 2D buffers and data files are typically still created offline in tools like Adobe Photoshop or colour palette creators. Our goal is standard file formats and extensibility (same language commands can work on any buffer type) over ease of use. We have automated some of these processes in general and especially when the source image and data comes from a CGI 3D source (Figure 34). We look to automate more pre-setup processes in the future. This is not a trivial area however, as most of these setup processes (e.g. auto regioning and labelling areas of the face/clothes/background in portraiture) are content semantics dependent, i.e. different for faces / portraiture than separating, labelling, and prioritizing objects in a landscape scene. We have begun conversations with Zhao (Zeng et al., 2009; Zhao & Zhu, 2011) about shared research in this semantic labelling area, including using semantic parse trees as an natural extension onto and above our blob trees.

We would like to extend the system with additional content knowledge semantics other than portraiture and faces. Nothing in the internal JAVA software code side is specific to portraiture, but many assumptions, knowledge buffers and scripting takes advantage of portrait knowledge. We need to understand other knowledge approaches in other painterly genres (still-life, figurative, landscapes, ...) , test how they do in our system and iterate on our toolkit approach to build a system that works best in all conditions.

Lastly, the field of Non Photorealistic Rendering (NPR) is a fast moving innovative area with many new techniques and strong active researchers. Our system as a general modelling toolkit is not able to incorporate all innovations in this area, nor do we claim it is better than the top systems. Instead it attempts to work as a toolkit for research that uses a computer modelling of artistic and cognitive knowledge as its main approach. In some ways, its main contribution is as an extensible and rigorous NPR parametric language rather than as a standalone system. We are working on making it as extensible as possible (hence working on any kind of knowledge data or buffers) to better incorporate new techniques from other researchers.

5.4 Conclusion

We have overviewed our phase one and phase two NPR toolkits. Appendix A documents all the current parameters available to our system with their syntax and usage for our main classes: Root Element, Pass Elements, Thinkers, Painters, End Conditions. In the upcoming chapter, we continue to document our toolkit and show how we make painterly output imagery using scripting, Concerns and other processes, plus detail our low level implementation.

6: Painterly System – Scripting, Output and Implementation

6.1 Scripting, Concerns and Examples

Figure 35 Test Jane Series Image which showing the Maps (inset) that affect its brush parameters across the canvas (see script) using a SmartPalette from a historically painting by Hopper for its tone mapped regional colour choices.

Below is an excerpt of one NPR toolkit XML script (in indented short form for readability) with several previously documented facilities demonstrated and commented (in red). The main processes start indented blocks in blue: a `pass` with a `thinker`, `painter`, and `concern` section. The Concern is a simple smart palette where content regions use a JCh tone

first colour mapping to their counterpart regions of a painting by the artist Hopper. Much of the commented areas show how stroke parameters can call simple floats variables or abstracted to full 2D map files that vary the stroke parameters of the canvas. Figure 35 shows these source maps created as layers in Adobe Photoshop to vary stroke curvature, opacity, stroke density, stroke detail and stroke thickness (cloud map). They also use math equation facility in the script to create new content. For instance, you will notice the brush curvature affected by the horizontal map, the opacity affected by the vertical map and the brush thickness affected by the cloud map.

```
% script demo with many available script constructs displayed
region-map % load 2d maps
  file resources/sourceImages/jane1/jane1.mat.psd
alpha-map
  file resources/testPatterns/tptest.psd
pass % begin a pass
  thinker class BlobThinker
    blur-size 1 relative FALSE
    blobs-from-regions TRUE
    leaf-blob-size 0.00066
    blob-blur-size 3
    blob-complexity 1
    density class AlphaChannel % Maps the stroke density float to a labelled horizontal stripe black-to-white map
      id horizontalStripe
    detail class AlphaChannel % Maps the same horizontal map to stroke detail parameter black-to-white map
      id horizontalStripe
  painter class StrokePainter
    grid-size class FloatMath relative TRUE % maps strokes to get thicker where 'clouds' map values are higher
    multiply -1 % Invert the alpha channel so it now goes from '-1' to '0' %note equation math system in toolkit
    add 1 % Add 1 so that it goes from '0' to '1' again but with black and white inverted
    multiply 0.03 % Rescale the alpha channel so it goes from '0' to '0.03'
    add 0.005 % Add '0.005' so it goes from '0.005' to '0.035'
    parameter class AlphaChannel
      id clouds
    brush-scale 1
    opacity class FloatMath % Map 'horizontalStripe' alpha map to stroke opacity
      multiply 0.8
      add 0.2 % Rescale alpha channel to go from '0.2' to '1'
    parameter class AlphaChannel
      id horizontalStripe
    curvature class AlphaChannel % Map 'verticalStripe' alpha map to curvature
      id verticalStripe
    min-stroke-length 1
    max-stroke-length 25
    seed 1
    gradient class RefGradient
      filter-size 0.05 relative TRUE
    error class RefError
    brush class eduPainter.toolkit.jogl.GLTextureBrush
      brush-file resources/brushes/spatter.gif
      brush-scale 1
    end-condition class ActionCountCondition
      max-actions 0
  concern class eduPainter.palette.PSDPortraitPalette
```

```
map-file resources/sourceImages/jane1/jane1mat.psd  
palette-folder resources/palettes/hopperself  
palette,class,eduPainter.palette.RelativeJChPalette,,
```

Figure 36 Black and white source image for the Jane Series. Since no colour exists in the source, all colour is generated through using tone to remap into a rescaled and regional art historical colour palette.

Using the same black and white source image (Figure 36) the Jane series output in Figure 37 uses four progressively refined passes with the final pass executing centre-of-focus details only for the eye, eye whites and lip region using an include-region command via labels from the Photoshop layer resource image:

```
pass  
include-regions, eyes eyewhites lips  
thinker,class,BlobThinker
```

It also uses a Concern to value-blend two smart palettes for its oil paint creamy skin effect.

Figure 37 Jane Series One: with details eye and Concern script to value-blend smart pallets.

‘Jane Series’ output in Figures 38 and Figure 39 started with the same basic Jane script but use Concerns to vary many stroke density parameters that affect the error, lightness and colorweight, which create unusual stroke density effects:

```
concern class ErrorConcern
  localization 1
  min-error 5
  max-error 35
  lightness-weight 1
  color-weight 0
```

Figure 38 Jane Series Two

Figure 39 Jane Series Three

Varying parameters a bit more gives a woodcut like style in this last painterly output (Figure 40) from the Jane Series group of paintings, which were accepted for submission as artwork in peer reviewed art gallery shows in Annecy Art Festival, France and at the TenderPixel gallery in London, UK. Annotated script printouts hung next to the art work.

Figure 40 Jane Series Four

6.1.1 Use in Industry Research: 3D Pipeline Automatic Control

As part of an NSERC Engage industry - research partnership grant, we collaborated with the long-time former cinematographer of Pixar, Jerrica Cleland and her pre-visualization company. We were interested to see if our scripts could automatically use knowledge from (and controlled by) a 3D authoring system (e.g. Autodesk Maya) in their design pipeline where Concerns were controlled within the authored 3D objects in the spline based key-frame facility designers typically use. The goal was to begin to give a vast array of artistic

style and look options (NPR styles) to their pre-visualization clients, but in a way that was controllable and part of their known and economical pipeline. Figures 41 and 42 show example output of 3D characters and objects where the content regions maps including non-portrait content like ‘playing cards’ and ‘table’ were created and labelled automatically by the 3D authoring system. We created stills and animation sequences that could interpolate NPR styles. While animation is possible with our toolkit (Figure 42), more development work is needed to build temporal parameters into its knowledge system. The figures are all stills from animated sequences with automatic regioning.

Figure 41 Images stills from industry / research partnership (NSERC) to control our NPR toolkit automatically from a 3D animation pipeline.

Figure 42 Animation stills from industry / research partnership (NSERC) to control our NPR toolkit automatically from a 3D animation pipeline.

6.1.2 Component Details and Implementation

We end this chapter with a more detailed software implementation description of the ThinkerPainter components, object oriented class structure and data flow as seen in, Figure 43.

Figure 43 Thinker painter main data flow.

6.1.3 The Top-Level Control Structure

TPRunner – the painterly NPR toolkit (called Painterly in our codebase) loads the reference image and converts it to JCh, loads the region map and alpha maps, and then creates the output canvas. It stores these entities in a new RunData instance that can be accessed by any component at any point during runtime.

TPPass - Painterly then locates the XML element corresponding to the first pass and instantiates a TPPass instance using the element to configure it. It creates new instances of the reference image and region map scaled according to the reference-resolution parameter; the region map may only contain a subset of regions if the include-regions parameter has been specified. Painterly then creates a mat defining the region on which this pass will operate, defined by the union of all included regions (which is then optionally intersected with a mat defined by the alpha-map parameter). All of these objects, as well as the previously-mentioned RunData object, are stored within a new PassData instance accessible by any component within the pass. Lastly the TPPass instantiates it is associated Thinker, Painter, End Condition and Concern components, then executes the pass.

After the first pass has been executed, Painterly will instantiate and execute each subsequent pass in the same fashion until no further passes remain, then save the configuration and output canvas to the specified output path.

6.1.4 BlobThinker, ActionPainter, RelativeJChPalette in Practice

Instantiation

BlobThinker

When BlobThinker is instantiated, it generates a hierarchical blob structure from the reference image. To do this, the ‘root’ blob is first created, which encompasses the entire canvas; if the blobs-from-regions parameter is true, this root blob is then divided into an initial batch of sub-blobs (one for each included region of the canvas). A new instance of the reference image is created and Gaussian blurred (kernel size is determined by BlobThinker’s blur-size parameter). This blurred instance is then passed to the blob generation algorithm.

Essentially, the blob algorithm subdivides parent blobs into child blobs by locating contiguous groups of pixels whose lightness values (that is, the ‘J’ in JCh) fall above the parent’s mean lightness as well as groups whose mean lightness falls below it. Thus, in principle, blobs represent different value regions within the reference image that may be perceived grossly (towards the root of the blob structure) or finely (towards the leaf blobs).

The algorithm proceeds iteratively, dividing each leaf blob into yet more leaf blobs until no leaf blob can be further subdivided without its children having an area smaller than the one defined indirectly by the num-leaf-blobs parameter. This parameter actually sets a ‘target’ number of blobs based on the total canvas area; for example, a value of 4500 would set the minimum allowable area for a blob to be the canvas’ total area divided by 4500 (Figure 44)

Figure 44 Left, the reference image; right, a visual representation of hierarchical leaf blobs in a structure derived from this reference.

ActionPainter

Although far less involved than BlobThinker’s, the instantiation of the ActionPainter object does involve the creation of gradient and error maps (‘error’ meaning the difference in lightness between the current output canvas and the reference image) for use during stroke placement. See Figure 45.

Figure 45 Left, a gradient map of the reference image; right, an error map derived from the reference.

PaletteConcern

PaletteConcern objects are a special kind of Concern in that they are usable simply by declaring a palette parameter within the pass configuration. They act as a wrapper around Painterly’s palette system; upon instantiation, these objects will in turn instantiate whichever Palette subclass has been specified in the configuration, using the Palette to generate a colour map of the reference image. It should be stated that Concerns are where high-level cognitive constructs can be implemented and tested. To date we have only a few and believe this is an area for great further investigation.

RelativeJChPalette

RelativeJChPalette accepts an image (typically though not necessarily a colour gradient) as input, each unique colour value from which is stored as an available palette colour for use by the painter. More importantly, the component also maps the distribution of lightness values across the pixels in this image. This distribution is maintained as a series of data points, one for each discovered lightness value, marking what percentage of pixels are of lower or equal lightness (effectively a mapping of absolute lightness to ‘relative lightness’, or

‘lightness relative to the image domain’). A second data structure of the same type is then created from the reference image.

When selecting the colour for a given pixel, the algorithm identifies the reference image’s relative lightness at that location (by referencing the second data structure) and then selects the data point from the first data structure that has the closest relative lightness. (Thus, for example, if 75% of the reference image is darker than the given pixel, the algorithm selects a palette colour with ~75% of the other palette colours darker than itself). From here, the algorithm identifies every palette colour having an absolute Lightness corresponding to the appropriate relative lightness, and selects whichever is closest in terms of absolute JCh colour difference to the reference image colour. See Figure 46.

Figure 46 A colour map generated by PSDPortraitPalette and RelativeJChPalette using the ‘Hopperself’ palette.

6.1.5 Algorithm – The Thinker/Painter Loop in Detail

The Blob Algorithm

Once every component has been instantiated, the Thinker/Painter loop begins. Each iteration of this loop begins by calling the ‘think’ method on BlobThinker. The first time ‘think’ is called, BlobThinker generates a list of every leaf blob in its structure. It selects the first blob in this list and creates a new PaintAction whose target region of the canvas is generated by smoothing the edges of the blob’s shape, first by applying a Gaussian blur with a 3-pixel radius and then applying a ramp function to harden the blurred edges somewhat. This new PaintAction’s ‘detail’ and ‘density’ attributes (which concern the style in which the targeted region should be painted) are calculated by finding the mean values of BlobThinker’s density and detail parameters across the targeted region. These parameters may be specified in the configuration as constant float values or as greyscale mats.

The Palette

BlobThinker returns this PaintAction to the loop, where it is then passed to the PaletteConcern component through the ‘evaluateConcern’ method. PaletteConcern calculates its colour map’s mean JCh value across the given PaintAction’s target region, and sets the PaintAction’s ‘colour’ attribute to the result of this calculation. It then returns the PaintAction, now with colour information in addition to shape and style, to the loop.

Figure 47 ThinkerPainter Flow in detail.

The Action Algorithm, in Detail

Next the PaintAction (see Figure 47) is passed to the ActionPainter component, which is tasked with rendering the PaintAction to the output canvas. The method this component employs borrows the basic methodology of Action's original technique (that is, superimposing a grid overtop the reference image and seeding strokes in each cell), though its particulars have been modified extensively. The algorithm begins by finding the mean grid-size value across the PaintAction's target region (grid size can in fact vary across the surface of the canvas); this value determines the size of the grid cells overlaying the target region. Next, the algorithm generates a rescaled instance of the target region such that one pixel corresponds precisely to one grid cell; this determines the alpha value of the target

region at each grid cell, providing a grid-level representation of the region's boundaries. A 'filled grid map' is then created to record at grid-level the degree to which each cell has been filled by brushstrokes.

The algorithm then uses the error map produced during instantiation to generate a list of seed points. One point from each grid cell, that which has the highest error value, is placed in the seed list.

Using the gradient map generated during instantiation, a 'default stroke direction' is calculated by finding the mean gradient value within the target region. All strokes begin travelling in this direction.

At this point, the algorithm iterates through the seed list, attempting to lay down one stroke per seed. No stroke is seeded in grid cells already sufficiently occupied by previous strokes. Should a stroke be seeded, however, its path is determined iteratively by laying control points to form a B-Spline that attempts to move perpendicularly along the reference gradient. Each stroke is assigned an 'alpha' value, which is initialized as the alpha value of the seed cell. As control points are placed, this alpha value is lowered to accommodate any target region boundaries or previously filled grid cells the stroke encounters along its path, thus theoretically preventing strokes from travelling outside the target region or overfilling one particular part of it. The PaintAction's detail attribute determines the extent to which previously filled or out-of-region cells diminish the stroke's alpha, where maximum precision ensures no stroke crosses these boundaries and minimum precision allows all strokes to do so at will. This implies that the detail parameter available in BlobThinker determines quite directly how accurately or grossly ActionPainter will fill its blobs.

Once a path has been determined, the algorithm checks the stroke's alpha value; should it be above 0 the stroke will be laid. After a stroke has been laid, the filled grid map is updated, each grid cell touched by the stroke becoming more filled; the extent to which these cells are filled is scaled by the PaintAction's density attribute, with high density causing no grid cell to report itself as more full and low density causing grid cells to add the stroke's entire alpha value to their fullness. This implies that the density parameter available in BlobThinker determines directly how densely or sparsely ActionPainter will fill its blobs.

Following this elaborate process, the algorithm renders each stroke to the canvas using OpenGL acceleration, making use of ActionPainter’s brush parameter to determine the style and scale of the stroke. A stroke’s thickness is, similarly to the old framework, determined firstly by grid size, then multiplied by the brush-scale parameter.

Subsequent PaintActions

At this point, the first PaintAction created by BlobThinker has now been rendered to the canvas. The pass loop, having completed one iteration, will then ask BlobThinker for a second PaintAction, a third, and so on; BlobThinker will iterate through its list of leaf blobs, returning each one in sequence. The PaintActions it generates will be given colour by the PaletteConcern and rendered by ActionPainter. The process continues until any component orders it to stop. Once every leaf blob has been painted, BlobThinker, having no more PaintActions to recommend, will order the stop, and the pass will at that point be complete. See Figure 37.

Figure 48 One pass, rendered to the canvas. This pass uses the same blob, gradient and error maps pictured above.

7: Cognitive Science Studies – Rembrandt’s Textural Agency and Eye Fixations

7.1 Introduction

A painted portrait differs from a photo in that only certain regions are selectively rendered in sharp detail. Artists have claimed that these choices guide the viewer’s gaze, but it has been difficult to test this because portrait gaze is typically also associated with greater meaning, stronger lighting and a more central location. Using our parameterized painterly NPR system, we report on four experiments in which we monitor viewer gaze using portraits where these factors are decoupled from relative detail. Portraits were rendered with our parameterized non-photorealistic system to mimic Rembrandt (DiPaola, 2007, 2009). Participants first viewed and then assigned artistic ratings to each portrait. Results showed that viewers’ gaze was attracted and held longer by regions of relatively finer detail (Experiments 1, 2), by textural highlighting (Experiment 3) and that artistic ratings increased when portraits strongly biased gaze (Experiments 1, 2, 4). These findings imply that successful portrait artists rely on an implicit understanding of how gaze is directed by relative detail. These studies give strong cognitive correlate evidence to the soft rules: [R3.1 Highlights - eye fixations], [R5.1 Lost & found edges - eye fixations], [R5.2 Centre of focus - eye fixations].

7.2 Background

A painted portrait is the product of a dynamic relationship between a sitter (the model in the portrait), an artist (the creator of the portrait) and a viewer (the audience intended or imagined by sitter and artist) (Berger Jr, 1998). The finished portrait can therefore be read as a narrative at several levels, including one of the intended self-presentation of the sitter, through their choice of posture, facial expression and costume. It is also a narrative about what the artist has chosen to emphasize, through a selection of painting style, colour palette and level of detail in the portrait. Finally, each portrait tells a narrative about how the sitter

and artist have chosen to communicate with viewers, who will look at the portrait with the intent of reading both the character of the sitter and the character of the artist in presenting the sitter to an audience.

As we have discussed in the previous chapter, an important tool wielded by the artist is the choice of which regions in the portrait to be rendered in greatest textural detail and which regions to be completed only in the imagination of the viewer. By directing the viewer’s gaze to selected regions, the artist is able to draw attention to certain character traits of the sitter and to provide a glimpse into the collaboration that occurred between sitter and painter in its creation (Berger Jr, 1998; Miall & Tchalenko, 2001; Nicholls et al., 1999). While there are several known and interrelated portrait artistic techniques for directing the viewer’s eye, such as caricature -- where areas are scaled or distorted in shape, chiaroscuro and other lighting emphasis as well as compositional emphasis, we will look at textural detail techniques. For example, well-known portraitist John Howard Sanden (2004) describes how he believes “centre of focus sharpening techniques” help to structure the experience of the viewer. In one of his portraits shown in Figure 49, the increased textural and color detail in the sitter’s left eye and eyebrow are intended to move the viewer’s gaze to these locations, thus helping to draw attention to the sitter’s intelligent, yet playful personality. Sanden further emphasizes this interpretation with his loose directional brush strokes under the eye, which he uses to guide the viewer’s gaze to the more sharply accentuated eye area, both by the implicit gestures of the spiraling large brushstrokes and by the repetition of large-scale asymmetric curvature around both the eyes and the mouth.

Figure 49 J.H. Sanden’s *Reverend Cole close-up*. From (Sanden & Sanden, 1999).

It is one thing to claim, as Sanden (2004) and other modern artists do, that textural detail in a portrait guides the gaze of the viewer, but it is quite another to find direct empirical evidence from a cognitive theory perspective (i.e. cognitive correlate), in support of this hypothesis. When we first began to examine the *detail-gaze hypothesis* from the perspective of vision science, we were surprised to learn that it had not yet been put to a direct test. On the one hand, there have been numerous previous studies examining the gaze patterns of viewers while they were inspecting works of original art – (reviews in (Livingstone, 2002; Molnar, 1981; Yarbus, 1967)) but in each case it is difficult to attribute the gaze patterns of viewers directly to the selective emphasis in the painting involving the degree of textural detail. The reasons for this lack of direct evidence are quite straightforward when considered from the perspective of visual fine art creation. When a painter selects one region of the canvas for increased detail over another region, these regions also invariably differ from one another in their meaningful content (i.e., they are usually of foreground rather than of background interest), in relative degree of lighting (i.e., textural detail is usually increased for

surfaces depicted as in direct light) and in relative spatial location (i.e., regions of increased detail are often at the centre of the composition). Of course, strong correlations such as these in an artwork — between the semantic level, the compositional level and at the level of textural detail — all likely conspire in a synergistic way to guide the gaze of the viewer to selected regions of the painting. But, for the purposes of putting the current claim — that textural variations in themselves guide the viewer’s eye — to a proper scientific test, these correlations in original portraits make it impossible to test by monitoring viewers’ gaze while they view portraits.

There is a considerable literature on how the eyes are guided when viewing scenes of various kinds (e.g., (Henderson, 2003) but this research also offers very little on the issue of how relative detail may guide the eyes during the exploration of a portrait. The research on scene viewing suggest that although eye fixations may initially be attracted to regions of high image salience (Itti & Koch, 2001) — defined as local regions that are markedly different from their surroundings in colour, orientation, motion and depth — once a scene has been viewed long enough so that its semantic content can be determined, the eye tends to explore those image regions deemed by the participant to be most informative with regard to the task they are performing (Birmingham et al., 2008). It would certainly seem possible, in principle, to incorporate relative detail into both the guidance of gaze through image salience (e.g., by including regions of sharp detail as one of the defining properties of salience) and into the guidance of gaze by task demands (e.g., when the task depends on high-resolution information). However, at the moment, relative detail is an image feature that has not yet been systematically explored with regard to its ability to influence gaze in either a bottom-up (image salience) or top-down (goal-directed) fashion.

Finally, we note that there is large applied literature in both photography and motion pictures that offers practical advice to those wishing to guide the viewer’s gaze. Photographers are advised to use blur and/or a large depth of field to direct the viewer’s gaze away from those regions of an image that they wish to de-emphasize (Langford, 2000, p. 81).

A portrait or close up still photograph might use a small depth of field to isolate the subject from a distracting background. The use of limited depth of field to emphasize one part of an image is known as selective focus or differential focus.

Cinematography, more since the late 1960s, began to rely on similar selective focus principles adding a dynamic or temporal dimension to its essentially story telling usage. For instance tight focus is typically shifted (rack focus) during a shot to lead the eye of the viewer from foreground to background, from one character to another during dialogue or to emphasize the narrative by selectively bringing a new detail (i.e. the gun on the table) into focus.

Figure 50 Source image (lower right), original Rembrandt *Hendrickje* close up (upper right) and NPR system output using style and palette from *Hendrickje*, in this case with manipulated ‘textural highlights’ supporting main eye for Exp 3.

7.3 Testing the Detail-Gaze Hypothesis

Our approach to testing the detail-gaze hypothesis involved generating portraits using the painterly NPR toolkit that were plausible works of art, and yet portraits under parametric control using our XML scripting language, which allow the relative level of detail to be

decoupled from other factors such as meaningful content, lighting and spatial layout. This was done through a three-step process. First, we photographed human models posing, dressed and lit in a similar way to four of Rembrandt’s most famous late portraits: *Self Portrait with Beret and Turned-Up Collar*, 1659 (hereafter *Beret*); *Man with a Magnifying Glass*, 1661 (hereafter *Man*); *Hendrickje Stoffels*, 1660 (hereafter *Hendrickje*) (Figure 50). and *Large Self-Portrait*, 1652 (hereafter *Large*) (Figure 51). Second, we rendered these photographs in the style of Rembrandt using our Painterly rendering system (DiPaola, 2009) where approximately 50 parameters of brush details, colour palette and other painterly attributes were matched as close as possible to the original Rembrandt portraits (DiPaola, 2007). Third, we selected four regions in each rendered portrait for selective manipulation with regard to textural detail: one region centred about each eye and one region centred on each side of the chin, where the material of the collar meets the skin of the neck, as illustrated in Figure 52. The variation in textural details for each eye was achieved using additional passes of progressively smaller brush strokes in the Painterly system algorithm as a base, with additional Gaussian blur and stroke manipulations where appropriate. The collar areas used similar methods but added edge detail (or loosening) to the parameterization to approximate what artist’s describe in their writings as the ‘lost and found edge’ technique, where the eye often prefers to follow tonal or colour edges, and when those edges purposely disappear (are ‘lost’) the eye can be guided into the artistic centre of focus.

Figure 51 Four original Rembrandts and our photographs of the human model analogues used in the eye tracking study.

This left us with the opportunity to compare gaze patterns of viewers examining the original photo of the models (where textural detail is more uniformly high for all image regions), with their gaze patterns when viewing the same models rendered as artistic portraits with systematic variation of detail in the chosen regions. Finally, to ensure that our results would be specific to the degree of relative detail — independent of relative location in the image — we presented viewers in Experiment 1 with both the original orientation and the mirrored image of the portraits.

A general discussion of the results and their implications will be presented following the detailed presentation of the design and results specific to each experiment.

7.4 Experiment 1

Our first experiment was designed to explore the effects of varying the degree of relative detail in selected regions of a portrait independent of other features, such as meaning, lighting and spatial location, which typically co-vary with detail. In participants’ first examination of each of the portraits, we tracked their eye position in order to assess the

detail-gaze hypothesis. Then, in a second inspection of each of the portraits, participants rated the artistic merit of the portraits, giving us the opportunity to look for correlations between gaze measures and the artistic ratings given to portraits.

The eye tracking record of each portrait was analysed with regard to four main questions: (1) Are there global differences in gaze pattern when participants view photos, critical portraits, and filler portraits? (2) Which of the selected regions of photos and critical portraits receive the most fixations? (3) How does the relative detail of a selected region influence gaze patterns? (4) Are eye tracking measures taken during the first viewing of a portrait associated with ratings of artistic merit made during the second viewing?

Figure 52 Top: Close-ups of 2 of 4 images for Beret (+Mirrors) where we vary the right or left detail of the eye (centre of focus) as well as the right and left lost and found ‘sharp’ edge. For all 4 Rembrandts. The original Rembrandts and the original sitter photographers were used for our studies. Bottom: *Hendrickje*.

7.4.1 Method (Experiment 1)

Participants. Thirty-two undergraduate students (11 male, 21 female) participated in return for extra-course credit in a one-hour testing session.

Apparatus. An EyeLink II tracker sampled eye position every 2 milliseconds. Saccades (eye movements) and fixations (periods of stable gaze) were assessed using the default settings, namely a saccade is a spatial shift with amplitude of 0.58 deg or more, with an acceleration threshold of 9,5008/s and a velocity threshold of 308/s. Images were shown on a Samsung 19 inch LCD screen with a resolution of 1024x768 and 24 bits per pixel.

Images. The images shown to participants consisted of 20 portraits specifically designed to test the detail-gaze hypothesis along with 20 filler portraits (selected from an assortment of fine art books covering noted artists from different periods and styles) intended to create a diverse context of portraiture in which the critical portraits could be judged. Each of the critical and filler images also appeared in both original and mirror image orientation, for a total of 80 images. All images were approximately 1 megapixel in resolution, with square or slightly rectangular aspect ratios.

The 20 critical portraits included the four original photos (of the models shown in Figure 51), along with four non-photorealistic computer-generated portraits based on each photo. These four portraits were created by varying the relative detail in the two eye regions, together with independent variation in the relative detail of the two collar regions below each eye. Detail was increased in a given region by requesting progressively finer brush strokes, combined with a slight Gaussian blurring of the other region. The four portraits from each photograph were therefore selectively more detailed in either (a) left eye/left collar, (b) left eye/right collar, (c) right eye/left collar, (d) right eye/right collar. Each of these four regions of interest (ROIs) was defined in advance of the viewing session as circular regions about 145 pixels in diameter.

Procedure. Following a brief period of eye tracking calibration, each participant viewed a series of 30 images, each for 5 seconds, and in a random order. Each trial began with a central fixation point that had to be fixated in order to trigger the presentation of a portrait. On the first pass through the images, participants were instructed to inspect each portrait in any way they saw fit, but not to make any overt responses or decisions. They were also told

that they would be given a second opportunity to view the images, at which time they would be asked to indicate how much each image appealed to them as ‘a work of art’. Prior to beginning this second viewing of the images, participants were instructed that they would again have the opportunity to view each image for 5 seconds, before indicating their judgment of the artistic merit of the image on an 8-point scale. Participants were instructed to use the entire range of the scale, reserving a rating of 1 for the ‘worst’ image in the set and a rating of 8 for the ‘best’ image.

The 30 images shown to each participant consisted of 10 of the 40 critical images shown once, along with 10 of the filler images shown twice. The filler images were shown twice in order to create a context in which the repeated presentations of a given model’s portrait (with slight variation in detail in the four regions of interests) would not seem unusual. Because each participant was shown only 20 of the 80 possible images in the experiment, four different sets of images were prepared for each of 8 different sets of participants.

Data Analyses. Several eye tracking measures tested for global differences in looking pattern as a function of portrait type (photos, critical, fillers), including the total number of saccades and the total time spent fixating during the first 5 second viewing period. Other eye tracking measures focused more specifically on the four regions of interests in the photos and critical portraits, including latency and duration of the first fixation to a region, the number and duration of discrete fixations directed to a region and the frequency of saccades from one region to another. During the second viewing of the set of portraits, the artistic rating assigned to each one was recorded with a key press (digit keys 1-8). Analyses of variance (ANOVA) were used to examine the eye gaze and rating measures as a function of portrait type.

7.4.2 Results (Experiment 1)

1. Rembrandt-renderings lead to a calmer eye in the viewer. Analysis of the number of saccades made in viewing each portrait type indicated that participants made fewer saccades when viewing critical portraits (mean = 14.1) than when viewing fillers (mean = 15.4) or photos (mean = 15.5), $F(1, 30) = 44.37$, $p < .001$, with no significant differences among fillers and photos, $F(1, 30) < 1.0$. Note that because viewing time was fixed at 5-seconds per

portrait, a smaller number of saccades implies a longer average fixation duration, which was approximately 30 ms longer for critical portraits than for fillers and photos.

2. *Eye regions of portraits attract most frequent fixations.* Figure 53 shows the superimposed fixations of all 32 participants on one of the critical portraits. When we examined the frequency of fixations in the four selected regions, the two eye regions received more than 52.1% of all fixations, the two collar regions received 4.3% of all fixations, and the remaining 43.6% of fixations were made outside the defined regions of interests. This strong bias to fixate on the eyes in a portrait is consistent with previous studies of gaze toward faces (Birmingham et al., 2008; Henderson, 2003). As also illustrated in Figure 53, for those fixations lying outside the selected regions, most lay in the nose region extending between eyes and the mouth, with the remainder coming to rest on various idiosyncratic locations on the face, hair and shoulders.

Figure 53 All fixations made for one of the Rembrandt-renderings during the 5 sec viewing period for all 32 participants in Experiment 1. Each participant’s fixations are coded in a different colour.

3. *An eye with greater detail attracts a first fixation in less time, for a longer duration, and attracts more repeat fixations.* When we examined the time taken to first fixate in one of the two eye regions in a critical portrait, we observed a strong influence of relative detail.

Figure 54 shows that the time taken to make a first fixation to either of the two eye regions in the photos is a little over 900 ms. However, in the critical photos the time of the first fixation depends much more on whether the eye has greater or lesser detail, with the first fixation to the more detailed eye occurring 150 ms earlier than the first fixation to the less detailed eye, $F(1,31) = 10.05$, $p < .01$. The first fixation in an eye region was also of longer duration (mean difference = 32 ms) when the regions of interest was rendered in greater detail, $F(1,31) = 4.42$, $p < .05$. Inspection of the total number of fixations made to the two eye regions showed that the tendency for an eye region of greater detail to attract looking continued for the entire viewing period, $F(1,31) = 11.47$, $p < .01$. Finally, when we examined the conditional probability of successive fixations to the same eye region versus moving to the other region, there was a stronger tendency for viewers’ gaze to move from an eye region of less detail to one of more detail in a portrait, than in the opposite direction, $F(1,31) = 5.84$, $p < .03$.

Figure 54 Mean time to first fixation (in milliseconds) in one of the two eye regions for photos and Rembrandt-renderings in Experiment 1. Error bars are +/- one standard error.

Figure 55 highlights an important interaction observed when the number of fixations to the two eye regions was examined as a function of the relative detail in the two collar regions. Recall that whereas the two eye regions received more than 50% of all fixations, the two collar regions in themselves received only 4% of all fixations. Yet, the repeated fixations made to the eye with greater detail were influenced by the location of the greater detail in the collar regions as well. In particular, the difference in repeated fixations to the more versus less detailed eye was larger when the more detailed collar was on the opposite side of the more detailed eye, $F(1,31) = 4.47$, $p < .05$. This same interaction effect was observed in several other measures, including proportion of total dwell time, $F(1,31) = 4.40$, $p < .05$ and time to first fixation, $F(1,31) = 3.86$, $p < .05$. This interaction is consistent with the more detailed eye attracting even more of the viewer’s gaze when it neighbours a region of even less detail (i.e., when the collar nearest the eye is also less detailed)

Figure 55 Mean number of fixations made to high and low detailed eye regions, separately as a function of whether the more detailed collar is on the same or opposite side of the more detailed eye region. Error bars are +/- one standard error.

4. *Artistic merit ratings during the second viewing of each portrait are predicted by looking patterns in the first view.* Figure 56 shows the mean artistic ratings for all 40 portraits in this study, along with some examples of the portraits rated most highly, at an intermediate level and those rated lowest. The relatively small standard error bars indicate that there was considerable reliability across participants in these ratings. On average, the studio photos fell in the low to moderate range (mean for photos A-D = 3.5 to 5.2) among all portraits, while the critical portraits as a group were in the moderate range (mean for A-D = 4.1 to 5.3). However, it is also clear that there was considerable variation in the mean ratings within the set of critical portraits, ranging from a low of 3.7 to a high of 5.4.

Figure 56 Mean artistic ratings for the 40 portraits in Experiment 1, ordered from highest to lowest. Selected sample portraits are shown below in miniature form.

We next examined whether ratings of individual portraits could be predicted by some aspect of the gaze pattern measured during the first viewing of the portraits. We discovered that the best predictor of a high rating was if the portrait tended to have a large difference in the frequency of fixations to one of the two eyes. This relationship is shown in Figure 57. Across the 40 photos and critical portraits in Experiment 1, there was a significant correlation

between the difference in fixation frequency for the two eye regions and mean artistic rating, $r(38) = .381$, $p < .02$. This indicates that guiding the participant’s eye strongly to one eye region rather than the other in a portrait is associated with increased artistic preference.

Figure 57 Scatterplot of the correlation between the difference in fixation frequency for the two eye regions and the mean artistic rating in Exp. 1. Each data point represents the mean rating of 8 participants on 1 of the 40 photos and critical portrait.

No differences were found between original and mirror orientations for any of the parameters analysed.

7.5 Experiment 2

This experiment addressed several questions left unanswered following Experiment 1. First, the design and interpretation of Experiment 1 were premised on the assumption that our Rembrandt-like renderings of the modern day photos were similar to original Rembrandt

portraits, both in the way they guided the viewer’s eyes and in their artistic merit ratings. Yet, this assumption was left untested, for the practical reason that we could not manipulate relative clarity in an original work of art, independent of lighting and spatial location. However, armed with the evidence from Experiment 1 that eye gaze is influenced by relative clarity, as artists have anticipated, we were now prepared to compare original Rembrandt portraits directly with renderings that had been designed to mimic Rembrandt’s style or to violate it.

A second change we made to the design of the experiment was to vary the origin of the participant’s eye position at the beginning of each trial. In Experiment 1 each portrait was preceded by a fixation at the centre of the image. Perhaps this contributed to the disproportionate number of fixations made to the eye regions, which were also generally at the centre, rather than to the collar regions, which were lower down. In Experiment 2 we therefore began each trial with fixation either in the lower right or left corner of the image viewing area. This forced the eye to move past the collar regions before arriving at the eye regions. It also more closely simulated the typical gallery experience of first encountering a portrait from one side or the other as one enters a room or moves from one frame to the next.

A third change concerned the way in which we obtained participant ratings of artistic merit. In Experiment 1 a rating was given to each portrait when it was viewed in isolation. With such a procedure it is possible that participant ratings were not sensitive to the subtle differences in the various versions of each portrait (e.g., distinguishing between a detailed left versus right eye region). This time we gave participants the opportunity to make a forced choice among all four versions of the critical portraits for a given model. That is, the four variants of a given portrait were randomly placed into each of the four quadrants of a large screen monitor and participants were asked to say which one was ‘the best.’

7.5.1 Method (Experiment 2)

Participants. Twenty undergraduate students (4 male, 16 female) participated in return for extra-course credit in a one-hour testing session.

Apparatus, Images, and Procedure. Methods were the same as in Experiment 1 with the following exceptions. Each of the participants viewed the same set of images in this

experiment, albeit in a different random order. The image types included (1) the four original Rembrandt portraits cropped in a comparable way to (2) the four photos of the modern models, and (3) two of the critical portraits for each of the four models, and (4) ten filler portraits each viewed twice, for a total of 36 images. The two critical portraits for each model included the one with a detailed eye and collar on the same side (i.e., corresponding to the original Rembrandt) and the variant with the detailed collar on the opposite side from the more detailed eye (referred to as the anti-Rembrandt for convenience). This time portraits were shown entirely in their original orientation (i.e., no mirror image presentation).

Each trial began with a fixation cross either in the lower left hand corner of the screen (for 1/2 of participants) or in the lower right-hand corner of the screen (for the other 1/2). Fixating this symbol for a period of 500 ms initiated the presentation of the next portrait in the series.

Following the first viewing of the 36 images, participants were shown the filler portraits and the original Rembrandts in a new random order so that they could give ratings to each image, as in Experiment 1. However, the photos and the critical portraits were not included in this series, so that the original Rembrandts could be rated in the context of the filler portraits without influence from the photos and critical renderings, which were similar.

Finally, the same participants were seated at another computer with a larger screen (24 inch Mac computer) so that they could make forced-choice preferences among the four variants of the critical portraits for each model. Each participant was shown 16 screens, in a random order, each screen containing the 4 portraits for a given model (one portrait in each quadrant, also randomly determined). Thus each participant was given the opportunity to select ‘the best’ variant four times for each of the four models.

7.5.2 Results (Experiment 2)

Rembrandt-renderings lead to a calmer eye. Figure 58 shows the mean number of saccades made in viewing each portrait type. Participants made fewer saccades (and therefore longer fixations) when viewing critical portraits, than when viewing original Rembrandts, photos, or filler portraits, $F(3, 57) = 4.73$, $p < .01$. This translates into an average fixation

duration that was 20 ms longer for the Rembrandt renderings than the other types, $F(1, 57) = 12.50$, $p < .01$.

Figure 58 Mean number of fixations made during the viewing of each type of portrait in Experiment 2. Error bars are +/- one standard error.

2. *Eye regions of portraits attract most frequent fixations.* Despite the initial fixation beginning below and to one side of the portrait, the two eye regions still received 45.1% of all fixations (compared to 52.1% in Experiment 1), the two collar regions received only 5.0% of fixations (compared to 4.3%), with the remaining 49.1% of fixations being made outside the four selected regions of interest (compared to 43.6%). This finding confirms that the strong bias to fixate the eye regions and the very weak tendency to fixate the collar regions immediately below the eyes is not a function of an initial central fixation position.

3. *An eye with greater detail attracts a first fixation in less time, for a longer duration, and attracts more repeat fixations.* Figure 59 shows the average time taken to make a first fixation to one of the two eye regions in the various types of portrait. In comparison to the

800 ms taken to fixate one of the two eye regions in a photo, the more detailed eye was fixated 200 ms earlier in the original Rembrandts, 110 ms earlier in the pro-Rembrandt renderings and only 30 ms earlier in the anti-Rembrandt renderings. ANOVA examining the factors of eye detail (sharp, coarse) and portrait type (original, pro-, anti-) indicated significant effects of both detail, $F(1, 19) = 40.98$, $p < .01$, and a detail x type interaction, $F(2, 38) = 3.87$, $p < .05$.

As in Experiment 1, an eye region with greater detail continued to attract more fixations for the entire 5 sec viewing period, $F(1, 19) = 39.90$, $p < .01$. Examining this tendency with respect to the pro- and anti-Rembrandt renderings indicated that it was even stronger for anti-Rembrandt portraits, where the more detailed collar was on the opposite side of the detailed eye, $F(1,38) = 7.28$, $p < .05$, making the detail in that eye appear even more distinct.

Figure 59 Mean time (in milliseconds) taken to make a first fixation to an eye region in Experiment 2. Error bars are +/- one standard error.

4. *Artistic merit ratings during the second viewing.* Figure 60 shows the mean artistic ratings given by participants to the four original Rembrandts and the ten fillers portraits. The relatively small standard error bars indicate once again that there was considerable stability across participants in these ratings. The original Rembrandts fell in the low to moderate range (mean for photos A-D = 3.5 to 4.5) in these ratings.

Figure 60 Mean artistic ratings for the 10 filler portraits and 4 Rembrandt original portraits in Experiment 1, rank from highest to lowest. Selected sample portraits are shown below in thumbnail form.

5. *Forced-Choice Preferences show bias for pro-Rembrandt renderings.* Since participants were given a choice among four portraits on every screen, having no preference (guessing) should result in any particular variant being selected 25% of the time. Across all four models, participants preferred the most Rembrandt-like rendering 42.5% of the time (136 of 320 selections), which was significantly greater than the guessing level, $\chi^2(1) = 7.04$, $p < .01$. Broken down for each model, the portrait mimicking *Beret* showed the strongest Rembrandt preference of 53.8%, $\chi^2(1) = 33.75$, $p < .01$, the portrait mimicking *Hendrickje* was 43.8%, $\chi^2(1) = 14.01$, $p < .01$, the portrait mimicking *Large* was 41.2%,

chi-sq (1) = 10.41, $p < .01$, and portrait mimicking the Man was 31.2%, chi-sq (1) = 1.35, $p < .24$. Interestingly, none of the participants were able to articulate a reason for their choices that referred specifically to our portrait manipulations, namely, to the fact that the eye and collar region depicted most directly in the light were also the eye and collar regions that were rendered in greatest detail in the portraits that were selected most often.

7.6 Experiment 3

In the previous two experiments, the detail in selected regions of a portrait was increased by using progressively finer brush strokes and it was decreased by using thicker brush strokes and Gaussian blur using the Concern mechanisms in our NPR toolkit discussed in Chapter 6 and 7. But this is not the only tool artists can wield to guide the eyes of the viewer. Another very effective technique is to manipulate the *textural highlights* of a painting, those light-valued brushstrokes that are interpreted by the viewer as specular highlights (i.e., shiny regions indicating maximum light being reflected from the viewed surfaces). Artists speak of highlights as finishing a portrait, and are usually applied last and with added agency with the final composition and effect of the painting in mind. Since (specular) highlights reflect the light source itself (and less the object), artists teach that there is some amount of play in how they can be applied, given some flexibility and agency to their application.

In Experiment 3 we selected the two Rembrandt portraits that had been given the highest artistic ratings by participants (i.e., *Beret*, *Hendrickje*), both in their original form and in our modern pro-Rembrandt rendered versions, and digitally manipulated the textural highlights using the Painterly system refinements as well as using added highlighting in Adobe Photoshop. Thus, relative detail was not varied in this experiment, but instead it was held constant at a level shown in previous experiments to bias looking to one eye region over another and to increase artistic ratings. What was varied was the relationship between the textural highlights applied to these portraits and the two eye regions.

As shown for *Beret* in Figure 61, and *Hendrickje* in Figure 62, textural highlights were applied to the rendered portraits so as to either support focus on the more detailed eye (support), to be scattered so as to divert emphasis from both eyes (scatter), to be reduced below the level of highlighting in previous experiments (absent), or to support focus on the

less detailed eye (opposite). Otherwise, the procedures were similar to Experiment 1, so that we were able to examine the influence of textural highlighting on both gaze patterns and artistic ratings.

The eye tracking record of each portrait was analysed with regard to two main questions: (1) Do textural highlights influence looking patterns in portraits over and above the already established influences of relative detail? (2) Are eye tracking measures taken during the first viewing of a portrait associated with ratings of artistic merit made during the second viewing?

Figure 61 Image showing *Beret* close ups of two different highlight variant, towards bad eye (left) and towards good eye (right), also no highlights and scattered highlights were used.

Figure 62 Image showing *Hendrickje* close ups of two different highlight variant, towards bad eye (left) and towards good eye (right), also no highlights and scattered highlights were used.

7.6.1 Method (Experiment 3)

Participants. Twenty-four undergraduate students (6 male, 18 female) participated in return for extra-course credit in a one-hour testing session.

Apparatus, Images, and Procedure. Methods were the same as in Experiment 2 with the following exceptions. The image types included (1) two original Rembrandts (*Beret*, *Hendrickje*) (2) eight renderings generated by the combination of two models and four variants of textural highlighting (support, scatter, absent, opposite), and the ten filler portraits used in Experiment 2, each viewed twice, for a total of 30 images.

Following the initial 5 sec viewing of each of the 30 images in a random order, participants were given the opportunity to view the images again, in a new random order, in order to provide ratings of artistic merit on an 8-point scale.

In part 2 all 20 images seen in part 1 were shown again in a different random order and participants were asked to rate them using the same guidelines as in experiment 1. The eye tracking record from the first viewing of each portrait was analysed with regard to two main questions: (1) Do textural highlights added to Rembrandt renderings contribute to looking patterns, over and above the role of relative detail already documented in previous

experiments? (2) Can ratings of artistic merit given on the second viewing be predicted by looking patterns in the first view?

7.6.2 Results (Experiment 3)

1. *Textural highlights influence looking patterns.* Figure 63 shows the mean number the fixations made during the first 5 s viewing period to the two eye regions in each of the five portrait types. As in previous experiments, a larger number of fixations were generally made to the sharp than to the coarse eye across all portraits, $F(1,23) = 61.64$, $p < .01$. However, this bias also interacted with portrait type, $F(4,92) = 5.98$, $p < .01$. Simple effects indicated that the difference in fixations made to the two eye regions was similar for the original Rembrandt (mean difference = 3.8), the rendering with supporting highlights (mean difference = 3.6), and the rendering with scattered highlights (mean difference = 3.0), $F(2, 92) = 1.14$, but that the difference in looking for renderings with scattered highlights was significantly greater than for rendering with no highlights (mean diff = 2.2 fixations) and for rendering with highlights emphasizing the coarse eye (mean diff = 1.2 fixations).

Figure 63 Mean number of fixations made to each eye region in Experiment 3, plotted as a function of five different portrait types. Error bars are +/- one standard error.

2. *Artistic merit ratings predicted by looking frequency to the biased eye.* Figure 64 shows the correlation between the difference in fixation frequency to the sharp versus coarse eye region and the mean artistic ratings for a portrait, separately for the two models based on *Beret* and *Hendrickje*. Data were separated for these two models because in the context of the present experiment, there was a significant difference in ratings given to portraits based on *Beret* (mean rating = 4.4) than on those based on *Hendrickje* (mean rating = 3.8), $F(1,23) = 5.77$, $p < .01$. Yet, once this baseline difference in artistic merit was accounted for, there was a strong positive relationship between the difference in fixation frequency to the sharp versus coarse eye and a portrait’s mean artistic rating, *Beret* $r(3) = .92$, $p < .01$; *Hendrickje*, $r(3) = .80$, $p < .01$.

Figure 64 Scatterplot showing each critical portrait’s mean artistic rating as a function of the difference in fixation frequency to the sharp versus coarse eye. Separate regression lines show baseline difference in artistic ratings for two portraits.

7.7 Experiment 4

Does the relationship we have found between artistic ratings and bias in looking more to the detailed eye than to the coarser eye hold only for our Rembrandt-like portraits, or is this a correlation that holds more generally for portraits of many artists and styles? We were not able to answer this question in our previous experiments for the simple reason that we had not defined the eyes in the filler portraits as regions of interest for our eye tracking analyses. In this Experiment, therefore, we defined these regions in advance and allowed a new set of participants to view and rate all of the 21 filler portraits we had used in previous experiments, in order to see if the correlation held.

7.7.1 Method (Experiment 4)

Participants. Twenty-one undergraduate students (7 male, 14 female) participated in return for extra-course credit in a one-hour testing session.

Apparatus, Images, and Procedure. Methods were the same as in Experiment 1 with the following exceptions. Each of the participants viewed the same set of 21 filler portraits in a different random order. These included the 20 filler portraits tested originally in Experiment 1, along with one additional portrait. Unlike previous experiments, this time we defined two selected regions of interest in advance for the filler portraits, 145 pixels in diameter centered on each eye, allowing us to see whether the correlation we had observed for critical portraits, between artistic ratings and differences in fixation frequency to the two eyes, also held for this assortment of portraits representing many different styles and artists.

7.7.2 Results (Experiment 4)

1. *Artistic merit ratings are predicted by looking frequency to the biased eye for a wide range of portrait styles and artists.* Figure 65 shows the correlation between the difference in fixation frequency to the sharp versus coarse eye region and the mean artistic ratings for the 21 filler portrait used in the previous experiments. This correlation was significant, $r(19) = .498$, $p < .02$, indicating that a stronger bias in looking to the detailed eye was associated with increased artistic preference even for the filler portraits we had used in previous experiments.

Figure 65 Scatterplot showing each filler portrait’s mean artistic rating as a function of the difference in fixation frequency to the sharp versus coarse eye. Several examples of the portraits are shown below the graph.

7.8 General Discussion of Experiments

7.8.1 Summary of Our Findings

These four eye tracking / NPR toolkit studies give strong cognitive correlate evidence to the soft rules: [R3.1 Highlights - eye fixations], [R5.1 Lost & found edges - eye fixations], [R5.2 Centre of focus - eye fixations]:

1. *Calm eye effect of selective detail.* Portraits with selective regions of both higher and lower levels of detail in their brushstrokes result in a viewing experience in which the viewer fixates longer in one location and makes fewer eye movements overall.

2. *Support for the detail-gaze hypothesis.* We found support for the detail-gaze hypothesis in all four experiments. Viewers look first to the more detailed of two eyes and their

subsequent fixations return to the more detailed eye more frequently than to the coarser eye (Exp 1-4).

Even regions of a portrait that are not fixated directly, nevertheless contribute to the pattern of looking that occurs between the two eye regions in a portrait. Specifically, if the collar region below an eye is less detailed, the participant will look more often at the eye above that region, even though fixations are only rarely made to the collar region itself.

3. *Selective use of detail affects viewer appreciation.* Selective use of detail in a portrait does not only guide the viewer’s looking pattern, but it has an influence on how much the viewer appreciates the portrait as a work of art (i.e. artistic merit). This was the most surprising finding in this study, for several reasons:

Artistic merit is likely based on numerous factors that were left uncontrolled in this study. As such, no guarantee that our somewhat subtle manipulations of relative detail would have an influence on artistic ratings.

Artistic merit would likely be similar to what was found here even if viewers were unable to move their eyes when examining the portraits. There is likely some ‘thin-slicing’ study (brief presentations compared to extended viewing) that is relevant here.

Therefore, to find a systematic relationship between the gaze-detail hypothesis and participants’ preferences for these portraits was a welcome surprise.

4. *Viewers can distinguish Rembrandt’s correct style.* Viewers can select a portrait at levels well above chance, distinguishing Rembrandt-like portraits from other very similar portraits that violate Rembrandt’s style; nevertheless these participants are unable to articulate the basis for their choices. Exp 2

5. *Textural highlights support and enhance texture detail.* Textural highlights added to a portrait have additional influences on looking patterns, in that they can be used to support looking to the eye that is rendered in greatest detail or they can be used to reduce the frequency of looking to that eye. Exp 3

When the participant’s looking pattern is influenced by textural highlights it has the same effect on artistic ratings as the influence coming from relative detail. Namely, portraits with textural highlights that increase looking to one eye over the other eye are judged to be better

works of art than portraits with textural highlights that lead to more balanced looking to the two eyes.

6. *The detail-gaze hypothesis is supported beyond Rembrandt portraits.* The relationship discovered between artistic ratings and the bias in looking toward one eye more than the other holds not only for Rembrandt-style portraits but also holds more generally for portraits of many artists and styles. Exp 4

7.9 Implications of These Findings

7.9.1 From a Vision Science Perspective

In 1890, William James proclaimed that it is human nature for many events to capture our “attention, including strange things, moving things, wild animals, bright things, pretty things, metallic things, words, [and] blood” (James, 1890). These experiments support that we would be wise to add relative detail to the list of visual attributes that humans find attractive as orienting cues, such as suddenly appearing objects and bright lights, eyes, ... As this detail-gaze relationship is better understood and quantified, it could have benefits in several fields including entertainment and industrial uses, and where guiding the user’s eye and attention to a particular area is important. For instance, in video games and computer animated movies virtual characters’ emotions could be emphasized by using similar textual detailing techniques under dynamic and program control. In computer human interfaces, especially where too much visual information needs to be displayed and analysed, textural detail techniques can be used as another constraint based tool along more traditional data visualization parameters.

7.9.2 From a Relationship for Art, Art History and Vision Science

Artists and scientists have different historical approaches to knowledge acquisition, usage and dissemination. This research work is one attempt to bridge these different fields from a knowledge transfer and verification point of view allowing aspects of art and aspects of vision science to both enhance and validate each other.

Supporting this notion that knowledge can flow both ways, we have presented a study as a case where artistic knowledge anticipates scientific knowledge, in that art provided the detail-gaze hypothesis, which was confirmed scientifically. In one aspect this work supports the specific hypothesis that Rembrandt, during the Renaissance, began to intuitively exploit central visioning via textual agency to influence the eye fixations of his viewers, long before central vision was scientifically understood. It begs the question, are there other techniques that artists are using but not always quantitatively describing that can be studied by scientist that will lead not to verification but to new insights in the vision sciences?

Looking at the knowledge flow from science to art, from an art historical side, this science research has been used to support art critics. While modern artists discuss their use of these textural detail techniques, from an art history perspective, it is typically not associated with or date back to the Renaissance period and late Rembrandt. This vision science work then is able to scientifically support the contention of some art critics such as Martin Jay and Harry Berger Jr. that the Renaissance application of science to art went well beyond the contribution of mathematics, perspective and geometry in the construction of a painted image. It may have also included an understanding, implicit or explicit, of the behavioural and experiential dynamics that occur when a human eye with limited spatial resolution is confronted with a large scene or image. Using similar techniques to our eye tracking and parameterized painting generation process, further work could be used, along with the art historical record, to pinpoint whether Rembrandt himself reacting to his Italian contemporaries first mastered the understanding of texture detail as a tool for guiding the viewer’s eye or whether its evolution occurred in art in a different way.

8: Conclusion

8.1 Discussion of Studies

A recent study of photograph viewing confirmed what artists have long believed, namely, that a viewer's gaze can be implicitly influenced by the selective use of image clarity and blur (Enns & MacDonald, 2012). The first fixation to a region is made sooner when the region is selectively sharper than when it is selectively blurred. Subsequent fixations to the same regions are also more frequent when it is selectively sharp versus blurred. In our four eye tracking studies using our painterly NPR system, we asked whether similar rules apply when viewers examine portraits as works of art, and whether these implicit viewing biases have an influence on the subjective appeal of the artwork.

Healthy adult participants first simply viewed and then later rated their artistic appreciation for a variety of portraits. The eye tracking data from each participant's first viewing of portrait indicated that viewers looked first to the more detailed of the two eyes, and then with their subsequent fixations, returned more often to look at the more detailed eye region than to the one that was rendered more coarsely. The first of these findings implies that the selective use of detail has an attention attraction quality. The eye rendered in greater detail draws the gaze of the viewer and thus their first impression of the artwork. The second finding points to an additional attention holding quality. Namely, regions of selective detail invite return glances over a sustained period of time.

In Experiment 2 of our studies, we documented that textural highlights in a portrait can have an influence on viewer gaze that is over and above the influence of relative detail. The results showed that either painted highlights can serve to support the tendency for a more detailed eye to attract looking, or they can work against it. Artists instructing younger artists, over the decades, have often suggested that highlights have great importance, both for directing the attention of the viewer and for establishing the narrative intended by the portraitist. They are most often the 'finishing touches' on a piece and can therefore help to

Chapter 8: Conclusion

either consolidate themes that have been established by other techniques (e.g., use of colour and edge) or to work against them to create intentional discord. Our findings are evidence that there is a measurable influence on viewer gaze when highlights are used in this way, such that the viewer's gaze pattern is reinforced when added textural highlights support the direction given by the use of relative detail, and that it is disrupted when the highlights are at odds with the layout of relative detail. For us 'textural highlights' are an example of a full knowledge gathering cycle. That is, we first gathered knowledge of their usage from our qualitative collection of artist knowledge, second we were able to build low level support quantitatively in our painterly NPR system and now thirdly, using our NPR systems parameterized output with eye tracking in experiment 2, produce scientific evidence that supports through cognitive and vision science, this artistic knowledge on texture highlights.

Our inspiration for studying the influence of textural highlights indeed originally came from the art instruction literature that we gathered in our collection phase, but more so from our critically reviewing late Rembrandt portraits, detailed in later chapters. We have already overviewed our textural highlight study findings using our parameterized NPR output using Rembrandt's *Beret* and *Hendrickje* portraits as source material. To underscore how significantly at times Rembrandt seemed to use textural highlights, as well as to see how our study findings could support art critical theory, let us look at the original *Hendrickje* portrait where highlights support a focal point, as well as another Rembrandt as an extreme in using highlights with no focal point.

In an effort to validate that, our source imagery (from the ARTstor academic database) was correct; we went to New York City where we were allowed to take close-ups photographs of the *Hendrickje* portrait in the Metropolitan Museum of Art. Our original close-up photographs taken in front of the painting in Figure 66 reveal how significant one form of texture highlights on the *Hendrickje's* cheek, which are both strictly out of place on the cheek therefore surely affecting the viewer's vision and head directly towards the dominate eye, that we later showed in this study has a very high attraction.

Figure 66 Close up of original, showing significant form of textural highlights on *Hendrickje*'s cheek from the Metropolitan Museum of Art, NYC.

When critically reviewing and doing secondary research on Rembrandt's own portrait, we also reviewed *Portrait of an Elderly Man 1667* (hereafter *ElderlyMan*), as illustrated in close-up in Figure 67. This portrait has been responsible for a lively debate among art historians, who have wondered about the purpose behind Rembrandt's use of the unusually scattered highlighted strokes in this painting (Graham-Dixon, 2007). Did Rembrandt

Chapter 8: Conclusion

purposefully scatter the highlights using a palette knife or butt of the brush to create a feeling of a man who was drunk, “unflattering” or in “disarray”? Graham-Dixon (2007) describes the figure in the *Elderly Man* as someone who has “no idea of who he is, or of where he might be going.” Because of these scattered highlights and other techniques, this portrait was once thought unfinished. But Graham-Dixon counters that “in fact the face, blurred though it appears, was painted by the artist with an immense degree of care ... in places, teased with palette knife or the butt end of the brush into forms that indicate the precise contours ... emphasizing his subject’s apparently frail hold on reality” (Graham-Dixon, 2007). While this debate may never be resolved, the *Elderly Man* portrait and the critical discourse that surrounds it, point to the possibly powerful influence that highlights can have in the viewing and appreciation of a portrait. It is one of the goals of this thesis work, that our NPR toolkit and its findings (those that we have done here and those done by us and others in the future) can support art scholars in debates like this, in addition to its use for vision science.

Figure 67 Rembrandt's *Elderly Man* with untypical scattered textural highlights.

Chapter 8: Conclusion

The most important contribution of this study, however, is that the selective use of detail can influence the viewer's appreciation of a portrait. Put simply, when viewers spent a disproportionate amount of time looking at only one of the two sitter's eyes in a portrait they tend to judge it as a better work of art. Of course, at present this is still a correlation. Reversing the direction then, it is possible that better works of portraiture influence viewers to rest their gaze on only one of the two eyes in the painting rather than looking equally at both. To put this differential gaze-appreciation correlation into perspective, it is important to be reminded of the many potential influences that play a role in whether a viewer finds a portrait to be appealing. These include factors relating to the viewer's experience, training, and past exposure to art, factors involving cultural norms and deemed societal approval, and factors involving the artist's skill in conveying aspects of the multi-layered narrative of the portrait experience to the viewer. These would include the sitter's character, intent and expression; the artist's skill and choices. It is against this backdrop that finding a consistent correlation between how strongly one of the two eyes in a portrait receives repeated looks and how highly the portrait is rated as a work of art is, simply put, remarkable!

Indeed, our initial finding of a relationship between differential gaze and portrait appreciation in the DiPaola et al (2010) data came as a surprise to us. In contrast to our a priori hypothesis that detailed regions of a portrait might attract greater looking than less detailed regions, the initial discovery of a correlation between differential looking and a portrait's judged merit was post hoc. As such, it was important to replicate this finding with independent groups of viewers and to check for its robustness, which we did in subsequent experiments. Experiment 2 showed that the relationship held when viewer gaze was manipulated through the use of textural highlights rather than only through relative detail, and the new group of participants tested in Experiment 3 on filler portraits showed that the relationship held when viewer's rated a collection of portraits that were originally outside the scope of our interest in the study and were not originally intended to test for this effect.

These considerations prompt us to speculate that the differential gaze-appreciation effect involves a causal link. Simply stated, a portrait that guides the viewer's gaze disproportionately to one of the two eyes will be considered by viewers, all else being equal, to be a better work of art than a portrait that results in a more equal distribution of gaze. We are sufficiently confident of these initial results to state this as a working hypothesis for

future research, although at the same time we acknowledge that it is technically still a statistical correlation. This means the direction of causality may yet be found to run in the reverse direction (i.e., when a portrait is judged to be especially good, the viewer's gaze will tend to look disproportionately on the more detailed of the two eyes in the image for reasons unrelated to textural detail and highlights) or that a third unknown variable is involved. Whatever the outcome, the discovery of this correlation is worth pursuing because it suggests a heretofore-unknown link between looking and liking.

8.1.1 Implications for Vision Science

The perception of paintings is possible because paintings evoke many of the same perceptual processes that are evoked by a real scene. However, paintings and natural perception also have important differences, including the inherent dual reality of paintings (Haber, 1980), the reduced dynamic range of pictorial images (Hochberg, 1980), and the fixed image resolution of paintings, which differs from natural perception where image resolution changes dynamically with every eye movement.

In natural perception, there is a tight coupling between image clarity and the viewer's spatial attention. One reason is lens accommodation, which causes image clarity to be correlated with image acuity. Objects at depth planes other than ones currently being accommodated are seen with lower resolution (Campbell & Westheimer, 1958; Fisher & Ciuffreda, 1988), making image clarity coincidental with the objects at the spatial centre of attention. A second reason is vergence, or the angular relation between the two eyes. Objects at the point of convergence will be represented with clearer images than objects not at that point, making image clarity again coincidental with objects that are currently attended. A third reason lies in foveal vision. Only those portions of a scene that are registered on the fovea are signalled with high resolution; objects in the periphery are less clear.

These visual mechanisms all conspire during natural perception to have the viewer's interest coincide with image clarity. Painting viewing, in contrast, offers the opportunity for the normal direction of causality in this relationship to be reversed. By rendering some regions of a painting in greater detail, the artist implicitly invites the viewer to interpret the gaze-clarity correlations that occur using the default assumptions of natural perception. Of

course, this high jacking of the natural clarity-attention relation in painting viewing is not immutable. Viewers can deliberately choose to gaze directly at the blurrier regions and to avoid the sharper regions, but the tendency to revert to the over-learned coupling of clarity and attention in natural vision allows paintings to be used as implicit communication devices, signalling viewers where they ‘ought’ to look, from the perspective of the artist who created the painting.

The present finding of a link between selective looking and the subjective evaluation of an artwork (i.e., the differential gaze-appreciation effect), takes this implicit guidance by the artist one step beyond simply providing a roadmap for viewing the artwork. The finding suggests that, at least in the realm of portraiture, there is a direct connection between how one looks and whether what one is looking at will be interpreted as being emotionally positive or negative. The specific finding is that looking longer and more often at an eye in a portrait that has been selectively highlighted by the artist leads to a more positive evaluation than looking more equally at both eyes. We offer as a tentative hypothesis for this finding, the proposal that viewers are sensitive to the implicit cues offered by the artist about where to look, and when these cues are highly consistent with one another the artist is credited by the viewer (again implicitly) for offering a clear roadmap to viewing the work. This sensitivity to the clarity of the goals of the artist may well operate to enhance the emotional appraisal of the artwork, in much the same way that perceptual clarity has a positive influence on the emotional appraisal of an otherwise neutral stimulus (Reber et al., 1998) and that attentional neglect of a stimulus has a negative influence on its emotional appraisal (Fenske & Raymond, 2006).

8.2 A Toolkit and Process for Visual Art and Science

Artists and scientists have historically taken different approaches to knowledge acquisition, usage and dissemination. This thesis is one attempt to build a bridge across these fields, both by transferring knowledge from one domain to another and by testing and validating claims in a way than is not usually considered conventional within a field. Our goal was to develop a NPR toolkit that can be used in such knowledge acquisition and transfer work, that has parameters and knowledge informed from both.

Chapter 8: Conclusion

In support of the idea that knowledge can flow both ways, we note first that our study began with a case where artistic knowledge clearly anticipated scientific knowledge. In particular, the collective wisdom of artists held that regions of increased detail in a portrait would generally attract more frequent and longer looking. This has now been verified for the first time, through the collection of objective data on where viewers look when first considering a portrait.

Moving in the other direction, the present study data can help build the art historical case that during the Renaissance period, Rembrandt was one of the first artists to begin exploring the consequences of varying relative textural detail in his artwork. This argument for precedence had been made previously by art historians, including Martin (1965) and Berger (1994; 1998), in the context of a much wider discussion of how the Renaissance application of science to art went well beyond the contribution of mathematics, perspective and geometry in the construction of a painted image. Specifically, these authors propose that the emerging understanding of science during the Renaissance included an understanding, implicit or explicit, of the perceptual-motor dynamics that occur when a human eye with limited spatial resolution is confronted with a large scene or image. Yet, prior to the present study, this argument had to rely on authors and readers sharing intuitive assumptions about how the gaze of an observer actually behaved when viewing works of art. Here we report data that bears directly on the issue. Looking to the future, it is possible that a similar approach to the one developed here, possibly using our categorized painterly soft rules and correlates data presented in Chapter 3 and/or our painterly NPR toolkit presented in Chapter 5 and 6, may be used to test other assumptions concerning the art historical record.

Another contribution to the interdisciplinary exchange made by this NPR toolkit and study concerns the specific experimental design and eye tracking employed to track the gaze of unbiased viewers in an objective way. When an argument is made about the relations between composition and eye gaze, based on a single, finished work of art, the argument depends on the untested assumption that gaze would be different if the painting was altered. In other words, a control or comparison condition is rarely possible when studying gaze and appreciation of specific artworks. Here we were able to test such comparisons for the first time, because we were able to use a knowledge-based NPR painterly toolkit system (DiPaola, 2007, 2009) to construct images that varied in systematic ways from one another, yet that

were judged by our participants to be plausible works of art in the same category as Rembrandt and other portraitists. The results using these images showed that looking and liking are indeed linked, and that one of the ways they are linked is through the detail-gaze hypothesis. Viewers tend to look first to regions of increased detail in an image and their gaze tends to return to these regions more often (Experiments 1 and 2). When their return fixations are such that one eye region in a portrait is examined in much greater detail than the other eye, then they also judge that portrait to be a better work of art (Experiment 3). These experiments and their initial findings therefore move the debate about eye fixations in viewing art away from the art critic's introspective reports of where they looked when viewing an artwork and place it squarely into the realm of empirical research.

The history of visual cognition is filled with examples of large gaps between meta-awareness (i.e., what participants thought they are doing) and what they are actually doing (i.e., measured performance). This has been documented in the realms of reasoning (Wilson, 2009), memory (Dunlosky & Bjork, 2008), and visual perception (Levin & Beck, 2004). Yet, nowhere is this gap larger than in the areas of visual attention and eye gaze when viewing scenes, as attested to by the well-known effects of change blindness (Simons & Rensink, 2005), inattention blindness (Mack & Rock, 1998; Mack, 2003), the attentional blink (Shapiro et al., 1997), and by the work of entertaining magicians over hundreds of years (Kuhn et al., 2008). Aside from the historical observation that vision scientists and artists have not been communicating with each other very much, there is no reason we can think of, why our modern understanding of the art experience should be lagging so far behind our more general understanding of human image and scene perception. Although these two disciplines may at times use specialized language and tools, it is undeniable that they share an interest in many of the same questions about the human experience. The work we have presented here in researching and implementing a NPR toolkit aims to combine collected and parameterized artistic knowledge with vision and cognitive science rigor, and the first set of study findings using our toolkit are offered as a small step toward improving this dialogue. Given that the painterly NPR toolkit is accessible, documented and open sourced, it is hoped that other researchers use and extend it for future work.

As discussed in Section 3.2, these soft rules or passed down painterly heuristics, are meant as a generalized norm of heuristics but given the complexity of art making have many

exceptions and deviations – in generally, most artists use an integrated and related mix of a subset of techniques that work holistically for a given authored narrative in a painting or style of paintings. What integrated mix or group of techniques work in what specific authored artistic situation is still beyond the scope of this research but an intriguing direction for future work. In the case of the Rembrandt paintings in our studies, this is probably the case as well. Therefore, it should not be assumed that the relative detailed eye technique we have examined for instance would always achieve the authored results we documented here, say in a wholly different style like an abstract painting or Seurat pointillism work. Therefore our study work presented in Chapter 7 should be considered specific to a related area of painterly techniques where our more generalized research presented in Chapter 3 and implemented in our NPR toolkit should be considered more wider – with the possibly of more specific research results to be performed in the future.

8.3 Thesis Summary and Future Work

Artists and scientists have different approaches to knowledge acquisition, usage and dissemination. This research work has been one attempt to bridge these different fields, through a multi-stage process involving the creation of a software toolkit for non-photorealistic rendering. Our domain of inquiry was the creation and viewing of fine art painting -- we were interested in elucidating cognitive and perceptual mechanisms or ‘cognitive correlates’ which correspond and relate to artists’ techniques and conceptions regarding fine art painting in general and portraiture as a content source in particular.

We analysed an extensive corpus of painterly heuristics within art-theory literature, to identify broadly accepted understandings and techniques, which might be relevant to human perception and cognition. It is our hope, that our categorized painterly soft rules are useful to other researcher in many fields.

We further condensed and categorized this artistic knowledge into a concise set of heuristics, which are suitable for parameterization and algorithmic implementation in Section 3.2 and examined findings from psychology and neuroscience, which correlate to each heuristic in Section 3.3. We presented our painterly NPR toolkit, which was informed by these heuristics within a suitable object-oriented, cognitively inspired architecture in Section

3.4 and Chapters 5 and 6. We have made available and open sourced our painter NPR toolkit with the hope that other researchers can utilize it for interdisciplinary research needs. It has been available at <http://painterly.costar.sfu.ca/>.

By interpreting artistic and cognitive science knowledge into a well-defined computational framework, we gained opportunities to formalize and test new hypotheses. We demonstrated the productive power of such an approach by examining in depth particular gaze pattern techniques (lost-and-found edges, textural highlights and varying fine detail level) used by a particular artist -- Rembrandt. We formulated four experiments based on eye tracking of human viewers, using our painterly NPR toolkit to generate example artworks with manipulated generation parameters. We then obtained findings suggesting that artists such as Rembrandt use techniques, which leverage perceptual and cognitive function to exert control over viewer's gaze patterns, which in turn influences the experienced artistic merit of a painting.

8.3.1 Future Work

We have begun a new Rembrandt / NPR toolkit, eye tracking study with a portable and mobile eye tracker; life size, un-cropped and virtually framed Rembrandt projected paintings and a Virtual Reality (VR) projection room that supports motion tracking. The main goal is to simulate a more reality based, real world art viewing experience – i.e. that of seeing a full sized framed art work in a gallery where a subject/viewer can approach the work from a distance, move close in and about the art work area while viewing. The head mounted portal eye-tracking system now allows this mobility and our cognitive science VR study room allows for motion tracking and full projection. In this way, subjects can enter the space with a virtual Rembrandt projected on a wall and move around freely as they would in a museum or gallery. Early results show that the texture-agency of eye detail placement can affect not only where and how a viewer looks at the painting but where a viewer stands. This research is ongoing.

We have interest in investigating using similar techniques presented in Chapter 7 (i.e. full psychology experiments using the NPR toolkit to vary dependent variables of the output) on

other cognitive correlates of artistic practice. Warm-cool color temperature usage and reception is high on our list.

We have begun to research specific and informed temporal coherence techniques for animating sequences with our Painterly NPR toolkit. In theory, given the blob and other knowledge data structures, we should be able to develop an informed NPR animation system. That work is still in front of us.

The system was always built with an Artificial Intelligence front end in mind, which generates scripts that the toolkit executes. We have begun research on a Genetic Algorithm based approach that tries to mimic human creativity, similar to our other work in this area (DiPaola & Gabora, 2009; DiPaola et al., 2013). We are also looking at a Deep Learning approach (Bengio et al., 2013). Our starting evolutionary technique is to use real world pairs of A) source photographs as a genetic starting point and B) the final human artist painting created from that source as the fitness function to be evolved towards. The system would then generate working scripts that try to achieve resemblance to the final painting staring with the source photograph. Since our toolkit parameters are based on source material and are generally rigorous, it is believed we would be evolving virtual painters not paintings – that is the scripts would be in a painterly style that would make unique but styled output based on different source photographs. Once this technique is perfected, many evolved style templates could be authored and re-evolved into a larger template space – creating an easy-to-use meta system front end to the painterly NPR toolkit.

References

General References (followed by Appendix B References)

- Alberti, L. B. (1966). *On Painting*, trans. John R. Spencer. New Haven: Yale University Press.
- Aristotle. (1961). *Aristotle's Poetics*. (S. H. Butcher, Trans.) (First Edition.). Hill and Wang.
- Barla, P. (2006, November 10). Representation and acquisition models for expressive rendering. Institut National Polytechnique de Grenoble - INPG.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation Learning: A Review and New Perspectives. *IEEE transactions on pattern analysis and machine intelligence*.
- Berger, H. (1994). Fictions of the Pose: Facing the Gaze of Early Modern Portraiture. *Representations*, (46), 87–120.
- Berger Jr, H. (1998). The System of Early Modern Painting. *Representations*, 62, 31–57.
- Birmingham, E., Bischof, W. F., & Kingstone, A. (2008). Gaze selection in complex social scenes. *Visual Cognition*, 16(2-3), 341–355.
- Bordwell, D. (2007). *Poetics of Cinema* (1st ed.). Routledge.
- Brown, H., & Lehrman, L. B. (2004). *Harley Brown's Eternal Truths for Every Artist*. International Artist Publishing.
- Byrne, A., & Hilbert, D. R. (1997). *Readings on color: The philosophy of color. Volume 1*. MIT Press.

References

- Campbell, P. (2004). Seeing and seeing: visual perception in art and science. *Physics Education*, 39(6), 473–479.
- Cavanagh, P. (2005). The artist as neuroscientist. *Nature*, 434(7031), 307, 301.
- Cavanagh, P. (2011). Visual cognition. *Vision Research*, 51(13), 1538–1551.
- Chatterjee, A. (2002). Portrait profiles and the notion of agency. *Empirical Studies of the Arts*, 20(1), 33–41.
- Chen, H., Liu, Z., Rose, C., Xu, Y., Shum, H.-Y., & Salesin, D. (2004). Example-based composite sketching of human portraits. 95–153. Annecy, France: ACM Press.
- Clark, S. K. (1966). *Rembrandt and the Italian Renaissance*. John Murray Pub.
- Clark, S. K. (1978). *An Introduction to Rembrandt*. John Murray.
- Colahan, C. (1919). *Max Meldrum: His Art and Views*. Alexander McCubbin.
- Collomosse, J. (2004). *Higher Level Techniques for the Artistic Rendering of Images and Video - Thesis*. University of Bath.
- Collomosse, J., & Hall, P. (2002). Painterly Rendering using Image Saliency. In *Eurographics UK Conference, 2002. Proceedings. The 20th*, 128-22.
- Collomosse, J., Kyprianidis, J. E., Wang, T., & Isenberg, T. (2012). State of the “Art”: A Taxonomy of Artistic Stylization Techniques for Images and Video. *IEEE Transactions on Visualization and Computer Graphics*.
- Collomosse, J. P., & Hall, P. M. (2006). Saliency-Adaptive Painterly Rendering using Genetic Search. *International Journal on Artificial Intelligence Tools* 551-575.
- Conway, B. R., & Livingstone, M. S. (2007). Perspectives on science and art. *Current Opinion in Neurobiology*, 17(4), 476–482.

References

- Crary, J. (1992). *Techniques of the Observer: On Vision and Modernity in the Nineteenth Century* (New Ed.). MIT Press.
- DiPaola, S. (1991). Extending the range of facial types. *The Journal of Visualization and Computer Animation*, 2(4), 131, 129.
- DiPaola, S. (2002). FaceSpace: a facial spatial-domain toolkit. In *Information Visualisation, 2002. Proceedings. Sixth International Conference on*. 105–109.
- DiPaola, S. (2007). Painterly rendered portraits from photographs using a knowledge-based approach. In *Human Vision and Electronic Imaging XII* (Vol. 6492). Presented at the SPIE.
- DiPaola, S. (2008). The Trace and the Gaze: Textural Agency in Rembrandt's Late Portraiture from a Vision Science Perspective. In *Proceedings of Electronic Visualisation and the Arts*, British Computer Society. London.
- DiPaola, S. (2009). Exploring a parameterised portrait painting space. *International Journal of Arts and Technology*, 2(1-2), 82–93.
- DiPaola, S., & Arya, A. (2006). Emotional remapping of music to facial animation, *Proceedings of the 2006 ACM SIGGRAPH symposium on Videogames*. 143–149. Boston, Massachusetts: ACM SIGGRAPH.
- DiPaola, S., & Arya, A. (2007a). Multi-Space Behavioral Model for Face-based Affective Social Agents. *International Journal of Image and Video Processing*.
- DiPaola, S., & Arya, A. (2007b). A framework for socially communicative faces for game and interactive learning applications. *Proceedings of the 2007 conference on Future Play*. 129–136. Toronto, Canada: ACM.

References

- DiPaola, S., & Gabora, L. (2009). Incorporating characteristics of human creativity into an evolutionary art algorithm. *Genetic Programming and Evolvable Machines*, 10(2), 97–110.
- DiPaola, S., McCaig, R., Carson, K., Salevati, S., & Sorenson, N. (2013). Adaptation of an Autonomous Creative Evolutionary System for Real-World Design Application Based on Creative Cognition. In *Proceedings of International Conference on Computational Creativity*.
- DiPaola, S., Riebe, C., & Enns, J. T. (2010). Rembrandt's Textural Agency: A Shared Perspective in Visual Art and Science. *Leonardo*, 43(2), 145–151.
- Du Buf, H., Rodrigues, J., & Nunes, S. (2006). Painterly Rendering Using Human Vision. *Portuguese Journal of Computer Graphics*.
- Dunlosky, J., & Bjork, R. A. (2008). *Handbook of metamemory and memory* (Vol. xiv). New York, NY, US: Psychology Press.
- Durand, F. (2002). An Invitation to Discuss Computer Depiction. Presented at the *Proceedings of NPAR 2002*.
- Elliott, V. (2007). *Traditional Oil Painting: Techniques and Concepts of Classical Realism from the Renaissance to the Past*. Watson-Guption.
- Enns, J. T. (2004). *The Thinking Eye, the Seeing Brain: Explorations in Visual Cognition* (1st ed.). W. W. Norton & Company.
- Enns, J. T. (2006). Mental schemata & the limits of perception. In *In the Mind's Eye : Julian Hochberg on the Perception of Pictures, Films, and the World: Julian Hochberg on the Perception of Pictures, Films, and the World*. Oxford University Press.

References

- Enns, J. T. (2009). Visual Experience and Immediate Memory. In Encyclopedia of Consciousness. Elsevier.
- Enns, J. T., & MacDonald, S. C. (2012). The Role of Clarity and Blur in Guiding Visual Attention in Photographs. *Journal of Experimental Psychology: Human Perception and Performance*.
- Findlay, J. M., & Gilchrist, I. D. (2003). *Active Vision: The Psychology of Looking and Seeing*. Oxford University Press, USA.
- Foster, H. (1988). Scopic Regimes of Modernity - Martin Jay. In *Vision and Visuality* (p. 152). New Press.
- Gabora, L. (2002). Cognitive mechanisms underlying the creative process. *Proceedings of the Fourth International Conference on Creativity and Cognition*, October 13-16, Loughborough University, UK, p. 126-133.
- Goguen, J. (1999). *Art and the Brain*. Imprint Academic.
- Goguen, J. A., & Myin, E. (2000). *Art And the Brain II: Investigations Into the Science of Art*. Imprint Academic.
- Gombrich, E. H. (2000). *Art and Illusion*. Princeton University Press.
- Gombrich, E. H. (2006). *Story of Art* (1st ed.). Phaidon.
- Gombrich, E. H., Black, M., & Hochberg, J. E. (1972). *Art, Perception and Reality*. Johns Hopkins University.
- Gooch, A. A., Long, J., Ji, L., Estey, A., & Gooch, B. S. (2010). Viewing progress in non-photorealistic rendering through Heinlein's lens. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*. 165–171. New York, NY, USA: ACM.

References

- Gooch, B., Coombe, G., & Shirley, P. (2002). Artistic Vision: painterly rendering using computer vision techniques. NPAR '02 Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering. (p. ff, 83). ACM.
- Gooch, B., & Gooch, A. (2001). Non-Photorealistic Rendering. AK Peters, Ltd.
- Graham, D. J., Friedenber, J. D., Rockmore, D. N., & Field, D. J. (2010). Mapping the similarity space of paintings: Image statistics and visual perception. *Visual Cognition*, 18(4), 559–573.
- Graham-Dixon, A. (2007). Dutch Portraits: The Age of Rembrandt and Frans Hals. *Sunday Telegraph Reviews*, (01-07.2007).
- Grossberg, S. [1]. (2008). The art of seeing and painting. *Spatial Vision*, 21, 463–486.
- Haber, R. N. (1980). How we perceive depth from flat pictures. *American Scientist*, 68, 370–380.
- Haeberli, P. (1990). Paint by numbers: abstract image representations. In *Proceedings of the 17th annual conference on Computer graphics and interactive techniques*. 207–214. Dallas, TX, USA: ACM.
- Hall, P., & Lehmann, A. (2013). Don't Measure—Appreciate! NPR Seen Through the Prism of Art History - Springer. In *Image and Video-Based Artistic Stylisation*. Springer.
- Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Hackett Publishing.
- Hardin, C. L. (2000). Red and yellow, green and blue, warm and cool: explaining colour appearance. *Journal of Consciousness Studies*, 7(8-9), 113–122.
- Hays, J., & Essa, I. (2004). Image and video based painterly animation. In *Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering*. 113–120. Annecy, France: ACM.

References

- Hegde, S., Gatzidis, C., & Tian, F. (2013). Painterly rendering techniques: a state-of-the-art review of current approaches. *Computer Animation and Virtual Worlds*, 24(1), 43–64.
- Henderson, J. M. (2003). Human gaze control during real-world scene perception. *Trends in Cognitive Sciences*, 7(11), 498–504.
- Hertzmann, A. (1998). Painterly rendering with curved brush strokes of multiple sizes, SIGGRAPH. 453–460. ACM.
- Hertzmann, A. (2001). Algorithms for Rendering in Artistic Styles. NYU PhD Thesis.
- Hertzmann, A. (2003). A Survey of Stroke-Based Rendering. *IEEE Computer Graphics and Applications*, 23(4), 70–81.
- Hochberg, J. (1968). In the mind's eye. *Contemporary theory and research in visual perception*, 309–331.
- Hochberg, J. (1980). Pictorial functions and perceptual structures. *The perception of pictures*, 2, 47–93.
- Horton, J. (1995). *How to Paint Skin Tones*. North Light Books.
- Huang, H., Fu, T.N., & Li, C.F. (2011). Painterly rendering with content-dependent natural paint strokes. *The Visual Computer*, 27(9), 861–871.
- Hubel, D. H. (1988). *Eye, Brain, and Vision* (First Printing.). W H Freeman & Co (Sd).
- Hyman, J. (2010). Art and neuroscience. In *Beyond mimesis and convention*. 245–261.
- Ione, A. (2000a). An inquiry into Paul Cezanne: The role of the artist in studies of perception and consciousness. *Journal of Consciousness Studies*, 7(8-9), 57–74.
- Ione, A. (2000b). Connecting the cerebral cortex with the artists eyes, mind and culture. *Journal of Consciousness Studies*, 7(8-9), 21–28.

References

- Isenberg, T. (2013). Evaluating and Validating Non-photorealistic and Illustrative Rendering. In *Image and Video-Based Artistic Stylisation*. Springer.
- Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, 2(3), 194–203.
- James, W. (1890). *The Principles of Psychology*. Harvard University Press.
- Kagaya, M., Brendel, W., Deng, Q., Kesterson, T., Todorovic, S., Neill, P. J., & Zhang, E. (2011). Video Painting with Space-Time-Varying Style Parameters. *IEEE Transactions on Visualization and Computer Graphics*, 17(1), 74–87.
- Kang, H. W., Chui, C. K., & Chakraborty, U. K. (2006). A unified scheme for adaptive stroke-based rendering. *The Visual Computer*, 22(9-11), 814–824.
- Kontsevich, L. L., & Tyler, C. W. (2004). What makes Mona Lisa smile? *Vision research*, 44(13), 1493–1498.
- Kuhn, G., Amlani, A. A., & Rensink, R. A. (2008). Towards a science of magic. *Trends in Cognitive Sciences*, 12(9), 349–354.
- Land, M. F. (1999). Motion and vision: why animals move their eyes. *Journal of Comparative Physiology A*, 185(4), 341–352.
- Land, M. F. (2006). Eye movements and the control of actions in everyday life. *Progress in Retinal and Eye Research*, 25(3), 296.
- Langford, M. (2000). *Basic Photography* (7th ed.). Focal Press.
- Levin, D. T., & Beck, M. R. (2004). Thinking about seeing: Spanning the difference between metacognitive failure and success. In *Thinking and Seeing: Visual Metacognition in Adults and Children*. MIT Press.
- Leyton, M. (1992). *Symmetry, causality, mind*. MIT Press.

References

- Lindsay, P., & Norman, D. (1977). *Human Information Processing*. Academic Press.
- Litwinowicz, P. (1997). Processing images and video for an impressionist effect. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 407–414. ACM Press/Addison-Wesley Publishing Co.
- Livingstone, M. S. (2002). *Vision and Art: The Biology of Seeing*. Harry N. Abrams.
- Mack, A. (2003). Inattentional Blindness Looking Without Seeing. *Current Directions in Psychological Science*, 12(5), 180–184.
- Mack, A., & Rock, I. (1998). *Inattentional blindness*. MIT press Cambridge, MA.
- Mamassian, P. (2008). Ambiguities and conventions in the perception of visual art. *Vision Research*, 48(20), 2143–2153.
- Martin, F. (1965). *Spiritual asymmetry in portraiture*. British Society of Aesthetics.
- Melcher, D., & Colby, C. L. (2008). Trans-saccadic perception. *Trends in Cognitive Sciences*, 12(12), 466–473.
- Miall, & Tchalenko. (2001). A Painter's Eye Movements: A Study of Eye and Hand Movement during Portrait Drawing. *Leonardo*, 34(1), 35–40.
- Molnar, F. (1981). About the role of visual exploration in aesthetics. *Advances in Intrinsic Motivation and Aesthetics*.
- Molnar, F. (1997). A Science of Vision for Visual Art. *Leonardo*, 30(3), 225–232.
- Myin, E. (2000). Two sciences of perception and visual art: editorial introduction to the Brussels Papers. *Journal of Consciousness Studies*, 7(8-9), 43–56.
- Nicholls, M. E., Clode, D., Wood, S. J., & Wood, A. G. (1999). Laterality of expression in portraiture: putting your best cheek forward. *Proceedings of the Royal Society B: Biological Sciences*, 266(1428).

References

- Noton, D., & Stark, L. (1971). Scanpaths in Eye Movements during Pattern Perception. *Science*, 171(3968), 308–311.
- Parke, F. I. (1982). Parameterized Models for Facial Animation. *IEEE Comput. Graph. Appl.*, 2(9), 61–68.
- Pinna, B. (2008). Art and perception. *VSP*.
- Prime, S. L., Tsotsos, L., Keith, G. P., & Crawford, J. D. (2007). Visual memory capacity in transsaccadic integration. *Experimental Brain Research*, 180(4), 609–628.
- Purves, D., Lotto, R. B., Williams, S. M., Nundy, S., & Yang, Z. (2001). Why we see things the way we do: evidence for a wholly empirical strategy of vision. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 356(1407), 285–297.
- Ramachandran, V., & Hirstein, W. (1999). The science of art: a neurological theory of aesthetic experience. *Journal of Consciousness Studies*, 6(6-7), 15–51.
- Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of Perceptual Fluency on Affective Judgments. *Psychological Science*, 9(1), 45–48.
- Redies, C. (2007). A universal model of esthetic perception based on the sensory coding of natural stimuli. *Spatial Vision*, 21(1-2), 97–117.
- Roberto, L., Rodrigues, J., & du Buf, H. (2006). Looking through the eyes of the painter: from visual perception to nonphotorealistic rendering. In *Proc. 14th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*.
- Rosand, D. (1981). Titian and the Eloquence of the Brush. *Artibus et Historiae*, 2(3), 85–96.

References

- Rosand, D. (1990). *Painting in Cinquecento Venice: Titian, Veronese, Tintoretto*. Yale University Press.
- Salesin, D. (2002). Non-photorealistic animation & rendering: 7 grand challenges. In Keynote talk at NPAR 2002.
- Sanden, J. (2004). *Portraits from Life in 29 Steps*. North Light Books.
- Sanden, J. H., & Sanden, E. (1999). *Portraits from Life*. North Light Books.
- Santella, A. (2005). *The art of seeing: visual perception in design and evaluation of Non-Photorealistic Rendering*. Rutgers, The State University of New Jersey.
- Santella, A., & Decarlo. (2002). Abstracted painterly renderings using eye-tracking data. NPAR '02 Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering. (p. ff, 75). ACM.
- Santella, A., & DeCarlo, D. (2004). Visual interest and NPR: an evaluation and manifesto. In Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering. 71–150. New York, NY, USA: ACM.
- Saper, C. (2001). *Painting Beautiful Skin Tones with Color & Light*. North Light Books.
- Seeley, W. P. (2006). Naturalizing aesthetics: art and the cognitive neuroscience of vision. *Journal of Visual Art Practice*, 5(3), 195–213.
- Seeley, W. P. (2011). What Is the Cognitive Neuroscience of Art... and Why Should We Care? *American Society for Aesthetics*, 31(2), 1–4.
- Setlur, V., & Gooch, B. (2004). Is that a smile?: gaze dependent facial expressions. In Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering. 79–151. New York, NY, USA: ACM.

References

- Shapiro, K. L., Arnell, K. A., & Raymond, J. E. (1997). The attentional blink: A view on attention and a glimpse on consciousness (1997). *Trends in Cognitive Science*, 1, 291–296.
- Shugrina, M., Betke, M., & Collomosse, J. (2006). Empathic painting: interactive stylization through observed emotional state. In *Proceedings of the 4th international symposium on Non-photorealistic animation and rendering*. 87–96. New York, NY, USA: ACM.
- Simons, D. J., & Rensink, R. A. (2005). Change blindness: past, present, and future. *Trends in Cognitive Sciences*, 9(1), 16–20.
- Solso, R. L. (1996). *Cognition and the Visual Arts*. The MIT Press.
- Solso, R. L. (2000). The cognitive neuroscience of art: a preliminary fMRI observation. *Journal of Consciousness Studies*, 7(8-9), 75–86.
- Stewart, E. (1997). Using edges and stroke quality to control space. *American Artist*, 61(659), 26–31.
- Valentine, T. (2001). Face-space models of face recognition. In *Computational, geometric, and process perspectives on facial cognition: Contexts and challenges*. 83–113.
- Vanderhaeghe, D., & Collomosse, J. (2013). Stroke Based Painterly Rendering. In P. Rosin & J. Collomosse (Eds.), *Image and Video-Based Artistic Stylisation*. 3–21. Springer London.
- Vidyasagar, T. R. (1999). A neuronal model of attentional spotlight: parietal guiding the temporal. *Brain research. Brain research reviews*, 30(1), 66–76.
- Wenrich, S. J. (2013). *All the Colors of Life: From the History and Mystery of Color! and Secrets of the Spectrum*. Balboa Press.
- Wilson, T. D. (2009). Know thyself. *Perspectives on Psychological Science*, 4(4), 384–389.

References

- Winkenbach, G., & Salesin, D. H. (1994). Computer-generated pen-and-ink illustration. In Proceedings of the 21st annual conference on Computer graphics and interactive techniques. 91–100. New York, NY, USA: ACM.
- Winnemöller, H., Olsen, S. C., & Gooch, B. (2006). Real-time video abstraction. In ACM SIGGRAPH 2006 Papers. 1221–1226. New York, NY, USA: ACM.
- Yarbus, A. L. (1967). Eye Movements and Vision. Plenum Press.
- Zeki, S. (2000). Inner Vision : An Exploration of Art and the Brain. Oxford University Press, USA.
- Zeki, S. (2001). ESSAYS ON SCIENCE AND SOCIETY: Artistic Creativity and the Brain. Science, 293(5527), 51–52.
- Zeng, K., Zhao, M., Xiong, C., & Zhu, S.C. (2009). From image parsing to painterly rendering. ACM Trans. Graph., 29(1), 2:1–2:11.
- Zhao, M., & Zhu, S.C. (2011). Portrait painting using active templates. In Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Non-Photorealistic Animation and Rendering . 117–124. New York, NY, USA: ACM.
- Zhao, M., & Zhu, S.C. (2013). Artistic Rendering of Portraits. In Image and Video-Based Artistic Stylisation. Springer.

References of Painterly Source Materials from Appendix B

- Albala, M. (2009). Landscape Painting: Essential Concepts and Techniques for Plein Air and Studio Practice (1st ed.). Watson-Guption.

References

- Aristides, J. (2006). *Classical Drawing Atelier: A Contemporary Guide to Traditional Studio Practice* (First Edition.). Watson-Guptill.
- Aristides, J. (2008). *Classical Painting Atelier: A Contemporary Guide to Traditional Studio Practice* (1st ed.). Watson-Guptill.
- Aristides, J. (2011). *Lessons in Classical Drawing: Essential Techniques from Inside the Atelier*. Watson-Guptill.
- Arnheim, R. (1986). *New Essays on the Psychology of Art*. University of California Press.
Retrieved from <http://www.amazon.com/New-Essays-Psychology-Rudolf-Arnheim>
- Arnheim, R. (1996). *The Split and the Structure: Twenty-Eight Essays*. University of California Press. Retrieved from <http://www.amazon.com/Split-Structure-Twenty-Eight-Essays>
- Arnheim, R. (2004a). *Art and Visual Perception: A Psychology of the Creative Eye*. University of California Press. Retrieved from
<http://www.amazon.ca/exec/obidos/redirect?tag=citeulike04-20&path=ASIN/0520243838>
- Arnheim, R. (2004b). *Visual Thinking* (1st ed.). University of California Press.
- Bahr, B. (2011). Paint Color and Light the Titian Way. *American Artist*, 75(819), 58–65.
- Baxter, K. (2000a). The true color of contrast. work of M. Padgett, 64(691), 28–31.
- Baxter, K. (2000b). A journey in light and layers. work of B. McEnroe, 64(698), 28–33.
- Baxter, K. (2001). Expanding on light versus dark. *American Artist*, 65(713), 28–33.
- Baxter, K. (2003). Sargent and Italy. Exhibit at the Los Angeles County Museum of Art, 67(727), 22–27.
- Berger Jr, H. (1998). The System of Early Modern Painting. *Representations*, 62, 31–57.

References

- Bimler, D. (2005). Are color categories innate or internalized? Hypotheses and implications. *International Journal of Clinical Monitoring and Computing*, 5(3), 265–292.
- Brady, P. (2008). *Rethinking Acrylic: Radical Solutions For Exploiting The World's Most Versatile Medium* (Spi.). North Light Books.
- Brooker, S. (2010). *Portrait Painting Atelier: Old Master Techniques and Contemporary Applications*. Watson-Guption.
- Brown, H. (2003). *Confessions of a Starving Artist: the Art and Life of Harley Brown* (1ST ed.). North Light Books.
- Brown, H., & Lehrman, L. B. (2004). *Harley Brown's Eternal Truths for Every Artist*. International Artist Publishing. Retrieved from <http://www.amazon.com/dp/1929834314>
- Buechner, T. S. (2005). From Life . . . or Photographs? *American Artist*, 69(760), 36–41.
- Carlyle, L. A. (2001). *The artist's assistant : oil painting instruction manuals and handbooks in Britain 1800-1900, with reference to selected Eighteenth-century sources*. Retrieved from <http://www.bcin.ca/Interface/openbcin.cgi?submit=submit&Chinkey=196133>
- Cateura, L. (1995). *Oil Painting Secrets From a Master*. Watson-Guption.
- Chapman, H. (1990). *Rembrandt's Self-Portraits: A Study in Seventeenth Century Identity*. Princeton Univ Press.
- Clark, T. J. (1984). Using soft and hard edges in watercolor. *American Artist*, 48(505), 48–51.

References

- Clarke, M., Hampton, A., Obermeyer, M., Short, K., Sonneman, A., & Swimm, T. (2012).
The Art of Painting Landscapes, Seascapes, and Skyscapes in Oil & Acrylic: Discover
simple step-by-step techniques for painting an array of outdoor scenes. Walter Foster.
- Colahan, C. (1919). Max Meldrum: His Art and Views. Alexander McCubbin.
- Collins, N. (2004). Dependable Rules for Using Complementary Colors. *American Artist*,
68(740), 34–39.
- Connolly, A. (2012). *Painting Portraits*. Crowood Press.
- Cooke, H. L. (1972). *Painting Techniques of the Masters: Painting Lessons from the Great
Masters, Revised and Enlarged Edition (Revised & enlarged.)*. Watson-Guptill
Publications.
- Creevy, B. (1999a). *The Pastel Book*. Watson-Guptill.
- Creevy, B. (1999b). *The Oil Painting Book: Materials and Techniques for Today's Artist
(New edition.)*. Watson-Guptill.
- Dews, P. (2003). *The Painter's Workshop: Creative Composition & Design (1st Edition.)*.
North Light Books.
- Dixon, H. (2009). *Paint Amazing Watercolors from Photographs*. North Light Books.
- Doherty, M. S. (2003). Using Charts to Achieve Color Harmony. *American Artist*, 67(736),
22–93.
- Doherty, M. S. (2007). Dissecting the Painting Process. *Work of C. Pugliese*, 71(782), 66–
71.
- Eagle, E. (2013). *Pastel Painting Atelier: Essential Lessons in Techniques, Practices, and
Materials*. Watson-Guptill.

References

- Eastlake, S. C. L. (2001). *Methods and Materials of Painting of the Great Schools and Masters*. Dover Publications.
- Ekperigin, N. (2008). Placing the Focal Point in a Painting. *American Artist*, 72(789), 8–9.
- Elliott, V. (2007). *Traditional Oil Painting: Techniques and Concepts of Classical Realism from the Renaissance to the Past*. Watson-Guptill. Retrieved from <http://www.amazon.ca/dp/0823030660>
- Faigin, G. (2008). *The Artist's Complete Guide to Facial Expression* (2nd ed.). Watson-Guptill.
- Fig, J. (2009). *Inside the Painter's Studio* (1st ed.). Princeton Architectural Press.
- Freeland. (2007). Portraits in painting and photography. *Philosophical Studies*, 135(1), 95–109.
- Freeland, C. (2002). *But Is It Art?: An Introduction to Art Theory* (New Ed.). Oxford University Press, USA.
- Friel, M. (2010). *Still Life Painting Atelier: An Introduction to Oil Painting*. Watson-Guptill.
- Gage, J. (2000). *Color and Meaning: Art, Science, and Symbolism* (1st ed.). University of California Press.
- Gombrich, E. H. (2000). *Art and Illusion*. Princeton University Press.
- Gombrich, E. H. (2006). *Story of Art* (1st ed.). Phaidon. Retrieved from <http://www.amazon.com/dp/0465004458>
- Gombrich, E. H., Hochberg, J., & Black, M. (1973). *Art, Perception, and Reality*. The Johns Hopkins University Press. Retrieved from <http://www.amazon.com/Art-Perception-Reality-Thalheimer-Lectures>

References

- Gombrich, E. H., Hochberg, J. E., & Black, M. (1972). *Art, Perception and Reality*: John Hopkins University Press, 1977. - x, 136 s. JHU Press.
- Gorst, B. (2004). *The Complete Oil Painter: The Essential Reference for Beginners to Professionals*. Watson-Guptill.
- Gottlieb, L., & Doherty, M. S. (2006). Painting the Same Subject to Learn More About Light & Color. *American Artist*, 70(770), 50–57.
- Graham-Dixon, A. (2007). Dutch Portraits: The Age of Rembrandt and Frans Hals. *Sunday Telegraph Reviews*, (01-07.2007).
- Gurney, J. (2009). *Imaginative Realism: How to Paint What Doesn't Exist* (1st ed.). Andrews McMeel Publishing.
- Gurney, J. (2010). *Color and Light: A Guide for the Realist Painter* (Original.). Andrews McMeel Publishing.
- Gury, A. (2009). *Alla Prima: A Contemporary Guide to Traditional Direct Painting* (1st ed.). Watson-Guptill.
- Haines, J. (2010). *Colour & Light*. Search Press.
- Hall, E. T. (1990). *The Hidden Dimension*. Anchor.
- Hall, M. B. (1992). *Color and Meaning: Practice and Theory in Renaissance Painting*. Cambridge University Press.
- Hampton, A., Loughlin, J., & al, et. (2003). *The Art of Oil Painting: Discover all the techniques you need to know to create beautiful oil paintings*. Walter Foster.
- Hampton, A., Loughlin, J., & Swimm, T. (2002). *Oil Painting Step by Step* (1St Edition.). Walter Foster.
- Hardy, D. (2002). Taking the Temperature. *American Artist*, 66(721), 54–59.

References

- Hecht, H., Schwartz, R., & Atherton, M. (2003). Looking into Pictures: An Interdisciplinary Approach to Pictorial Space. The MIT Press. Retrieved from <http://www.amazon.com/Looking-into-Pictures-Interdisciplinary-Pictorial>
- Hermens, E. (1998). Looking through paintings : the study of painting techniques and materials in support of art historical research. Retrieved from <http://www.bcin.ca/Interface/openbcin.cgi?submit=submit&Chinkey=172172>
- Hockney, D. (2001). Secret Knowledge: Rediscovering the Lost Techniques of the Old Masters. Studio. Retrieved from <http://www.amazon.com/Secret-Knowledge-Rediscovering-Techniques-Masters>
- Hockney, D. (2009). Hockney on Art: Conversations with Paul Joyce (First Edition Thus.). Little, Brown Book Group.
- Hockney, D., & Falco, C. (2005). Quantitative Analysis of Qualitative Images. Retrieved from <http://csdl2.computer.org.proxy.lib.sfu.ca/persagen/DLAbsToc.jsp?resourcePath=/dl/proceedings/&toc=comp/proceedings/dicta/2005/2467/00/2467toc.xml&DOI=10.1109/DICTA.2005.67>
- Hockney, D., & Falco, C. M. (2000). Optical Insights into Renaissance Art. Optics and Photonics News, 11(7), 52–59.
- Hodge, S. (2008). The Complete Guide To Painting Pictures From Photos (2nd Revised edition.). David & Charles.
- Hodge, S. (2010). Great Artists. Quercus.
- Howard, K. (1999). The Complete Artist. Clb.
- Januszczak, W. (1982). Understanding art. Macdonald.

References

- Januszczak, W. (1993). Techniques of the World's Great Painters. (W. Kanusezak & W. Kanuszczak, Eds.). Book Sales.
- Jasper, C. (2005). Powercolor: Master Color Concepts for All Media. Watson-Guptill.
Retrieved from <http://www.amazon.com/Powercolor-Master-Color-Concepts-Media>
- Kandinsky, W. (1979). Point and Line to Plane. Dover Publications.
- Kandinsky, W. (2011). Concerning the Spiritual in Art. Empire Books.
- Kessler, M. (1992). Painting Better Landscapes: Specific Ways to Improve Your Oils (Reprint.). Watson-Guptill.
- Kessler, M. (2012). Color Harmony in your Paintings. North Light Books.
- Kinstler, E. R. (1981). Painting Faces, Figures and Landscapes (First Edition.). Watson-Guptill Publications Inc., U.S.
- Kinstler, E. R. (1987). Painting Portraits. (S. E. Meyer, Ed.) (Revised.). Watson-Guptill Publications.
- Klee, P. (1959). The inward vision; watercolors, drawings, writings. (First edition.). H. N. Abrams.
- Klee, P. (1961). Paul Klee: The Thinking Eye (1st ed.). G. Wittenborn, NY/Lund Humphries, London.
- Klee, P., & (editor), F. K. (1968). The Diaries of Paul Klee, 1898-1918. University of California Press.
- Konecni, V. J. (1991). Portraiture: An Experimental Study of the Creative Process. Leonardo, 24(3), 325–328.
- Kreutz, G. (1997). Problem Solving for Oil Painters: Recognizing What's Gone Wrong and How to Make it Right. Watson-Guptill.

References

- Langford, M. (2000). *Basic Photography* (7th ed.). Focal Press.
- Lawrence-Lightfoot, S. (2005). Reflections on Portraiture: A Dialogue Between Art and Science. *Qualitative Inquiry*, 11(1), 3–15.
- Lawrence-Lightfoot, S., & Davis, J. H. (1997). *The Art and Science of Portraiture* (1st ed.). Jossey-Bass.
- Leffel, D. A. (2009). *Self-Portraits* (1st Edition.). Bright Light Publishing.
- Leffel, D. A. (2012). *An Artist Teaches: Reflections on the Art of Painting* (4th ed.). Bright Light Publishing.
- Leland, N. (1998). *Exploring Color: How to Use and Control Color in Your Painting* (Revised.). North Light Books.
- Leland, N. (2000). *Exploring Color Coloring Book (Workbook.)*. Moonflower Books.
- Leland, N. (2006). *The New Creative Artist* (Revised.). North Light Books.
- Leland, N. (2008). *Confident Color: An Artist's Guide To Harmony, Contrast And Unity* (Spi.). North Light Books.
- Leyton, M. (2006). *The Structure of Paintings* (1st ed.). Springer. Retrieved from <http://www.amazon.com/Structure-Paintings-Michael-Leyton>
- Lindsay, K. C., & Vergo, P. (1994). *Kandinsky: Complete Writings On Art*. Da Capo Press.
- Lopes, D. M. (2003). The Aesthetics of Photographic Transparency. *Mind*, 112(447), 433–448.
- Lopes, D. M. (2004). Directive Pictures. *Journal of Aesthetics and Art Criticism*, 62(2), 189–196.
- Lopes, D. M. (2005). *Sight and Sensibility: Evaluating Pictures*. Oxford University Press, USA. Retrieved from <http://www.amazon.com/gp/product/0199277346>

References

- Macpherson, K. (2000). *Fill Your Oil Paintings with Light & Color*. North Light Books.
- Macpherson, K. (2009). *Landscape Painting Inside & Out*. North Light Books.
- Mann, H. (1933). *The Technique of Portrait Painting: A Complete & Detailed Guide to the Handling, Composition & Lighting of Portraits in Oils*. J.B. Lippincott company.
- Monahan, P., Seligman, P., & Clouse, W. (2005). *Art School: A Complete Painter's Course*. Hamlyn.
- Ocvirk, O., Stinson, R. E., Wigg, P. R., Bone, R. O., & Cayton, D. L. (2001). *Art Fundamentals with Core Concepts CD-ROM (9th ed.)*. McGraw-Hill Humanities/Social Sciences/Languages. Retrieved from <http://www.amazon.com/Art-Fundamentals-Core-Concepts-CD-ROM>
- Paul, T. (2003). *How to Mix & Use Color: The Artist's Guide to Achieving the Perfect Color in Every Medium*. North Light Books.
- Paul, T. (2005). *How to Create Texture in Your Paintings: The Artist's Guide to Using Paint Effectively (Second Impression.)*. New Holland.
- Paul, T. (2006). *How to Paint from Photographs: The Artist's Guide to Techniques and Materials*. New Holland.
- Perricelli, L. M. (2006). Teaching Values. *American Artist*, 70(767), 12–19.
- Powell, W. F. (2004). *Color Mixing Recipes: Mixing recipes for more than 450 color combinations*. Walter Foster.
- Powell, W. F. (2006). *Color Mixing Recipes for Portraits: More than 500 Color Combinations for skin, eyes, lips & hair*. Walter Foster.

References

- Powell, W. F. (2012). 1,500 Color Mixing Recipes for Oil, Acrylic & Watercolor: Achieve precise color when painting landscapes, portraits, still lifes, and more (Spi.). Walter Foster.
- Price, L. S. (2002a). The portrait process. *American Artist*, 66(714), 26–75.
- Price, L. S. (2002b). Using Technology to Create Classical Paintings. *American Artist*, 66(725), 18–25.
- Price, L. S. (2004). Creating Dynamic Tension. *Work of M. Mason*, 68(744), 44–49.
- Price, L. S. (2006). Painting Loose and Free. *Watercolor*, 12(47), 116–127.
- Price, L. S. (2007). Attracting the Viewer's Eye With Skilful Simplicity. *Work of D. Higgins*, 71(782), 50–57.
- Pumphrey, R. (1996). *Elements of Art*. Prentice Hall College Div.
- Quiller, S. (1999). *Painter's Guide to Color*. Watson-Guptill.
- Quiller, S. (2002a). *Color Choices: Making Color Sense Out of Color Theory* (New edition.). Watson-Guptill.
- Quiller, S. (2002b). *Acrylic Painting Techniques* (New edition.). Watson-Guptill.
- Quiller, S., & Whipple, B. (1983). *Water Media Techniques* (First Printing.). Watson-Guptill.
- Reid, C. (1985). *Pulling Your Paintings Together* (First Edition.). Watson-Guptill.
- Reid, C. (1987). *Painting What You Want to See*. Watson-Guptill.
- Reid, C. (1989). *Portrait Painting in Watercolor*. Watson-Guptill.
- Reid, C. (1991). *Painting by Design: Getting to the Essence of Good Picture-Making* (1st ed.). Watson-Guptill.
- Reid, C. (2000). *The Natural Way to Paint: Rendering the Figure in Watercolor Simply and Beautifully*. Watson-Guptill.

References

- Rosand, D. (1981). Titian and the Eloquence of the Brush. *Artibus et Historiae*, 2(3), 85–96.
- Rosand, D. (1990). *Painting in Cinquecento Venice: Titian, Veronese, Tintoretto*. Yale University Press.
- Rothko, M. (2006a). *Writings on Art*. (M. Lopez-Remiro, Ed.) (First Edition.). Yale University Press.
- Rothko, M. (2006b). *The Artist's Reality: Philosophies of Art*. (C. Rothko, Ed.). Yale University Press.
- Sallis, J. (1998). *Shades - Of Painting at the Limit* (First Edition.). Indiana University Press.
- Sallis, J. (Ed.). (2012). *Paul Klee: Philosophical Vision- From Nature to Art*. McMullen Museum Of Art.
- Sanden, J. (2004). *Portraits from Life in 29 Steps*. North Light Books.
- Sanden, J. H., & Sanden, E. (1999). *Portraits from Life*. North Light Books. Retrieved from <http://www.amazon.com/dp/0891349022>
- Saper, C. (2001). *Painting Beautiful Skin Tones with Color & Light*. North Light Books.
- Saper, C. (2012). *Classic Portrait Painting in Oils: Keys to Mastering Diverse Skin Tones*. North Light Books.
- Saunders, D., Spring, M., & Meek, A. (Eds.). (2013). *The Renaissance Workshop: The Materials and Techniques of Renaissance Art*. Archetype Books.
- Schmid, R. (1973). *Richard Schmid Paints the Figure; Advanced Techniques in Oil*. (First Edition.). Watson-Guption Pubns.
- Schmid, R. (1975). *Richard Schmid Paints Landscapes: Creative Techniques in Oil* (First Edition.). Watson-Guption Publications.
- Schmid, R. (2003). *Richard Schmid A Retrospective 2003*. Butler inst. of Art stove prairie.

References

- Schmid, R. (2004). *Alla Prima: Everything I Know About Painting*. Stove Prairie Press.
- Schreyer, B. (1997). Making shadows important. *American Artist*, 61(654), 50–53.
- Seligman, P. (2002a). *Oils: Step-by-Step Art School*. Hamlyn.
- Seligman, P. (2002b). *Step By Step - Oils*. Bounty Books.
- Seligman, P. (2003). *History & Techniques of the Great Masters Pissarro*. Eagle Editions.
- Seligman, P. (2004). *Painting Light & Shade*. North Light Books.
- Sheppard, J. (1983). *How to Paint Like the Old Masters (First Edition.)*. Watson-Guption.
- Shirley, R. (2004). *Oil Paintings From the Landscape: A Guide for Beginners*. Guild of Master Craftsman.
- Shirley, R. (2012a). *Portrait Painting in Oil: 10 Step by Step Guides from Old Masters*. CreateSpace Independent Publishing Platform.
- Shirley, R. (2012b). *Why do My Skin Tones Look Lifeless? Plus 25 Solutions to Other Portrait Painting Peeves*. CreateSpace Independent Publishing Platform.
- Solomon, S. J. (1910). *The Practice of Oil Painting*. Seeley Service.
- Speed, H. (1987). *Oil Painting Techniques and Materials*. Dover Publications.
- Spring, M. (2011). *Studying Old Master Paintings: Technology and Practice*. Archetype Books.
- Stewart, E. (1997). Using edges and stroke quality to control space. *American Artist*, 61(659), 26–31.
- Stinson, R. E., Wigg, P. R., Bone, R. O., Cayton, D. L., & Ocvirk, O. G. (1997). *Art Fundamentals: Theory and Practice*. (O. G. Ocvirk, Ed.) (8th ed.). McGraw-Hill College.
- Thompson, D. (2010). Maintaining Master Practices. *American Artist*, 74(810), 38–47.

References

- Thompson, D. V. (1936a). *The Materials of Medieval Painting* (First Edition.). George Allen & Unwin Ltd.
- Thompson, D. V. (1936b). *The Practice of Tempera Painting* (Later Printing.). Yale University Press.
- Thompson, D. V. (1956a). *The materials and techniques of medieval painting*. Courier Dover Publications.
- Thompson, D. V. (1956b). *The Materials and Techniques of Medieval Painting*. Dover Publications.
- Thompson, D. V. (1962). *The Practice of Tempera Painting: Materials and Methods*. Dover Publications.
- Walls, W. M. T. (2004). The Value of Shadows. *American Artist*, 68(739), 40–43.
- Weber, M. C. (2009). *Bold Strokes: Dynamic Brushwork In Oils And Acrylics*. North Light Books.
- Weber, M. C. (2010). *Brushwork Essentials* (1st ed.). North Light Books.
- Whyte, M. (2005). *An Artist's Way Of Seeing*.
- Whyte, M. (2011). *Painting Portraits and Figures in Watercolor* (Original.). Watson-Guptill.
- Willard, C. (1999). Paintings from photographs. technique of H. Gordon, 63(680), 12.
- Willard, C. (2000a). Mixing it up: palettes. *American Artist*, 64(698), 10–13.
- Willard, C. (2000b). Copying paintings at museums. *American Artist*, 64(697), 62–70.
- Willard, C. (2001a). The extended possibilities of a limited palette. *American Artist*, 65(709), 10–13.
- Willard, C. (2001b). Splitting the color wheel for luminosity and energy. work of C. Elliott, 65(704), 12–15.

References

Willard, C. (2001c). Color lurking in the shadows. *American Artist*, 65(712), 12–15.

Willard, C. (2003). Painting Skin Tones. *American Artist*, 67(737), 12–15.

Willard, C. (2004). Methods & Materials: Simply Exciting Compositions. *American Artist*, 68(746), 12–16.

Wyk, H. V. (1970). *Acrylic Portrait Painting (First Edition.)*. Watson-Guption.

Wyk, H. V. (2000). *Color Mixing the Van Wyk Way: A Manual for Oil Painters*. Art Instruction Associates.

Appendix A: Painterly System Documentation and XML language

Thinker Painter Documentation of Parameters

This document specifies every available parameter for every component of the ThinkerPainter framework, including the parameter's name, its XML representation, its default value (enclosed within the XML representation if applicable), and some extra information specified in brackets:

Parameter Name: `<xml-representation attribute="defaultAtr">defaultValue</xml-representation>` (extraInfo, moreExtraInfo) **ParameterType** [min, max]

If the word “optional” appears in the extra information section, then the parameter need not be explicitly specified and when not specified will supply its default value; if the symbol “+” appears, more than one instance of this parameter is permitted.

Root Element: `<painting>`

Reference: `<reference>`

Specifies how the reference image should be read.

Max Width: `<max-width>0</max-width>` (optional)

Specifies a maximum width for the reference image; if the reference image is wider than this value, it will be scaled down to match it (maintaining its aspect ratio) at the beginning of the run. A value of ‘0’ (the default) imposes no limits on reference image size.

Max Height: `<max-height>0</max-height>` (optional)

Specifies a maximum width for the reference image; if the reference image is wider than this value, it will be scaled down to match it (maintaining its aspect ratio) at the beginning of the run. A value of ‘0’ (the default) imposes no limits on reference image size.

Base Resolution: <base-resolution>1.0</base-resolution> (optional)

Specifies a float value corresponding to the resolution of the reference image. All other resolutions in the config file act as multipliers to this base resolution; for example, if the reference image is 250x250 pixels and is assigned a base resolution of 1.0 (the default), then specifying a resolution of 2.0 somewhere else in the config corresponds to 500x500 pixels.

Output: <output>

Specifies how the final painting should be written to a file.

Resolution: <resolution>baseResolution</resolution> (optional)

The resolution at which the painting should be rendered and written, specified as a multiple of the base resolution.

File Type: <file-type>png</file-type> (optional)

The encoding format used to write the output image to a file. Supports PNG and JPEG, among others.

File Extension: <file-extension>png</file-extension> (optional)

The extension appended to the end of the output image to define its file format. Should generally be the same as the file type.

Image Debug: <image-debug>>false</image-debug> (optional)

Specifies whether or not various debug features should be turned on during the run, which output useful information about the run to a debug folder in the same directory as the output entitled outputname_debug.

Background Color: <background>1.0 1.0 1.0</background> (optional)

Specifies the initial color of the output canvas as a floating point RGB triplet. White by default.

Region Map: <region-map> (optional)

Specifies a PSD file containing layers that define different sections of the reference image, labeled based on the name of each layer. Higher layers are superimposed on top of lower ones, meaning white pixels in a higher layer will cancel out white pixels in every layer below it and only the highest layer's region will contain the pixel.

As with every other sort of alpha map in the program, layers should use greyscale colors; black pixels are not in the region, while white pixels are fully in the region. Note that while layers should only use greyscale colors, they must be stored within the Photoshop file as RGB 8 bits/channel (the default) for compatibility reasons.

Region maps are automatically scaled to the resolution and aspect ratio of the reference image.

Map File: <file>path/to/file.psd</file>

Specifies the path at which the region map file is located.

Alpha Map: <alpha-map> (optional)

Specifies any number of greyscale maps that define different arbitrary regions of the image, to be referenced by various parameters elsewhere in the config file.

Alpha maps are automatically scaled to the resolution and aspect ratio of the reference image.

Image Files: <file id="nameOfMap">path/to/file.extension</file> (optional, +)

Specifies an image file to load as an alpha map, and a unique identifier to assign to this alpha map so that it can be referenced by other parameters.

PSD Files: <file>path/to/file.psd</file> (optional, +)

Specifies a PSD file containing layers that contain different alpha maps. Unlike the region map, these layers are NOT superimposed, and therefore their ordering is unimportant and pixels may overlap. The ID of each layer is assigned based on the layer's name.

Pass Elements: <pass> (+)

One painting can have any number of passes; each lays successive layers of paint on the output canvas.

Reference Resolution: <reference-resolution>baseResolution</reference-resolution> (optional)

Setting this to some resolution other than the base resolution will generate a resized instance of the source image for use during this pass only. This allows very expensive passes to improve their running time by reducing the number of pixels on which they need to operate.

Mat: <mat> (optional)

A pass's mat selects a specific area of the canvas to which the pass will be applied; other areas should be ignored. In this way passes can act only on small, arbitrarily-defined areas of the canvas to add extra detail.

File: <file>path/to/file.extension</file>

With this element, a mat may be specified by loading a greyscale alpha map from a file. See 'Alpha Map' for more information about alpha maps.

File: <id>idOfAlphaMap</id>

With this element, a mat may be specified using the ID of a map already specified in the Alpha Map element. See 'Alpha Map' for more information about alpha maps.

Include Regions: <include-regions>firstRegion secondRegion</include-regions> (optional)

This element creates a mat for the pass by specifying any number of regions from the region map to include in the pass. If a Mat element is also specified, the resulting pass mat is defined as the intersection between the Mat and the included regions.

Thinker: <thinker class="ThinkerType">

Analyzes the current canvas as well as the source image and creates PaintAction objects representing its high-level intent to paint on a specific region of the canvas in a certain way. See the Thinker section for details about specific Thinker classes.

Painter: <painter class="PainterType">

Receives a PaintAction object from a Thinker, and decides how to render it to the canvas. See the Painter section for details about specific Painter classes.

Concern: <concern class="ConcernType"> (optional)

A modular, lightweight object that can be plugged into any Thinker to suggest modifications to the PaintActions it produces. See the Concern section for details about specific Concern classes.

Palette: <palette class="package.name.PaletteType"> (optional)

Palettes are actually a type of Concern, and behave as such; this special syntax allows them to be specified conveniently within a Pass element.

End Condition: <end-condition class="EndConditionType"> (optional)

Checks after every stroke to decide whether or not the painting is finished. By default there is no end condition, leaving the Thinker to decide when to stop painting. See the EndCondition section for details about specific EndCondition classes.

Thinkers

BlobThinker: <thinker class="BlobThinker">

This class implements a lightness-based subdivision technique to produce hierarchical blob maps of the reference image, using these blobs to define the areas of the PaintActions it creates. Because this particular Thinker never refers to color values in the reference image, it is perfectly permissible to use greyscale ones with this class.

Blur Size: <blur-size relative="false">2</blur-size> (optional)

ConstantFloat Range: [0.0, unlimited]

The size of the Gaussian blur applied to a copy of the reference image, to which this Thinker will refer for information about said reference image. Blurring helps to remove noise, making the edges of blobs more regular (but also reducing the accuracy of edges in the reference image, like the boundary between a face and the background). If the “relative” attribute is set to false, this size is specified in absolute pixels; otherwise, it is specified as a fraction of the size of the reference image’s width (ie: If the reference image is 500 pixels wide and blur-size is 0.2, it would apply a Gaussian Blur with a kernel of size 100).

Blobs From Regions: <blobs-from-regions>>false</blobs-from-regions> (optional)

ConstantBoolean

If set to true, the initial blobs this Thinker creates will be identical in shape to the regions of the reference image defined by the Region Map element. Otherwise, all blobs are generated procedurally.

Number of Leaf Blobs: <num-leaf-blobs>1000.0</num-leaf-blobs> (optional)

ConstantFloat Range: [0.0, unlimited]

The target number of leaf blobs that should exist at the end of the subdivision process. This is a theoretical value derived from the assumption that every leaf blob will be of the same size; in practice this is not the case, and so the actual number of leaf blobs will be lower.

Very high numbers (25000+) will yield many very small blobs and therefore many small, accurate PaintActions. Very low numbers will yield few larger blobs and provide a small number of large, more gross PaintActions.

Density: `<density>0</density>` (optional)

FloatParameter **Range:** [0.0, 1.0]

How densely strokes should be laid within an area. This value is inserted directly into PaintActions without processing, leaving its implementation to the Painter.

Detail: `<density>0</density>` (optional)

FloatParameter **Range:** [0.0, 1.0]

How precisely strokes should conform to PaintAction areas. This value is inserted directly into PaintActions without processing, leaving its implementation to the Painter.

Painters

StrokePainter: `<painter class="StrokePainter">`

This Painter object uses the specified GLBrush object to render strokes to the canvas. These strokes are created based on the Hertzmann algorithm, which in general draws strokes that run perpendicular to the image gradient.

This particular implementation attempts to keep strokes within the area of the PaintAction by limiting the opacity of each stroke to the smallest alpha value that stroke passes over; it also attempts to avoid laying too many strokes in the same place by limiting the total opacity of each stroke passing through a given grid tile to the value of the opacity parameter.

Grid Size: `<grid-size relative="true">0.05</grid-size>` (optional)

FloatParameter **Relative** **Range:** (0.0, 1.0], **Absolute** **Range:** (0, min(referenceWidth, referenceHeight)]

The size of the grid tiles into which the reference image is divided; one stroke is laid per grid tile. If relative is set to true, this value represents a fraction of the reference image's dimensions; otherwise, it represents an absolute pixel size.

Brush Scale: `<brush-scale">1.0</brush-scale>` (optional)

FloatParameter **Range:** (0.0, unlimited]

The size of the strokes laid by this Painter, relative to gridSize. A value of 1.0 lays strokes with a diameter equal to the grid size.

Opacity: <opacity>1.0</opacity> (optional)

FloatParameter **Range:** [0.0, 1.0]

The maximum opacity for strokes laid by this painter.

Curvature: <curvature>0.5</curvature> (optional)

FloatParameter **Range:** [0.0, 1.0]

The extent to which strokes should change direction along their length (0 meaning not at all, 1.0 representing perfect conformance to the image gradient).

Minimum Stroke Length: <min-stroke-length>1</min-stroke-length> (optional)

ConstantInt **Range:** [1, maxStrokeLength]

The minimum possible number of control points per stroke (and therefore the minimum number of grid tiles the stroke should traverse).

Maximum Stroke Length: <max-stroke-length>5</max-stroke-length> (optional)

ConstantInt **Range:** [minStrokeLength, unlimited]

The maximum possible number of control points per stroke (and therefore the maximum number of grid tiles the stroke should traverse).

Seed: <seed>1</seed> (optional)

ConstantLong **Range:** Unlimited

Value used to generate pseudo-random numbers in instances where multiple possible seeds exist for one grid tile.

Gradient: <gradient> (optional)

GradientParameter

Image gradient used to determine the direction of strokes; in general, strokes attempt to run perpendicular to this gradient. By default, this is generated procedurally.

Error: <error> (optional, FloatParameter)

Range: [0.0, 1.0]

Alpha channel representing the difference in relative lightness between the current canvas and the reference image, used to locate seed points for strokes. By default, this is generated procedurally.

Brush: <brush class=“package.name.BrushClass”> (optional)

The GLBrush object used to render each stroke. See the older documentation for more information about brushes.

FillPainter: <painter class=“FillPainter”>

This Painter object fills the areas of the PaintAction objects it receives in their entirety using a solid color specified by these PaintAction objects ‘getColor’ method.

End Conditions

ActionCountCondition: <end-condition class=“ActionCountCondition”>

This EndCondition permits the specification of a maximum number of PaintActions to make; once the number of actions taken in one pass meets or exceeds this maximum, the pass will end.

This class may also serve as an unlimited end condition by specifying a max action count of 0; when no end condition is specified, ActionCountCondition is used in this way.

Maximum Actions: <max-actions>0</max-actions> (optional)

ConstantInt Range: [0-unlimited]

The maximum number of paint actions to permit. A value of 0 indicates an unlimited number of paint actions.

Appendix B: Table of Artistic Practice Source Material

We detail our source material that we used to systematically categorize our Painterly Soft Rules in Section 3.2. Materials came from approximately 200 modern and historical art practice documents including from artists such as Klee, Kandinsky and Hockney. Reference cites are followed by tags depicting where they were used with the Chapter 3.2 sections including:

- 3.2.1 Tonal
- 3.2.2 Colour
- 3.2.3 Textural Highlights (Highlights)
- 3.2.4 Shadows and halftone 'core' (Shadow)
- 3.2.5 Shapes, Edges and Centre of Interest (Shapes)
- 3.2.6 Brush Strokes (Strokes)
- 3.2.8 BackGrounds (BG)
- 3.2.7 Working from Photograph as Source Image (Photo),
- 3.2.9 Process Plan (Process)

Source Material (Alphabetical with Section Tags):

(Albala, 2009) Tonal, Colour, Shape, Shadow, Stroke, Process
(Aristides, 2006, 2008, 2011) Tonal, Colour, Shape, Shadow, Stroke, Process
(Arnheim, 1986, 1996, 2004a, 2004a, 2004b)
(Brown, 2003) Tonal, Colour, Shapes, Strokes, Process
(Bahr, 2011) Tonal, Colour
(Baxter, 2000a, 2000b, 2001, 2003) Tonal, Colour, Shape
(Berger Jr, 1998) Shapes, Process
(Bimler, 2005) Tonal, Colour
(Brady, 2008) Colour, Shape, Highlights, Process
(Brooker, 2010) Tonal, Strokes, Backgrounds
(Brown, 2003) Tonal, Colour, Shadow, Shapes, Strokes, Process
(Brown & Lehrman, 2004) Tonal, Colour, Highlights, Shadow, Shapes, Strokes, Process
(Buechner, 2005) Photo
(Carlyle, 2001) Colour, Stroke, Shape, Process
(Chapman, 1990) Shadow, Shapes, Process
(Clark, 1984) Shape
(Colahan, 1919) Tonal, Colour, Shape
(Collins, 2004) Colour

Appendix B: Table of Artistic Practice Source Material

(Connolly, 2012) Highlights, Shapes, BG, Process
(Cooke, 1972) Colour, Shape, Process
(Creevy, 1999a, 1999b) Colour, Process
(Thompson, 2010) Colour
(Thompson, 2010, 1956a) Tonal, Colour, Shadow
(Thompson, 1936a, 1936b, 1956a, 1956b, 1962) Tonal, Colour, Shape, Shadow, Stroke, Process
(Dews, 2003) Shape, Process
(Dixon, 2009) Colour, Shape, Photo, Process
(Doherty, 2003, 2007; Gottlieb & Doherty, 2006) Tonal, Colour, Process
(Gombrich, 2000, 2006; Gombrich et al., 1972, 1973) Tonal, Colour, Shape, Shadow, Process
(Eagle, 2013) Tonal, Colour, Shape
(Eastlake, 2001) Colour, Shape, Highlights, Process
(Ekperigin, 2008) Shape
(Elliott, 2007) Tonal, Colour, Shapes
(Faigin, 2008) Shapes
(Fig, 2009) Tonal, Colour, Shape, Shadow, Stroke, Process
(Freeland, 2007, 2002) Tonal, Colour, Shapes, Process
(Friel, 2010) Tonal, Colour, Shape
(Gage, 2000) Tonal, Colour
(Gorst, 2004) Colour, Shape, Process
(Graham-Dixon, 2007) Shape, Shadow, Stroke, Process
(Gurney, 2009, 2010) Tonal, Colour
(Gury, 2009) Colour, Strokes, Shapes, Process
(Haines, 2010) Tonal, Colour, Shadow
(Clarke et al., 2012; Hampton et al., 2002, 2003) Colour, Shape, Process
(Hardy, 2002) Colour
(Hecht et al., 2003) Colour, Shapes, Process
(Hermens, 1998) Shape, Stroke, Process
(Hockney & Falco, 2005, 2000; Hockney, 2001, 2009) Shape, Stroke, BG, Process
(Hodge, 2008, 2010) Photo, Stroke, Shape, Process
(Howard, 1999) Colour, Shape, Strokes, Process
(Sanden & Sanden, 1999; Sanden, 2004) Tonal, Colour, Shape. Photo, Process
(Januszczak, 1982, 1993) Tonal, Colour, Shape, Process
(Jasper, 2005) Colour
(Kandinsky, 1979, 2011) Colour, Stroke, Shape, Process
(Kessler, 1992, 2012) Colour, Shape, Process
(Kinstler, 1981, 1987) Tonal, Colour, Shape, Stroke, Process
(Klee & (editor), 1968; Klee, 1959, 1961; Sallis, 2012) Tonal, Colour, Shape, Stroke, Process
(Konecni, 1991) Shape, Stroke, Highlight, Process
(Kreutz, 1997) Colour, Shape, BG, Process
(Langford, 2000) Tonal, Shape
(Lawrence-Lightfoot & Davis, 1997; Lawrence-Lightfoot, 2005) Colour, Shape, Process
(Cateura, 1995; Leffel, 2009, 2012) Shape, Shadow, Stroke, Process

Appendix B: Table of Artistic Practice Source Material

(Leland, 1998, 2000, 2006, 2008) Tonal, Colour, Process
(Leyton, 2006) Shape
(Lindsay & Vergo, 1994) Colour, Stroke, Shape, Process
(Lopes, 2003, 2004, 2005) Process
(Hall, 1990, 1992) Tonal, Colour, Shapes, Process
(Macpherson, 2000, 2009) Colour, Shape
(Mann, 1933) Colour, Shape, Highlights
(Monahan et al., 2005) Tonal, Colour, Photo, Stroke, Process
(Ocvirk et al., 2001) Shape, Stroke
(Paul, 2003, 2005, 2006) Tonal, Colour, Photo, Stroke, Process
(Perricelli, 2006) Tonal
(Powell, 2004, 2006, 2012) Colour
(Price, 2002a, 2002b, 2004, 2006, 2007) Shape, Stroke, Process
(Pumphrey, 1996) Colour, Shapes, Stroke, Highlights, Process
(Quiller & Whipple, 1983; Quiller, 1999, 2002a, 2002b) Tonal, Colour
(Reid, 1985, 1987, 1989, 1991, 2000) Tonal, Colour, Shape, Stroke, Process
(Rosand, 1981, 1990) Tonal, Colour, Stroke, Process
(Rothko, 2006a, 2006b) Colour, Shape, Process
(Sallis, 1998, 2012) Colour, Shape, Process
(Saper, 2001, 2012) Colour, Shape, Photo, Stroke, Process
(Saunders et al., 2013) Tonal, Colour, Shape, Process
(Schmid, 1973, 1975, 2003, 2004) Colour, Shape, Stroke, Process
(Schreyer, 1997) Shadow
(Seligman, 2002a, 2002b, 2003, 2004) Tonal, Colour, Shape
(Sheppard, 1983) Colour, Shape, Stroke
(Shirley, 2004, 2012a, 2012b) Colour, Shape, Highlights, Process
(Solomon, 1910) Tonal, Colour, Shape, Process
(Speed, 1987) Tonal, Colour, Shape
(Spring, 2011) Colour, Shape, Process
(Stewart, 1997) Shape, Stroke
(Stinson et al., 1997) Colour, Shape, Highlight
(Walls, 2004) Shadow
(Weber, 2009, 2010) Strokes, Highlights, Shape
(Whyte, 2005, 2011) Tonal, Colour, Shape, Stroke, Process
(Willard, 1999, 2000a, 2000b, 2001a, 2001b, 2001c, 2003, 2004) Colour, Shadow, Shape, Shadow, Strokes, Photo, Process
(Wyk, 1970, 2000) Colour, Process

See the second section in the References for these specific Appendix B references.