

STYLIZED 2D FABRICATION OF NON-PHOTOREALISTIC IMAGES

A Thesis
Submitted to the Faculty
in partial fulfillment of the requirements for the
degree of

Doctor of Philosophy

in

Computer Science

by Athina Panotopoulou

Guarini School of Graduate and Advanced Studies
Dartmouth College
Hanover, New Hampshire

August, 2020

Examining Committee:

(chair) Emily Whiting, Ph.D.

Xing-Dong Yang, Ph.D.

Wojciech Jarosz, Ph.D.

Sylvain Paris, Ph.D.

F. Jon Kull, Ph.D.

Dean of the Guarini School of Graduate and Advanced Studies

ProQuest Number:28089689

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent on the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 28089689

Published by ProQuest LLC (2020). Copyright of the Dissertation is held by the Author.

All Rights Reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

Abstract

A current trend in computer graphics is the use of programmable tools that allow non-experts to engage in the design of physical prototypes. Within fabrication, one area of research focuses on non-photorealistic images which are stylized to depict a particular aesthetic quality or convey key information. In cases where authenticity is demanded or the images need to be manipulated, fabrication is necessary. Non-photorealistic image fabrication involves two challenges: identifying and abstracting key information during design and considering material restrictions during fabrication. This thesis showcases two examples for fabricating new types of non-photorealistic images, the first involving watercolors, and the second producing tactile images on microcapsule paper for improved shape understanding for the blind. The watercolor project is the first to fabricate watercolorized images and produces color-gradients using a wet-on-wet technique by limiting the usually unpredictable outcome of this medium. The tactile project implements a novel multi-projection line drawing design for the depiction of everyday 3D shapes, different from currently used designs. For the fabrication of watercolor images, we propose an image analysis pipeline which inputs an image, extracts key information elements such as color, gradients, and details, and combines them to design wood stamps produced by a laser cutter. We then paint the stamps with watercolors and press them on water-sprayed paper, following a carefully designed fabrication process. For the fabrication of the tactile images, we gather characteristics that convey shape information through touch from studies in cognitive psychology, human-computer interaction, and our own formative studies. We process a 3D pre-segmented object, compute a multi-projection line drawing of the shape's cross sections, and laser engrave the drawings as tactile images. We run user studies that show the usefulness of stylization on both image designs. These two pipelines highlight the importance of image and geometry processing and analysis, key approaches, in the two applications, one of an artistic and another of a utilitarian nature. These graphics approaches play the role of “painter” and “designer of tactile graphics” and allow non-experts accessibility to these challenging fabrication design spaces.

Acknowledgments

This work includes copyrighted images reprinted with permissions, as well as public domain images. All of them are acknowledged in the following subsection. Also, in the next subsection, people and organizations are acknowledged.

0.1 Images

- Figure 1.1

(a): Giovanni Battista Piranesi (1770), “etched line drawing,” Retrieved from factum-arte.com;

(b)–(d): ©Factum Arte, Madrid, Spain, (2010), “coffeepot,” Retrieved from factum-arte.com.

- Figure 1.2

(a)–(b): Georg Grosz (1923), Museo Thyssen-Bornemisza, Madrid, Spain, “Twilight,” Retrieved from museothyssen.org;

(c): Java (1911), Peabody Museum, Cambridge, MA, US, Photograph posted by Daderot, “shadow puppet,” Retrieved from wikimedia.org

- Figure 1.3

(a): Harold Cohen, 1984, Victoria and Albert Museum, London, UK, “computer generated ink,” Retrieved from collections.vam.ac.uk;

(b): Lindemeier (2011), “e-David: Non-Photorealistic Rendering using a Robot and Visual Feedback,” [245];

(c): Tresset and Leymarie, (2012), “Sketches by Paul the Robot,” [415].

- Figure 1.4
 - (a): Raffaello Sanzio da Urbino (1505), Kunsthistorisches Museum, Wien, Austria, Photograph posted by DcoetzeeBo, “Madonna in the Meadow,” Retrieved from [wikimedia.org](https://www.wikimedia.org);
 - (b)–(c): Reichinger et al. (2011), “High-Quality Tactile Paintings,” [346].
- Figure 1.5
 - (a)–(c): Galea et al. (2016), “Stippling with areal robots,” [138].
- Figure 1.6
 - Isenberg et al. ©IEEE (2013), “A systematic review on the practice of evaluating visualization,” [189].
- Figure 1.8
 - ©Cate Parr, “Rooney Mara, Oscars in Givenchy,” Retrieved from cateparr.com.
- Figure 1.9
 - ©Braille Authority of North America (2010), “Guidelines and Standards for Tactile Graphics,” Retrieved from brailleauthority.org.
- Figure 2.1
 - (a) - top: Karl Joseph Stieler (1820), Beethoven-Haus, Bonn, Germany, Photograph posted by Blight55, “Beethoven,” Retrieved from [wikimedia.org](https://www.wikimedia.org);
 - (a) - bottom: Iseringhausen et al. (2020), “Computational Parquetry,” [190];
 - (b) - top: Raffaello Sanzio da Urbino (1505), Kunsthistorisches Museum, Wien, Austria, posted Photograph by DcoetzeeBo “Madonna in the Meadow,” Retrieved from [wikimedia.org](https://www.wikimedia.org);
 - (b) - bottom: Reichinger et al. (2011), “High-Quality Tactile Paintings,” [346];
 - (c) - top: the Stanford Computer Graphics Laboratory (1994), “Stanford bunny,” [306];
 - (c) - bottom: Narayanan et al. (2019), “Visual Knitting Machine Programming,” [306].

- Figure 2.2
 - (a)–(c): ©David Bull (2015), “woodblock printing,” Retrieved from patreon.com/mokuhankan ;
 - (d): David Bull (2015), original by Hokusai, “The Great Wave off Kanagawa,” Retrieved from mokuhankan.com.
- Figure 2.3
 - (a): Photograph posted by UCT MOOCs (2020), “map of Africa,” Retrieved from wikimedia.org;
 - (b): ©Paul Bartlett Ré, “A book and two traveling exhibits for the blind and sighted,” paulre.org;
 - (c): ©Perkins School for the Blind Archives, Watertown, MA, USA, M. Kunz (1902), Retrieved from perkins.org;
 - (d): ©American Printing House for the Blind (2004), “World Maps,” aph.org.
- Figure 2.5

Winkenbach and Salesin, ©ACM (1994), “Computer-Generated Pen-and-Ink Illustration,” [451].
- Figure 2.6

Lindemeier (2018), “e-David: Non-Photorealistic Rendering using a Robot and Visual Feedback,” [245].
- Figure 3.1(a), Figure 3.2(a), Figure 3.11 Input first row, Figure 3.12 Input

©Joshua Greene (2018), Milton H. Greene, “Marilyn Monroe,” Retrieved from archiveimages.com.
- Figure 3.6
 - (a) first, fourth: Zhang et al. (2016), “Spatially variant defocus blur map estimation and deblurring from a single image” data set, [478];

(a) second: Free Pixels Talk (2018), “Zebra Backgrounds,” Retrieved from Wallpaperget.com;

(a) third: Didier Descouens (2012), “Church of Santa Maria Formosa,” Retrieved from wikimedia.org.

- Figure 3.11 input second row and Figure 3.13 input

Petr Kratochvil (2018), “Hibiscus Flower,” Retrieved from publicdomainpictures.net.

- Figure 3.11 input third row and Figure 3.14 input

Themis Nasopoulou (2018), “European Robin,” Retrieved from floraolympus.blogspot.com.

- Figure 3.11 input fourth row and Figure 3.15 input

Pexels (2018), “Lake Cottage House,” Retrieved from pexels.com.

- Figure 4.1

(a)–(b): ©Museum of Fine Arts, Boston, MA, US, “Pat Lyon at the Forge,” John Neagle, 1826–27, Accession Number: 1975.806, Henry H. and Zoe Oliver Sherman Fund, Retrieved from mfa.org;

(c)–(d): ©Museum of Fine Arts, Boston, MA, US, “Coffee pot,” Paul Revere, Jr., 1791, Accession Number: 95.1359, Bequest of Buckminster Brown, M.D., Retrieved from mfa.org.

- target 3D objects for User study 3 displayed in all Figures in Chapters 4.3– 4.4

Yu et al. (2019), “Partnet: A recursive part decomposition network for fine-grained and hierarchical shape segmentation,” [468].

- variations of 3D objects for User study 3 displayed in all Figures in Chapters 4.3– 4.4

designed by Colleen Kim, Boston University, (2019).

0.2 People – Organizations

The projects in this thesis were supported by: Adobe, National Science Foundation under Grant No. #1813319 and No. #1464267, and the Alfred P. Sloan Foundation: Sloan Research Fellowship.

I would like to thank the people of the committee. My advisor, Emily Whiting (BU, Professor), urged me to get involved in exciting research topics, shared her resources and ideas, and provided feedback for the thesis and defense writing. Wojciech Jarosz (Dartmouth, Professor) shared helpful writing resources and provided feedback during the thesis writing. Xing-Dong Yang (Dartmouth, Professor), the initiator of the tactile project, shared his ideas during the project meetings. Sylvain Paris (Adobe, Senior Principle Scientist), the external committee member, guided me with his expertise during the watercolor project.

Also, I would like to thank my non-committee collaborators during my graduate studies, including: Elena Balashova (Princeton University, PhD Student), Josh Urban Davis (Dartmouth, PhD Student), İlke Demir (DeepScale, Research Scientist), Bastien Durix (University of Toulouse, PhD Student), Amelie Fondevilla (Université Grenoble Alpes, PhD Student), Camilla Hahn (Bergische Universität Wuppertal, PhD Student), Evelyne Hubert (INRIA, Senior Research Scientist), Adam Kortylewski (Johns Hopkins University, PostDoc), Xinyi Le (The Chinese University of Hong Kong, PhD Student), Hanyi Lui (Dartmouth), Kathryn Leonard (Occidental College, Professor), Géraldine Morin (University of Toulouse, Professor), Tammy Qiu (BU Student), Ryan Schmidt (Research Scientist, Professor, University of Toronto), Bo Shi (Dartmouth), Dana Rahbani (University of Basel, PhD Student), Elissa Ross (Research Scientist, MESH Consultants), Nobuyuki Umetani (University of Tokyo, Professor), Charlie C.L. Wang (University of Manchester, Professor), Kathrin Welker (Helmut-Schmidt-Universität Hamburg, Professor), Te-Yen Wu (Dartmouth, PhD Student), Xiaoting Zhang (BU, PostDoc). Your diverse and unique research backgrounds and experience offered a lot to my development as a researcher.

I would like to thank all the labs and groups who supported me. First, the people in the Visual Computing Lab (VCL) at Dartmouth College: Shruti Agarwal, Benedikt Bitterli, Shaojie Jiao,

Liane Makatura, Christos Mousas, Srinath Ravichandran, and Gurprit Singh, for the memories we shared. Also, the people in the Shape Lab at Boston University: Xingjian Jessie Han, Anthea Yichen Li, Zezhou Sun, Benjamin Verdier, and Xiaoting Zhang, for being present whenever I asked their help. Finally, the people in the Image and Video Computing group (IVC) at Boston University for making me feel welcome.

I would like to thank people who I met during the work of this thesis. All of them in different ways contributed to the progress of it. Amber Percy (BU, Graduate Student, MFA Boston Access Consultant, National Braille Press Proofreader) with her experience in tactile graphics, gave feedback for our tactile image in all the stages of its development and tested the images before we run the studies. Elizabeth Tremmel (Dartmouth College Multilingual Graduate Student and Postdoctoral Writing Consultant) improved my academic writing and helped me write the tactile project and this thesis. Sarah M. Smith (Dartmouth Book Arts Workshop, Program Manager) and Katie Baldwin (Bookbinder and Printmaker) provided printing knowledge and resources. Colleen Kim (BU, Student), created the objects for the User Study 3 of the tactile project. Ayan Azmat (Dartmouth, Student) experimented with machinery for the fabrication of woodblocks. Forrester Cole (Google, Research Scientist) provided ideas on line drawing and perception. Joy Ding (BU, Visiting Student) for letting me guide her in her research project. Sunil Hadap (Adobe, Research Scientist) provided insight about watercolor techniques and suggested the Hibiscus flower to watercolorize. Oliver Wang (Adobe, Research Scientist) helped debug my defocus map estimation code. Alex Barnett (Dartmouth, Professor Dept. of mathematics) provided a review of the Printone paper before its submission. Margrit Betke (BU, Professor Dept. of Computer Science) highlighted the importance of tactile diagrams in books. Luis Carvalho (BU, Professor Dept. of Mathematics and Statistics) advised on statistical testing. Sam Ling (BU, Professor Dept. of Psychological and Brain Sciences) pointed us to the Recognition-by-components theory. James O'Malley (Dartmouth, Professor Dept. of Biomedical Data Science of Dartmouth Institute) advised on statistical testing and helped format the linear mixed model. The volunteer Dartmouth Art Students: Phoebe, Chicaiza, Iliana, Karen, gave me feedback for the watercolorized images. Also, the volunteer BU, PhD Student, Sofia Maria Nikolakaki, helped test the format of study 1

for the tactile project.

Also, I would like to thank the organizations who contributed in the progress of this thesis projects. The people from the Perkins School for the Blind: Jennifer Arnott (Research Librarian) and Kevin Hartigan (Museum tour-guide) provided access to resources, Kim Charlson (Executive Director of Perkins Library) and Tanja Milojevic (Braille Production Specialist) provided help in realizing study 3. The people from the Museum of Fine Arts Boston: Hannah Goodwin (Manager of Accessibility) and Ronit Minchom (Accessibility Coordinator) showed me how tactile images are used in real life. The Thayer Machine Shop at Dartmouth College and its crew (Kevin Baron, Pete Fontaine, Jason Downs, and Dan Denauw) provided laser cutting resources The Engineering Product Innovation Center (EPIC) at Boston University and its crew (Kara Mogensen and Caroline) provided 3D printing and laser cutting resources. The Visually Impaired and Blind User Group (VIBUG) helped me find participants for the tactile project.

Closing, I would like to thank my friends and family who were always by my side and supported me.

Published Papers

The following peer-reviewed papers are described in this thesis, here contributions of each author and copyrights are discussed.

Watercolor Prints (Chapter 3)

This article was published under the title “Watercolor Woodblock Printing with Image Analysis” in Eurographics 2018. The main contributions of each author were the following:

- Athina Panotopoulou: conducted field work to develop research project, gathered data set, organized and fabricated objects, implemented algorithms, main contributor in the shading map, minor contributor on the color map, and video editing.
- Sylvain Paris: contributed in the conceptualization of the project, main contributor in the color map, semantic map, overlapping woodblocks map, and in the written text.
- Emily Whiting: contributed in the conceptualization of the project, main contributor in conceptualizing the idea of printing watercolors with woodblocks, and the details map, and in the written text.

Statement of compliance with copyright requirements: Computer Graphics Forum '18. Copyright ©2018 The Eurographics Association and John Wiley & Sons Ltd. Published by John Wiley & Sons Ltd. DOI: <https://doi.org/10.1111/cgf.13360>

Tactile Images (Chapter 4)

Multiple authors – Copyrights This article was published under the title “Tactile Line Drawings for Improved Shape Understanding in Blind and Visually Impaired Users” in, 2020.

The main contributions of each author were the following:

- Athina Panotopoulou: led the project, contributed to the conceptualization of the project, conducted field work to develop research project, gathered the data set for the studies, organized 3D object and laser cut fabrication, implemented all algorithms (other than that implemented by Xiaoting Zhang), ran statistical analysis.
- Xiaoting Zhang: implemented skeleton and cross section extraction, contributed in laser cut fabrication, contributed in the written text.
- Tammy Qiu: designed guideline images, tested multiple materials for tactile image fabrication, contributed in laser cut fabrication.
- Xing-Dong Yang: contributed in the conceptualization of the project, main contributor in ideas for user study 1, and the statistical analysis of the data, contributed in the written text.
- Emily Whiting: contributed in the conceptualization of the project, main contributor in ideas for user study 3, the use of semantic segmentation, multi-perspective rendering and freestyle, the viewpoint selection of master camera, and in the written text.

Statement of compliance with copyright requirements: SIGGRAPH '20 Technical Paper, July 2020, online. Copyright ©2020 Association for Computing Machinery, 0730-0301/2020/7-ART89, DOI: <https://doi.org/10.1145/3386569.3392388>

Viewpoint Selection (Appendix A)

This article was published under the title “ Perceptual models of preference in 3D printing direction” in SIGGRAPH Asia, 2015. The main contributions of each author were the following:

- Xiaoting Zhang: led the project, contributor in the conceptualization of the project, implemented all algorithms (other than that implemented by Athina Panotopoulou and

Xinyi Le), data gathering to feed the network, fabricated objects, gathered data set, contributed in the written text.

- Xinyi Le: implemented the machine learning algorithm.
- Athina Panotopoulou: implemented the viewpoint preference metrics, video design and editing, rendered images except for Figure 6 and Figure 16.
- Emily Whiting: contributor in the conceptualization of the project, contributor in the video, contributor in the written text
- Charlie C. L. Wang: contributor in the conceptualization of the project, main contributor in the written text.

Statement of compliance with copyright requirements: SIGGRAPH Asia '15 Technical Paper, November 02 – 05, 2015, Kobe, Japan. Copyright ©2015 ACM 978–1–4503–3931–5/15/11, DOI: <https://doi.org/10.1145/2816795.2818121>

From Shape to Skeleton (Appendix A)

This article was published under the title “Printone: Interactive Resonance Simulation for Free-form Print-wind Instrument Design” in SIGGRAPH Asia, 2016. The main contributions of each author were the following:

- Nobuyuki Umetani: led the project, main contributor in the conceptualization of the project, found the fast approximation sound simulation algorithm, implemented all algorithms (other than that implemented by Athina Panotopoulou), fabricated objects, conducted testing, contributed in the written text, contributor in the video.
- Athina Panotopoulou: conducted work on how to fabricate objects and how to do the sound recording, identified how to compose the simulated sound and implemented it, implemented algorithm to design freeform instruments, video design and editing.

- Ryan Schmidt: contributed in the conceptualization of the project, the video, the written text.
- Emily Whiting: contributed in the conceptualization of the project, the data set, the written text.

Statement of compliance with copyright requirements: SIGGRAPH Asia '16 Technical Papers, December 05–08, 2016, Macao. Copyright ©2016 ACM ISBN: 978–1–4503–4514–9/16/12, DOI: <https://doi.org/10.1145/2980179.2980250>

From Skeleton to Shape (Appendix A)

This article was published under the title “Scaffolding a Skeleton” in Research in Shape Analysis, Association for Women in Mathematics Series, 2018. The main contributions of each author were the following:

- Athina Panotopoulou: contributed in the solution, the written text, and the image editing, implemented mesh generation and frame propagation.
- Elissa Ross: contributed in the solution, the written text, and the image editing, gathered data set.
- Kathrin Welker: contributed in the solution, the written text.
- Evelyne Hubert: led the project, contributed in the conceptualization of the project, the solution, the written text.
- Géraldine Morin: led the project, contributed in the conceptualization of the project, the solution, the written text, implemented branch point segmentation to generate initial frames.

Statement of compliance with copyright requirements: Copyright ©The Author(s) and the Association for Women in Mathematics 2018 A. Genctav et al. (eds.), Research in Shape

Analysis, Association for Women in Mathematics Series 12, https://doi.org/10.1007/978-3-319-77066-6_2. This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Contents

Abstract	ii
Acknowledgments	iii
0.1 Images	iii
0.2 People – Organizations	vii
Published Papers	x
Contents	xv
List of Figures	xix
List of Abbreviations	xxviii
1 Introduction	1
1.1 Computational Fabrication	2
1.2 Non-Photorealistic Rendering	3
1.3 Non-Photorealistic Rendering in Computational Fabrication	4
1.4 Contributions and Thesis Overview	9
2 Related Work	14
2.1 Traditional Methods for Creating Images	16
2.1.1 Watercolor Paintings	17
2.1.2 Woodblock Prints	18
2.1.3 Tactile Images	18
2.2 Computational Design of 2D Images	20
2.2.1 Non-Photorealistic Images	20

2.2.2	Paintings	22
2.3	Computational Fabrication of 2D images	24
2.3.1	Paintings	25
2.3.2	Prints	26
2.4	Tactile Images	27
2.5	Summary	29
3	Watercolor Prints	31
3.1	Background	33
3.1.1	Color	33
3.1.2	Variation on Color Intensity	34
3.1.3	Edges	35
3.1.4	Boundaries	35
3.1.5	Overview	38
3.2	Image Analysis	38
3.2.1	Overview of Our Image Analysis Pipeline	38
3.2.2	Color Map - M_c	39
3.2.3	Dry and Wet Areas	40
3.3	Fabrication Method	44
3.4	Implementation Details	46
3.4.1	Parameters of the Image Analysis	46
3.4.2	Design of the Woodblocks	48
3.4.3	Fabrication Details	48
3.5	Results	48
3.5.1	Results Summary	51
3.5.2	Informal User Study	54
3.6	Exploring Alternative Paths	56
3.6.1	Masking the Paper	56
3.6.2	Carving the Blocks	58

3.6.3	Dampening the Paper	60
3.6.4	Stretching the Paper	62
3.6.5	Modifying the Image Analysis Pipeline	63
3.7	Conclusion	64
4	Tactile Illustrations	66
4.1	Background Material	69
4.1.1	Skeleton	70
4.1.2	Curvature	71
4.1.3	Viewpoints	71
4.1.4	Shape Segmentation	73
4.2	Formative Studies	73
4.2.1	User Study 1: 3D Shape Replication	74
4.2.2	User Study 2: Assessment of Existing Illustration Guidelines	76
4.3	Tactile Illustration Method	81
4.3.1	Overview	81
4.3.2	Master Camera Placement	82
4.3.3	Part Camera Placement	83
4.3.4	Cross Section Textures	83
4.3.5	Part Camera Up-Vector Alignment	85
4.3.6	Compositing Part Renderings	86
4.3.7	Designs for User Study 3	89
4.4	User Study 3: Evaluation	91
4.5	Fabrication and Implementation Details	98
4.6	Discussion	99
4.6.1	Limitations & Future Work	99
4.6.2	Conclusion	100
5	Summary and Future Work	102
5.1	Changes in Pipeline	103

5.2	Changes in Design	105
5.3	Exploring New Designs	108
5.4	Interdisciplinary Collaborations	110
5.5	Conclusion	111
A	Related Research	112
A.1	Viewpoints	112
A.2	Skeleton	120
A.2.1	From Shape to Skeleton	120
A.2.2	From Skeleton to Shape	123
A.3	Key-Points	126
B	Supplemental Material	127
B.1	Watercolor Woodblock: Shading Algorithm	127
B.1.1	Details	128
B.2	Ring Location Sampling	128
B.3	Negative y Location Restriction	129
B.4	Canonical Coordinate Frame	129
B.5	Canonical Bounding Box	129
B.6	Part Viewing Direction	130
B.7	Up-Vector Initialization	130

List of Figures

1.1	Coffeepot that was materialized by Factum Arte (2010) based on the line drawing (a) of Piranesi (1770). The digital design (b) was 3D printed (c) to a cast, before it was made with silver (d). Fast prototyping tools make fabrication, but not design, easier. [©Factum Arte, Madrid, Spain (factum-arte.com)]	2
1.2	In some cases images need to be fabricated: to display an authentic quality, such as in the case of the Twilight watercolor painting (a) of Georg Grosz (1923) [©Museo Thyssen-Bornemisza, Madrid, Spain (museothyssen.org)], or to be interacted with such as in the case of a shadow puppet, Java (1911) Figure (c) [Daderot, Peabody Museum, Cambridge, USA (wikimedia.org)]	4
1.3	Representative examples of machine-made images. In these first works, the images were black and white. In later works, more colors and a greater variety of media, such as oils, dyes, and acrylic paints, were added. The composition in (a) is generated by the computer, while in (b) and (c) the composition is stylized from the input. [Harold Cohen (1984), Victoria and Albert Museum, London, UK, (collections.vam.ac.uk)] [Lindemeier [245] (2010), “e-David: Non-Photorealistic Rendering using a Robot and Visual Feedback”] [Tresset and Leymarie [415] (2012), “Sketches by Paul the Robot”]	6
1.4	Image analysis approach to design. The input is Raphael’s Madonna in the Meadow, painting 1505 (a) and edges(b) and depth values(c) are extracted. [Photograph DcoetzeeBo, Kunsthistorisches Museum, Wien, Austria, (wikimedia.org)], [Images Reichinger et al. [346] (2011), “High-Quality Tactile Paintings”]	7

1.5	Non-standard approach to fabricate images. Here Galea et al. use drones (a) to create a stippled image (b). However, they face problems that in virtual design are not present, such as the dripping of the paints (c). [Galea et al. [138] (2016), “Stippling with aerial robots”]	7
1.6	In many cases, testing with human participants is omitted. Review in IEEE visualization is showing that less than 20% of works conduct user studies (red box). [Isenberg et al. [189] (2013), “A systematic review on the practice of evaluating visualization”, IEEE]	9
1.7	The two types of stylized images discussed in this thesis. The first, a watercolorized image (a), is fabricated with watercolor paints for the first time in graphics and displays an authentic style. The second, a tactile image (b), conveys shape information to individuals with blindness.	11
1.8	A characteristic watercolor painting (left) displays the gradients which are distinct in this medium (right, closeup). [Painting ©Cate Parr., “Rooney Mara, Oscars in Givenchy”, (cateparr.com)]	12
1.9	Graphic designers use their experience and existing extensive guidelines to transcribe a visual image (a) to a tactile image (b), as seen here. [©Braille Authority of North America, Guidelines and Standards for Tactile Graphics (2010), (bana.org)]	13
2.1	Stylized fabrication. (Top) input; (bottom) stylized output. This research area includes different production techniques and artifacts, e.g., (a) laser cut 2D portraits [Image Iseringhausen et al. [190] (2020), “Computational Parquetry”]; (b) 2.1D plastic paintings [Reichinger et al. [346] (2011), “High-Quality Tactile Paintings”]; or (c) 3D yarn bunnies [Narayanan et al. [306] (2019) “Visual Knitting Machine Programming”]. Input photographs: [Karl Joseph Stieler (1820), “Beethoven”, Beethoven-Haus, Bonn (wikimedia.org)]; [Raphael, “Madonna in the Meadow”, Kunsthistorisches Museum, Wien, Austria, (wikimedia.org)]; [Narayanan (2019), “Stanford bunny”]	16

2.2	An example of the traditional woodblock printing process. First, the negative of a single colored image (a) is carved on a flat wood surface (woodblock) (b). Next, ink is applied on the woodblock before is pressed on a paper with the baren (c). The final print is a succession of prints from many woodblocks (d). [©David Bull (2015), (patreon.com/mokuhankan) and (mokuhankan.com)] . . .	19
2.3	Examples of tactile images made with a variety of different materials: a map of Africa (a), an art exhibit (b), a picture (c), and a map (d). [(2020) (wikimedia.org)]; [©Paul Bartlett Ré, “A book and two traveling exhibits for the blind and sighted”, (paulre.org)]; [M. Kunz (1902), ©Perkins school for the blind archives, Watertown, MA, USA (perkins.org)] [American Printing House for the Blind (2004), “World Maps” (aph.org)]	21
2.4	Stylized fabrication of non-photorealistic images. The input can be a 3D geometry or a 2D image. The task for either input can be utilitarian or aesthetic. Finally, the image can be distinguished in terms of fabrication technique to 2D, 2.5D, and 3D.	22
2.5	The pipelines for non-photorealistic images are often semi-automatic. Here, the input 3D model (a) is annotated with lines to indicate where texture (b) appears in the output (c). [Winkenbach and Salesin [451] (1994), “Computer-Generated Pen-and-Ink Illustration”, ACM]	23
2.6	Examples of computational design (a), and computational fabrication (b) of stylized images for the same input. [Lindemeier [245] (2018), “e-David: Non-Photorealistic Rendering using a Robot and Visual Feedback”]	24
3.1	A Marilyn Monroe photograph was used as an input to our pipeline (a), with our hybrid fabrication process constituted by watercolor paint and laser cut woodblocks (b). We created the final painting (c) through layered prints of the individual woodblocks. [Photograph by Milton H. Greene ©(2018) Joshua Greene (archiveimages.com)]	31

3.2	Image analysis Pipeline. The inputs are the original photograph ¹ I_p (a), and the user semantic annotation I_s (b). The colors in the output images represent print colors depicted in (d), the white and the chess pattern in (g)-(j) correspond to non-printed areas.	37
3.3	The user can control the level of abstraction by varying the number of colors, k , in the color palette extraction.	39
3.4	Images given as input to the laser cutter, corresponding to the woodblocks in Figure 3.1.	45
3.5	An example of our printing sequence. (a) We start by masking the “dry” areas without color diffusion (white) and printing the “wet” colors with water diffusion first. (b) We then remove the masking fluid and print the dry colors without water diffusion. (c) Lastly we print the details blocks (minima and edges). . . .	46
3.6	4 out of 20 images we used to identify the values for the parameters for the smoothness map algorithm	47
3.7	Removing stages in our pipeline negatively affects the results of our algorithm. In this figure, non-skin areas are digitally masked out for clarity.	49
3.8	We let users choose whether the print has fine details, e.g., the feathers in this case. The difference in the color gradient on the bird body and wing comes from the inherent variations that occur between prints.	49
3.9	Two examples of repeated prints to illustrate the variations inherent in our approach. The top row is the Hibiscus flower and the bottom is the Robin bird. All the results look stylistically similar. However, due to the liquidity of watercolor unique variations emerge.	50
3.10	Print of two color blocks. Different strategies to create color gradients. The top row shows actual prints and the bottom row shows a visualization of the two blocks that we fabricated. Overlapping blocks (c) produce the smoothest gradients. This is the option that we use in our technique.	51

3.11	Examples of watercolor prints with a variety of input photographs. (Row 1) Portrait of Marilyn Monroe [‡] . (Row 2) The flower is a representative example of creating a color gradient. (Row 3) The bird has diffusion of more than two colors (Row 4) The landscape contains detailed textures which are produced with the detail woodblocks printed on top of the dry blocks.	
	[‡] Photograph by Milton H. Greene ©(2018) Joshua Greene www.archiveimages.com	52
3.12	Marilyn Monroe’s portrait. The shading shows the curvature on the face, the neck, and the shoulders; Marilyn is immediately recognizable even with the absence of the blond hair and the red lips.	53
3.13	The Hibiscus flower. The shading has two color gradations and non-painted areas on the petals that improve image quality.	53
3.14	The Robin bird. This example, entirely made of watercolor, has a three-color gradation on the middle of the feathers.	53
3.15	The house on the lake. Detailed texture was achieved by the use of two rather than one dry woodblocks for each color, making the image to appear less flat.	54
3.16	The blocks for the Hibiscus flower. The background is not printed, and three colors are used: yellow, pink and brown. In some of the woodblocks, more than one color is carved to minimize the consumption in fabrication, e.g., the first shading woodblock prints both yellow and brown.	54
3.17	The blocks for the Robin bird. The background is not printed, and three colors are used: light blue, orange, and gray. Also, there are no dry woodblocks, except from the details woodblocks which in some of our prints were printed with water.	55
3.18	The blocks for the house on the lake. This large number of woodblocks was a result of the more extensive color palette (light gray, blue, green, and dark gray) and made the printing process more challenging. The wet woodblock printing in particular required a quick process for the paper to remain wet.	56
3.19	Choosing materials for masking. The left image shows a previously masked area with masking fluid. The middle and right image show the results of experimentation with stencil & tapes.	57

3.20	Woodblocks carved by different machines and the prints produced by them . . .	58
3.21	Prints by blocks carved on different materials: woodblock Vs Plexiglas block . .	59
3.22	Carving different types of woods	60
3.23	Color expansion when dampening paper with brushes, the color gradients are not sufficiently big.	60
3.24	Different number of spraying after the printing of the block	61
3.25	Result of diffusion with excessive water application.	61
3.26	The same woodblocks for pink and yellow colors printed with different amount of spraying. Here, we adjusted the images to better match the input photo. . . .	62
3.27	Result in diffusion by not well stretched watercolor.	63
3.28	Different approaches than ours. It is a matter of personal taste which is preferable.	64
3.29	Adding eastern style lines, produces a less watercolor look.	64
4.1	Tactile illustrations (a,c) displayed alongside exhibits in the Museum of Fine Arts Boston (b,d) for accessibility purposes (www.mfa.org). Black areas in illustrations indicate raised regions on the paper surface that can be perceived by touch.	67
4.2	Generating tactile illustrations to improve shape understanding in blind individ- uals: (a) Physical 3D printed object; (b) Input to our pipeline: a pre-partitioned object with colors indicating segments; (c) Local part cameras; (d) Resulting stylized illustration with curvature indicative texture; (e) Evaluation of tactile illustrations on microcapsule paper in a user study with 20 blind participants. . .	68
4.3	Sculptures from User Study 1. (Top row) Reference objects. (Middle row) Participant 1 sculptures. (Bottom row) Participant 2 sculptures.	74
4.4	3D printed reference objects for User Study 2 (row 1). Corresponding illus- trations: non-canonical views (row 2), canonical views (row 3), and textured canonical views(row 4).	77

4.5	Overview of our pipeline for the design of tactile illustrations. (a) Input model, colors represent different parts of the faucet. (b) Master camera placement (inset box shows camera view). (c) Local cameras are assigned to each part (inset boxes show camera views). (d) The image layers are composited; positioning of each layer is guided by the master camera and the part connectivities (d). (e) Infill textures are generated based on cross sections of each part (black represents raised areas). (f) The fabricated illustration created with microcapsule paper.	80
4.6	Different metrics considered for finding the master camera placement. The viewpoints maximize: (A) Silhouette Curvature Extrema, (B) Silhouette Curvature, (C) Projected Area, (D) Silhouette Length, (E) Surface Visibility, (F) Maximum Depth, (G) Depth Distribution, (H) Viewpoint Entropy. The maximum silhouette length (D) decreases part occlusion.	82
4.7	The chair displaying the bounding box and skeleton for each part (a). Cross sections sampled uniformly along the skeleton. The cross sections of each part as seen from their corresponding local cameras (c), and after the flattening operation (d).	85
4.8	Alignment of the part camera up-vector. (a) Input object as seen from the master camera. (b) Center and connections for the tabletop from the part camera view (top) and master camera view (bottom). Angle computed so that horizontal distances of connections, w , match the master camera view. (c) Final upright orientation of the local camera, $\hat{\mathbf{u}}'_j$	86
4.9	Aligning part camera image layers. (a) The input model as seen from the master camera. (b) If the centroids in the master camera view are used for aligning part renderings, the resulting image is disconnected. (c) Instead, the connection points are used for aligning image layers.	87
4.10	Photograph of tactile illustration of the monitor, including a closeup view of the base. Fabricated with microcapsule paper.	88

4.11	Objects 1–4 in User Study 3: Evaluation. For each object: (top) 3D reference model with variations; (bottom left) tactile illustration using BANA [315] guidelines as baseline; (bottom right) tactile illustration using our technique. In each case the illustrations depict the middle reference object.	90
4.12	Objects 5–8 in User Study 3: Evaluation. For each object: (top) 3D reference model with variations; (bottom left) tactile illustration using BANA [315] guidelines as baseline; (bottom right) tactile illustration using our technique. In each case the illustrations depict the middle reference object. The jug was used as a training example at the beginning of the study.	91
4.13	For each 3D object included in User Study 3: Evaluation, we created three versions with minor geometric variations. Beige regions are identical among the three versions; green regions have been modified. The models were 3D printed so that participants could compare (by touch) the 3D objects depicted in the 2D tactile illustrations. The tactile illustrations shown in Figure 4.14 correspond to the middle version in each case.	92
4.14	Resulting designs for tactile illustrations using our technique, with 3D input models shown.	93
4.15	Example materials for User Study 3: Evaluation. Participants are given three variations of the object (a), produced by 3D printing (b). They are given one tactile illustration style – either our new multi-projection style (c, left) or the BANA baseline style (c, middle-right). Illustrations are produced using microcapsule paper (d). Participants are asked to identify which of the three variations matches the illustration. (Supplemental Materials for all materials used in study 3).	94
4.16	User Study 3 results, shows mean success rate for matching tactile illustration to correct 3D object. Results shown for average over all objects, and per object success rates.	95
4.17	User Study 3 results, shows mean success rate for matching tactile illustration to correct 3D object. Grouped by age of vision loss (self-identified).	96

5.1	Manual combination of two (here) or more colors per block, result in a smaller number of blocks and the automation is left as a direction for future work. . . .	106
5.2	Multi-projection exploded line drawings	108
5.3	Interdisciplinary collaborations allow to integrate ideas from different domains. In this thesis, knowledge was shared by artists, such as print-makers (a) and mathematicians (b), among others. [Photograph Sarah H. Smith ©(2018)] [Workshop, Women in shape modeling, İzmir, Turkey, (2016)]	111
A.1	Perceptual attributes for the bunny model.	114
A.2	Interactive modification of viewpoint preference – (top left) automatic computation $VP_5(\mathbf{d})$, (top right) Harmonic field $a_{14}(\mathbf{d})$ according to user specification and (bottom) the modified field by average and normalization. Blue denotes <i>zero</i> and red denotes <i>one</i> (highest preference). To display the overall distribution, the spheres are also unfolded onto a rectangular region.	119
A.3	An example of skeleton extraction.	121
A.4	Designed musical instruments which were designed based on the skeleton extracted from a 2D closed curve.	122
A.5	The two main steps of the algorithm: Partition (top) and Sleeving (bottom). . .	124
A.6	Scaffold meshes generated (right) by input skeletons (left).	126
B.1	The vector field, green arrows, we use to initialize the up-vector for the part cameras. It gives reasonable initial renderings.	131

List of Abbreviations

List of abbreviations used in this dissertation, with page numbers where the abbreviation is first used.

CNC computer numerical control machines	2
BANA Braille Authority of North America	12
PCA principal component analysis	72

1

Introduction

Throughout history, humans have conceptualized and materialized hand-made artifacts [164]. With materials like stone, wood, metal, and clay, humans constructed building tools [1, 2, 5, 8] that they could use manually to make more complex, beautiful, and elaborate objects such as garments [11], ceramics [8], jewelry [6], musical instruments [10], and books [12]. However, up to the 19th century, the artifacts were mainly an outcome of the experience and talent of artisans, and the process to create high-quality objects was time-consuming [3].

During the industrial revolution (18th – 19th centuries), new paths opened with the invention of motors, energy-powered machines, and heavy duty tools for construction and material processing. These machines [9] made production of objects faster and easier, allowing the needs of the rapidly growing population to be met. As these machines were created for mass production in industrial environments, they did not allow for customization. Also, during this period, the small-scale machines, such as electric saws or drills, did not allow for much creative capacity to non-experts. However, with the recent development of small-scale easily programmable fabrication tools, such as laser cutters (1965-1970) and 3D printers (1985-1990), expert designers can fabricate with less effort high-quality artifacts which are useful and aesthetically pleasing. Such an example is the coffeepot in Figure 1.1(d). This silver coffeepot is casted, meaning that a form is molded to allow for fabrication of multiple copies. The manual fabrication of its cast requires a highly experienced craftsman; the process takes several weeks and mistakes cannot be easily corrected. Here, as can be seen from Figure 1.1(c), the cast of the coffeepot is 3D-printed.




Figure 1.1: Coffeepot that was materialized by Factum Arte (2010) based on the line drawing (a) of Piranesi (1770). The digital design (b) was 3D printed (c) to a cast, before it was made with silver (d). Fast prototyping tools make fabrication, but not design, easier. [©Factum Arte, Madrid, Spain (factum-arte.com)]

This type of research, namely computational fabrication, engages computational capabilities for the realization of expert domain artifacts.

1.1 Computational Fabrication

Researchers from different fields, such as material scientists, mechanical engineers, computer scientists, mathematicians, and artists, have been using fast prototyping tools for a variety of creative and functional tasks. Particularly, the graphics community, over the past 15 years, has noticeably increased its research into computational fabrication with the development of new hardware (e.g., position correcting tool [351], stylus [334]), new software (e.g., 3D-modeling [352], functional mechanisms [216], multi-material printing [272, 383]), and new algorithmic approaches (e.g., winding number computation [37]). This research has resulted in making computer numerical control machines (CNC) or rapid prototyping devices more accessible and more versatile.

Despite the advantages of these technologies, the design process is still incredibly challenging; while it is somehow straightforward to 3D print from a digital design, e.g., from the digital design in Figure 1.1(b) the cast in Figure 1.1(c), it is not as obvious how to design such a 3D (three-dimensional) shape. Therefore, another branch of computational fabrication research focuses on providing advanced design capabilities [48], either by targeting high-level features

(e.g., light propagation [331], stress [365], elasticity [282, 366], sturdiness [460], and material consumption [424]), or by accessing expert design domains (e.g., furniture making [367, 439], knitting [306], articulated characters design [34], and spray painting [380]). These improvements make it possible to design personalized objects with great attention to detail, as well as to replicate the process of fabrication. Improving the design in these ways is also important because it is more sustainable and provides a larger number of people access to design; amateurs gain access to accurate design, fast prototyping, and, in general, less labored processes [487].

1.2 Non-Photorealistic Rendering

A different area of design, closely related to the works I discuss here, focuses on the creation of images. However, the focus is not images that faithfully resemble photographic results, namely photorealistic images [15, 133, 296, 333], like the rendering of the digital model in Figure 1.1(b). Instead, the focus is on stylized images, often called non-photorealistic images [225, 248]. Stylized images were traditionally developed by hand in arts and sciences, and later in computer graphics computationally in the sub-field of non-photorealistic rendering. A representative example of non-photorealistic imagery is the line-drawing in Figure 1.1(a); lines have been used instead of smooth shades which make the image less realistic but the shape more pronounced and the drawing more aesthetically notable. Researchers in non-photorealistic rendering provide computational approaches for the design of stylized images, to either convey information (e.g., [244]) or create an artistic depiction (e.g., [46, 162]). The question then becomes: why is it necessary to fabricate images with real materials and not only design virtual ones? Especially when considering that by introducing real materials we introduce further difficulties in the process.

Real materials are necessary when authenticity is a required aesthetic quality, such as the use of oil-paints, pens, and pencils for creating artistic productions. For example, as one moves closer to the watercolor painting of George Grosz (Figure 1.2(a)), one can notice delicate details, including the texture of the paper and the difference in granularity between the pigments (Figure 1.2(b)). This interplay of real materials defines one quality component of a real painting,

something that a virtual drawing lacks. There might be psychological reasons for people being interested in creating and buying paintings [29,318]. For example, research suggests that images made of real materials amplify the emotional response of the observer [415]. Real materials are not only used to create authentic works, but they are also used when images need to be handled, such as the example of the shadow theater figure shown in Figure 1.2(c). Other examples falling into this category are puzzle pieces or self-shadowed and caustic images. Because of this, researchers are motivated to fabricate non-photorealistic images for different tasks and from different materials. The next section describes key works that bridge the fields of computational fabrication and non-photorealistic rendering, which is the subject of this thesis.




Figure 1.2: *In some cases images need to be fabricated: to display an authentic quality, such as in the case of the Twilight watercolor painting (a) of Georg Grosz (1923) [©Museo Thyssen-Bornemisza, Madrid, Spain (museothyssen.org)], or to be interacted with such as in the case of a shadow puppet, Java (1911) Figure (c) [Daderot, Peabody Museum, Cambridge, USA (wikimedia.org)]*

1.3 Non-Photorealistic Rendering in Computational Fabrication

This thesis provides the means for the fabrication of two types of images: images made with watercolor paints, i.e., watercolor woodblock prints, and graphics made with microcapsule paper, i.e., tactile illustrations. To identify and combine key design information for the watercolor woodblock prints, we examine hand-made watercolor pieces of art and tutorials on watercolor painting, and we enumerate a set of information elements (i.e., colors, semantic boundaries,

smooth intensity gradients, details, and color gradients) and how these elements should be maintained. In the second work regarding tactile illustrations, we examine tutorials on tactile images and prior research in human cognition, and we design a pipeline that clearly presents a set of information elements which are identified to be key in shape understanding by individuals with blindness (i.e., surface geometry, topology/connections, viewpoints, and shape segmentation).

While researchers that work across computational fabrication and non-photorealistic rendering use a variety of real-world materials to manufacture images (e.g., [52, 132, 190, 222, 237, 355, 456, 471, 481]), as far as we know, graphics research has not yet explored images made with either watercolors or microcapsule paper. The fabrication for these two projects is important because it allows, in the first project, for the production of an authentic art image, and in the second project, for the tactile exploration of the depicted 3D shape. Watercolor is a difficult painting medium including advanced techniques where the wet paint is placed on a wet surface and requires training to produce a somewhat predictable outcome. Tactile microcapsule paper is a material that is used extensively in the blind community to convey pictorial information. In all cases, stylized images are first designed, then fabricated, and finally tested.

Gaining control over the labored and time-consuming process of design is hard, as is shown by the many years researchers have spent in designing new image making approaches. Harold Cohen [95–97, 97–99], back in 1979, created some of the first works in fabrication of non-photorealistic images, producing novel algorithmic compositions (Figure 1.3(a)). In one of his articles he mentions that it took him 15 years before he introduced colors in the process and 20 years to identify and fine-tune the rules of color composition [100]. Similarly, Lindemeier et al. [246–249], over a number of years, explored a range of media, including black and white ink and acrylic paints and algorithmic methods (Figure 1.3(b)).

One approach of image design is by abstracting key input information and combining it. When the input is an image (e.g., Figure 1.4(a)), extracted elements can include edges (Figure 1.4(b)) or depth values (Figure 1.4(c)). We often collect this type of information using image processing and analysis, as in the watercolor project. On the other hand, when the input is a 3D shape we use geometry processing and analysis, as in the tactile project. While key elements can be extracted with the same techniques for different image fabrication projects,



(a) Computer-generated ink [245]



(b) Pen drawing [245]



(c) Ink with brush [415]

Figure 1.3: Representative examples of machine-made images. In these first works, the images were black and white. In later works, more colors and a greater variety of media, such as oils, dyes, and acrylic paints, were added. The composition in (a) is generated by the computer, while in (b) and (c) the composition is stylized from the input. [Harold Cohen (1984), *Victoria and Albert Museum, London, UK*, (collections.vam.ac.uk)] [Lindemeier [245] (2010), “*e-David: Non-Photorealistic Rendering using a Robot and Visual Feedback*”] [Tresset and Leymarie [415] (2012), “*Sketches by Paul the Robot*”]

the way this information is used and combined differs because of the different media and tasks. For example, in the pen-drawn portrait by a process designed by Tresset and Leymarie [415], Figure 1.3(c), edges and shading were extracted, similar to our work in watercolor, but the process differs due to the different media. While Tresset and Leymarie [415] chose the use of limited edge orientations for the design of a discrete level of shading based on personal observation of pen art pieces, we fabricate continuous gradients, which result from the diffusion of paints in the water.

In terms of fabrication, researchers have become unconventional in order to achieve the artistic styles they wish. Stippling, a depiction technique made of dots, has been researched for a couple of decades as a virtual rendition in graphics [278], and has only lately been explored in fabrication. Galea et al. [138] used a creative fabrication approach; drones flew along a canvas holding sponges saturated with black ink and touched the canvas lightly, leaving a soft spattering of ink (Figure 1.5(a)–(c)). While their work focuses on how to control the drones, physical limitations such as inconsistent ink absorption or ink running (Figure 1.5(c)), prevent error-free fabrication, a difficulty that only researchers working with real materials face.

As it has been pointed out, computational design allows for absolute freedom, while com-




Figure 1.4: Image analysis approach to design. The input is Raphael’s *Madonna in the Meadow*, painting 1505 (a) and edges (b) and depth values (c) are extracted. [Photograph DcoetzeeBo, Kunsthistorisches Museum, Wien, Austria, (wikimedia.org)], [Images Reichinger et al. [346] (2011), “High-Quality Tactile Paintings”]

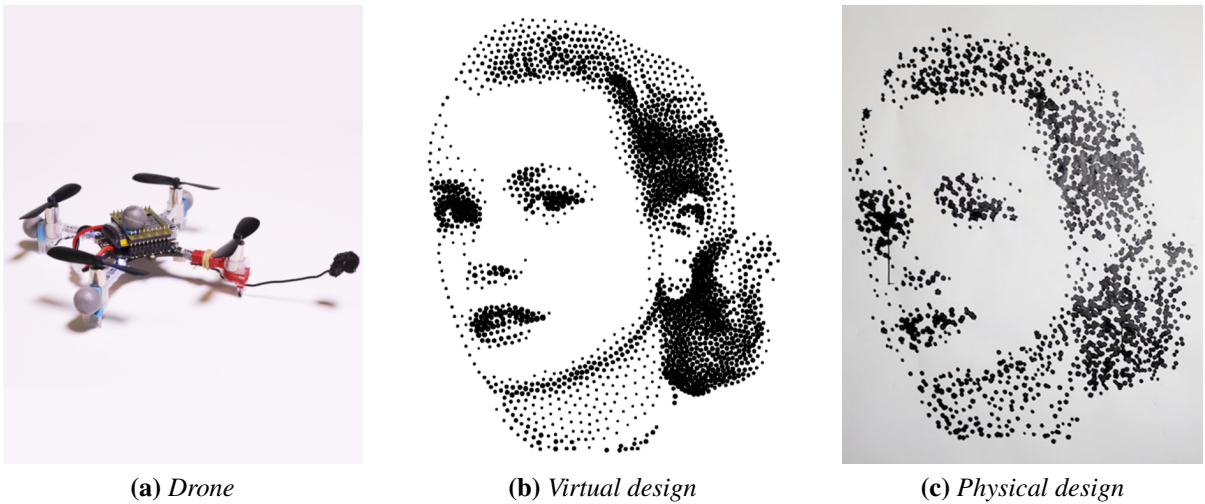


Figure 1.5: Non-standard approach to fabricate images. Here Galea et al. use drones (a) to create a stippled image (b). However, they face problems that in virtual design are not present, such as the dripping of the paints (c). [Galea et al. [138] (2016), “Stippling with aerial robots”]

putational fabrication implies deriving, operating, and putting in place strategies to realize the physical design. Algorithms have been developed which work exceptionally well for specific classes of materials, but unfortunately, not all media effects fall within these narrow confines. In contrast to other research areas, in which parameter tuning would result in a technique suitable for different materials (e.g., rendering differently textured materials such as leather and wood

in a unified way, rendering scenes with different participating media such as fog and clouds, or modeling solids and liquids such as sand and gelatin [20, 193, 463]), that is not the case for this type of work. Fabrication needs to be fine-tuned to the specific material and task. While simulation techniques have succeeded in predicting the behavior of a range of materials and tasks, not all image making results can be accurately predicted (e.g., reflectance microfacets for projected caustic images and refractive multi-layered transparent filters for light-source dependent multi-images [36, 447]). Many materials, like watercolors, are also some of the most difficult to replicate, since the environment is not fully controllable. Researchers have pointed out such difficulties and used visual feedback and made simplifications in order to achieve high quality results. For example, visual feedback was needed in the work of Tresset and Leymarie [414] (Figure 1.3(c)); they mention that if the simulation of both the drawing process and the robot action was sufficiently precise, this external visual feedback would not be necessary. Similarly, Lindemeier et al. [248] applied black ink, with a robotic arm, and allowed to dry before painting in the same place to simplify the process, and noted in his later work that “incorporating wet-on-wet paint interaction will be a very challenging problem, but we think that this results in much more aesthetic paintings and new opportunities of creating artistic imagery.” [246]. In the first project of this thesis, we implement the wet-on-wet technique by applying the paints with woodblocks and the water with a water-spray and make for the first time watercolor multi-color gradients. We accept the variations of this challenging and unstable material, the watercolor, and try to limit the uncontrollable effects of diffusion in places they were needed to replicate effects created by highly experienced artists. Similarly, the use of a printing press for the creation of watercolor images and the laser cutter for the creation of the tactile images are both non-standard fabrication approaches.

Unlike research on non-stylized images [272, 325], in non-photorealistic imagery it is still unclear how to measure the success of the design. One way to test the effectiveness of stylized images is conducting studies that involve human subjects. User studies are suitable in visualization, because images are judged using human criteria. However, because working with human subjects is a time and resource intensive process, in most cases it is omitted. As mentioned in the review of Isenberg et al. regarding how often user studies are conducted, “it

is interesting to see that their prevalence was relatively low” [189], as shown in Figure 1.6 in which the percentage of human participation is less than 20%. The evaluation, in some cases, can be done quantitatively by using the difference between the filtered and input images [415]. Aesthetically, though, this does not test the image quality as assessed by professionals with expert knowledge. Lindemeier et al. [248] highlight that an: “interesting question for the future would be a numerical evaluation of the painting results. Appropriate statistical means have to be found to compute the quality of different artistic representations.” Therefore, we follow the approach of human testing in the works of this thesis. This is addressed in the watercolor project by using visual inspection and conducting an informal user study with art students. In the tactile work, the success of the design space was tested by running a comparative user performance study populated by blind individuals to test the fabricated prototypes.

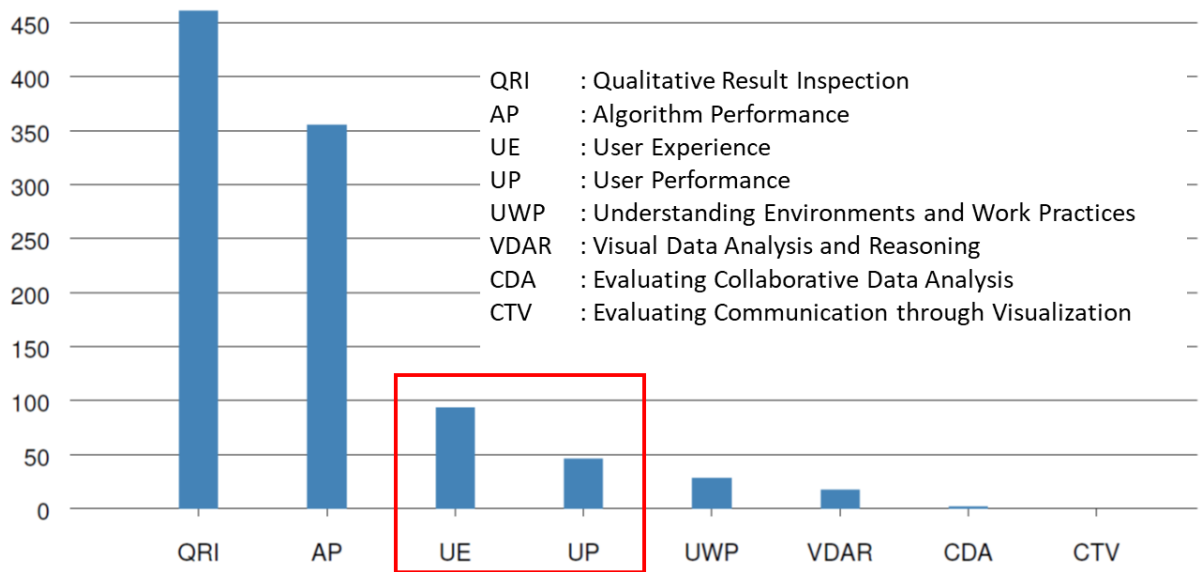


Figure 1.6: *In many cases, testing with human participants is omitted. Review in IEEE visualization is showing that less than 20% of works conduct user studies (red box). [Isenberg et al. [189] (2013), “A systematic review on the practice of evaluating visualization”, IEEE]*

1.4 Contributions and Thesis Overview

In sum, the goal of this dissertation is to develop techniques for the computational fabrication of two types of images that have not been previously examined in graphics, specifically wa-

tercolorized and tactile images for shape understanding. These two non-photorealistic images are made of real materials. For the design we extract key information elements, for the fabrication we use non-conventional approaches, and for the testing we conduct user studies. Our techniques not only enable anyone with the necessary resources to produce watercolorized and tactile images, using minimal guidance as input, but also provide other researchers baselines to compare and improve upon.

The work in this dissertation builds upon artistic elements of watercolor painting, as well as the sense of touch in human perception. The major contributions are:


- a fabrication-aware pipeline for the design of watercolorized images including an approach for the design of single and multi-color gradients;
- a fabrication procedure for creating watercolorized images (Figure 1.7(a));
- an algorithm for the design of tactile images with improved shape understanding for the blind; and
- a format for a user study to test tactile images on shape understanding (Figure 1.7(b)).

Recent approaches in graphics have focused on how to fabricate images using a variety of materials. While there is much prior work in the fabrication of images in graphics (Chapter 2), works with watercolor and microcapsule paper have not drawn much attention.

In watercolor painting, similar to the painting in Figure 1.8, most often a limited number of paints and water are applied with brushes on a paper, and the interactions between the paints and water create unique diffusion patterns. Another common material is masking fluid, which prevents the paints and water from touching the paper, and allows distinction between wet and dry areas which change the diffusion of paints. For example, in the painting in Figure 1.8, there is much diffusion around the model, but much less in it. Unfortunately, the abstract watercolor painting, with the delicate color gradients and finely detailed contours can be challenging for non-experts to effectively create. Specifically, the challenge lies on deciding where the abstractions, gradients, and details should be placed and executing that placement. We take inspiration from color gradations and propose a novel pipeline for the design and fabrication of



(a) Watercolorized image



(b) Tactile Image

Figure 1.7: *The two types of stylized images discussed in this thesis. The first, a watercolorized image (a), is fabricated with watercolor paints for the first time in graphics and displays an authentic style. The second, a tactile image (b), conveys shape information to individuals with blindness.*

watercolorized images (Chapter 3). We first study artists working with the medium as well as watercolor images themselves and identify prominent characteristics. We also experiment with watercolor to identify a way to resolve its uncontrollability, resulting in the use of woodblock printing. To design the woodblocks, we extract key elements using image analysis to identify and guide the placement of pigments and masking fluid on paper. Using a semi-automatic fabrication process, the pigments are transferred first to laser cut wooden stamps (woodblocks) and along with sprayed water, are then transferred on the paper. The woodblocks can be reused to produce multiple copies but each print is unique due to the physical process. The final piece depicts the desired scene while exhibiting the distinctive features of watercolor. The images are non-photorealistic and convey a particular abstract style, created by smooth gradations of a limited number or colors, similar to Figure 1.7(a).

In contrast to this earlier work, the goal of which is artistic, the next image type is utilitarian. Members of the blind and visually impaired community rely heavily on tactile graphics (or tangible graphics) – informational displays that are made in relief to be touched [119] – to understand geometric ideas in school textbooks, depict a story in children’s books, or conceptualize



Figure 1.8: A characteristic watercolor painting (left) displays the gradients which are distinct in this medium (right, closeup). [Painting ©Cate Parr., “Rooney Mara, Oscars in Givenchy”, (cateparr.com)]

exhibits in museums. However, these images often fail to achieve their goals, in large part due to the lack of understanding in how 3D shapes can be represented in 2D projections. Existing work in tactile image design explored ways to improve these images, and many guidelines based on these findings have been disseminated. The image in Figure 1.9 shows the transcription of the visual image to the corresponding tactile, by manually following the Braille Authority of North America (BANA) guidelines. This document includes a large number of rules. For example, it is recommended that textures are used to indicate semantic areas, e.g., sepals have one texture and petals have a different one. It is also suggested that a minimum number of images taken from canonical viewpoints is used to convey the necessary information, e.g., the use of a single side view of the flower. We take an entirely different approach in our design, introducing a curvature indicative line texture, and a single multi-projection line drawing (Chapter 4). We fabricate tactile images using microcapsule paper and a laser cutter. Taking inspiration from research in cognitive science and findings from our own user studies, we improve upon the existing guidelines [315] on the design of tactile images to increase shape understanding. Successful

image design of 3D objects presupposes identification and combination of important information in topology and geometry. The images differ significantly from photorealistic images, e.g., they use raised lines and textures, instead of color variations, similar to Figure 1.7(b). Finally, we develop and implement a user study in order to test the images.

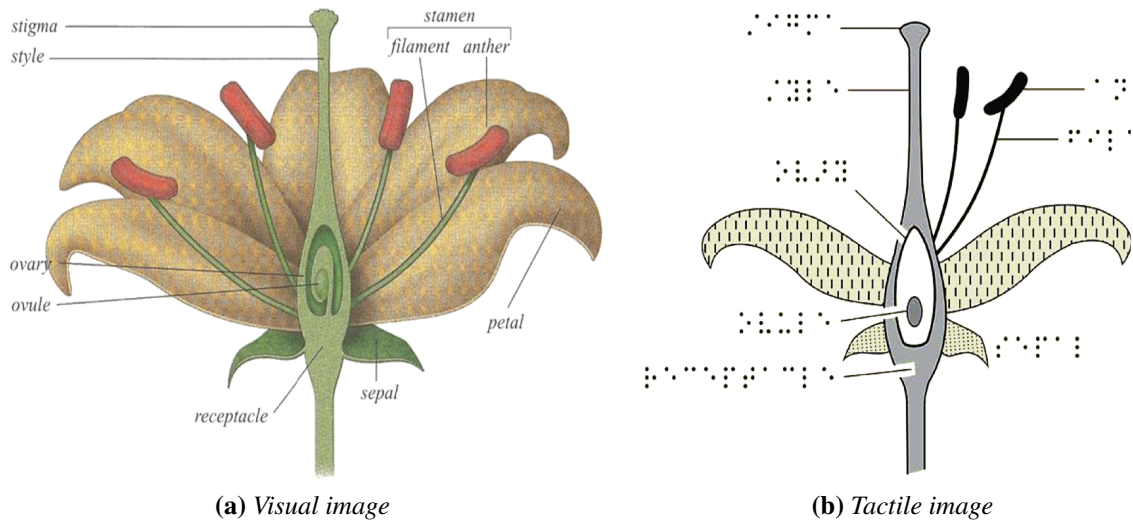


Figure 1.9: Graphic designers use their experience and existing extensive guidelines to transcribe a visual image (a) to a tactile image (b), as seen here. [©Braille Authority of North America, *Guidelines and Standards for Tactile Graphics* (2010), (bana.org)]

2

Related Work

Works in computational fabrication [302] provide anyone the means to both achieve precisely designed artifacts by customizing object properties and express creativity [50] through computer programming. While machines have been used to customize object properties such as levitation [291, 312], mechanics [187], shape, volume, and stiffness [319, 320], among others [157, 284, 299, 354, 418, 419, 434], another avenue of research in computational fabrication focuses on the creative aspect of fabrication and has allowed users to customize and stylize the design. Examples of stylized design are shown in Figure 2.1, with the input in the top row used to design and fabricate the artifact in the bottom row. Stylized fabrication involves a variety of materials, artifacts, and production techniques, including 2D fabrication of laser cut wooden portraits (Figure 2.1(a)), 2.1D or 2.5D CNC milled relief paintings (Figure 2.1(b)), or 3D bunnies knitted from yarn (Figure 2.1(c)). While 3D is commonly known, 2.1D refers to discretized bas-reliefs, layered depth diagrams, or flat layered bas-relief, and 2.5D refers to bas-reliefs [347]. All of these works are heterogeneous, e.g. they process different materials, they follow different fabrication techniques, or they produce different artifacts. However, at the same time, they are all stylized fabrication works because they create stylized versions of the given input.

This variety of materials includes yarn, paper, leather, metal, plastic, resin, and wires [214, 252, 255, 305, 406, 456]. Materials can be man-made (e.g., plastic in Figure 2.1(b)) or minimally processed natural materials (e.g., wood in Figure 2.1(a)). They can be recyclable (e.g., cotton yarn), reusable (e.g., leather), or neither (e.g., majority of resins). For environmental reasons

materials of this first category are preferable, but in many cases the material selection is dictated by the fabricated artifact.

These artifacts cover different tasks, e.g., soft robots [49] and 3D puzzles [258] used as toys, wire sculptures [292] used as decor, and furniture [21, 251] used in readying an area for occupancy. One way to group the artifacts is by the dimensionality of the output. For example, the wooden portrait in Figure 2.1(a) is 2D, the relief painting in Figure 2.1(b) is 2.1D, and the yarn bunny in Figure 2.1(c) is 3D. Another category that is not present is 2.5D. Another way to group the task into categories is by defining a spectrum between utilitarian and artistic. For example, in the above Figures, the first is the most artistic since it is meant to be admired for its aesthetic quality, the second the most utilitarian since it is meant to be used as a tool for painting exploration by the blind, while the third lies somewhere in the middle, since it could potentially be used for garments, which are functional items.

A third categorization comes from the fabrication process. 3D fabrication is often used to create 3D artifacts. Commonly used approaches are: 3D printers [56, 115, 376, 381], robots [251], CNC milling [482], and casting [26, 385]. We can also include the knitting machine in this category, which built the bunny in Figure 2.1(c). However, researchers have also used 2D fabrication to make higher-dimensionality shapes [41, 91, 159, 208, 214, 215, 273, 303, 332, 393, 417, 427, 433, 434], or to replicate and transfer image properties on 3D shapes [32, 71, 84, 221, 228, 323, 324, 348, 364]. One example that uses 2D fabrication to design a 2.1D image can be seen in Figure 2.1(b). More rarely, 3D fabrication has also been used for the making of lower dimensional output such as wire jewelry [412], or images [180]. Finally, more closely related to this thesis research, is the 2D fabrication of 2D output, and particularly images similar to the portrait in Figure 2.1(a).

This section discusses traditional and computational techniques to design and fabricate images. The benefit of computational approaches, as discussed in Chapter 1, is allowing non-experts access to different design spaces, which I showcase in this thesis by presenting two techniques for the computational design and fabrication of non-photorealistic images. Here, fabrication techniques of images are briefly discussed, specifically painting, printing, and the generation of tactile images. I will first discuss traditional fabrication, then computational



(a) Wooden portrait [190]

(b) Tactile painting [346]

(c) Knitted bunny [306]

Figure 2.1: Stylized fabrication. (Top) input; (bottom) stylized output. This research area includes different production techniques and artifacts, e.g., (a) laser cut 2D portraits [Image Iseringhausen et al. [190] (2020), “Computational Parquetry”]; (b) 2.1D plastic paintings [Reichinger et al. [346] (2011), “High-Quality Tactile Paintings”]; or (c) 3D yarn bunnies [Narayanan et al. [306] (2019) “Visual Knitting Machine Programming”]. Input photographs: [Karl Joseph Stieler (1820), “Beethoven”, Beethoven-Haus, Bonn (wikimedia.org)]; [Raphael, “Madonna in the Meadow”, Kunsthistorisches Museum, Wien, Austria, (wikimedia.org)]; [Narayanan (2019), “Stanford bunny”]

design, and finally computational fabrication of images.

2.1 Traditional Methods for Creating Images

Painting [350] is one of the oldest visual arts in human history. Colored pigments are applied on a surface with different methods, such as hands, brushes, and sprays. Artists have used

different painting media such as oil and watercolor and have mastered different styles such as expressionism, pointillism, and Baroque. In contrast to this type of pigment application which is freeform, printing [142] came much later in human history and restricts pigment application to predetermined flat surfaces. While printing techniques exist that allow artists to make unique artistic pieces, such as monotype and monoprint, printing fulfilled the need to easily create replicas of text and image using methods such as metal etching, rubber stamps, silk screen, stencil, thermal, inkjet, and xerography. More recently, images have been constructed to be touched rather than seen. Tactile images are defined as raised graphics that are felt by hand and have been made manually with a variety of different craft materials, such as glue, yarn, wax, and foil (Figure 2.3(a)).

The first project, detailed in Chapter 3, builds upon two well-established techniques in watercolor painting and woodblock printing. The second project, detailed in Chapter 4 describes a novel way to design tactile illustrations.

2.1.1 Watercolor Paintings

Compared to other painting techniques, watercolor [444] stands out by its use of pigments that are suspended in water. This enables characteristic visual effects based on the flow of the pigments in the water and on the diffusion of the water on the paper (Figure 1.8). One way painters control the flow of water is with *masking fluid*, a paint-like liquid that stops the water flow. Artists also use pigment diffusion to create color gradients. They brush a thick layer of water on the paper and deposit paint of different colors at different locations. The progressive diffusion of the pigments in water creates a smooth color gradation and, at the same time, produces subtle patterns due to the interplay with the paper. The project in Chapter 3 is inspired by these two techniques, and describes a systematic approach for deciding where to put masking fluid and for creating color gradients using water and pigment diffusion.

Watercolor painting follows a subtractive color model, similar to standard CMYK printing. Watercolor painting uses freeform strokes and varies the concentration of paint in water to achieve different transparency levels and colors. Unlike standard printing that relies on

halftoning to generate a large spectrum of colors from a small set of inks, in this project the inks are continuously mixed using water diffusion, which also generates the gradient patterns characteristic of watercolor.

2.1.2 Woodblock Prints

The technique of woodblock printing originated from East Asia as a form of relief printing [72] and has been used for well-known pieces like *The Great Wave off Kanagawa* by Hokusai, which has been replicated by many artists, including by David Bull (Figure 2.2(d)). The artist carves the negative of the desired image (Figure 2.2(a)) on the surface of a flat piece of wood called a *woodblock* (Figure 2.2(b)). Ink is then applied to the raised area and the woodblock is pressed on a paper canvas (Figure 2.2(c)). The ink used in the traditional East Asia technique is water-based, while western woodblock printing inks are often oil-based. The work in this project builds upon the former and extends it to generate watercolor prints. This work proposes a computational approach to the design of the woodblocks and describes how to generate color gradients using water diffusion instead of painting them manually on the woodblocks as done traditionally.

Woodblock printing, similar to watercolor painting, uses different concentrations of paint in water to achieve different transparency. However, it does not allow for the pigments to flow in water; instead, rice paste is used to help pigment attach to the paper. Woodblocks also differ from painting in that the final application of colors on the paper surface is not free but predetermined by the woodblock geometry, which allows for easier reproducibility of the image. Also, in contrast to standard printing, which uses halftoning, woodblock printing uses freeform strokes on the woodblocks to create color gradations.

2.1.3 Tactile Images

One of the many applications of printing is to fabricate 2D tactile images. One way to create tactile images on a paper is using a standard inkjet printer. A black and white image is printed on a piece of a special type of paper, called microcapsule, which later, using light, expands from



Figure 2.2: An example of the traditional woodblock printing process. First, the negative of a single colored image (a) is carved on a flat wood surface (woodblock) (b). Next, ink is applied on the woodblock before is pressed on a paper with the baren (c). The final print is a succession of prints from many woodblocks (d). [©David Bull (2015), (patreon.com/mokuhankan) and (mokuhankan.com)]

the heat in areas where the black ink is placed. Tactile graphics [119] are currently made using microcapsule (swell) paper using infrared heater machines [238], and plastic using vacuum-forming (e.g., Figure 2.3(d)), 3D printing [210, 391, 432], or milling. However, images made with thermoform plastic or microcapsule paper are static. Therefore, to design reconfigurable images, another technology is being developed, namely Braille displays made of pins [317].

Braille displays have drawn considerable attention, but still suffer from low resolution. In addition to this technology, another innovative tool are sketchpads, such as inTACT [80], that allow the blind to be the “designers” of graphics.

Tactile graphics can also be designed by artists [304,343], but are most often made by expert designers in a time consuming process involving identification and application of principles from existing guidelines (e.g., Braille Authority of North America [315]). Rules relate to the use of points, lines, textures, and labels, as well as clutter avoidance, e.g., by elimination, simplification, or separation of graphical elements. However, previous studies have still shown poor 3D shape perception from traditional 2D tactile diagrams. Similar to our formative study (Section 4.2.2), Klatzky et al. [212] cite success rates at the level of 30% for identifying everyday objects, with some improvement observed for textured images compared to outline images [409]. Thompson and Chronicle [410] introduced the TaxyForm system and discuss the need for re-evaluation of design methods. The project in Chapter 4 aims to improve the perception of 3D objects by introducing a new approach that automates the design of tactile illustrations.


2.2 Computational Design of 2D Images

While in the previous section, manual techniques for the design of images were discussed, in this section, computational techniques for the design of images are covered.


2.2.1 Non-Photorealistic Images

Non-photorealistic images are images that do not “attempt to synthesize the field of light intensities” [133] and are extensively used in modern art, technical, and artistic illustrations in design and visualization, e.g., [39, 167, 285].


Artists, crafters, and scientists have extensively used strategies such as simplification, abstraction, and distortion on their designs to create emotional responses in viewers, or to present an idea more clearly. This can be seen from the use of multiple perspectives in the mathematically-inspired woodcuts of Escher [17], the colorful urgent strokes in the paintings of Vincent van Gogh [19], and the simplified Tin-tin of Hergé [18], in which artists create




(a) Craft materials



(b) Plastic



(c) Paper



(d) Plastic

Figure 2.3: Examples of tactile images made with a variety of different materials: a map of Africa (a), an art exhibit (b), a picture (c), and a map (d). [(2020) (wikimedia.org)]; [©Paul Bartlett Ré, “A book and two traveling exhibits for the blind and sighted”, (paulre.org)]; [M. Kunz (1902), ©Perkins school for the blind archives, Wattertown, MA, USA (perkins.org)] [American Printing House for the Blind (2004), “World Maps” (aph.org)]

emotions and tell a story; to the minimalistic human figure sketches for the garments of Gianfranco Ferré [14], the angular design of Alfa Romeo by Marcello Gandini [16, 139], and the technical designs for the machines of Nicolas Tesla [144], which crafters use to understand and fabricate objects; to, finally, the planar and schematic pictures of floating bodies in Archimedes Palimpsest [308, 309], the statistically informed “pie charts” of William Playfair [389], and the simplified illustration of DNA from Crick and Watson [441], which are used to clarify or support an idea.

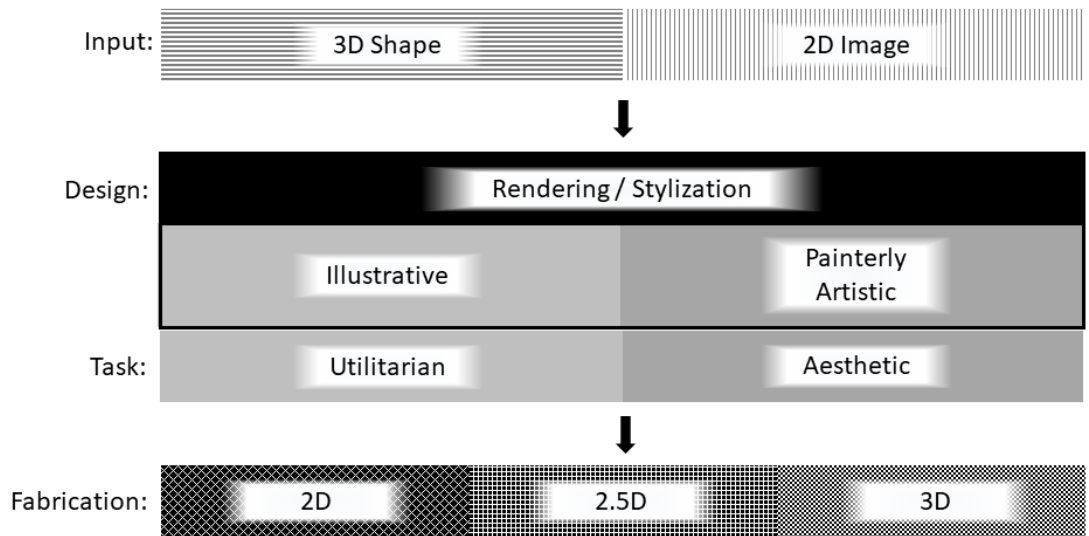


Figure 2.4: *Stylized fabrication of non-photorealistic images. The input can be a 3D geometry or a 2D image. The task for either input can be utilitarian or aesthetic. Finally, the image can be distinguished in terms of fabrication technique to 2D, 2.5D, and 3D.*

Inspired by this type of non-realistic depiction, the discipline of computer graphics developed the field of non-photorealistic rendering (or expressive rendering) in the 1990s [147, 229]. Non-photorealistic rendering [362], used in technical or artistic illustrations in design and visualization [59, 61, 88, 104, 129, 154, 167, 167, 217, 226, 250, 270, 276, 338, 375, 436, 453–455, 467], diverges from reality by modifying the appearance of colors, lines, shading, and views presented [69, 79, 92, 107, 149, 357]. Non-photorealistic rendering is also referred to as stylization. The input is either an image or an object [102] (Figure 2.4 Input). Following, I discuss work relevant to non-photorealistic image creation that is of interest for the projects in this thesis. Moreover, extensive reviews of the developed techniques [147, 175, 229, 394] exist.

2.2.2 Paintings

Related to paintings, there are many works on non-photorealistic rendering that generate stylized compositions that reflect artistic intent [58, 59, 104, 108, 113, 143, 202, 226, 259, 267, 316, 328, 360, 363, 371, 375, 436, 450]. This type of rendering is referred to as artistic rendering (or painterly rendering), and as artistic stylization [102, 225] (Figure 2.4 Design and Task).

Many of these works can be applied to any type of image input and can be adapted to design aesthetically different images only by changing parameters of the underlying algorithms [158,

160, 174, 363], or by providing examples of the desired style [46, 131, 262, 266]. However, this is not always the case. Works in this field have also focused on rendering specific styles on a case-by-case basis by inferring direct rules for its rendering [253, 316, 462]. To allow for changes in the aesthetic style, e.g., level of abstraction, works in this research area are often semi-automatic, similar to the watercolor project in this thesis. One such work is Winkenbach and Salesin [451] where the user specifies where detail should appear in the drawing, in which a 3D input is rendered using black strokes (Figure 2.5). In other cases (e.g., [359]) the input is 2D, similar to the watercolor project in this thesis. One work of particular interest focuses on painting plants [265], in which the output is based not only on images but also on 3D geometric information, and the plants are watercolorized with smooth gradients, similar to our work in watercolors. Going beyond static images, painterly rendering has been used in animation [60, 61, 167, 264, 314]. In the first project of this thesis, instead of creating a virtual watercolorization, we aim to fabricate an artistic version of the input photograph.




Figure 2.5: *The pipelines for non-photorealistic images are often semi-automatic. Here, the input 3D model (a) is annotated with lines to indicate where texture (b) appears in the output (c). [Winkenbach and Salesin [451] (1994), “Computer-Generated Pen-and-Ink Illustration”, ACM]*

Techniques also allow for “deleting” mistakes in already painted areas. For example, strokes can be modified or removed [369]. These techniques, usual for rendering, are obviously not yet realized for computational fabrication. Works have tried to circumvent these issues by using diluted paints and a large number of robotic painted strokes, which, due to the inaccuracy of the results, are guided by visual feedback. This is expensive in terms of time and machinery. For the watercolor project, the “painting” pipeline does not require heavy machinery, and both dense and diluted paints are carefully placed using woodblocks.

Projects about simulation of different painting media, including watercolor, create virtual paintings, investigate the properties of the media, or provide virtual painting tools for non-experts [43,44,86,89,94,118,204,261,423]. In the watercolor project, then, we try to understand the properties of the media we use in order to use them efficiently.

These types of techniques have allowed for the more “formal” and accurate design of images (Figure 2.6(a)). However, in some cases, due to physical manipulation or authenticity requirements, the images also need to be fabricated (Figure 2.6(b)). Select techniques that computationally design and fabricate images are discussed in the following section.



(a) Image design [245]



(b) Image fabrication [245]

Figure 2.6: Examples of computational design (a), and computational fabrication (b) of stylized images for the same input. [Lindemeier [245] (2018), “e-David: Non-Photorealistic Rendering using a Robot and Visual Feedback”]

2.3 Computational Fabrication of 2D images

Image fabrication works (similar to stylized fabrication works, Section 2) can be categorized in terms of the production technique (Figure 2.4 Fabrication), in addition to the input (Section 2.2.1) and the task (Sections 2.2.2, 2.4) as we saw earlier. Additional related works include 3D fabrication techniques that display 2D images, such as wire-sculptures that are meant to be viewed from a discrete set of directions and perceived as 2D images [180] and 3D volumes that cast different shadows in different directions [293]. More relevant to our work is 2.5D- or 2D-fabrication of transformable 2D-images. Some examples are images that encode multiple images [36, 325], that cast different shadows when lit from different directions [47], and

images that transform when viewed from different directions [330] or under different light conditions [30].

Techniques have also been explored for the fabrication of 2D images that remain unchanged. In some of the works, fabricated forms need to interact with light to display images. One category are images that result from shadows of a surface [28], or from shadows to a surface [288, 480]. A different type of image that makes use of interactions with the light is caustic images, which include images of controlled light refraction [269, 326, 368, 472] and controlled light reflection [211, 447]. This type of work falls under the category of inverse design [327].

Other works in fabrication of 2D-images try to stylize or provide user interfaces using specific material. Wood pieces have been used in mechanical art [355], as well as in the computational design of parquetry, which is a traditional type of image made of woods [190]. Images made of strings [52], paper-collage [456], Lego bricks [222], resin sticks [471], strawberry jam [481], and traditional painting media, such as oil paints [132] and pens [237], show the diversity of materials that has already been explored regarding how to fabricate computationally.

2.3.1 Paintings

Various techniques have been developed for the computational fabrication of paintings. Researchers have focused on reproducing existing paintings by 3D printing them [378]. Also, they have used a robotic arm to paint with acrylic ink, oil-paint, and pen strokes for a given input image [75, 112, 247, 248, 413, 415, 461]. Galea et al. [138] created stylized black and white stippling images from a given photograph with acrylic ink applied by drones. Iarussi et al. [183] presented a guidance drawing tool. Yao et al. [465] focused on robot-controlled Japanese brush painting techniques specific to bamboo imagery. Prevost et al. [337] proposed a system for painting large-scale images using spray paint based on a photograph. Our technique is on a smaller scale and aims to create effects specific to watercolors such as diffusion between colors. Flagg et al. [132] propose another system of guided painting, this time for oil painting, which incorporates a color mixing model. Their system is intended as a learning tool and its results rely upon the painter. Instead, we propose a technique that produces results purely contingent

upon the procedure itself. Also, as discussed in the introduction, some of the works do not create a painting based on a given input, e.g., Cohen [100], but composite images, instead of that, we use the watercolor medium to stylize the input photograph, similar to the works of researchers with other painting media, e.g., Lindemeier et al. [247].

2.3.2 Prints

Printing and pressing techniques have also been used in computational fabrication research for broader applications. Standard printers have been used to print ink in a way that controls the specular highlights [274], which results in images with different shadows when lit from different directions, or color inks printed on transparency film to allow for colored holograms [345]. They also have been used in controlling the reflection of images, for example metallic inks to allow for different reflectance [172, 283], or printing color images on silver paper that display different images in different viewing modes [335]. (For a similar result, UV-printers have been used [358].) Fluorescent inks have also been printed to become visible only when lighted with UV and allow for the reproduction of the original color hues of the input image [173]. More recently, 3D printers have also been used for color reproduction [32, 122].

In addition to modern printers, other techniques have also been used to transfer images. Hydrographic printing has been used to physically map the texture on to free-form 3D objects [324, 479]. Similarly, thermoforming transfers pre-distorted images onto plastic sheets, which are then deformed into 3D shapes [364]. In contrast to these techniques, which aim to accurately transfer a given image onto a three-dimensional surface, we focus on transferring a specific pigment type, that is watercolor pigments, to achieve characteristic visual affects. While techniques such as hydrographic and thermoforming have material limitations that cannot be applied directly to watercolor pigments, woodblock printing has traditionally been used to apply water-dilutable pigments. Related to woodblock printing, graphics researchers [405] have looked into how a given woodblock printed image could design a set of woodblocks that would reproduce it. The suitability of using woodblocks to apply watercolors is because wood can attract water, and therefore, pigments diluted in water, in contrast to plastic surfaces. However,

while in traditional woodblock printing color gradients have to manually be painted on the surface of the woodblocks (e.g., black sky gradient in Figure 2.2(d)) and then transferred on the paper, we use the diffusion caused by the water itself to create approximate color gradients without needing to paint them.

2.4 Tactile Images

Visual images, as we show previously, often exist as virtual renditions and do not need to be materialized. In contrast, tactile images are always made with real materials, since otherwise they cannot serve their utilitarian task. Related to tactile images is rendering for illustrative “visualization” [232, 428], which includes technical illustrations and abstracted drawings that clearly convey information (Figure 2.4 Design and Task). While a subset of works inputs images, another subset of works inputs 3D objects. Relevant works focus on enhancing images e.g., modifying sharpness for vision problems [182, 297] and investigating how well line drawings represent 3D shapes [101], for more details [45]. In general, tactile images are self-explanatory without the use of colors, even though in certain cases distinct colors (e.g., yellow and red) are present to annotate information for individuals with low vision. Tactile images can be visualized as greyscale images, with black indicating raised areas of the relief and white indicating non-raised surfaces (Figure 1.9(b)). Therefore, graphics research that uses color shades to convey information (e.g., [169, 349]) is not of relevance in our project. On the other hand, textures are key features for their design and lines remain of great importance, similar to their visual counterparts. Another key element in conveying shape information in rendering is camera placement: their viewing directions, their type of projections, their number and how they are placed and composited together.

When representing 3D shapes in 2D media, choice of camera projections vary significantly depending on the case. For example, orthographic and oblique projections are used widely in CAD and technical drawing design [87, 219], while linear convergent perspective projection is found more commonly, e.g., in video games and film. Linear divergent perspective is used in artistic drawing [179, 196], and non-linear perspective (e.g., curvilinear perspective [395])

projection characterizes wide angle photography and some artistic stylizations, e.g., Escher [69].

Past work in multi-projection rendering [23, 469] shows how images from different camera viewpoints can be combined into a single image. In the tactile images project, we draw from the multi-projection approach of Agrawala et al. [23] and extend their work by using local cameras for segmented parts within the same input 3D shape, the goal of which is to fabricate a tactile illustration with improved shape understanding by individuals with blindness.

Extensive work has been done in computer graphics on visual line drawings that improve how shape and topology are conveyed [155], e.g., using suggestive contours [106], emulating pen strokes on parametric surfaces [452], and using silhouettes and hatching [177, 200, 336, 473]. Studies show the effectiveness of these techniques in depicting shape [101]. However, these methods were designed for visual perception, and properties such as shading due to illumination are not relevant to tactile perception. Further, the cutaneous system for touch is limited in spatial resolution compared to the visual system [53, 233]. High-frequency details are lost when perceiving haptically raised lines. Way and Barner [443] identified size restrictions for the low resolution sense of touch which are enforced in most guidelines for tactile images [166, 287, 315].

Geometric modifications can provide stylization [254], or help visualize internal structure with approaches such as cutaways [242], splitting objects [191], or using deformations [286]. Geometric modifications also help display the global shape of forms, by abstracting the geometry [289], or improving visibility in exploded diagrams [239, 404] (e.g., complex mechanical assemblies). Non-geometric modifications include material changes, such as the use of transparency and ghosting [114]. While most previous works for the design of tactile images do not modify the geometry of a shape, our work geometrically modifies the shape, i.e., uses geometry processing to extract and annotate on the image.

New ideas for the production of tactile graphics covers a range of applications such as maps [67, 152, 397], scientific diagrams [67, 194], computer user interfaces [238], children's books [209, 210, 210, 391, 392, 432], and portraits [438]. A general survey on relevant tools is written by Levesque [236]. Efforts have been made to provide blind individuals with tools for the design and fabrication of tactile graphics [151, 153, 223, 398]. Researchers have also enhanced tactile tools beyond touch, using air [384], force feedback [66, 342, 476], and sound [65, 134],

and have focused on automation for the design of relief paintings [78, 346] and math [70], among other images (Paneels and Roberts have written a review [321]).

Automatic tactile illustration generation has been studied for 2D input images, resulting in improved edge extraction [171], filters capable of removing details too small to be perceived by touch [443], addition of segmentation [128, 230, 440], and tool-specific fabrication, i.e., using a Braille printer [220]. In contrast, our proposed method operates on 3D geometric input and introduces a new non-photorealistic rendering procedure. We were inspired by two lines of work that explored non-standard illustration techniques:

- First, the procedure of “unfolding” or “flattening” 3D shapes to obtain 2D depictions [224], based on drawings made by blind individuals. We expand on this idea with a multi-projection strategy for general 3D objects so that flattened illustrations can be created automatically. Our approach considers viewpoint selection, known to play a key role in 3D shape recognition [126].
- Second, we expand on the use of line curvature and orientation, which have been shown to be salient features [181, 466], used successfully in 2D tactile Braille and Moon alphabets [298]. Thompson et al. [411] studied the use of curvature lines in shape perception. However, examples were constructed case by case and did not provide a system for application on novel objects.

We develop a systematic approach to the design of tactile illustrations.

2.5 Summary

Computational techniques for stylized 2D fabrication of images is an area of computer graphics currently in development. These techniques aim to create an authentic artistic image, or to convey information more clearly. The fabricated images are stylized either to display common elements of artistic composition for the given style, media, or artist or to serve a utilitarian task. Because of the dependence of the fabrication process on the physical materials, techniques that fall under this umbrella are time intensive. Material appearance and handling need to be taken

into account and the design process needs to be tuned (i.e., made aware) for the corresponding fabrication task.

Our main contribution over previous techniques is that we aim to find solutions to design spaces not previously explored in computational fabrication. These design spaces – watercolored and tactile images for shape understanding – are detailed in the next chapters.

It may not be entirely obvious to people in the engineering disciplines, but in the art-making game we do not think about finding the “best” solution to a problem; we think about generating the widest possible range of excellent solutions.

Harold Cohen, 2006

3

Watercolor Prints

Watercolor is a popular form of painting that mixes delicate color gradients with finely detailed contours and diffusion patterns that are immediately recognizable. This unique, abstract look is created by diluting the pigment in water and using a brush to deposit the paint on absorbent paper. Several complex physical phenomena occur during the painting process such as the diffusion of the pigments through the wet paper and the progressive absorption and drying of the water, which combined together generate the visual features characteristic of watercolor. Unfortunately, mastering this technique is challenging and typically out of reach of most people. Further, even for a trained artist, producing several copies of a painting is a tedious task because one needs to manually reproduce all the steps.

Several options have been proposed to overcome this challenge. For instance, one can make



Figure 3.1: A Marilyn Monroe photograph was used as an input to our pipeline (a), with our hybrid fabrication process constituted by watercolor paint and laser cut woodblocks (b). We created the final painting (c) through layered prints of the individual woodblocks. [Photograph by Milton H. Greene ©(2018) Joshua Greene (archiveimages.com)]

copies by scanning an actual painting and using a standard 2D printer. However, this approach creates identical replicas that have lost the characteristic uniqueness of actual paintings. Further, one still has to create the first painting using the standard process. Another approach is to rely on watercolor simulation software, which brings the convenience and controllability of digital authoring to the creation process but sacrifices its physicality. The latter may limit the ability to reproduce all the features of watercolor and, even if one can introduce variation through randomization, each copy is unique only insofar as one does not print each piece more than once. Closer to our work, lithography is a process where the artist first creates an etched stone that is later used to print several copies. Since the print process is manual, each print is subtly different from the others, thereby preserving the unique nature of the work. However, lithography does not apply to watercolor that is the focus of our paper.

In this paper, we propose a new solution in this design space. We enable users to create watercolor prints using a photograph and minimal guidance as input. We first introduce an image analysis technique to determine the main color regions of the image and their level of sharpness. We design this step so that users can easily provide guidance if needed, e.g., to control the level of abstraction of the final print. Based on these data, we then fabricate woodblocks and define the paints and the presence of water on each of them. Our use of woodblocks replaces brushes in the traditional watercolor painting technique, providing a higher level of automation. Finally, we describe a protocol to realize the final print. In our prototype, this last printing step is done by the users themselves but it is fully specified, i.e., it does not involve subjective judgment during the process, and we envision that it could be performed by a robot in the future. We provide a variety of examples created with our approach to demonstrate that it successfully reproduces the characteristic look of watercolor without requiring any painting skills from the users. Recognizable characteristics of watercolor paintings are present in our results: the paper texture, the variability of the thicknesses on the deposited pigments, and the broad gamut of the produced colors.

3.1 Background

In this project, we focus on watercolor paintings, which is a particular artistic style of visual images. In visual images, we can access information that we could not sense without our sight, including color, variation in color intensity, and access to distant information. Also, when viewing visual images, we can accurately and quickly extract edges and, with their help, semantic boundaries. Artists have used this type of information to convey form and emotion in their paintings. In this project, we also extract this type of information to computationally design watercolorized versions of a given photograph. Before explaining our process, we discuss graphics research that has taken this information into consideration.

3.1.1 Color

Color is a salient characteristic which we use to visually distinguish objects [53]. Artists have used color to indicate different objects and to provoke emotions. To achieve high aesthetic quality, artists carefully select the number and the combination of colors they use. Too many colors, even though they might result in a more accurate photorealistic image, might also create a less artistic, less harmonic, and less aesthetically pleasing result.

Color and color perception have been used in painterly rendering [108,316] and researched in computer graphics. The most relevant type of work is the techniques that have been developed to extract palettes of representative colors and segment the image in several uniformly colored regions. While Morse et al. [301] select colors based on color hues, saturations, and lightness values, O'Donovan et al. [313] use a database which includes rated color themes, and Lin et al. [243] create color palettes based on a data set of human-extracted color themes. We use the palette extraction algorithm of Chang et al. [81], which in contrast to other techniques, does not rely on existing color databases and is a simple algorithm to implement, which extracts a palette based on pixel color clustering. Palette extraction for painterly rendering has been looked into prior to our work (e.g., [264]), especially for the case of watercolor, for which traditional painting techniques use a limited number of colors. Color palettes have been used to recolor watercolor images [24]. Color mixing based on limited color-palettes has been recently

examined in fabrication but has not been researched for watercolors [246].

3.1.2 Variation on Color Intensity

Variation of intensity, similar to shading in painting [213, 341], helps shape understanding. One example of this phenomenon is that two-colored images have been shown to fail in recognition tasks, where gray-scale images succeed [53]. Artists have used shading, i.e., gradual color intensity variations, using media such as pen, pencil, and watercolor, to display different materials, shadows, and 3-dimensional forms on the 2-dimensional space of a flat piece of paper. In this subsection, we focus on how we capture the gradual variation of color intensity on an image for the purpose of reproducing it with watercolors.

While computation of gradients has been used in graphics, and especially relevant to our work in painterly rendering, for example to place outlines in line drawings [201], to decide size of brush-strokes [263], or to represent gray-scale gradients using lines and visual feedback [415], identifying where gradual color variations occur (smooth gradients) in order to fabricate them, has not been addressed in the field of graphics. In a similar line of work, Luft and Deussen [265], noticing that smooth color transitions were desirable for watercolors, proposed a way to design plants with smooth gradients by considering the underlying 3D geometry. Similar observations were made for general purpose scenes [60]. Our project adopted and extended the following three techniques in order to identify where the places of gradual color variations occur on an image. One of these techniques, defocus estimation, was used to identify regions of constant blurriness in an image [127, 478, 485]. Depending on how blurry or out of focus these regions are determines whether they are shaded or not; that is, they are shaded as areas of variation of gradual color intensity are identified (i.e., smooth region). The second technique is the identification of relatively higher or lower color intensity areas within a uniformly colored smooth region. These areas can be found using image morphological analysis, such as the Valued Regional Maxima Image Filter provided by ITK [7] used for medical imaging, and are further expanded using our technique explained in Appendix B.1. The third technique is the constant overlaps we define for the woodblocks in order to create smooth transitions between

different colors.

3.1.3 Edges

Image edges, discontinuities on the color intensities of an image, help people understand object boundaries and provide information about shapes [31]. Artists have displayed edges as lines or color discontinuities in their paintings and drawings.

Similarly, many computer image processing techniques have used [108] and looked into how to extract edges from images. A commonly used technique is the Canny edge detector [76], which involves computations based on gray-scale image gradients. Techniques of edge extraction can approximate the gradient, matching predefined templates, or fit local intensities with edge models according to Ballard and Brown [35]. In addition to these, recent techniques have employed neural networks [121] for edge computation. Due to its existence in many graphic packages, we also use the Canny edge detector, which gives satisfactory results for the purpose of our work.

3.1.4 Boundaries

Along with the use of the two perceptually lower level characteristics discussed previously, color and edges, a higher level characteristic relates to semantic regions, enclosed within what are called boundaries, and helps identify objects on a higher level. Boundaries are created where edges occur and signal semantically different parts. Watercolors introduce edges at the painted boundaries as they dry. Therefore, artists, when they do not want to introduce edges within a semantic area, finish painting a semantic area quickly before they move to the next area.

Separation of images into semantic areas has previously been used in painterly rendering [111, 247, 474], and has also been investigated in graphics. Some of the segmentation techniques use a seed as input which later grows the regions, taking into consideration local or global characteristics. One example is the work [218], in which a global model defined on a graph of image pixels propagates the regions to the entire image using Bayesian probabilities. This is also the algorithm that we use in our work. More recent techniques have resolved the

need of adding a user given scribble (e.g., [25]). An extensive review of existing techniques can be found in the survey of Vantaram et al. [425].




Figure 3.2: Image analysis Pipeline. The inputs are the original photograph I_p (a), and the user semantic annotation I_s (b). The colors in the output images represent print colors depicted in (d), the white and the chess pattern in (g)-(j) correspond to non-printed areas.

3.1.5 Overview

The input to our technique is a digital image, I_p , and user-provided annotations, I_s , of the process. We designed a pipeline having in mind the abstraction of watercolor paintings. First, we analyze the input image to determine how each region should be represented in the final print. We compute a palette composed of a few paint colors and determine where each color will be applied. We also estimate the local smoothness of the image and use it to locate where to put water. We combine these pieces of information to define a set of regions where we will place paint or masking fluid. We give the shape of these regions as input to the laser cutter to carve the woodblocks. Finally, we use these blocks to apply the paints and masking fluid onto white paper. While this step is manual, it is fully deterministic and could be performed by a robot. However, the design of such robot is beyond the scope of our work. Instead, we describe a step by step procedure, in which no subjective evaluation is required at any point

3.2 Image Analysis

This section describes our image analysis methodology to generate the components of our watercolor prints. We later use the output to fabricate the woodblocks for printing.

3.2.1 Overview of Our Image Analysis Pipeline

The input to our technique is a digital image, I_p , and user-provided annotations, I_s , indicating the main semantic regions that the user would like the final print to retain. For example, in Figure 3.2(a), the user selects distinct regions for the background, hair, skin, eyes and lips. Using these inputs, the pipeline of our image processing algorithm consists of the following steps (Figure 3.2):

1. Decomposition of the input image into a small set of color regions that we use to determine the color of the paint at each location.
2. Estimation of the local smoothness of the input image.

¹Photograph by Milton H. Greene ©2018 Joshua Greene www.archiveimages.com




Figure 3.3: *The user can control the level of abstraction by varying the number of colors, k , in the color palette extraction.*

3. Decomposition of the image into semantic regions based on the user’s annotation. We use this information and the local smoothness to locate where to put water.
4. Detailed analysis of the variations of color and intensity in the water regions, which we use to generate watercolor diffusion patterns that approximate the content of the input image.
5. Optional: characterization of the fine details not represented by the previous step so that we can print them over the diffusion areas once the water has dried.

Throughout this process, we create *maps* that store intermediate data, e.g., a *color map* for the representation with a small color palette and a binary *water map* to indicate where to put water. The output of our image analysis is a set of 2D shapes in the image plane describing the woodblocks that we will use for creating the print. We will send these shapes to a laser cutter to fabricate the woodblocks.

3.2.2 Color Map - M_c

While watercolor exhibits a spectrum of colors, they are all produced with a small set of paints – the color diversity comes from mixing, diffusion, and dilution. The first stage of our image analysis pipeline determines this small set of paint colors, *the color palette*. We use the palette extraction algorithm of Chang et al. [81] which takes the photograph bitmap file and the desired number k of colors in the palette as input, and outputs the actual palette colors. This variant of the k -means algorithm uses an acceleration data structure, where k is the number of colors in the output palette. In our context, the number of colors k controls the level of abstraction of the

final watercolor print. A low k value represents the input image with few colors and generates a more abstract output, whereas a high k value preserves more details. We typically set k between 3 and 5 (Figure 3.3). Next, we compute the *color map* M_c by assigning each pixel to its closest color in the palette in the CIE-Lab color space. This map defines the color of the paint at each location of the final print. The next step is to determine where we will apply water.

3.2.3 Dry and Wet Areas

Now that we have computed the color of the paint at each point, we decide where to apply water to create diffusion gradients. In “dry” areas, we use watercolor paint with a minimal amount of water so that no diffusion occurs. This is the desired behavior for printing fine details like the eyes of a portrait. In comparison, “wet” areas favor diffusion and are well suited for smooth color gradients. Because of the diffusion process, controlling the output in these regions is more challenging and we introduce a dedicated multi-step process for that purpose. In the rest of this section, we first describe how we determine where to put water, then briefly discuss how we handle the dry areas, and finally discuss in detail the more challenging case of the wet areas.

Where to Put Water

We wish to use water diffusion in smooth regions of the input. To do so, we first estimate a *smoothness map*, M_{sm} , that quantifies the local image smoothness and then describe how we guide users to decide where to put water based on this information. We call the output of this step the *water map*, M_w .

Smoothness Map – M_{sm} We use the algorithm of Zhang et al. [478] to estimate the local smoothness in the input image. This technique has been designed for characterizing defocus blur and we found that it generalizes well to locate smooth regions in general, e.g., it successfully identifies the Marilyn Monroe’s skin as smooth (Figure 3.2c, brighter values are for smoother regions). The algorithm analyzes the gradients in the image and outputs a continuous sharpness value at each pixel, we call this result the *smoothness map* M_{sm} . We discuss the parameter settings in the section about implementation details.

For our application, we seek to convert these continuous values into binary ones that indicate where to put water. This is a critical decision because the transition between a dry and a wet region creates a visible edge that, if misplaced, generates unsightly results. We experimented with options like per-pixel thresholding and superpixel-based thresholding, and found that they do not prevent such unsightly edges. These early tests also suggested that the location of the diffusion gradients is in part a subjective choice and that some degree of user controllability is desirable. The next section describes our solution motivated by these observations.

Semantic Map – M_s To ensure that wet–dry edges appear at visually adequate locations, we let the user define the meaningful regions of the input image, e.g., the sky and the trees in a landscape, and later only allow dry–wet edges to appear between two such regions. In the previous example, this guarantees that no edge can appear in the middle of the sky for instance, but allows for an edge between the sky and the trees. For this purpose, we use the DenseCRF algorithm [218]. This technique takes the input image and a set of user-provided scribbles, one for each region they wish to define, and it returns a dense map where each pixel is assigned to a region according to criteria like color similarity and image distance between pixels. In our experiments, we found that regions representing semantic image elements such as skin, eyes, and mouth for a portrait, or sky, trees, and ground for a landscape yields satisfying results. Based on this observation, we call this output of the DenseCRF algorithm the *semantic map*. We refer to the original article for the details of the algorithm and to Section 3.4.1 for our parameter settings. Semantic image segmentation has also been incorporated in previous work in rendering [483] and physical painting [249].

Water Map – M_w Using the semantic map, we next define a binary map that indicates where to put water, the *water map*. To ensure that dry–wet edges only appear along semantic boundaries, we assign a single smoothness value to each semantic region by averaging the smoothness map M_{sm} over each region. Without the Smoothness Map, there would be no indication on which semantic regions to put water and users would have to provide this information. On complex images, this may be a challenging task for novices. We let the users define the water map with

a threshold on this per-region value. By varying the threshold, one can choose a look from a sharp rendition with details everywhere to an abstract print composed mostly of water diffusion gradients.

Dry Woodblocks

Now that we have computed the color and water maps, we define the shape of the woodblocks that we will later fabricate with a laser cutter. We first discuss the “dry woodblocks,” that is, the blocks that we will use to deposit paint on paper that is only lightly moistened (in comparison to the areas where it is fully wet). We create one such block for each color of the color map by computing the intersection between the corresponding color region and the inverse of the water map (i.e., the dry regions).

Wet Woodblocks

In this section, we discuss the woodblocks that we will use on wet paper. These require more sophistication because of the water diffusion. First, we create a *masking woodblock* by inverting the water map M_w and adding edges (Section 3.2.3). We will use this block to apply masking fluid on the regions that should remain dry. Then, we create woodblocks representing each color in the color map. We first describe how to design the blocks to produce a smooth color gradation and later refine this process to also handle brightness variations.

Color Woodblocks – B_{wc} In our early experiments, we created a wet block for each color by intersecting the corresponding color region in M_c with the water map M_w . This applies paints of different colors next to each other and, with water diffusion, creates a smooth transition between them. However, the extent of this transition is limited. We tested different amounts of water and paint, but the result never exhibited the large gradients typical of watercolor. We tried to shrink the shape of the woodblocks before fabricating them. We hoped that this would leave a gap between the colors that diffusion would fill, thereby creating larger gradients, but instead this created gradients with a white band in their middle. The opposite solution succeeded. We dilate the color regions of the color map with a disk of radius 1.75 mm, before computing the

intersection with the water map. This creates areas where colors overlap, essentially introducing an intermediate color between any two adjacent colors. Coupled with water diffusion, this produces large smooth gradients. By varying the amount of dilation, the size of the overlap changes, which in turn affects the extent of the gradient. The resulting color regions correspond to B_{wc} as shown in Figure 3.2(h), the red areas indicate where the adjacent colors overlap. See Figure 3.10 for a comparison of gradient effects.

Shading Woodblocks. Overlapping blocks allow us to create a smooth gradation between two colors. We also observed that gradients in the input often exhibit an intensity variation, i.e., they go from dark to bright to suggest shadows and highlights. To reproduce this effect, we refine our woodblock design. First, we compute the block shape as previously described using the *water map* (Section 3.2.3) and consider the corresponding image region in I_p . We convert the image data to gray-scale and generate a *grayscale map* in the same way as we have created the color map, i.e., using the k -means algorithm [81]. Since we operate within a region of the color map, we seek a finer-grain decomposition and generate a larger number of clusters for $k=50$ grayscale values. We call superpixels the groups of adjacent pixels with the same grayscale value. First, we identify the brightest superpixels using non-maxima suppression, that is, we select the superpixels that are brighter than all their adjacent superpixels. We modify the block shape by removing these bright superpixels from it. As a consequence, these regions will not receive paint directly and color pigments will only reach them in a diluted form via diffusion, which will reveal more of the underlying white paper and create a brighter spot as desired. These brightest superpixel regions are denoted by the checkerboard pattern in B_d in Figure 3.2(j).

For the shadow regions, we identify a set of darkest superpixels using non-minima suppression, on the gray-scale k -means map for $k=30$ gray-scale values and for each color block using a procedure that we describe in Appendix B.1 we expand them. These shadow regions correspond to the colored areas in B_m in Figure 3.2(i). We will use these blocks to apply a darker shade of paint that will smoothly blend because of the water diffusion.

Detail Woodblocks

So far, we have worked with large and medium sized superpixels, and have not captured fine details. While it is common for watercolor imagery not to exhibit such details, artists sometimes add them back using a fine brush. We provide a similar option. From dry areas, we allow users to add details by extracting the edges of the input image using the Canny edge detector and printing them on top of the color regions. For this purpose, we create a dedicated woodblock for each color region and use it to apply a darker paint of the same color. Water areas exhibit even less details because of the water diffusion. First, we add the Canny edges to the masking fluid block to create discontinuities in the diffusion gradient, thereby making these edges visible in the print. Also, the watercolor block expansion stops wherever an edge is present. Then, we extract small image features. We use k -means clustering [81] on a gray-scale version of the image for $k=50$ gray-scale values to create small image regions and we select the dark regions using non-minima suppression. We create a woodblock with these small regions in addition to the edges of the dry region and use it to apply a darker paint of the same color once the water has dried.

3.3 Fabrication Method

In the previous section, we have computed the shape of the various woodblocks needed to create the final print. We now explain how to fabricate the actual woodblocks and how to create the print.

Woodblock Cut Pattern Given the image analysis output from Section 3.2, we translate the output images (Figure 3.2(g)–(j)) into woodblock cut patterns that we send to a laser cutter (Figure 3.4). The images indicate the carved regions in black. While a laser cutter can carve arbitrary patterns [235], thin wood features are fragile and likely to break under the action of the printing press that we use in the final stage. To prevent breakage, we shrink the carved regions by 0.05 mm, to account for the laser spot radius of 0.05 mm.

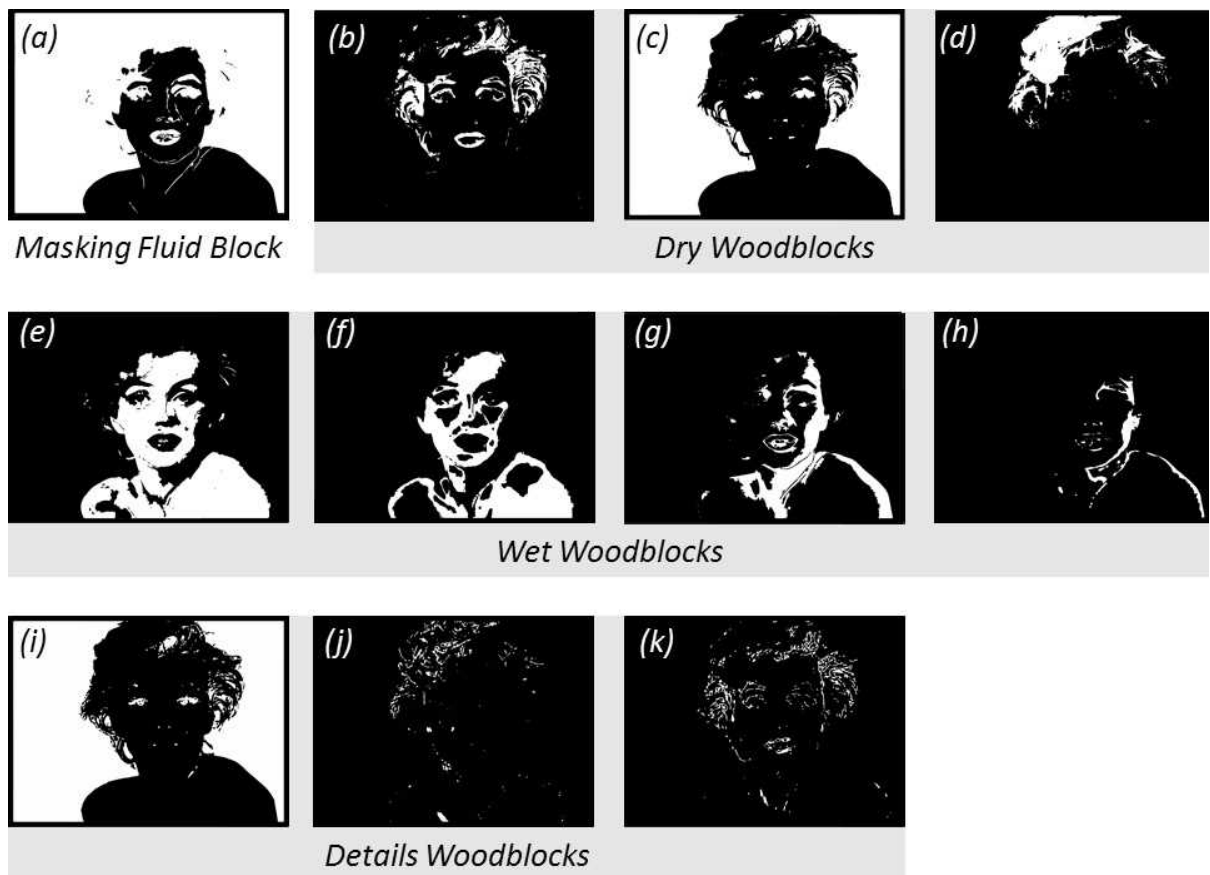


Figure 3.4: Images given as input to the laser cutter, corresponding to the woodblocks in Figure 3.1.

Creating the Print We now describe how to use the woodblocks to create the final print. First, we affix the paper onto a flat plywood board to ensure no distortion when we apply water to the paper. Then, we apply the masking fluid using the corresponding woodblock, i.e., cover the woodblock with the fluid, position the block on the paper, and press it against the paper. Once the masking fluid is dry, we spray water on the paper and we print the wet color regions. Then we spray water again to allow diffusion. Once the print is dry, we remove the masking fluid and print the dry color regions and the details. The printing sequence is illustrated in Figure 3.5 and we reference our equipment in the Implementation section.

Discussion: We observed that the order in which we print the wet color influences the final result. The last color we apply tends to “dominate,” e.g., red over yellow gives a redder orange than yellow over red. For consistency, in our experiments, we always applied the colors from brightest to darkest.



Figure 3.5: *An example of our printing sequence. (a) We start by masking the “dry” areas without color diffusion (white) and printing the “wet” colors with water diffusion first. (b) We then remove the masking fluid and print the dry colors without water diffusion. (c) Lastly we print the details blocks (minima and edges).*

3.4 Implementation Details

We describe the implementation details of our algorithm and list materials and machinery used for fabrication. Before we finalized them we looked into different other implementation and fabrication ideas, some major can be found in Section 3.6.

3.4.1 Parameters of the Image Analysis

To promote a straightforward process in the image analysis stage, the parameters that the user controls are limited to the threshold for the Water Map, the k number of colors in the Color Map, and the Semantic annotation scribble.

For the k -means palette extraction algorithm of Chang et al. [81], we made modifications so that dark colors are not discarded and shadows are not ignored. We use 10 bins for every color channel.

We use the publicly available implementation [4] of the DenseCRF algorithm [218]. In all our examples, we set the weight and the standard deviation for the smoothness kernel (color independent terms) as $w^{(2)} = 3$ (unitless) and $\theta_\gamma = 3$ pixels, respectively. For the appearance kernel (color dependent terms) the weight, standard deviation and color standard deviation are: $w^{(1)} = 10$ (unitless), $\theta_\alpha = 80$ pixels, and $\theta_\beta = 19$ (RGB channels $\in [0, 255]$).

The algorithm for estimating the local image smoothness [478] required more adaptation. We observed that, on our data, producing consistent results was challenging with the settings

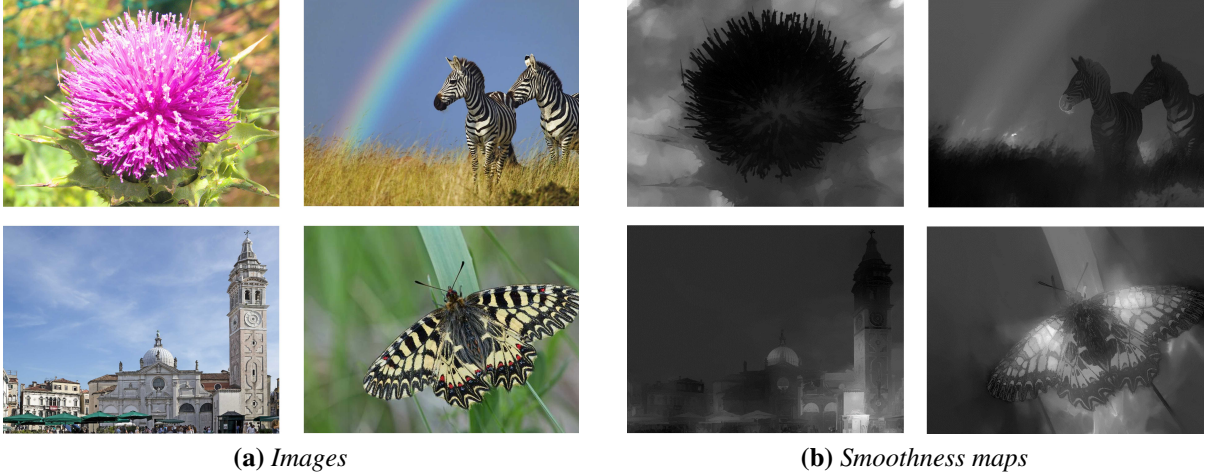


Figure 3.6: 4 out of 20 images we used to identify the values for the parameters for the smoothness map algorithm

suggested in the original article. We conducted a parameter search using 19 test images and picked values that performed well while being robust to perturbations, i.e., results produced with slightly different parameters were also acceptable. We set the standard deviation for blur to $\sigma = 10$ pixels, the Canny [76] thresholds to 20 and 100 (for pixel intensities $\in [0, 255]$), and $\epsilon = 10^{-6}$ to compute the defocus value. For the gradient we use the Sobel function of OpenCV. For the guided filter [168] we set $\epsilon = 10^{-9}$ and the radius $r = 20$ pixels. For the KNN matting [85] we set $\lambda = 0.2$, which is a constant controlling the user’s confidence on the scribble. We set the level to $l = 5$ which is a unitless constant that controls the amount of spatial coherence, we use 10 neighbors (pixels) with default spatial coherence and 1 neighbor with weak spatial coherence. We consider all the values of the user annotation with a defocus threshold of $t = 10^{-5}$. Figure 3.6 shows a few examples of smoothness maps computed with these parameters.

We use the OpenCV implementation of the Canny edge detector with parameters 34 and 45 (channel $\in [0, 255]$). We applied a morphological dilation of radius 0.5 mm followed by an erosion of the same size (a.k.a. closure) to fill in the small holes that may exist in the edges. We discarded the edges between the wet and dry regions because they are already visible as discussed previously.

3.4.2 Design of the Woodblocks

The images from the Image Analysis were vectorized using path trace in Adobe Illustrator. We resized them in accordance with the dimensions of the machine press. Marilyn Monroe was printed in 24 cm × 16 cm blocks, while all the other results in 20 cm × 15 cm blocks.

3.4.3 Fabrication Details


To carve the woodblocks we used a Laser Pro Spirit GLS with 170 pulses/cm or dots/cm, in raster, with 90% power and 35% speed. The maximum laser processing speed and power are: 203 cm/s and 100 W respectively.

We used Arches hot pressed paper 300 g/m². For the paint materials we used Winsor and Newton Colorless masking fluid and Winsor and Newton Cotman paint. In many cases we were able to find paint tints that closely matched the color palette, otherwise we performed color mixing using a guidebook by Wilcox [449]. Light colors used a low ink-to-water ratio and darker parts used a higher ratio (e.g., 1 : 4 for light areas and 1 : 1 for dark areas).

To ensure a flat surface for printing, the paper was stretched on plywood. Before printing the wet color regions, the paper must also be moistened to allow absorption of the watercolor paints and permit diffusion effects. We used an aerosol sprayer, and performed a series of spraying and absorption iterations (see the supplemental material for further details). To perform the printing we used an Albion press, which is an early iron hand printing press that operates with a simple toggle action. We experimented with different amounts of pressure. Low pressure was used for the water areas for improved diffusion gradients. We found that high pressure gave sharper results in the non-water areas.

3.5 Results

We first illustrate the effects of key elements of our approach. In Figure 3.7, we show that without our complete pipeline, either spurious edges appear, shading is poorly represented, and/or diffusion patterns are inadequately shaped. Figure 3.8 gives an example of a print with



(a) Without semantic regions, spurious edges appear.

(b) Without edges masked, diffusion covers face features.

(c) Without shading blocks, shadows are weak.

(d) Our result.

Figure 3.7: Removing stages in our pipeline negatively affects the results of our algorithm. In this figure, non-skin areas are digitally masked out for clarity.



(a) Print without details



(b) Print with details

Figure 3.8: We let users choose whether the print has fine details, e.g., the feathers in this case. The difference in the color gradient on the bird body and wing comes from the inherent variations that occur between prints.

and without using detail woodblocks. Figure 3.9 shows three prints made with the same process. The variability inherent in our process makes each print unique and, at the same time, all the

prints remain qualitatively similar to the input. Figure 3.10 illustrates the benefits of overlapping woodblocks for creating smooth color gradients.




Figure 3.9: *Two examples of repeated prints to illustrate the variations inherent in our approach. The top row is the Hibiscus flower and the bottom is the Robin bird. All the results look stylistically similar. However, due to the liquidity of watercolor unique variations emerge.*

Figure 3.11 shows a set of diverse prints. We chose the portrait because artifacts on faces can quickly look unpleasant; this result shows that our technique can accurately combine fine details and smooth gradients. The Hibiscus flower illustrates our ability to produce smooth gradients with both color and intensity gradations. We picked the Robin bird for the 3-color gradient on its body and its thin legs. The lake cabin demonstrates that we can also produce prints with extensive details. We masked out the background region for the bird and flower examples.




Figure 3.10: *Print of two color blocks. Different strategies to create color gradients. The top row shows actual prints and the bottom row shows a visualization of the two blocks that we fabricated. Overlapping blocks (c) produce the smoothest gradients. This is the option that we use in our technique.*

3.5.1 Results Summary

Here, we present for completeness all material for the watercolor prints: Marilyn’s portrait, Hibiscus flower, the Robin bird, and the House on the lake. The output of the Image Analysis pipeline are placed in Figures 3.12, 3.13, 3.14, and 3.15. The image sequence from top left to bottom right is: input photograph, color palette, user semantic annotation, semantic map, defocus map, smoothness map, water map, overlap watercolors, local minima, local maxima, details, and edges. The designs of the woodblocks are placed in Figures 3.16, 3.17, and 3.18, and the woodblocks of Marilyn’s portrait in Figure 3.4.




Figure 3.11: *Examples of watercolor prints with a variety of input photographs. (Row 1) Portrait of Marilyn Monroe[‡]. (Row 2) The flower is a representative example of creating a color gradient. (Row 3) The bird has diffusion of more than two colors (Row 4) The landscape contains detailed textures which are produced with the detail woodblocks printed on top of the dry blocks.*

[‡] Photograph by Milton H. Greene ©(2018) Joshua Greene www.archiveimages.com

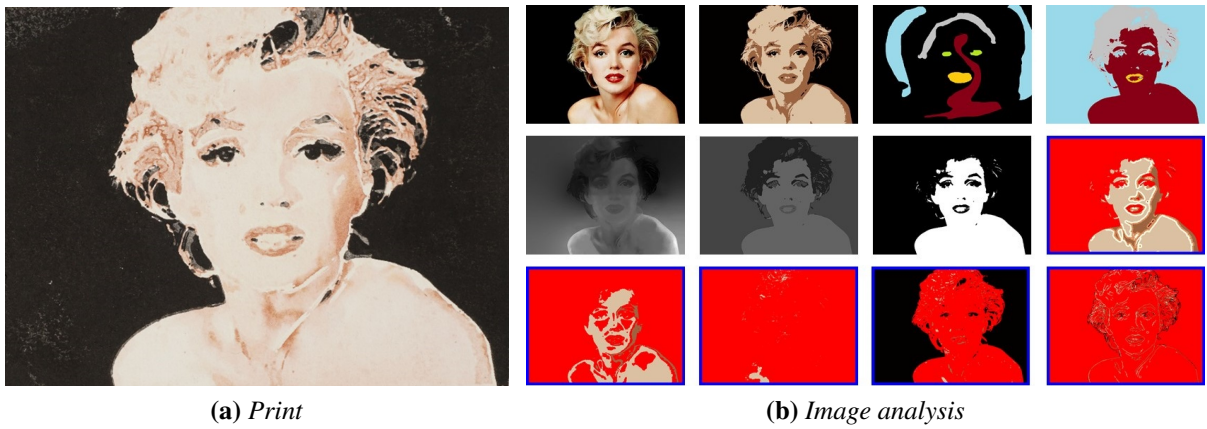


Figure 3.12: Marilyn Monroe's portrait. The shading shows the curvature on the face, the neck, and the shoulders; Marilyn is immediately recognizable even with the absence of the blond hair and the red lips.




Figure 3.13: The Hibiscus flower. The shading has two color gradations and non-painted areas on the petals that improve image quality.

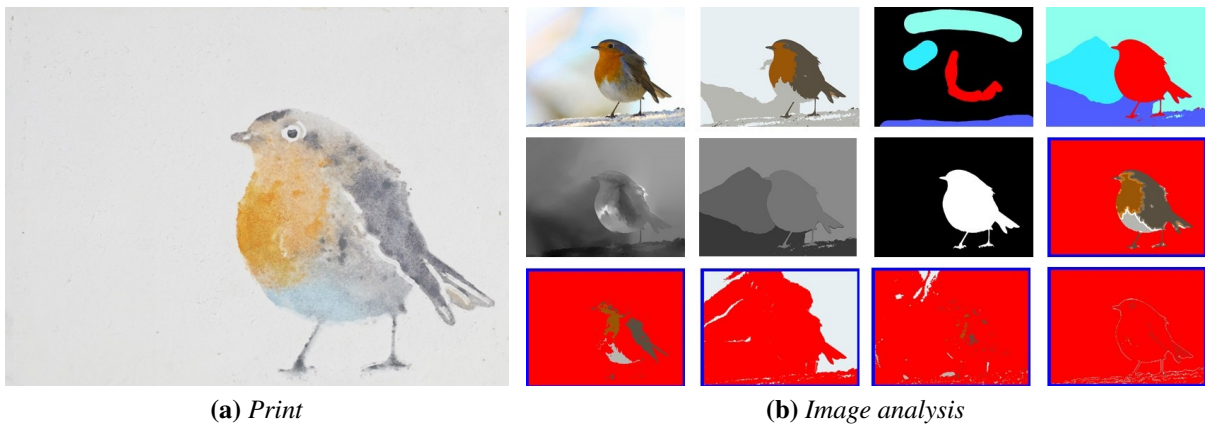


Figure 3.14: The Robin bird. This example, entirely made of watercolor, has a three-color gradation on the middle of the feathers.

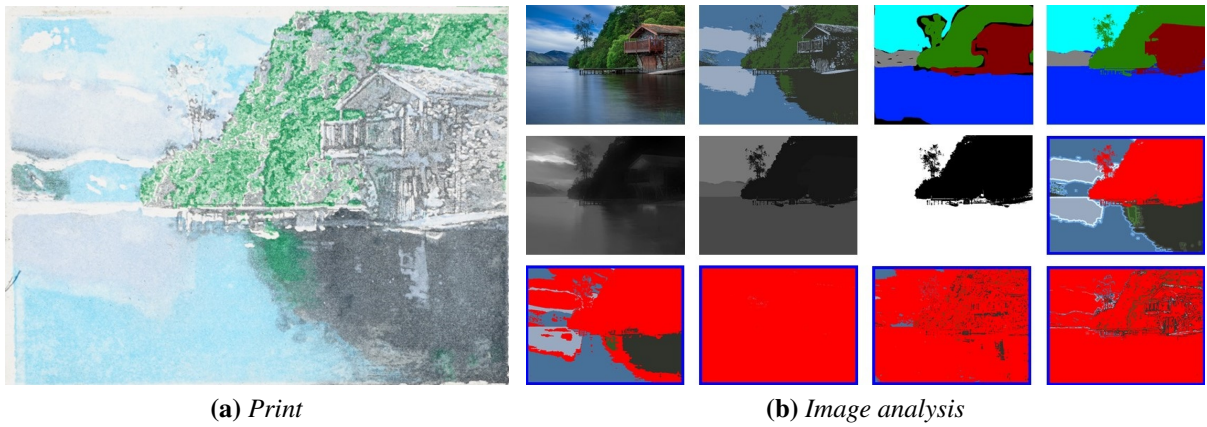


Figure 3.15: *The house on the lake. Detailed texture was achieved by the use of two rather than one dry woodblocks for each color, making the image to appear less flat.*

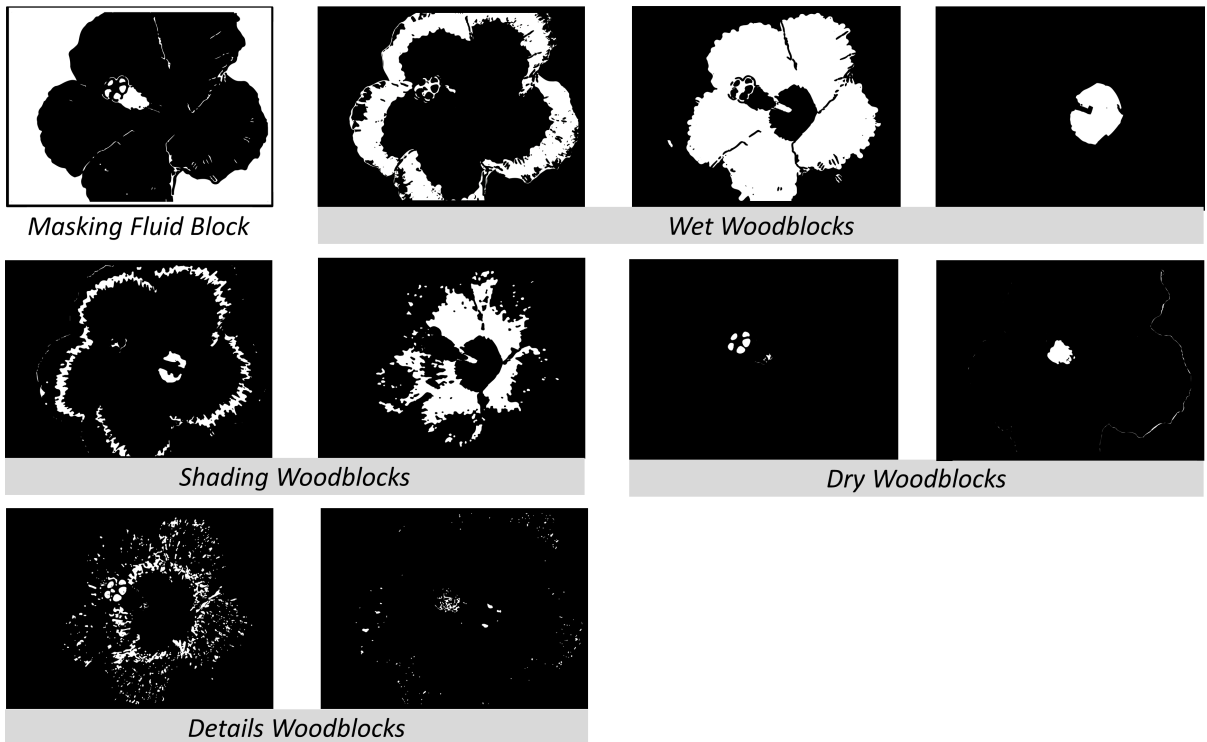


Figure 3.16: *The blocks for the Hibiscus flower. The background is not printed, and three colors are used: yellow, pink and brown. In some of the woodblocks, more than one color is carved to minimize the consumption in fabrication, e.g., the first shading woodblock prints both yellow and brown.*

3.5.2 Informal User Study

We performed an informal user study with four participants from the local art community, where we interviewed college-level art majors, and students that had taken at least one class on printmaking or painting. We asked for feedback on three results including the Marilyn portrait,

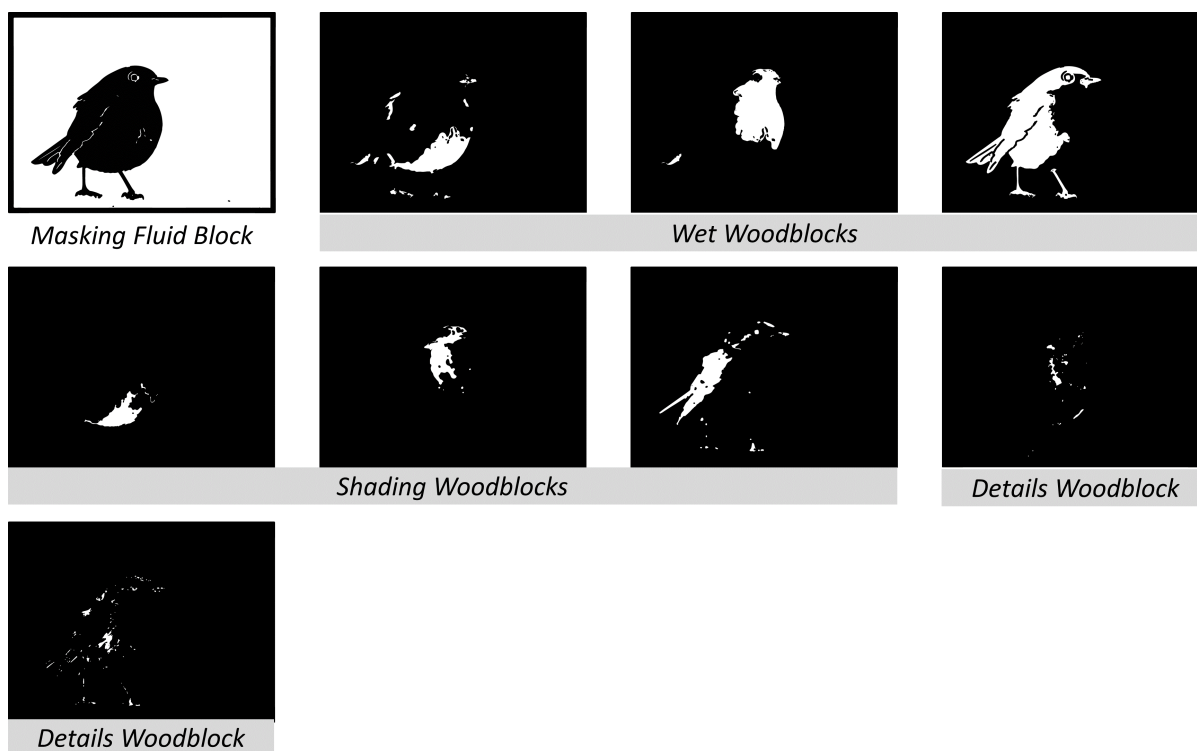


Figure 3.17: *The blocks for the Robin bird. The background is not printed, and three colors are used: light blue, orange, and gray. Also, there are no dry woodblocks, except from the details woodblocks which in some of our prints were printed with water.*

the flower, and the bird. We did not describe how the paintings were made.

All participants had positive reactions and specifically commented on the color diffusion effects and the shading. They remarked on the “creation of roundness,” “illusion of depth,” “nice washes of color,” “color blending,” and “interesting color overlapping effects.” All participants liked the color palette selections with two specific exceptions: students would have preferred a darker brown color in the flower print, and in the Marilyn print they would have preferred the black regions in the hair to be lighter since it appears to be background rather than shading. Participants also commented on the sharp edges and use of white lines between color regions (results of the masking fluid). An art major commented on the flower: “the sharp border is very cleverly placed because it makes the stem pop out.” Another art major liked the negative space created by the borders. A third student noted “the borders add to the abstraction and simplicity in a good way.”

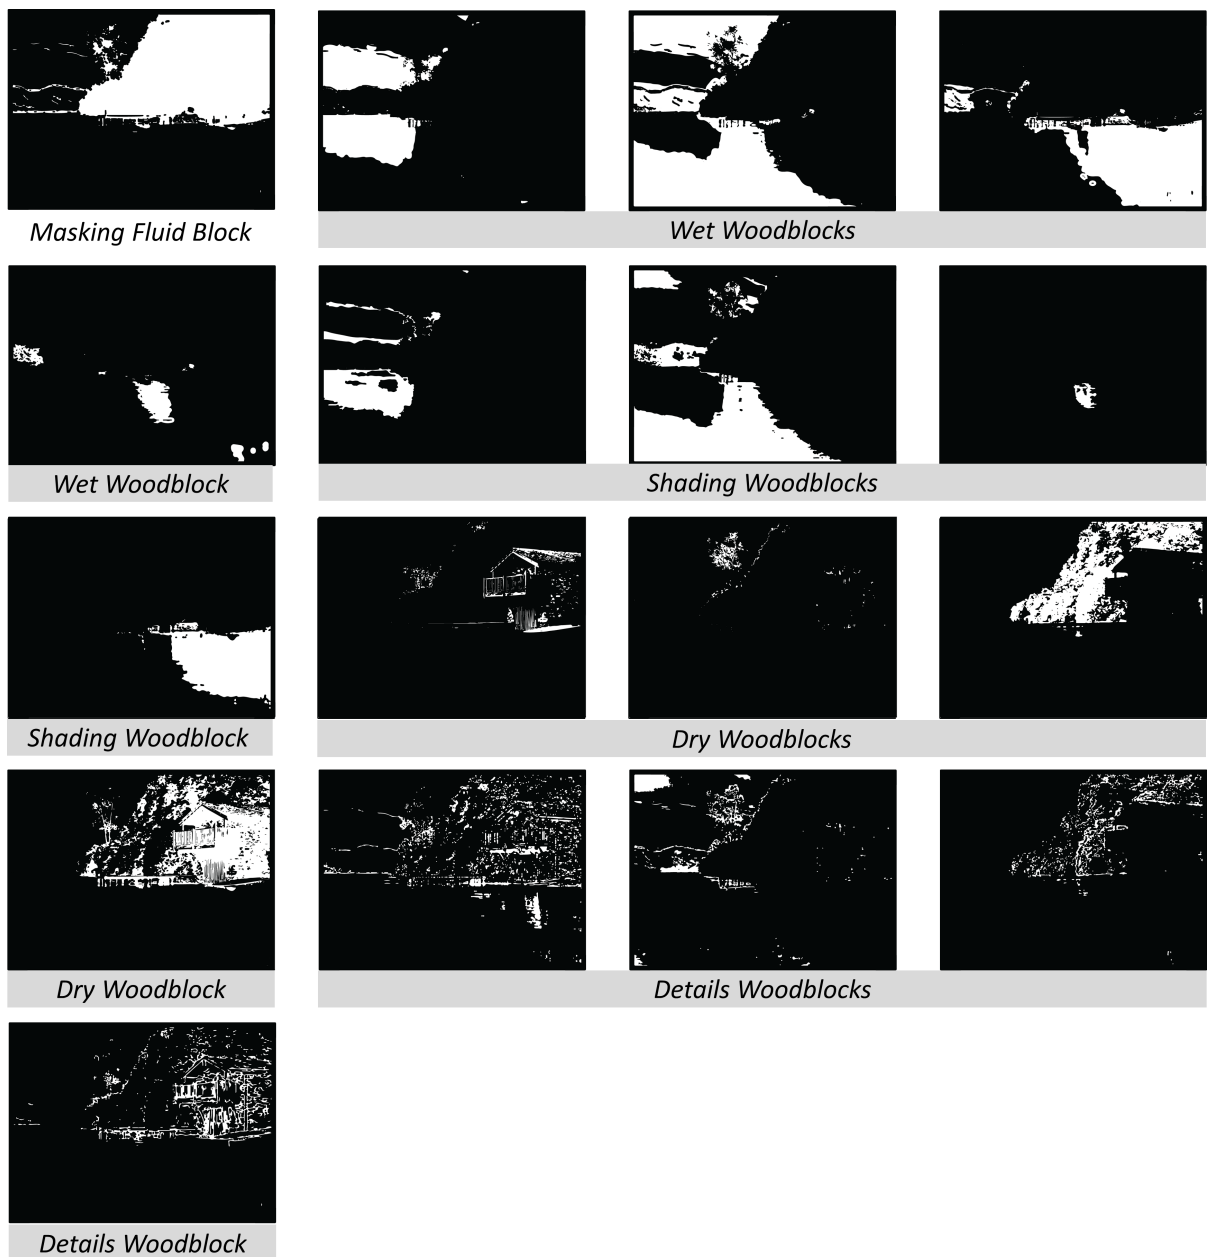


Figure 3.18: *The blocks for the house on the lake. This large number of woodblocks was a result of the more extensive color palette (light gray, blue, green, and dark gray) and made the printing process more challenging. The wet woodblock printing in particular required a quick process for the paper to remain wet.*

3.6 Exploring Alternative Paths

3.6.1 Masking the Paper

The paper needs to withstand heavy dampening with water without tearing apart, to allow for the diffusion of the colors and to permit some kind of masking technique. To meet our low budget

[!h]



Figure 3.19: *Choosing materials for masking. The left image shows a previously masked area with masking fluid. The middle and right image show the results of experimentation with stencil & tapes.*

objective, we experimented by using conventional papers for printing and cheap watercolor papers. These papers were not suitable for our fabrication pipeline. Specifically, some paper types diffuse color more easily than watercolor paper, but they are destroyed when we apply masking or water. We use Arches hot pressed paper 300 g/m² which met all the above criteria.


For masking the paper, we experimented with ideas borrowed from watercolor painting. We used different leveling on the woodblocks that could simulate the touch or non-touch of a brush on the paper, but all the outcomes resulted in failures. Thus, we looked into how we could combine masking fluid or tapes and stencils.

All the tapes and stencils we tried seemed to visually disturb the surface of our paper (Figure 3.19). To cut the tape or the stencil to a desired shape, we used a vinyl cutter. The benefit of this technique is that the resulting masking area has a sharp and accurate shape. A negative is that the exact alignment of the stencils/tapes on the printing in progress becomes complicated. Also, we had to manually remove the negative areas of the stencils/tapes once glued on the paper. This process becomes more difficult when the complexity of the shape of the masked area and the number of the disconnected pieces of the shape increases.

For all the reasons mentioned above, we decided to use masking fluid. Compared to the tapes, the masking fluid disrupts the paper less. We confirmed that we can apply the masking fluid using a regular woodblock. Therefore, deviations from our printing process were not necessary. We brush it and press the woodblock on the paper three consecutive times. However, the masking fluid creates shapes which are less sharp-accurate, compared to the shapes created by the stencils. The masking fluid is always applied on the white paper before application of

other colors and is removed after it dries using rubber cement.

3.6.2 Carving the Blocks



(a) *Unsuccessful print, carved by milling machine*

(b) *Successful print by milling machine thickening the lines*

(c) *Successful print, carved by laser cutter*

Figure 3.20: *Woodblocks carved by different machines and the prints produced by them*

Our goal was to create multiple prints from the same woodblocks. In other words, we wanted the woodblocks to withstand the pressure from the machine press without tearing apart. Also, we wanted the carved area of the woodblocks to be as close as possible to the black-white images given as input to the machine. For that reason we first looked into a variety of machinery to carve the woodblocks. We tried two laser cutters Laser Pro Spirit GLS, C180II and a CNC (computer numerical control) milling machine, ShopBot Desktop. The woodblocks carved by the milling machine were not as detailed as the woodblocks carved by the laser cutter. Specifically in Figure 3.20, one can see three woodblocks and their corresponding prints on the right. The first two, Figures 3.20a, 3.20b, were carved by the milling machine, while the third, Figure 3.20c, by the laser cutter. Many of the small areas that were correctly carved by the laser cutter were destroyed by the milling machine, which is evident on the prints. To successfully print an image from a woodblock carved by the milling machine, we substantially thickened the lines of the block Figure 3.20b. Therefore, the milling machine is not appropriate for our small scaled woodblock printing.

We chose the Laser Pro Spirit for all the reasons mentioned above and because the smaller laser cutter C180II was slower and could not cut through our 13 mm thick woodblocks.

The time to carve one plywood woodblock of size 200 mm × 150 mm × 13 mm using the

Laser Pro Spirit GLS, was approximately 20 min, depending on the input image.

Small elevated pieces in the woodblock are easy to break with multiple prints. To address this issue we explored multiple materials.

Specifically, in Figure 3.21 there are two prints of the same image. The laser cutter carved the image onto a woodblock and on a Plexiglas block. The print resulting from the woodblock is placed on the left of the Figure, while the print from the Plexiglas block is placed on the right of the Figure. Plexiglas produced a low quality print with watercolors. We also tried MDF (Medium-density fiberboard), which also gave bad quality prints. Therefore, we rejected them.




Figure 3.21: *Prints by blocks carved on different materials: woodblock Vs Plexiglas block*

We tested pine (Figure 3.22(c)), cherry (Figure 3.22(b)), and plywood (Figure 3.22(a)) and compared their results. Cherry gave strong woodblocks with sharp details but it was expensive compared to pine and plywood. The pine wood appeared to be the softest, and to be destroyed easier after burnt by the laser cutter. Therefore we settled for plywood, which although less strong than the cherry, it was also more affordable. We could not print features thinner than 0.7 mm because the wood broke. For that reason for the details blocks we used the offset path tool of Adobe Illustrator with value 0.2 mm. We tried 6 mm thick plywood, but we noticed that the surface warped. The prints were not acceptable, because the pressure was not uniform on the entire surface of the block. We switched to 13 mm thick plywood, which in the majority of the cases remained flat. To avoid warping of the low thickness plywood, waterproofing could be effective.



Figure 3.22: *Carving different types of woods*

3.6.3 Dampening the Paper



Figure 3.23: *Color expansion when dampening paper with brushes, the color gradients are not sufficiently big.*

We want a substantial uniform water application to successfully print predictable watercolor diffusion. We explored two methods for dampening the surface of the paper: brushing and spraying. For the brush watering experiments we used two different sizes of flat watercolor brushes, one of size 19 mm (Neptune Watercolor 4750 Aquarelle 75) and one of size 38 mm (Neptune Watercolor 4750 Mottler 150). We covered the paper with multiple parallel strokes one after the other. It became difficult to have uniform dampening of the paper because of strokes overlapped. Also, due to the time difference between the first and the last brush stroke, the paper dampening was inconsistent. Moreover, the control of the dampening was dependent on the person that was doing it, and for that reason the amount of water was difficult to be measured. In addition the diffusion was not sufficiently large. To distinguish between the two different brushes we used different colors in Figure 3.23. The yellow represents the small brush

and the black the big brush. The size of the expansion was approximately 0.1 cm for the small brush and 0.2 cm for the large brush (Figure 3.23).




Figure 3.24: *Different number of spraying after the printing of the block*

For all the reasons mentioned above we switched in using a spray which gave us better results. If we sprayed from a fixed position, we could not succeed in a uniform and fully covering the paper water application. Therefore, we moved around the spray trying to succeed an approximately uniform dampening. Here we noticed that if the spray sprays too fast or if it creates big droplets, it is difficult for a person to maintain a uniform spray. We tried different sprays and the spray that gave the most satisfying results was a hand pump aerosol spray. The uniformity of the manual spraying using this spray was sufficient for the purposes of our fabrication goal. To have a more accurate and uniform liquid deposition we could use a patented technique [344,442]

Finding a balance between the desirable abstraction and the undesirable uncontrollability of the watercolors was a main challenge of our work. Small water presence removed the abstract watercolor quality of the print, while excessive water gave unpredictable results (Figure 3.25). We looked into how we could controllably expand the area of color diffusion. When spraying happened before



Figure 3.25: *Result of diffusion with excessive water application.*

the printing, it resulted in big uncontrollable diffusions. When the paper is overly wet, it adheres

to the woodblock. Subsequently, the woodblock's removal from the paper is not smooth and it disturbs the water and the paints on the paper. On the other hand, if the surface is not wet enough, then the diffusion of the colors is small. To solve this problem, we apply the water in small quantities until we ensure the paper is uniformly and sufficiently wet. We then print the woodblock and finally we spray again to drive a wide controllable diffusion. The amount of this last spray influences the width of diffusion. The sub-figures of Figure 3.24 show the color expansion after 1 spray, 4 sprays and 7 sprays. The prints made with the same woodblocks, but by applying different amounts of water are in Figure 3.26. On the left the image sprayed with less water, at the right an image with our spray setting.




Figure 3.26: *The same woodblocks for pink and yellow colors printed with different amount of spraying. Here, we adjusted the images to better match the input photo.*

3.6.4 Stretching the Paper

To have uniform color diffusion we want the paper to be well stretched and attached to a plywood of 13 mm thickness. On Figure 3.27 is a not well stretched paper and the colors diffuse in an unpredictable way. Before the plywood, we had used boards specifically designed for stretching watercolor papers, but after a couple of iterations inside the machine press, the boards were destroyed. To stretch the paper we used a stapler which we found more successful than gummed tape. Sometimes it happened that even after we had stretched the paper it did not remain stretched when we wet it. In those cases we wet on the opposite side of the print and

re-stapled the paper before we continued printing. To align the woodblocks we designed and assembled a corner using the laser cutter. Please watch accompanying video to see the printing process and tools used.



Figure 3.27: *Result in diffusion by not well stretched watercolor.*

In our fabrication process we use wet-on-wet to print the colors in both wet and dry areas. For the dry area the paper is moderately wet, as in traditional woodblock printing. For the wet area the paper is more wet and we apply additional water after the print. When we print the masking fluid we use wet-on-dry.

3.6.5 Modifying the Image Analysis Pipeline

If the user does not want to have white lines, she can segment manually the woodblocks. We tried that for the flower, the corresponding print is placed in Figure 3.28(a).



(a) Printing without white areas on the edges.



(b) Printing by finding minima in the entire image.

Figure 3.28: Different approaches than ours. It is a matter of personal taste which is preferable.

Another idea that we looked into, was to create the blocks for local minima found on the entire image and not found in every color separately, similar to how we found the local maxima. One can look the resulting print when we changed the minima pink block for the Hibiscus flower in Figure 3.28(b). The comparison with our decision of the minima block is up to the aesthetic interpretation of the designer.

Some of the ideas we experimented with created prints with “reduced” watercolor look. For example, placing the prints closer to the look of a traditional woodblock print, one such example is placed in Figure 3.29, this example has edges printed on the wet area.




Figure 3.29: Adding eastern style lines, produces a less watercolor look.

3.7 Conclusion

We have presented a novel technique of fabricating original watercolor prints given an input photograph using image processing. Our technique generates woodblocks based on the photograph’s segmentation according to color values, user-annotated semantics, smoothness values, and luminosity. We provide a pipeline to fabricate the prints using watercolor paints by

taking advantage of the color diffusion in the water. Our technique fabricates prints that match different aesthetic preferences. We hope that it will stir the interest of people from different backgrounds, who will consider using it as a tool to create their own designs driving our idea of water-woodblock to other directions, or experimenting and creating other innovative art media.

“Touch is the sense through which life is defined, because without touch is impossible to be.”

Aristotle, D’anima

4

Tactile Illustrations

Understanding the 3D geometry of everyday objects via 2D media is essential to live, learn, and work. Designers communicate the shape of a product (e.g., a chair) via sketches and renderings. Consumers understand the design by browsing images of the product in a catalog. Students learn physics, such as why a plane is shaped in a certain way for aerodynamics, from figures in a textbook. Visual perception of complex 3D geometry from a 2D projection is taken for granted by people with healthy vision. However, for people with near or total blindness, understanding 3D geometry of daily objects from current media is extremely challenging.

In this paper we introduce a novel approach to generating *tactile line drawings* that aid 3D shape understanding in users with near or total blindness, enabling the blind community to perceive complex 3D objects from 2D media. Members of the blind community rely heavily on tactile illustrations, defined as raised graphics (e.g., on paper) that are felt by hand. Common uses of tactile illustrations are to make visual information accessible in textbooks, maps [67,150], scientific diagrams [70], children’s literature [93,391,392], or museum exhibits. The tactile illustration in Figure 4.1 shows two representative samples, both used to give access to artifacts displayed in the Museum of Fine Arts Boston.

To design these types of tactile illustrations, designers rely on guidelines provided by a number of associations, such as the Braille Authority of North America [315]. Hundreds of pages of rules make manual design time consuming. For example, rules related to size, placement, variety, and form of different elements (e.g., lines, textures, points, labels, titles, and



Figure 4.1: *Tactile illustrations (a,c) displayed alongside exhibits in the Museum of Fine Arts Boston (b,d) for accessibility purposes (www.mfa.org). Black areas in illustrations indicate raised regions on the paper surface that can be perceived by touch.*

captions) permit easier exploration. Rules related to simplification, separation, elimination, and distortion of the illustration avoid clutter. Rules related to semantic segmentation, keying of the design, and selection of viewpoints permit easier understanding.

Unfortunately, studies have shown that shape understanding is flawed in traditional tactile diagrams, which can be seen from low object identification scores [212]. We aim to improve shape understanding by re-evaluating graphic design considerations and accounting for the unique needs of blind individuals. For example, while current tactile illustrations are created from a standard single viewpoint, a blind individual who has never had access to visual information might find it unnatural to explore a visual image that includes perspective distortions. This is supported by drawings created by blind individuals [205], in which illustrations of 3D objects appear to use a “fold-out” method to arrange faces onto a flat surface [224].

In this paper we first present two formative studies that motivated our approach to tactile illustration. We investigate salient shape properties in tactile perception, and assess the strengths and weaknesses of existing tactile illustration styles. Next, we propose a novel method to render tactile graphics with a focus on improving shape understanding of 3D objects. We design a pipeline introducing a multi-projection rendering approach combined with texture infills that communicate surface geometry. Compared to a single-viewpoint illustration, our multi-projection method carefully places a local camera for each semantic part of the input

object(Figure 4.2(c)), so that each part may be rendered from its optimal viewpoint. A compositing stage aligns image layers created from each camera to maintain proper connectivity of the input object. Our texturing technique then uses closely spaced lines generated from cross sections to give cues about surface geometry of the object(Figure 4.2(d)). In summary, we make the following contributions:

- We present results of a formative study involving sculpting 3D objects, aiming to discover salient information for blind users in haptic shape understanding (Section 4.2.1).
- We present results of a second formative study on the success of existing 2D tactile illustration guidelines for shape understanding (Section 4.2.2).
- We introduce a new design methodology for tactile illustrations based on the outcome of our formative studies, applying techniques in multi-projection rendering and geometry-aware textures (Section 4.3).
- We design and implement a user study evaluating the ability for blind users to understand 3D shape with our new illustration style (Section 4.4).




Figure 4.2: *Generating tactile illustrations to improve shape understanding in blind individuals: (a) Physical 3D printed object; (b) Input to our pipeline: a pre-partitioned object with colors indicating segments; (c) Local part cameras; (d) Resulting stylized illustration with curvature indicative texture; (e) Evaluation of tactile illustrations on microcapsule paper in a user study with 20 blind participants.*

4.1 Background Material

Our approach to tactile graphics, in addition to the design principles of stylized graphics illustration presented in Section 2, is also motivated by cognitive principles in tactile perception as discussed below.

There are still many unanswered questions about human perception: how perceived information is transformed into meaningful representations for the brain, how it is stored, and how it is accessed and retrieved at a later time. Cognitive science focuses on giving answers to these questions. Important findings in cognitive science have revolutionized other areas, such as computer vision, marketing, and information visualization. These findings play an important role on computer graphics research, especially related to perception, which is fundamental to the work of this project.

Human perception has been a concept underlying many works in graphics-related research [38], e.g., visual [62, 280, 329, 421] and audiovisual perception of materials [279] and in fabrication-oriented research, e.g., perception of tools in virtual drawing [334]. Perception has also been researched in relevant research regarding touch, e.g., vibrations [40], mid-air haptics [135], perception of epidermal devices [310], tactile saliency [231], and tactile roughness [416]. Finally, more relevant to our work concerning shape understanding, perceptual studies have been conducted related to, e.g., enhancement of 3D objects with texture for better shape understanding [458] or sketches of objects [120, 188].

Touch While vision has been well studied, touch has not been as extensively researched. In this project, we focus on the sense of touch, which allows humans to perceive information of objects that surround them. The skin is our largest organ, and the perception of touch arises not only through the skin, but also by the positions and movements of muscles, tendons, and joints, which makes tactile perception complicated. Also, skin resolution is much lower compared to the resolution that the eye can capture, and the pathway of signals is different from vision, making its corresponding illustrations different in terms of design and fabrication.

Tactile Perception of Shape Mental representations in shape understanding are an active area of research in psychology and neurobiology, particularly concerning similarities in how humans perceive shapes visually and through touch [126, 192, 227]. A prominent theory on shape understanding is “recognition by components” [51, 53], suggesting that our visual system processes shape information by representing an object as a spatial arrangement of 3D parts called geons (generalized cones). Our visual system extracts geons and their arrangement from an image to recognize the object. Visual translation theory hypothesizes that haptic perception is translated to the same visual representations [103, 125, 163, 233], while other work suggests vision and haptics are functionally overlapping but do not necessarily have equivalent representations of 3D shapes [126, 192, 227]. Our proposed work on illustration style abides by visual shape recognition theory, where pictorial representations should clearly show subcomponents and their arrangement in space [411].

The following subsections discuss four key elements on 3D shape perception: shape skeleton, shape curvature, camera viewpoint, and shape segmentation.

4.1.1 Skeleton

A skeleton, like the skeleton of vertebrates, synthetically defines a shape and makes intelligible its general organization while helping to visualize its possible deformations. Mathematically, a skeleton is a union of lower dimensional manifolds that is “homotopy equivalent” to the shape. As such it characterizes the topology of the shape. Skeletons, and more particularly medial representations, are computed for analyzing and identifying shapes, with a purpose of automatic recognition of objects, or of a human eye diagnosis. While such models lend themselves to identification and matching, and are therefore used for retrieval [145, 396], other researchers looked into the role skeletons play in perceptually identifying shape similarity [260]. In non-photorealistic rendering, skeletons have been used to display trees as cartoons [129] and to decide direction of lines in line drawings [110] among others.

Computer graphics also use skeletons as a starting point for shape design [184]. Skeletal animation (rigging) is an established technique to animate moving characters. Ultimately, shapes

will be represented by their boundary surfaces, as either an implicit or parametric surface, or as a fine mesh. Yet the overall design of the shape and its topology is more easily entered as a skeleton. This correspondence between shapes and skeletons arose during the projects briefly described in Appendices A.2.1 A.2.2, and helped to develop the approach used to extract lines that indicate curvature. Of the many techniques that have been developed to extract skeletons [368, 399, 401], we use the freely available implementation of CGAL [407]. Skeletons have been previously used in painterly rendering techniques to parametrize a 3D shape, e.g., similar to [110].

4.1.2 Curvature

Curvature, tactily, is an important geometric characteristic of 3D shapes [178, 466]. Because of its importance in the human perception of 3D shapes, we would like to communicate this information in our tactile illustrations. For this purpose, curvature at different parts of a 3D model is extracted in the form of cross section lines following its model's skeletal structure.

Curvature [356] has been used extensively for visualization, e.g., for vector fields [408], for maps [361], for line depiction [186, 353], for watercolor paintings [267], for shapes [107, 185], and for shape analysis, used in decomposition [117] and fabrication [402]. It has also been used for the design of textures in simplified 3D shapes [458]. For the design of curvature lines we sample the skeleton in equidistant points (with constant radius spheres) and we intersect with planes that pass through the lines and are perpendicular to the skeleton. The points are the centers of the cross sections and are used to lay the cross sections parallel to the viewing directions that result in curvature lines.

4.1.3 Viewpoints

The selection of viewpoint strongly influences object recognition accuracy [126, 170]. An experiment, testing how viewpoints affect haptics, showed that on recognition of 3D shapes, viewpoint plays a more pronounced role in haptics compared to vision [126]. Given that viewpoints affect perception, we want to have a measure of the suitability of a specific viewpoint.

This type of concern has been investigated by computer graphics researchers. A review of works about computation of viewpoints is written by Bonaventura et al. [57]. Preferred viewpoints have been investigated with different optimization criteria in mind, such as scene understanding [426], organizing and displaying shape collections [486], displaying object collisions [141], and cartoon rendering [339].

As concerns viewpoint preference, we follow an approach borrowed by Secord et al. [372] which measures the preference of a viewpoint based on human perception. We follow the approach of Secord et al. and restrict the camera up-vector to the z-axis of the object. While the computation of up-vector is a research area in the graphics community [136], in the used data set is provided [294]. In vision, natural viewpoints, i.e., most typical viewpoints [54], follow the up-right position of an object, but are not usually aligned with its coordinate frame. The importance of viewpoint preference was illustrated during one of the first projects I was involved during my PhD. In this project, viewpoints helped to improve the aesthetic quality of 3D printed objects; more details about the Viewpoint Preference Computation can be found in Appendix A.1.

Canonical Viewpoints

Illustrators often choose domain-specific canonical views. In some cases, projections in which the main axis of the object is not aligned with the viewing direction (i.e., non-canonical viewpoints) are used in advertisements and in instructional assembly materials, e.g., IKEA furniture directions. However, in other cases, views are aligned with the object axis (i.e., canonical viewpoints), such as human anatomy illustrations. Usually, more than one axis aligned views are used, e.g., floor-plan, section, elevation in architecture.

It has been shown that, for touch, canonical viewpoints [74] are easier to interpret. We use canonical viewpoints for the rendering of each object part. Different graphics techniques for extracting canonical viewpoints have been proposed, with many of them relying on principal component analysis (PCA) [82, 161, 256, 429], which is also the technique we use in the current project. We allow for different viewing directions for each part by using a multi-projection rendering approach.

Camera Projections

Non-standard perspectives were used to improve visibility, e.g., to remove occlusions [123] or to create a particular style. For example, these can be found in technical drawing, in the form of linear perspective with 0 vanishing points (parallel projection), or in Byzantine drawings in the form of divergent (or inverse) perspective [179]. Similarly, curvilinear perspective [277] (e.g., fisheye views) can be used in photographic cameras and maps. In addition to the use of a single non-standard perspective with multiple number of vanishing points, multiple viewpoints have been used to create a single image with a large angle of view.

Some works in multi-perspective rendering allow for broad visibility and create images with space distortions [290, 340, 470], which we want to avoid. Therefore, we build upon a work from computer graphics research [23] which can create a distortion-free image, and we extend this technique to be a better fit for the purposes of our application. The positions of the cameras are defined from PCA, while the up-vectors are initialized using a vector field [13] and later aligned according to connectivity between parts. We also use orthographic cameras, similar to technical drawings, to avoid perspective distortions.

4.1.4 Shape Segmentation

In graphics several segmentation techniques exist [146, 165, 198, 199, 382, 437, 468, 484]. For the purpose of this work we use a readily segmented data set [294]. Shape segmentation is a hard open problem actively researched in graphics. When we were working in the second project of this thesis and tried publicly available implementations for shape segmentation, the results were not satisfactory, and therefore we decided to use pre-segmented data set, which allowed experimentation, without the need of segmentation.

4.2 Formative Studies

In this section we present two formative studies that guided our new design approach for tactile illustrations. The first investigates what information is considered salient in tactile shape

understanding. The second study assesses how successfully existing illustration guidelines convey 3D shape information.

4.2.1 User Study 1: 3D Shape Replication

We performed an object replication task to develop a deeper understanding of shape characteristics blind people consider salient. The experiment involved a sculpting task where the aim was to identify shape characteristics that participants focus on.



Figure 4.3: *Sculptures from User Study 1. (Top row) Reference objects. (Middle row) Participant 1 sculptures. (Bottom row) Participant 2 sculptures.*

Apparatus and Participants.

We selected a variety of man-made and natural objects with different levels of shape complexity. Two congenitally blind participants (1 female) participated in the study, where they were given an hour to sculpt five shapes, including a pyramid, dolphin, cup, chair, and human in a neutral standing pose (Figure 4.3, top row) using 4 oz of polymer clay for each shape.

Task and Procedure.

Participants were instructed to sculpt replicas of the shapes. They were also asked to identify each object and describe its shape and any difficulties encountered while sculpting.

Observations.

The results of the study (Figure 4.3) suggest a number of important findings. First, the participants segmented the shapes into distinct geometric parts (**F1**). For example, they first created the four chair legs and brace segments, then proceeded to connect the pieces together. Participants were also able to correctly follow the topology of the reference shapes (**F2**). An example of this is sculpting the correct number of pieces to construct the legs in the chair model. They were also able to sense simple shapes of the objects or the cross section of a part (e.g., square base for the pyramid, rounded cross section for the body of the dolphin, or cylindrical cross section for the cup). Participants were also able to notice differences when cross sections change within the same part (**F3**). For example, they noticed the change in cross section from the front to the back of the dolphin body. Participants also pay attention to the curvature properties of the surface geometry (**F4**). For example, they flattened the facets of the pyramid and created distinct edges, or sculpted a rounded shape for the dolphin's body.

Participants encountered a number of challenges during the study. For example, participants were unable to accurately replicate positions of parts (**F5**). Sometimes they even misidentified relative positions (e.g., failing to make the dolphin's pectoral flippers symmetric). The participants also experienced difficulties in accurately capturing proportions (**F6**), e.g., the height-to-width ratio of the pyramid was less important than replicating a sharp point at the apex.

Our findings suggest that to communicate 3D shape information effectively to blind users, the following important characteristics should be conveyed through the illustration:

- Distinct parts within an object (F1)
- Topology of the shape (i.e., preserve part connectivity) (F2)

- Cross section shape and propagation (F3)
- Curvature of surface geometry (F4)

Our findings also suggest that the following characteristics have low importance in shape understanding:

- Accurate placement of the connection points (F5)
- Accurate proportions (F6)

Note that F5-F6 were independently observed, not inferred from F1-F4. The preliminary findings in our object replication study align well with existing theory on shape understanding, in particular the “recognition by components” theory [51,53] of how any view of an object can be represented as a spatial arrangement of 3D parts (called geons, or generalized cones).

4.2.2 User Study 2: Assessment of Existing Illustration Guidelines

Apparatus and Participants.

Seven congenitally blind participants (3 female, aged between 30 and 72) participated in our study. We chose five reference objects familiar to our participants: teapot, lamp, eyeglasses, chair, and rocking horse. We created tactile illustrations for each object in three different styles (Figure 4.4):

1. Line drawing with object in non-canonical view (also known as perspective or 3/4 view).
2. Line drawing with object in canonical view (e.g., front or side).
3. Line drawing with canonical view and texture-infills to indicate different semantic regions in the illustration.

In case (2) we used two viewpoints (i.e., 2 corresponding illustrations) if important information was contained in both, following recommendations by the Braille Authority of North America (BANA) guidelines [315]. In case (3) the separation into semantic regions and selection of textures were following BANA guidelines. All illustrations were made by a trained graphic

designer. The physical tactile diagrams were created by laser engraving commercially available capsule paper.




Figure 4.4: 3D printed reference objects for User Study 2 (row 1). Corresponding illustrations: non-canonical views (row 2), canonical views (row 3), and textured canonical views (row 4).

Task and Procedure.

During the study, participants were asked to use the three tactile illustrations to recognize the tested objects, identify parts of the objects (e.g., locate the spout of the teapot), explain the perceived shape of the parts, explain the connection between the parts (e.g., where the spout connects to the main vessel), and count the number of the parts. We showed the objects in the same sequence to all users: (1) non-canonical, (2) canonical, (3) textured canonical.

Observations

Several important findings were observed from our study regarding the key characteristics related to tactile line-drawings in communicating 3D shape information, including the number

of viewpoints, viewpoint selection, occlusion, and shape understanding.

Multiple viewpoints According to the Braille Authority of North America (BANA) guidelines for tactile graphics [315], multiple views should be provided if they contain information important for the intended task (guidelines: 10.1.4; 7.1.1.5; 3.6.2). However, we found people could not successfully identify correspondences between the different illustrations (**F7**). Instead, participants often chose to explore their preferred viewpoint to perform a task, even if the essential information was contained in the other view. It was considered cumbersome to explore more than one diagram at a time.

Viewpoint Selection Selection of viewpoints has been shown to influence object recognition accuracy [170]. We also found that participants' ability to identify an object from an illustration had higher success rates for canonical viewpoints (21%) vs. non-canonical viewpoints (6%) (**F8**). However, participants were generally not satisfied with any of the canonical viewpoints provided. For example, in the eyeglasses illustrations participants commented that the front view depicted the correct size for the lenses but the wrong size for the frame arms, and the opposite was true for the side view. A participant suggested combining viewpoints: showing the lenses from the front and the frame arms from the side. That is, create a combined multi-projection image in which each part will be shown in its optimal view (**F9**).

Occlusions One of the most common complaints from participants was the absence of parts from the diagrams, which was a consequence of occlusion (**F10**). For example, in the chair illustrations, front and side views were provided (Figure 4.4) and a participant indicated that the chair seemed to have 2 legs instead of 4. This is consistent with an existing study by Moringen et al. [300] that showed occlusions affect tactile shape recognition.

Shape understanding Illustrations rendered in canonical views were generally better for understanding surface geometry compared to non-canonical views (**F8**). When asked about parts being more rectangular versus cylindrical, users were 71% correct for canonical illustrations versus 48% for non-canonical views. However, in some cases the non-canonical views improved

shape understanding. For example, the non-canonical view of the teapot was more successful for recognizing the spherical shape of the main container, whereas in the canonical view it was perceived as cylindrical or cubical. Possible reasons are the curved edge along the bottom of the teapot and ellipsoidal shape of the lid (Figure 4.4). This is in agreement with a study by Thompson et al. [411] which showed that curved lines as textures in tactile illustrations give indication about the geometry of the represented shape. Further, a study by Yau et al. [466] showed that line orientations and the level of curvature and curvature direction of lines changes are identifiable by touch (F11).

Foundations for the New Design

Based on User Study 2, we propose a number of additional design guidelines to improve 3D shape understanding using tactile illustrations:

- Contain all information in a single illustration (F7)
- Depict canonical viewpoints (F8)
- Combine multiple viewpoints within a single illustration so that each part of the object can be depicted optimally (F9)
- Choose a viewpoint of the object to reduce occlusions (F10)
- Use curved lines to communicate surface geometry (F11)

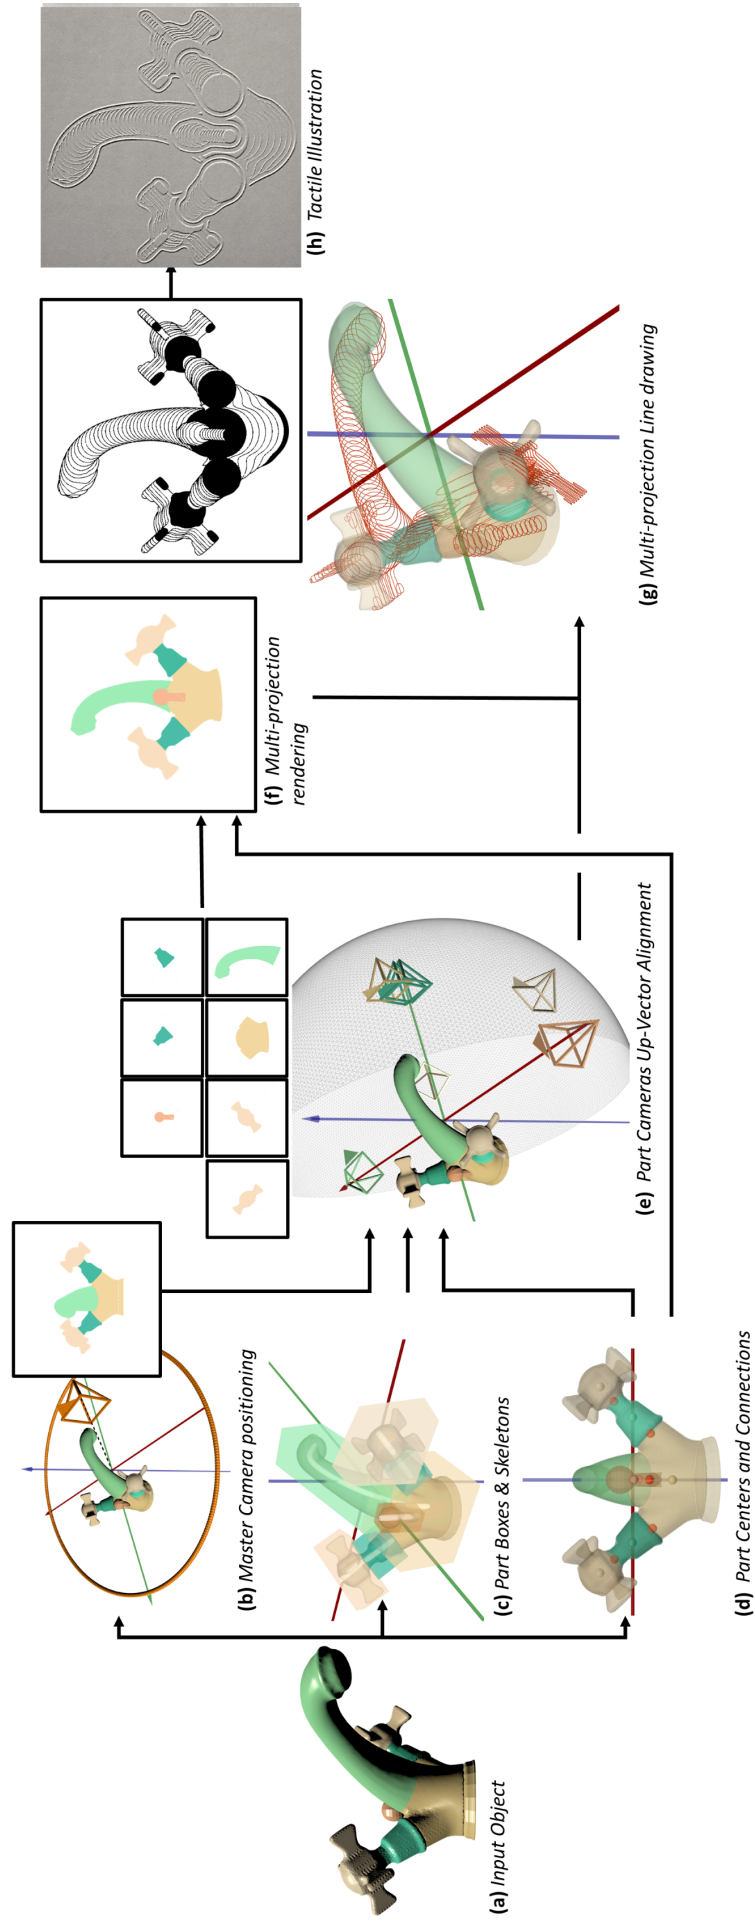


Figure 4.5: Overview of our pipeline for the design of tactile illustrations. (a) Input model, colors represent different parts of the faucet. (b) Master camera placement (inset box shows camera view). (e) Local cameras are assigned to each part (inset boxes show camera views). (f) The image layers are composited; positioning of each layer is guided by the master camera and the part connectivities (d). (g) Infill textures are generated based on cross sections of each part (black represents raised areas). (h) The fabricated illustration created with microcapsule paper.

4.3 Tactile Illustration Method

4.3.1 Overview

Given the design foundations derived from our formative studies, we present a new approach to creating tactile illustrations of 3D objects. Our method combines images from different viewpoints, applying techniques in multi-projection rendering. We also introduce a new approach to generating tactile infill textures that communicate the surface geometry of the object.

Our illustration approach draws inspiration from the artistic multi-projection rendering stylization introduced by Agrawala et al. [23]. Similar to their work, we create a set of local cameras which render images from different viewpoints. We then employ a master camera to combine the local renderings into a single layered image. A key difference to Agrawala et al. [23] is that our local cameras are assigned to segmented parts within the same object. This means our method must constrain relative positions of part-level renderings in order to maintain correct connectivity in the final illustration. Further, our approach automates camera positioning by defining a notion of optimal viewpoint for each part.

The input to our pipeline is a partitioned 3D object. Users can either manually partition the object or use an automated technique (e.g., [484]). We use the PartNet data set [295] consisting of pre-partitioned 3D objects in a wide range of common object categories, and provided in up-right orientation. Our pipeline outputs a stylized multi-projection rendering, which can then be fabricated to produce a tactile illustration. The main steps of our algorithm are as follows:

1. Determine local camera poses. A *part camera* is assigned to generate a rendering for each decomposed part of the object, with position and orientation chosen according to visibility objectives (Section 4.3.3). A *master camera* is employed to combine the multiple renderings into a unified illustration (Section 4.3.2).
2. Generate texture infills for each part rendering. We integrate surface geometry information using cross sections extracted along the object skeleton (Section 4.3.4).
3. Compositing the final illustration from local part renderings. The combination of image



Figure 4.6: *Different metrics considered for finding the master camera placement. The viewpoints maximize: (A) Silhouette Curvature Extrema, (B) Silhouette Curvature, (C) Projected Area, (D) Silhouette Length, (E) Surface Visibility, (F) Maximum Depth, (G) Depth Distribution, (H) Viewpoint Entropy. The maximum silhouette length (D) decreases part occlusion.*

layers accounts for relative up-vector alignment (Section 4.3.5), part connectivity within the 3D object, and occlusion handling (Section 4.3.6).

4.3.2 Master Camera Placement

In our multi-projection approach a master camera is used to organize the rendered images from each local part camera. As discussed in Section 11 (design foundation (F10)), viewpoints that result in non-coincidence of parts on the image (i.e., not more than one part appearing on each pixel) are preferred. The view from the master camera helps to define the local part camera frames to assess part visibility in the final composition.

To determine the master view direction $\hat{\mathbf{d}}_M$ we follow the approach of Secord et al. [373], where several attributes are defined for identifying preferred visual viewpoints. We notice that viewpoints that maximize the silhouette length result in decreased part coincidence. Silhouette length is defined as the overall length of the object’s silhouette in the camera’s image plane. For example, in Figure 4.6(D), overlap between the handle and main container is small when viewed from a direction that maximizes silhouette length.

Next, we align the up-vector of the camera $\hat{\mathbf{u}}_M$ with the up-vector of the object. This results in a top to bottom layout in the final illustration (i.e., the base of the object is positioned at the bottom of the illustration), which is a natural choice and follows the Braille Authority of North America (BANA) guidelines and standards for tactile graphics [315] (6.11). Note the object up-vector is not considered in Secord et al. [373] for viewpoint selection. We fix the up-vector

and constrain the camera location to a ring centered in the object (Figure 4.5(b)). More details about the master camera computation are in Appendix B.2–B.3.

4.3.3 Part Camera Placement

Next, we determine a local frame for each part camera. As discussed in Sections 4.2.2 and 11 (design foundation (F8)), canonical views tend to improve shape understanding for blind users. To identify the canonical viewpoint for each part we first extract a coordinate frame $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3\}$ using PCA, similar to Vranic et al. [430], more details in Appendix B.4.

We next extract an one-dimensional skeleton for each part using the mean curvature skeleton by Tagliasacchi et al. [399] (white lines in Figure 4.5(c)). We reject the PCA axis that most closely aligns with the skeleton, which will be important later when creating infill textures.

Between the remaining axes we choose the direction corresponding to the largest projected area. For simplicity this is done heuristically by aligning a bounding box (Appendix B.5) to the frame (Figure 4.5(c)), and selecting the axis corresponding to the largest face on the bounding box (i.e., face with normal $\hat{\mathbf{e}}_i$). Between the $\pm\hat{\mathbf{e}}_i$ directions, we choose the one that gives a viewing direction in the same half-space as the master camera, where local camera viewing direction is $\hat{\mathbf{d}}_P = \mp\hat{\mathbf{e}}_i$. More details are in Appendix B.6. The part camera location is given by $\mathbf{p}_P = \mathbf{c}_P - r\hat{\mathbf{d}}_P$, where $\mathbf{c}_P \in \mathbb{R}^3$ is the centroid of part P and r is the distance to the camera. The part P is centered in the local camera viewport.

4.3.4 Cross Section Textures

The next step is to create the infill textures. Compared to the BANA standard [315] which uses uniform textures to denote semantic regions, we choose to incorporate properties of surface geometry motivated by design foundations (F3, F4, F11). Many methods have been explored for creating interior lines and textures in non-photorealistic rendering, such as hatching and cross hatching [200] or apparent ridges [197]. However, these approaches are based on illumination effects of shading or view-dependent properties which are not relevant to blind users for tactile perception of 3D forms. Similar to Deussen et al. [110] we choose to create texture infills using

cross section boundaries to illustrate the geometry of each part.

As noted in Section 4.3.3, for each part we extract a 1D skeleton using the method of Tagliasacchi et al. [399]. We initialize the cross sections to be uniformly sampled along the skeleton, with normals, $\hat{\mathbf{n}}_{cs}$, aligned with the local skeleton direction. Next, to allow for visibility of cross sections on the image and remove perspective distortion, each cross section is flattened by rotating to lie parallel to its corresponding part camera view plane (Figure 4.5(g) and Figure 4.7). The skeleton sample point remains fixed with the rotation axis given by $\hat{\mathbf{d}}_P \times \hat{\mathbf{n}}_{cs}$, where $\hat{\mathbf{d}}_P$ is the view plane normal for local part camera P . We then render a line for each cross section silhouette from the corresponding part camera. The layering of cross section silhouettes in the image follows the sampling ordering along the skeleton (pseudocode in Algorithm 1).

The spacing of cross section lines considers limits on resolution for tactile perception. The cutaneous system has poor spatial resolution compared to vision [53, 233, 443], such that high-frequency details are lost when haptically perceiving raised lines. We follow recommendations from the BANA guidelines [315] of 2 mm to 3 mm gaps between elements to be distinguishable by hand.

ALGORITHM 1: FLATTENEDCROSSSECTIONS ($Sk, Shape_P, \hat{\mathbf{d}}_P$)

Input: $Sk = [\mathbf{sk}_1 \dots \mathbf{sk}_n]$ a sequence of regularly sampled points on the skeleton. The geometry of the part, $Shape_P$. View plane normal $\hat{\mathbf{d}}_P$ for camera of part P .
Output: Cross sections $CS = [cs_1 \dots cs_n]$ parallel to view plane and layered. Each cs defined by sample point from skeleton \mathbf{b}_{cs} , normal $\hat{\mathbf{n}}_{cs}$, and $Boundary_{cs}$.

```

1 for  $sk_i \in Sk$  do                                     // Extract cross sections along skeleton
2    $\mathbf{b}_{cs_i} \leftarrow \mathbf{sk}_i$ 
3    $\hat{\mathbf{n}}_{cs_i} \leftarrow$  local tangent direction at  $\mathbf{sk}_i$ 
4    $Boundary_{cs_i} \leftarrow Shape_P \cap CUTTINGPLANE(\mathbf{b}_{cs_i}, \hat{\mathbf{n}}_{cs_i})$ 
5    $cs_i \leftarrow CROSSSECTION(\mathbf{b}_{cs_i}, \hat{\mathbf{n}}_{cs_i}, Boundary_{cs_i})$ 
6 end
7 for  $cs_i \in CS$  do                                     // Align cross sections with view plane
8   if  $\hat{\mathbf{d}}_P \neq \hat{\mathbf{n}}_{cs_i}$  then
9      $RotAxis \leftarrow \hat{\mathbf{d}}_P \times \hat{\mathbf{n}}_{cs_i}$ 
10     $PivotPt \leftarrow \mathbf{b}_{cs_i}$ 
11     $RotAngle \leftarrow ANGLE(\hat{\mathbf{d}}_P, \hat{\mathbf{n}}_{cs_i})$ 
12     $cs_i \leftarrow ROTATECROSSSEC(cs_i, RotAxis, PivotPt, RotAngle)$ 
13  end
14   $\mathbf{b}_{cs_i} \leftarrow \mathbf{b}_{cs_i} - SmallStep * i * \hat{\mathbf{d}}_P$ 
15 end
16 return  $CS$ 

```

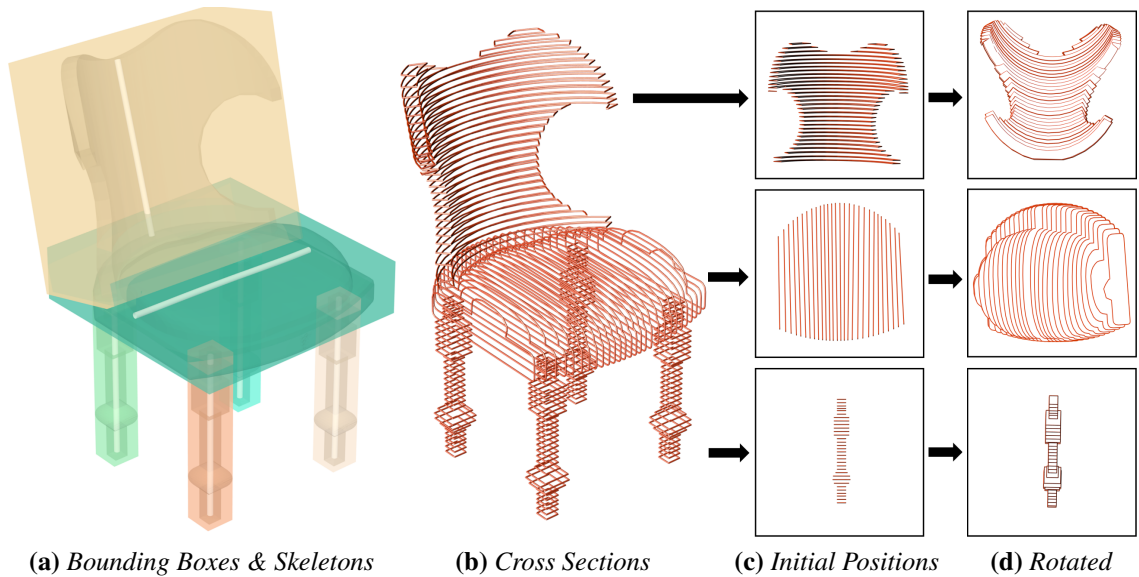


Figure 4.7: The chair displaying the bounding box and skeleton for each part (a). Cross sections sampled uniformly along the skeleton. The cross sections of each part as seen from their corresponding local cameras (c), and after the flattening operation (d).

4.3.5 Part Camera Up-Vector Alignment

We next align the orientation of the local part cameras. In Section 4.3.3 we determined the view direction for each part camera, $\hat{\mathbf{d}}_P$, prioritizing the direction of max projected area for the part. However, the up-vector orientation of the camera, $\hat{\mathbf{u}}_P$, was left ambiguous. Our approach takes into account the master camera viewpoint and connectivity of parts within the 3D object. As an initialization we use the vector field provided by Lopes et al. [13] – the vertical axis is the master camera up-vector, $\hat{\mathbf{u}}_M$, the normal vector corresponds to the local camera view direction, $\hat{\mathbf{d}}_P$, and the tangent given by the vector field provides an initial guess for $\hat{\mathbf{u}}_P$ (Appendix B.7).

The next step is to identify connection points between adjacent parts. For each part P_j , we collect the set of adjacent parts \mathcal{A}_j . We assign a representative point $\mathbf{a}_{ij} \in \mathbb{R}^3$ to each connecting part P_i , for $i \in \mathcal{A}_j$, computed as the centroid of the contact surface between P_j and P_i . The part centroids \mathbf{c}_j and connection points \mathbf{a}_{ij} are shown for the faucet in Figure 4.5(d).

We make a slight adjustment for the reference “center” point of the part, $\bar{\mathbf{c}}_j$. If there is only one connection point, the center of the part $\bar{\mathbf{c}}_j$ coincides with the mesh centroid. Otherwise the average of all connection points is used. We then compute a rotation for each part camera’s up-vector, $\hat{\mathbf{u}}_j$, such that the relative *horizontal* positioning between $\bar{\mathbf{c}}_j$ and \mathbf{a}_{ij} is the same when




Figure 4.8: Alignment of the part camera up-vector. (a) Input object as seen from the master camera. (b) Center and connections for the tabletop from the part camera view (top) and master camera view (bottom). Angle computed so that horizontal distances of connections, w , match the master camera view. (c) Final upright orientation of the local camera, $\hat{\mathbf{u}}'_j$.

viewed in the master and part camera image planes (Figure 4.5(f), 4.8). To achieve this we use the following steps:

Recall that we are using orthographic cameras, where the object size in the image plane does not diminish with depth. We define the vector $\mathbf{v} = \mathbf{a}_{ij} - \bar{\mathbf{c}}_j$ in world space. The horizontal distance in the master camera image is given by $w = \mathbf{v} \cdot \hat{\mathbf{x}}_M$, where $\hat{\mathbf{x}}_M$ is the x -axis of the master camera viewport. We then determine the rotation angle θ_{ij} of the camera frame around the view direction $\hat{\mathbf{d}}_j$. The angle is chosen such that $\mathbf{v} \cdot \hat{\mathbf{x}}'_j = w$, where $\hat{\mathbf{x}}'_j$ is the x -axis of the viewport for part camera j after rotation (Figure 4.8).

We repeat the same process for all connection points of the part and use the average rotation, $\bar{\theta}_j = \text{avg}(\theta_{ij})_{i \in \mathcal{A}_j}$ as the final alignment of the up vector, $\hat{\mathbf{u}}'_j = \mathbf{R}(-\bar{\theta}_j)\hat{\mathbf{u}}_j$. The rotation axis is the viewing direction $\hat{\mathbf{d}}_j$.

4.3.6 Compositing Part Renderings

The final stage is to combine the images from each local part camera to create the composited illustration. We compute a depth value for each part, calculated as the distance from the master camera location, \mathbf{p}_M , to the furthest vertex of the part's bounding box (in the master camera

coordinate frame). The order controls the layering with which images appear on the multi-projection illustration. The default is to start from the farthest and move to the closest part. The user may choose to change the ordering manually since this does not handle many occlusion cases, e.g., where parts close to the camera completely occlude other parts.




Figure 4.9: *Aligning part camera image layers. (a) The input model as seen from the master camera. (b) If the centroids in the master camera view are used for aligning part renderings, the resulting image is disconnected. (c) Instead, the connection points are used for aligning image layers.*

Initially all image layers are stacked such that their local frames align (i.e., image centers coincide). We traverse the graph of all parts starting with the farthest part, where nodes are the parts P_j , and edges are their part connectivity \mathcal{A}_j . When we traverse the edge ji , i.e., a connection from part P_j to its adjacent part P_i , we update the position of the image layer from part camera i . The image layer is translated to enforce the part renderings to align at their connection point (Figure 4.9). The translation vector is given by $\mathbf{t}_j \in \mathbb{R}^2$:

$$\mathbf{t}_j = \text{proj}_j(\mathbf{a}_{ij}) - \text{proj}_i(\mathbf{a}_{ij})$$

where $\text{proj}_i(\mathbf{a}_{ij})$ is the projection of connection point \mathbf{a}_{ij} in the image plane of part camera i . After applying the transformations, the connection points from different camera renderings align on the image plane. Note, we assume no cycles exist, this procedure does not handle cases where the connectivity of parts forms a multiply-connected graph rather than a tree

As a final post-processing step, we create an outline around each of the part textures of 1 mm thickness, and enforce a 3 mm wide gap between each part. This helps to differentiate




Figure 4.10: *Photograph of tactile illustration of the monitor, including a closeup view of the base. Fabricated with microcapsule paper.*

between parts, following BANA guidelines [315]. Renderings for the tactile illustrations of a selection of objects can be seen in Figure 4.5(g) and Figure 4.14. Finally, we send the renderings to the laser cutter to fabricate the tactile illustrations, examples can be seen in Figure 4.5(h) and Figure 4.10.

4.3.7 Designs for User Study 3

This section contains all materials for User Study 3: Evaluation. Materials include:

- Seven (7) 3D reference objects. Each object has three geometric variations for use in the study. Modified parts are highlighted in green. All three versions were physically fabricated with 3D printing.
- Tactile illustrations generated with our technique. Fabricated with microcapsule paper.
- Tactile illustrations designed following BANA guidelines [315] as a baseline. Fabricated with microcapsule paper.
- An eighth example (jug) is provided which was used for training.

During the study the objects were presented in the following order:

1. Monitor 1
2. Table 1
3. Headphones
4. Chair
5. Monitor 2
6. Table 2
7. Faucet




Figure 4.11: Objects 1–4 in User Study 3: Evaluation. For each object: (top) 3D reference model with variations; (bottom left) tactile illustration using BANA [315] guidelines as baseline; (bottom right) tactile illustration using our technique. In each case the illustrations depict the middle reference object.

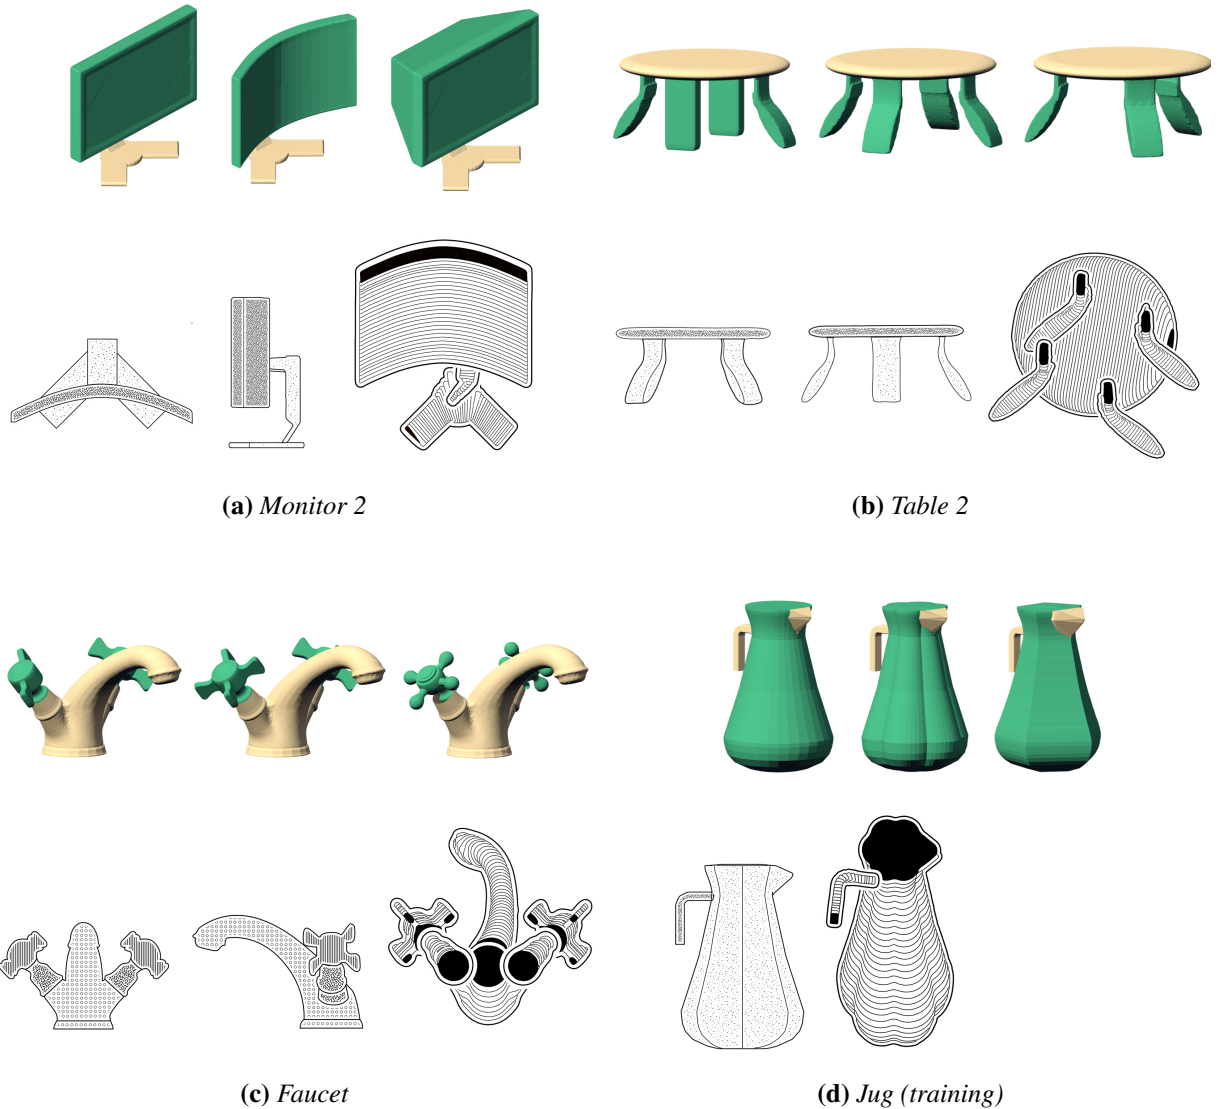


Figure 4.12: Objects 5–8 in User Study 3: Evaluation. For each object: (top) 3D reference model with variations; (bottom left) tactile illustration using BANA [315] guidelines as baseline; (bottom right) tactile illustration using our technique. In each case the illustrations depict the middle reference object. The jug was used as a training example at the beginning of the study.

4.4 User Study 3: Evaluation

We conducted a user study to evaluate the effectiveness of our novel tactile diagrams. The study included an object recognition task, in which participants with near or total blindness were invited to a controlled lab environment to identify a set of daily objects through the proposed tactile illustrations.




Figure 4.13: For each 3D object included in User Study 3: Evaluation, we created three versions with minor geometric variations. Beige regions are identical among the three versions; green regions have been modified. The models were 3D printed so that participants could compare (by touch) the 3D objects depicted in the 2D tactile illustrations. The tactile illustrations shown in Figure 4.14 correspond to the middle version in each case.

Apparatus and Participants.

Our study consisted of 20 blind participants: 12 female, aged between 26 and 72, and 14 congenitally blind. Seven objects were used in the study, we chose common household objects that would be familiar to the participants (in the order presented): monitor, square table, headphones, chair, curved monitor, round table, and faucet (Figure 4.14). For each object, the study materials included:

1. Three variations of the 3D reference object (Figure 4.13). All three versions were physically fabricated using 3D printing.
2. One tactile illustration generated with our technique (Figure 4.14). Fabricated with microcapsule paper (Figure 4.10).
3. One tactile illustration designed following BANA guidelines as a baseline. Fabricated with microcapsule paper.

Including variations of the 3D objects allowed us to measure how well our illustration technique helps participants precisely perceive the 3D geometry information and recognize subtle differences between the reference objects and their variations. To create variations of the 3D




Figure 4.14: *Resulting designs for tactile illustrations using our technique, with 3D input models shown.*

reference objects, each model was modified by either a replacement operation, subtraction, or addition of parts. The variations of the objects are as follows (Figure 4.13):

- Monitor 1: changes in base design
- Table 1: leg design (rectangular vs. cylindrical cross section, and cylindrical with connectors)
- Headphones: size of ear cushion and outer shape of earpieces
- Chair: curved vs. flat back, and presence of cushion
- Monitor 2: curved vs. flat screen and CRT style
- Table 2: number of legs and straight vs. tilted shape
- Faucet: number and shape of lever arms on handles
- Jug (training example): cross section shape of main container

Example materials for the chair are shown in Figure 4.15. Materials for all objects in the study are included in the Appendix 4.3.7.

Task and Procedure.

The tactile illustrations were presented to our participants, whose task was to find the corresponding 3D reference object or its variation. A baseline was included to allow us to measure the

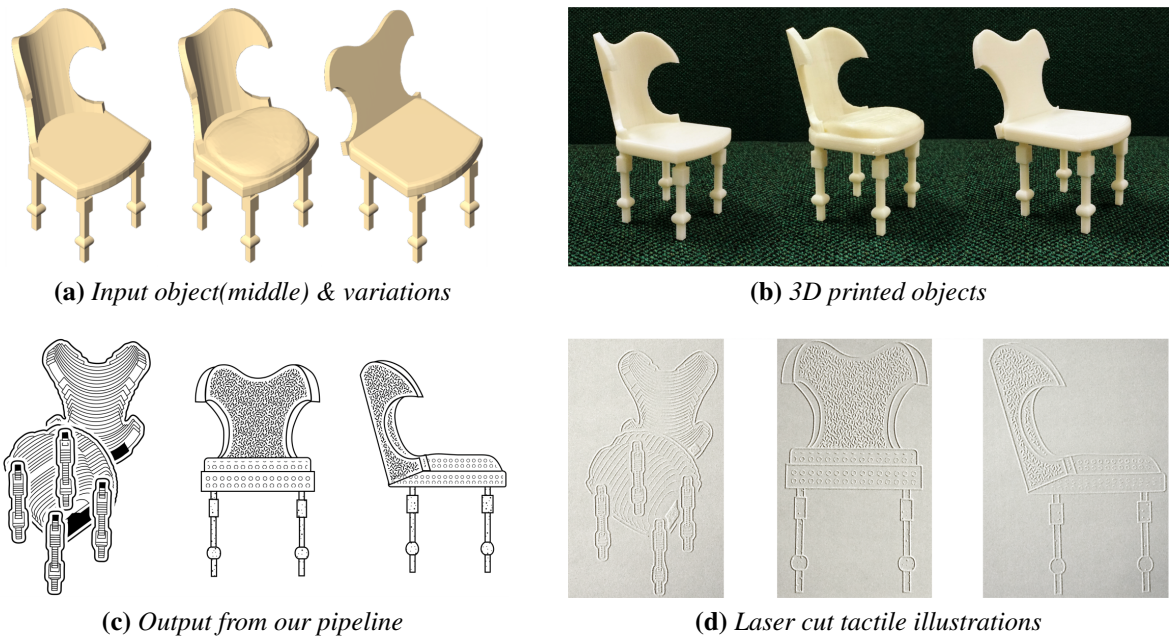


Figure 4.15: *Example materials for User Study 3: Evaluation. Participants are given three variations of the object (a), produced by 3D printing (b). They are given one tactile illustration style – either our new multi-projection style (c, left) or the BANA baseline style (c, middle-right). Illustrations are produced using microcapsule paper (d). Participants are asked to identify which of the three variations matches the illustration. (Supplemental Materials for all materials used in study 3).*

performance of our technique over the state-of-the-art. The baseline was produced by following the Braille Authority of North America (BANA) guidelines [315]. Sample materials for the chair object can be found in Figure 4.15. Participants were trained in how to interpret the tactile illustrations using the jug example. The procedure was as follows:

1. Participants were given materials for one object at a time consisting of three variations of the object (3D printed), and one illustration of “type A.”
2. The participant was asked to choose which one of the three object variations most closely matched the illustration.
3. Steps (1) and (2) were repeated for all seven objects.
4. Steps (1)–(3) were then repeated using illustrations of “type B.”

For half of the participants “type A” was the baseline (BANA guidelines) and “type B” was our illustrations. For the other half of participants the illustration type was reversed. We chose to




Figure 4.16: *User Study 3 results, shows mean success rate for matching tactile illustration to correct 3D object. Results shown for average over all objects, and per object success rates.*

use only three variations of the reference object to avoid cognitive overload, this was determined after running a small pilot study.

Our quantitative measures included matching rate, which measures the number of times participants were able to correctly match a tactile illustration with the corresponding 3D object. Qualitative observations were also collected to assess how easy it was for participants to understand the 3D shape information shown in the tactile illustrations and identify the subtle differences between them.

Observations.

We analyzed the data using a generalized linear mixed model for a comparison between the baseline BANA illustrations and our technique, considering each response an individual sample. The network includes two fixed effects: design and object, and a random effect: participant, and is built on a family of binomial distributions. The variable object has seven levels (monitor 1, table 1, headphones, chair, monitor 2, table 2, and faucet) and the design has two (BANA, our technique). To compare success rates within each object we do Post-Hoc pairwise comparisons on estimated marginal means and apply Bonferroni corrections [234].

The mean success rate in matching the correct 3D reference object to the tactile illustration

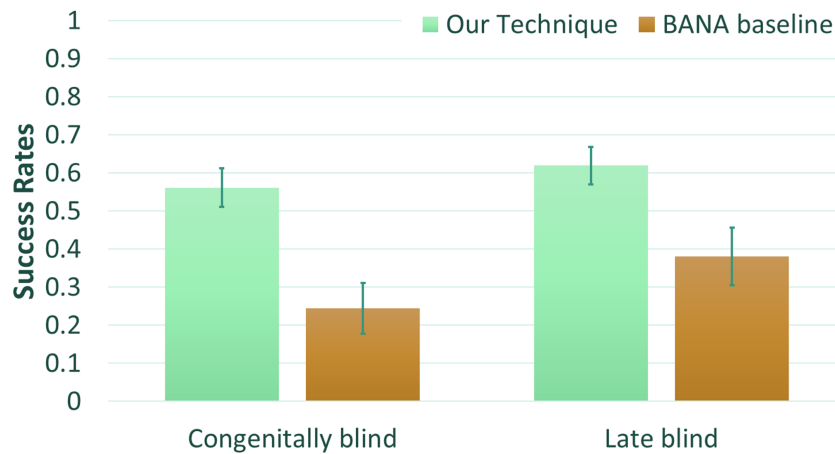


Figure 4.17: User Study 3 results, shows mean success rate for matching tactile illustration to correct 3D object. Grouped by age of vision loss (self-identified).

was 58% for our technique versus 29% for the BANA baseline. We found that there exists a significant difference in shape understanding among the two types of illustrations ($\chi^2(1, N = 140) = 22.77, p < 0.005$). Figure 4.16 shows the mean success rate over all objects, as well as per object success rates. Confidence intervals are computed by standard error. There is also a significant effect of Object ($\chi^2(6, N = 140) = 31.55, p < 0.005$). Post-hoc analysis shows that higher success rates were observed for our technique for the majority of objects, with monitor 1, chair, and faucet remarkably better, with an improvement of $\geq 45\%$ in mean success rate (all $p < 0.005$). Table 1, monitor 2, and table 2 also have large improvement (e.g., improvement of $\geq 15\%$ in mean success rate), but without statistical significance. An exception to this trend is the headphones where participants responded more accurately using the BANA baseline ($p < 0.05$). Figure 4.17 compares results between congenitally (14) and late (6) blind participants. Both groups show better success rate with our technique.

Although our results are promising, there were still misconceptions caused by our technique. One observation was that overlapping parts were sometimes overlooked. For example, in our illustration for table 1 (square top, Figure 4.14) one of the legs is entirely contained by the tabletop. Several participants were unable to locate this leg, likely since the silhouette was not easily discoverable.

The headphones were the one object with poorer mean success rate than the baseline. We

found that participants were not using the infill textures to understand the shape of the earpieces, but they were judging shape just by the outline which may have indicated a spherical instead of cylindrical shape. Another consideration was the lack of symmetry in our illustration – although identical, the two earpieces were shown from different views which produced different silhouettes and seemed to confuse participants.

To understand user perception of our new illustration style, we conducted a usability survey based on the System Usability Scale (SUS) [68]. With a maximum score of 100 we found similar scores on average of 65 for our technique and 62 for the BANA baseline. This result also indicates that participants tended to be open to new ideas and solutions for the improvement of tactile illustrations.

At the end of the study we collected additional unstructured feedback from participants, summarized here:

Infill textures Participants had positive feedback on the infill textures for both our technique and the BANA baseline approach. Participants found the BANA textures useful for differentiating between different parts (2 users). In contrast, our textures were useful for understanding shape details (2 users).

Participants generally liked the textures in our illustration technique, with two users saying it is a consistent way to represent shapes, and four users saying the shape could be felt clearer/easier. Other positive remarks were that the designs were “more detailed,” “interesting,” and that the textures were “more tactile.” Specific use cases were suggested, including art objects in a museum (e.g., statues), and for depicting architecture (e.g., Cape vs. Victorian roof styles or different pillar styles). Three users raised concerns, mostly related to the clutter that the lines created, which made the illustration feel “complicated” and “everywhere the same,” so that they could not easily identify “where one piece starts and where it stops.”

Multiple viewpoints Participants generally appreciated that our technique created a single unified illustration, “there was only one shape to look at.” Two participants commented that our technique represents the shape more effectively, is more simplified, and conveys the object more

efficiently. One participant commented: “If I had seen the diagrams [our technique] and not the objects, but I was told that the diagram was a chair, then I could have interpreted it. But once we moved to different angled views [BANA baseline] that became less and less something I would be confident about.” However, three users found difficulties in interpreting our illustrations, saying they looked more like an “artist drawing” and “more 3-dimensional,” and not knowing the view was problematic.

Regarding the BANA baseline illustrations, four participants commented that it was difficult to explore multiple views (e.g., top and side) which they later needed to match and compare. E.g.: “in some of the views you did not see all the object, you did not understand what you were looking at, you had to use both of them and kind of put them together or use what information you could get from the one that made the most sense.” This response highlights the benefits of a single illustration.

Fabrication issues One participant noted that he would like not only thickness differences in lines, but also height differences, while another commented that he would prefer sharper lines.

4.5 Fabrication and Implementation Details

Implementation Details The skeleton extraction uses the CGAL implementation of the mean curvature skeleton [399]. To create the renderings for User Study 2, we extracted lines using Freestyle [156], making modifications manually when needed in Adobe Illustrator. We created lines of 1 mm thickness with 3 mm wide gaps between different object parts and then modified according to BANA guidelines [315]. To create the BANA baseline renderings for User Study 3, we follow the same approach, using five pre-made textures tested with a blind individual for their identifiability. When creating illustrations with our technique for User Study 3, we created an outline around each part texture using Adobe Illustrator. We set the texture line thickness at 0.5 mm, and the outlines and gaps are the same as User Study 2. As some manual steps were used in creating the outlines, this post-processing stage prevented large-scale evaluation of our pipeline.

3D Reference Objects The 3D objects in User Study 2 were selected from online repositories: archibaseplanet.com, archive3d.net, and www.thingiverse.com. For User Study 3, the 3D reference objects were selected from the PartNet data set [295]. PartNet objects are pre-partitioned and provided in up-right orientation. Variations of the objects were created using Blender 2.79.

Fabrication Details The objects in User Study 2 were fabricated with a Form 3 SLA 3D printer. The objects in User Study 3 were fabricated with a UPrint FDM 3D printer and Form 3 SLA 3D printer. The same printer was used for all variations of an object for consistency.

To fabricate the tactile illustrations we used microcapsule paper. The capsules expand under heat application, which was applied with laser engraving. We used an Epilog laser cutter with settings: 100% speed, and power ranging between 12% and 15%, with maximum machine power being 60 W. Alternatively: 100% speed, and power ranging between 16% and 20%, with maximum machine power being 40 W. Each image takes 20 min to 30 min minutes to complete.

4.6 Discussion

4.6.1 Limitations & Future Work

We have presented a new methodology for generating tactile illustrations. We introduced a preliminary algorithm that integrates design considerations on tactile graphics motivated by a set of formative studies. However, many opportunities exist for improving the technique and experimenting with alternative stylizations. For example, further investigation could be done into selecting optimal viewpoints for parts, and layering images appropriately to handle occlusions. We also hope to scale up to more complex objects while maintaining a usable level of detail. Future work should investigate how shapes that are not approximated well by 1D skeletons can be represented (e.g., objects with concavities such as bowls). Future work can also develop treatment for objects with cycles. Currently, each part is positioned to coincide with its parent connection, and a part that closes a cycle is not taken into account.

Furthermore, symmetry has been shown to be a key characteristic of tactile shape perception [42], as well as from our own observation, and should be integrated into the design process in future work. Similarly, our method could incorporate special handling of repetitions within an object to communicate identical parts. Additionally, future work could incorporate other tactile saliency properties, such as the functional tactile saliency metrics developed by Lau et al. [231] which was based on grasping.

Our illustrations aim to convey shape characteristics that make designs distinct from one another. Future studies could also test the effects of our illustration style on object recognition tasks (e.g., recognizing ‘chair’ as the object category, rather than differentiating between variations of an object). Further, an ablation study was beyond the scope of our study resources but would be valuable as future work, for example, to study the effect of multi-projection rendering with standard textures.

Our technique relies on a semantically segmented object. Skeleton-based segmentation could be applied, e.g., Livesu et al. [257], Zhou et al. [484], Tagliasacchi et al. [400]. The segmentation could also be obtained with machine-learning approaches, which have recently shown improved performance over traditional approaches [199, 457].

Finally, other stylizations could be explored, such as exploded diagrams [239] which might allow for better visibility of occluding parts, abstraction using characteristic curves [289] or geometric primitives to handle more complex geometry, or abstractions such as skeletons to emphasize topological relationships. We also see opportunities for integration with interactive tactile graphics, e.g., combined with interactive audio feedback, dynamic markers [397], or refreshable tactile graphics displays where low resolution pin-based stimuli will present new challenges [63].

4.6.2 Conclusion

We have successfully built a general pipeline for the design of tactile illustrations to improve 3D shape understanding for blind individuals. We have also designed a user study to evaluate the effectiveness of the illustrations, which we implemented on a selection of common household

objects. We are the first to create a systematic approach that can generate tactile illustrations for objects with complex 3D geometry. We demonstrate promising results, showing significant improvement compared to baseline guidelines for tactile illustration design.

Providing a tool that creates tactile illustrations is integral as a resource for the blind community to allow them to gain access to visual information. For example, our illustrations could be used as a tool to depict products, from furniture in a store catalog, to art pieces in a museum, to object repositories for 3D printing. An exciting area of future work would be to use our illustration algorithm as a language to build a design tool for the blind. As one of our study participant mentioned: “One of the biggest difficulties blind people have is getting an idea out of my head and into your head.”

5

Summary and Future Work

In this thesis I present two works for 2D fabrication of images that cover two different tasks. For the project discussed in Chapter 3, we implemented a fabrication-aware pipeline for the design of watercolorized images, including an approach for the design of single and multiple continuous color transitions (i.e., gradients). We also designed a repeatable fabrication procedure that is coupled with the design process and succeeds in creating images with real watercolor pigments. For the project discussed in Chapter 4, we proposed an algorithm for the design of tactile images with improved shape understanding for the blind, and we tested the success of the images by formatting a user study about shape understanding.

The first work proposes an image analysis pipeline for the design of woodblocks, which, coupled with a fabrication process, results in controllable watercolor images. The input image is broken down through our image analysis pipeline in a series of maps that reflect the placement of paints. This pipeline consists of a series of extraction steps including: color palette, smoothness values, relatively darker and lighter in color-brightness areas, and edges. The fabrication process was the outcome of an experimentation with different materials, tools, and fabrication procedures. The images succeed in resembling watercolor paintings, including multi- and single-color gradients, while at the same time clearly conveying their content, as shown from our informal user study.

The second work proposes a multi-projection rendering that more clearly depicts 3D shapes through touch compared to images designed following standards in tactile illustration. The 3D

object input is geometrically processed into sequences of cross sections, taking into consideration its semantic segmentation. An analysis of the geometry of its segments and its whole determines the set of camera positions. The cross sections are aligned to these cameras and rendered. The complete image is composited using a multi-projection approach which aligns the multiple camera renderings following the information of connectivity between segments and using rotations and translations. The characteristics that such tactile images should possess are gathered through our formative studies as well as reading material in fields related to touch perception, such as cognitive psychology. The success of our tactile images is shown by a user study we design and run in the local blind community, with respect to shape understanding.

Our works, being initial steps in our field that focus on watercolor and tactile fabrication, can guide and inform future research. First, they will provide baselines. In watercolor fabrication, a set of images that are commonly watercolor painted themes similar to those used in the paper's data set (e.g., a landscape, a portrait, an animal, and a flower), could be fabricated with other approaches and compared in terms of effects with our approach. Also, for fabrication of tactile images, a new design technique for illustration for shape understanding could compare their results with our own, by running a study similar to the one presented here.

Two other main ways that our work could guide future research are: building upon our pipelines to improve their efficiency and the quality of the designs, or exploring new design spaces and a wider variety of designs. Section 5.1 gives examples of ways that the pipeline can be improved to increase efficiency, Section 5.2 details different design approaches, Section 5.3 discusses novel design ideas, while Section 5.4 broadens the contributions of these works outside their initial tasks.

5.1 Changes in Pipeline

Watercolorized Images While our image processing algorithm is simple to implement and successfully extracts the corresponding maps and elements, other techniques may be able to improve the results. For instance, Lin et al. [243] create color palettes based on a data set of human-extracted color themes, human-extracted themes could embody color selection

based on artists preference. Also, other pixel-level defocus techniques [478,485] might lead to more accurate smoothness maps. Similarly, several other options exist for computing a semantic map [311,425] which would lead in different image segmentations. The use of other image morphology algorithms for edge extraction and vectorization [271,275,445] could create differences in the details map. Specifically, for the edge guided polygon expansion, a clipping algorithm such as [281] guiding a curve offset [459] which takes into account self-intersections could be used.

Tactile Images In tactile images, in terms of fabrication, taking into account the way micro-capsule expands and introducing half-toning [307], could lead to clearer illustration design. In terms of design, we could experiment with other vector fields for the initialization of the up-vector cameras, as well as with other skeleton extraction algorithms. Also, we could sample different in place cross section rotations around the corresponding skeletal sections before their layering, and choose a rotation that results in better visibility of the outline curves.

A common line of work would be to make both of these pipelines more automatic. For example, regarding the fabrication of watercolor images, new image analysis techniques [25] could exclude the need for user given input, e.g., a scribble for semantic segmentation. Implementing a vectorization algorithm would also remove the need of using an independent software to process the rasterized input images. Similarly, in the tactile pipeline, our technique relies on a semantically segmented object. Skeleton-based segmentation could be applied, e.g., [257,400,484]. The segmentation could also be obtained with machine-learning approaches, which have recently shown improved performance over traditional approaches [199,457]. It might be possible to design a vector field [130,475] that would align the cameras without the need for the alignment operation we followed.

5.2 Changes in Design

Another direction for future work would be to investigate how the existing designs could be improved.

Watercolorized Images Key work in watercolor design can address color management, simulation, and abstraction and optimization.

Key directions regarding color focus on color mixing, color gradients, and color reproducibility in different input sizes. As revealed in the informal user study, as well as other works, e.g., Flagg et al. [132] which provided a color mixing tool for oil paints, color selection is an important feature. Future work could include an automatic mixing tool for watercolor paints. To account for the transparency of the medium and layering effects (e.g., glazing) that produce new color hues when pigments overlap, the field can build upon the insights of Babaei et al. [33] and Tan et al. [403] on color layering. Regarding color gradients, we used more or less diluted paints. However, artists, in order to design gradients, use different color hues (cool hues or complementary colors), which lead to the creation of gradients. This use of complementary color paints could, for example, allow for the creation of even darker gradients, if the paints are placed by a carefully designed woodblock. Regarding reproducibility, one area for future work is the achievement of more detailed color reproduction if the goal is realism. A small set of paints limits our ability to reproduce color details, which might negatively impact some texture, e.g., the trees in the landscape. This can be mitigated with an increase in paint colors at the expense of a longer printing process and more detailed annotation that is more complex to produce. Finally, regarding fabrication of color gradients, we did not experiment with various sizes in input images. We expect that parameters like the dilation radius need to be scaled proportionally to the print size, but more thorough investigation is necessary.

Other key directions for future research that also need to be fabrication aware include simulation, abstraction, inclusion of more water levels, and minimizing number of woodblocks. One direction could be predicting the faithfulness of the fabricated prototype as a function of the number of the colors, the amount of water, and the printing order, such as by using a simulation

interface. In this way, we could also use natural watercolor diffusion not only to stylize blurred areas with gradients, but also to approximate them.

Additionally, augmenting the abstraction of the print is of interest. For example, simplification of the input image using non-photorealistic rendering or watercolorization techniques can modify the abstraction and uniqueness of the print. An interesting avenue for future work could be to identify characteristic elements of watercolor automatically, similar to the work of [116]. Another direction for future work is refining our diffusion model to support more levels between our current ‘dry’ and ‘wet’ options, which would require a finer-grain analysis of the local image smoothness. Finally, we have experimented with using woodblocks for more than one color to simplify the printing process. Our early experiments done by hand are promising (Figure 5.1) and suggest that it should be possible to design an algorithmic solution.



Figure 5.1: *Manual combination of two (here) or more colors per block, result in a smaller number of blocks and the automation is left as a direction for future work.*

Tactile Images In tactile illustrations, many opportunities exist for improving the technique and experimenting with alternative stylizations.

For example, further investigation could be done into selecting optimal viewpoints for parts, and layering images appropriately to handle occlusions. Also, future work could scale up to more complex objects while maintaining a usable level of detail. Future work should investigate how shapes that are not approximated well by 1D skeletons can be represented (e.g., objects

with concavities such as bowls). Future work can also develop treatment for objects with cycles. Currently, each part is positioned to coincide with its parent connection, and a part that closes a cycle is not taken into account.

Furthermore, symmetry has previously been shown to be a key characteristic of tactile shape perception [42], as our user study also indicates, and should be integrated into the design process in future work. Similarly, our method could be developed to incorporate special handling of repetitions within an object to communicate identical parts. Additionally, future work could incorporate other tactile saliency properties, such as the functional tactile saliency metrics developed by Lau et al. [231] which was based on grasping.

Future design directions for tactile illustrations could be related to how other visual cues can be substituted with tactile cues. For example, in vision, cues related to shadows, perspective distortion, and occlusions play a role in shape understanding. Specifically, as occlusions help understand depth, ideal approaches might not minimize or remove occlusions [124]. Approaches that tactily annotate different occlusion types (e.g., containment, contact, and blocking) using textures or a series of images that reveal where occlusions occur [22, 109, 377] might be preferable. Similarly, connections might be sufficiently annotated with symbols and not with actual connectivity [387].

Finally, other stylizations could be explored, such as exploded diagrams [239–241, 404] which might allow for better visibility of occluding parts (Figure 5.2), abstraction using characteristic curves [289] or geometric primitives to handle more complex geometry, or abstractions such as skeletons to emphasize topological relationships. We also see opportunities for integration with interactive tactile graphics, e.g., combined with interactive audio feedback, dynamic markers [397], or refreshable tactile graphics displays where low resolution pin-based stimuli will present new challenges [63].

Another direction of future work, related to both watercolor and tactile illustrations, is the careful placement of distortions that draw attention. For example, distortions can create new styles by augmentation of semantic areas or color changes. Another type of research could focus on drawing attention to parts of the fabricated image, similar to research in virtual

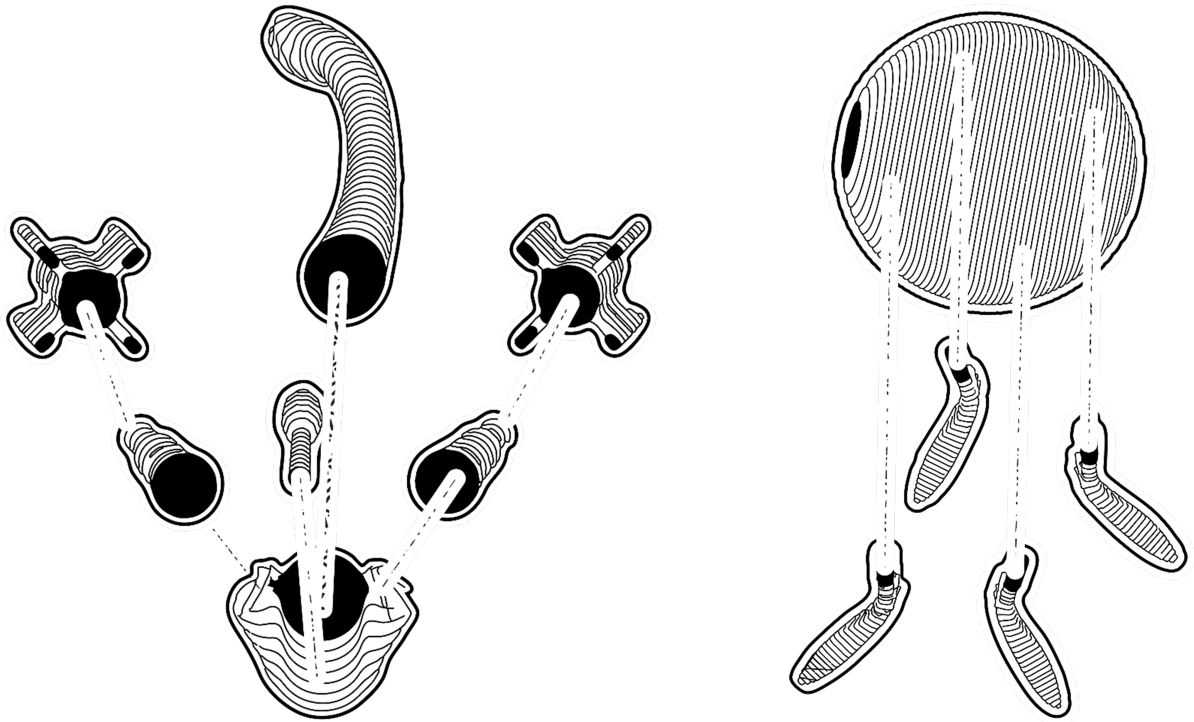


Figure 5.2: *Multi-projection exploded line drawings*

images [431, 435], in which color hues and image distortions were used to highlight results, and improve visibility. In tactile images, we noticed that proportions need not be accurate. Therefore, resizing might be an important distortion to allow for better visibility of parts.

The next two sub-sections discuss ideas for future design in each of these two design spaces.

5.3 Exploring New Designs

Future work could also investigate new designs and design spaces.

While considering that 2D line-drawings have low success rates for object recognition, a closely related design space which has not been researched is 2.5D models, which define a height field over an image and resemble sculptural reliefs [446]. 2.5D models have been used in visual pieces of art [27, 90, 206, 386, 388] and researched for that purpose in graphics [207, 386, 446]. They have also been used as tactile illustrations for the blind [73, 137, 137, 210, 346, 391]. Applying ideas from our project, such as annotating different parts with texture or curvature and selecting camera viewpoints, might improve the quality of illustrations.

Another direction of future work, related to both watercolor and tactile illustrations, could help explore new designs. For each design-oriented project, the first step is for the researchers to gain knowledge of the corresponding field. Experts need to understand the underlying design procedure, identify common design and fabrication patterns, understand how the output results from a given input, and think abstractly to adapt the design and fabrication processes for new inputs. The process of acquiring skills in new design spaces is often slow and difficult. Because of this, a more systematic way to identify good solutions without the need for researchers to understand the design space would be extremely helpful.

One benefit would be the ability to create designs that are significantly different from already created designs. For example, while the watercolor pipeline has parameters that change the design, such as number of colors and the smoothness threshold, it cannot create significantly different results. To this end, recently, neural network techniques try to approach previously unexplored spaces. Generative networks can infer rules for the creation of new designs. This has been applied to the design of, for example, caricatures [77], furniture [268], or artistic portraits [464].

A second positive outcome would be a systematic exploration of the design space that allows for improvement of the future designs. As Hertzmann discusses, non-photorealistic rendering might allow scientific theories of art [176] to be inferred. For example, this may occur by describing the outcome in terms of an objective function that will “maximize the desired visual response in the viewer, subject to the constraints of the artistic medium” [176]. Accurate reproduction of an image is not, by definition, the desired outcome of non-photorealistic rendering, e.g., artists can exaggerate in terms of colors [374]; therefore, an evaluation of design preference is still an unanswered question. At the moment, because of the unpredictable nature of the watercolor material and the need for tactile images to be tested by blind individuals, the design improvement is a demanding process.

One possible approach could be the use of networks, which would input and score designs. Then, by slowly changing designs, new designs could be generated, e.g., by changing colors in watercolor woodblocks, or by changing camera positions by rotations in tactile illustrations, then scored, and later tested in real world settings. For example, evaluation for tactile images

could happen in collaboration with a fine arts museum, while evaluation of watercolors could occur in collaboration with a school of fine arts. This exploration of different design spaces could lead us to the implementation of a more general framework for fabrication of different types of images and inform us on how to analyze different illustration styles [140].

5.4 Interdisciplinary Collaborations

Fabrication works, being interdisciplinary by nature, greatly benefits by collaborations with other researcher areas, as shown by the works in this thesis (Figure 5.3).

Artists could help us abstract ideas and rules for better image composition [148]. For example, when scoring the training data set for the neural network, it would be useful to consult with domain experts in order to create a data set that would be artistically complete, or to help us separate different styles. Artists could also inform us about techniques that exaggerate characteristics of shapes and might be of use for tactile images. And they could suggest new artistic ways for the use of materials, such as the microcapsule paper which can detailed and accurate.

For the continued development of tactile illustrations, researchers in cognitive science and psychology could collaborate with graphics researchers to further improve upon the designs. For example, studies for which the response of the individuals could be measured in terms of neural firing could give more objective indications as to which illustrations are more successful in shape understanding. Another avenue for inter-disciplinary collaboration could identify human perceptual limitations. For example, researchers might be able to help in designing studies to identify how many different parts can be conveyed accurately within a tactile illustration. Ideally, this type of research could make tactile images of shapes more intuitive to understand. The tactile imagery could be then used as a depiction of drafting tools for the blind, and in this way, individuals with blindness could become, on their own, designers of artifacts.



(a) Book arts workshop



(b) Nesin mathematics village

Figure 5.3: *Interdisciplinary collaborations allow to integrate ideas from different domains. In this thesis, knowledge was shared by artists, such as print-makers (a) and mathematicians (b), among others. [Photograph Sarah H. Smith ©(2018)] [Workshop, Women in shape modeling, İzmir, Turkey, (2016)]*

5.5 Conclusion

In this thesis, I presented two works in expert-domain fabrication of 2D images, one pertaining to the use of watercolors and another to tactile illustrations. The approaches introduce, respectively for the two tasks, systematic approaches for the image design using graphics techniques and the materialization of these images using fast prototyping tools. For the first project, an image analysis pipeline infers key information from an input image and designs stamps that apply the watercolors on paper. For the second project, a mixed geometric and image analysis pipeline extracts important information from the input 3D model, which is used to render, from multiple viewpoints, cross sections that are transferred on low-relief paper. In both projects, the fabrication, i.e., the transfer of the 2D images on the wood-stamps and on the microcapsule paper, is achieved using a laser cutter. To test the performance of these new designs, we interviewed people in the local art and blind communities, and we received positive feedback regarding watercolor depiction and shape understanding from tactile images. We hope that these two works can initiate more research in the corresponding fields, i.e., artistic image and tactile image fabrication and in non-photorealistic image fabrication in large.



Related Research

Here, I highlight three projects I was involved in that were published during my PhD. These works, related to geometric modeling and analysis, helped inspire my approach to the tactile project presented in the previous chapter. The first project, related to improving the quality in 3D printed objects, demonstrated the importance of viewpoint on the object display, while the other two projects, on 3D printing of wind musical instruments and designing of 3D shapes based on 1D skeletons, showed the importance of skeletons in shape understanding.

A.1 Viewpoints

In the first project, Zhang et al. [477], we computed viewpoints that improved the quality of 3D printed models by selecting printing direction. Printing direction affects the location of support structures, which leave artifacts on the model after their removal. The quality of the final model is then compromised when support structures are placed in visually important regions. Our model for preference in 3D printing direction was formulated as a combination of metrics, including area of support, visual saliency, smoothness preservation, and viewpoint preference.

As concerns the viewpoint preference, we followed the approach developed by Secord et al. [372], which measures viewpoint preference based on human perception. A series of attributes were computed and were then linearly blended. The attributes were: projected area (a_1), surface visibility (a_2), viewpoint entropy (a_3), silhouette length (a_4), silhouette curvature

(a_5), silhouette curvature extrema (a_6), max depth (a_7), depth distribution (a_8), mean curvature (a_9), Gaussian curvature (a_{10}), mesh saliency (a_{11}), above preference (a_{12}), and eye placement (a_{13}). From these metrics, I implemented the algorithms relevant to *Area attributes*(a_1 - a_3), *Silhouette attributes*(a_4 - a_6), *Depth attributes*(a_7 - a_8), and *Semantic attributes* (a_{12} - a_{13}), but not *Geometric attributes* (a_9 - a_{11}) since they were already available. Of these metrics I focused on shape characteristics that can be computed, and therefore perceived, directly on the image, using simple line drawings and without prior knowledge on the object geometry(a_1 - a_8).

We then used these attributes to score a printing direction. To do so, the final score of viewpoint preference was obtained by combining the five most dominant attributes: projected area, surface visibility, silhouette length, max depth, and above preference. To choose attributes, we tested different attribute combinations and the weights provided by Secord et al. [372]. The five most dominant attributes gave us more satisfactory results compared to other combinations. In addition to the already existing attributes, one new attribute was implemented named user preference(a_{14}). This allowed us to modify this automatically computed score by considering the viewpoint that the user preferred.

Below, I define each of the previously mentioned attributes and provide a short description for their use. The input 3D objects are manipulated to be centered in the origin and fit in a unit cube.

We use the depth buffer [448], $DEPTH[] \in (0, 1)^{Image_Size}$, a matrix which contains a depth value for each image pixel. Also, let b_{pixel} be the set of pixels that occupy the background $\{pixel | DEPTH[pixel] = 1\}$, and o_{pixel} be the set of pixels that occupy the object on the image plane $\{pixel | DEPTH[pixel] < 1\}$, while $total_surface$ is the total surface area of the image plane.

Projected area (a_1) is the number of pixels on the image plane that display the object, over the total number of image pixels for a given viewpoint. This metric increases as more image surface is occupied by the input object and is maximized when all image pixels are covered by the object. To give some insight into why this metric matters for viewpoints, we can simply say that this attribute penalizes images for which only a small portion of the image displays the

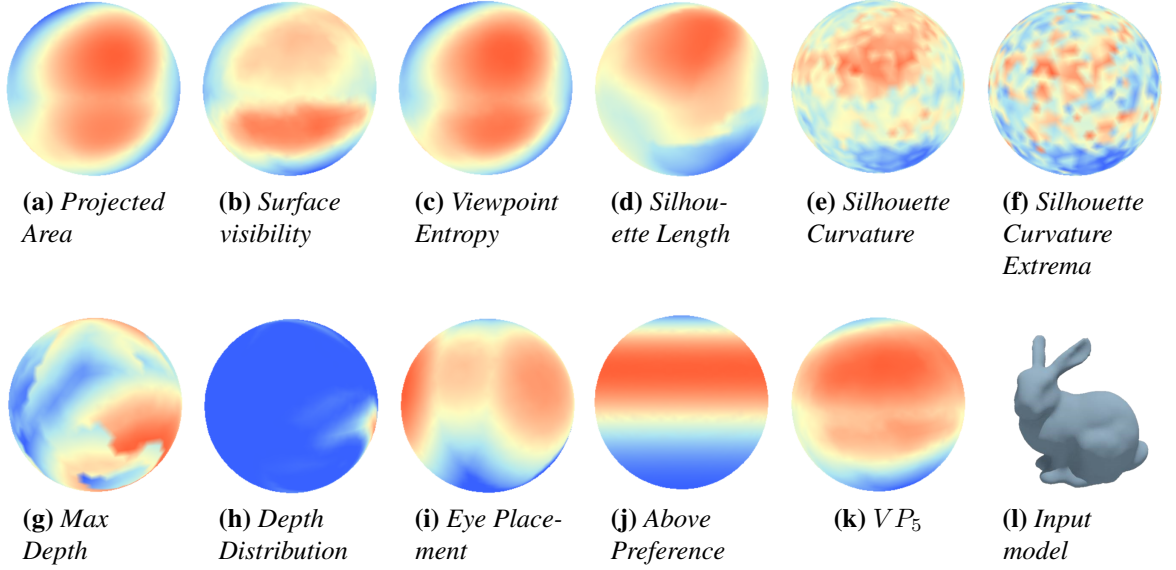


Figure A.1: Perceptual attributes for the bunny model.

object of interest.

The metric is defined as $a_1 = \frac{|o_{pixel}|}{total_surface}$, i.e., the number of pixels that display the input object (Figure A.1(a)).

Surface visibility (a_2) is the surface area of all visible object polygons over the entire object surface area. This metric increases as more surface of the input object is displayed on the image, and is maximized when all object surface is visible. By containing a large object surface area in the image, we succeed in achieving higher visibility for the object we are interested in.

For each triangle of the mesh, we compute its surface area $AREA^{tri}$ using Heron’s formula and its centroid $centr^{tri}$, and we project the centroid to the image plane as seen from the viewpoint which results in an approximate pixel location $pixel_{pr}^{tri}$ and its depth value on the image plane d_{pr}^{tri} . To be more precise, by saying “as seen from the viewpoint,” we are referring to projection on the image plane applying the camera projection matrix where the camera is placed at the input viewpoint. The metric is defined as $a_2 = \frac{\sum_{tri \in VIS} AREA^{tri}}{\sum_{tri} AREA^{tri}}$, where VIS is the set of visible triangles defined as $VIS = \{tri : |d_{pr}^{tri} - DEPTH[pixel_{pr}^{tri}]| < Threshold\}$ (Figure A.1(b)).

Viewpoint entropy(a_3) defines a distribution over the visible polygons on the image plane. It increases when more space is occupied by the object on the image and the rendered polygons are evenly visible, and it is maximized once every visible polygon and the unoccupied area have equal visibility. By using this measure, not only does projected area increase, but polygons, that could be ignored in the surface visibility metric, are also taken into account for the metric computation.

Here, for each visible triangle a probability is defined, which shows how much each visible triangle occupies the total image space $P(tri) = \frac{AREA^{tri}}{total_surface}$. Next, for each visible triangle and the background, we add the above scores to compute the metric, i.e., $\sum_{tri \in VIS} P(tri) \times \log_2 P(tri) + P(background) \times \log_2 P(background)$ (Figure A.1(c)).

Silhouette length (a_4) defines the outline of the object on the image plane. It increases when more details are uncovered and the corresponding parts become visible. While it is not known when it is maximized, from our observations, it increases as more parts occupy the image plane and do not coincide within it. Therefore, we use silhouette length to minimize part occlusions.

To measure the object silhouette length, we identify all limit pixels or l_{pixels} , which are pixels rendering the object on the image plane that have at least one neighboring b_{pixel} . In a row-by-row fashion, we mark sequences of consecutive l_{pixels} . Let the number of pixels for a consecutive sequence be l . We consider only the sequences that have $l > 1$: if only the first or only the last sequence pixel has a single b_{pixel} neighbor, then the sequence contributes to the object silhouette, length equal to $\sqrt{(l^2 + 1)}$; if not only the first or last sequence pixel has a single b_{pixel} neighbor, then the sequence contributes length l . The contributing pixels are marked as background pixels, and the same procedure is repeated in a column-by-column fashion, without the limitation of $l > 1$. The metric a_4 is computed by summation of the l values (Figure A.1(d)). While there are already-implemented techniques for silhouette extraction, as seen in the next attribute, we used our approach because it was more accurate.

Silhouette curvature (a_5) defines the average curvature of the outline of the object on the image plane. It increases as more curves of the object appear on the image plane. For a closed

oriented silhouette the absolute curvature is greater than the 2π [83], given that the input mesh is a single well-defined shape. The silhouette curvature is a number that depends on the 3D input object curvature and genus, as well as the viewpoint selection. Large silhouette curvature allows the image to depict not only large outlines, but parts that are curved and are represented by a succession of hills and valleys.

For computation of curvature we use the silhouette, which is extracted using “findContours” in OpenCV [379]. The reason for not using the technique of the previous attribute, is that here the silhouette can be easily retrieved as n contours, i.e., n point sequences of size s : $\forall i = 1 \dots n, C^i = [c_1^i, c_2^i, \dots, c_{s_i}^i]$ which create the outline of the object. The computation of curvature, based on [203], by integration over all the points, is then defined as the summation of the clockwise positive angles for all consecutive three points in the contours (c_i^-, c_i, c_i^+) , i.e., $angle_i = CLOCKWISE_ANGLE^+(c_i - c_{i-1}, c_{i+1} - c_i)$ (Figure A.1(e)).

Silhouette curvature extrema (a_6) defines the extremity of curvature over the entire silhouette. It increases when points on the silhouette on the image plane have high sharpness. Silhouette curvature extrema is maximized when consecutive contour segments with large absolute angles and short distance are visible on the image plane. It gives higher importance in the depiction of the highest curved protrusions and craters of the model.

For computation of the silhouette curvature extrema we use the same definition as in a_5 . The integration over all contour segments results in summing for all i the quantity $angle_i^2 / arc_length_i$, where $arc_length_i = 0.5(|c_i - c_{i-1}| + |c_{i+1} - c_i|)$ (Figure A.1(f)). Here, we note that we have preprocessed the contour points for the computations of a_5 and a_6 , so that every angle is strictly positive (nonzero) for every i .

Max depth (a_7) defines the depth value of the pixel that appears to be the farthest away from the camera. It increases as the distance between the closest and the farthest away visible points on the image screen increases, given that the object is centered in the origin. The max depth value is limited by the diameter of the object’s convex hull, but it cannot always approach that limit, since the farthest away point, when looking from a viewpoint close to the diameter, is not

always visible. This metric allows for depiction of a deep side of the model, and its computation is straightforward: $\text{DEPTH}[\arg \max_{px, \text{DEPTH}(px) \neq 1} \text{DEPTH}(px)]$ (Figure A.1(g)).

Depth distribution (a_8) defines a distribution over the depth values of the visible pixels. It increases as depth values of all pixels become more evenly distributed, and it is maximized when exactly the same number of pixels is set to all different pockets of the histogram, i.e., uniform distribution of depths. Similarly to the previous metric, it allows for consideration of depth of view, but not only with the appearance of one extreme value, but for the appearance of wide range of depths.

Depth distribution is defined over a constant number of bins, $|B|$. The depth values, that range from 0 to 1, fill the $|B|$ bins. Specifically, the i th bin has the value $val_i = \frac{1}{|o_{pixel}|} \sum_{px \in o_{pixel}} \mathbb{1}_{\text{DEPTH}[px] \in [\frac{i-1}{|B|}; \frac{i}{|B|}]}$. The metric of depth distribution is then defined as $a_8 = 1 - \sum_i |B|^{-1} val_i^2$ (Figure A.1(h)).

Mean curvature (a_9), **Gaussian curvature** (a_{10}), **Mesh saliency** (a_{11}) These three attributes are geometric characteristics of the object and are more focused on curvature. Mean and Gaussian curvature are two different ways to compute how curved or flat the surface is at different areas of a 3D model, while saliency is described as a measure that finds salient characteristics through computing mean curvature at different resolutions of the model. However, I will not go into details here since these geometric attributes did not inform the work in this thesis.

Above preference (a_{12}) measures how much the view deviates from a visual canonical viewpoint. To compute this metric, the model needs to be placed in its up-right position. The attribute value is the same for all positions on any circle parallel to the model base. It is maximized at viewpoints that are placed on the circle of 112.5° -latitude, which is 22.5° above the equator, the equator has 90° -latitude. This attribute captures findings from psychological studies about human vision preferences, indicating that people prefer to view objects at a higher angle from their canonical position.

The up-right orientation of the object coincides with the z-axis. To compute the latitude for a given viewpoint $v_i = (x_i, y_i, z_i)$ we define $vf_i = (x_i, 0, z_i)$ and $angle_i = 180 - \text{ANGLE}(v_i, vf_i)$,

resulting in the latitude being $\phi = (90 + angle_i) \cdot \mathbb{1}_{y_i < 0} + (90 - angle) \cdot \mathbb{1}_{y_i \geq 0}$. For a generic latitude ϕ , $-90^\circ \leq \phi \leq 90^\circ$, the attribute of above preference is computed as: $a_{12} = \exp\{-(\phi - \mu)^2 / \sigma^2\}$, with $\mu = 67.5^\circ$ and $\sigma = 45^\circ$ (Figure A.1(j)). Please note that compared to previous metrics, this metric does not take into account the object or its image depiction.

Eye placement (a_{13}) measures how much the eyes are visible from different viewpoints. To compute this metric along with the model, the area of the eyes is indicated on the model by the user, meaning that the model has to display a face. The eye-placement attribute increases as viewpoints approach full front views and is maximized when the entire eye area is visible. This attribute again captures psychological findings that in contrast to inanimate object views, animate object views need to consider the eyes in the illustration.

This metric is computed similarly to a_1 . The only difference is that instead of counting the pixels that are object pixels, the pixels that display eyes are counted. To do so, eye surface is color-coded with red, and therefore, the corresponding pixels can be retrieved by accessing the color buffer. The metric is defined as $a_{13} = \frac{|eyes_pixel|}{total_surface}$ (Figure A.1(i)).

User preference (a_{14}) defines a harmonic field on the viewing sphere influenced by a user preferred viewpoint. For that purpose, in our paper, an interactive tool was developed to further incorporate the viewing preference of a user into the perceptual preference score. The attribute is maximized for the user preferred viewpoint where its value is 1. One example of viewpoint preference computation can be seen in Figure A.2, where the user had a strong preference for the front viewing direction of the bunny, which resulted in giving a boost to the values depicted on the right side of the image (red means higher values).

By selecting a strongly intended viewing direction \mathbf{v}_a , a Harmonic field $H(\mathbf{d})$ is computed on the sphere as $\nabla^2 a_{14}(\mathbf{d}) = 0$, s.t. $a_{14}(\mathbf{v}_a) = 1$, $a_{14}(-\mathbf{v}_a) = 0$ (Figure A.2 top right). Here, note the attribute does not depend on the geometry of the model or the rendered image.

Viewpoint preference (VP_5) In the 3D printing paper, this metric measures the preference of a viewpoint based on the perceptual model introduced by Secord et al. [372]. The five most dominant attributes were employed in our approach, which include projected area (a_1), surface




Figure A.2: *Interactive modification of viewpoint preference – (top left) automatic computation $VP_5(\mathbf{d})$, (top right) Harmonic field $a_{14}(\mathbf{d})$ according to user specification and (bottom) the modified field by average and normalization. Blue denotes zero and red denotes one (highest preference). To display the overall distribution, the spheres are also unfolded onto a rectangular region.*

visibility (a_2), silhouette length (a_4), max depth (a_7) and above preference (a_{12}). The final score of VP_5 (Figure A.1(k)) is obtained by linear blending of these five attributes using the weights provided [372]. When integrating User Preference, the metric was then updated by the average of $VP_5(\mathbf{d})$ and $a_{14}(\mathbf{d})$ and then normalized.

After having explained the key attributes for the computation of viewpoint preference, I now turn to a description of how this project helped me choose preferred viewpoints in my thesis work.

Project Take-away This project helped inform the tactile illustration project (Chapter 4), particularly, the identification of preferred viewpoints for tactile shape understanding. To create tactile illustrations (Chapter 4), the viewpoint preference was obtained by using the silhouette length attribute. This time the goal was not a user preferred visual viewpoint, but a viewpoint that would result in less occlusions on the image plane, trying to identify how to compute a canonical view that would allow for multiple part visibility. To identify which attributes were

more relevant, we computed the above metrics for a set of objects, and we noticed that for all of these metrics, the one that consistently gave us the best outcome was the maximization of the silhouette length. This project also informed us that we could gain information about a model, not only from its geometric-based characteristics but also from image-based characteristics, which led to a fully image-based project, the watercolor, and a partially image-based project, the tactile paper.

A.2 Skeleton

The following two projects emphasized the importance of skeletons for shape design, as well as the ability to extract a skeleton from a given shape. Based on this experience, skeleton extraction was a natural step for me to take when conveying 3D shape information with tactile images.

A.2.1 From Shape to Skeleton

Computer graphics researchers have developed numerous techniques for the extraction of skeletons from a given shape. In this project, Umetani et al. [420], we designed and fabricated 3D objects with hollowed cavities that could be used as musical instruments. To provide a tool to non-experts for the design of wind-musical instruments, Umetani proposed a fast simulation algorithm for sound of free-form cavities. To design the free-form objects in Figure A.4, we followed an approach from Igarashi et al. [184]. With this approach, the user first designs a closed free-form curve on the computer screen, and the skeleton of this boundary curve is then extracted. The extracted skeleton can be further used to manipulate the shape, e.g., by inflation.

Here, the process of a 3D free-form object is described. An outline is drawn on the 2D image plane. If the outline is not intersecting, it is closed by drawing a straight line from the starting point to the ending point. Points are then uniformly sampled on the closed outline. The skeleton is then extracted according to the following process (a detailed description can be found in Igarashi et al.):

1. Sample the closed curve in equally spaced points (outline of Figure A.3)




Figure A.3: *An example of skeleton extraction.*

2. Perform a Delaunay triangulation on the closed curve area (blue lines of Figure A.3)
3. Compute the chordal axis of the shape (also referred to as the spine or the skeleton) (non-blue lines of Figure A.3)
4. Fan the terminal non-significant triangles and prune the skeleton (non-pure red chordal axis of Figure A.3)

After performing the skeleton extraction, the triangulation is refined, the skeleton is elevated, and the internal edges of the triangles are elevated to fit quarter ovals. Specifically, the elevated spine edges have height equal to the distance from the boundary, i.e., for every spine vertex, the mean distance to its neighboring outline vertices is used to elevate the spine toward the positive z -direction. The spine is then smoothed by adding the height of each spine vertex to the two neighboring vertices and then dividing with the total number of vertex heights added, i.e., 3 for internal vertices and 2 for external. Next, each edge connecting a spine $s = (s_x, s_y, s_z)$ and a non-spine $i = (e_x, e_y, e_z)$ vertex are sampled in 4 equally-spaced segments, with internal vertices $v^i, i = 1, 2, 3$ with z -coordinates $v_z^i = \frac{4-i}{4}(s_z + e_z)$, y -coordinates: $v_y^i = m_y + \text{sign}(e_y - m_y) \cdot \sqrt{\frac{5+2i}{16}} \cdot |m_y - e_y|$, and x -coordinates $v_x^i = m_x + \text{sign}(e_x - m_x) \cdot \sqrt{\frac{5+2i}{16}} \cdot |m_x - e_x|$. The newly created vertices are then triangulated by matching non-spine/non-outline edges in the triangles. A symmetric shape on the negative z -half-space is then constructed, and the positive

and negative z -shapes are glued on their seams. The shape is smoothed and can be further manipulated, e.g., resized or cut. Finally, it is given as input to the simulation interface with the aim of fabricating a 3D freeform wind instrument, e.g., Figure A.4. Using the fast simulation approach of Umetani, we then predicted the resonance frequency. This frequency was used to compose sound, using inverse FFT, which we provided as feedback to the instrument-design interface. Finally, we fabricated the instruments, using 3D printers, and compared the resonance frequency between the simulated and the recorded sound. Now, having finished a description of the technical information about the wind musical instrument project, I explain how this project informed my approach to 3D shape handling for my thesis research.




Figure A.4: *Designed musical instruments which were designed based on the skeleton extracted from a 2D closed curve.*

Project Take-away This project informed me about how important the meshing can be on the results of an algorithm. Specifically, the proposed algorithm is influenced by how well the object is meshed, and does not allow modifications which change the number of vertices on the original mesh. Specifically, while in the beginning we planned to incorporate the ability to cut objects in the user interface, we later noticed this would not be possible. Similar to this work, I noticed in later projects, such as the tactile illustration project, that a well-defined and with nice flow as input achieves more accurate results for many algorithms, including skeleton extraction.

Also, the importance of skeleton extraction was highlighted. During the wind musical instrument project, the skeletons were extracted from a 2D rather than a 3D shape. As this technique could not be directly applied to the new project for which the shapes were 3D, we

used a completely different algorithm [407] that was already available. However, one approach to apply this technique to 3D shapes could be using multiple views from a part, extracting skeletons for each 2D view and constructing the skeleton in the 3D space using some averaging or heuristic step, e.g., the maximum silhouette from the previous paper.

A.2.2 From Skeleton to Shape

Computer graphics have also used skeletons as a starting point for shape design [184]. The approach we took in the third project, Panotopoulou et al. [322], thus allows for the creation of a 2-manifold mesh of surfaces of high genus, starting from a schematic description of the shape, the skeleton (Figure [390]).

The aim of our project was to construct a quadrilateral mesh around a skeleton that was as coarse as possible, similar to a scaffold on which a surface is built. Such a mesh provided the first step for defining a joint surface representation for a skeleton-based shape; the scaffold quads can be refined by subdivision to fit a regular quad mesh to a surface shape, as in [422]. Our construction was optimal in the sense that it only consists of quads, in contrast to other methods of constructing meshes from skeletons which also use triangles [64, 195, 422]. In addition, our method was minimal in the sense that it is not possible to construct a coarser quad mesh of the same skeleton with uniform sections. In other words, every section consists of the same type of polygon (i.e., uniform sections), such as triangle, quad, or hexagon. Finally, we presented a selection of examples of skeletons for which we generated a scaffold with the described method (Figure A.6).

A *skeleton* is defined by its topology and its geometry. The topology is given by a simple, unweighted, undirected and connected graph $G = (V, E)$. Here, V denotes the set of vertices, and E denotes the set of edges. Each vertex v of V is associated with a point p_v in \mathbb{R}^3 and a radius $r_v > 0$, which defines a sphere centered at p_v . In other words, a *skeleton* is defined as the triple (G, p, r) , where

- $G = (V, E)$ denotes the graph defined above,
- $p: V \rightarrow \mathbb{R}^3$,

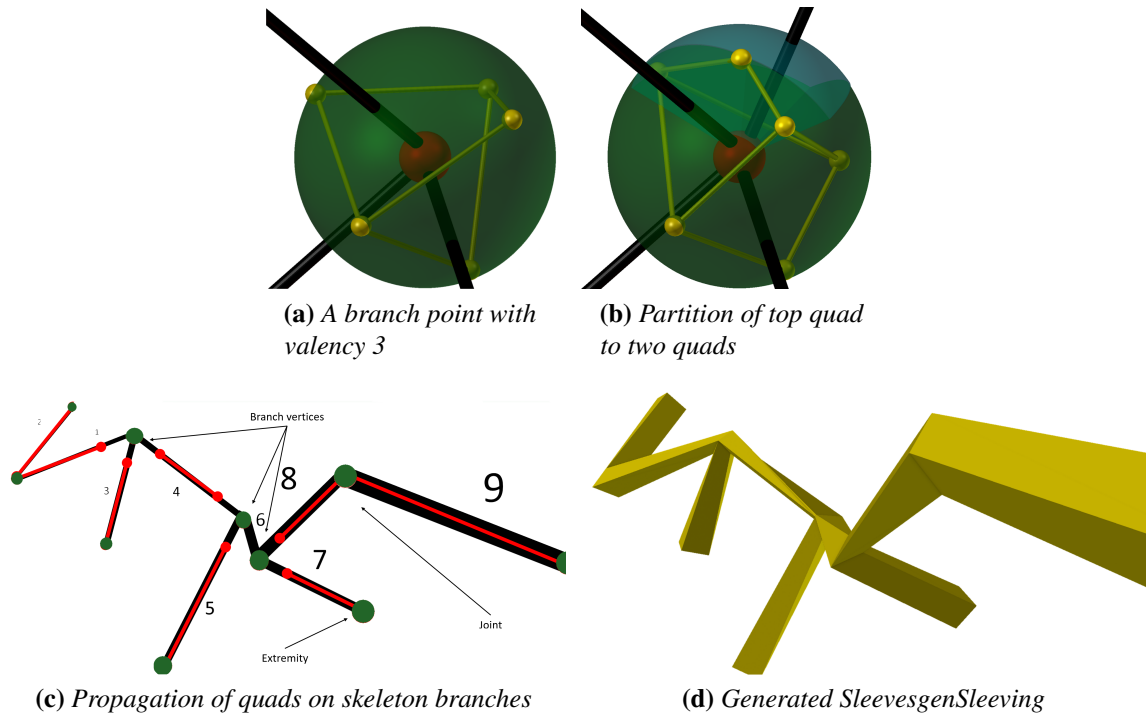


Figure A.5: The two main steps of the algorithm: *Partition (top) and Sleeving (bottom)*.

$v \mapsto p_v$ is the graph embedding, i.e., an assignment of unique positions from \mathbb{R}^3 to each vertex of the graph, and

- $r: V \rightarrow \mathbb{R}^+$,
- $v \mapsto r_v$ with r_v denoting the radius associated to the vertex $v \in V$.

Consequently, any edge of the graph G is embedded as a line segment of \mathbb{R}^3 . Edges are assumed to intersect only at endpoints. We furthermore assume that, outside the spheres centered at the vertices, the line segments are far enough from each other for the scaffold not to intersect. In this project, we also assume that the spheres defined on the embedded graph do not intersect. There does not seem to be a large obstacle to treating this case, as shown in some of the examples.

We distinguish between the so-called joints, branch points, and extremities: The image in \mathbb{R}^3 of a vertex of valency 2 (respectively 1) is called a *joint* (respectively an *extremity*). A *branch point* is the image p_v of a vertex v of valency 3 or more. Together, we call the joints, the extremities, and the branch points *skeleton nodes* (or simply *nodes*, where it is clear from context).

For a given skeleton (V, E, p, r) , a *scaffold* is a closed orientable polyhedral surface enclos-

ing the union of line segments defined by the skeleton, while being homotopically equivalent to it. The scaffold we construct is a quad mesh, the faces of which we call *scaffold quads*. Each scaffold quad is associated to a single skeleton edge and to each skeleton edge are associated exactly 4 scaffold quads (a *sleeve*). Our construction of a *minimal* scaffold for a skeleton consists of two main steps:

Partition: For all branch points p_v , partition the sphere S_v centered at p_v and with radius r_v , according to the issuing branches. In this step, the spheres are partitioned into quadrilateral faces, called *spherical quads* (yellow quads in Figures A.5(a), A.5(b)). Each spherical quad face is intersected by a single line segment joining the branch point to a neighbor vertex.

Sleeve: With a depth first search order, starting from any branch point or extremity, *sleeve* all the branches. In this step, we compute a propagated spherical quad, for every joint (i.e., skeleton node of degree-2) following the DFS order (Figure A.5(c)). We use a heuristic alignment for the joints or extremities connected to a branch point, and the rotation minimization algorithm for all remaining skeleton nodes. We use these quadrilaterals surrounding the nodes of the skeleton to mesh every edge of the skeleton using four quad faces. Finally, we close the mesh at extremities (Figure A.5(d)).

Having explained the technical information regarding the scaffolding of a skeleton, I next explain how this project impacted my decision in using skeletons for the tactile project in this thesis.

Project Take-away Compared to the musical instrument project, this project since it was used to describe 3D shapes presented a stronger insight on why a relationship between a surface model and its skeleton can be one to one. The project strengthened the idea that topological information, and specifically skeletons, is of great importance in graphics research in general, and more particularly for the concise description of 3 dimensional models. This concise graphical description was the goal of the tactile project, since touch, compared to vision, has a much more limited bandwidth to process information on paper.




Figure A.6: Scaffold meshes generated (right) by input skeletons (left).

A.3 Key-Points

To summarize, my previous graduate work in computer graphics, even though not directly applied to my thesis, gave me insight into the fabrication process, and most importantly, influenced my approach to the watercolor and tactile projects. As seen in this chapter, three key takeaways are: the importance of information that a 2D image conveys, the complexity of viewpoint selection for 3D models, and the importance of skeletons in shape analysis.

B

Supplemental Material

B.1 Watercolor Woodblock: Shading Algorithm

We build an associated neighboring graph where nodes represent superpixels and edges indicate pixel boundaries on the image with weights that represent color intensity closeness between superpixels. To build the shadow block corresponding to a color region, we proceed as follows:

1. Compute superpixels, sets of neighboring pixels in the gray-scale map image that have the same gray intensity value. Next, Process every consecutive color region of the color map separately.
2. Initialize the set of shadow superpixels with the locally darkest superpixels, are all the superpixels that all their neighbors have paler or equal intensity values.
3. Build the set of candidate superpixels as the superpixels not selected yet with brightness less than 1.1 of the brightness of the color region.
4. Define a graph over the superpixels as following: Every superpixel represents a node in the graph. There is an edge between any two nodes iff the corresponding superpixels have at least one neighboring pixel. The edges are weighted with the difference of the brightness values between the two superpixels.
5. Expand the shadow superpixels to the closest neighbor superpixel in the graph. Update

the graph by merging the two nodes. Update the edges and their corresponding weights, according to the brightest value of all the superpixels included in the expanded superpixel.

6. For every superpixel if there is no candidate or if the number of pixels of the superpixel is greater or equal to 20 pixels, stop expanding.

B.1.1 Details

To remove small details, we perform a series of morphological operations on the images with a disk of specified radius. The radius r is used to define a disk of total size $2 \times r + 1$ in OpenCV. Given that all our images are of size 800 pixels \times 600 pixels, 1 pixel radius corresponds to approximately 0.25 mm radius.

We smooth the input image with an open followed by a close operation of radius 0.25 mm. For the darkest superpixels (Wet Woodblocks - Shading Woodblocks - Local Minima) to compute the gray-scale map we use k -means for $k = 30$ gray values on the gray-scale input photograph. We then smooth the image with an open followed by a close operation of radius 0.5 mm. We shrink the water area of the water map with an erode operation of radius 0.25 mm and we remove the dry area. To compute the gray-scale map for the details superpixels (Dry Woodblocks - Details minima) and the brightest superpixels (Wet Woodblocks - Shading Woodblocks - Local maxima), we use k -means for $k = 50$ gray values. We smooth the gray-scale map with a close operation of radius 0.25 mm. We process the entire image, differently than algorithm's Step 1. For the brightest superpixels, we consider the local maxima, which have color values greater than 220 (gray channel $\in [0, 255]$), Steps 2 - 3. Also, we do not expand, eliminating Steps 4-6.

B.2 Ring Location Sampling

We restrict the *master* camera up-vector to the up-vector of the up-right origin placed object, $\hat{\mathbf{u}}_M = (0, 0, 1)$, and sample master camera locations *loc* that lie on the origin centered *ring*

$$\forall (x, y, z) \in \mathbb{R}^3 [(x^2 + y^2 = 10 \wedge z = 0) \implies loc(x, y)]$$

B.3 Negative y Location Restriction

After translating the *master* camera at the computed location, found from the maximum silhouette computation Section A.1, we point the camera at the origin [55, 105] using a standard track-to operation. Following we rotate the object and the master camera the same horizontal angle. The angle makes the master camera coincide with the negative y axis and is $r_M = \text{SignedAngle}(loc, \text{axis-}y^-)$, while axis- z is the rotation axis.

B.4 Canonical Coordinate Frame

To compute the canonical coordinate frame for each *part*, we first translate the centre of the *part* to the origin. We then perform eigen-decomposition on the covariance matrix of the vertices, which results in three eigenvectors, the axes of the local canonical coordinate frame.

Given the *part* vertices $v_i, i \in \{1 \dots n\} = I$, we define the covariance matrix $Cov = Trv^T Trv$, where the matrix $Trv \in \mathbb{R}^{n \times 3}$, with $Trv_{i,*} = t_i, t_i \in \mathbb{R}^3, t_i = v_i - c_p$, and $c_p = \frac{(\sum_1^n v_i)}{n}$. Next, using a standard package [370] we find the eigenvectors of the square matrix C , named $\vec{e}v_1, \vec{e}v_2, \vec{e}v_3 \in \mathbb{R}^3$, which constitute the three axes of the canonical frame for the corresponding *part*.

B.5 Canonical Bounding Box

To compute the bounding box for the origin centered *part*, with vertices T , aligned with the canonical frame $\vec{e}v_1, \vec{e}v_2, \vec{e}v_3$, as defined in Appendix B.4, we compute the *box* vertices.

The eight *box* vertices are

The eight *box* vertices are

$$box := \left\{ \min_R(x), \max_R(x) \right\} \times \left\{ \min_R(y), \max_R(y) \right\} \times \left\{ \min_R(z), \max_R(z) \right\}$$

with $R := \{ (x_i, y_i, z_i) \mid (x_i, y_i, z_i) = Rot * t_i, t_i \in rows(Trv), i \in I \}$ being the set of rotated vertices, where $Rot = \begin{pmatrix} \vec{e}v_1 & \vec{e}v_2 & \vec{e}v_3 \end{pmatrix}^T \in \mathbb{R}^{3 \times 3}$ is the rotation matrix, $rows()$ the set of

row vectors of the given matrix, and $\max_R, \min_R : \{x, y, z\} \mapsto \mathbb{R}$, functions that return the largest and the smallest coordinates from the set of vertices R , e.g., $\max_R(x) = \max\{x_i \mid \forall i \in I, (x_i, y_i, z_i) \in R\}$.

B.6 Part Viewing Direction

To select which *box* side to use as the *part* camera viewing direction, we take into account the skeleton, the area of the side, and *master* camera placement.

For a given *part* centered in the origin, let $B := \{b_1, b_2, b_3\}$ be the set of the three perpendicular *box* sides with normals $\{\vec{e}v_1, \vec{e}v_2, \vec{e}v_3\}$, and $S := \{\vec{s}_i \mid i \in \{1 \dots k\}\}$ the directions of the skeletal segments. We choose two of the sides that align well with the skeleton, by discarding the less orthogonal to the skeleton side $b_d \in B$:

$$d = \operatorname{argmax}_{i \in \{1,2,3\}} \frac{\sum_{s \in S} \min(|\operatorname{angle}(\vec{s}, \vec{e}v_i)|, |\operatorname{angle}(-\vec{s}, \vec{e}v_i)|)}{k}$$

where $\operatorname{angle}(\vec{v}_1, \vec{v}_2)$ returns the angle between two vectors. Next, the normal of the selected box side, ev_s , is restricted so that the *part* camera is placed on the negative y half-space, resulting in the *part* camera viewing direction $\hat{\mathbf{d}}_P$. The index is computed by $s = \operatorname{argmax}_{i \in \{1,2,3\} - \{d\}} \operatorname{area}(b_i)$, where area returns the surface area of the given rectangle. To implement the correct camera direction, we rotate the camera and make its viewing direction coincide with the $\hat{\mathbf{d}}_P$ using a standard track-to operation.

B.7 Up-Vector Initialization

The part camera location is a point on the surface of the sphere centered at the corresponding part, with radii direction from the point to the sphere center $\hat{\mathbf{d}}_P$. The corresponding up-vector will then be a tangent vector on the surface at this point. The up-vectors are defined for all points on the negative y hemisphere, should change smoothly from point to point, and should follow the up-right orientation of the object. In other words, these vectors constitute a vector field on the sphere. For our initial guess for the up-vector field, we use the vector field provided by Lopes

et al. [13], Figure B.1. Specifically, the part camera up-vector is $\vec{u}_P = -(2 * \hat{\mathbf{h}}.y * \hat{\mathbf{h}}.z), -(-2 * \hat{\mathbf{h}}.x * \hat{\mathbf{h}}.z), 1 - (2 * \hat{\mathbf{h}}.z * \hat{\mathbf{h}}.z)$, where $\vec{h} = \left(\max\{\hat{\mathbf{d}}_P.y - 1, \hat{\mathbf{d}}_P.y + 1\}, \hat{\mathbf{d}}_P.x, \hat{\mathbf{d}}_P.z \right)$.




Figure B.1: The vector field, green arrows, we use to initialize the up-vector for the part cameras. It gives reasonable initial renderings.

To make the current up-vector $\hat{\mathbf{u}}$ of the origin centered *part* camera coincide with the found $\hat{\mathbf{u}}_P$, we rotate the *part* camera around the viewing direction $\hat{\mathbf{d}}_P$.

We translate the *part* from the origin back to its original position v_i , with $v_i = t_i + c_p$ (Appendix B.4), and translate the camera at the center of the part c_p . We move the *part* camera away from the *part*, distance equal to $dst = \|p_M - c_p\|$ so that the position of the *part* camera is $c_p - dst * \hat{\mathbf{d}}_p$.

Bibliography

- [1] Axes. <https://www.thoughtco.com/mousterian-definition-167233>. Accessed: 2020-02-27. Citation on page 1.
- [2] A brief history of writing. <https://www.thoughtco.com/brief-history-of-writing-4072560>. Accessed: 2020-02-27. Citation on page 1.
- [3] Craft specialization. <https://www.thoughtco.com/craft-specialization-167073>. Accessed: 2020-02-27. Citation on page 1.
- [4] DenseCRF implementation. <https://github.com/cvlab-epfl/densecrf>. Accessed: 2017-05-21. Citation on page 46.
- [5] Hand tools. <https://www.britannica.com/technology/hand-tool>. Accessed: 2020-02-27. Citation on page 1.
- [6] The history of jewelry design. <https://www.britannica.com/art/jewelry/Egyptian>. Accessed: 2020-02-27. Citation on page 1.
- [7] Insight segmentation and registration toolkit. <https://itk.org/>. Accessed: 2017-05-21. Citation on page 34.
- [8] The invention of the wheel. <https://www.thoughtco.com/the-invention-of-the-wheel-1992669>. Accessed: 2020-02-26. Citation on page 1.
- [9] Machine. <https://www.britannica.com/technology/machine-tool/Modifications-of-basic-machines>. Accessed: 2020-02-27. Citation on page 1.
- [10] Musical instrument. <https://www.britannica.com/art/musical-instrument>. Accessed: 2020-02-27. Citation on page 1.

- [11] Spindle. <https://www.thoughtco.com/spindle-whorls-ancient-tool-for-weavers-172908>. Accessed: 2020-02-26. Citation on page 1.
- [12] Thistory of publishing. <https://www.britannica.com/topic/publishing/The-age-of-early-printing-1450-1550>. Accessed: 2020-02-28. Citation on page 1.
- [13] Tangent vectors to a 3-d surface normal: A geometric tool to find orthogonal vectors based on the householder transformation. *Computer-Aided Design*, 45(3):683 – 694, 2013. Citation on pages 73, 85, and 131.
- [14] Gianfranco Ferré Foundation, Milan, IT. Gianfranco ferré designs. http://www.fondazionegianfrancoferre.com/home/collezioni_dhtc.php?lang=en&t=disegni&sta=ai&y=1988, 2015. Accessed: 2020-02-17. Citation on page 21.
- [15] Metro exodus: Geforce rtx real-time ray traced global illumination, 2018. Online at <https://youtu.be/Ms7d-3Dprio>. Citation on page 3.
- [16] Alfa Romeo Museum. Masterpieces from the alfa romeo museum. https://www.museoalfaromeo.com/en-us/news/Pages/23012019_MasterpiecesFromAlfaRomeoMuseum.aspx, 2019. Accessed: 2020-02-18. Citation on page 21.
- [17] National Gallery of Art, Washington, DC. M.c. escher — life and work. https://www.nga.gov/features/slideshows/mc-escher-life-and-work.html?view=grid#slide_2, 2019. Accessed: 2020-02-17. Citation on page 20.
- [18] Hergé Museum. The adventures of tintin. <http://en.tintin.com/albums>, 2020. Accessed: 2020-02-17. Citation on page 20.
- [19] Van Gogh Museum, Amsterdam, NL. Van gogh’s life, 1853–1890. <https://www.vangoghmuseum.nl/en/vincent-van-gogh-life-and-work/van-goghs-life-1853-1890>, 2020. Accessed: 2020-02-17. Citation on page 20.

- [20] Adobe, Substance. Tools for creating 3d textures and materials. <https://www.substance3d.com/>, 2020. Accessed: 2020-03-04. Citation on page 8.
- [21] H. Agrawal, U. Umapathi, R. Kovacs, F. Johannes, T. Chen, S. Mueller, and P. Baudisch. Protopiper: Physically sketching room-sized objects at actual scale. 11 2015. DOI [10.1145/2807442.2807505](https://doi.org/10.1145/2807442.2807505). Citation on page 15.
- [22] M. Agrawala, D. Phan, J. Heiser, J. Haymaker, J. Klingner, P. Hanrahan, and B. Tversky. Designing effective step-by-step assembly instructions. *ACM Trans. Graph.*, 22:828–837, 07 2003. Citation on page 107.
- [23] M. Agrawala, D. Zorin, and T. Munzner. Artistic multiprojection rendering. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*, pages 125–136, London, UK, UK, 2000. Springer-Verlag. Citation on pages 28, 73, and 81.
- [24] E. Aharoni-Mack, Y. Shambik, and D. Lischinski. Pigment-based recoloring of watercolor paintings. In *Proceedings of the Symposium on Non-Photorealistic Animation and Rendering*, NPAR '17, New York, NY, USA, 2017. Association for Computing Machinery. DOI [10.1145/3092919.3092926](https://doi.org/10.1145/3092919.3092926). Citation on page 33.
- [25] Y. Aksoy, T.-H. Oh, S. Paris, M. Pollefeys, and W. Matusik. Semantic soft segmentation. *ACM Trans. Graph.*, 37(4), July 2018. DOI [10.1145/3197517.3201275](https://doi.org/10.1145/3197517.3201275). Citation on pages 36 and 104.
- [26] T. Alderighi, L. Malomo, D. Giorgi, N. Pietroni, B. Bickel, and P. Cignoni. Metamolds: Computational design of silicone molds. *ACM Trans. Graph.*, 37(4):136:1–136:13, July 2018. DOI [10.1145/3197517.3201381](https://doi.org/10.1145/3197517.3201381). Citation on page 15.
- [27] M. Alexa and W. Matusik. Reliefs as images. *ACM Trans. Graph.*, 29(4), July 2010. DOI [10.1145/1778765.1778797](https://doi.org/10.1145/1778765.1778797). Citation on page 108.
- [28] M. Alexa and W. Matusik. Images from self-occlusion. In *Proceedings of the International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*,

CAe '11, page 17–24, New York, NY, USA, 2011. Association for Computing Machinery. DOI [10.1145/2030441.2030445](https://doi.org/10.1145/2030441.2030445). Citation on page 25.

- [29] P. V. Arman. Cloudpainter. <http://www.cloudpainter.com/blog?category=Artistic%20Robots>, 2020. Accessed: 2020-03-05. Citation on page 4.
- [30] S. Arpa, T. Ritschel, K. Myszkowski, T. Çapın, and H.-P. Seidel. Purkinje images: Conveying different content for different luminance adaptations in a single image. *Computer Graphics Forum*, 34(1):116–126, 2015. DOI [10.1111/cgf.12463](https://doi.org/10.1111/cgf.12463). Citation on page 25.
- [31] F. Attneave. Some informational aspects of visual perception. *Psychological Review*, 61(3):183–193, 1954. DOI [10.1037/h0054663](https://doi.org/10.1037/h0054663). Citation on page 35.
- [32] V. Babaei, K. Vidimče, M. Foshey, A. Kaspar, P. Didyk, and W. Matusik. Color contoning for 3d printing. *ACM Trans. Graph.*, 36(4), July 2017. DOI [10.1145/3072959.3073605](https://doi.org/10.1145/3072959.3073605). Citation on pages 15 and 26.
- [33] V. Babaei, K. Vidimče, M. Foshey, A. Kaspar, P. Didyk, and W. Matusik. Color contoning for 3D printing. *ACM Trans. Graph.*, 36(4):124:1–124:15, July 2017. DOI [10.1145/3072959.3073605](https://doi.org/10.1145/3072959.3073605). Citation on page 105.
- [34] M. Bäcker, B. Bickel, D. L. James, and H. Pfister. Fabricating articulated characters from skinned meshes. *ACM Trans. Graph.*, 31(4), July 2012. DOI [10.1145/2185520.2185543](https://doi.org/10.1145/2185520.2185543). Citation on page 3.
- [35] D. Ballard and C. Brown. *Computer Vision*. Prentice Hall, Upper Saddle River, New Jersey 07458, USA, 1982. Citation on page 35.
- [36] I. Baran, P. Keller, D. Bradley, S. Coros, W. Jarosz, D. Nowrouzezahrai, and M. H. Gross. Manufacturing layered attenuators for multiple prescribed shadow images. *Comput. Graph. Forum*, 31:603–610, 2012. Citation on pages 8 and 24.
- [37] G. Barill, N. G. Dickson, R. Schmidt, D. I. W. Levin, and A. Jacobson. Fast winding numbers for soups and clouds. *ACM Trans. Graph.*, 37(4), July 2018. DOI [10.1145/3197517.3201337](https://doi.org/10.1145/3197517.3201337). Citation on page 2.

- [38] D. Bartz, D. Cunningham, J. Fischer, and C. Wallraven. The Role of Perception for Computer Graphics. In T. Theoharis and P. Dutre, editors, *Eurographics 2008 - State of the Art Reports*. The Eurographics Association, 2008. DOI [10.2312/egst.20081045](https://doi.org/10.2312/egst.20081045). Citation on page 69.
- [39] G. Bathgate. Openscad user manual/3d to 2d projection, 2010. Online at <https://gilesbathgate.com/2010/06/23/extracting-2d-mendel-outlines-using-openscad/>. Citation on page 20.
- [40] O. Bau and I. Poupyrev. Revel: Tactile feedback technology for augmented reality. *ACM Trans. Graph.*, 31(4), July 2012. DOI [10.1145/2185520.2185585](https://doi.org/10.1145/2185520.2185585). Citation on page 69.
- [41] P. Baudisch, T. Roumen, A. Silber, Y. Kommana, M. Gruner, L. Wall, K. Reuss, L. Heilman, R. Kovacs, and D. Rechlitz. Kyub: A 3d editor for modeling sturdy laser-cut objects. pages 1–12, 05 2019. DOI [10.1145/3290605.3300796](https://doi.org/10.1145/3290605.3300796). Citation on page 15.
- [42] C. Bauer, L. Yazzolino, G. Hirsch, Z. Cattaneo, T. Vecchi, and L. Merabet. Neural correlates associated with superior tactile symmetry perception in the early blind. *Cortex*, 63:104–117, 02 2015. Citation on pages 100 and 107.
- [43] B. Baxter, V. Scheib, M. C. Lin, and D. Manocha. Dab: Interactive haptic painting with 3d virtual brushes. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '01, page 461–468, New York, NY, USA, 2001. Association for Computing Machinery. DOI [10.1145/383259.383313](https://doi.org/10.1145/383259.383313). Citation on page 24.
- [44] W. Baxter, J. Wendt, and M. C. Lin. Impasto: A realistic, interactive model for paint. In *Proceedings of the 3rd International Symposium on Non-Photorealistic Animation and Rendering*, NPAR '04, page 45–148, New York, NY, USA, 2004. Association for Computing Machinery. DOI [10.1145/987657.987665](https://doi.org/10.1145/987657.987665). Citation on page 24.
- [45] P. Bénard and A. Hertzmann. Line drawings from 3d models. *CoRR*, abs/1810.01175, 2018. Online at <http://arxiv.org/abs/1810.01175>. Citation on page 27.

- [46] I. Berger, A. Shamir, M. Mahler, E. Carter, and J. Hodgins. Style and abstraction in portrait sketching. *ACM Trans. Graph.*, 32(4), July 2013. DOI [10.1145/2461912.2461964](https://doi.org/10.1145/2461912.2461964). Citation on pages 3 and 23.
- [47] A. Bermano, I. Baran, M. Alexa, and W. Matusk. Shadowpix: Multiple images from self shadowing. *Comput. Graph. Forum*, 31(2pt3):593–602, May 2012. DOI [10.1111/j.1467-8659.2012.03038.x](https://doi.org/10.1111/j.1467-8659.2012.03038.x). Citation on page 24.
- [48] A. H. Bermano, T. Funkhouser, and S. Rusinkiewicz. State of the art in methods and representations for fabrication-aware design. *Computer Graphics Forum*, 36(2):509–535, 2017. DOI [10.1111/cgf.13146](https://doi.org/10.1111/cgf.13146). Citation on page 2.
- [49] J. M. Bern, G. Kumagai, and S. Coros. Fabrication, modeling, and control of plush robots. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3739–3746, Sep. 2017. DOI [10.1109/IROS.2017.8206223](https://doi.org/10.1109/IROS.2017.8206223). Citation on page 15.
- [50] B. Bickel, P. Cignoni, L. Malomo, and N. Pietroni. State of the art on stylized fabrication. *Computer Graphics Forum*, 37(6):325–342, 2018. DOI [10.1111/cgf.13327](https://doi.org/10.1111/cgf.13327). Citation on page 14.
- [51] I. Biederman. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2):115–147, 1987. Citation on pages 70 and 76.
- [52] M. Birsak, F. Rist, P. Wonka, and P. Musialski. String art: Towards computational fabrication of string images. *Computer Graphics Forum*, 37, 04 2018. DOI [10.1111/cgf.13359](https://doi.org/10.1111/cgf.13359). Citation on pages 5 and 25.
- [53] R. Blake and R. Sekuler. *Perception*. McGraw-Hill, Boston, MA, USA, 2006. Citation on pages 28, 33, 34, 70, 76, and 84.
- [54] V. Blanz, M. J. Tarr, and H. H. Bülthoff. What object attributes determine canonical views? *Perception*, 28(5):575–599, 1999. Citation on page 72.

- [55] blender.org. Blender 2.78 - api documentation: bpy.types.tracktoconstraint, 2019. Online at https://docs.blender.org/api/blender_python_api_current/bpy.types.TrackToConstraint.html#bpy.types.TrackToConstraint. Citation on page 129.
- [56] Z. Bo, L. Lu, A. Sharf, Y. Xia, O. Deussen, and B. Chen. Printable 3d trees. *Computer Graphics Forum*, 36(7):29–40, 2017. DOI [10.1111/cgf.13269](https://doi.org/10.1111/cgf.13269). Citation on page 15.
- [57] X. Bonaventura, M. Feixas, M. Sbert, L. Chuang, and C. Wallraven. A survey of viewpoint selection methods for polygonal models. *Entropy*, 20:370, 05 2018. DOI [10.3390/e20050370](https://doi.org/10.3390/e20050370). Citation on page 72.
- [58] R. Bosch. Opt art. *Math Horizons*, 13(3):6–9, 2006. Citation on page 22.
- [59] I. Botkin. Painting with polygons: Non-photorealistic rendering using existing tools. In *SIGGRAPH 2009: Talks*, SIGGRAPH '09, pages 22:1–22:1, New York, NY, USA, 2009. ACM. DOI [10.1145/1597990.1598012](https://doi.org/10.1145/1597990.1598012). Citation on page 22.
- [60] A. Bousseau, M. Kaplan, J. Thollot, and F. X. Sillion. Interactive watercolor rendering with temporal coherence and abstraction. In *International Symposium on Non-Photorealistic Animation and Rendering (NPAR)*, Annecy, France, 2006. ACM. Online at <https://hal.inria.fr/inria-00510223>. Citation on pages 23 and 34.
- [61] A. Bousseau, F. Neyret, J. Thollot, and D. Salesin. Video watercolorization using bidirectional texture advection. In *ACM SIGGRAPH 2007 Papers*, SIGGRAPH '07, New York, NY, USA, 2007. ACM. DOI [10.1145/1275808.1276507](https://doi.org/10.1145/1275808.1276507). Citation on pages 22 and 23.
- [62] A. Bousseau, J. P. O'shea, F. Durand, R. Ramamoorthi, and M. Agrawala. Gloss perception in painterly and cartoon rendering. *ACM Trans. Graph.*, 32(2), Apr. 2013. DOI [10.1145/2451236.2451244](https://doi.org/10.1145/2451236.2451244). Citation on page 69.
- [63] D. Brauner. Refreshable tactile graphics display: Making on-screen images accessible!, 2016. Accessed: 2019-08-17, Online at <https://www.perkinselearning.org/technology/>

posts/refreshable-tactile-graphics-display-making-screen-images-accessible. Citation on pages 100 and 107.

- [64] J. A. Bærentzen, M. K. Misztal, and K. Welnicka. Converting skeletal structures to quad dominant meshes. *Computers & Graphics*, 36(5):555 – 561, 2012. DOI <https://doi.org/10.1016/j.cag.2012.03.016>. Citation on page 123.
- [65] S. Brewster and L. M. Brown. Tactons: structured tactile messages for non-visual information display. In *Proceedings of the fifth conference on Australasian user interface-Volume 28*, pages 15–23. Australian Computer Society, Inc., 2004. Citation on page 28.
- [66] S. Brewster, A. M. Burton, and G. Dimigen. Multivis: A multimodal visualisation system for blind students using virtual reality. Citation on page 28.
- [67] A. M. Brock, P. Truillet, B. Oriola, D. Picard, and C. Jouffrais. Interactivity improves usability of geographic maps for visually impaired people. *Human–Computer Interaction*, 30(2):156–194, 2015. Citation on pages 28 and 66.
- [68] J. Brooke. Sus-a quick and dirty usability scale. 194:189–194, 1996. Citation on page 97.
- [69] J. Brosz, F. F. Samavati, M. S. T. Carpendale, and M. C. Sousa. Single camera flexible projection. In *Proceedings of the 5th International Symposium on Non-photorealistic Animation and Rendering*, NPAR '07, pages 33–42, New York, NY, USA, 2007. ACM. Citation on pages 22 and 28.
- [70] C. Brown and A. Hurst. Viztouch: Automatically generated tactile visualizations of coordinate spaces. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*, TEI '12, pages 131–138, New York, NY, USA, 2012. ACM. Citation on pages 29 and 66.

- [71] A. Brunton, C. A. Arikan, and P. Urban. Pushing the limits of 3d color printing. *ACM Transactions on Graphics*, 35(1):1–13, Dec 2015. DOI [10.1145/2832905](https://doi.org/10.1145/2832905). Citation on page 15.
- [72] D. Bull, Seseragi Studio, Ome Tokyo. Encyclopedia of woodblock printmaking. http://woodblock.com/encyclopedia/topics/toc_process.html. Accessed: 2017-09-18. Citation on page 18.
- [73] F. Buonomici, R. Furferi, L. Governi, and Y. Volpe. Making blind people autonomous in the exploration of tactile models: A feasibility study. In *International Conference on Universal Access in Human-Computer Interaction*, pages 82–93. Springer, 2015. Citation on page 108.
- [74] D. Burch. Development of a multiple contact haptic display with texture-enhanced graphics. 2012. Citation on page 72.
- [75] S. Calinon, J. Epiney, and A. Billard. A humanoid robot drawing human portraits. *5th IEEE-RAS International Conference on Humanoid Robots, 2005.*, pages 161–166, 2005. Citation on page 25.
- [76] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, Nov 1986. DOI [10.1109/TPAMI.1986.4767851](https://doi.org/10.1109/TPAMI.1986.4767851). Citation on pages 35 and 47.
- [77] K. Cao, J. Liao, and L. Yuan. Carigans: Unpaired photo-to-caricature translation. *ACM Trans. Graph.*, 37(6), Dec. 2018. DOI [10.1145/3272127.3275046](https://doi.org/10.1145/3272127.3275046). Citation on page 109.
- [78] M. Carfagni, R. Furferi, L. Governi, Y. Volpe, and G. Tennirelli. Tactile representation of paintings: An early assessment of possible computer based strategies. In *Progress in Cultural Heritage Preservation*, pages 261–270, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. Citation on page 29.
- [79] I. Carlbom and J. Paciorek. Planar geometric projections and viewing transformations. *ACM Comput. Surv.*, 10(4):465–502, Dec. 1978. Citation on page 22.

- [80] G. F. Carrillo. Empowering blind creators. *Human-Technology*, 2018. Citation on page 20.
- [81] H. Chang, O. Fried, Y. Liu, S. DiVerdi, and A. Finkelstein. Palette-based photo recoloring. *ACM Trans. Graph.*, 34(4):139:1–139:11, July 2015. DOI [10.1145/2766978](https://doi.org/10.1145/2766978). Citation on pages 33, 39, 43, 44, and 46.
- [82] M. Chaouch and A. Verroust-Blondet. Alignment of 3d models. *Graph. Models*, 71(2):63–76, Mar. 2009. DOI [10.1016/j.gmod.2008.12.006](https://doi.org/10.1016/j.gmod.2008.12.006). Citation on page 72.
- [83] B.-Y. Chen. Riemannian submanifolds. *Handbook of differential geometry*, 1:360, 2000. Citation on page 116.
- [84] D. Chen, D. I. W. Levin, P. Didyk, P. Sitthi-Amorn, and W. Matusik. Spec2fab: A reducer-tuner model for translating specifications to 3d prints. *ACM Trans. Graph.*, 32(4), July 2013. DOI [10.1145/2461912.2461994](https://doi.org/10.1145/2461912.2461994). Citation on page 15.
- [85] Q. Chen, D. Li, and C. K. Tang. KNN matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(9):2175–2188, Sept 2013. DOI [10.1109/TPAMI.2013.18](https://doi.org/10.1109/TPAMI.2013.18). Citation on page 47.
- [86] Z. Chen, B. Kim, D. Ito, and H. Wang. Wetbrush: GPU-based 3D painting simulation at the bristle level. *ACM Trans. Graph.*, 34(6):200:1–200:11, Oct. 2015. DOI [10.1145/2816795.2818066](https://doi.org/10.1145/2816795.2818066). Citation on page 24.
- [87] F. D. K. Ching and J. S. P. *Design Drawing*. WILEY, Indianapolis, IN 46256, 2019. Citation on page 27.
- [88] H.-K. Chu, W.-H. Hsu, N. J. Mitra, D. Cohen-Or, T.-T. Wong, and T.-Y. Lee. Camouflage images. In *ACM SIGGRAPH 2010 Papers*, SIGGRAPH '10, New York, NY, USA, 2010. Association for Computing Machinery. DOI [10.1145/1833349.1778788](https://doi.org/10.1145/1833349.1778788). Citation on page 22.

- [89] N. S.-H. Chu and C.-L. Tai. Moxi: Real-time ink dispersion in absorbent paper. *ACM Trans. Graph.*, 24(3):504–511, July 2005. DOI [10.1145/1073204.1073221](https://doi.org/10.1145/1073204.1073221). Citation on page 24.
- [90] P. Cignoni, C. Montani, and R. Scopigno. Computer-assisted generation of bas- and high-reliefs. *J. Graph. Tools*, 2(3):15–28, Dec. 1997. DOI [10.1080/10867651.1997.10487476](https://doi.org/10.1080/10867651.1997.10487476). Citation on page 108.
- [91] P. Cignoni, N. Pietroni, L. Malomo, and R. Scopigno. Field-aligned mesh joinery. *ACM Trans. Graph.*, 33(1), Feb. 2014. DOI [10.1145/2537852](https://doi.org/10.1145/2537852). Citation on page 15.
- [92] P. Cignoni, R. Scopigno, and M. Tarini. A simple normal enhancement technique for interactive non-photorealistic renderings, 2005. Citation on page 22.
- [93] P. Claudet, G. Chougui, and P. Richard. The typhlo and tactus guide for children’s books with tactile illustrations, 2000-2008. Citation on page 66.
- [94] M. T. Cockshott. Wet and sticky : a novel model for computer-based painting. 1991. Citation on page 24.
- [95] H. Cohen. What is an image? In *Proceedings of the 6th International Joint Conference on Artificial Intelligence, Volume 2, IJCAI’79*, page 1028–1057, San Francisco, CA, USA, 1979. Morgan Kaufmann Publishers Inc. Citation on page 5.
- [96] H. Cohen. How to draw three people in a botanical garden. In *AAAI*, 1988. Citation on page 5.
- [97] H. Cohen. The further exploits of aaron, painter. *Stanford Hum. Rev.*, 4(2):141–158, jul 1995. Citation on page 5.
- [98] H. Cohen. Colouring without seeing: a problem in machine creativity. *AISB quarterly*, 102:26–35, 1999. Citation on page 5.

- [99] H. Cohen. A self-defining game for one player: On the nature of creativity and the possibility of creative computer programs. *Leonardo*, 35(1):59–64, 2002. Online at <https://doi.org/10.1162/002409402753689335>. Citation on page 5.
- [100] H. Cohen, UCSD. Aaron, colorist: from expert system to expert. <http://www.aaronshome.com/aaron/publications/index.html>, 2006. Accessed: 2020-03-05. Citation on pages 5 and 26.
- [101] F. Cole, K. Sanik, D. DeCarlo, A. Finkelstein, T. Funkhouser, S. Rusinkiewicz, and M. Singh. How well do line drawings depict shape? *ACM Trans. Graph.*, 28(3):28:1–28:9, 2009. Citation on pages 27 and 28.
- [102] J. P. Collomosse and J. E. Kyprianidis. Artistic stylisation of images and video. In *Eurographics (Tutorials)*, 2011. Citation on page 22.
- [103] T. Cooke, F. Jäkel, C. Wallraven, and H. H. Bühlhoff. Multimodal similarity and categorization of novel, three-dimensional objects. *Neuropsychologia*, 45(3):484 – 495, 2007. Advances in Multisensory Processes. Citation on page 70.
- [104] C. J. Curtis, S. E. Anderson, J. E. Seims, K. W. Fleischer, and D. H. Salesin. Computer-generated watercolor. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97*, pages 421–430, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co. DOI 10.1145/258734.258896. Citation on page 22.
- [105] J. de Vries. Learn opengl, 2019. Online at <https://learnopengl.com/Getting-started/Camera>. Citation on page 129.
- [106] D. DeCarlo, A. Finkelstein, S. Rusinkiewicz, and A. Santella. Suggestive contours for conveying shape. *ACM Trans. Graph.*, 22(3):848–855, July 2003. Citation on page 28.
- [107] D. DeCarlo and S. Rusinkiewicz. Highlight lines for conveying shape. In *Proceedings of the 5th International Symposium on Non-Photorealistic Animation and Rendering*, NPAR

- '07, page 63–70, New York, NY, USA, 2007. Association for Computing Machinery. DOI [10.1145/1274871.1274881](https://doi.org/10.1145/1274871.1274881). Citation on pages 22 and 71.
- [108] D. DeCarlo and A. Santella. Stylization and abstraction of photographs. *ACM Trans. Graph.*, 21(3):769–776, July 2002. DOI [10.1145/566654.566650](https://doi.org/10.1145/566654.566650). Citation on pages 22, 33, and 35.
- [109] J. D. Denning, W. B. Kerr, and F. Pellacini. Meshflow: Interactive visualization of mesh construction sequences. *ACM Trans. Graph.*, 30(4), July 2011. DOI [10.1145/2010324.1964961](https://doi.org/10.1145/2010324.1964961). Citation on page 107.
- [110] O. Deussen, J. Hamel, A. Raab, S. Schlechtweg, and T. Strothotte. An illustration technique using hardware-based intersections and skeletons. In *Proceedings of the 1999 Conference on Graphics Interface '99*, pages 175–182, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc. Citation on pages 70, 71, and 83.
- [111] O. Deussen, S. Hiller, C. Van Overveld, and T. Strothotte. Floating points: A method for computing stipple drawings. *Computer Graphics Forum*, 19(3):41–50, 2000. DOI [10.1111/1467-8659.00396](https://doi.org/10.1111/1467-8659.00396). Citation on page 35.
- [112] O. Deussen, T. Lindemeier, S. Pirk, and M. Tautzenberger. Feedback-guided stroke placement for a painting machine. In *Proceedings of the Eighth Annual Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*, CAe '12, pages 25–33, Aire-la-Ville, Switzerland, Switzerland, 2012. Eurographics Association. Online at <http://dl.acm.org/citation.cfm?id=2328888.2328894>. Citation on page 25.
- [113] O. Deussen and T. Strothotte. Computer-generated pen-and-ink illustration of trees. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, page 13–18, USA, 2000. ACM Press/Addison-Wesley Publishing Co. DOI [10.1145/344779.344792](https://doi.org/10.1145/344779.344792). Citation on page 22.
- [114] J. Diepstraten, D. Weiskopf, and T. Ertl. Transparency in interactive technical illustrations. *Computer Graphics Forum*, 21:317 – 325, 05 2003. Citation on page 28.

- [115] H. Q. Dinh, F. Gelman, S. Lefebvre, and F. Claux. Modeling and toolpath generation for consumer-level 3d printing. In *ACM SIGGRAPH 2015 Courses*, SIGGRAPH '15, pages 17:1–17:273, 2015. Citation on page 15.
- [116] C. Doersch, S. Singh, A. Gupta, J. Sivic, and A. A. Efros. What makes paris look like paris? *ACM Trans. Graph.*, 31(4), July 2012. DOI [10.1145/2185520.2185597](https://doi.org/10.1145/2185520.2185597). Citation on page 106.
- [117] C. Dorai and A. K. Jain. Cosmos-a representation scheme for 3d free-form objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(10):1115–1130, 1997. Citation on page 71.
- [118] P. J. Doran and J. Hughes. Expressive rendering with watercolor, 2010. Citation on page 24.
- [119] P. Edman. *Tactile graphics*. American Foundation for the Blind, 1992. Citation on pages 11 and 19.
- [120] M. Eitz, J. Hays, and M. Alexa. How do humans sketch objects? *ACM Trans. Graph.*, 31(4), July 2012. DOI [10.1145/2185520.2185540](https://doi.org/10.1145/2185520.2185540). Citation on page 69.
- [121] M. A. El-Sayed, Y. A. Estaitia, and M. A. Khafagy. Automated edge detection using convolutional neural network. *International Journal of Advanced Computer Science and Applications*, 4(10), 2013. DOI [10.14569/IJACSA.2013.041003](https://doi.org/10.14569/IJACSA.2013.041003). Citation on page 35.
- [122] O. Elek, D. Sumin, R. Zhang, T. Weyrich, K. Myszkowski, B. Bickel, A. Wilkie, and J. Křivánek. Scattering-aware texture reproduction for 3d printing. *ACM Trans. Graph.*, 36(6), Nov. 2017. DOI [10.1145/3130800.3130890](https://doi.org/10.1145/3130800.3130890). Citation on page 26.
- [123] N. Elmqvist and P. Tsigas. A taxonomy of 3d occlusion management for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1095–1109, 2008. Citation on page 73.
- [124] N. Elmqvist and M. E. Tudoreanu. Evaluating the effectiveness of occlusion reduction techniques for 3d virtual environments. In *Proceedings of the ACM Symposium on Virtual*

- Reality Software and Technology*, VRST '06, page 9–18, New York, NY, USA, 2006. Association for Computing Machinery. DOI [10.1145/1180495.1180500](https://doi.org/10.1145/1180495.1180500). Citation on page 107.
- [125] G. Erdogan, I. Yildirim, and R. A. Jacobs. Transfer of object shape knowledge across visual and haptic modalities. In *In Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 2014. Citation on page 70.
- [126] J. Farley Norman, H. F Norman, A. Marie Clayton, J. Lianekhammy, and G. Zielke. *The visual and haptic perception of natural object shape*, volume 66. Springer Berlin Heidelberg, Berlin, Heidelberg, 03 2004. Citation on pages 29, 70, and 71.
- [127] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graph.*, 25(3):787–794, July 2006. DOI [10.1145/1141911.1141956](https://doi.org/10.1145/1141911.1141956). Citation on page 34.
- [128] T. J. Ferro and D. T. Pawluk. Automatic image conversion to tactile graphic. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–2, 2013. Citation on page 29.
- [129] F. D. Fiore, W. V. Haevre, and F. V. Reeth. Rendering artistic and believable trees for cartoon animation. In *Computer Graphics International*, 2003. Citation on pages 22 and 70.
- [130] M. Fisher, P. Schröder, M. Desbrun, and H. Hoppe. Design of tangent vector fields. *ACM Trans. Graph.*, 26(3):56–es, July 2007. DOI [10.1145/1276377.1276447](https://doi.org/10.1145/1276377.1276447). Citation on page 104.
- [131] J. Fišer, O. Jamriška, M. Lukáč, E. Shechtman, P. Asente, J. Lu, and D. Sýkora. StyliT: Illumination-guided example-based stylization of 3d renderings. *ACM Trans. Graph.*, 35(4), July 2016. DOI [10.1145/2897824.2925948](https://doi.org/10.1145/2897824.2925948). Citation on page 23.
- [132] M. Flagg and J. M. Rehg. Projector-guided painting. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, UIST '06, pages 235–244,

- New York, NY, USA, 2006. ACM. DOI [10.1145/1166253.1166290](https://doi.org/10.1145/1166253.1166290). Citation on pages 5, 25, and 105.
- [133] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics: Principles and Practice (2Nd Ed.)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1990. Citation on pages 3 and 20.
- [134] N. C. for Accessible Media. *Image Description Guidelines*. Carl and Ruth Shapiro Family National Center, 2015. Citation on page 28.
- [135] W. Frier, D. Pittner, D. Ablart, M. Obrist, and S. Subramanian. Sampling strategy for ultrasonic mid-air haptics. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, page 1–11, New York, NY, USA, 2019. Association for Computing Machinery. DOI [10.1145/3290605.3300351](https://doi.org/10.1145/3290605.3300351). Citation on page 69.
- [136] H. Fu, D. Cohen-Or, G. Dror, and A. Sheffer. Upright orientation of man-made objects. *ACM Trans. Graph.*, 27(3):1–7, Aug. 2008. DOI [10.1145/1360612.1360641](https://doi.org/10.1145/1360612.1360641). Citation on page 72.
- [137] R. Furferi, L. Governi, Y. Volpe, L. Puggelli, N. Vanni, and M. Carfagni. From 2d to 2.5 d ie from painting to tactile model. *Graphical Models*, 76(6):706–723, 2014. Citation on page 108.
- [138] B. Galea, E. Kia, N. Aird, and P. G. Kry. Stippling with aerial robots. In *Proceedings of the Joint Symposium on Computational Aesthetics and Sketch Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering*, Expressive '16, page 125–134, Goslar, DEU, 2016. Eurographics Association. Citation on pages iv, xx, 6, 7, and 25.
- [139] M. Gandini, Bertone. Alfa romeo carabo. http://www.carstyling.ru/de/car/1968_alfa_romeo_carabo/images/35865/, 1968. Accessed: 2020-02-18. Citation on page 21.
- [140] E. Garces, A. Agarwala, D. Gutierrez, and A. Hertzmann. A similarity measure for illustration style. *ACM Trans. Graph.*, 33(4), July 2014. DOI [10.1145/2601097.2601131](https://doi.org/10.1145/2601097.2601131). Citation on page 110.

- [141] A. Garg, A. Jacobson, and E. Grinspun. Computational design of reconfigurables. *ACM Transactions on Graphics (TOG)*, 35(4), 2016. Citation on page 72.
- [142] B. Gascoigne. *How to Identify Prints*. Thames and Hudson, 2004. Citation on page 17.
- [143] Y. J. Gi, Y. S. Park, S. H. Seo, and K. H. Yoon. Mosaic Rendering using Colored Paper. In M. Ioannides, D. Arnold, F. Niccolucci, and K. Mania, editors, *VAST: International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*. The Eurographics Association, 2006. DOI [10.2312/VAST/VAST06/025-030](https://doi.org/10.2312/VAST/VAST06/025-030). Citation on page 22.
- [144] J. Glenn. Complete patents of nicola tesla. Marboro Books, 1994. Citation on page 21.
- [145] W.-B. Goh. Strategies for shape matching using skeletons. *Similarity Matching in Computer Vision and Multimedia*, 110(3):326–345, June 2008. DOI [10.1016/j.cviu.2007.09.013](https://doi.org/10.1016/j.cviu.2007.09.013). Citation on page 70.
- [146] A. Golovinskiy and T. Funkhouser. Consistent segmentation of 3d models. *Computers & Graphics*, 33(3):262–269, 2009. Citation on page 73.
- [147] A. A. Gooch, J. Long, L. Ji, A. Estey, and B. S. Gooch. Viewing progress in non-photorealistic rendering through heinlein’s lens. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, NPAR ’10, page 165–171, New York, NY, USA, 2010. Association for Computing Machinery. DOI [10.1145/1809939.1809959](https://doi.org/10.1145/1809939.1809959). Citation on page 22.
- [148] B. Gooch, E. Reinhard, C. Moulding, and P. Shirley. Artistic composition for image creation. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, page 83–88, Berlin, Heidelberg, 2001. Springer-Verlag. Citation on page 110.
- [149] B. Gooch, P.-P. J. Sloan, A. Gooch, P. Shirley, and R. Riesenfeld. Interactive technical illustration. In *Proceedings of the 1999 Symposium on Interactive 3D Graphics*, I3D ’99, pages 31–38, New York, NY, USA, 1999. ACM. Citation on page 22.

- [150] T. Götzelmann. Lucentmaps: 3d printed audiovisual tactile maps for blind and visually impaired people. In *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '16, page 81–90, New York, NY, USA, 2016. Association for Computing Machinery. DOI 10.1145/2982142.2982163. Citation on page 66.
- [151] T. Götzelmann. Autonomous selection and printing of 3d models for people who are blind. *ACM Transactions on Accessible Computing (TACCESS)*, 11(3):1–25, 2018. Citation on page 28.
- [152] T. Götzelmann. Visually augmented audio-tactile graphics for visually impaired people. *ACM Transactions on Accessible Computing (TACCESS)*, 11(2):1–31, 2018. Citation on page 28.
- [153] T. Götzelmann, L. Branz, C. Heidenreich, and M. Otto. A personal computer-based approach for 3d printing accessible to blind people. In *Proceedings of the 10th International Conference on Pervasive Technologies Related to Assistive Environments*, PETRA '17, page 1–4, New York, NY, USA, 2017. Association for Computing Machinery. DOI 10.1145/3056540.3064954. Citation on page 28.
- [154] S. Grabli, F. Durand, and F. X. Sillion. Density measure for line-drawing simplification. In *12th Pacific Conference on Computer Graphics and Applications, 2004. PG 2004. Proceedings.*, pages 309–318, 2004. Citation on page 22.
- [155] S. Grabli, E. Turquin, F. Durand, and F. Sillion. Programmable style for npr line drawing. In *Rendering Techniques 2004 (Eurographics Symposium on Rendering)*. ACM Press, 2004. Citation on page 28.
- [156] S. Grabli, E. Turquin, F. Durand, and F. Sillion. Freestyle, 2008. Online at <http://freestyle.sourceforge.net/>. Citation on page 98.
- [157] D. Groeger and J. Steimle. Lasec: Instant fabrication of stretchable circuits using a laser cutter. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing*

- Systems*, CHI '19, New York, NY, USA, 2019. Association for Computing Machinery. DOI [10.1145/3290605.3300929](https://doi.org/10.1145/3290605.3300929). Citation on page 14.
- [158] M. Grundland, C. Gibbs, and N. A. Dodgson. Stylized rendering for multiresolution image representation. In *Human Vision and Electronic Imaging X*, volume 5666, pages 280–292. International Society for Optics and Photonics, 2005. Citation on page 22.
- [159] R. Guseinov, E. Miguel, and B. Bickel. Curveups: Shaping objects from flat plates with tension-actuated curvature. *ACM Trans. Graph.*, 36(4), July 2017. DOI [10.1145/3072959.3073709](https://doi.org/10.1145/3072959.3073709). Citation on page 15.
- [160] P. Haeberli. Paint by numbers: Abstract image representations. *SIGGRAPH Comput. Graph.*, 24(4):207–214, Sept. 1990. DOI [10.1145/97880.97902](https://doi.org/10.1145/97880.97902). Citation on page 22.
- [161] P. M. Hall and M. Owen. Simple canonical views. In *BMVC*, pages 839–848, 2005. Citation on page 72.
- [162] C. Han, Q. Wen, S. He, Q. Zhu, Y. Tan, G. Han, and T.-T. Wong. Deep unsupervised pixelization. *ACM Trans. Graph.*, 37(6), Dec. 2018. DOI [10.1145/3272127.3275082](https://doi.org/10.1145/3272127.3275082). Citation on page 3.
- [163] G. Handjaras, E. Ricciardi, A. Leo, A. Lenci, L. Cecchetti, M. Cosottini, G. Marotta, and P. Pietrini. How concepts are encoded in the human brain: A modality independent, category-based cortical organization of semantic knowledge. 135, 04 2016. Citation on page 70.
- [164] Y. N. a. Harari. *Sapiens : a brief history of humankind*. First U.S. edition. New York : Harper, [2015], 2015. Online at <https://search.library.wisc.edu/catalog/9910419687402121>. Citation on page 1.
- [165] J. Hassberger, J. Lee, and P. Gmytrasiewicz. Fault tree based diagnostics using fuzzy logic. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 19(11):1115–1119, nov 1990. DOI [10.1109/34.61713](https://doi.org/10.1109/34.61713). Citation on page 73.

- [166] L. Hasty. Design principles for tactile graphics, 1999. Online at <http://www.tactilegraphics.org/readability.html>. Citation on page 28.
- [167] J. Hays and I. Essa. Image and video based painterly animation. In *Proceedings of the 3rd International Symposium on Non-photorealistic Animation and Rendering*, NPAR '04, pages 113–120, New York, NY, USA, 2004. ACM. DOI [10.1145/987657.987676](https://doi.org/10.1145/987657.987676). Citation on pages 20, 22, and 23.
- [168] K. He, J. Sun, and X. Tang. Guided image filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(6):1397–1409, June 2013. DOI [10.1109/TPAMI.2012.213](https://doi.org/10.1109/TPAMI.2012.213). Citation on page 47.
- [169] C. G. Healey. Formalizing artistic techniques and scientific visualization for painted renditions of complex information spaces. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 1*, IJCAI'01, page 371–376, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. Citation on page 27.
- [170] M. Heller, T. M. Riddle, E. K. Fulkerson, L. Wemple, A. D. M. Walk, S. M. Guthrie, C. Kranz, and P. Klaus. The influence of viewpoint and object detail in blind people when matching pictures to complex objects. *Perception*, 38 8:1234–50, 2009. Citation on pages 71 and 78.
- [171] S. E. Hernandez and K. E. Barner. Tactile imaging using watershed-based image segmentation. In *Proceedings of the Fourth International ACM Conference on Assistive Technologies*, Assets '00, pages 26–33, New York, NY, USA, 2000. ACM. Citation on page 29.
- [172] R. D. Hersch, F. Collaud, and P. Emmel. Reproducing color images with embedded metallic patterns. *ACM Trans. Graph.*, 22(3):427–434, July 2003. DOI [10.1145/882262.882288](https://doi.org/10.1145/882262.882288). Citation on page 26.
- [173] R. D. Hersch, P. Donzé, and S. Chosson. Color images visible under uv light. *ACM Trans. Graph.*, 26(3):75–es, July 2007. DOI [10.1145/1276377.1276471](https://doi.org/10.1145/1276377.1276471). Citation on page 26.

- [174] A. Hertzmann. Painterly rendering with curved brush strokes of multiple sizes. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, page 453–460, New York, NY, USA, 1998. Association for Computing Machinery. DOI [10.1145/280814.280951](https://doi.org/10.1145/280814.280951). Citation on page 22.
- [175] A. Hertzmann. A survey of stroke-based rendering. *IEEE Computer Graphics and Applications*, (4):70–81, 2003. Citation on page 22.
- [176] A. Hertzmann. Non-photorealistic rendering and the science of art. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, NPAR '10, page 147–157, New York, NY, USA, 2010. Association for Computing Machinery. DOI [10.1145/1809939.1809957](https://doi.org/10.1145/1809939.1809957). Citation on page 109.
- [177] A. Hertzmann and D. Zorin. Illustrating smooth surfaces. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, pages 517–526, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co. Citation on page 28.
- [178] B. Horst and A. Kappers. Using curvature information in haptic shape perception of 3d objects. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 190:361–7, 08 2008. DOI [10.1007/s00221-008-1478-6](https://doi.org/10.1007/s00221-008-1478-6). Citation on page 71.
- [179] I. Howard and R. Allison. Drawing with divergent perspective, ancient and modern. *Perception*, 40:1017–33, 01 2011. Citation on pages 27 and 73.
- [180] K.-W. Hsiao, J.-B. Huang, and H.-K. Chu. Multi-view wire art. *ACM Trans. Graph.*, 37(6), Dec. 2018. DOI [10.1145/3272127.3275070](https://doi.org/10.1145/3272127.3275070). Citation on pages 15 and 24.
- [181] S. Hsiao. Central mechanisms of tactile shape perception. *Current Opinion in Neurobiology*, 18(4):418 – 424, 2008. Sensory systems. Citation on page 29.
- [182] F.-C. Huang, D. Lanman, B. A. Barsky, and R. Raskar. Correcting for optical aberrations using multilayer displays. *ACM Trans. Graph.*, 31(6), Nov. 2012. DOI [10.1145/2366145.2366204](https://doi.org/10.1145/2366145.2366204). Citation on page 27.

- [183] E. Iarussi, A. Bousseau, and T. Tsandilas. The drawing assistant: Automated drawing guidance and feedback from photographs. In *ACM Symposium on User Interface Software and Technology (UIST)*. ACM, 2013. Online at <http://www-sop.inria.fr/revs/Basilic/2013/IBT13>. Citation on page 25.
- [184] T. Igarashi, S. Matsuoka, and H. Tanaka. Teddy: A sketching interface for 3d freeform design. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '99*, page 409–416, USA, 1999. ACM Press/Addison-Wesley Publishing Co. DOI [10.1145/311535.311602](https://doi.org/10.1145/311535.311602). Citation on pages 70, 120, and 123.
- [185] T. Inglis, S. C. Inglis, and C. S. Kaplan. Op art rendering with lines and curves. *Comput. Graph.*, 36:607–621, 2012. Citation on page 71.
- [186] V. Interrante. Illustrating surface shape in volume data via principal direction-driven 3d line integral convolution. Technical report, 1997. Citation on page 71.
- [187] A. Ion, D. Lindlbauer, P. Herholz, M. Alexa, and P. Baudisch. Understanding metamaterial mechanisms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19*, New York, NY, USA, 2019. Association for Computing Machinery. DOI [10.1145/3290605.3300877](https://doi.org/10.1145/3290605.3300877). Citation on page 14.
- [188] T. Isenberg. Evaluating and validating non-photorealistic and illustrative rendering. In *Image and video-based artistic stylisation*, pages 311–331. Springer, 2013. Citation on page 69.
- [189] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A systematic review on the practice of evaluating visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2818–2827, 2013. Citation on pages iv, xx, and 9.
- [190] J. Iseringhausen, M. Weinmann, W. Huang, and M. B. Hullin. Computational parquetry: Fabricated style transfer with wood pixels. *ACM Trans. Graph.*, 39(2), Feb. 2020. DOI [10.1145/3378541](https://doi.org/10.1145/3378541). Citation on pages iv, xx, 5, 16, and 25.

- [191] S. Islam, S. Dipankar, D. Silver, and M. Chen. Spatial and temporal splitting of scalar fields in volume graphics. In *Proceedings of the 2004 IEEE Symposium on Volume Visualization and Graphics, VV '04*, pages 87–94, Washington, DC, USA, 2004. IEEE Computer Society. Citation on page 28.
- [192] T. W. James, G. Humphrey, J. S. Gati, P. Servos, R. S. Menon, and M. A. Goodale. Haptic study of three-dimensional objects activates extrastriate visual areas. *Neuropsychologia*, 40(10):1706 – 1714, 2002. Citation on page 70.
- [193] W. Jarosz, C. Donner, M. Zwicker, and H. W. Jensen. Radiance caching for participating media. *ACM Transactions on Graphics (Presented at SIGGRAPH)*, 27(1):7:1–7:11, mar 2008. DOI [10/cwnw78](https://doi.org/10.1111/j.1467-8659.2010.01805.x). Citation on page 8.
- [194] C. Jayant, M. Renzelmann, D. Wen, S. Krisnandi, R. Ladner, and D. Comden. Automated tactile graphics translation: In the field. In *ACM SIGACCESS. Conference on Computers and Accessibility, Assets '07*, pages 75–82, New York, NY, USA, 2007. ACM. Citation on page 28.
- [195] Z. Ji, L. Liu, and Y. Wang. B-mesh: A modeling system for base meshes of 3d articulated shapes. *Computer Graphics Forum*, 29(7):2169–2177, 2010. DOI [10.1111/j.1467-8659.2010.01805.x](https://doi.org/10.1111/j.1467-8659.2010.01805.x). Citation on page 123.
- [196] JMS. True reverse perspective, 2010. Online at <https://vimeo.com/12518619>. Citation on page 27.
- [197] T. Judd, F. Durand, and E. H. Adelson. Apparent ridges for line drawing. *ACM Trans. Graph.*, 26(3):19, 2007. Citation on page 83.
- [198] O. V. Kaick, N. Fish, Y. Kleiman, S. Asafi, and D. Cohen-OR. Shape segmentation by approximate convexity analysis. *ACM Trans. Graph.*, 34(1), Dec. 2015. DOI [10.1145/2611811](https://doi.org/10.1145/2611811). Citation on page 73.

- [199] E. Kalogerakis, A. Hertzmann, and K. Singh. Learning 3d mesh segmentation and labeling. *ACM Trans. Graph.*, 29(4):102:1–102:12, July 2010. Citation on pages 73, 100, and 104.
- [200] E. Kalogerakis, D. Nowrouzezahrai, S. Breslav, and A. Hertzmann. Learning hatching for pen-and-ink illustration of surfaces. *ACM Trans. Graph.*, 31(1):1:1–1:17, Feb. 2012. Citation on pages 28 and 83.
- [201] H. Kang, S. Lee, and C. K. Chui. Coherent line drawing. In *Proceedings of the 5th International Symposium on Non-Photorealistic Animation and Rendering, NPAR '07*, page 43–50, New York, NY, USA, 2007. Association for Computing Machinery. DOI [10.1145/1274871.1274878](https://doi.org/10.1145/1274871.1274878). Citation on page 34.
- [202] C. S. Kaplan, R. Bosch, et al. Tsp art. In *Renaissance Banff: Mathematics, music, art, culture*, pages 301–308. Bridges Conference, 2005. Citation on page 22.
- [203] M. Kazhdan. *Differential Geometry: Curves, Tangents, and Curvature*. Course: Advanced Topics in Computer Graphics: Discrete Differential Geometry, 2009. Online at <http://www.cs.jhu.edu/~misha/Fall09/1-curves.pdf>. Citation on page 116.
- [204] D. F. Keefe, D. A. Feliz, T. Moscovich, D. H. Laidlaw, and J. J. LaViola. Cavepainting: A fully immersive 3d artistic medium and interactive experience. In *Proceedings of the 2001 Symposium on Interactive 3D Graphics, I3D '01*, page 85–93, New York, NY, USA, 2001. Association for Computing Machinery. DOI [10.1145/364338.364370](https://doi.org/10.1145/364338.364370). Citation on page 24.
- [205] J. M. Kennedy. *Drawing and the Blind: Pictures to Touch*. Yale Univ. Press, New Haven, CT 06520-9040, USA, 1993. Citation on page 67.
- [206] J. Kerber, A. Belyaev, and H.-P. Seidel. Feature preserving depth compression of range images. In *Proceedings of the 23rd Spring Conference on Computer Graphics, SCCG '07*, page 101–105, New York, NY, USA, 2007. Association for Computing Machinery. DOI [10.1145/2614348.2614363](https://doi.org/10.1145/2614348.2614363). Citation on page 108.

- [207] J. Kerber, M. Wang, J. Chang, J. J. Zhang, A. Belyaev, and H.-P. Seidel. Computer assisted relief generation—a survey. *Comput. Graph. Forum*, 31(8):2363–2377, Dec. 2012. DOI [10.1111/j.1467-8659.2012.03185.x](https://doi.org/10.1111/j.1467-8659.2012.03185.x). Citation on page 108.
- [208] M. Kilian, A. Monszpart, and N. J. Mitra. String actuated curved folded surfaces. *ACM Trans. Graph.*, 36(3), May 2017. DOI [10.1145/3015460](https://doi.org/10.1145/3015460). Citation on page 15.
- [209] J. Kim, H. Oh, and T. Yeh. A study to empower children to design movable tactile pictures for children with visual impairments. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*, TEI '15, pages 703–708, New York, NY, USA, 2015. ACM. Citation on page 28.
- [210] J. Kim and T. Yeh. Toward 3d-printed movable tactile pictures for children with visual impairments. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 2815–2824, New York, NY, USA, 2015. ACM. Citation on pages 19, 28, and 108.
- [211] T. S. Kiser, M. Eigensatz, M. M. Nguyen, P. Bompas, and M. Pauly. Architectural caustics - controlling light with geometry. In *AAG*, 2012. Citation on page 25.
- [212] R. L. Klatzky, J. M. Loomis, S. J. Lederman, H. Wake, and N. Fujita. Haptic identification of objects and their depictions. *Perception & Psychophysics*, 54:170–8, 09 1993. Citation on pages 20 and 67.
- [213] D. A. Kleffner and V. S. Ramachandran. On the perception of shape from shading. *Perception & Psychophysics*, 52(1):18–36, Jan. 1992. DOI [10.3758/BF03206757](https://doi.org/10.3758/BF03206757). Citation on page 34.
- [214] M. Konakov, K. Crane, B. Deng, S. Bouaziz, D. Piker, and M. Pauly. Beyond developable: Computational design and fabrication with auxetic materials. *ACM Trans. Graph.*, 35(4), July 2016. DOI [10.1145/2897824.2925944](https://doi.org/10.1145/2897824.2925944). Citation on pages 14 and 15.

- [215] B. Koo, J. Hergel, S. Lefebvre, and N. J. Mitra. Towards zero-waste furniture design. *IEEE Transactions on Visualization and Computer Graphics*, 23(12):2627–2640, Dec 2017. DOI [10.1109/TVCG.2016.2633519](https://doi.org/10.1109/TVCG.2016.2633519). Citation on page 15.
- [216] B. Koo, W. Li, J. Yao, M. Agrawala, and N. J. Mitra. Creating works-like prototypes of mechanical objects. *ACM Trans. Graph.*, 33(6), Nov. 2014. DOI [10.1145/2661229.2661289](https://doi.org/10.1145/2661229.2661289). Citation on page 2.
- [217] J. Kopf, M. Agrawala, D. Barger, D. Salesin, and M. Cohen. Automatic generation of destination maps. *ACM Trans. Graph.*, 29(6), Dec. 2010. DOI [10.1145/1882261.1866184](https://doi.org/10.1145/1882261.1866184). Citation on page 22.
- [218] P. Krähenbühl and V. Koltun. Efficient inference in fully connected CRFs with gaussian edge potentials. In *Advances in Neural Information Processing Systems 24*, pages 109–117. Curran Associates, Inc., 2011. Citation on pages 35, 41, and 46.
- [219] J. Krikke. Axonometry: a matter of perspective. *IEEE Computer Graphics and Applications*, 20(4):7–11, July 2000. Citation on page 27.
- [220] S. E. Krufka, K. E. Barner, and T. C. Aysal. Visual to tactile conversion of vector graphics. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 15(2):310–321, 2007. Citation on page 29.
- [221] T. Kuipers, E. Doubrovski, and J. Verlinden. 3d hatching: Linear halftoning for dual extrusion fused deposition modeling. In *Proceedings of the 1st Annual ACM Symposium on Computational Fabrication, SCF '17*, New York, NY, USA, 2017. Association for Computing Machinery. DOI [10.1145/3083157.3083163](https://doi.org/10.1145/3083157.3083163). Citation on page 15.
- [222] M.-H. Kuo, Y.-E. Lin, H.-K. Chu, R.-R. Lee, and Y.-L. Yang. Pixel2brick: Constructing brick sculptures from pixel art. *Computer Graphics Forum (Proc. Pacific Graphics)*, 34, 2015. Citation on pages 5 and 25.

- [223] M. Kurze. Tdraw: a computer-based tactile drawing tool for blind people. In *Proceedings of the second annual ACM conference on Assistive technologies*, pages 131–138, 1996. Citation on page 28.
- [224] M. Kurze. Rendering drawings for interactive haptic perception. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '97*, pages 423–430, New York, NY, USA, 1997. ACM. Citation on pages 29 and 67.
- [225] J. E. Kyprianidis, J. P. Collomosse, T. Wang, and T. Isenberg. State of the “art”: A taxonomy of artistic stylization techniques for images and video. *IEEE Transactions on Visualization and Computer Graphics*, 19:866–885, 2013. Citation on pages 3 and 22.
- [226] C. Lü and X. Chen. Image watercolorization based on visual weight-map. In *2014 7th International Congress on Image and Signal Processing*, pages 233–237, Oct 2014. DOI [10.1109/CISP.2014.7003783](https://doi.org/10.1109/CISP.2014.7003783). Citation on page 22.
- [227] S. Lakatos and L. E. Marks. Haptic form perception: Relative salience of local and global features. *Perception & Psychophysics*, 61(5):895–908, Jan 1999. Citation on page 70.
- [228] Y. Lan, Y. Dong, F. Pellacini, and X. Tong. Bi-scale appearance fabrication. *ACM Trans. Graph.*, 32(4), July 2013. DOI [10.1145/2461912.2461989](https://doi.org/10.1145/2461912.2461989). Citation on page 15.
- [229] J. Lansdown and S. Schofield. Expressive rendering: a review of nonphotorealistic techniques. *IEEE Computer Graphics and Applications*, 15(3):29–37, May 1995. DOI [10.1109/38.376610](https://doi.org/10.1109/38.376610). Citation on page 22.
- [230] D. Lareau and J. Lang. Haptic rendering of photographs. In *2012 IEEE International Workshop on Haptic Audio Visual Environments and Games (HAVE 2012) Proceedings*, pages 107–112. IEEE, 2012. Citation on page 29.
- [231] M. Lau, K. Dev, W. Shi, J. Dorsey, and H. Rushmeier. Tactile mesh saliency. *ACM Trans. Graph.*, 35(4):52:1–52:11, July 2016. Citation on pages 69, 100, and 107.

- [232] K. Lawonn, I. Viola, B. Preim, and T. Isenberg. A survey of surface-based illustrative rendering for visualization. *Computer Graphics Forum*, 37(6):205–234, 2018. DOI [10.1111/cgf.13322](https://doi.org/10.1111/cgf.13322). Citation on page 27.
- [233] S. Lederman and R. Klatzky. Haptic perception: A tutorial. *Attention, Perception & Psychophysics*, 71:1439–59, 10 2009. Citation on pages 28, 70, and 84.
- [234] R. Lenth. emmeans: Estimated marginal means, aka least-squares means. r package version 1.4. 3.01, 2019. Citation on page 95.
- [235] C. Leone, V. Lopresto, and I. D. Iorio. Wood engraving by Q-switched diode-pumped frequency-doubled Nd:YAG green laser. *Optics and Lasers in Engineering*, 47(1):161 – 168, 2009. DOI <https://doi.org/10.1016/j.optlaseng.2008.06.019>. Citation on page 44.
- [236] V. Levesque. Blindness, technology, and haptics. Technical report, Ecole de technologie supérieure, 10 2005. Citation on page 28.
- [237] J. Li, J. Jacobs, M. Chang, and B. Hartmann. Direct and immediate drawing with cnc machines. In *Proceedings of the 1st Annual ACM Symposium on Computational Fabrication*, SCF ’17, New York, NY, USA, 2017. Association for Computing Machinery. DOI [10.1145/3083157.3096344](https://doi.org/10.1145/3083157.3096344). Citation on pages 5 and 25.
- [238] J. Li, S. Kim, J. A. Miele, M. Agrawala, and S. Follmer. Editing spatial layouts through tactile templates for people with visual impairments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI ’19, pages 206:1–206:11, New York, NY, USA, 2019. ACM. Citation on pages 19 and 28.
- [239] W. Li, M. Agrawala, B. Curless, and D. Salesin. Automated generation of interactive 3d exploded view diagrams. *ACM Trans. Graph.*, 27(3):101:1–101:7, Aug. 2008. Citation on pages 28, 100, and 107.
- [240] W. Li, M. Agrawala, B. Curless, and D. Salesin. Automated generation of interactive 3d exploded view diagrams. *ACM Trans. Graph.*, 27(3):1–7, Aug. 2008. DOI [10.1145/1360612.1360700](https://doi.org/10.1145/1360612.1360700). Citation on page 107.

- [241] W. Li, M. Agrawala, and D. Salesin. Interactive image-based exploded view diagrams. In *Proceedings of Graphics Interface 2004*, GI '04, page 203–212, Waterloo, CAN, 2004. Canadian Human-Computer Communications Society. Citation on page 107.
- [242] W. Li, L. Ritter, M. Agrawala, B. Curless, and D. Salesin. Interactive cutaway illustrations of complex 3d models. *ACM Trans. Graph.*, 26(3), July 2007. Citation on page 28.
- [243] S. Lin and P. Hanrahan. Modeling how people extract color themes from images. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 3101–3110, New York, NY, USA, 2013. ACM. DOI [10.1145/2470654.2466424](https://doi.org/10.1145/2470654.2466424). Citation on pages 33 and 103.
- [244] Y.-E. Lin, Y.-L. Yang, and H.-K. Chu. Scale-aware black-and-white abstraction of 3d shapes. *ACM Trans. Graph.*, 37(4), July 2018. DOI [10.1145/3197517.3201372](https://doi.org/10.1145/3197517.3201372). Citation on page 3.
- [245] T. Lindemeier. *e-David : Non-Photorealistic Rendering using a Robot and Visual Feedback*. PhD thesis, Universität Konstanz, Konstanz, 2018. Citation on pages iii, v, xix, xxi, 6, and 24.
- [246] T. Lindemeier, J. Gülzow, and O. Deussen. Painterly rendering using limited paint color palettes. EG VMV '18, pages 135–145. Eurographics Association, 2018. DOI [10.2312/vmv.20181263](https://doi.org/10.2312/vmv.20181263). Citation on pages 5, 8, and 34.
- [247] T. Lindemeier, J. Metzner, L. Pollak, and O. Deussen. Hardware-Based Non-Photorealistic Rendering Using a Painting Robot. *Computer Graphics Forum*, 2015. DOI [10.1111/cgf.12562](https://doi.org/10.1111/cgf.12562). Citation on pages 5, 25, 26, and 35.
- [248] T. Lindemeier, S. Pirk, and O. Deussen. Image stylization with a painting machine using semantic hints. *Comput. Graph.*, 37(5):293–301, aug 2013. DOI [10.1016/j.cag.2013.01.005](https://doi.org/10.1016/j.cag.2013.01.005). Citation on pages 3, 5, 8, 9, and 25.

- [249] T. Lindemeier, M. Spicker, and O. Deussen. Artistic composition for painterly rendering. The Eurographics Association, 2016. DOI [10.2312/vmv.20161350](https://doi.org/10.2312/vmv.20161350). Citation on pages 5 and 41.
- [250] T. Lindemeier, M. Spicker, and O. Deussen. Artistic composition for painterly rendering. In M. Hullin, M. Stamminger, and T. Weinkauff, editors, *Vision, Modeling and Visualization*. The Eurographics Association, 2016. DOI [10.2312/vmv.20161350](https://doi.org/10.2312/vmv.20161350). Citation on page 22.
- [251] J. I. Lipton, A. Schulz, A. Spielberg, L. Trueba, W. Matusik, and D. Rus. Robot assisted carpentry for mass customization. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8, 2018. Citation on page 15.
- [252] W. Lira, C.-W. Fu, and H. Zhang. Fabricable eulerian wires for 3d shape abstraction. *ACM Trans. Graph.*, 37(6), Dec. 2018. DOI [10.1145/3272127.3275049](https://doi.org/10.1145/3272127.3275049). Citation on page 14.
- [253] P. Litwinowicz. Processing images and video for an impressionist effect. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97*, page 407–414, USA, 1997. ACM Press/Addison-Wesley Publishing Co. DOI [10.1145/258734.258893](https://doi.org/10.1145/258734.258893). Citation on page 23.
- [254] H.-T. D. Liu and A. Jacobson. Cubic stylization. *ACM Transactions on Graphics*, 2019. Citation on page 28.
- [255] L. Liu, D. Ceylan, C. Lin, W. Wang, and N. J. Mitra. Image-based reconstruction of wire art. *ACM Trans. Graph.*, 36(4), July 2017. DOI [10.1145/3072959.3073682](https://doi.org/10.1145/3072959.3073682). Citation on page 14.
- [256] Y.-S. Liu and K. Ramani. Robust principal axes determination for point-based shapes using least median of squares. *Comput. Aided Des.*, 41(4):293–305, Apr. 2009. DOI [10.1016/j.cad.2008.10.012](https://doi.org/10.1016/j.cad.2008.10.012). Citation on page 72.

- [257] M. Livesu, M. Attene, G. Patane, and M. Spagnuolo. Explicit cylindrical maps for general tubular shapes. *Computer-Aided Design*, 05 2017. Citation on pages 100 and 104.
- [258] K.-Y. Lo, C.-W. Fu, and H. Li. 3d polyomino puzzle. *ACM Trans. Graph.*, 28(5):1–8, Dec. 2009. DOI [10.1145/1618452.1618503](https://doi.org/10.1145/1618452.1618503). Citation on page 15.
- [259] Y.-H. Lo, R.-R. Lee, and H.-K. Chu. Generating color scribble images using multi-layered monochromatic strokes dithering. *Computer Graphics Forum*, 38(2):265–276, 2019. DOI [10.1111/cgf.13636](https://doi.org/10.1111/cgf.13636). Citation on page 22.
- [260] A. Lowet, C. Firestone, and B. Scholl. Seeing structure: Shape skeletons modulate perceived similarity. *Attention, perception & psychophysics*, 80, 03 2018. DOI [10.3758/s13414-017-1457-8](https://doi.org/10.3758/s13414-017-1457-8). Citation on page 70.
- [261] J. Lu, C. Barnes, S. DiVerdi, and A. Finkelstein. Realbrush: Painting with examples of physical media. *ACM Trans. Graph.*, 32(4):117:1–117:12, July 2013. DOI [10.1145/2461912.2461998](https://doi.org/10.1145/2461912.2461998). Citation on page 24.
- [262] J. Lu, C. Barnes, C. Wan, P. Asente, R. Mech, and A. Finkelstein. Decobrush: Drawing structured decorative patterns by example. *ACM Trans. Graph.*, 33(4), July 2014. DOI [10.1145/2601097.2601190](https://doi.org/10.1145/2601097.2601190). Citation on page 23.
- [263] J. Lu, P. V. Sander, and A. Finkelstein. Interactive painterly stylization of images, videos and 3d animations. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D '10*, page 127–134, New York, NY, USA, 2010. Association for Computing Machinery. DOI [10.1145/1730804.1730825](https://doi.org/10.1145/1730804.1730825). Citation on page 34.
- [264] T. Luft and O. Deussen. Interactive watercolor animations. 2005. Citation on pages 23 and 33.
- [265] T. Luft and O. Deussen. Real-time watercolor illustrations of plants using a blurred depth test. In *Proceedings of the 4th International Symposium on Non-Photorealistic Animation*

- and Rendering*, NPAR '06, page 11–20, New York, NY, USA, 2006. Association for Computing Machinery. DOI [10.1145/1124728.1124732](https://doi.org/10.1145/1124728.1124732). Citation on pages 23 and 34.
- [266] M. Lukáč, J. Fišer, J.-C. Bazin, O. Jamriška, A. Sorkine-Hornung, and D. Sýkora. Painting by feature: Texture boundaries for example-based image creation. *ACM Trans. Graph.*, 32(4), July 2013. DOI [10.1145/2461912.2461956](https://doi.org/10.1145/2461912.2461956). Citation on page 23.
- [267] E. B. Lum and Kwan-Liu Ma. Non-photorealistic rendering using watercolor inspired textures and illumination. In *Proceedings Ninth Pacific Conference on Computer Graphics and Applications. Pacific Graphics 2001*, pages 322–330, 2001. Citation on pages 22 and 71.
- [268] Z. Lun, E. Kalogerakis, R. Wang, and A. Sheffer. Functionality preserving shape style transfer. *ACM Trans. Graph.*, 35(6), Nov. 2016. DOI [10.1145/2980179.2980237](https://doi.org/10.1145/2980179.2980237). Citation on page 109.
- [269] L. Maguire, M. Papas, W. Jarosz, P. Fox, G. Dicoski, and M. Olivares. The modelling of caustics to produce a projection image. 02 2018. Citation on page 25.
- [270] R. Maharik, M. Bessmeltsev, A. Sheffer, A. Shamir, and N. Carr. Digital micrography. In *ACM SIGGRAPH 2011 Papers*, SIGGRAPH '11, New York, NY, USA, 2011. Association for Computing Machinery. DOI [10.1145/1964921.1964995](https://doi.org/10.1145/1964921.1964995). Citation on page 22.
- [271] M. T. Mahmoudi, M. Sharifi, and M. Fathy. A classified and comparative study of edge detection algorithms. In *Information Technology: Coding and Computing, International Conference on(ITCC)*, volume 00, page 0117, 04 2002. DOI [10.1109/ITCC.2002.1000371](https://doi.org/10.1109/ITCC.2002.1000371). Citation on page 104.
- [272] H. T. Maia, D. Li, Y. Yang, and C. Zheng. Layercode: Optical barcodes for 3d printed shapes. *ACM Trans. Graph.*, 38(4), July 2019. DOI [10.1145/3306346.3322960](https://doi.org/10.1145/3306346.3322960). Citation on pages 2 and 8.

- [273] L. Malomo, J. Pérez, E. Iarussi, N. Pietroni, E. Miguel, P. Cignoni, and B. Bickel. Flexmaps: Computational design of flat flexible shells for shaping 3d objects. *ACM Trans. Graph.*, 37(6), Dec. 2018. DOI [10.1145/3272127.3275076](https://doi.org/10.1145/3272127.3275076). Citation on page 15.
- [274] T. Malzbender, R. Samadani, S. Scher, A. Crume, D. Dunn, and J. Davis. Printing reflectance functions. *ACM Trans. Graph.*, 31(3), June 2012. DOI [10.1145/2167076.2167078](https://doi.org/10.1145/2167076.2167078). Citation on page 26.
- [275] P. Maragos. Morphological signal and image processing. 1998. Citation on page 104.
- [276] L. Markosian, B. J. Meier, M. A. Kowalski, L. S. Holden, J. D. Northrup, and J. F. Hughes. Art-based rendering with continuous levels of detail. In *Proceedings of the First International Symposium on Non Photorealistic Animation and Rendering (NPAR) for Art and Entertainment*, June 2000. To be held in Annecy, France. Citation on page 22.
- [277] M. Marrazzo. Curvilinear perspective: how to draw and examples, 04 2017. DOI [10.13140/RG.2.2.21979.36649](https://doi.org/10.13140/RG.2.2.21979.36649). Citation on page 73.
- [278] D. Martín, G. Arroyo, A. Rodríguez, and T. Isenberg. A survey of digital stippling. *Computers & Graphics*, 67:24–44, 2017. Citation on page 6.
- [279] R. Martín, J. Iseringhausen, M. Weinmann, and M. B. Hullin. Multimodal perception of material properties. In *ACM SIGGRAPH Symposium on Applied Perception, SAP '15*, pages 33–40, New York, NY, USA, Sept. 2015. ACM. DOI [10.1145/2804408.2804420](https://doi.org/10.1145/2804408.2804420). Citation on page 69.
- [280] R. Martín, M. Weinmann, and M. B. Hullin. Digital transmission of subjective material appearance. *Journal of WSCG*, 25(2):57–66, June 2017. Online at http://wscg.zcu.cz/WSCG2017/!_2017_Journal_WSCG-No-2.pdf. Citation on page 69.
- [281] F. MartíNez, C. Ogayar, J. R. JiméNez, and A. J. Rueda. A simple algorithm for boolean operations on polygons. *Adv. Eng. Softw.*, 64:11–19, Oct. 2013. DOI [10.1016/j.advengsoft.2013.04.004](https://doi.org/10.1016/j.advengsoft.2013.04.004). Citation on page 104.

- [282] J. Martínez, S. Hornus, H. Song, and S. Lefebvre. Polyhedral voronoi diagrams for additive manufacturing. *ACM Trans. Graph.*, 37(4), July 2018. DOI 10.1145/3197517.3201343. Citation on page 3.
- [283] W. Matusik, B. Ajdin, J. Gu, J. Lawrence, H. P. A. Lensch, F. Pellacini, and S. Rusinkiewicz. Printing spatially-varying reflectance. In *ACM SIGGRAPH Asia 2009 Papers*, SIGGRAPH Asia '09, New York, NY, USA, 2009. Association for Computing Machinery. DOI 10.1145/1661412.1618474. Citation on page 26.
- [284] P. Mayencourt and C. V. Mueller. Structural optimization and digital fabrication of timber beams. 2017. Citation on page 14.
- [285] M. McEachern. Fox/blue sky bring a flood of innovation to ice age: The meltdown, 2006. Online at <http://www.cgw.com/Publications/CGW/2006/Volume-29-Issue-4-April-2006-/Thawsome.aspx>. Citation on page 20.
- [286] M. J. McGuffin, L. Tancau, and R. Balakrishnan. Using deformations for browsing volumetric data. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, VIS '03, pages 53–, Washington, DC, USA, 2003. IEEE Computer Society. Citation on page 28.
- [287] N. McLennan, F. Hayden, T. Poppe, and W. Pierce. The good tactile graphic. Louisville, KY 40206-0085, USA, 1998. Citation on page 28.
- [288] V. Megaro, B. Thomaszewski, D. Gauge, E. Grinspun, S. Coros, and M. Gross. Chacra: An interactive design system for rapid character crafting. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '14, page 123–130, Goslar, DEU, 2015. Eurographics Association. Citation on page 25.
- [289] R. Mehra, Q. Zhou, J. Long, A. Sheffer, A. Gooch, and N. J. Mitra. Abstraction of man-made shapes. *ACM Transactions on Graphics*, 28(5):#137, 1–10, 2009. Citation on pages 28, 100, and 107.
- [290] C. Mei, V. Popescu, and E. Sacks. The Occlusion Camera. *Computer Graphics Forum*, 2005. DOI 10.1111/j.1467-8659.2005.00858.x. Citation on page 73.

- [291] G. Memoli, M. Caleap, M. Asakawa, D. R. Sahoo, B. W. Drinkwater, and S. Subramanian. Metamaterial bricks and quantization of meta-surfaces. *Nature Communications*, 8(1):14608, feb 2017. DOI [10.1038/ncomms14608](https://doi.org/10.1038/ncomms14608). Citation on page 14.
- [292] E. Miguel, M. Lepoutre, and B. Bickel. Computational design of stable planar-rod structures. *ACM Trans. Graph.*, 35(4), July 2016. DOI [10.1145/2897824.2925978](https://doi.org/10.1145/2897824.2925978). Citation on page 15.
- [293] N. J. Mitra and M. Pauly. Shadow art. *ACM Trans. Graph.*, 28(5):1–7, Dec. 2009. DOI [10.1145/1618452.1618502](https://doi.org/10.1145/1618452.1618502). Citation on page 24.
- [294] K. Mo, S. Zhu, A. X. Chang, L. Yi, S. Tripathi, L. J. Guibas, and H. Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. *CoRR*, abs/1812.02713, 2018. Online at <http://arxiv.org/abs/1812.02713>. Citation on pages 72 and 73.
- [295] K. Mo, S. Zhu, A. X. Chang, L. Yi, S. Tripathi, L. J. Guibas, and H. Su. PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. Citation on pages 81 and 99.
- [296] K. Moltenbrey. Creature feature: 'the shape of water', 2018. Online at <http://www.cgw.com/Press-Center/In-Focus/2018/Creature-Feature-The-Shape-of-Water.aspx>. Citation on page 3.
- [297] C. Montalto, I. Garcia-Dorado, D. Aliaga, M. M. Oliveira, and F. Meng. A total variation approach for customizing imagery to improve visual acuity. *ACM Trans. Graph.*, 34(3), May 2015. DOI [10.1145/2717307](https://doi.org/10.1145/2717307). Citation on page 27.
- [298] D. Moore. *Disabled Students in Education: Technology, Transition, and Inclusivity: Technology, Transition, and Inclusivity*. IGI Global, Hershey, PA 17033, USA, 2011. Citation on page 29.

- [299] E. Moore, M. Porter, I. Karamouzas, and V. Zordan. Precision control of tensile properties in fabric for computational fabrication. In *Proceedings of the 2nd ACM Symposium on Computational Fabrication*, SCF '18, New York, NY, USA, 2018. Association for Computing Machinery. Online at <https://doi.org/10.1145/3213512.3213514>. Citation on page 14.
- [300] A. Moringen, K. Krieger, R. Haschke, and H. Ritter. Haptic search for complex 3d shapes subject to geometric transformations or partial occlusion. In *2017 IEEE World Haptics Conference (WHC)*, pages 299–304, June 2017. Citation on page 78.
- [301] B. Morse, D. Thornton, Q. Xia, and J. Uibel. Image-based color schemes. *International Conference on Image Processing*, 2007. Citation on page 33.
- [302] S. Mueller, P. Punpongsanon, D. Turakhia, and M. D. Dogan, HCI Engineering Group, MIT CSAIL. Personal fabrication research in hci and graphics: An overview of related work. <https://hcie.csail.mit.edu/fabpub/>, 2017. Accessed: 2020-02-18. Citation on page 14.
- [303] A. Muntoni, S. Nuvoli, A. Scalas, A. Tola, L. Malomo, and R. Scateni. Mill and fold: Shape simplification for fabrication. *Computers & Graphics*, 80:17–28, May 2019. DOI [10.1016/j.cag.2019.03.003](https://doi.org/10.1016/j.cag.2019.03.003). Citation on page 15.
- [304] A. Myers. Screw tactile pictures. <http://www.andrewmyersart.com/read-me>, 2019. Accessed: 2020-03-31. Citation on page 20.
- [305] V. Narayanan, L. Albaugh, J. Hodgins, S. Coros, and J. McCann. Automatic machine knitting of 3d meshes. *ACM Trans. Graph.*, 37(3):35:1–35:15, Aug. 2018. DOI [10.1145/3186265](https://doi.org/10.1145/3186265). Citation on page 14.
- [306] V. Narayanan, K. Wu, C. Yuksel, and J. McCann. Visual knitting machine programming. *ACM Trans. Graph.*, 38(4), July 2019. DOI [10.1145/3306346.3322995](https://doi.org/10.1145/3306346.3322995). Citation on pages iv, xx, 3, and 16.

- [307] A. Nayak and K. E. Barner. Optimal halftoning for tactile imaging. *IEEE Transactions on neural systems and rehabilitation engineering*, 12(2):216–227, 2004. Citation on page 104.
- [308] R. Netz, The Archimedes Palimpsest. The diagrams as floating bodies. <http://www.archimedespalimpsest.org/about/scholarship/floating-bodies.php>, 2000. Accessed: 2020-02-18. Citation on page 21.
- [309] R. Netz, The Archimedes Palimpsest. Part four: The archimedes codex. <https://www.gresham.ac.uk/lectures-and-events/part-four-the-archimedes-codex>, 2011. Accessed: 2020-02-17. Citation on page 21.
- [310] A. S. Nittala, K. Kruttwig, J. Lee, R. Bennewitz, E. Arzt, and J. Steimle. Like a second skin: Understanding how epidermal devices affect human tactile perception. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, New York, NY, USA, 2019. Association for Computing Machinery. DOI 10.1145/3290605.3300610. Citation on page 69.
- [311] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. *CoRR*, abs/1505.04366, 2015. Online at <http://arxiv.org/abs/1505.04366>. Citation on page 104.
- [312] Y. Ochiai, T. Hoshi, and J. Rekimoto. Pixie dust: Graphics generated by levitated and animated objects in computational acoustic-potential field. *ACM Trans. Graph.*, 33(4), July 2014. DOI 10.1145/2601097.2601118. Citation on page 14.
- [313] P. O'Donovan, A. Agarwala, and A. Hertzmann. Color compatibility from large datasets. *ACM Trans. Graph.*, 30(4):63:1–63:12, July 2011. DOI 10.1145/2010324.1964958. Citation on page 33.
- [314] P. O'Donovan and A. Hertzmann. Anipaint: Interactive painterly animation from video. *IEEE Transactions on Visualization and Computer Graphics*, 18(3):475–487, 2012. DOI <http://doi.ieeecomputersociety.org/10.1109/TVCG.2011.51>. Citation on page 23.

- [315] B. A. of North America. *Guidelines and Standards for Tactile Graphics*. Braille Authority of North America, Pittsburgh, PA, USA, 2010. Citation on pages xxvi, 12, 20, 28, 66, 76, 78, 82, 83, 84, 88, 89, 90, 91, 94, and 98.
- [316] S. C. Olsen, B. A. Maxwell, and B. Gooch. Interactive vector fields for painterly rendering. In *Proceedings of Graphics Interface 2005*, GI '05, page 241–247, Waterloo, CAN, 2005. Canadian Human-Computer Communications Society. Citation on pages 22, 23, and 33.
- [317] W. M. Oosting. Giving visual impaired people an extra dimension: Designing a tactile graphics handling braille display. In *1st Twente Student Conference on IT*. Citeseer, 2004. Citation on page 19.
- [318] W. H. Oskay and L. M. Edman, Evil Mad Scientist Laboratories. Diy and open source hardware for art, education, and world domination. <https://www.evilmadscientist.com/about/>, 2006. Accessed: 2020-02-29. Citation on page 4.
- [319] J. Ou, M. Skouras, N. Vlavianos, F. Heibeck, C.-Y. Cheng, J. Peters, and H. Ishii. Aeromorph - heat-sealing inflatable shape-change materials for interaction design. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, UIST '16, page 121–132, New York, NY, USA, 2016. Association for Computing Machinery. DOI 10.1145/2984511.2984520. Citation on page 14.
- [320] J. T. B. Overvelde, T. A. de Jong, Y. Shevchenko, S. A. F. Becerra, G. M. Whitesides, J. C. Weaver, C. Hoberman, and K. Bertoldi. A three-dimensional actuated origami-inspired transformable metamaterial with multiple degrees of freedom. In *Nature communications*, 2016. Citation on page 14.
- [321] S. Paneels and J. C. Roberts. Review of designs for haptic data visualization. *IEEE Transactions on Haptics*, 3(2):119–137, 2009. Citation on page 29.
- [322] A. Panotopoulou, E. Ross, K. Welker, E. Hubert, and G. Morin. *Scaffolding a Skeleton*, pages 17–35. Springer International Publishing, 2018. Citation on page 123.

- [323] D. Panozzo, O. Diamanti, S. Paris, M. Tarini, E. Sorkine, and O. Sorkine-Hornung. Texture mapping real-world objects with hydrographics. *Computer Graphics Forum*, 34(5):65–75, 2015. DOI [10.1111/cgf.12697](https://doi.org/10.1111/cgf.12697). Citation on page 15.
- [324] D. Panozzo, O. Diamanti, S. Paris, M. Tarini, E. Sorkine, and O. Sorkine-Hornung. Texture mapping real-world objects with hydrographics. *Computer Graphics Forum (proceedings of EUROGRAPHICS Symposium on Geometry Processing)*, 34(5):65–75, 2015. Citation on pages 15 and 26.
- [325] M. Papas, T. Houit, D. Nowrouzezahrai, M. Gross, and W. Jarosz. The magic lens: Refractive steganography. *ACM Transactions on Graphics (TOG)*, 31, 11 2012. DOI [10.1145/2366145.2366205](https://doi.org/10.1145/2366145.2366205). Citation on pages 8 and 24.
- [326] M. Papas, W. Jarosz, W. Jakob, S. Rusinkiewicz, W. Matusik, and T. Weyrich. Goal-based caustics. *Computer Graphics Forum*, 30(2):503–511, 2011. DOI [10.1111/j.1467-8659.2011.01876.x](https://doi.org/10.1111/j.1467-8659.2011.01876.x). Citation on page 25.
- [327] G. Patow and X. Pueyo. A survey of inverse surface design from light transport behavior specification. *Computer Graphics Forum*, 24(4):773–789, 2005. DOI [10.1111/j.1467-8659.2005.00901.x](https://doi.org/10.1111/j.1467-8659.2005.00901.x). Citation on page 25.
- [328] H. Pedersen and K. Singh. Organic labyrinths and mazes. In *Proceedings of the 4th International Symposium on Non-Photorealistic Animation and Rendering, NPAR '06*, page 79–86, New York, NY, USA, 2006. Association for Computing Machinery. DOI [10.1145/1124728.1124742](https://doi.org/10.1145/1124728.1124742). Citation on page 22.
- [329] F. Pellacini, J. A. Ferwerda, and D. P. Greenberg. Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '00*, page 55–64, USA, 2000. ACM Press/Addison-Wesley Publishing Co. DOI [10.1145/344779.344812](https://doi.org/10.1145/344779.344812). Citation on page 69.

- [330] T. Pereira, C. L. A. P. Leme, S. Marschner, and S. Rusinkiewicz. Printing anisotropic appearance with magnetic flakes. *ACM Trans. Graph.*, 36(4), July 2017. DOI [10.1145/3072959.3073701](https://doi.org/10.1145/3072959.3073701). Citation on page 25.
- [331] T. Pereira, S. Rusinkiewicz, and W. Matusik. Computational light routing: 3d printed optical fibers for sensing and display. *ACM Trans. Graph.*, 33(3), June 2014. DOI [10.1145/2602140](https://doi.org/10.1145/2602140). Citation on page 3.
- [332] J. Pérez, M. A. Otaduy, and B. Thomaszewski. Computational design and automated fabrication of kirchhoff-plateau surfaces. *ACM Trans. Graph.*, 36(4), July 2017. DOI [10.1145/3072959.3073695](https://doi.org/10.1145/3072959.3073695). Citation on page 15.
- [333] M. Pharr, W. Jakob, and G. Humphreys. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, 2016. Citation on page 3.
- [334] M. Piovračić, D. I. W. Levin, D. M. Kaufman, and P. Didyk. Perception-aware modeling and fabrication of digital drawing tools. *ACM Trans. Graph.*, 37(4), July 2018. DOI [10.1145/3197517.3201322](https://doi.org/10.1145/3197517.3201322). Citation on pages 2 and 69.
- [335] P. Pjanic and R. D. Hersch. Color imaging and pattern hiding on a metallic substrate. *ACM Trans. Graph.*, 34(4), July 2015. DOI [10.1145/2766944](https://doi.org/10.1145/2766944). Citation on page 26.
- [336] E. Praun, H. Hoppe, M. Webb, and A. Finkelstein. Real-time hatching. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01*, pages 581–, New York, NY, USA, 2001. ACM. Citation on page 28.
- [337] R. Prévost, A. Jacobson, W. Jarosz, and O. Sorkine-Hornung. Large-scale painting of photographs by interactive optimization. *Computers & Graphics*, 55:108–117, Apr. 2016. DOI [10.1016/j.cag.2015.11.001](https://doi.org/10.1016/j.cag.2015.11.001). Citation on page 25.
- [338] G. Puglisi and S. Battiato. Artificial mosaic generation. In *Image and video-based artistic stylisation*, pages 189–209. Springer, 2013. Citation on page 22.

- [339] P. Rademacher. View-dependent geometry. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, page 439–446, USA, 1999. ACM Press/Addison-Wesley Publishing Co. DOI [10.1145/311535.311612](https://doi.org/10.1145/311535.311612). Citation on page 72.
- [340] P. Rademacher and G. Bishop. Multiple-center-of-projection images. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, page 199–206, New York, NY, USA, 1998. Association for Computing Machinery. DOI [10.1145/280814.280871](https://doi.org/10.1145/280814.280871). Citation on page 73.
- [341] V. S. Ramachandran. Perception of shape from shading. *Nature*, 331(6152):163–166, jan 1988. DOI [10.1038/331163a0](https://doi.org/10.1038/331163a0). Citation on page 34.
- [342] S. Rasool and A. Sourin. Tangible images. *SIGGRAPH Asia 2011 Sketches, SA'11*, 12 2011. DOI [10.1145/2077378.2077430](https://doi.org/10.1145/2077378.2077430). Citation on page 28.
- [343] P. Ré. On my drawings and paintings: An extension of the system of their classification. *Leonardo*, 14(2):106–113, 1981. Citation on page 20.
- [344] B. Reeves, C. Crawford, B. Jones, G. Rousseau, and J. Wright. Removable applicator assembly for applying a liquid layer, 1998. US Patent 5,808,645, Online at <https://www.google.com/patents/US5808645>. Citation on page 61.
- [345] C. Regg, S. Rusinkiewicz, W. Matusik, and M. Gross. Computational highlight holography. *ACM Trans. Graph.*, 29(6), Dec. 2010. DOI [10.1145/1882261.1866192](https://doi.org/10.1145/1882261.1866192). Citation on page 26.
- [346] A. Reichinger, S. Maierhofer, and W. Purgathofer. High-quality tactile paintings. *J. Comput. Cult. Herit.*, 4(2), Nov. 2011. DOI [10.1145/2037820.2037822](https://doi.org/10.1145/2037820.2037822). Citation on pages iv, xix, xx, 7, 16, 29, and 108.
- [347] A. Reichinger, M. Neumüller, F. Rist, S. Maierhofer, and W. Purgathofer. Computer-aided design of tactile models. In K. Miesenberger, A. Karshmer, P. Penaz, and W. Zagler, edi-

- tors, *Computers Helping People with Special Needs*, pages 497–504, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. Citation on page 14.
- [348] T. Reiner, N. Carr, R. Měch, O. Št'ava, C. Dachsbacher, and G. Miller. Dual-color mixing for fused deposition modeling printers. *Comput. Graph. Forum*, 33(2):479–486, May 2014. DOI [10.1111/cgf.12319](https://doi.org/10.1111/cgf.12319). Citation on page 15.
- [349] P. Rheingans and D. Ebert. Volume illustration: nonphotorealistic rendering of volume models. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):253–264, 2001. Citation on page 27.
- [350] S. Richard and S. Katie. *Alla prima: Everything i know about painting*. Stove Prairie Press, 2013. Citation on page 16.
- [351] A. Rivers, I. E. Moyer, and F. Durand. Position-correcting tools for 2d digital fabrication. *ACM Trans. Graph.*, 31(4), 2012. DOI [10.1145/2185520.2185584](https://doi.org/10.1145/2185520.2185584). Citation on page 2.
- [352] E. Rosales, J. Rodriguez, and A. SHEFFER. Surfacebrush: From virtual reality drawings to manifold surfaces. *ACM Trans. Graph.*, 38(4), July 2019. DOI [10.1145/3306346.3322970](https://doi.org/10.1145/3306346.3322970). Citation on page 2.
- [353] C. Rossl and L. Kobbelt. Line-art rendering of 3d-models. In *Proceedings the Eighth Pacific Conference on Computer Graphics and Applications*, pages 87–96, Oct 2000. Citation on page 71.
- [354] T. Roumen, J. Shigeyama, J. Rudolph, F. Grzelka, and P. Baudisch. Springfit: Joints and mounts that fabricate on any laser cutter. 10 2019. DOI [10.1145/3332165.3347930](https://doi.org/10.1145/3332165.3347930). Citation on page 14.
- [355] D. Rozin. Interactive art, mechanical mirrors. <http://www.smoothware.com/danny/>. Accessed: 2020-03-28. Citation on pages 5 and 25.
- [356] S. Rusinkiewicz. Estimating curvatures and their derivatives on triangle meshes. In *Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and*

- Transmission, 2004. 3DPVT 2004.*, pages 486–493, Sep. 2004. DOI [10.1109/TDPVT.2004.1335277](https://doi.org/10.1109/TDPVT.2004.1335277). Citation on page 71.
- [357] T. Saito and T. Takahashi. Comprehensible rendering of 3-d shapes. *SIGGRAPH Comput. Graph.*, 24(4):197–206, Sept. 1990. Citation on page 22.
- [358] K. Sakurai, Y. Dobashi, K. Iwasaki, and T. Nishita. Fabricating reflectors for displaying multiple images. *ACM Trans. Graph.*, 37(4), July 2018. DOI [10.1145/3197517.3201400](https://doi.org/10.1145/3197517.3201400). Citation on page 26.
- [359] M. P. Salisbury, S. E. Anderson, R. Barzel, and D. H. Salesin. Interactive pen-and-ink illustration. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '94*, page 101–108, New York, NY, USA, 1994. Association for Computing Machinery. DOI [10.1145/192161.192185](https://doi.org/10.1145/192161.192185). Citation on page 23.
- [360] M. P. Salisbury, M. T. Wong, J. F. Hughes, and D. H. Salesin. Orientable textures for image-based pen-and-ink illustration. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97*, page 401–406, USA, 1997. ACM Press/Addison-Wesley Publishing Co. DOI [10.1145/258734.258890](https://doi.org/10.1145/258734.258890). Citation on page 22.
- [361] T. Samsonov, S. Koshel, D. Walther, and B. Jenny. Automated placement of supplementary contour lines. *International Journal of Geographical Information Science*, 33:2072–2093, 05 2019. DOI [10.1080/13658816.2019.1610965](https://doi.org/10.1080/13658816.2019.1610965). Citation on page 71.
- [362] R. Sayeed and T. Howard. State of the art non-photorealistic rendering (npr) techniques. In *TPCG, 2006*. Citation on page 22.
- [363] S. Schlechtweg, T. Germer, and T. Strothotte. Renderbots—multi-agent systems for direct image generation. *Computer Graphics Forum*, 24(2):137–148, 2005. DOI [10.1111/j.1467-8659.2005.00838.x](https://doi.org/10.1111/j.1467-8659.2005.00838.x). Citation on page 22.

- [364] C. Schüller, D. Panozzo, A. Grundhöfer, H. Zimmer, E. Sorkine, and O. Sorkine-Hornung. Computational thermoforming. *ACM Trans. Graph.*, 35(4):43:1–43:9, July 2016. DOI [10.1145/2897824.2925914](https://doi.org/10.1145/2897824.2925914). Citation on pages 15 and 26.
- [365] A. Schulz, J. Xu, B. Zhu, C. Zheng, E. Grinspun, and W. Matusik. Interactive design space exploration and optimization for cad models. *ACM Transactions on Graphics*, 36(4), July 2017. Citation on page 3.
- [366] C. Schumacher, B. Bickel, J. Rys, S. Marschner, C. Daraio, and M. Gross. Microstructures to control elasticity in 3d printing. *ACM Trans. Graph.*, 34(4), July 2015. DOI [10.1145/2766926](https://doi.org/10.1145/2766926). Citation on page 3.
- [367] Y. Schwartzburg and M. Pauly. Fabrication-aware design with intersecting planar pieces. *Comput. Graph. Forum*, 32:317–326, 2013. Citation on page 3.
- [368] Y. Schwartzburg, R. Testuz, A. Tagliasacchi, and M. Pauly. High-contrast computational caustic design. *ACM Trans. Graph.*, 33(4), July 2014. DOI [10.1145/2601097.2601200](https://doi.org/10.1145/2601097.2601200). Citation on pages 25 and 71.
- [369] M. Schwarz, T. Isenberg, K. Mason, and S. Carpendale. Modeling with rendering primitives: An interactive non-photorealistic canvas. In *Proceedings of the 5th International Symposium on Non-Photorealistic Animation and Rendering*, NPAR '07, page 15–22, New York, NY, USA, 2007. Association for Computing Machinery. DOI [10.1145/1274871.1274874](https://doi.org/10.1145/1274871.1274874). Citation on page 23.
- [370] Scipy.org. `numpy.linalg.eig`, 2018. Online at <https://het.as.utexas.edu/HET/Software/Numpy/reference/generated/numpy.linalg.eig.html>. Citation on page 129.
- [371] A. Secord. Weighted voronoi stippling. In *Proceedings of the 2Nd International Symposium on Non-photorealistic Animation and Rendering*, NPAR '02, pages 37–43, New York, NY, USA, 2002. ACM. Citation on page 22.

- [372] A. Secord, J. Lu, A. Finkelstein, M. Singh, and A. Nealen. Perceptual models of viewpoint preference. *ACM Trans. Graph.*, 30(5):109:1–109:12, Oct. 2011. Citation on pages 72, 112, 113, 118, and 119.
- [373] A. Secord, J. Lu, A. Finkelstein, M. Singh, and A. Nealen. Perceptual models of viewpoint preference. *ACM Trans. Graph.*, 30(5):109:1–109:12, Oct. 2011. Citation on page 82.
- [374] J. Seims. Putting the artist in the loop. In *COMG*, 1999. Citation on page 109.
- [375] A. Semmo, D. Limberger, J. E. Kyprianidis, and J. Döllner. Image stylization by interactive oil paint filtering. *Comput. Graph.*, 55(C):157–171, Apr. 2016. DOI [10.1016/j.cag.2015.12.001](https://doi.org/10.1016/j.cag.2015.12.001). Citation on page 22.
- [376] A. Shamir, B. Bickel, and W. Matusik, ACM SIGGRAPH. Computational tools for 3d printing, course program. <http://computational-fabrication.com/2016/>, 2016. Accessed: 2020-02-23. Citation on page 15.
- [377] T. Shao, D. Li, Y. Rong, C. Zheng, and K. Zhou. Dynamic furniture modeling through assembly instructions. *ACM Trans. Graph.*, 35(6), Nov. 2016. DOI [10.1145/2980179.2982416](https://doi.org/10.1145/2980179.2982416). Citation on page 107.
- [378] L. Shi, V. Babaei, C. Kim, M. Foshey, Y. Hu, P. Sitthi-Amorn, S. Rusinkiewicz, and W. Matusik. Deep multispectral painting reproduction via multi-layer, custom-ink printing. *ACM Trans. Graph.*, 37(6), Dec. 2018. DOI [10.1145/3272127.3275057](https://doi.org/10.1145/3272127.3275057). Citation on page 25.
- [379] R. Shilkrot and D. M. Escrivá. *Mastering OpenCV 4: A comprehensive guide to building computer vision and image processing applications with C++*. Packt Publishing Ltd, 2018. Citation on page 116.
- [380] R. Shilkrot, P. Maes, J. A. Paradiso, and A. Zoran. Augmented airbrush for computer aided painting (cap). *ACM Trans. Graph.*, 34(2), Mar. 2015. DOI [10.1145/2699649](https://doi.org/10.1145/2699649). Citation on page 3.

- [381] M. Shugrina, A. Shamir, and W. Matusik. Fab forms: Customizable objects for fabrication with validity and geometry caching. *ACM Trans. Graph.*, 34(4), July 2015. DOI [10.1145/2766994](https://doi.org/10.1145/2766994). Citation on page 15.
- [382] O. Sidi, O. van Kaick, Y. Kleiman, H. Zhang, and D. Cohen-Or. Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. *ACM Trans. Graph.*, 30(6):1–10, Dec. 2011. DOI [10.1145/2070781.2024160](https://doi.org/10.1145/2070781.2024160). Citation on page 73.
- [383] P. Sitthi-Amorn, J. E. Ramos, Y. Wangy, J. Kwan, J. Lan, W. Wang, and W. Matusik. Multifab: A machine vision assisted platform for multi-material 3d printing. *ACM Trans. Graph.*, 34(4), July 2015. DOI [10.1145/2766962](https://doi.org/10.1145/2766962). Citation on page 2.
- [384] R. Sodhi, I. Poupyrev, M. Glisson, and A. Israr. Aireal: Interactive tactile experiences in free air. *ACM Trans. Graph.*, 32(4), July 2013. DOI [10.1145/2461912.2462007](https://doi.org/10.1145/2461912.2462007). Citation on page 28.
- [385] L. Somm, D. Hahn, N. Kumar, and S. Coros. Expanding foam as the material for fabrication, prototyping and experimental assessment of low-cost soft robots with embedded sensing. *IEEE Robotics and Automation Letters*, 4(2):761–768, April 2019. DOI [10.1109/LRA.2019.2893416](https://doi.org/10.1109/LRA.2019.2893416). Citation on page 15.
- [386] W. Song, A. Belyaev, and H.-P. Seidel. Automatic generation of bas-reliefs from 3d shapes. In *Proceedings of the IEEE International Conference on Shape Modeling and Applications 2007, SMI '07*, page 211–214, USA, 2007. IEEE Computer Society. DOI [10.1109/SMI.2007.9](https://doi.org/10.1109/SMI.2007.9). Citation on page 108.
- [387] H. Sonnet, S. Carpendale, and T. Strothotte. Integrating expanding annotations with a 3d explosion probe. In *Proceedings of the Working Conference on Advanced Visual Interfaces, AVI '04*, page 63–70, New York, NY, USA, 2004. Association for Computing Machinery. DOI [10.1145/989863.989871](https://doi.org/10.1145/989863.989871). Citation on page 107.

- [388] A. Sourin. Functionally based virtual computer art. In *Proceedings of the 2001 Symposium on Interactive 3D Graphics*, I3D '01, page 77–84, New York, NY, USA, 2001. Association for Computing Machinery. DOI [10.1145/364338.364369](https://doi.org/10.1145/364338.364369). Citation on page 108.
- [389] I. Spence. William playfair and the psychology of graphs. 08 2006. Citation on page 21.
- [390] V. Srinivasan, E. Akleman, and Jianer Chen. Interactive construction of multi-segment curved handles. In *10th Pacific Conference on Computer Graphics and Applications, 2002. Proceedings.*, pages 429–430, Oct 2002. DOI [10.1109/PCCGA.2002.1167888](https://doi.org/10.1109/PCCGA.2002.1167888). Citation on page 123.
- [391] A. Stangl, J. Kim, and T. Yeh. 3d printed tactile picture books for children with visual impairments: A design probe. In *Proceedings of the 2014 Conference on Interaction Design and Children, IDC '14*, pages 321–324, New York, NY, USA, 2014. ACM. Citation on pages 19, 28, 66, and 108.
- [392] A. J. Stangl. *Tactile Media Consumption and Production for and by People Who Are Blind and Visually Impaired: A Design Research Investigation*. PhD thesis, University of Colorado at Boulder, 2019. Citation on pages 28 and 66.
- [393] G. Stoychev, M. J. Razavi, X. Wang, and L. Ionov. 4d origami by smart embroidery. *Macromolecular Rapid Communications*, 38(18):1700213, 2017. DOI [10.1002/marc.201700213](https://doi.org/10.1002/marc.201700213). Citation on page 15.
- [394] T. Strothotte and S. Schlechtweg. *Non-Photorealistic Computer Graphics: Modeling, Rendering, and Animation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002. Citation on page 22.
- [395] N. Sudarsanam, C. Grimm, and K. Singh. Interactive manipulation of projections with a curved perspective. In *Eurographics*, 2005. Citation on page 27.
- [396] H. Sundar, D. Silver, N. Gagvani, and S. Dickinson. Skeleton based shape matching and retrieval. In *2003 Shape Modeling International.*, pages 130–139, May 2003. DOI [10.1109/SMI.2003.1199609](https://doi.org/10.1109/SMI.2003.1199609). Citation on page 70.

- [397] R. Suzuki, A. Stangl, M. D. Gross, and T. Yeh. Fluxmarker: Enhancing tactile graphics with dynamic tactile markers. In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility, ASSETS '17*, pages 190–199. ACM, 2017. Citation on pages 28, 100, and 107.
- [398] S. Swaminathan, T. Roumen, R. Kovacs, D. Stangl, S. Mueller, and P. Baudisch. Linespace: A sensemaking platform for the blind. *Informatik-Spektrum*, 40, 09 2017. DOI [10.1007/s00287-017-1070-1](https://doi.org/10.1007/s00287-017-1070-1). Citation on page 28.
- [399] A. Tagliasacchi, I. Alhashim, M. Olson, and H. Zhang. Mean curvature skeletons. *Comput. Graph. Forum*, 31(5):1735–1744, Aug. 2012. Citation on pages 71, 83, 84, and 98.
- [400] A. Tagliasacchi, T. Delame, M. Spagnuolo, N. Amenta, and A. Telea. 3d skeletons: A state-of-the-art report. *Computer Graphics Forum*, 35, 05 2016. Citation on pages 100 and 104.
- [401] A. Tagliasacchi, H. Zhang, and D. Cohen-Or. Curve skeleton extraction from incomplete point cloud. In *ACM Transactions on Graphics (TOG)*, volume 28, page 71. ACM, 2009. Citation on page 71.
- [402] M. Takezawa, T. Imai, K. Shida, and T. Maekawa. Fabrication of freeform objects by principal strips. *ACM Trans. Graph.*, 35(6), Nov. 2016. DOI [10.1145/2980179.2982406](https://doi.org/10.1145/2980179.2982406). Citation on page 71.
- [403] J. Tan, J.-M. Lien, and Y. Gingold. Decomposing images into layers via RGB-space geometry. *ACM Trans. Graph.*, 36(1):7:1–7:14, Nov. 2016. DOI [10.1145/2988229](https://doi.org/10.1145/2988229). Citation on page 105.
- [404] M. Tatzgern, D. Kalkofen, and D. Schmalstieg. Compact explosion diagrams. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering, NPAR '10*, pages 17–26, New York, NY, USA, 2010. ACM. Citation on pages 28 and 107.

- [405] T. Terai, S. Mizuno, and M. Okada. A Color Decomposition Method for Preserving Ukiyo-e Woodblocks. In M. Alexa and E. Galin, editors, *Eurographics 2004 - Short Presentations*. Eurographics Association, 2004. DOI [10.2312/egs.20041006](https://doi.org/10.2312/egs.20041006). Citation on page 26.
- [406] R. Testuz, Y. Schwartzburg, and M. Pauly. Automatic Generation of Constructable Brick Sculptures. In M.-A. Otaduy and O. Sorkine, editors, *Eurographics 2013 - Short Papers*. The Eurographics Association, 2013. DOI [10.2312/conf/EG2013/short/081-084](https://doi.org/10.2312/conf/EG2013/short/081-084). Citation on page 14.
- [407] The CGAL Project. *CGAL User and Reference Manual*. CGAL Editorial Board, 4.14 edition, 2019. Online at <https://doc.cgal.org/4.14/Manual/packages.html>. Citation on pages 71 and 123.
- [408] H. Theisel. *Vector Field Curvature and Applications*. PhD thesis, Arizona State University, 1995. Citation on page 71.
- [409] A. Theurel, A. Witt, P. Claudet, Y. Hatwell, and E. Gentaz. Tactile picture recognition by early blind children: The effect of illustration technique. *Journal of experimental psychology. Applied*, 19:233–40, 09 2013. Citation on page 20.
- [410] L. Thompson and E. Chronicle. Beyond visual conventions: Rethinking the design of tactile diagrams. *British Journal of Visual Impairment*, 24:76–82, 05 2006. Citation on page 20.
- [411] L. J. Thompson, E. P. Chronicle, and A. F. Collins. Enhancing 2-d tactile picture design from knowledge of 3-d haptic object recognition. *European Psychologist - EUR PSYCHOL*, 11:110–118, 01 2006. Citation on pages 29, 70, and 79.
- [412] C. Torres, W. Li, and E. Paulos. Proxyprint: Supporting crafting practice through physical computational proxies. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, DIS '16, page 158–169, New York, NY, USA, 2016. Association for Computing Machinery. DOI [10.1145/2901790.2901828](https://doi.org/10.1145/2901790.2901828). Citation on page 15.

- [413] P. Tresset and O. Deussen. Artistically skilled embodied agents. In *AISB 2014, 1st - 4th April 2014, Goldsmiths, University of London, UK*, 2014. Citation on page 25.
- [414] P. Tresset and F. Fol Leymarie. Portrait drawing by paul the robot. *Computers & Graphics*, 37(5):348 – 363, 2013. DOI <https://doi.org/10.1016/j.cag.2013.01.012>. Citation on page 8.
- [415] P. A. Tresset and F. F. Leymarie. Sketches by Paul the Robot. In *Proceedings of the Eighth Annual Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*, CAe '12, pages 17–24, Aire-la-Ville, Switzerland, Switzerland, 2012. Eurographics Association. Online at <http://dl.acm.org/citation.cfm?id=2328888.2328892>. Citation on pages iii, xix, 4, 6, 9, 25, and 34.
- [416] C. Tymms, E. P. Gardner, and D. Zorin. A quantitative perceptual model for tactile roughness. *ACM Trans. Graph.*, 37(5), Oct. 2018. DOI [10.1145/3186267](https://doi.org/10.1145/3186267). Citation on page 69.
- [417] U. Umapathi, H.-T. Chen, S. Mueller, L. Wall, A. Seufert, and P. Baudisch. Laserstacker: Fabricating 3d objects by laser cutting and welding. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, UIST '15, page 575–582, New York, NY, USA, 2015. Association for Computing Machinery. DOI [10.1145/2807442.2807512](https://doi.org/10.1145/2807442.2807512). Citation on page 15.
- [418] N. Umetani, Y. Koyama, R. Schmidt, and T. Igarashi. Pteromys: Interactive design and optimization of free-formed free-flight model airplanes. *ACM Trans. Graph.*, 33(4):65:1–65:10, July 2014. DOI [10.1145/2601097.2601129](https://doi.org/10.1145/2601097.2601129). Citation on page 14.
- [419] N. Umetani, J. Mitani, T. Ibaraki, and T. Igarashi. Designing custom-made metallophone with concurrent eigenanalysis. *NIME++*, 2010. Citation on page 14.
- [420] N. Umetani, A. Panotopoulou, R. Schmidt, and E. Whiting. Printone: Interactive resonance simulation for free-form print-wind instrument design. *ACM Trans. Graph.*, 35(6):184:1–184:14, Nov. 2016. DOI [10.1145/2980179.2980250](https://doi.org/10.1145/2980179.2980250). Citation on page 120.

- [421] P. Urban, T. M. Tanksale, A. Brunton, B. M. Vu, and S. Nakauchi. Redefining A in RGBA: towards a standard for graphical 3d printing. *ACM Trans. Graph.*, 38(3):21:1–21:14, 2019. DOI [10.1145/3319910](https://doi.org/10.1145/3319910). Citation on page 69.
- [422] F. Usai, M. Livesu, E. Puppo, M. Tarini, and R. Scateni. Extraction of the quad layout of a triangle mesh guided by its curve skeleton. *ACM Trans. Graph.*, 35(1), Dec. 2016. DOI [10.1145/2809785](https://doi.org/10.1145/2809785). Citation on page 123.
- [423] T. Van Laerhoven, J. Liesenborgs, and F. Van Reeth. Real-time watercolor painting on a distributed paper model. In *Proceedings Computer Graphics International, 2004.*, pages 640–643, 2004. Citation on page 24.
- [424] J. Vanek, J. A. G. Galicia, and B. Benes. Clever support: Efficient support structure generation for digital fabrication. *Comput. Graph. Forum*, 33:117–125, 2014. Citation on page 3.
- [425] S. R. Vantaram and E. Saber. Survey of contemporary trends in color image segmentation. *Journal of Electronic Imaging*, 21(4):040901–1–040901–28, 2012. DOI [10.1117/1.JEI.21.4.040901](https://doi.org/10.1117/1.JEI.21.4.040901). Citation on pages 36 and 104.
- [426] P.-P. Vázquez, M. Feixas, M. Sbert, and W. Heidrich. Viewpoint selection using viewpoint entropy. In *VMV*, 2001. Citation on page 72.
- [427] J. Vekhter, J. Zhuo, L. F. G. Fandino, Q. Huang, and E. Vouga. Weaving geodesic foliations. *ACM Trans. Graph.*, 38(4), July 2019. DOI [10.1145/3306346.3323043](https://doi.org/10.1145/3306346.3323043). Citation on page 15.
- [428] I. Viola, M. C. Sousa, D. Ebert, B. Andrews, B. Gooch, and C. Tietjen. Illustrative Visualization for Medicine and Science. In N. Magnenat-Thalmann and K. Bühler, editors, *Eurographics 2006: Tutorials*. The Eurographics Association, 2006. DOI [10.2312/egt.20061068](https://doi.org/10.2312/egt.20061068). Citation on page 27.

- [429] D. V. Vranic, D. Saupe, and J. Richter. Tools for 3d-object retrieval: Karhunen-loeve transform and spherical harmonics. In *2001 IEEE Fourth Workshop on Multimedia Signal Processing (Cat. No. 01TH8564)*, pages 293–298. IEEE, 2001. Citation on page 72.
- [430] D. V. Vranic, D. Saupe, and J. Richter. Tools for 3d-object retrieval: Karhunen-loeve transform and spherical harmonics. In *2001 IEEE Fourth Workshop on Multimedia Signal Processing (Cat. No.01TH8564)*, pages 293–298, Oct 2001. Citation on page 83.
- [431] M. Waldner, A. Karimov, and E. Gröller. Exploring visual prominence of multi-channel highlighting in visualizations. In *Proceedings of the 33rd Spring Conference on Computer Graphics, SCCG '17*, New York, NY, USA, 2017. Association for Computing Machinery. DOI [10.1145/3154353.3154369](https://doi.org/10.1145/3154353.3154369). Citation on page 108.
- [432] B. Walsh. Texture, buttons, sound and code: Modal preference and the formation of expert identities. In *Proceedings of the 7th Annual Conference on Creativity and Fabrication in Education, FabLearn '17*, pages 19:1–19:4, New York, NY, USA, 2017. ACM. Citation on pages 19 and 28.
- [433] C. C. L. Wang, Yunbo Zhang, and Hoi Sheung. From designing products to fabricating them from planar materials. *IEEE Computer Graphics and Applications*, 30(6):74–85, 2010. Citation on page 15.
- [434] L. Wang and E. Whiting. Buoyancy optimization for computational fabrication. *Comput. Graph. Forum*, 35(2):49–58, May 2016. DOI [10.1111/cgf.12810](https://doi.org/10.1111/cgf.12810). Citation on pages 14 and 15.
- [435] L. Wang, Y. Zhao, K. Mueller, and A. Kaufman. The magic volume lens: An interactive focus+ context technique for volume rendering. In *VIS 05. IEEE Visualization, 2005.*, pages 367–374. IEEE, 2005. Citation on page 108.
- [436] M. Wang, B. Wang, Y. Fei, K. Qian, W. Wang, J. Chen, and J. H. Yong. Towards photo watercolorization with artistic verisimilitude. *IEEE Transactions on Visualization and*

Computer Graphics, 20(10):1451–1460, Oct 2014. DOI [10.1109/TVCG.2014.2303984](https://doi.org/10.1109/TVCG.2014.2303984).
Citation on page 22.

- [437] Y. Wang, M. Gong, T. Wang, D. Cohen-Or, H. Zhang, and B. Chen. Projective analysis for 3d shape segmentation. *ACM Trans. Graph.*, 32(6), Nov. 2013. DOI [10.1145/2508363.2508393](https://doi.org/10.1145/2508363.2508393). Citation on page 73.
- [438] Z. Wang and B. Li. A bayesian approach to automated creation of tactile facial images. *IEEE transactions on multimedia*, 12(4):233–246, 2010. Citation on page 28.
- [439] Z. Wang, P. Song, and M. Pauly. Desia: A general framework for designing interlocking assemblies. *ACM Trans. Graph.*, 37(6), Dec. 2018. DOI [10.1145/3272127.3275034](https://doi.org/10.1145/3272127.3275034). Citation on page 3.
- [440] Z. Wang, X. Xu, and B. Li. Enabling seamless access to digital graphical contents for visually impaired individuals via semantic-aware processing. *EURASIP Journal on Image and Video Processing*, 2007:1–14, 2007. Citation on page 29.
- [441] J. D. WATSON and F. H. C. CRICK. Molecular Structure of Nucleic Acids: A Structure for Deoxyribose Nucleic Acid. *Nature*, 171(4356):737–738, Apr. 1953. DOI [10.1038/171737a0](https://doi.org/10.1038/171737a0). Citation on page 21.
- [442] K. Watt. Method and apparatus for applying thin coatings of fluid droplets, 1994. US Patent 5,314,119, Online at <https://www.google.com/patents/US5314119>. Citation on page 61.
- [443] T. P. Way and K. E. Barner. Automatic visual to tactile translation. *IEEE Transactions on Rehabilitation Engineering*, 5(1):95–105, March 1997. Citation on pages 28, 29, and 84.
- [444] D. Webb. *Painting in watercolor: The indispensable guide*. Firefly Books, 2016. Citation on page 17.
- [445] L. Wenyin and D. Dori. *A survey of non-thinning based vectorization methods*, pages 230–241. Springer Berlin Heidelberg, Berlin, Heidelberg, 1998. DOI [10.1007/BFb0033241](https://doi.org/10.1007/BFb0033241). Citation on page 104.

- [446] T. Weyrich, J. Deng, C. Barnes, S. Rusinkiewicz, and A. Finkelstein. Digital bas-relief from 3D scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 26(3):32, 07 2007. Citation on page 108.
- [447] T. Weyrich, P. Peers, W. Matusik, and S. Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. In *ACM SIGGRAPH 2009 Papers*, SIGGRAPH '09, New York, NY, USA, 2009. Association for Computing Machinery. DOI [10.1145/1576246.1531338](https://doi.org/10.1145/1576246.1531338). Citation on pages 8 and 25.
- [448] O. Wiki. Depth test — opengl wiki., 2018. Online at http://www.khronos.org/opengl/wiki/opengl/index.php?title=Depth_Test/&oldid/=14331. Citation on page 113.
- [449] M. Wilcox. *Blue and yellow don't make green*. North light books, 1994. Citation on page 48.
- [450] B. Wilson and K.-L. Ma. Rendering complexity in computer-generated pen-and-ink illustrations. In *Proceedings of the 3rd International Symposium on Non-Photorealistic Animation and Rendering*, NPAR '04, page 129–137, New York, NY, USA, 2004. Association for Computing Machinery. DOI [10.1145/987657.987674](https://doi.org/10.1145/987657.987674). Citation on page 22.
- [451] G. Winkenbach and D. H. Salesin. Computer-generated pen-and-ink illustration. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '94, page 91–100, New York, NY, USA, 1994. Association for Computing Machinery. DOI [10.1145/192161.192184](https://doi.org/10.1145/192161.192184). Citation on pages v, xxi, and 23.
- [452] G. Winkenbach and D. H. Salesin. Rendering parametric surfaces in pen and ink. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, pages 469–476, New York, NY, USA, 1996. ACM. Citation on page 28.

- [453] F. J. Wong and S. Takahashi. A graph-based approach to continuous line illustrations with variable levels of detail. *Computer Graphics Forum*, 30(7):1931–1939, 2011. DOI [10.1111/j.1467-8659.2011.02040.x](https://doi.org/10.1111/j.1467-8659.2011.02040.x). Citation on page 22.
- [454] J. Xu and C. S. Kaplan. Image-guided maze construction. *ACM Trans. Graph.*, 26(3):29–es, July 2007. DOI [10.1145/1276377.1276414](https://doi.org/10.1145/1276377.1276414). Citation on page 22.
- [455] J. Xu and C. S. Kaplan. Artistic thresholding. In *Proceedings of the 6th International Symposium on Non-Photorealistic Animation and Rendering*, NPAR '08, page 39–47, New York, NY, USA, 2008. Association for Computing Machinery. DOI [10.1145/1377980.1377990](https://doi.org/10.1145/1377980.1377990). Citation on page 22.
- [456] J. Xu, C. S. Kaplan, and X. Mi. Computer-generated papercutting. In *15th Pacific Conference on Computer Graphics and Applications (PG'07)*, pages 343–350, Oct 2007. DOI [10.1109/PG.2007.10](https://doi.org/10.1109/PG.2007.10). Citation on pages 5, 14, and 25.
- [457] K. Xu, V. G. Kim, Q. Huang, and E. Kalogerakis. Data-driven shape analysis and processing, 2015. Citation on pages 100 and 104.
- [458] S. Xue, X. Chen, J. Dorsey, and H. Rushmeier. Printed patterns for enhanced shape perception of papercraft models. *Computer Graphics Forum*, 29(2):625–634, 2010. DOI [10.1111/j.1467-8659.2009.01632.x](https://doi.org/10.1111/j.1467-8659.2009.01632.x). Citation on pages 69 and 71.
- [459] U. Yaman and M. Dolen. A gradient-based morphological method to produce planar curve offsets. *The International Journal of Advanced Manufacturing Technology*, 80(1):255–274, Sep 2015. DOI [10.1007/s00170-015-6971-z](https://doi.org/10.1007/s00170-015-6971-z). Citation on page 104.
- [460] X. Yan, C. Rao, L. Lu, A. Sharf, H. Zhao, and B. Chen. Strong 3d printing by tpms injection. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2019. DOI [10.1109/TVCG.2019.2914044](https://doi.org/10.1109/TVCG.2019.2914044). Citation on page 3.
- [461] Yan Lu, J. H. M. Lam, and Y. Yam. Preliminary study on vision-based pen-and-ink drawing by a robotic manipulator. In *2009 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, pages 578–583, 2009. Citation on page 25.

- [462] H.-L. Yang and C.-K. Yang. A non-photorealistic rendering of seurat’s pointillism. In *ISVC*, 2006. Citation on page 23.
- [463] T. Yang, J. Chang, M. C. Lin, R. R. Martin, J. J. Zhang, and S.-M. Hu. A unified particle system framework for multi-phase, multi-material visual simulations. *ACM Trans. Graph.*, 36(6), Nov. 2017. DOI [10.1145/3130800.3130882](https://doi.org/10.1145/3130800.3130882). Citation on page 8.
- [464] J. Yaniv, Y. Newman, and A. Shamir. The face of art: Landmark detection and geometric style in portraits. *ACM Trans. Graph.*, 38(4), July 2019. DOI [10.1145/3306346.3322984](https://doi.org/10.1145/3306346.3322984). Citation on page 109.
- [465] F. Yao and G. Shao. Painting brush control techniques in chinese painting robot. In *ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005.*, pages 462–467, Aug 2005. DOI [10.1109/ROMAN.2005.1513822](https://doi.org/10.1109/ROMAN.2005.1513822). Citation on page 25.
- [466] J. Yau, S. S. Kim, P. Thakur, and S. Bensaïa. Feeling form: The neural basis of haptic shape perception. *Journal of neurophysiology*, 115:jn.00598.2015, 11 2015. Citation on pages 29, 71, and 79.
- [467] C.-K. Yeh, H.-K. Chu, M.-J. Chang, and T.-Y. Lee. Anamorphic image generation using hybrid texture synthesis. *Journal of Information Science and Engineering*, 34(1):123–134, 2018. Citation on page 22.
- [468] F. Yu, K. Liu, Y. Zhang, C. Zhu, and K. Xu. Partnet: A recursive part decomposition network for fine-grained and hierarchical shape segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9491–9500, 2019. Citation on pages vi and 73.
- [469] J. Yu and L. McMillan. A framework for multiperspective rendering. In *Proceedings of the Fifteenth Eurographics Conference on Rendering Techniques, EGSR’04*, pages 61–68, Aire-la-Ville, Switzerland, Switzerland, 2004. Eurographics Association. Citation on page 28.

- [470] J. Yu, L. McMillan, and P. Sturm. Multiperspective modeling, rendering, and imaging. In *ACM SIGGRAPH ASIA 2008 Courses*, SIGGRAPH Asia '08, pages 14:1–14:36, New York, NY, USA, 2008. ACM. Citation on page 73.
- [471] Y. Yue, K. Iwasaki, B.-Y. Chen, Y. Dobashi, and T. Nishita. Pixel art with refracted light by rearrangeable sticks. *Computer Graphics Forum*, 31(2pt3):575–582, 2012. DOI [10.1111/j.1467-8659.2012.03036.x](https://doi.org/10.1111/j.1467-8659.2012.03036.x). Citation on pages 5 and 25.
- [472] Y. Yue, K. Iwasaki, B.-Y. Chen, Y. Dobashi, and T. Nishita. Poisson-based continuous surface generation for goal-based caustics. *ACM Trans. Graph.*, 33(3), June 2014. DOI [10.1145/2580946](https://doi.org/10.1145/2580946). Citation on page 25.
- [473] J. Zander, T. Isenberg, S. Schlechtweg, and T. Strothotte. High quality hatching. *Computer Graphics Forum*, 23(3):421–430, 2004. Citation on page 28.
- [474] K. Zeng, M. Zhao, C. Xiong, and S.-C. Zhu. From image parsing to painterly rendering. *ACM Trans. Graph.*, 29(1), Dec. 2009. DOI [10.1145/1640443.1640445](https://doi.org/10.1145/1640443.1640445). Citation on page 35.
- [475] E. Zhang, K. Mischaikow, and G. Turk. Vector field design on surfaces. *ACM Trans. Graph.*, 25(4):1294–1326, Oct. 2006. DOI [10.1145/1183287.1183290](https://doi.org/10.1145/1183287.1183290). Citation on page 104.
- [476] T. Zhang, B. S. Duerstock, and J. P. Wachs. Multimodal perception of histological images for persons who are blind or visually impaired. *ACM Transactions on Accessible Computing (TACCESS)*, 9(3):1–27, 2017. Citation on page 28.
- [477] X. Zhang, X. Le, A. Panotopoulou, E. Whiting, and C. C. L. Wang. Perceptual models of preference in 3d printing direction. *ACM Trans. Graph.*, 34(6):215:1–215:12, Oct. 2015. Citation on page 112.
- [478] X. Zhang, R. Wang, X. Jiang, W. Wang, and W. Gao. Spatially variant defocus blur map estimation and deblurring from a single image. *Journal of Visual Communication and*

- Image Representation*, 35:257 – 264, 2016. DOI <https://doi.org/10.1016/j.jvcir.2016.01.002>. Citation on pages v, 34, 40, 46, and 104.
- [479] Y. Zhang, C. Yin, C. Zheng, and K. Zhou. Computational hydrographic printing. *ACM Trans. Graph.*, 34(4):131:1–131:11, July 2015. DOI [10.1145/2766932](https://doi.org/10.1145/2766932). Citation on page 26.
- [480] H. Zhao, L. Lu, Y. Wei, D. Lischinski, A. Sharf, D. Cohen-Or, and B. Chen. Printed perforated lampshades for continuous projective images. *ACM Trans. Graph.*, 35(5), June 2016. DOI [10.1145/2907049](https://doi.org/10.1145/2907049). Citation on page 25.
- [481] H. Zhao, J. Wang, X. Ren, J. Li, Y.-L. Yang, and X. Jin. Personalized food printing for portrait images. *CAD/Graphics 2017*, 70:188–197, Feb. 2018. DOI [10.1016/j.cag.2017.07.012](https://doi.org/10.1016/j.cag.2017.07.012). Citation on pages 5 and 25.
- [482] H. Zhao, H. Zhang, S. Xin, Y. Deng, C. Tu, W. Wang, D. Cohen-Or, and B. Chen. Dscarver: Decompose-and-spiral-carve for subtractive manufacturing. *ACM Transactions on Graphics (Special Issue of SIGGRAPH)*, 37(4):Article No. 137, 2018. Citation on page 15.
- [483] M. Zhao and S.-C. Zhu. Sisley the abstract painter. In *NPAR '10: Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, pages 99–107, New York, NY, USA, 2010. ACM. DOI <http://doi.acm.org/10.1145/1809939.1809951>. Citation on page 41.
- [484] Y. Zhou, K. Yin, H. Huang, H. Zhang, M. Gong, and D. Cohen-Or. Generalized cylinder decomposition. *ACM Trans. Graph.*, 34(6):171:1–171:14, Oct. 2015. Citation on pages 73, 81, 100, and 104.
- [485] S. Zhuo and T. Sim. Defocus map estimation from a single image. *Pattern Recogn.*, 44(9):1852–1858, Sept. 2011. DOI [10.1016/j.patcog.2011.03.009](https://doi.org/10.1016/j.patcog.2011.03.009). Citation on pages 34 and 104.

- [486] L. Zi, X. Cong, and Y. Peng. An evolution perception mechanism for organizing and displaying 3d shape collections. In J. Yang, Q. Hu, M. Cheng, L. Wang, Q. Liu, X. Bai, and D. Meng, editors, *Computer Vision - Second CCF Chinese Conference, CCCV 2017, Tianjin, China, October 11-14, 2017, Proceedings, Part I*, volume 771 of *Communications in Computer and Information Science*, pages 3–13. Springer, 2017. DOI [10.1007/978-981-10-7299-4_1](https://doi.org/10.1007/978-981-10-7299-4_1). Citation on page 72.
- [487] A. Zoran. Hybrid craft: Showcase of physical and digital integration of design and craft skills. In *ACM SIGGRAPH Art Gallery, SIGGRAPH '15*, page 384–398, New York, NY, USA, 2015. Association for Computing Machinery. DOI [10.1145/2810185.2810187](https://doi.org/10.1145/2810185.2810187). Citation on page 3.