Perception-Aware Computational Fabrication

Increasing The Apparent Gamut of Digital Fabrication

Doctoral Dissertation submitted to the
Faculty of Informatics of the Università della Svizzera Italiana
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

presented by
Michal Piovarči

under the supervision of
Piotr Didyk

July 2020

Dissertation Committee

Piotr Didyk	Università della Svizzera Italiana, Switzerland
Kai Hormann	Università della Svizzera Italiana, Switzerland
Evanthia Papadopoulou	Università della Svizzera Italiana, Switzerland
David I.W. Levin	University of Toronto, Canada
Bernd Bickel	Institute of Science and Technology, Austria
Fabio Pellacini	Sapienza University of Rome, Italy

Dissertation accepted on 1 July 2020

Research Advisor	PhD Program Director
Piotr Didyk	**Prof. Walter Binder / Prof. Silvia Santini**

I certify that except where due acknowledgement has been given, the work presented in this thesis is that of the author alone; the work has not been submitted previously, in whole or in part, to qualify for any other academic award; and the content of the thesis is the result of work which has been carried out since the official commencement date of the approved research program.

Michal Piovarči
Lugano, 1 July 2020

To my beloved

Abstract

Haptic and visual feedback are important for assessing objects' quality and affordance. One of the benefits of additive manufacturing is that it enables the creation of objects with personalized tactile and visual properties. This personalization is realized by the ability to deposit functionally graded materials at microscopic resolution. However, faithfully reproducing real-world objects on a 3D printer is a challenging endeavor. A large number of available materials and freedom in material deposition make exploring the space of printable objects difficult. Furthermore, current 3D printers can perfectly capture only a small amount of objects from the real world which makes high-quality reproductions challenging. Interestingly, similar to the manufacturing hardware, our senses of touch and sight have inborn limitations given by biological constraints.

In this work, we propose that it is possible to leverage the limitations of human perception to increase the *apparent* gamut of a 3D printer by combining numerical optimization with perceptual insights. Instead of optimizing for exact replicas, we search for perceptually equivalent solutions. This not only simplifies the optimization but also achieves prints that better resemble the target behavior. To guide us towards the desired behavior, we design perceptual error metrics. Recovering such a metric requires conducting costly experiments. We tackle this problem by proposing a likelihood-based optimization that automatically recovers a metric that relates perception with physical properties. To minimize fabrication during the optimization we map new designs into perception via numerical models. As with many complex design tasks modeling the governing physics is either computationally expensive or we lack predictive models. We address this issue by applying perception-aware coarsening that focuses the computation towards perceptually relevant phenomena. Additionally, we propose a data-driven fabrication-in-the-loop model that implicitly handles the fabrication constraints. We demonstrate the capabilities of our approach in the contexts of haptic and appearance reproduction. We show its applications to designing objects with prescribed compliance, and mimicking the haptics of drawing tools. Furthermore, we propose a system for manufacturing objects with spatially-varying gloss.

Acknowledgements

First of all, I would like to thank my advisor Piotr Didyk for taking me on the awesome journey called doctoral studies. I learned from you how to approach research from identifying interesting research questions, through rigorously searching for possible solutions, to the writing of well motivated, easy to understand scientific papers, and presenting the results to a wider audience in an engaging yet informative way. A unique part of your supervision was the encouraging personal touch that put me into the correct mindset for wanting to push each project to a successful finish. Finally, your eye for details showed me that going the extra mile pays off in not only having results but having results one can be proud of. Thanks for shaping me into the researcher that I am today.

I would also like to thanks David I. W. Levin and Danny M. Kaufman for our discussions on numerical simulation and fabrication. Your ability to explain difficult concepts on an intuitive level is indispensable. Thanks to you I saw the beauty in mathematical formulations that capture the intricate details of our world in equations.

Finally, I would also like to thank my parents Oco and Mama for their unbounded love and care that accompanied me even thousands of kilometers away from home. A special thanks goes also to my brother Ivo who was there for me at all times (and oddly late hours) whether it would be random chatting, medical advice, or our weekend gaming sessions. Thanks to you and my parents I always felt there is a place I can call my home. Last but not least I would like to thank the person closest to my hearth, Veronka. You helped me to not get stuck in local minima of thinking and to look at the world through a new perspective. Even during the hardest parts of my life you stayed by my side and always remained positive and encouraging. The moments we shared together made all the problems in the world disappeared and all that remained was us and our love.

Contents

Contents **ix**

List of Figures **xv**

List of Tables **xxix**

1 Introduction **1**
1.1 Contributions 4
1.1.1 List of Publications 5
1.1.2 Overview 6

2 Perceptual Background **9**
2.1 Investigating Perception 9
2.2 Haptic Perception 15
2.2.1 Surface Exploration Via Direct Touch 16
2.2.2 Surface Exploration Via A Probe 18
2.2.3 Compliance Perception 20
2.3 Visual Perception 23
2.3.1 Visual Acuity 23
2.3.2 Color Perception 24
2.3.3 Gloss Perception 25

3 Related Work **29**
3.1 Haptic Reproduction 29
3.1.1 Commercial Haptic Styli 31
3.1.2 Research Prototypes of Haptic Styli 32
3.2 Computational Fabrication 33
3.2.1 Numerical Simulation of Haptic Feedback 34
3.2.2 Optimization Algorithm 36
3.2.3 Parametrization 37

3.3 Perception in Computer Graphics 38
3.4 Appearance Fabrication 39
3.4.1 Color Reproduction 40
3.4.2 Translucency Reproduction 40
3.4.3 Reflectance Fabrication 41

4 Perception-Aware Computational Fabrication for Haptics 45

5 An Interaction-Aware, Perceptual Model for Non-Linear Elastic Objects 49
5.1 Overview 51
5.2 Perceptual Space of Stiffness 51
5.2.1 Stimuli 52
5.2.2 Methodology 53
5.2.3 Data Reliability 54
5.2.4 Non-metric MDS 54
5.2.5 Analysis 54
5.3 Computational Model for Stiffness 56
5.3.1 Force Distribution 57
5.3.2 Models 59
5.3.3 Analysis 61
5.4 Complex Geometries 63
5.4.1 Computational Model for Complex Geometries 64
5.4.2 Evaluation 67
5.5 Application to 3D Printing 70
5.5.1 Accuracy 70
5.5.2 Material Selection for Prototyping 72
5.5.3 Improving Fabrication Algorithms 73
5.6 Limitations and Future Work 74
5.7 Conclusions 76

6 Perception-Aware Modeling and Fabrication of Digital Drawing Tools 77
6.1 Overview 80
6.2 Physical Measurements 81
6.2.1 Measuring setup 81
6.2.2 Defining measurement parameters 82
6.2.3 Measurements 82
6.3 Perceptual Space Optimization 85
6.3.1 Recovering The Perceptual Space 85

6.3.2 Experiment Design 88
6.3.3 Experiment 88
6.3.4 Perceptual Space of Drawing Tools 90
6.3.5 Accuracy 91
6.4 Physical Simulation 93
6.4.1 Perception-Aware Coarsening for Exponential Euler Integrator 93
6.4.2 Recovering the Forcing Signal 101
6.4.3 Results 103
6.5 Application For Drawing Tool Design 104
6.5.1 Perceptual Space Exploration 104
6.5.2 Optimizing Digital Styluses 107
6.6 Limitations and Future Work 109
6.7 Conclusion 110

7 Fabrication-in-the-Loop Co-Optimization of Surfaces and Styli for Drawing Haptics 111
7.1 Overview 114
7.2 Problem Modeling 114
7.2.1 Stylus Parametrization 114
7.2.2 Surface Parametrization 115
7.3 Efficient Fabrication-In-The-Loop Optimization 116
7.3.1 Haptic Feedback Similarity 117
7.3.2 Gaussian Process Surrogate Model 118
7.3.3 Acquisition Function 122
7.4 Application To Stylus-Surface Design 125
7.4.1 Reliability Test 125
7.4.2 Optimizing Haptic Feedback 126
7.4.3 Gamut of Haptic Feedback 129
7.4.4 Universal Drawing Surface 130
7.5 Comparison and Evaluation 132
7.5.1 Comparison With Geometry-Inspired Baseline 133
7.5.2 Measurements of Tool-Surface Combinations 133
7.5.3 Optical Properties of Generated Patterns 134
7.6 Experimental Validation 135
7.6.1 Study 1: Acquiring Vocabulary 135
7.6.2 Study 2: Distinguishing Drawing Tools 136
7.6.3 Study 3: Comparison with State of the Art 137
7.6.4 Study 4: Realism of Reproductions 139

7.6.5 Summary 140
7.7 Survey With Professional Artists 140
7.8 Limitations And Future Work 142
7.9 Conclusion 144

8 Towards Spatially Varying Gloss Reproduction for 3D Printing 145
8.1 Overview 147
8.2 Hardware Apparatus 148
8.3 Varnish Jetting 149
8.3.1 Effects of Printing Parameters on Droplet Shape 149
8.3.2 Effect of Varnish Spacing 151
8.3.3 Optimized Jetable Varnishes 152
8.4 Varnish Selection 153
8.4.1 Measurement Setup 153
8.4.2 BRDF Model Fitting 154
8.4.3 Reflectance Gamut 155
8.5 Gloss Reproduction 157
8.5.1 Simplex-Interpolation for Halftone Reflectance Prediction 157
8.5.2 Predicting Halftone Pattern Quality 159
8.5.3 Prescribed Reflectance Reproduction 160
8.6 Results 162
8.6.1 Flat Samples 162
8.6.2 Height-field Samples 164
8.7 Limitations and Future Work 165
8.8 Conclusion 168

9 Conclusions And Future Work 171

Bibliography 175

A Compliance Measurements 199

B Drawing Tools Measurements 203

C Stylus Simulation Validation 209

D Survey With Professional Artists 213
D.1 Artist 1 213
D.2 Artist 2 217
D.3 Artist 3 221

D.4 Artist 4 224
D.5 Artist 5 228

Figures

1.1 Stages of computational fabrication. The input geometry is subdivided and materials are assigned. Using numerical simulation, we estimate the physical properties of the assignment. Finally, we compare the current state with target properties and use the error estimate to drive the optimization. 2

2.1 Three designs of detection threshold experiments from left to right: Method of Limits, Method of Constant Stimuli, and Adaptive Method. 10

2.2 The plot shows the Just Noticeable Difference as a function of stimulus intensity. These two quantities are correlated through the Weber fraction. 11

2.3 Example of investigating the magnitude difference of appearance. The pairwise experiment presents participants with two samples and a slider to quantify the perceived difference. In contrast, the triplet experiments shows participants a reference and two possible reproductions from which they have to pick one. 12

2.4 Perceptual space of gloss recovered by [Wills et al., 2009]. Since the axes of the space are apriory unknown additional experimental effort is required to identify them. 14

2.5 Diagram of human skin structure showing the mechanoreceptors responsible for tactile perception. Image taken from [Boundless, 2013]. 15

2.6 Stimuli used during the psychophysical experiments. From left to right a pin array display [Wang and Hayward, 2008], sinusoidal gratings [Skedung et al., 2013], vibratory shaker [Israr, Tan and Reed, 2006]. 17

2.7 Perceptual space of tactile feedback recovered by Hollins et al. [2000]. 18

2.8 Israr, Choi and Tan [2006] used a shaker with attached stylus to investigate the sensitivity to vibration in a pen-holding posture. . 19

2.9 Perceptual space of haptic feedback induced by a probe recovered by Yoshioka et al. [2007]. 20

2.10 Stimuli used for investigating compliance ranged from real materials [Tiest and Kappers, 2009], through robotic arms to visual stimuli [Kuschel et al., 2010]. 21

2.11 Different interaction modes used to investigate compliance perception: tapping with a finger (A), tapping with a stylus (B), applying force through a lever (C), and passive pressing of a finger (D), Friedman et al. [2008]. 22

2.12 The experimental setup of Leškowský et al. [2006] used both real and virtual samples resulting in a two-dimensional perceptual space. 22

2.13 Diagram of human eye structure showing the components involved in visual perception. Image taken from [Didyk, 2012]. 23

2.14 Optical image of the foveal cone mosaic [Curcio et al., 1990]. The distance between two cones defines the maximal distinguishable frequency of visual stimuli or the so-called contrast-sensitivity function (right). 24

2.15 The response of human photoreceptors to different wavelengths (left) and the CIEL*a*b* perceptually uniform color space where ellipsis mark perceptually equivalent colors [*Pantone*, 2018] (right). 25

2.16 Renders of materials from the Merl database [Matusik et al., 2003] manifesting various perceived gloss differences. 26

2.17 Various geometrical shapes used to investigate the perception of gloss [Fleming et al., 2004; Fores et al., 2012; Havran et al., 2016]. 27

3.1 Haptic reproduction can help during education, tele-operation of robots [Peshkin et al., 2001], or training for laparoscopic surgery [Romanelli and Earle, 2009]. 30

3.2 Exemplars of haptic feedback devices ranging from haptic phantoms [*3D Systems*, 2020], through wearable displays [Perez et al., 2017], to static displays based on electrovibration [Bau et al., 2010]. 31

3.3 Commercial approaches to vary the haptic feedback range from altering the stylus nib [*Wacom*, 2020], through use of rough foils [*Paperlike*, 2020], to digitizing the strokes of traditional drawing instruments [*Repaper*, 2020]. 32

3.4 Research prototypes rely on active modulation trough vibrations induced by voice coils [Romano and Kuchenbecker, 2012], linear resonant actuators [Cho et al., 2016], or solenoids [Lee et al., 2004]. 32

3.5 Fabricated metamaterials can be used to modify the compliance [Schumacher et al., 2015], texture, or weight of an object [Torres et al., 2015]. 33

3.6 Numerical simulation of haptic feedback ranges from relatively simple contact of rigid bodies [Otaduy et al., 2004] to simulating the frictional interaction between a deformable finger and a surface [Tada and Pai, 2008]. 34

3.7 Robust handling of frictional contact allows us to predict the behavior of an elastic object during interaction [Chen et al., 2017]. Such predictions can be later used to optimize the flight path of the object. 35

3.8 Black-box function optimization using the CMA-ES method from [Fujii et al., 2018]. The algorithm starts by randomly sampling from a uniform distribution. The samples are then filtered based on fitness value. From the filtered samples a new covariance matrix and sample distribution is generated which is then used to sample the next generation. 36

3.9 Selecting an appropriate parametrization can significantly facilitate the optimization problem [Chen et al., 2013]. Modelling the object as smooth bezier curves is advantageous for caustic optimization, whereas for texture design it is more appropriate to model individual voxel elements. 37

3.10 Experimental designs investigating the perception of gloss [Wills et al., 2009] (left) and translucency [Gkioulekas et al., 2013] (right). 38

3.11 Applications of perception in creating more intuitive user interfaces for material [Serrano et al., 2016] and cloth simulation editing [Sigal et al., 2015]. 39

3.12 Differen strategies for color optimization for inkjet printers. From left to right: halftone dithering [Brunton et al., 2015], color contoning [Babaei et al., 2017], optimization using rendering predictions [Elek et al., 2017; Sumin et al., 2019]. 40

3.13 Reproducing target translucent properties with an inkjet printer [Hašan et al., 2010; Urban et al., 2019] and by optimizing colored silicone mixtures [Papas et al., 2013]. 41

3.14 Different approaches to reflectance manufacturing. From left to right: reflectance reproduction via ink deposition [Matusik et al., 2009], microfacet optimization of target gloss [Rouiller et al., 2013], and bi-scale approach combining ink and geometry modification for appearance reproduction [Lan et al., 2013]. 43

5.1 3D printing allows us to print objects with varying deformation properties. The question that we want to answer is: Given a set of printing materials and a 3D object with desired elasticity properties, which material should be used to print the object? For example, given sample ducks (left) with desired elasticity properties (e.g., measured), our system considers several candidate materials that can be used for replicating the ducks (right), and chooses materials that will best match compliance properties when examined by an observer (red and green outlines). Moreover, we can sort all possible materials by their perceived compliance as predicted by our model. The measured compliance is indicated with colors ranging from stiff (blue) to soft (red). 49

5.2 The 12 stimuli that were used in our experiments. The plot on the right shows corresponding force-displacement curves obtained via uniaxial test. The lines above blocks indicate the corresponding line styles. For a reference, we also include the force-displacement curve for a solid block (black line). The gray area roughly marks maximum applied force to each sample. 52

5.3 Left: Design of our samples. Equally spaced out cylindrical holes are defined using two radii. Right: Parameter of our samples. The two axes correspond to two radii of the cylinders (in mm). 12 green dots correspond to the samples used for our experiment. The red lines indicate the lower limits for the values of the radii, due to the printer resolution, and the upper limit for radii to guarantee that the cylinders do not intersect. Block 4 was introduced as a counterpart of block 8 to verify whether compliance properties are symmetric with respect to the choice of radii. 53

5.4 In a single trial participants were presented with three different samples. Their task was to compare softness difference between the leftmost and the rightmost pair. Additionally, we placed a force sensor beneath the middle block in order to collect force data throughout the entire experiment. 54

5.5 Left: Testing and training error for non-linear MDS as a function of λ parameter. Right: Variance across each embedding dimension (spread) for $\lambda = 5$. 55

5.6 Perceptual space for our 12 metamaterials found using non-metric multi-dimensional scaling. The error bars visualize the confidence obtained using the bootstrapping test. The intervals marked on the axis correspond to 1 JND. 56

5.7 The plots visualize the data collected in our force experiment for blocks 3 and 8. Plots on the left visualize the raw data captured from using our setup: the x-axis corresponds to samples (time), and the y-axis corresponds to the force applied by subjects. The red circles show the peak forces which are used to compute peak-force histograms on the right. As the deformation properties of blocks 3 and 8 differ significantly, the forces applied by users are also different. 58

5.8 The plots show how different computational models correlate with the perceptual space obtained using non-metric multi-dimensional scaling. Each chart's X and Y axis correspond to computational model and perceptual space, respectively. Blue points correspond to our stimuli. The red lines are the results of the linear regression and the green dashed lines indicate prediction intervals with a confidence level of 0.95. A good correlation between computational models and experimental data was found for several models. 62

5.9 The two plots visualize the distance prediction between our stimuli as predicted by the LogLS model and the L2 norm with respect to the observed distances in the perceptual space. LogLS outperforms the L2 norm by providing a much closer match to the observed distances. 63
5.10 The three models used for the simulation and their corresponding tetrahedral mesh. 65
5.11 To compute compliance properties of an object we start with simulating the interaction performed by an observer for a given location on the surface (a). This allows us to extract force-displacement information. Next, the force histogram is obtained using our database of blocks with ground-truth force information (b). To this end, we interpolate histograms corresponding to the two closest curves in the database. Given the force-displacement curve together with the corresponding force histogram, compliance properties are computed according to our model (c). The procedure is repeated for a set of locations on the surface (d). To obtain a dense, spatially-varying compliance map, we interpolate the compliance values on the object's surface (e). 66
5.12 To compute difference between two versions of the same object made of two different materials, we first compute a compliance map for each of them. Then, we subtract the values corresponding to the same locations. Here, Material A and Material B correspond to silicones with Young's modulus values of 0.0964 MPa and 0.0973 MPa, respectively. 67
5.13 The results of the evaluation of our compliance model. The plots visualize the compliance as judged by the participants (top row), as well as our prediction (bottom row). Additionally, the schematic figures visualize the mode of interaction. Letters A-I correspond to different silicon mixtures ordered by Shore hardness. For full simulations of the used materials please refer to Appendix A. . . . 69
5.14 From left to right: a rendering of the original seahorse model. Seahorse models used for the study: Dragon Skin 30, Ecoflex 00-30, TangoBlack+, Flexible resin, TPU 92A-1. 70
5.15 Left: a rendering of the original octopus model. Right: all printed octopuses ordered by stiffness. 72
5.16 Left: force-displacement curves obtained from Spec2Fab optimization. Right: Slopes of the force-displacement curves along with corresponding histogram of forces. 75

6.1 We propose a system for fabricating digital drawing tools that mimic the feel of real tools. To this end, we measure properties of different real drawing tools, study their perception, and design a perception-aware space of drawing tools. We later develop a simulation technique which allows us to embed new designs into the space, evaluate the pairwise similarity between them and the tools we want to replicate. This drives the design process of different digital tools. 77
6.2 Measurements and perceptual user study setup. Turntable on which a mechanical arm is used to capture friction and vibration of drawing tools. There is a holder for a human arm (during user studies) to provide consistent grip and orientation. In the middle, we can see normalized friction measurements recovered using our setup. On the right is the accelerometer response for a fixed velocity. . . 81
6.3 Sample images drawn during our preliminary user study. 82
6.4 Our user study considers 5 drawing tools: 2H and 8B pencils, fineliner, ballpoint pen, and charcoal and 3 drawing substrates: 80-gram office paper, rough artist paper, and stone paper. 83
6.5 Velocity and pressure histograms recorded from free drawings during our pilot study. 84
6.6 Acceleration measurement setup. The arm holds a drawing instrument, and the drawing substrate is attached via a magnetic plate. The vibrational response of the drawing tool is captured using an accelerometer. 85
6.7 Friction measurement setup. The arm is enhanced with a lever that amplifies the frictional force excreted on the force sensor. . . 86
6.8 Measurements of a ballpoint pen on oiled acrylic (left), and 80-gram office paper (right). 87
6.9 Turntable setup for user studies. Participants' hands rest in a holder and an arm-wrap is used to limit wrist motion. 89
6.10 The plot shows the match of our optimized perceptual spaces to the experimental data. Different colors correspond to spaces of different dimensionality, while points with the same color correspond to different values of λ. 91
6.11 The perceptual space obtained using our data and optimization. The axis are correlated with vibration and friction measurements. Confidence intervals obtained by bootstrapping visualize 95% and 68% regions. 92

6.12 Synthetic results of optimizing perceptual spaces using our experiment design as a function of space dimensionality. The bars show mean error and standard deviation. 93
6.13 Human sensitivity to vibration as a function of frequency. We are most sensitive at around 240 Hz and regarded as insensitive bellow 50 and above 500 Hz. 95
6.14 Oscilating 3D printed bar captured with accelerometer (blue) and physical simulation of the bar (red). 101
6.15 Fit of testing dataset at a fixed velocity. We compare physical measurements (blue) to our data-driven simulation (red). We can see a good match of estimated vibration for the testing dataset. . . . 102
6.16 Material (left) and tip (right) scaling factor. Both factors were optimized jointly. 103
6.17 An interpolated design was measured (blue line) and simulated using our pipeline (red line). We also compare to results of our simulation if we would not interpolate tip and material scaling parameters (gray lines). 104
6.18 Perceptual space grouped by means of fabrication: real tools (red), 3D printed tools on paper (green), and artificial substrates (blue), and commercial styli (yellow). 105
6.19 The results of the free drawing user study. For each out 24 triplets, we plot preference obtained in our validation experiment next to the prediction coming from the perceptual space. The red background indicates the failure cases where the majority vote did not agree with our prediction. 107
6.20 Details of our perceptual space for optimization of digital styli. We show two cases: interpolation of 3D-printed styli on a plastic screen protector (left), and interpolation of 3D-printed styli on the glass surface in comparison to the Apple Pencil (right). Number indicate subject preference when compared to real pencil. 108

7.1 We propose a data-driven method for mimicking haptic feedback of drawing tools. Our method uses fabrication-in-the-loop design enabled by our data-driven surrogate model which automatically handles exploration-exploitation trade-offs and minimizes the amount of printed samples. The final stylus-surface combinations are manufacturable on commonly available hardware and can be directly integrated into current digital drawing solutions. 111
7.2 Frictional contact 112

7.3 The modular design used for our styli (left). The replaceable tip can be made of different material m and have a different radius r. A picture of a manufactured stylus (right). 115

7.4 Gelsight scans of drawing substrates used in our experiments. . . 116

7.5 Sample surfaces from our parametrization quantized to one layer. 117

7.6 A one-dimensional example of information transfer between multiple categories. The estimate of category 1 (left) is used to improve prediction in uncertain regions of category 2 (middle). With more samples the prediction is weighted towards observations of the new category (right). 121

7.7 Visualization of acquisition function. 123

7.8 Example of parallel execution of stylus and surface manufacturing. In the first iteration, we manufacture a surface and enhance our data-driven model. The model then predicts which tool we should print to best match the target. We prepare the print for the tool and while the tool is printing we optimize for a new surface. The new surface causes an update to the surrogate model which results in a new suggestion for a stylus to print. Since the printer is currently occupied we compare the tools based on their expected improvement over time. We find that the expected improvement of the new tool outweighs the difference in printing time. Therefore, we stop the current print and prepare the new tool. In order to verify that the model prediction was correct, we manufactured and measured both styli. We can observe that the newly predicted stylus indeed achieved better improvement. 127

7.9 Pen and surface designs used to initialize the surrogate model. . 128

7.10 Initial samples measurements are shown on the left. We can observe that initial sampling can capture a ballpoint pen on office paper. We verify our algorithm by removing the surface from initial sampling and optimize for the same tool, right. 128

7.11 Results of optimizing haptic feedback using our surrogate model. Black samples mark original perceptual space. Blue samples are sample surfaces used to initialize the model. The red color is the position of a tool we would like to reproduce. Dark green dots mark the path our optimizer took in the perceptual space to find the final reproduction (light green). 129

7.12 Four traditional drawing tools (top) and their optimized digital counterparts (bottom). We show both the optimized stylus and the Gelsight scan of the optimized surface. Please note that our surfaces do not match the original substrates yet manifest the same haptic response. 130

7.13 The gamut of haptic feedback achievable by our different manufacturing processes: laser-jet on transparency sheet (blue), laser-jet on transparency sheet with fixative (red), 3D printer (green) compared with traditional drawing tools (black). 131

7.14 Four iteration of optimal surface optimization. 132

7.15 Results of surface reproduction of a ballpoint pen. The scanned height-field (left) and measured haptic feedback (right) form paths in perceptual space parametrized by the scaling factor. 134

7.16 Vibration measurements of original drawing instruments and their 3D printed counterparts. 134

7.17 Visual properties of our surface evaluated by taking photos of an LCD screen displaying a checkerboard pattern. 135

7.18 Bar-plot showing the similarity of our tools when compared to other digital styli. For each reference, indicated on the top, and an alternative solution, indicated below the bars, the plot shows the percentage of participants finding our replica more similar than the alternative commercial or research solution. 138

7.19 Bar-plot showing the similarity of our tools when compared to traditional drawing tools. For each reference, indicated on the top, and an alternative traditional tool, indicated below the bars, the plot shows the percentage of participants finding our replica more similar than the alternative tool. 140

7.20 Two-dimensional perceptual space computed for tools in our experiments. The Euclidean distance between the samples can be interpreted as a perceived distance between different drawing tools combinations. 141

8.1 The input to our system is a diffuse color and spatially-varying gloss. We first reproduce the color using commercial ink-jet printers (left halves). Next, as a post-processing step we use our custom printer to jet varnishes that match the input reflectance (right halves). 145

8.2 Varnish printing apparatus (left) consists of needle jetting valves (middle). To jet with the valve, the needle rises allow pressurized material to flow. By quickly shutting the valve the material is jetted onto the substrate. 149
8.3 Gelsight captures of droplets created by varying the jetting parameters: valve stroke, valve open duration, valve close duration, and air pressure. 150
8.4 Gelsight (top) and optical (bottom) scans of varnish droplets for a selection of off-the-shelf varnishes. From left to right, 1:1 ratio of Golden MSA gloss varnish and Golden MSA solvent, 1:1 ratio of Golden MSA satin varnish and Golden MSA solvent, 1:1 ratio of Golden MSA matte varnish and Golden MSA solvent, Schmincke 611 matte varnish, Schmincke 610 gloss varnish, Amsterdam 115 Matte Varnish, 1:1 ratio of Golden matte heavy gel and Golden gloss extender. 151
8.5 Effect of varying spacing. In the observed range increasing the spacing leads to thinning of the film created by the deposited varnish. 152
8.6 Photos of our printed varnishes on a cylindrical setup with a line light source. 153
8.7 Our measurement setup and sample capture of matte and satin varnish. 154
8.8 Fitted raw measurements of our varnishes (top) and corresponding rendering with the fitted model parameters (bottom). 155
8.9 The reflectance gamut achievable by our varnishes (orange). The gamut contains various hand-made samples for comparison: varnished paper (blue), varnished plastics (yellow), oil colors (green), and transparencies (purple). The varnishes we selected for printing are in bold. 156
8.10 Test of linearity of varnish halftoning. We start by generating halftone patterns and applying them on a substrate. The samples are then measured and compared with linear prediction. 157
8.11 A two-dimensional simplex created to predict the reflectance of gloss halftones. To predict a new sample we locate the enclosing simplex and use barycentric coordinates to interpolated our measurements. 158
8.12 Cross-validation of our simplex-based reflectance prediction. Leaving out multiple samples, creating a new tesselation, and predicting the missing samples. 159

8.13 Predicting dithering pattern visibility for different varnish mixtures of matte-glossy-satin. 160

8.14 Prescribed roughness for a target reflectance (left) is dithered using our simplex model with (middle) and without (right) dithering pattern visibility optimization. 161

8.15 A gradation of gloss formed by jetting uniform patches of various varnish mixtures. The images are captured by placing the samples against a display showing a checkerboard pattern. 162

8.16 Spatially-varying gloss is defined with two varnishes: matte (black) and glossy (white); and dithered using our device. To validate the fabrication we capture the fabricated samples with an area light source and compare with rendered predictions. 163

8.17 Albedo, roughness, and metallic textures for our objects with their corresponding generated halftone patterns. 164

8.18 Differently slanted 3D printed surfaces covered with the same coverage of Golden gloss varnish, captured at specular configuration. The angles below the pictures indicate the deviation of the surface from the 3D printer's tray. The appearance start to show significant deviations from the flat surface at approximately 20°. 166

8.19 Appearance manufacturing with separate fabrication of color (inkjet printer) and gloss (our device). The final combined appearance manifests both high-resolution color and spatially-varying gloss. Photos are captured at two exposure values (+0, -1) with a still camera and a moving light source. For a full visualization please see the supplementary video. 169

8.20 Manufactured height-fields without varnish (top) and with halftoned varnish using our system (bottom). Photos are captured with a still camera and moving light source. We plot the luminance of two locations as the light rotates around to showcase the gloss variation achieved by our system. For the full capture please see the supplementary video. 170

A.1 Left: peak-force data recovered for each block during psychophysical experiment. Right: histograms stored in our database. 200

A.2 Force-displacement data recovered for interaction with all fabricated objects. 201

B.1 Vibration spectrograms of real tools recorded using our setup. Each title of the plot consist of name of the tool and the surface. The first column provides measurements of the tools on oiled surface in order to demonstrate the lack of vibrations coming from the measurement setup. 204
B.2 Vibration spectrograms of 3D printed tools recorded using our setup. The titles of the plots consist of the name of the material, size of the tip, and the name of the surface. 205
B.3 Vibration spectrograms of 3D printed and commercial tools recorded using our setup. The titles of the plots consist of the name of the material/tool (for commercial products), size/type of the tip, and the name of the surface. 206
B.4 Friction measurement of drawing tools on various kinds of paper recorded using our setup. 207
B.5 Friction measurement of 3D printed and commercial styli recorded using our setup. 208

C.1 Full leave-one-out cross-validation of our data-driven forcing term. 210
C.2 Interpolated forcing term evaluation. We predicted the vibration response of an interpolated forcing term using our simulator (red line) and compare it to measurements on our turn-table (blue line). We also compare it to the simulations without interpolation of forcing terms (gray lines). 211

Tables

5.1 Results of our JND estimation experiment. The first two columns provide block numbers (Figure 5.2) that were compared by the participants. The third column provides a percentage of people that answered that SAMPLE 1 was softer than SAMPLE 2. The fourth column provides the corresponding difference in our initial perceptual space. 57

5.2 Statistics from experiments evaluating each computational model. High values marked in bold indicate good candidates for computational models. Each of them does a good job of predicting not only the distances evaluated using the perceptual space found using MDS (above 90%), but also the data collected in our main experiment. 62

5.3 Correlation between distances in the perceptual space and the prediction provided by the computational models and the L2 norm. . 63

5.4 Silicon mixture with their corresponding IDs. 68

5.5 Correlation between distances predicted by our method and the L2 norm. 69

5.6 Results of the perceptual study. From each trial the triplet is shown as well as the corresponding participants' picks. Our predictions always correspond to the participants' choices. Cases when the L2 norm did not agree with our prediction are highlighted in red. . . 71

5.7 Results of the perceptual study. For both trials of picking two and three samples we present the people's choice and its corresponding probability. Sets predicted by our model and the L2 norm are underlined in green and red, respectively. 73

Chapter 1

Introduction

Additive manufacturing is a powerful technology that through spatial deposition of functionally graded materials enables us to control the haptic and visual impression an object creates. This is of broad importance to everyone, from a shopper buying shoes to a doctor palpating a growth, that uses their sense of touch to learn about the world. The haptic and visual feedback convey different characteristics of each object that lend themselves to different applications, e.g., texturing surfaces to mark handles. Thus there is an intimate connection between the way an object feels and how we use it. As a result, reproducing target haptic and visual behavior is critical for both replicating existing, real-world constructs and designing novel ones.

To design an object with customized haptic and visual properties a designer needs to take full advantage of the new capabilities [Schmidt and Ratto, 2013]. However, the extreme amount of degrees of freedom makes an efficient exploration of the design space difficult. Therefore, it is necessary to design new algorithms that help the designer to fully exploit the potential of additive manufacturing. This idea is similar to enhancing camera and display devices with computational techniques to push them to the hardware limits [Lukac, 2010; Masia et al., 2013]. In the context of manufacturing, we collectively call computational fabrication the set of techniques that use computation to leverage the full capabilities of manufacturing devices.

Computational fabrication is a goal-driven design process [Chen et al., 2013] (Figure 1.1). The input is a description of the target behavior. The output of the technique is a representation of a printable design for the manufacturing device which will produce an object closely matching the prescribed behavior. This goal is achieved by a parametrize-simulate-evaluate optimization loop. The algorithm is initialized with a parametric design from a sufficiently expressive design space.

Next, numerical simulation is employed to estimate the properties of the digital design. Finally, a difference between desired and simulated behavior is evaluated, which causes an update to the parametric model. Such techniques were successfully employed to optimize objects with target reflectance [Matusik et al., 2009], elastic behavior [Schumacher et al., 2015], or sound [Li et al., 2016].

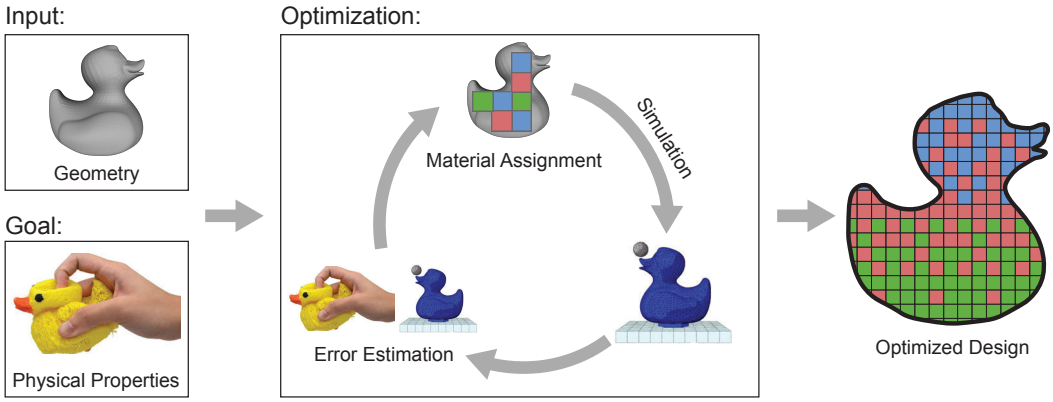

Figure 1.1. Stages of computational fabrication. The input geometry is subdivided and materials are assigned. Using numerical simulation, we estimate the physical properties of the assignment. Finally, we compare the current state with target properties and use the error estimate to drive the optimization.

Directly applying computational fabrication for the reproduction of target behavior is a challenging task. The optimization process hinges on efficient estimation of physical properties via numerical simulation to explore the design space. This, in general, leads to a highly non-convex and derivative-free optimization procedure. The problem is further complicated in our setting. The perceived haptic and visual sensations depend on the fine-scale interaction between objects surface and our fingers or incoming light rays. As with many complex design tasks, the governing physics are not currently predictively modeled to the complexity of the physical phenomenons which govern the haptic feedback, the scale at which they occur, and the limitations of the fabrication processes [Myshkin et al., 1998; Cirac and Zoller, 2012]. Moreover, the hardware limits constraint the space of objects that can be reproduced on a particular printer or the so-called gamut of the device. The gamut is limited by the selection of available materials, the need for supporting structures to print overhangs, or the resolution at which the print can be realized. While computational techniques can bring us closer to the edge of what is possible, they fundamentally cannot overcome the hardware limitations and are always bound to operate within the gamut of a particular device. To address these problems, we make an interesting

observation that similarly to hardware limitations, human sensorial system has its own set of imperfections, e.g., our fingers processes applied stimuli in a non-linear fashion [Skedung et al., 2011], we are sensitive to vibration in a relatively narrow range [Israr, Choi and Tan, 2006], our eyes have limited visual accuity [Ferwerda et al., 1996; Larson et al., 1997; Mannos and Sakrison, 1974], etc. [Nisky et al., 2011; Skedung et al., 2013; Goodman and Bensmaia, 2017].

In this thesis, we propose that it is possible to exploit the limitations of the human sensorial system to enhance the fabrication process. By combining computation with numerical models of perception, we mask the computational and hardware limits of fabrication, effectively increasing the *apparent* gamut of a 3D printer. Instead of optimizing for exact replicas, we search for perceptually equivalent solutions that are manufacturable on current hardware. To achieve this goal, we enhance the standard parametrize-simulate-evaluate optimization loop by replacing the evaluation step with a perceptual error metric. To recover this metric, we rely on designing psychophysical experiments [Fechner, 1860] that investigate how perceived haptic properties relate to physical attributes of manufactured digital designs. Such experiments require large amounts of exemplars and participants, which makes the experiments challenging to use in the context of fabrication where stimuli cannot be digitally distributed. We address this issue by proposing new experimental designs that minimize the number of participants. We support the experiments with numerical optimization that automatically generates a computational error metric that relates to perceived quantities to physical properties. To evaluate the error metric during optimization, we need to estimate the physical attributes of newly generated designs. Since such a simulation is often computationally intractable, we propose two alternatives. The first option is to combine numerical simulation with perceptual insights to perform perception-aware coarsening that focuses the computational effort on relevant phenomena. The second options is a purely data-driven fabrication-in-the-loop model that implicitly handles the fabrication constraints and the exploration-exploitation trade-off. These two approaches enable efficient numerical estimation of physical properties that are otherwise challenging to simulate, and often open research problems. We demonstrate the application of the proposed methodology to two problems: the design of objects with prescribed compliance and replicating the haptic feedback of traditional drawing instruments. In the last part of the dissertation, we extend our investigation beyond haptics – more specifically to appearance reproduction. We propose a full system for modifying objects' surface gloss by spatial halftoning of varnishes. The core of our approach is a predictive model that can estimate both the appearance and the perceived quality of varnish halftone patterns.

1.1 Contributions

The main contribution of this dissertation is a new methodology that combines perception with computational fabrication to design objects with prescribed haptic feedback. We validate this concept in three different scenarios and present directions for future work.

We demonstrate our methodology in the case of manufacturing objects with prescribed elastic behavior [Piovarči et al., 2016]. We propose to evaluate the similarity of elastic objects by building a perceptual space of compliance. We show that the perception depends on the desired behavior, objects geometry, and most critically on the interaction with the sample. To capture all of these effects, we combine numerical simulation that handles the shape and material properties with a data-driven predictor of interaction based on the local compliance. We demonstrate that the proposed model can predict the similarity of objects, provide material replacements that reduce manufacturing cost, and generate higher-quality reproductions than state-of-the-art methods.

In the second scenario, we focus on manufacturing drawing tools with prescribed haptic feedback [Piovarči et al., 2018]. To model the human perception of drawing tools, we propose a custom likelihood-based optimization of perceptual error metrics. We show that by jointly optimizing for a metric and its governing physical attributes, we maximize the probability of explaining the haptic perception. The recovered error metric relies on efficient estimates of physical attributes. We can not employ off-the-shelf methods since due to the scale of the interaction and printing artifacts an accurate numerical simulation of the governing physical phenomena is intractable. To this end, we propose to accelerate the simulation using perception-aware numerical coarsening. The key observation is that we need to simulate the haptic interaction only up to the resolution appreciable by the humans. We demonstrate this by designing an Exponential Euler simulator that computes the elastodynamic equations only in the range perceptible to a human holding a drawing tool. We combine the simulator with a data-driven forcing term to model different paper substrates and show an excellent match to measured data. Finally, we demonstrate an application of the simulator and error metric in the design of digital styli that are more similar to traditional drawing tools than commercial alternatives.

In a followup work, we push the haptic reproduction of drawing tools to the limits by reproducing the haptic feedback of traditional drawing instruments [Piovarči, Kaufman, Levin and Didyk, 2020]. Due to the lack of resolution and printing materials, it is not possible to 3D print carbon copies of drawing tools. Instead, we cast the optimization into a perceptual space where we search for a

stylus-surface combination that is printable and mimics the target haptic feedback. We show that efficiently estimating the interaction is critical for high-quality haptic reproduction. To tackle such a challenging problem, we propose to use fabrication-in-the-loop optimization. The key ingredient of our method is the ability to predict the physical attributes with confidence bounds. This enables us to design an exploration strategy that automatically handles the exploration-exploitation trade-off. At each iteration, it proposes a digital design that maximizes the improvement towards target haptic feedback. We show the benefits of our method by reproducing several traditional drawing instruments. The generated instruments are fully 3D printed yet manifest haptic feedback similar to materials that lie well outside of the printing gamut. The quality of the reproductions is demonstrated in a blind user study with casual users and a survey with professional artists.

Finally, we investigate future possibilities in applying perception-aware fabrication to appearance reproduction. Appearance reproduction is a growing field and as such there is yet no standardized method to reproduce the full objects' appearance, i.e., the surface reflectance (gloss), diffuse color, and translucency. To address this issue, we propose a custom printing hardware capable of jetting highly-viscous materials [Piovarči, Foshey, Babaei, Rusinkiewicz, Matusik and Didyk, 2020]. We use our novel printer to jet off-the-shelf varnishes that modify only the gloss of an object without modifying its sub-surface reflectance. This allows us to combine our printer with commercial devices and simultaneously reproduce the sub-surface color and surface gloss. We characterize the proposed design by describing a calibration procedure and by identifying the effects of jetting parameters on droplet shape. Next, using our optimized parameters and a full reflectance measurement device, we quantify the gamut of gloss achievable by our setup. Based on the recovered gamut we select three varnishes as basis materials for halftoning. To predict the appearance of halftoned varnishes we propose a simple yet effective predictive model based on simplex interpolation that can estimate the appearance of a dithered varnish mixture. Additionally, we enhance the model with a dithering pattern quality predictor inspired by the properties of the human visual system and physical varnish mixing. We validate the model with numerical simulation and by manufacturing several 2D and 3D examples with prescribed spatially-varying gloss.

1.1.1 List of Publications

The main contributions of this dissertation were published in three technical papers at ACM SIGGRAPH conference (and ACM Transaction on Graphics Journal):

- **Michal Piovarči**, David I. W. Levin, Jason Rebello, Desai Chen, Roman Ďurikovič, Hanspeter Pfister, Wojciech Matusik, Piotr Didyk, **An Interaction-Aware, Perceptual Model For Non-Linear Elastic Objects**, ACM Transactions on Graphics 35(4) (Proc. SIGGRAPH 2016, Anaheim, California, USA)

- **Michal Piovarči**, David I. W. Levin, Danny M. Kaufman, Piotr Didyk, **Perception-Aware Modeling and Fabrication of Digital Drawing Tools**, ACM Transactions on Graphics 37(4) (Proc. SIGGRAPH 2018, Vancouver, Canada)

- **Michal Piovarči**, Danny M. Kaufman, David I. W. Levin, Piotr Didyk, **Fabrication-in-the-Loop Co-Optimization of Surfaces and Styli for Drawing Haptics**, ACM Transactions on Graphics 39(4) (Proc. SIGGRAPH 2020, Washington DC, USA)

- **Michal Piovarči**, Michael Foshey, Vahid Babaei, Szymon Rusinkiewicz, Wojciech Matusik, Piotr Didyk, **Towards Spatially Varying Gloss Reproduction for 3D Printing**, ACM Transactions on Graphics 39(6) (Proc. SIGGRAPH ASIA 2020, Daegu, South Korea), *(to appear)*

Additionally, not included in this dissertation, I have a separate stream of work on geometry-based reflectance reproduction where as future work we would like to combine this approach with varnish deposition to improve manufacturing of objects with prescribed color and gloss:

- **Michal Piovarči**, Michael Wessely, Michał Jagielski, Marc Alexa, Wojciech Matusik, and Piotr Didyk. **Directional Screens**. Proceedings of Symposium on Computational Fabrication 2017, Cambridge, MA, USA

- **Michal Piovarči**, Michael Wessely, Michał Jagielski, Marc Alexa, Wojciech Matusik, and Piotr Didyk. **Design and Analysis of Directional Front Projection Screens**. Computers & Graphics Journal 2018, Volume 74

1.1.2 Overview

The dissertation is structured as follows. Chapter 2 provides a background on haptic and visual perception where we start by describing the common methods for conducting psychophysical experiments and how they were applied to investigate the limitations of our senses. We continue with an overview of related work (Chapter 3) in graphics for reproducing haptic properties in both virtual worlds and real environments. Furthermore, we discuss related work in

appearance reproduction. In Chapter 4, we formalize the goal of perception-aware fabrication and provide an overview of how the rest of the dissertation contributes to achieving this goal. The remaining chapters are organized based on the projects. We start with a case study on compliance (Chapter 5) where we show the basic building blocks of perception-aware fabrication and how we can use the method to improve numerical optimization. In the following chapter (Chapter 6), we revisit the methodology and address the three main challenges by proposing algorithms to conduct more efficient psychophysical experiments, automatically generate computational perceptual models, and accelerate physical simulation via perception-aware coarsening of the numerical model. We show the application of these improvements to the design of digital drawing tools. In the next project (Chapter 7), we extend the previous optimization to design digital drawing tools that mimick the haptic feedback of real instruments. The key ingredient is our fabrication-in-the-loop optimization algorithm that efficiently utilizes a black-box objective function to optimize the digital design. In the last project (Chapter 8), we investigate a potential avenue of future work in applying perception-based techniques to appearance reproduction. We present a complete workflow for editing the gloss of an object with special care to minimize the visibility of dithering artifacts inherent to juxtaposed materials. Finally, in Chapter 9, we conclude the dissertation and discuss the potential of future work.

Chapter 2

Perceptual Background

In this chapter, we start by introducing the field of psychophysics and describe a systematic experimental approach to investigating the human sensorial system. Afterward, we focus on the biological and psychological capabilities of the sense of touch and sight that motivate how leveraging perceptual limitations can mask the constraints of digital fabrication.

2.1 Investigating Perception

To manufacture objects with desired haptic or appearance properties we seek to associate their physical attributes with perceived sensations. The quantitative study of these effects is called psychophysics and the main investigative tool is a psychophysical experiment [Fechner, 1860]. During psychophysical experiments, the stimuli are varied in a controlled and systematic manner and the perceived sensations are observed. More specifically, we want to know if a subject can detect stimuli, describe their magnitude, differentiate between different stimuli, or associate a stimulus with physical properties.

During detection experiments, we are interested in finding the so-called detection threshold, i.e., the weakest stimulus that our sense of touch can detect. To recover such a threshold we can employ the Method of Limits [Gescheider, 2013], (Figure 2.1 left). The basic idea is to present the participants with a series of stimuli with progressively increasing or decreasing magnitude and observe at which intensity the detection changes. The main drawback of this method is that it introduces a dependency on the order of presented stimuli that can bias the participants. An alternative option is to use the Method of Constant Stimuli [Laming and Laming, 1992], (Figure 2.1 middle), where participants are presented with stimuli at random and the stimulus that is detected and not detected roughly

50% of the time is considered to be the detection threshold. This prevents the observer from guessing the next stimuli from previously observed samples at the cost of longer studies and less precise threshold estimation. The detection estimate can be improved by using an Adaptive Method [Treutwein, 1995], (Figure 2.1 right). Here the participants are presented with a detectable stimulus that is progressively reduced in fixed steps until it is not detected. At this point the direction is inverted, step size reduced, and the stimulus is increased in magnitude until it is detectable. These steps can then be repeated again to obtain more precise estimates of the detection threshold. From these experiments, we

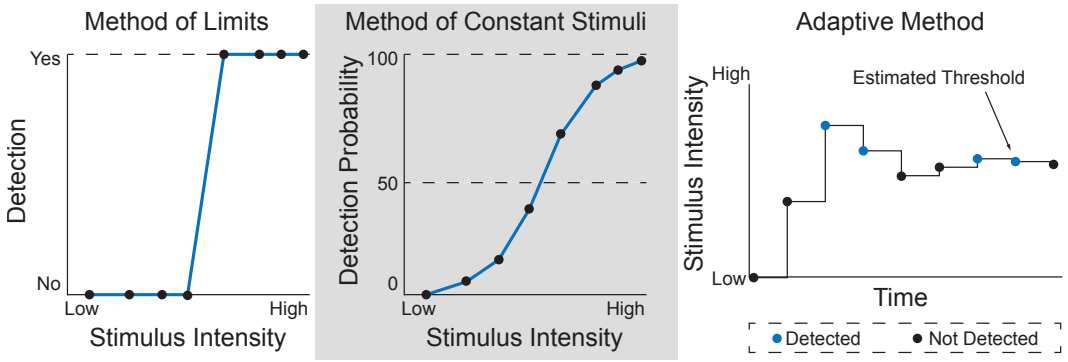

Figure 2.1. Three designs of detection threshold experiments from left to right: Method of Limits, Method of Constant Stimuli, and Adaptive Method.

can recover two thresholds. The *absolute detection threshold* that tells us the lowest magnitude at which a stimulus is perceivable; and the *differential detection threshold* which describes the magnitude increase necessary to detect a change in a stimulus. The differential threshold often obeys the so-called Weber's law which states that the perceived change in stimuli is proportional to the initial stimuli and this proportion is the so-called Weber fraction [Fechner, 1860], (Figure 2.2), or formally:

$$\mathrm{d}S = kS, \tag{2.1}$$

where dS is the Just Noticeable Difference (JND), k is the Weber fraction, and S the stimulus intensity. The JND is particularly important in the context of haptic reproduction as it is an objective estimate of similarity that can be used as the stopping criterion in an optimization. This is crucial since the exact reproduction of haptics is often unachievable on a 3D printer. However, by employing the JND it is possible to define a subspace of acceptable solutions that significantly simplifies the optimization task.

In magnitude estimation experiments the participants are presented with a stimulus and asked about numerical representation of its intensity on an arbitrary

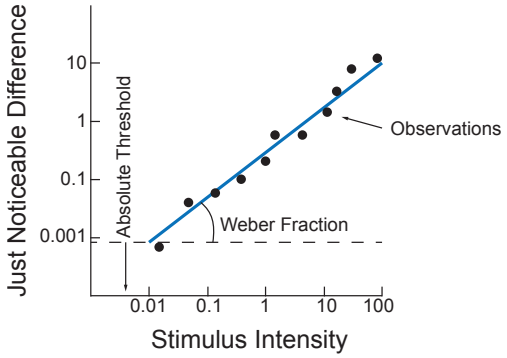

Figure 2.2. The plot shows the Just Noticeable Difference as a function of stimulus intensity. These two quantities are correlated through the Weber fraction.

continuous scale [Stevens, 2017], (Figure 2.3 left). Since such experiments are often challenging for the participants [Wills et al., 2009; Piovarči et al., 2016] they can be facilitated by presenting pairs of stimuli and asking about their dissimilarity. The experiment can be further simplified by reducing the continuous scale to a so-called Likert scale. The Likert scale is discrete and has between five to eleven magnitudes [Joshi et al., 2015]. The outcome of magnitude estimation studies are distances that describe the relations between the presented stimuli. However, even the simplified versions of the experiments are noisy [Wills et al., 2009; Piovarči et al., 2016]. This noise is attributed to the fact that observers do not have a consistent scale on which the samples are judged and judgements evolve during the experiment as new stimuli are presented. An alternative option to sidestep these issues is to study the difference of perceived distances rather than the magnitudes themselves. This leads to a subjectively easier experiment design called the two options forced-choice [Wills et al., 2009], (Figure 2.3 right). The experiments consist of individual trials during which the participants are presented with a reference stimulus and two possible reproductions. Their task is to select which reproduction is more similar to the reference. Since each trial depends only on the samples presented and does not require a numerical estimate of the difference the participants are more consistent. The outcome of such an experiment are the differences between distances. By further analyzing either the magnitudes or the differences of perceived distances we can recover semantic meaning and associate perception with physical attributes of the presented samples.

To model the semantic meaning and associate perceived stimuli with physical phenomena it is possible to compute a so-called *perceptual space*. A *perceptual*

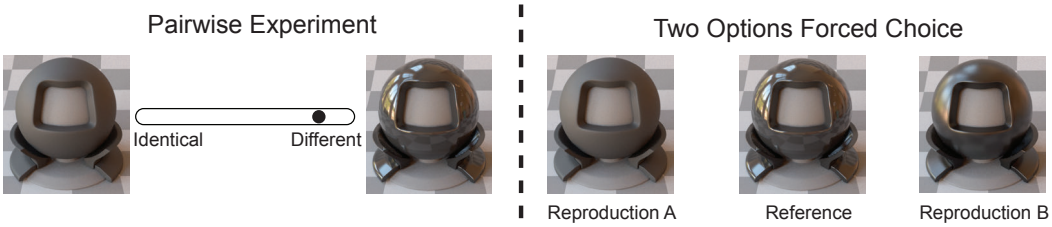

Figure 2.3. Example of investigating the magnitude difference of appearance. The pairwise experiment presents participants with two samples and a slider to quantify the perceived difference. In contrast, the triplet experiments shows participants a reference and two possible reproductions from which they have to pick one.

space is a multidimensional embedding of stimuli in which the perceived dissimilarity corresponds with Euclidean distance. The computational method to recover a perceptual space depends on whether we estimated magnitude or differences of perceived distances. In the following text we will derive both methods. Throughout the derivation we use the following notation. Lower case letters $i, j, k...$ are used to indicate indices of stimuli. Bold lower case letters $\mathbf{x}$ indicate vectors, and upper case letters $P, Q, R...$ are used for matrices. Columns of the matrix X denote the embedded coordinates of the samples, and matrix K denotes the Gramian $K = X^T X$, which is positive semi-definite ($K \succeq 0$).

Transforming the magnitudes of perceived distances into an Euclidean embedding can be realized by applying Multi-Dimensional Scaling (MDS) [Mead, 1992]. The input to MDS is a matrix of perceived pairwise distances. The algorithm minimizes a loss function usually denoted as stress a name adopted from physics due to individual distances pulling each stimulus in a different direction. The output is a M-dimensional perceptual space. In the special case of a dense distance matrix, the perceptual space can be recovered using Eigen decomposition. Let us set the distance matrix $D = [d_{ij}]$ where d_{ij} is the distance between sample i and sample j. To recover the perceptual space we perform the following calculation:

$$\begin{aligned} D^{(2)} &= [d_{ij}^2], \\ B &= -0.5 J D^{(2)} J, \\ B &= \Lambda E \Lambda^T, \end{aligned} \tag{2.2}$$

where $J = \mathbb{I} - \frac{1}{N}\mathbf{1}\mathbf{1}^T$ is the centering matrix, N is the number of samples, $\mathbf{1}$ is a column vector of ones, Λ is the vector of recovered eigenvalues, and E is the matrix of eigenvectors. The use of Eigen decomposition guarantees that the

axis of the recovered space are orthogonal. Additionally, the contribution of each axis to the distance can be directly estimated using the magnitude of the Eigenvalue. To compute a perceptual space of dimensionality M we pick the first M largest eigenvalues Λ_M and their corresponding eigenvectors E_M. The resulting perceptual space X is then:

$$X = E_M \Lambda_M^{0.5} \tag{2.3}$$

To recover a perceptual space from differences of perceived distance we use Non-Metric Multi-Dimensional Scaling (NMDS) [Wills et al., 2009; Silverstein and Farrell, 2001]. Let us represent the outcome of the psychophysical experiment as triplets (i, j, k), which indicates that the distance between samples i and k was smaller than between j and k. This is denoted as $d_{ik} < d_{jk}$. The goal of NMDS is to find a Euclidean embedding that satisfies as many of the constraints as possible. One can reformulate the NMDS problem using the positive semi-definite Gramian matrix K as follows. We can write:

$$d_{ik}^2 = ||\mathbf{x}_i - \mathbf{x}_k||_2^2 = K_{ik} - 2K_{ik} + K_{kk}, \tag{2.4}$$

where $\mathbf{x}_i$ and $\mathbf{x}_k$ is the position of sample i and k in the recovered perceptual space, respectively. As the distances are always non-negative, we can replace the constraints $d_{ik} < d_{jk}$ with $d_{ik}^2 < d_{jk}^2$. For such constraints our problem of finding the embedding is equivalent to finding a matrix K that satisfies all these constraints. The constraint that K is a Gramian matrix is sufficient and necessary for the matrix to be an inner product for some set of coordinates $\mathbf{x}$. Therefore, the solution of such a problem will always lead to a valid embedding:

$$X = E\Lambda^{0.5}, \tag{2.5}$$

for a given Eigen decomposition $K = \Lambda E \Lambda^T$ [Wills et al., 2009].

In the raw form, the above formulation is insufficient. First of all, the constraints do not uniquely define matrix K, as the constraints are invariant to scaling and translation. In particular, the optimization can collapse the entire space to one point. Another problem is that we want to minimize the dimensionality of the embedding; therefore, we must modify the cost function. We insert an additional term that minimizes the rank of X. This is equivalent to minimizing the rank of K. However, the rank of a matrix is a non-convex term, which complicates the optimization. Lastly, the above formulation seeks a matrix K that satisfies all the constraints. This, apart from the case where all participants are consistent with each other, is impossible.

Wills et al. [2009] accounted for all these problems and proposed the following optimization problem:

$$\min_{K,\xi} \sum_{(i,j,k)} \xi_{ijk} + \lambda tr(K), \tag{2.6}$$
$$\forall(i,j,k) \quad d_{ik}^2 + 1 \leq d_{jk}^2 + \xi_{ijk},$$
$$\xi_{ijk} \geq 0, \quad \sum_{ij} K_{ij} = 0, \quad K \succcurlyeq 0$$

Instead of reducing the rank of the matrix K, the optimization reduces the trace, as it can be used as an approximation of rank. In order to deal with possible constraint violations, they introduce slack variables ξ_{ijk}. Furthermore, to prevent collapsing the embedding to one point and to remove translation ambiguity, the distance between points is enforced to be at least one and the center of the embedding must be at the origin (second constraint in the third line). In order to solve the above optimization problem, we use the SeDuMi 1.05 Optimization Toolbox [Sturm, 1999].

Figure 2.4. Perceptual space of gloss recovered by [Wills et al., 2009]. Since the axes of the space are apriory unknown additional experimental effort is required to identify them.

The result of either MDS or NMDS analysis is a perceptual space X valid for the stimuli included in the psychophysical experiments, (Figure 2.4). To include

a new sample into the space it is necessary to conduct additional experiments with the new samples and compute a new space. However, such approach is impractical in the context of computational fabrication where we need to quickly evaluate the quality of newly generated designs. To this end, it is possible to design a computational approximation of the perceptual space where the axes of the space are correlated with measurable physical attributes. Unfortunately, this is typically a manual and tedious process that leaves the tradeoff between good correlation and explanation of experimental data in the hands of the experimentator. As an alternative in this thesis we propose a likelihood-based optimization that automatically computes a perceptual space correlated with physical attributes (Chapter 6). Our algorithm jointly maximizes the probability of observing the experimental data and correlation with various physical properties.

2.2 Haptic Perception

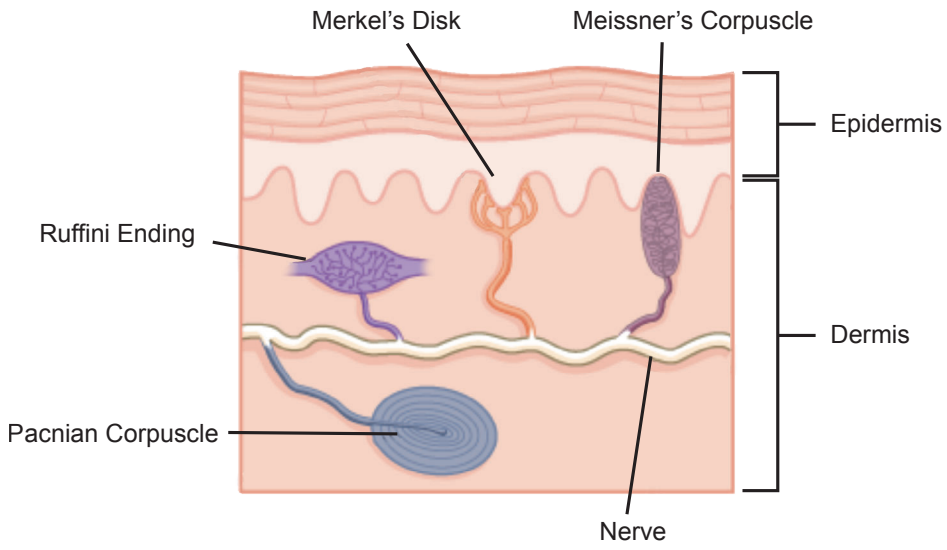

Figure 2.5. Diagram of human skin structure showing the mechanoreceptors responsible for tactile perception. Image taken from [Boundless, 2013].

The sense of touch is vital in interacting with real and virtual environments [Robles-De-La-Torre, 2006]. Using our fingers we can dextrously manipulate objects even without directly looking at them. The object's shape, weight, orientation, and other properties are perceived by the cutaneous deformation that is captured via mechanoreceptors [Johnson, 2001]. We recognize four kinds of

mechanoreceptors (Figure 2.5): two slowly adapting mostly sensitive to static stimuli, and two fast adapting sensitive to transient information. The Merkel disks are slowly adapting mechanoreceptors which main function is to transmit contact pressure to the brain and encode curvature and edges of the explored surface. The second slowly adapting mechanoreceptors are Ruffini endings which are sensitive to lateral stretching of the skin, motion direction, finger shape, and contact force. The Meissners corpuscles are fast adapting mechanoreceptors that responds to low-frequency vibration and transfer the image of skin motion. Finally, Pacinian corpuscles respond to higher frequency vibrations and are highly sensitive to transient signals. The information gathered by the mechanoreceptors is integrated to form the overall haptic sensation which depends on the mode of interaction. In this section, we explore recent findings in perception of surfaces via both direct and indirect touch mediated with a probe. We focus on the incredibly high discriminatory power of our fingers and how the impulses from mechanoreceptors combine into higher order concepts such as roughness, or softness. Finally, we perform a case study on the perception of compliance that can be perceived in both direct and indirect touch conditions.

2.2.1 Surface Exploration Via Direct Touch

The human finger can interact with surfaces either by passively touching them or through active exploration. For the static contact it was shown that the most significant attribute is roughness perceived by the spatial variation of mechanoreceptor firing rates [Connor et al., 1990; Yoshioka et al., 2001; Goodman and Bensmaia, 2017] and that the overall shape is mostly explained by the perceived curvature which depends on the local gradient of the surface [Kim et al., 2013]. During active exploration the most salient information is lateral forces [Kim et al., 2013] perceived as hysteretic signals [Ding and Bhushan, 2016] and high-frequency vibration [Romano and Kuchenbecker, 2012].

The discriminative power of the finger was studied by Wang and Hayward [2008]. They found that the minimal two-point discrimination is roughly 800 microns and humans can distinguish embossed dots with 550 microns diameter and 3 microns height. Active exploration increases our sensitivity. When fingerprint ridges slide over a surface the effect of pattern beating amplifies the perceived surface differences. Skedung et al. [2013] demonstrated that with a moving finger the smallest distinguishable bump is as tall as 13 nanometers. Moreover, participants could distinguish sinusoidal patterns with wavelengths down to 760 nanometers. The roughness of the surface also induces high-frequency vibration to which our sensitivity has a U shape in the range from 50-500 Hz with peak

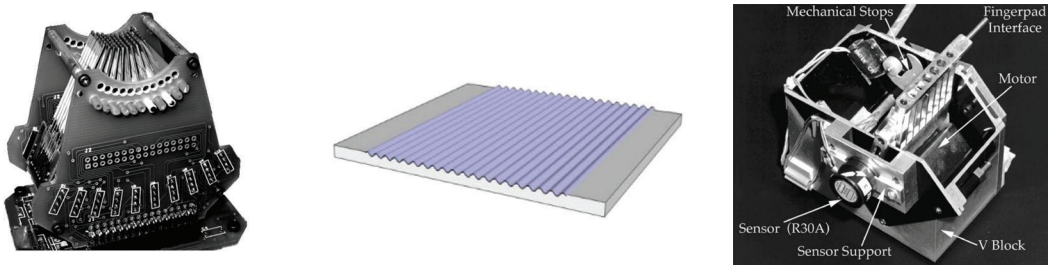

Figure 2.6. Stimuli used during the psychophysical experiments. From left to right a pin array display [Wang and Hayward, 2008], sinusoidal gratings [Skedung et al., 2013], vibratory shaker [Israr, Tan and Reed, 2006].

sensitivity around 240 Hz [Verrillo and Gescheider, 1975; Israr, Tan and Reed, 2006].

Roughness has been identified as the most significant attribute of objects surface for tactile exploration [Holliins et al., 1993]. Hollins et al. [2000] performed one of the first larger studies of texture perception on 17 different samples. They discovered three important perceptual components: roughness, compliance, stickiness, (Figure 2.7). [Bergmann Tiest and Kappers, 2006] extended the studies to over 120 textured surfaces. They recovered a four-dimensional perceptual space and found that the two axis correlate with roughness and compliance. Chen et al. [2009] explored how participants perceive different packaging during free exploration. They used six adjectives that describe haptic feedback perceived through a finger: warm-cold, slippery-sticky, smooth-rough, hard-soft, bumpy-flat, wet-dry. The primary identified dimensions are roughness and compliance with a secondary sticky dimension [Chen et al., 2009; Skedung et al., 2013]. Besides the physical properties of the surface [Cooke et al., 2006, 2010] demonstrated that our judgement of perceived roughness is influenced by the interaction mode.

The high precision of our fingers in both active and passive surface exploration provides challenges in haptic reproduction. With a resolution of commercial 3D printers at the order of microns [Sitthi-Amorn et al., 2015] it is impractical to mimic features down to the nanometer resolution of our fingerpads. However, the exploration of tactile properties governing the perception of materials provides a path in the possibility to trick our sensorial system. Even with its high precision humans reduce the tactile exploration to simple concepts such as roughness, compliance, or stickiness. Therefore, by focusing on reproduction of these higher-order concepts it is possible to mimic the haptics of real world even

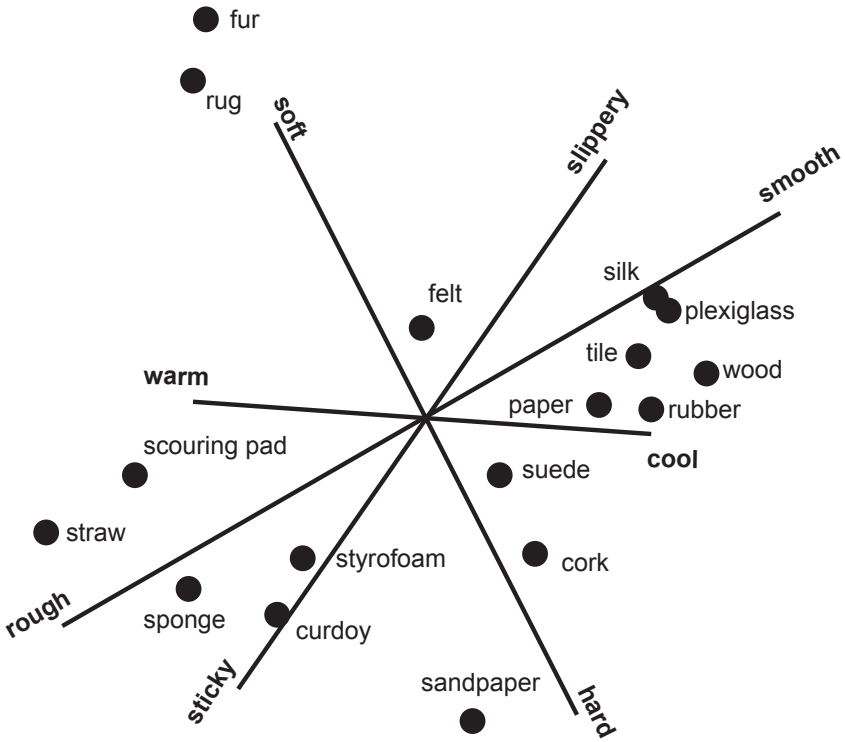

Figure 2.7. Perceptual space of tactile feedback recovered by Hollins et al. [2000].

with reproduction facilities that have orders of magnitude coarser resolution.

2.2.2 Surface Exploration Via A Probe

Researchers also considered the perception of texture exploration via a probe. In this context, the most dominant attribute is roughness – the perception of which is primarily governed by vibrations of the tool [Klatzky and Lederman, 2008]. Vibrations are sensed by Pacinian corpuscles [Hollins et al., 2006; Klatzky et al., 2003; Yoshioka and Zhou, 2009] with a peak sensitivity at 160-320 Hz [Israr, Choi and Tan, 2006], (Figure 2.8). The vibrations sensed through the tool depend on the properties of the tool and the surfaces. Klatzky and Lederman [2002] studied surfaces composed of raised pins. They demonstrated that along with surface properties, i.e., spacing of pins, tool size, exploratory speed, and applied force influence roughness judgments [Klatzky et al., 2003]. Furthermore, their results suggest that probe size and pin spacing are primary factors influencing the vibrations and so perceived roughness. Bensmaïa et al. [2005] proposed to

quantify the aggregated effect of the above factors on roughness perception using the spectral power of the vibrations produced by the probe.

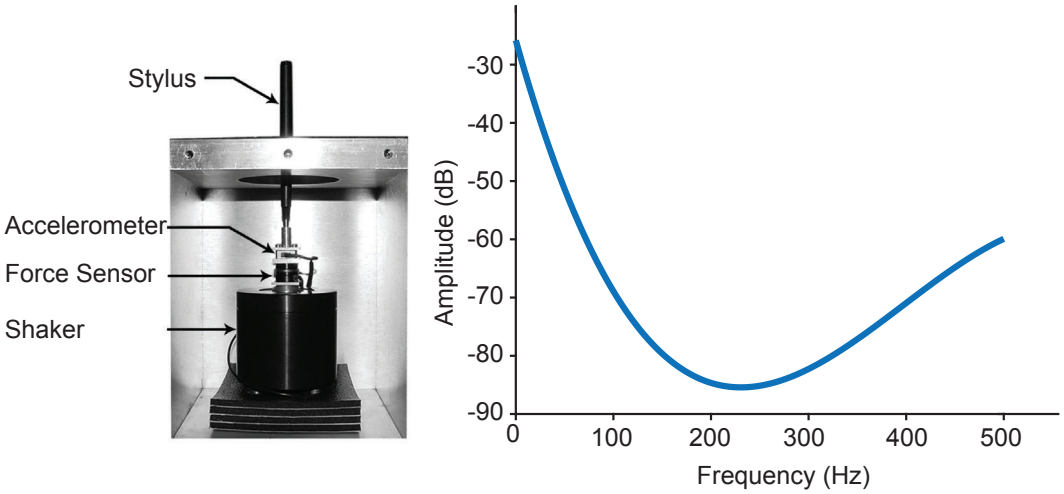

Figure 2.8. Israr, Choi and Tan [2006] used a shaker with attached stylus to investigate the sensitivity to vibration in a pen-holding posture.

Yoshioka et al. [2007] conducted user experiments in which subjects explored a wide range of real materials using a probe. They found that primary driving attributes for texture perception are roughness, stickiness, and compliance, (Figure 2.9). They also proposed to model these attributes using logarithms of vibratory power, the coefficient of friction, and relative compliance, respectively. These findings are supported by demonstrating that the sense of touch obeys the Webers's law. The Weber fractions were estimated for the perception of compliance [Jones and Hunter, 1990], vibration [Pongrac, 2008], and force [Nisky et al., 2011].

The perception of surfaces via a probe provides interesting insights into how our brain can compensate for lack of certain cues. Even when limited to half the mechanoreceptors and being presented only with vibratory and lateral forces our perception of roughness, stickiness, and compliance remains virtually unaffected. On the other hand, the use of a probe simplifies the haptic reproduction problem by masking phenomena that occur at the nanoscale: pattern-beating, and heat transfer. To this end, the reproduction of haptic properties explored via a probe is an excellent testbed that provides a controllable environment yet does not significantly limit our discriminatory power.

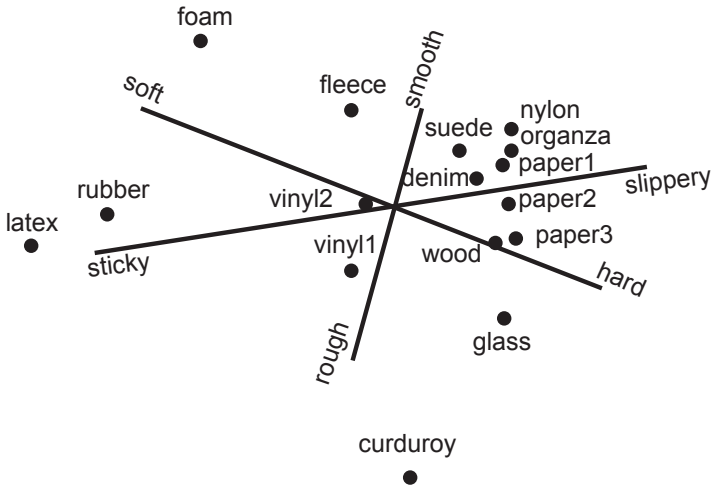

Figure 2.9. Perceptual space of haptic feedback induced by a probe recovered by Yoshioka et al. [2007].

2.2.3 Compliance Perception

Compliance perception is a complex phenomenon and many different cues are used by humans to judge material softness, (Figure 2.10). They range from purely kinesthetic, e.g., applied force, displacement, pressure distribution, surface deformation [Friedman et al., 2008; Tiest and Kappers, 2009], to visual [Wu et al., 1999; Kuschel et al., 2010]. As a result, many factors affect perceived compliance. For example, it has been demonstrated that softness perception can vary depending on the mode of interaction or velocity with which the stimuli is compressed [Friedman et al., 2008]. Even such things as texture, sample size, and location can influence the perceived compliance [Wu et al., 1999]. Moreover, compliance perception is closely related to cutaneous tactile information such as fingerpad contact area spread rate [Ambrosi et al., 1999].

Human ability to discriminate between linear materials of varying stiffness has been extensively studied and it has been shown that compliance perception obeys Weber's law. For example, Tan et al. [1992] showed that (using a vernier caliper) the Weber fraction is 8% when the displacement applied to the object is fixed and 22% when the displacement is allowed to vary. These experiments corroborate those of Jones and Hunter [1990], who reported that the Weber fraction is stable from 670 N/m to 6260 N/m, and equal to approximately 23%. Compliance discrimination can also be greatly affected by lack of certain cues. For example, lack of a work cue and minimized force cue can increase the Weber fraction to 99% [Tan et al., 1993], which suggests that these cues might be cru-

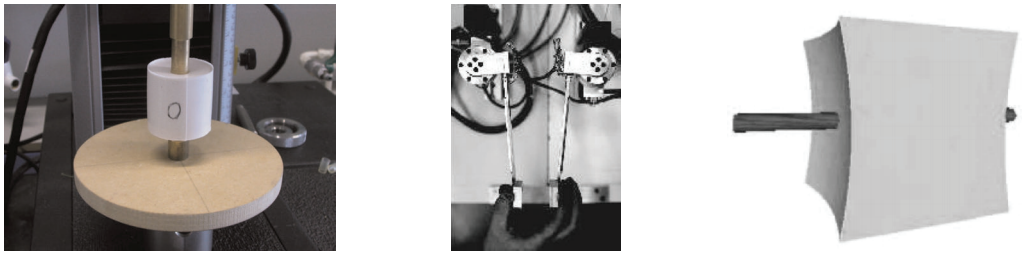

Figure 2.10. Stimuli used for investigating compliance ranged from real materials [Tiest and Kappers, 2009], through robotic arms to visual stimuli [Kuschel et al., 2010].

cial for discriminating between the softness of different materials. Interestingly, it has been demonstrated that the Weber fraction for compliance is much higher than for force and displacement [Jones and Hunter, 1990]. Koçak et al. [2011] reported that there is also a memory influence on compliance discrimination. Their experiments, performed using a haptic device, found that participants are more sensitive to softness changes during continuous presentation, i.e., when investigation of one sample is followed directly by another sample. More recently, Weber fraction has been also measured for a haptic jamming display to be 16% [Genecov et al., 2014]. Difference Magnitudes Estimations Studies on compliance discrimination give us insights into human ability to detect different levels of softness. To better understand the characteristics of perceived compliance, it is necessary to estimate the magnitude of perceived differences. One such experiment was conducted by Harper and Stevens [1964]. They tested subjective hardness of nine compliant specimens and calculated the relationship to physical stiffness. Magnitude estimation tests were also performed more recently by Friedman et al. [2008], (Figure 2.11). In that work, the main goal was to investigate how different modes of interaction affect the compliance perception.

The estimated magnitudes of differences can be used to construct a perceptual space of compliance. Leškowský et al. [2006] used MDS to investigate the difference between real and virtual compliant objects and found that a one-dimensional space is sufficient to describe real samples. However, when virtual samples are added, two dimensions are required, (Figure 2.12). Misra et al. [2009] used MDS to quantify the role of the Poynting effect (i.e., a presence of normal forces during shearing) on material perception. Several computational models explaining the perceived compliance have been proposed in the context of haptic telemanipulation systems. Pressman et al. [2007] found that delay between force and displacement can significantly affect the softness perceived by

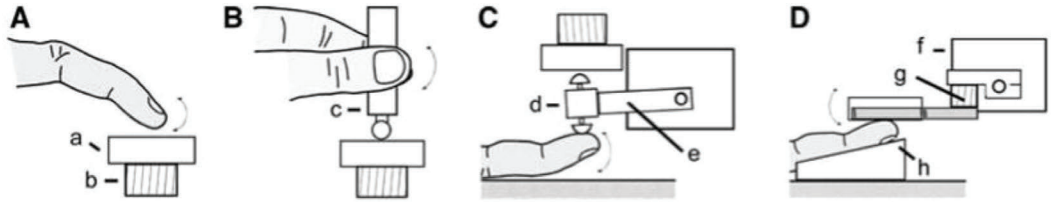

Figure 2.11. Different interaction modes used to investigate compliance perception: tapping with a finger (A), tapping with a stylus (B), applying force through a lever (C), and passive pressing of a finger (D), Friedman et al. [2008].

an operator. To predict the modified perception they investigated several computational models and showed that some of them, e.g., peak-force to displacement ratio, local stiffness, and maximum force can predict the phenomenon well. Later work [Pressman et al., 2011] showed that these models are powerful predictors of the perceptual effect on a compliance sensation induced when the position of an object is changed without visual feedback. These models were tested on linearly elastic materials, but delay added to haptic interaction introduced velocity-dependent, non-linear behavior. Leib et al. [2010] considered a similar situation, but this time the relationship between position and force was pairwise linear and independent of velocity. In their work, an additional computational model was considered in which compliance was estimated as the inverse of the slope of a line fitted to the force-displacement curve. Subsequent work by [Nisky et al., 2011] considered the effect of a delay in the context of the nonlinear behavior of a haptic device.

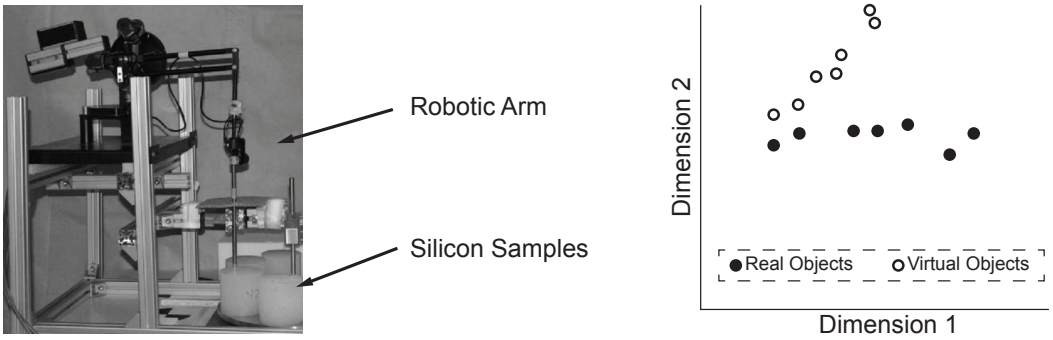

Figure 2.12. The experimental setup of Leškowský et al. [2006] used both real and virtual samples resulting in a two-dimensional perceptual space.

2.3 Visual Perception

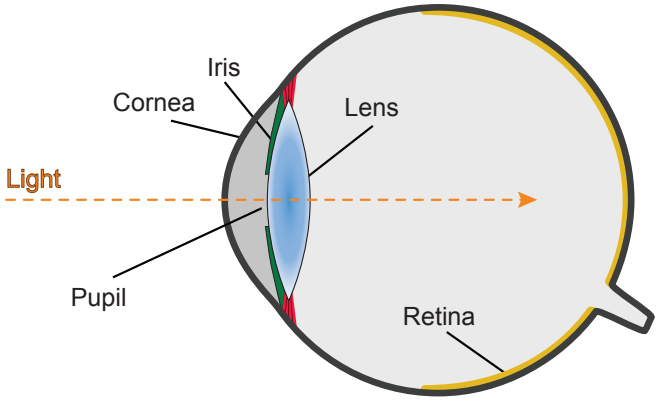

Figure 2.13. Diagram of human eye structure showing the components involved in visual perception. Image taken from [Didyk, 2012].

Vision is considered as the most important and most complex of our senses [Hutmacher, 2019]. When light enters our eyes it first passes through a dome-shaped cornea that helps us to focus the light toward the pupil. To control how much light enters our eyes the iris adjusts the diameter of the pupil similarly to a camera. The light is then focused by the lens directly onto the retina where photoreceptors convert the incoming light into electrical impulses that are further transferred by the nerve fibers into our brain. There are two kinds of photoreceptors: rods and cones each with a specialized function. On the one hand, the rods are sensitive to brightness and are responsible for perceiving objects' size, shape, and speed. On the other hand, the cones are sensitive to light at different wavelengths and are responsible for perceiving objects' color [Wandell and Thomas, 1997]. In this section, we discuss the visual acuity of human eyes and how these capabilities translate to the perception of color and gloss.

2.3.1 Visual Acuity

Eyes' visual acuity is limited by an optical low-pass filter and the spatial density of cones. To avoid spatial aliasing artifacts the optical filtering perfectly matches the density of the cones [Wandell and Thomas, 1997] and as a result, the limit of spatial acuity is defined by the spacing between cones in the fovea, see Figure 2.14 left. For the average observer this corresponds to roughly 28 arc seconds which translates to the ability to distinguish contours that are 1.75 mm apart at 6 meters [Curcio et al., 1990] or the so-called 6/6 vision [Britannica et al., 1957].

The visual acuity closely relates to the contrast-sensitivity function that tells us how sensitive our visual system is to various frequencies of visual stimuli, (Figure 2.14 right). If the frequency of the stimuli is too high the individual patterns cannot be recognized anymore and are interpreted as being uniform [Larson et al., 1997].

The high visual acuity of human eyes is a significant challenge to overcome for appearance reproduction. During close inspection, the human eye can distinguish details as fine as 50 microns which are at the level of commercial inkjet printers [Sitthi-Amorn et al., 2015]. Fortunately, when the contrast is lower or an object is moved further away the optical averaging of visual features can be used to our advantage as it will mask the presence of high-frequency materials deposited on the object's surface and perceive them as uniform.

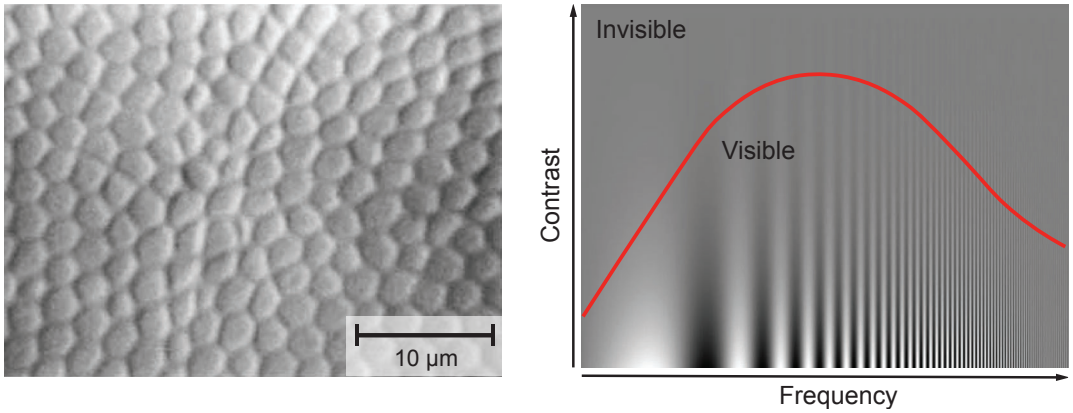

Figure 2.14. Optical image of the foveal cone mosaic [Curcio et al., 1990]. The distance between two cones defines the maximal distinguishable frequency of visual stimuli or the so-called contrast-sensitivity function (right).

2.3.2 Color Perception

Color is an optical property that defines the wavelength of light entering our eyes [Wyszecki and Stiles, 1982; Hunt and Wright, 1995]. To identify the color the eye uses a trichromatic system i.e., three different photoreceptors which signal is integrated into a single color. These receptors are sensitive to short (400-500 nanometers), medium (450-630 nanometers), and long (500-700 nanometers) wavelengths or more commonly referenced as blue, green, and red color, respectively, (Figure 2.15 left). Due to the trichromatic nature of color perception different spectral colors can appear the same. This effect is called metamerism

and is a result of the photoreceptors responding to cumulative energy received from a broader range of wavelengths. As a result, color reproduction is usually coupled with a predefined light source as changing the spectral distribution of the illuminant directly affects the perceived colors.

The perception of color is governed by three primary attributes: lightness that describes the overall intensity of the color, chroma or saturation that defines the strength of the color, and hue that defines the shade of the color. These primary attributes also define an intuitive space for color editing. Unfortunately, such a space is not perceptually consistent and small changes in individual attributes can lead to large perceived differences. To address this issue Hunter [1958] developed a perceptually uniform color space called *Lab*, (Figure 2.15 right). Here *L* stands for lightness and ranges from white to black, *a* ranges from green to red, and *b* has a range from blue to yellow. The perceptually uniform color space and its updated variant the *CIEL*a*b** [ICC 1:2004-10, 2003] are currently used as the industry standard for color calibration in printing devices.

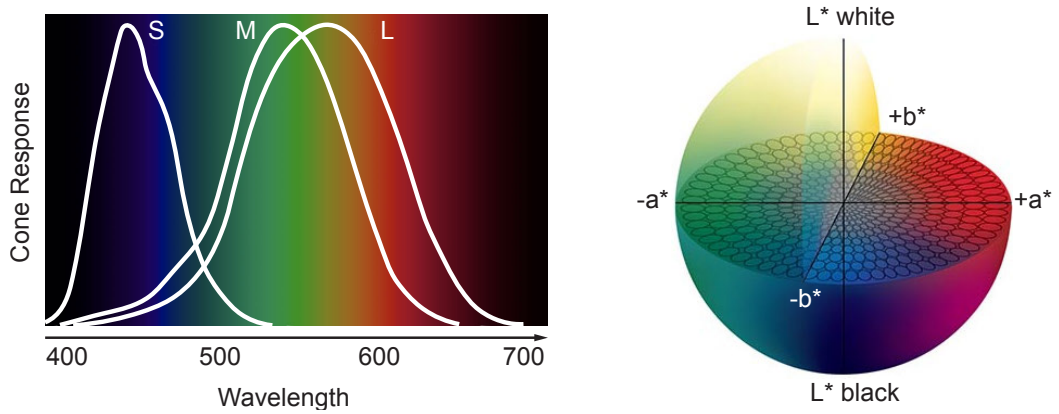

Figure 2.15. The response of human photoreceptors to different wavelengths (left) and the CIEL*a*b* perceptually uniform color space where ellipsis mark perceptually equivalent colors [Pantone, 2018] (right).

2.3.3 Gloss Perception

Gloss, or surface reflectance, is an optical property that defines how much light is reflected from an object into the specular direction. Initial research focused on quantifying gloss as a one-dimensional quantity correlated with physical measurements [Pfund, 1930]. However, in their seminal work Hunter et al. [1937] proposed a gestalt principle for gloss perception and identified six perceptual types of perceived reflectance (Figure 2.16):

- *Specular gloss* - defined as the reflection at specular angles and perceived as the magnitude of brightness.
- *Sheen* - defined as the specular reflection at grazing angles.
- *Contrast Gloss* - defined as the difference between highlighted glossy areas and adjacent material.
- *Haze* - defined as the spread of the specular reflection and associated with "cloudiness" or "milkiness" of the surface.
- *Distinctness-of-Image* - defined as the clarity of an image reflected from the surface.
- *Surface Uniformity* - defined as the absence of visible texture defects and often related to the smoothness of the surface.

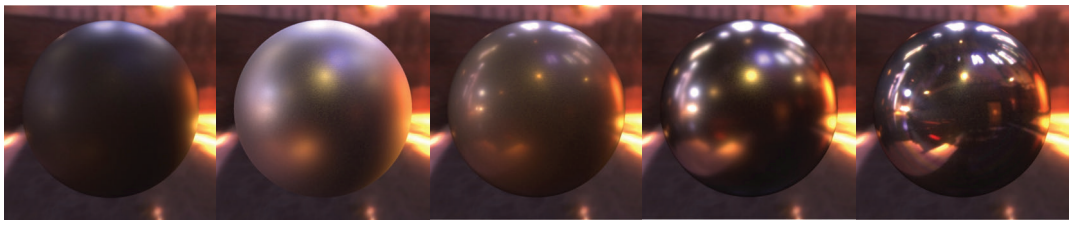

Figure 2.16. Renders of materials from the Merl database [Matusik et al., 2003] manifesting various perceived gloss differences.

Human observers identify gloss based on statistical models of reality [Fleming et al., 2003]. A realistic illumination is required for achieving perceptual constancy [Hartung and Kersten, 2002] and changing the statistical characteristics of the light source directly affects the perceived gloss [Motoyoshi and Matoba, 2012]. As a result, to properly identify gloss a number of properties need to be present. Some of the most prominent ones are relative brightness [Beck and Prazdny, 1981], gloss contrast [Marlow et al., 2012], and the orientation and shape of the glossy highlight [Fleming et al., 2004]. To enhance their sensitivity observers rely on motion [Lichtenauer et al., 2013], the binocular disparity [Wendt et al., 2008], and investigation of complex shapes [Fleming et al., 2004; Fores et al., 2012; Havran et al., 2016] (Figure 2.17). Gloss perception further depends on the diffuse color of the material where a darker color leads to higher perceived glossiness [Pellacini et al., 2000]. Conversely, the gloss of the material has an effect on the perceived color as shiny materials appear to have

more saturated colors [Dalal and Natale-Hoffman, 1999]. For a recent review on cues involved in gloss perception please refer to the survey of Chadwick and Kentridge [2015].

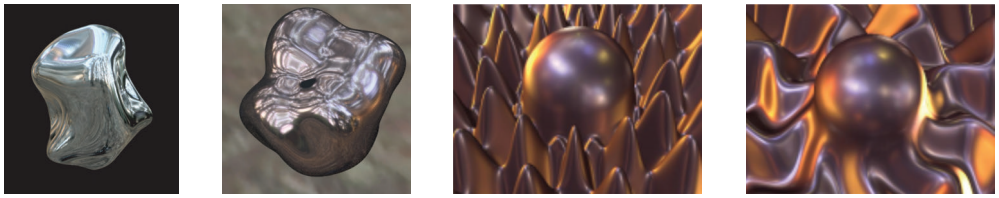

Figure 2.17. Various geometrical shapes used to investigate the perception of gloss [Fleming et al., 2004; Fores et al., 2012; Havran et al., 2016].

The many coupled cues that govern gloss perception provide a challenge for manufacturing hardware. Current fabrication devices operate only with a limited set of materials, leading to a relatively small gamut of gloss reproducible on the device. However, even though previous work showed the multidimensionality of gloss perception, observers are still able to integrate the individual cues into a single gloss rating. As a results attempts to recover a perceptual space of gloss often reveal a single dominant dimension [Pellacini et al., 2000; Wills et al., 2009]. Transferring these concepts from virtual world to the context of manufacturing is key in achieving faithful gloss reproduction.

Chapter 3

Related Work

In this work, we combine insights from psychophysics with numerical optimization to design novel algorithms for haptic reproduction that improve the apparent gamut of a 3D printer. We start by introducing the field of haptic reproduction and the challenges in creating faithful haptic experiences. Next, we link haptic reproduction with computational fabrication and focus on how optimization and numerical simulation can be jointly used to reproduce desired haptic behavior. Afterward, we motivate the feasibility of combining perception with computational fabrication with a review of related work in computer graphics on applying perception for reproducing visual and tactile properties. Finally, we introduce the field of appearance manufacturing and focus on the preliminary work towards full appearance reproduction.

3.1 Haptic Reproduction

The goal of haptic reproduction, or also referred to as haptic rendering, is to enable the user to manipulate and feel virtual objects investigated by touch or through a probe [Hale and Stanney, 2014]. The applications of haptic reproduction range from medicine where doctors can train operations in virtual environments [Chae et al., 2015], through education to teach physics and control [Martinez, Morimoto, Taylor, Barron, Pultorak, Wang, Calasanz-Kaiser, Davis, Blikstein and Okamura, 2016], to the operation of robotic devices where the operator can better control a robotic arm by perceiving the force acting on it [Peshkin et al., 2001; Salisbury et al., 2004], (Figure 3.1).

Some of the first devices that enabled haptic reproductions were interactive Phantoms [Massie et al., 1994], (Figure 3.2 left). To operate such a device the user holds a hand-held probe onto which forces are applied through a set

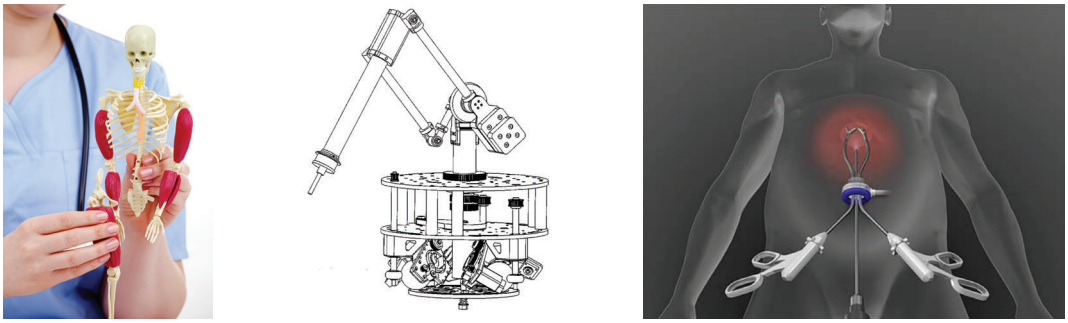

Figure 3.1. Haptic reproduction can help during education, tele-operation of robots [Peshkin et al., 2001], or training for laparoscopic surgery [Romanelli and Earle, 2009].

of pulleys operated by motors. Since a device handles all interactions, its limited force output leads to a poor reproduction of solid surfaces [Diolaiti et al., 2006; Walker et al., 2016]. To create an impression of hard surface so-called haptic displays could be used. The basic idea is to actuate a surface and spatially change its geometry or physical properties to create an illusion of interacting with a real object. The approaches vary from stationary and wearable pin-based displays [Hayward and Cruz-Hernandez, 2000; Mazzone and Kunz, 2005; Klare and Peer, 2015; Perez et al., 2017], (Figure 3.2 middle), to those that rely on micro-vibrations introduced to the surface [Winfield et al., 2007; Bau et al., 2010; Kim et al., 2013], (Figure 3.2 right). To allow interaction without physically touching a display Hoshi et al. [2010] proposed to use inaudible ultrasound speakers. By optimizing the pressure pattern of each speaker a local concentration can be created which can be perceived as a pressure on the hand. This concept can be extended to target multiple locations simultaneously [Carter et al., 2013] or to render three-dimensional shapes [Long et al., 2014]. Recent advancements in virtual reality reinvigorated the development of wearable tactile displays [Leonardis et al., 2015; Schorr and Okamura, 2017b; Schorr and Okamura, 2017a; Kato et al., 2018; Withana et al., 2018], gloves [Stetten et al., 2011; Polygerinos et al., 2015], and exoskeletons [Wehner et al., 2013]. For a more detailed review of recent haptic devices see the surveys of [Chouvardas et al., 2008], [Pacchierotti et al., 2017] and [Perret and Vander Poorten, 2018]. Unfortunately, the mechanical components required to produce faithful haptic feedback also introduce significant latency that deteriorates the overall haptic experience [Massie et al., 1994; Annett et al., 2014].

In this work, we do not tackle the general problem of actively reproducing

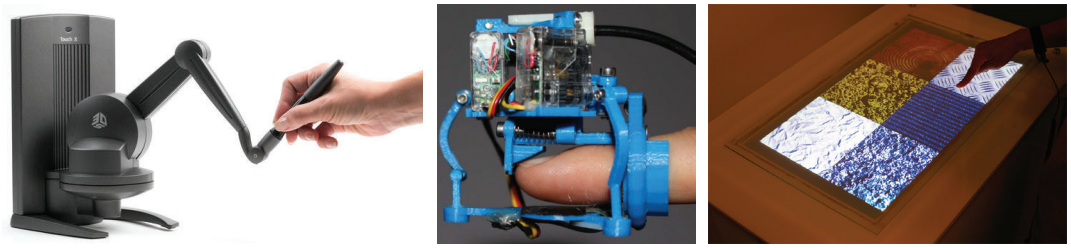

Figure 3.2. Exemplars of haptic feedback devices ranging from haptic phantoms [3D Systems, 2020], through wearable displays [Perez et al., 2017], to static displays based on electrovibration [Bau et al., 2010].

haptic feedback. We focus on *passive* systems that manifest the desired haptic properties. More specifically, we focus on the sensations generated by a surface in combination with a probe that, in our case, mimics a drawing tool. The haptic feedback induced by traditional drawing instruments is vital to our ability to make controlled drawing marks [Danna and Velay, 2015]. Different methods and purposes of the application result in drawing papers with different materials, thickness (weight), and surface quality. Each of which has typical haptic properties. Unfortunately, with average paper roughness of 3 microns [Fischer et al., 2017] one hits limits of current fabrication hardware. Moreover, even the perfect reproduction of the surface does not imply faithful haptic feedback due to complex frictional properties of the interaction between the manufactured substrate and the drawing tool.

3.1.1 Commercial Haptic Styli

The combination of a drawing tool and a paper creates a unique, intimate coupling that produces a characteristic haptic response. Several commercial solutions directly address the challenge of creating digital drawing tools which replicate the sensation and experience provided by the traditional materials, (Figure 3.3). Products such as Microsoft's Surface Pen, Wacom's Intuos, and reMarkable, offer nibs that are designed to replicate different drawing materials. Despite these efforts, users notoriously report lack of proper feel and dissimilarity to the traditional materials [Choi and Tan, 2005; Annett et al., 2014]. Some products try to improve these technologies. The PaperLike cover is a screen protector that significantly enhances the interaction of a stylus with the tablet surface. Other products, such as the Wacom Bamboo, focus on digitization by allowing users

Figure 3.3. Commercial approaches to vary the haptic feedback range from altering the stylus nib [Wacom, 2020], through use of rough foils [Paperlike, 2020], to digitizing the strokes of traditional drawing instruments [Repaper, 2020].

to draw on paper atop a force sensitive tablet. iSkn provides a similar solution, but applies an additional ring, attached to traditional drawing tools that enables digitization of a wider range of instruments. Although these solutions often provide good haptic experience, they do not offer the full range of the advantages of digital drawing.

3.1.2 Research Prototypes of Haptic Styli

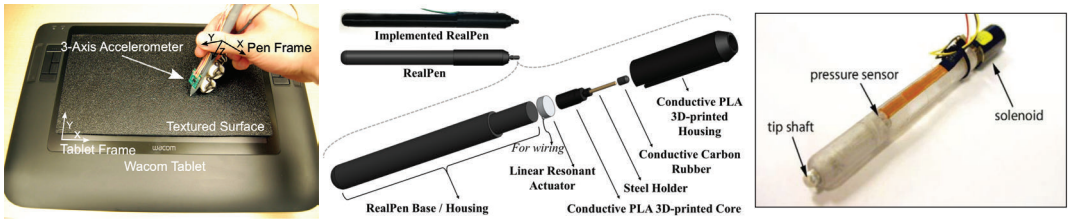

Figure 3.4. Research prototypes rely on active modulation trough vibrations induced by voice coils [Romano and Kuchenbecker, 2012], linear resonant actuators [Cho et al., 2016], or solenoids [Lee et al., 2004].

Researches took an alternative approach and have sought to improve haptic feedback of styli via active pens that are equipped with motors for creating artificial vibrations [Arasan et al., 2013; Lee et al., 2004; Poupyrev et al., 2004], (Figure 3.4). To reproduce the haptic feedback via a vibrational motor it is possible to record the forces during interaction and setup a playback loop [Takeuchi et al., 2012; Saga and Raskar, 2012; Pacchierotti et al., 2016; Cho et al., 2016].

Such a playback loop can manifest a perceivable seam that can be avoided by exploiting the stochastic nature of vibrations. Guruswamy et al. [2009] proposes to capture the haptic feedback by using a generative model based on infinite impulse response filters. Besides creating a seamless playback loop the proposed approach can be adapted to alter the haptic feedback based on the interaction [Romano and Kuchenbecker, 2012; Culbertson et al., 2014]. Romano and Kuchenbecker [2012] used this concept and proposed a fully data-driven system where interactions are recorded via accelerometer, and then, reproduced using a pair of vibration motors based on speed and pressure data provided by the stylus. This system could recreate the haptic sensation of a wide range of different surfaces that can be used to generate pallets of haptic feedback [Meyer et al., 2016]. However, active solutions require additional electronics which complicate the design, restrict the usage of these devices, and introduce significant latency [Burdea, 2000; Annett et al., 2014; Helps and Helps, 2016]. Our system for designing a passive stylus mitigates the problem of additional electronics and provides instantaneous adaptation of feedback to changes in stylus orientation and applied pressure. Additionally, our tools can be easily manufactured using available 2D and 3D printers.

3.2 Computational Fabrication

Figure 3.5. Fabricated metamaterials can be used to modify the compliance [Schumacher et al., 2015], texture, or weight of an object [Torres et al., 2015].

Reproducing desired target behavior is a common goal of computational fabrication methods. Typically such problems are solved by formulating an optimization procedure consisting of three steps: material assignment, numerical simulation, and error estimation [Chen et al., 2013]. Such design loops have been successfully applied across many domains, e.g., generation of caustics [Schwartzburg et al., 2014], elastic properties [Bickel et al., 2010], appearance

[Elek et al., 2017], and sound [Li et al., 2016]. In the context of haptic reproduction the method was applied to create the so-called metamaterials [Bickel et al., 2010], (Figure 3.5). A metamaterial is created by dithering, i.e., precisely depositing, a limited set of available printing materials. The introduced microscopic geometry variation manifests as a macroscopic bulk behavior of a single material [Allaire et al., 1997]. This technique was successfully applied in the design of compliant objects [Schumacher et al., 2015; Panetta et al., 2015; Martínez, Dumas and Lefebvre, 2016; Zhu et al., 2017], surface roughness [Ion et al., 2018; Degraen et al., 2019], and even objects weight [Torres et al., 2015]. Inspired by the previous work on metameterials we similarly do not seek to reproduce the haptic feedback by creating carbon copies of materials. Instead, we rely on formulating the optimization process in the perceptual-space. This allows us to focus the computational effort in designing materials that *feel* similar to target haptic properties but manifest drastically different geometries that are fabricable on 3D printers. Successful application of this idea hinges on the ability to efficiently simulating the haptic behavior. Since estimating haptic feedback is computationally expensive it has to be supported with efficient optimization and parametrization techniques.

3.2.1 Numerical Simulation of Haptic Feedback

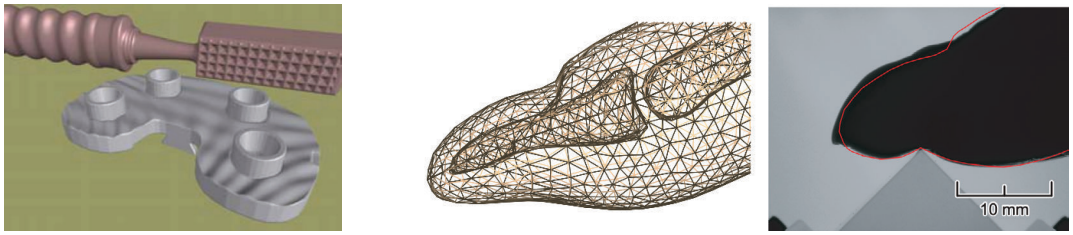

Figure 3.6. Numerical simulation of haptic feedback ranges from relatively simple contact of rigid bodies [Otaduy et al., 2004] to simulating the frictional interaction between a deformable finger and a surface [Tada and Pai, 2008].

To reproduce haptic feedback of virtual objects it is possible to use numerical simulation, (Figure 3.6). The general idea is to first simulate the contact in the virtual world and then use the appropriate forces during optimization [Kry and Pai, 2006]. Such techniques can be used to simulate both interactions with a finger [Tada and Pai, 2008], as well as, artificial probes [Otaduy et al., 2004]. The accuracy of the simulation can be increased by using more detailed biological

models of human tissue [Perez et al., 2015; Tymms et al., 2018], and by including the hardware [Perez et al., 2017] and perceptual [Zhang et al., 2017] limitations into the design loop. Unfortunately, the scale and complexity of the governing physical phenomena, coupled with fabrication constraints inhibit an accurate and efficient numerical solution [Myshkin et al., 1998; Cirac and Zoller, 2012]. In our work we build upon these ideas. To reduce the simulation complexity we leverage the limitations of the human sensorial system and coarsen the numerical model. This allows us to accelerate the simulation speed while maintaining the perceived accuracy.

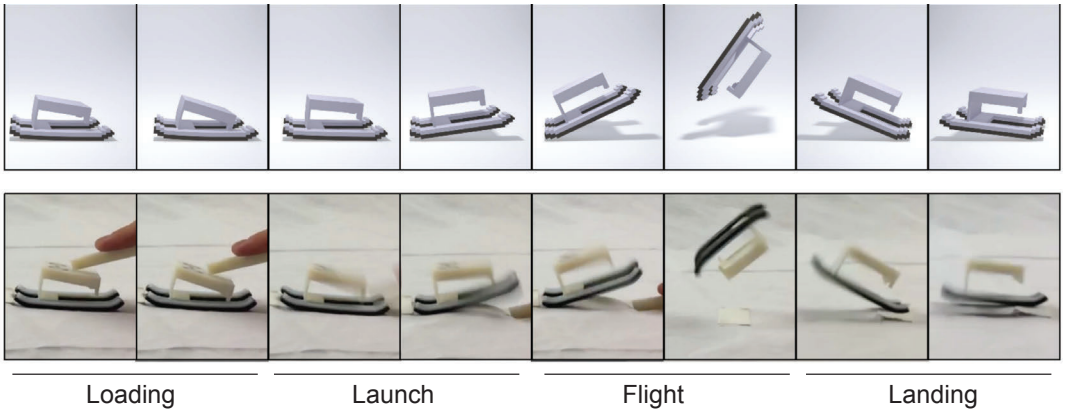

Figure 3.7. Robust handling of frictional contact allows us to predict the behavior of an elastic object during interaction [Chen et al., 2017]. Such predictions can be later used to optimize the flight path of the object.

The above simulations are based on modelling the contact between an object and a finger/probe. Standard contact models generally begin with Coulomb's friction model [Harnoy et al., 2008] – an assumption of frictional resistance proportional to normal load and in opposition to velocity. These assumptions are often reasonable for solid-to-solid frictional contact. In the presence of lubrication, however, viscosity effects become important. The Stribeck effect [Harnoy et al., 2008] describes frictional behavior in lubricated contact. Initially, with increased speed, friction forces drop. As speed is further increased the Stribeck model assumes an increase in friction linearly proportional to speed. Frictional forces can generate nonsmooth stick-slip behaviors. Karnopp [1985] models a regularized stick-slip behavior that removes nonsmoothness from the Coulomb model, simplifying numerical integration. A wide range of more complex frictional models [Armstrong-Hélouvry, 1991; Canudas de Wit et al., 1995; Dahl, 1976] are available and range in suitability depending on modeling needs. In

our work the particular problem, that of stylus-surface interaction is further complicated by the effects of material wear [Stachowiak, 2006], material deposition [Archard, 1953] on the surface, and the coupling between friction and contact with elasticity and viscous damping [Chen et al., 2017]. Applying a standard frictional contact modeling is not sufficient to obtain accurate numerical estimates of stylus behavior. Our simulator builds on impulse-based frictional contact modeling [Mirtich and Canny, 1995]. This remains an active research area in simulation [Kaufman et al., 2005, 2008; Bertails-Descoubes et al., 2011] and many of the additional complicating factors described above are not yet well understood in this context. To overcome this difficulty, we opt for an exponential integrator [Hochbruck and Ostermann, 2010; Michels et al., 2014] with a data-driven surface-interaction model. In combination, our exponential integration gives us an analytical solution to the linear elastodynamic equations of motion while achieving excellent match with experimental data captured during stylus-surface interaction.

3.2.2 Optimization Algorithm

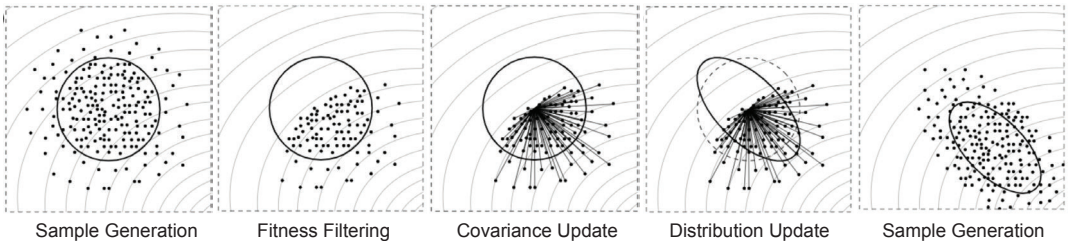

Figure 3.8. Black-box function optimization using the CMA-ES method from [Fujii et al., 2018]. The algorithm starts by randomly sampling from a uniform distribution. The samples are then filtered based on fitness value. From the filtered samples a new covariance matrix and sample distribution is generated which is then used to sample the next generation.

Achieving target behavior requires minimizing an expensive, possibly non-convex and derivative-free function that involves numerical simulation. Finding the global optima in such a setting is an open and challenging problem [Nocedal and Wright, 2006]. Researchers are often forced to resort to simulated annealing [van Laarhoven and Aarts, 1987] or CMA-ES style of algorithms [Hansen et al., 2003], (Figure 3.8). An alternative solution is to replace the expensive function with a cheaper but effective surrogate, see e.g., [Jones et al., 1998]. A particularly popular option of the surrogate model is Gaussian Processes [Rasmussen

and Williams, 2005], which not only provides derivative information but can also be used to estimate the certainty of the prediction. This certainty estimate can be exploited by formulating an acquisition function that samples the surrogate model to maximize the information gain with each sample [Frazier et al., 2009; Hennig and Schuler, 2012; Mockus, 1989]. Our work builds upon these ideas and adopts them in the context of computational fabrication by proposing a custom acquisition function suitable for haptic reproduction.

Our fabrication-in-the-loop approach draws inspiration from active learning, where the learning process is enhanced by engaging users to label new datapoints iteratively. Also here, surrogate models like Gaussian Processes are often used to generate the datasets presented to the participants [Akrour et al., 2011; Dudley et al., 2019; Koyama et al., 2017]. Our method can be seen as an instance of active learning where the user queries are replaced by an oracle based on physical manufacturing and measurements of samples generated based on a surrogate model formulated using Gaussian Processes.

3.2.3 Parametrization

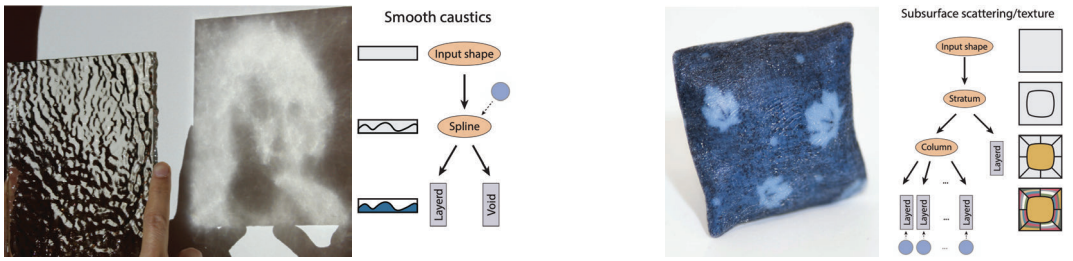

Figure 3.9. Selecting an appropriate parametrization can significantly facilitate the optimization problem [Chen et al., 2013]. Modelling the object as smooth bezier curves is advantageous for caustic optimization, whereas for texture design it is more appropriate to model individual voxel elements.

The resulting optimization problems pose significant computational challenges as the use of numerical simulation often prevents direct computation of analytical gradients of the objective function. Consequently, many computational fabrication methods rely on numerical approximations or more costly stochastic optimization procedures. To make the optimizations tracktable, the choice of suitable parametrization of the design space becomes critical, (Figure 3.9). A good parametrization can significantly facilitate the solution of high-dimensional

non-convex problems [Bharaj et al., 2015] or even cast them to convex subspaces [Piovarči et al., 2017]. Additionally, it may influence the durability of the manufactured objects [Panetta et al., 2017] and improve the expressiveness of the design space [Martínez et al., 2019]. Similarly, in this work, our problem leads to a non-trivial optimization procedure, where good parametrization becomes critical for capturing a sufficient range of haptic feedback.

3.3 Perception in Computer Graphics

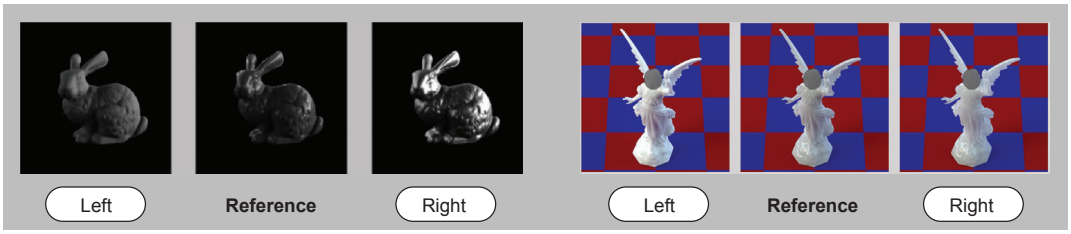

Figure 3.10. Experimental designs investigating the perception of gloss [Wills et al., 2009] (left) and translucency [Gkioulekas et al., 2013] (right).

In computer graphics, perceptual spaces have been widely applied to analyze haptic perception, e.g., roughness [Bergmann Tiest and Kappers, 2006; Tymms et al., 2018], compliance [Tiest and Kappers, 2009], and shape and texture [Cooke et al., 2006, 2010]. Beyond haptics, the methodology has been applied to understand the perception of reflectance properties, such as gloss [Pellacini et al., 2000; Wills et al., 2009], and translucency [Gkioulekas et al., 2013]. Pellacini et al. [2000] use standard multi-dimensional scaling (MDS), which is based on the distances between different images reported by a subject, to learn a two-dimensional space of bi-direction reflectance distribution functions (BRDF). Wills et al. [2009] use a non-metric multi-dimensional scaling (NMDS) to build a similar space but for a larger range of BRDFs, (Figure 3.10 left). In contrast to MDS, NMDS does not require the magnitude of the dissimilarities between BRDFs, but only their ordering. This helps in reducing the inter- and intra-subject variance observed in user studies. In their work on transparency perception, Gkioulekas et al. [2013] use both MDS – to build a space of translucent materials based on an image metric, and NMDS – to obtain a similar space from user data, (Figure 3.10 right). Like these works we also build perceptual spaces for haptic feedback using non-metric MDS. To find the space, we adapt a likelihood maximization method [Silverstein and Farrell, 2001] which was proposed for scaling

one-dimensional data. In contrast to previous work which treats correlating the perceptual axes with physical properties as a separate step, we propose to modify multi-dimensional scaling to jointly optimize for a perceptual space which both explains the experimental data and correlates with a set of physical properties.

Perceptual models can be also used to facilitate the production process. Serrano et al. [2016] proposes a perceptual encoding of reflectance functions to enable material change along intuitive dimensions, e.g., from plastic-like to metallic-like. Lagunas et al. [2019] developed a similarity measure of material appearance that can be used to populate a scene with visually distinct materials. Sigal et al. [2015] proposed to replace the physical parameters of garment simulation with perceptual alternatives. Each perceptual attribute operates on multiple physical ones to achieve a desired intuitive change in material properties e.g., changing the wrinkles of the cloth. Besides more intuitive user interfaces, the auditory perception of different materials has been investigated to improve sound rendering techniques [Ren et al., 2013]. The stream of related work where perception is used to improve the algorithmic side of computation is a motivation for our approach. Similarly to previous work, we seek to include perception into the computational model and bypass both numerical and hardware limitations.

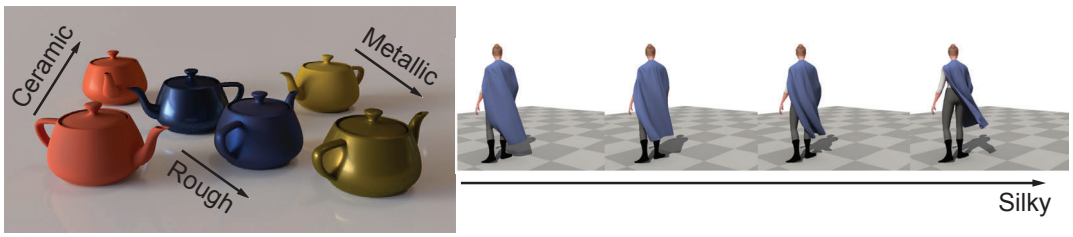

Figure 3.11. Applications of perception in creating more intuitive user interfaces for material [Serrano et al., 2016] and cloth simulation editing [Sigal et al., 2015].

3.4 Appearance Fabrication

Fine control of objects' appearance is the desired functionality of both display devices [Hullin et al., 2011] and fabrication processes [Hullin et al., 2013]. Recent developments in 3D printing and computational fabrication have enabled the fabrication of objects with prescribed color, translucency, and reflectance. Still, many limitations persist. In this section, we provide an overview of recent techniques for appearance fabrication with a focus on reflectance and gloss

properties.

3.4.1 Color Reproduction

Color is the primary appearance property of an object. The state-of-the-art color 3D printing solutions rely on UV curable ink-jet printers [Sitthi-Amorn et al., 2015]. Similar to 2D printers, they also mix several base materials, usually CMYKW, to achieve full-color. There is, however, a much broader choice in terms of materials and mixing strategies. The most common techniques perform halftoning of semi-opaque printing materials [Brunton et al., 2015], (Figure 3.12 left). Color printing is also possible using only transparent materials that are stacked to produce the desired color [Babaei et al., 2017], (Figure 3.12 middle). Recently, such a technique was also proposed for spectral color reproduction [Shi et al., 2019]. Due to significant scattering properties, special strategies were proposed to compensate for the effect [Elek et al., 2017; Sumin et al., 2019], (Figure 3.12 right). Besides ink-jet printers, color can be produced using different technologies such as paper lamination, powder-binder, or fused filament fabrication, but the quality does not match that of ink-jet printing. Color of an object can be also modified in a post-processing step where the ink is transfered onto the object surface through a thin water soluble film [Panozzo et al., 2015; Zhang et al., 2015] or a thermoformed plastic sheet [Schüller et al., 2016].

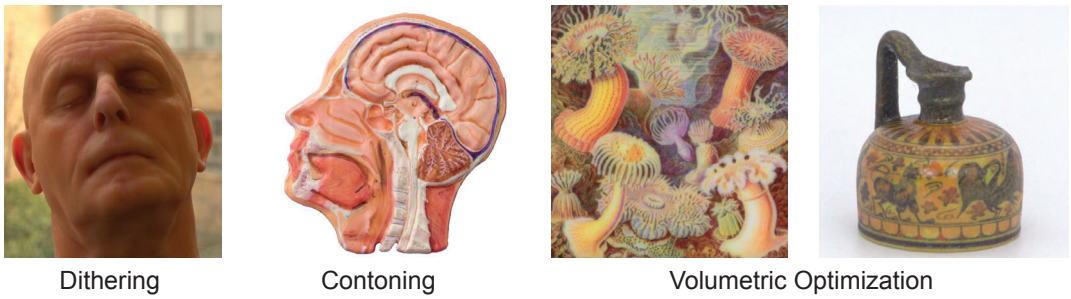

Figure 3.12. Differen strategies for color optimization for inkjet printers. From left to right: halftone dithering [Brunton et al., 2015], color contoning [Babaei et al., 2017], optimization using rendering predictions [Elek et al., 2017; Sumin et al., 2019].

3.4.2 Translucency Reproduction

Apart from color and reflectance properties, translucency is another factor influencing the appearance of objects. Early works try to optimize the material dis-

tributions withing the printing volume to match the subsurface scattering properties [Hašan et al., 2010; Dong et al., 2010], (Figure 3.13, right). For homogeneous silicon objects, Papas et al. [2013] determine a mixture of pigments for a given color and translucency reproduction, (Figure 3.13, left). Most recent work extends a color printing pipeline to translucency by incorporating transparent material [Brunton et al., 2018]. By reinterpreting the alpha channel from additive blending to subtractive mixing the translucency of the material can be encoded in conventional RGBA textures [Urban et al., 2019]. Our work does not address color or translucency printing but assumes that the substrate of the objects contains color information produced by one of the existing techniques. We solely focus on adding a layer of a material to alter the gloss of the object.

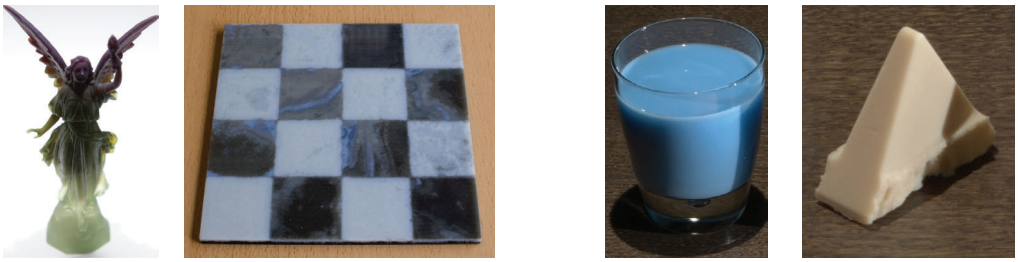

Figure 3.13. Reproducing target translucent properties with an inkjet printer [Hašan et al., 2010; Urban et al., 2019] and by optimizing colored silicone mixtures [Papas et al., 2013].

3.4.3 Reflectance Fabrication

General techniques tackling the problem of appearance try to reproduce BRDF of a given surface accurately. One class of methods try to achieve this goal by modifying the microgeometry of the surface, (Figure 3.14 left). The idea follows the microfacets theory [Cook and Torrance, 1982], where a surface is assumed to be a composition of tiny reflecting planes. By controlling their normal distributions, different BRDFs can be achieved. Several techniques realize this idea [Weyrich et al., 2009; Rouiller et al., 2013; Piovarči et al., 2017]. Unfortunately, the most prominent limitation is the scale at which the BRDF properties can be modified and varied. An interesting approach was proposed by Levin et al. [2013]. Using a high-precision fabrication process realizing features at $2-3$ μm, they demonstrated a reproduction of high-resolution spatially varying BRDF properties. Despite the high-quality results, the technique requires special fabrication

facilities, and therefore, is not suitable for combining with 3D printing.

More similar to our approach is the work of Matusik et al. [2009], (Figure 3.14 middle). Instead of modifying the local geometry, they proposed to use a broad set of inks that span a wide range of different reflectance properties. Their work presents a complete system for printing digital mixtures of inks to achieve desired, spatially-varying 2D surface appearance. Mixing different inks on top of the object's surface can be combined with surface geometry optimization [Malzbender et al., 2012; Lan et al., 2013], (Figure 3.14 right). At each spatial location, the desired BRDF is approximated with a height-field coated by dithered inks. The coupled optimization enables the generation of anisotropic samples and widens the range of available reflectance properties. The main drawback of these methods comes from the coupling between color and gloss fabrication. The appearance reproduction has to be formulated as a joint optimization, which leads to a large set of base inks, and therefore, complicates the hardware design. Additionally, since materials required for reproducing matte appearance contain larger particles and are more viscous, they cannot be deposited at very high resolutions. While this is not critical for the gloss reproduction where high resolution is not of very high importance, the problem reduces the quality of color reproduction. In contrast, our method decouples the manufacturing of color and gloss by depositing varnishes on top of a colored surface. The color can be produced using a commercial ink-jet printer at a resolution of 600 DPI or higher. We can then modify the gloss at a lower resolution, i.e., 79 DPI using our hardware.

Reflectance variation can be also achieved by modifying the printing parameters of the available materials. In the context of 2D printing, plotting with glossy varnishes has been a subject of investigation. Such a system was presented by Baar et al. [2014]. They used a 2D plotting system for multi-pass printing. By varying the order of inks, drying time, and varnish coverage, they were able to introduce a roughness to the surface, which resulted in different gloss levels. The idea can be abstracted to stereolitography where the surface roughness can be achieved by using sub-voxel growth of the printing resin [Luongo et al., 2019]. Unfortunately, achieving matte finish using such techniques with a glossy varnish is challenging [Samadzadegan et al., 2015]. Within the same body of work Elkhuizen et al. [2019] proposed to print multiple layers of UV-curable transparent material on top of glossy surface finish. The geometry created during this process introduces roughness to the surface, which results in a more matte appearance. They demonstrate capabilities of achieving reflectance properties ranging from 85 to 4 gloss units when measured with a glossimeter at 60-degrees angle. Although the method pushes the hardware capabilities to its

limits, the authors note that for a full gloss reproduction of paintings an even more matte finish is required. Comparing to this work, our technique enables the reproduction of a large range of surface finishes with continuous variation from highly matte (0.9 gloss units at 60-degrees) to highly glossy (87 gloss units at 60-degrees) without any surface modification. We achieve this by proposing a system that is capable of precise deposition of matte varnishes.

Figure 3.14. Different approaches to reflectance manufacturing. From left to right: reflectance reproduction via ink deposition [Matusik et al., 2009], micro-facet optimization of target gloss [Rouiller et al., 2013], and bi-scale approach combining ink and geometry modification for appearance reproduction [Lan et al., 2013].

Chapter 4

Perception-Aware Computational Fabrication for Haptics

The goal of haptic fabrication is to manufacture objects that manifest the desired haptic *feel*. To achieve this goal we are interested in investigating complex perceptual sensations governed by multiple potentially coupled physical phenomena. Imagine we would like to reproduce the feeling of touching a patch of leather which is a common problem in the automotive industry [Stoll and Cavalcante, 2011]. The feeling of leather is not governed by a single attribute but rather multiple cues (e.g., roughness, stickiness, warmth) combine to form the overall percept. How the individual stimuli combine is apriori unknown and attempting to fully reproduce leather-like material on a 3D printer is technologically challenging since many of the physical phenomena involved in haptic exploration occur at the atomic scale. Similarly to standard parametrize-simulate-evaluate approach we tackle this problem numerically by formulating a minimization:

$$\arg\min_{\mathbf{p}} P(f(\mathbf{p}), T), \tag{4.1}$$

where $\mathbf{p}$ are the design parameters governing the problem domain, $f()$ is a function mapping the design parameters to physical measurements e.g., estimating the roughness of a procedurally generated texture, T is the desired target haptic behavior, and $P()$ is a perceptual error metric.

The key challenge lies in the formulation of the error metric. To this end, we propose to build a so-called perceptual space. Perceptual space is a multi-dimensional embedding in which the Euclidean distance between samples corresponds to perceived difference. We investigate how to build such a space in Chapter 5 for the case of compliance perception. To observe how participants

perceive elastic samples we rely on psychophysical experiments (Chapter 5.2). More specifically, we propose to use the two options forced-choice experiment design. The participants are presented with a reference sample and two possible reproductions. Their task is to pick the reproduction that feels more similar to the reference. To turn the output of such a study into a perceptual space we employ Non-Metric Multidimensional Scaling (NMDS) (Chapter 5.2.4). After the numerical analysis with NMDS, we obtain a one-dimensional perceptual space for compliance. However, the axis of the recovered perceptual space is unknown. We test a set of candidate computational models that relate the physical measurements of elastic samples with their perceived compliance (Chapter 5.3). From the proposed models we select the best performing candidate and use numerical simulation to extend it to handle arbitrary 3D models (Chapter 5.4). Finally, we show applications of our perceptual error metric to improve the accuracy in material selection (Chapter 5.5.1), to create more intuitive design tools (Chapter 5.5.2), and to achieve better reproductions of objects with prescribed compliance than state-of-the-art methods (Chapter 5.5.3).

Studying the application of perception-aware fabrication for the case of compliance reveals three additional challenges in applying the methodology that are crucial to address in order to tackle more complex problems. We investigate these issues in Chapter 6 to design styli with prescribed haptic behavior. The first challenge lies in the psychophysical experiment design. By considering triplets of samples we run into a combinatorial increase in complexity for each new stimulus included in the study. To reduce the number of samples we propose an exploration-exploitation strategy for the experiment design that focuses the data-collection to samples that provide more information about the problem (Chapter 6.3.2). The second challenge is recovering the physical properties that define the axis of a perceptual space. Even a relatively simple one-dimensional space required significant experimental effort. We address this issue by formulating a probabilistic NMDS that optimizes for a perceptual space that both explains the experimental data and correlates with measurable physical attributes (Chapter 6.3.1). Lastly, to map new designs into the perceptual space we rely on a numerical simulation that when modeling more complex physical phenomena quickly becomes intractable. To improve the speed of numerical simulation we propose to use perception-aware coarsening of the underlying numerical model. We exploit the limitations of the human sensorial system to construct an efficient integrator optimized for simulation of a drawing tool that interacts with a paper substrate. Due to the complexity of the governing physics, we numerically model only the stylus and rely on data-driven techniques to model the paper substrate (Chapter 6.4). These improvements enable us to construct a perceptual space of

drawing tools that we use to design new digital styli which feel more similar to traditional drawing instruments than currently available commercial solutions (Chapter 6.5).

In Chapter 7 we demonstrate the possibility of increasing the *apparent* gamut of a 3D printer by manufacturing digital replicas of traditional drawing instruments. The haptic feedback of a drawing tool depends on the coupled intimate interaction between a drawing tool and the underlying substrate. This interaction occurs at such a fine scale that it is not possible to directly reproduce on a commercial 3D printer. Here we utilize the perceptual space of drawing tools from Chapter 6. The perceptual space reveals that drawing tools can be compared based on their perceived friction and perceived vibration. We exploit this insight to design a parametric space of stylus-surface combinations that capture a wide range of haptic behavior (Chapter 7.2). Unfortunately, to explore the design space we cannot rely on our coarsened numerical model from Chapter 6 since now the drawing surface is an unknown and it is not feasible to build a large enough dataset. Interestingly, the underlying physical problem is so complex that its numerical simulation approaches the manufacturing time. We utilize this insight to formulate a fabrication-in-the-loop optimization (Chapter 7.3). Such optimization requires us to take special care in sampling the design space due to the slow evaluation of the objective function that includes fabricating and measuring physical samples. To this end, we propose to use a surrogate model that for a design from the parametrization predicts its perceived friction and vibration (Chapter 7.3.2). A key property of the model is that it gives us confidence bounds on the predictions. We use these confidence bounds to formulate an acquisition function that efficiently samples the designs space by maximizing the expected improvement towards target behavior (Chapter 7.3.3). We show how such a model can be used to design stylus-surface pairs with prescribed haptic behavior (Chapter 7.4). We validated the quality of the reproductions in blind experiments with casual participants (Chapter 7.6), as well as, in a survey with professional artists (Chapter 7.7).

Perception-aware fabrication has applications beyond haptic reproduction. Our approach can be employed in other areas of fabrication where the resulting quality is judged by a human observer. The most prominent example in this context is the field of appearance reproduction that focuses on controlling the full visual impression an object creates. Unfortunately, there is not yet an established technique capable of reproducing the full appearance of a 3D object, (Chapter 3.4). Commercial inkjet printers are capable of manufacturing models with high-quality color and translucency information but are limited to two gloss finishes: glossy and matte. This limitation is a result of the hardware design

of inkjet nozzles that can process only a narrow range of viscosities and particle sizes effectively rendering material based gloss modification infeasible. We tackle the problem of full appearance reproduction in Chapter 8 by proposing a novel printing hardware capable of jetting off-the-shelf varnishes, i.e., viscous materials with relatively large particles. To fabricate an object with prescribed appearance we first reproduce objects' color on a commercial inkjet printer and then modify its gloss as a post-processing step with our varnish printer. To characterize the proposed printing hardware we investigate the effect of printing parameters on generated droplet quality, (Chapter 8.3). Using optimized printing parameters we jet a number of off-the-shelf varnishes and quantify the gloss gamut achievable by our device using physical measurements, (Chapter 8.4). Based on the measured gamut we identify three basis varnishes for our hardware setup. To achieve a larger range of gloss properties we spatially halftone the varnishes. Since the appearance of a halftoned varnish can not be explained as a simple linear mixture we propose a data-driven simplex-interpolation scheme. We enhance the halftoning with a dithering pattern quality predictor based on observed varnish mixing that minimizes the undesired halftoning artifacts, (Chapter 8.5). We validate the quality of our printing setup in simulation and by manufacturing several 2D and 3D objects with spatially-varying gloss, (Chapter 8.6). The proposed hardware is a step towards full appearance reproduction that will allow us to quantitatively and qualitatively investigate how we perceive materials in the real world under natural illumination.

Chapter 5

An Interaction-Aware, Perceptual Model for Non-Linear Elastic Objects

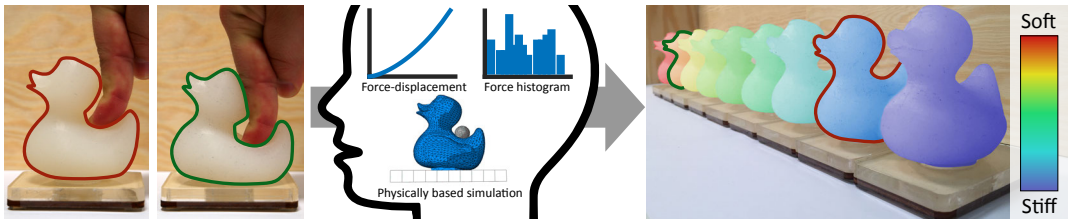

Figure 5.1. 3D printing allows us to print objects with varying deformation properties. The question that we want to answer is: Given a set of printing materials and a 3D object with desired elasticity properties, which material should be used to print the object? For example, given sample ducks (left) with desired elasticity properties (e.g., measured), our system considers several candidate materials that can be used for replicating the ducks (right), and chooses materials that will best match compliance properties when examined by an observer (red and green outlines). Moreover, we can sort all possible materials by their perceived compliance as predicted by our model. The measured compliance is indicated with colors ranging from stiff (blue) to soft (red).

The elasticity of objects is imporant for everyday interactions from a shopper buying a new pillow, to a doctor selecting shoe inlets for his patients. Unfortunately, the material gamut of individual printing platforms is limited and, just as a 2D printer with a limited color gamut can introduce artifacts into an image, material limits can cause unintended changes in the haptic properties of a 3D printed design. Minimizing these artifacts requires a method for optimally pro-

jecting a design onto the material space of a given printer, and this projection requires a suitable metric. Currently, people use standard metrics to quantify differences in material mechanical properties. The typical approach is comparing L2 norms of stress-strain curves [Bickel et al., 2010]. However, such metrics ignore a major component of haptic interaction, the users themselves. Previous work which explores similar problems for visual output devices has found that perceptual metrics can outperform standard measures [Morovic and Luo, 2001; Reinhard et al., 2010]. In this work, we seek to find a computationally efficient and accurate perceptual model for compliance of 3D printed objects that will allow us to compare them in a meaningful way.

Newer 3D printers, such as the Objet500 Connex Series from Stratysys, are not limited to printing rigid materials. Previous work [Bickel et al., 2010; Schumacher et al., 2015; Panetta et al., 2015] has demonstrated that, by varying the internal structure of an object, so-called metamaterials can be created. These metamaterials increase the gamut of printable, non-linearly elastic materials available. Correspondingly, we do not limit our investigations to rigid or linearly elastic materials but instead develop a metric suitable for nonlinearly elastic materials for which the internal forces are an arbitrary function of the applied deformation.

In order to design a suitable metric for comparing compliance of different materials, as perceived by a human, we first compute a perceptual space – a space wherein the Euclidean distances between samples correspond to magnitudes of perceived differences. Inspired by [Pellacini et al., 2000; Wills et al., 2009], we construct a pointwise approximation by applying multi-dimensional scaling to a predefined set of compliant 3D printed metamaterials. This, however, does not provide us with a closed-form transformation between different materials and the perceptual space, which is desired for many applications. Therefore, we evaluate a number of candidate computational models, and propose one that provides a mapping that best describes the empirical data. While the psychophysical experiments for evaluating the models are performed on simple-shape stimuli, we also demonstrate how to apply our model to complex geometries such that it accounts for geometrical variations in designs. Finally, we present several applications of our model in the context of 3D printing. To summarize, our contributions are:

- a psychophysical experiment evaluating compliance of different metamaterials produced using a 3D printer together with a corresponding perceptual space,
- new computational models for perceived compliance that incorporate force information,

- evaluation of these and previously proposed models based on our empirical experiment,
- an extension of the model derived for a simple shape to complex geometries,
- several examples of applying the model for 3D printing complex objects with different compliance properties, and
- a validation of our results in a series of user experiments.

5.1 Overview

The goal of this work is to relate the physical compliance of an object to its perceived compliance when examined by a human. Since physical stiffness can be expressed using measured force-displacement data, we consider this as the main cue for compliance. Consequently, we seek a relation between the force-displacement characteristic and the *feeling* of compliance. To determine this relationship we first perform a psychophysical experiment with simple 3D-printed stimuli (small cubes), for which we can easily vary the force-displacement characteristic (Section 5.2). Based on the perceptual space of the stimuli derived using non-metric multi-dimensional scaling, we evaluate various computational models that predict perceived compliance from the force-displacement characteristic (Section 5.3). Once a suitable one-dimensional relationship has been identified, we show how numerical simulation can be used to extend the model to complex geometries with which users interact via poking (Section 5.4). To do this we exploit the fact that applying a virtual point load to a simulated object yields a local, 1D, force-displacement curve, to which the computational model is applied. As a result, a spatially-varying map of the perceived compliance is obtained. Finally, we present several applications of our perceptual model in the context of 3D printing and rapid prototyping (Section 5.5).

5.2 Perceptual Space of Stiffness

Our perceptual compliance space is computed using non-metric multi-dimensional scaling [Wills et al., 2009] using data of paired comparisons between samples. Below we describe, in detail, the stimuli, experiment, and the analysis used for deriving the space.

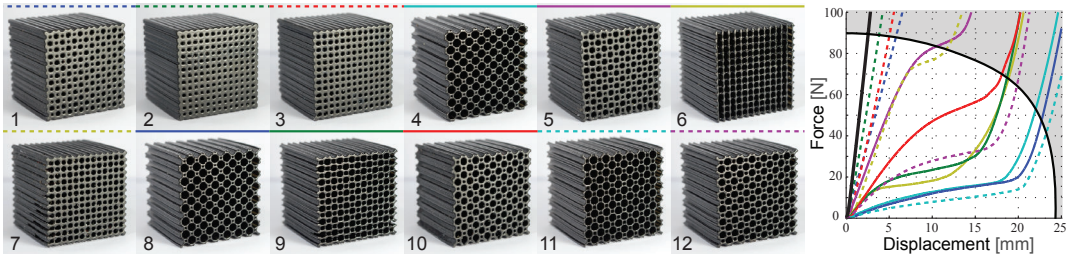

Figure 5.2. The 12 stimuli that were used in our experiments. The plot on the right shows corresponding force-displacement curves obtained via uniaxial test. The lines above blocks indicate the corresponding line styles. For a reference, we also include the force-displacement curve for a solid block (black line). The gray area roughly marks maximum applied force to each sample.

5.2.1 Stimuli

Current 3D printers use a relatively small number of printing materials, which restricts the spectrum of deformation properties that can be produced. Our experiment requires material samples with diverse mechanical properties. To counteract the limitation of 3D prints, we fabricate metamaterials with different internal structures. As presented in Figure 5.3 (left), our samples were designed as cubes with cylindrical holes of two sizes, spaced regularly across the sample. Each sample was defined by 4 parameters: block size, distance between centers of the cylinders, and two radii for the cylinders. For our experiment, we limited the design space to 42 mm blocks where centers of circles were kept at a constant distance of 3 mm. Different stiffness was achieved by varying the two radii parameters, which were kept between 0.5 mm and 2 mm due to manufacturing limitations. After accounting for the fact that the cylinders cannot intersect, we sampled our design space uniformly (Figure 5.3, right). We designed 12 different blocks in total, and printed them using TangoBlack+ material on a Stratasys Objet500 3D printer. Similar, foam-like metamaterials were used by [Bickel et al., 2010] for printing objects with desired deformation properties.

To characterize the material properties of each block, we perform uniaxial load testing. An increasing force (up to 100 N) was applied to each block at constant speed of 0.1 mm/s, and the corresponding deformation was recorded. Figure 5.2 shows all the samples used in the experiment along with their corresponding force-displacement curves.

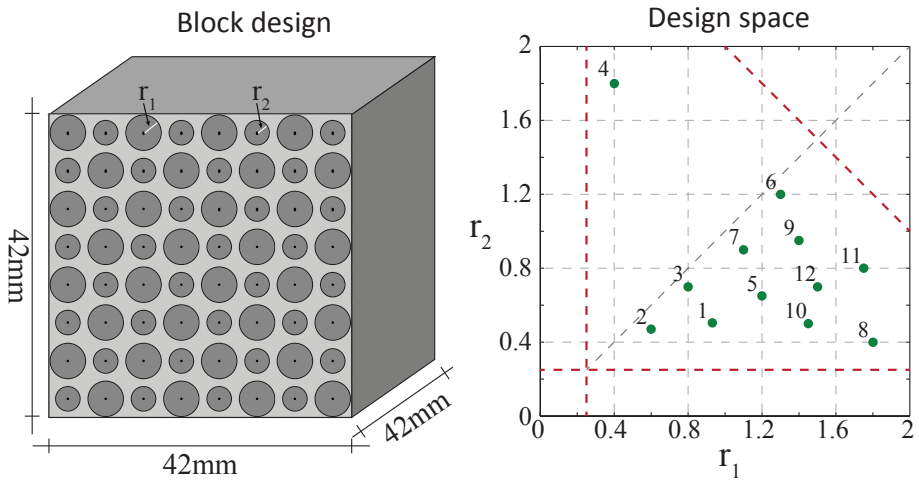

Figure 5.3. Left: Design of our samples. Equally spaced out cylindrical holes are defined using two radii. Right: Parameter of our samples. The two axes correspond to two radii of the cylinders (in mm). 12 green dots correspond to the samples used for our experiment. The red lines indicate the lower limits for the values of the radii, due to the printer resolution, and the upper limit for radii to guarantee that the cylinders do not intersect. Block 4 was introduced as a counterpart of block 8 to verify whether compliance properties are symmetric with respect to the choice of radii.

5.2.2 Methodology

During each trial of our experiment, participants were presented with one reference block and two test blocks arranged in a row. They were asked to judge which test block was more similar to the reference in terms of softness. Participants were asked to interact with the samples in a direction perpendicular to their surface. This interaction mode removes the effect of anisotropy present in the fabricated cubes. No limitations on the number of interactions during a single trial were imposed. Visual feedback was avoided by placing all blocks under a cover. The experimental setup is presented in Figure 5.4. The 12 blocks considered in our experiment resulted in 660 different possible trial combinations. The study was performed by 20 subjects. Each subject was presented with 78 trials which resulted in a total of 1,560 comparisons. The average duration of the experiment was about one hour. Participants were allowed to take a break halfway through the experiment.

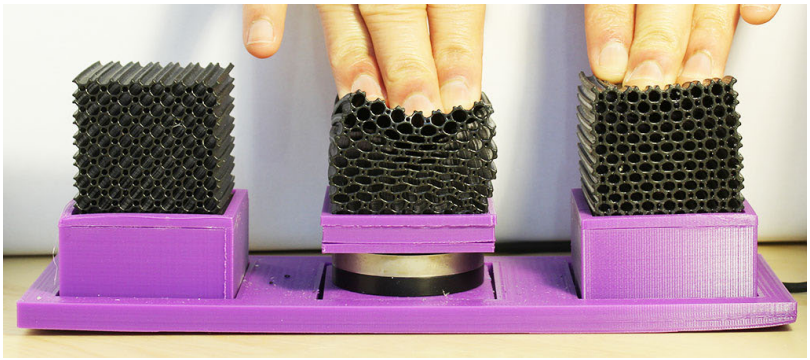

Figure 5.4. In a single trial participants were presented with three different samples. Their task was to compare softness difference between the leftmost and the rightmost pair. Additionally, we placed a force sensor beneath the middle block in order to collect force data throughout the entire experiment.

5.2.3 Data Reliability

We evaluated the quality of our data using three additional experiments. First, we evaluated the intra-subject consistency. We presented a random set of 9 triples, 4 times each to every participant. On average, in 72.22% of the cases subjects gave the same answer. Next, to test the inter-subject variability, we asked all participants to perform one trial on the same set of randomly chosen 36 triplets. On average, 93.88% of all responses were consistent with majority votes. In the last step, we checked whether subjects could provide consistent ordering of materials, i.e., we sought to avoid orderings such as $d_{ij} < d_{jk} < d_{ik} < d_{ij}$, where d_{ab} is the perceived difference between samples a and b. We randomly chose a set of 12 triples. Subjects were asked to evaluate each of the 3 distance comparisons. We found no cycles in this data.

5.2.4 Non-metric MDS

In order to find a Euclidean embedding of our metamaterials, we employed non-metric multi-dimensional scaling (NMDS). We use the version proposed by Wills et al. [2009] for building a perceptual space for gloss. For more details please refer to Chapter 2.

5.2.5 Analysis

The NMDS optimization searches for a valid embedding with minimal dimensionality. This minimization is driven by a parameter λ that is used to tune the

tradeoff between the number of dimensions and explanation of experimental data. To find a space that explains both the observed and the unobserved data well, we perform a cross-validation [Hastie et al., 2009]. At each step we divide the constraints collected in the experiments into two sets: training and testing. Next, we use the training set to optimize for the embedding. In order to find out how well the trained model explains the training data, we compute a training error by counting how many pairwise comparisons from the training set are violated. Similarly, we compute the testing error by counting violations in the testing set. The goal of cross-validation is to find the parameter λ which gives the smallest testing error. We run 10-fold cross-validation. The data set was divided into 10 sets and in each trial, one of them was treated as a testing set. The errors were averaged across all trials. Figure 5.5 (left) presents the outcome of this validation as a function of λ. The best trade-off was obtained for $\lambda = 5$ (training error equal to 8.96% and testing error 9.78%). To analyze the dimensionality of this embedding we computed a spread, which describes the variance along each embedding dimension (Figure 5.5, right). As seen on the plot, almost all variance is contained in the first coordinate. This suggests that the space can be approximated well using only one dimension. The space for the 12 materials from our experiment is shown in Figure 5.6. We also plot their confidence intervals and scale in just-noticeable units, both described below.

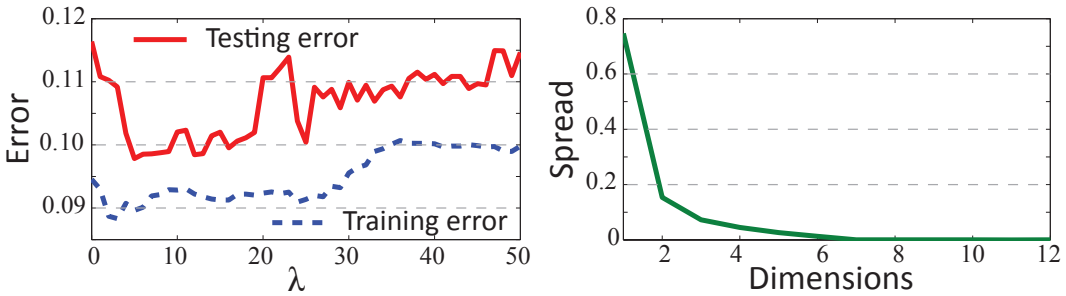

Figure 5.5. Left: Testing and training error for non-linear MDS as a function of λ parameter. Right: Variance across each embedding dimension (spread) for $\lambda = 5$.

Confidence of the Prediction The data used to compute our perceptual space is influenced by the inherent uncertainty of our psychophysical experiment. The question that arises is how confident we are in said perceptual space. To measure this, we applied a bootstraping technique [Davison and Hinkley, 1997] which allows us to compute confidence intervals for coordinates of our stimuli in the

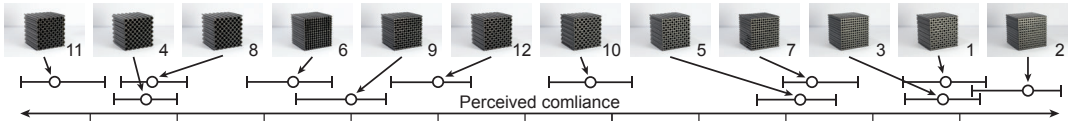

Figure 5.6. Perceptual space for our 12 metamaterials found using non-metric multi-dimensional scaling. The error bars visualize the confidence obtained using the bootstrapping test. The intervals marked on the axis correspond to 1 JND.

perceptual space. To this end, we computed our perceptual space 1000 times, each time performing resampling of our data with replacement, i.e., the number of measurements that were used in the NMDS was always the same but each time the measurements were chosen randomly with repetitions from our original data. Using this technique, for each stimulus we obtained an estimate of the distribution function of the sample position in the perceptual space. We then compute confidence intervals as the 5th and 95th percentile of this distribution.

JND Estimation The perceptual space allows us to compare differences between samples. However, the units in which they are expressed are undefined. For better understanding of the differences, we scale our prediction in JND (just-noticeable difference) units. A difference of one JND unit is detectable with a probability around 75%. Normally, to obtain the value of 1 JND a threshold estimation experiment needs to be conducted. Such experiments require a continuous change of the stimulus during the experiment. This is unsuitable in our scenario, where modifying stimuli requires reprinting them. Therefore, we decided to use the samples that we already had. We chose four pairs which lie close together in our perceptual space and performed a simple discrimination experiment in which 16 participants were asked to decide which sample in each pair was softer. The results of this experiment (Table 5.1) suggest that 1 JND unit in our initial perceptual space can be assumed to be around 0.15. Interestingly, as can be observed in Figure 5.6, 1 JND corresponds approximately to the size of the confidence intervals, which suggests that the confidence is on the level of human error.

5.3 Computational Model for Stiffness

The perceptual space derived in the previous section holds for the set of the 12 samples that were used in our experiment. It is unclear how new materials can

# Sample 1	# Sample 2	Answers	Diff. MDS
7	5	53 %	0.0117
9	12	76 %	0.1561
2	3	86 %	0.1738
8	11	95 %	0.2533

Table 5.1. Results of our JND estimation experiment. The first two columns provide block numbers (Figure 5.2) that were compared by the participants. The third column provides a percentage of people that answered that Sample 1 was softer than Sample 2. The fourth column provides the corresponding difference in our initial perceptual space.

be embedded in this space. To address this problem, we seek an analytical model that can be evaluated for an arbitrary material. To this end, we tested a number of computational models and compared them with our space. The models we chose are inspired by [Pressman et al., 2011], who proposed several models for discrimination tasks of delayed stiffness. Although this is different from compliance of nonlinear materials, the delayed feedback considered in their work also results in a nonlinear relationship between displacement and exerted force. Their models rely directly on data recorded using a haptic device, i.e., force and penetration applied by a user to a virtual material sample. This is crucial because compliance perception depends on the way people interact with objects [Friedman et al., 2008]. To account for the user interaction, we propose to characterize the interaction with materials using peak-force distribution functions. In the following part of this section, we first describe how such information can be acquired. Then, we describe and analyze models that can be used as an explanation of perception of compliant materials.

5.3.1 Force Distribution

We characterize human interaction with a material using probability distributions of peak-force magnitudes. More precisely, we seek a function $P(f) : \mathbb{R} \rightarrow \mathbb{R}$ which defines the probability that during probing the sample a user will perform a penetration with maximum force equal to f. To this end, we first collected force data as a function of time for each block (Figure 5.7, left). We did this by placing a force sensor under the reference sample in our experiment (Figure 5.4). To seek a unified model which predicts perceived compliance for an average observer, we combined force data from all subjects. To approximate function P_i for each

material block i, we computed a histogram of peak-force values (Figure 5.7, red circles) $H_i = (f_{ij}, n_{ij})$, where n_{ij} corresponds to the number of samples in the j-th bin of the histogram centered around force f_{ij}. In this work, we used 16 bins to compute the histograms ($j \in 1, 2, 3 \ldots 16$). After normalizing the histograms by the total number of samples, they can be used as an approximation of P_i: $P_i(f_{ij}) = n_{ij} / \sum_{j=1}^{16} n_{ij}$. An example of normalized histograms for two blocks is presented in Figure 5.7 (right). To see histograms for all blocks, please, refer to Appendix A.

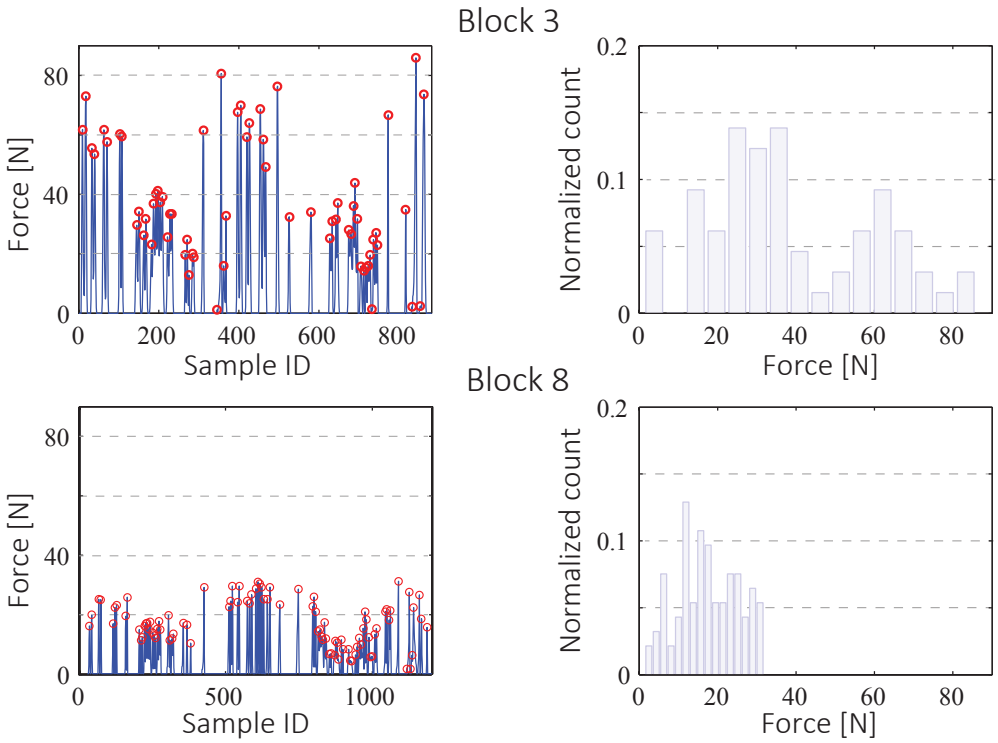

Figure 5.7. The plots visualize the data collected in our force experiment for blocks 3 and 8. Plots on the left visualize the raw data captured from using our setup: the x-axis corresponds to samples (time), and the y-axis corresponds to the force applied by subjects. The red circles show the peak forces which are used to compute peak-force histograms on the right. As the deformation properties of blocks 3 and 8 differ significantly, the forces applied by users are also different.

5.3.2 Models

Here, we describe several models that we consider candidates for approximating the perceptual space of compliance found using MDS. We rely on models presented by [Pressman et al., 2007] for delayed stiffness, but augment them with the force data. Additionally, we propose modifications which improve the performance of the models in the context of materials considered in this paper.

Global Stiffness (GS) This hypothesis states that people judge the compliance of a material by estimating the local slope of the force-displacement curve, which corresponds to local physical stiffness. In this case, perceived compliance k_i for material i can be estimated using linear regression on the force-displacement data c_i. The assumption is that the slope of the fitted line corresponds to the perceived compliance magnitude. Before performing the regression, the data is truncated to the maximum force applied in the force estimation experiment, i.e., the regression is performed on corresponding c_i but only for forces smaller than maximum force applied f_{max}, where $P_i(f) < 0$ for each $f < f_{max}$.

$$k_i = \arg\min_k \int_0^{c_i^{-1}(f_{max})} (kx + b - c_i(x))^2 dx \quad k, b \in \mathbb{R}.$$

Here, $c_i^{-1}(f)$ corresponds to the penetration distance for a given force f. Additionally, we introduce an alternative version of this model where the linear function fitted to the data passes through the origin GSF (*Global Stiffness Fixed*), i.e., $b = 0$. Good performance of this model could suggest that people, while estimating stiffness of a given material, account for the fact that zero-force should result in no displacement. Therefore, the estimated slopes should correspond to a linear function intersecting the origin.

Local Stiffness (LS) This model also estimates the compliance as a slope of the linear fit to the force-displacement curve, but to further include the force data, we use weighted least-square regression, where weights correspond to values of peak-force probability, i.e., the error of the fit is weighted by $\mathrm{P}_i(f)$ for each force f:

$$k_i = \arg\min_k \int_0^{c_i^{-1}(f_{max})} \mathrm{P}_i(c_i(x)) \cdot (kx + b - c_i(x))^2 dx \quad k, b \in \mathbb{R}.$$

This formulation requires a continuous function P. For the purpose of this work, it is approximated using a discrete summation over the peak-force histograms.

In the case of 16 histogram bins, the compliance can be computed as:

$$k_i = \arg\min_k \sum_{j=1}^{16} n_{ij} \cdot (kx + b - f_{ij})^2 dx \quad k, b \in \mathbb{R}.$$

Analogously to GSF model, we also define *Local Stiffness Fixed* (LSF) by enforcing $b = 0$.

Maximum Displacement (MD) This model, as proposed in [Pressman et al., 2007], estimates compliance using maximum displacement during one penetration. We propose to compute an expected value of maximum displacement. Hence, we define it as follows:

$$k_i = \int_0^\infty \mathrm{P}_i(f) \cdot c_i^{-1}(f)\, df.$$

Maximum Force (MF) Analogous to the previous model, this estimated the compliance as an expected value of force exerted on a given material:

$$k_i = \int_0^\infty \mathrm{P}_i(f) \cdot f\, df.$$

Peak Force/Penetration ratio (PFP) As originally proposed, this model estimates the compliance as a ratio between peak-force during a single interaction and the corresponding displacement. Analogously to the previous model, we define this model as an expected value of this quantity:

$$k_i = \int_0^\infty \mathrm{P}_i(f) \cdot \frac{f}{c_i^{-1}(f)}\, df.$$

Work (W) Inspired by the experiment of [Jones and Hunter, 1990], which showed that work is an important cue when discriminating stiffness, we added work as another hypothetical model for compliance. In this case, we define it as expected work that is done during a single penetration motion:

$$k_i = \int_0^\infty \mathrm{P}_i(f) \cdot \int_0^{c_i^{-1}(f)} c_i(x)\, dx\, df.$$

Additional Variations As suggested by [Leib et al., 2010] the compliance of nonlinear materials may relate to an inverse slope of the linear function fitted to the force-displacement curve. To investigate this hypothesis, we considered models INVGS and INVLS which are inverses of the previously defined models GS and LS.

It has been previously observed that for linear materials, the mechanical stiffness obeys the Weber law [Tan et al., 1992]. This suggests that the same may hold for the above models. In order to test this hypothesis, we added models which measure compliance as a logarithm of previously described quantities: LOGGS, LOGLS, etc.

5.3.3 Analysis

In order to validate which of the above computational models best explains compliance perception, we compare the prediction of each model with the perceptual space found in Section 5.2. To this end, for each computational model, we compute the coordinates of all 12 blocks and check how these coordinates correlate with those from the perceptual space. Our assumption is that a computational model is good only if its outcome has good linear correlation with the perceptual space. To measure the correlation we use linear regression. Figure 5.8 visualizes the results of this experiment for all computational models. The goodness of the fit for each model was measured using R-squared. The values are presented in Table 5.2. As we are not only interested in a good prediction of perceived compliance, but also perceived distances, we also check how computational models predict differences in pairwise distances between materials. To this end, we computed what percentage of experimental data (Section 5.2) is explained by each model. We performed the same test assuming that all triplets from the main experiment were evaluated using our perceptual model. For validating distances we also include the L2 norm on displacements evaluated on force-displacement curves proposed by [Bickel et al., 2010]. The results of these experiments are presented in Table 5.2.

The results indicate that there are five models that provide a good match with both the data obtained in the experiment and the perceptual spaced found using MDS. The fact that all these models are logarithmic supports the Weber-law hypothesis. Surprisingly, a simple L2 norm can also provide a good prediction of differences in pairwise distances (see last column of Table 5.3). To investigate this further we checked how well the distances in the perceptual space are explained by our models and the L2 norm. To this end, we considered all pairwise distances between our 12 stimuli and computed their magnitude using the per-

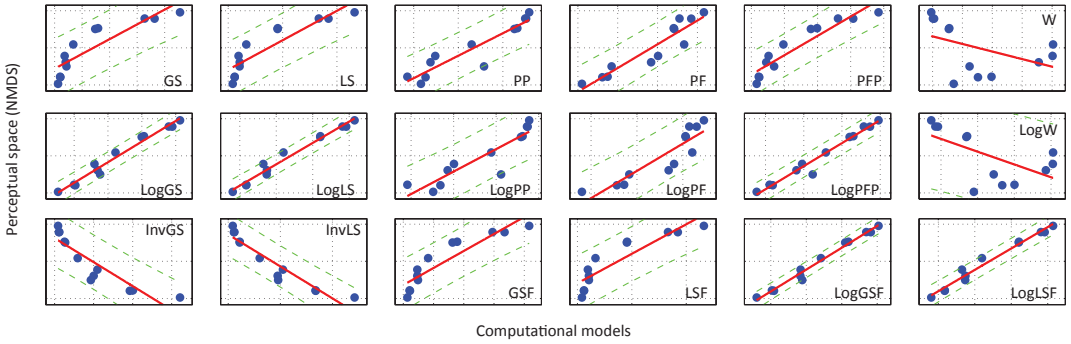

Figure 5.8. The plots show how different computational models correlate with the perceptual space obtained using non-metric multi-dimensional scaling. Each chart's X and Y axis correspond to computational model and perceptual space, respectively. Blue points correspond to our stimuli. The red lines are the results of the linear regression and the green dashed lines indicate prediction intervals with a confidence level of 0.95. A good correlation between computational models and experimental data was found for several models.

	GS	LS	PP	PF	PFP	W	LogGS	LogLS	LogPP	LogPF	LogPFP	LogW	InvGS	InvLS	GSF	LSF	LogGSF	LogLSF	L2
R^2	0.786	0.813	0.866	0.871	0.883	0.206	**0.973**	**0.975**	0.820	0.812	**0.969**	0.368	0.813	0.892	0.841	0.813	**0.984**	**0.975**	–
vs. data	83.1%	85.3%	83.8%	82.8%	87.7%	63.5%	**90.3%**	**91.2%**	83.5%	82.7%	**91.1%**	67.1%	86.1%	90.3%	86.2%	85.3%	**92.2%**	**92.3%**	91.1%
vs. MDS	83.8%	86.9%	84.4%	82.5%	88.0%	65.5%	**90.6%**	**92.9%**	84.1%	82.3%	**92.4%**	69.0%	85.7%	91.2%	86.5%	86.9%	**92.5%**	**92.9%**	92.1%

Table 5.2. Statistics from experiments evaluating each computational model. High values marked in bold indicate good candidates for computational models. Each of them does a good job of predicting not only the distances evaluated using the perceptual space found using MDS (above 90%), but also the data collected in our main experiment.

ceptual space, our models, and the L2 norm. We then computed a correlation between the observed distances in the perceptual space and the predicted ones. Table 5.3 shows the correlation for the five best performing models evaluated in this paper as well as the L2 norm. The correlation for the L2 norm turned out to be worse than for our perceptually-motivated models. In Figure 5.9, we also show plots of predicted versus observed distances for LogLSF and L2.

To summarize, our analysis suggests that five of the models tested in our experiments (LogGS, LogLS, LogPFP, LogGSF, LogLSF) perform well in predicting the perceptual space found in Section 5.2. Three of these models (LogGS, LogLS, LogPFP) were also found to perform well for perception of delayed stiffness of linear materials [Pressman et al., 2007]. However, in our case where we seek a perceptual space, we had to consider the logarithm of the corresponding quantities; otherwise the prediction was poorer. Furthermore, to account for dif-

LogGS	LogLS	LogPFP	LogGSF	LogLSF	L2
0.96	0.96	0.96	0.97	0.98	0.89

Table 5.3. Correlation between distances in the perceptual space and the prediction provided by the computational models and the L2 norm.

ferences in user interaction, we augmented these models with additional force information. The best performance was provided by LogLSF, which was not considered by Pressman et al. Surprisingly, the non-perceptual L2 norm performed well in predicting the relation between pairwise distances; however, it performed worse in predicting the magnitude of the differences.

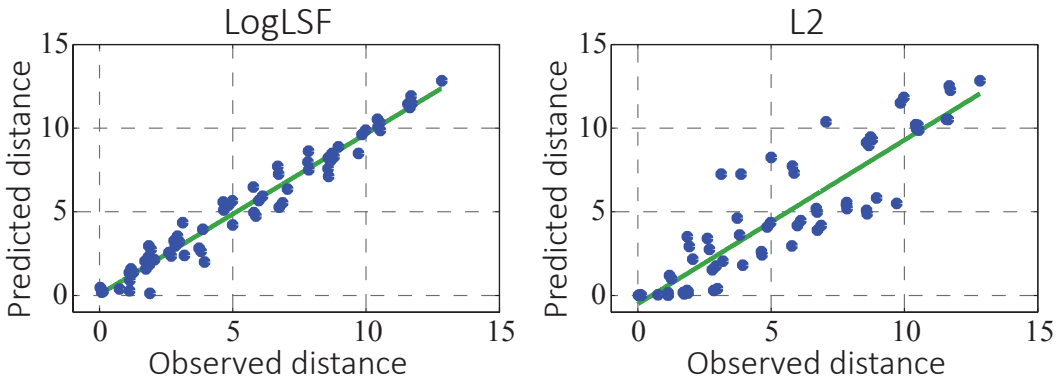

Figure 5.9. The two plots visualize the distance prediction between our stimuli as predicted by the LogLS model and the L2 norm with respect to the observed distances in the perceptual space. LogLS outperforms the L2 norm by providing a much closer match to the observed distances.

5.4 Complex Geometries

In the previous section, we considered different models for modeling compliance perception. Their performance was analyzed based on the data collected for simple shapes, i.e., blocks made of different metamaterials. In this section, we demonstrate how such models can be applied to complex geometries. In the rest of the chapter, we will use the LogLSF model as it performed best in our analysis in Section 5.3.

5.4.1 Computational Model for Complex Geometries

Compliance perception depends on the force-displacement relationship at the point of contact with an object. As such, it is influenced not just by the material an object is made of, but by its shape as well. Figure 5.11 outlines our method for extending our model to complex geometries using numerical simulation of elastic objects.

Compliance Computation To predict the perceived compliance at a given contact location $\mathbf{x}$ on the object's surface, we first have to determine the force-displacement curve. In the case of real objects, such information can be obtained using a uniaxial test with force applied at the point $\mathbf{x}$. For fabrication applications we estimate this information before the object is produced; to do this, we rely on numerical simulations. We simulate poking an object with a finger-like contact geometry in the direction of the contact surface normal, (Figure 5.11a). Recording both the force and displacement of the contact patch produces a holistic force-displacement curve for the entire object. More specifically, we employ a finite element method implemented in Abaqus[1] with an implicit integration scheme with the following settings.

Material Model For each material the force-displacement curves obtained by an uniaxial compression test were extracted and converted to nominal stress-strain curves. Next, using this data the following hyperelastic material models were analyzed: Mooney-Rivlin, Ogden 3th order model and Reduced polynomial models from 1st to 6th order [Dorfmann and Muhr, 1999]. After the analysis unstable material models were excluded and from the remaining models the model best explaining experimental data with minimal degree was used. For silicon materials we used Reduced-polynomial model of 4th degree. Objet materials were simulated using Reduced-polynomial model of 2nd degree.

Model Setup For all simulations in Abaqus quadratic tetrahedral elements with hybrid formulation were used to accommodate for nearly incompressible materials such as silicons. In our validation studies we consider three 3D models. The rubber duck model[2] (Figure 5.10 left) was simulated using 22803 elements. The stand of the model was fixed to simulate that it was glued to its base and ground was added to represent the contact of the model with its base. The octopus model

[1] https://www.3ds.com/products-services/simulia/products/abaqus/

[2] https://www.thingiverse.com/THING:139894

(Figure 5.10 middle) was simulated using 3282 elements. Only half of the model was modeled in the software. The other half was represented using symmetric boundary condition. The contact of the octopus with ground was modeled using a high friction coefficient to represent the interaction between rubber and a wooden plate. The seahorse model[3] (Figure 5.10 right) was discretized using 7429 elements. Only half of the model was used for simulation as the used materials (Ecoflex 00-30, Dragon skin 30, TangoBlack+) were soft enough to not affect the rest of the model. For the remaining two materials (flexible resin, TPU 92A-1) due to the lack of material samples we directly measured the interaction with a uniaxial device, instead of simulating it.

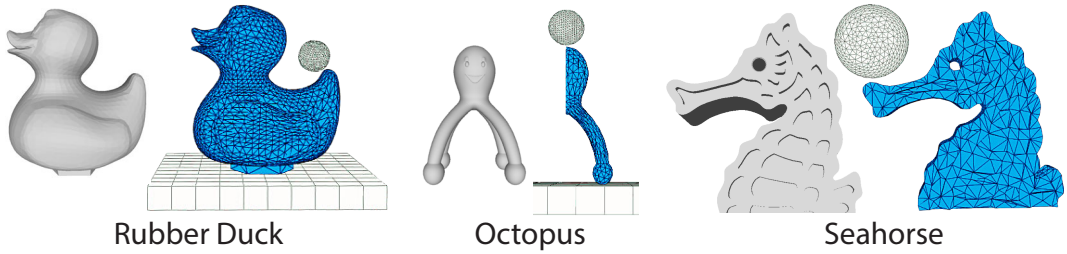

Figure 5.10. The three models used for the simulation and their corresponding tetrahedral mesh.

Interaction Simulation To represent the human interaction with an object a spherical indenter with diameter of 1.5 cm was used. In all simulations we modeled this indenter as a rigid body made from aluminum. During the simulation we extracted the reaction forces of the indenter to recover the force required to deform each specimen. The overall time duration of the simulation was set to 2 seconds to account for the human interaction speed.

Peak Force Histogram In addition to the force-displacement information, our model requires a force histogram that describes the interaction of a subject with the object. Here, we rely on the prediction based on the data collected in our experiments with blocks (Section 5.2). We assume that subjects interact in a similar way with objects that have similar force-displacement characteristics. Consequently, to predict the histogram for a new force-displacement curve, we first find two blocks that have the most similar force-displacement curves in our database (Figure 5.11b). The distance between curves is measured as the area between

[3] https://www.thingiverse.com/THING:561147

them, though any metric could be used. Then, we interpolate the two corresponding force histograms using the distances as weights. For the interpolation we use the method proposed by [Bonneel et al., 2011].

Compliance Metric Given the force-displacement curve and the force histogram, we can directly apply our model from Section 5.3 to obtain the prediction of perceived compliance at location $\mathbf{x}$ (Figure 5.11c). In order to compute a dense compliance map for the entire object, we sample its surface using Poisson disk sampling [Cook, 1986] and compute local compliance for each location separately (Figure 5.11d). Later, we interpolate the results using inverse distance weighting. The distance between vertices is measured with an approximation of geodesic distance using the Floyd-Warshall algorithm. A result of such computation is presented in Figure 5.11e.

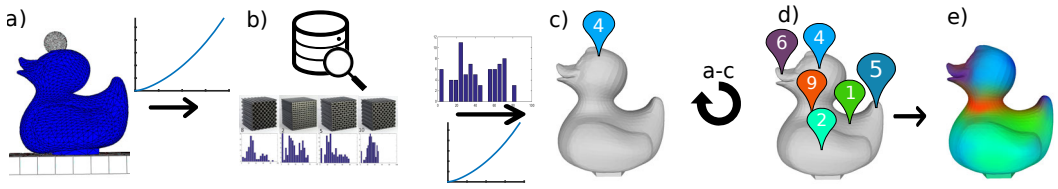

Figure 5.11. To compute compliance properties of an object we start with simulating the interaction performed by an observer for a given location on the surface (a). This allows us to extract force-displacement information. Next, the force histogram is obtained using our database of blocks with ground-truth force information (b). To this end, we interpolate histograms corresponding to the two closest curves in the database. Given the force-displacement curve together with the corresponding force histogram, compliance properties are computed according to our model (c). The procedure is repeated for a set of locations on the surface (d). To obtain a dense, spatially-varying compliance map, we interpolate the compliance values on the object's surface (e).

Dense compliance maps allow us to formulate a metric that predicts compliance differences between two objects with the same geometry but made of different materials. To this end, we first compute a dense compliance map for both objects which are then subtracted. As the compliance prediction provides information in perceptually uniform units, the direct differences are perceptually meaningful and their magnitude can be interpreted as perceived differences. An example of such a map is presented in Figure 5.12.

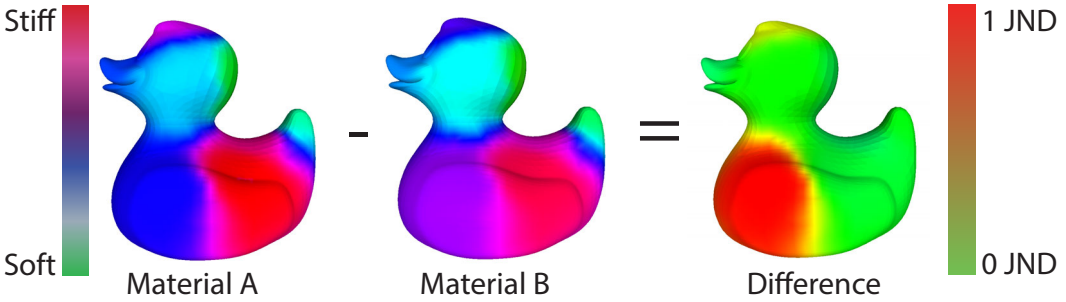

Figure 5.12. To compute difference between two versions of the same object made of two different materials, we first compute a compliance map for each of them. Then, we subtract the values corresponding to the same locations. Here, Material A and Material B correspond to silicones with Young's modulus values of 0.0964 MPa and 0.0973 MPa, respectively.

5.4.2 Evaluation

To evaluate our strategy for computing compliance properties of complex geometries, we conducted a user experiment, in which participants investigated objects with different mechanical properties and were asked to judge the differences between them. The outcome of the experiment was compared to predictions provided by our model and the L2 norm.

Stimuli For the purpose of this experiment, we used a duck-toy[4] shape for which we printed a mold and cast the shape with 9 different mixtures of silicones. For the base silicon material we have selected materials with different shore hardness [AST, 2010]: Dragon skin 30, Dragon skin 10, Ecoflex 00-50, and Ecoflex 00-30. The different mixtures as well as naming convention are shown in Table 5.4. This way we obtained shapes with different mechanical properties. The size of each duck was approximately 7×6.5×4.5 cm. Additionally, for each silicone we cast a small cube that was used to measure the nominal stress-strain curves for each material, which were later used for physically-based simulation of the shapes.

Experiment Ten participants took part in this experiment. Each of them was presented with all the stimuli at once and asked to order the objects according to compliance. Participants were allowed to palpate the object at a prescribed location in a direction perpendicular to the surface. To acquire information about perceived differences between the objects, we asked participants to space the

[4] https://www.thingiverse.com/THING:139894

ID	Silicon mixture
A	100% Dragon skin 30
B	50% Dragon skin 30 + 50% Ecoflex 00-30
C	100% Dragon skin 10
D	50% Dragon skin 10 + 50% Ecoflex 00-50
E	100% Ecoflex 00-50
F	87.5% 100% Ecoflex 00-50 + 12.5% Ecoflex 00-30
G	75% 100% Ecoflex 00-50 + 25% Ecoflex 00-30
H	50% 100% Ecoflex 00-50 + 50% Ecoflex 00-30
I	100% Ecoflex 00-30

Table 5.4. Silicon mixture with their corresponding IDs.

objects according to the differences between them. For this purpose they were asked to arrange the objects along a 2 meter scale. We chose four different locations which subjects were asked to evaluate: head, tip of the beak, body, and tail. Each participant performed the experiment in four separate sessions. This allowed us to reduce the bias introduced by previous examinations of different locations. To make the scenario realistic, we did not prevent visual feedback.

Results Each subject provided four orderings. For each of them the positions were recorded, scaled to $0-1$ range, and averaged across the participants. As a result, we obtained four orderings together with relative positions. We also predicted the perceived compliance using our model and compared them to the subjective evaluation. The results of this analysis are presented in Figure 5.13. To evaluate how good our prediction is, we computed Spearman's rank correlation and Pearson's correlation coefficients between the positions from the subjective experiment and the prediction. They are equal to 0.9617 and 0.9269, respectively. For computational fabrication techniques, such as [Bickel et al., 2010; Chen et al., 2013] more relevant information is the differences between objects. To evaluate our model in this context, we compare its prediction to the L2 norm which is often used in practice. To this end, we computed all pairwise distances between the objects according to our model, the L2 norm, and the positions provided by the participants. Table 5.5 presents coefficients for Pearson's correlations between the user data and our prediction as well as the user data and the L2 norm for each location separately.

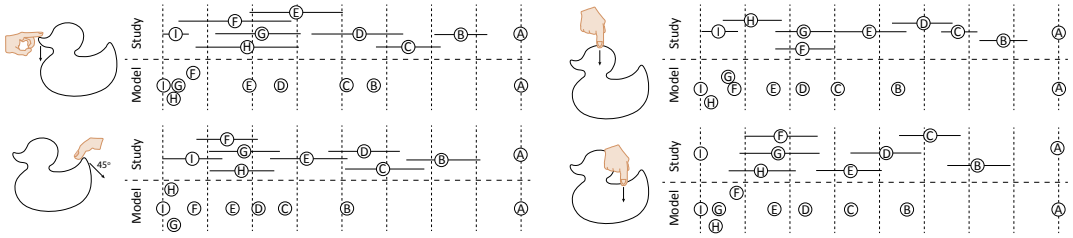

Figure 5.13. The results of the evaluation of our compliance model. The plots visualize the compliance as judged by the participants (top row), as well as our prediction (bottom row). Additionally, the schematic figures visualize the mode of interaction. Letters A-I correspond to different silicon mixtures ordered by Shore hardness. For full simulations of the used materials please refer to Appendix A.

	Beak	Head	Tail	Body
L2 vs. Exp. data	0.6541	0.6378	0.7657	0.6092
Ours vs. Exp. data	0.9090	0.8952	0.8959	0.8490

Table 5.5. Correlation between distances predicted by our method and the L2 norm.

Discussion The evaluation shows that our prediction correlates well with the positions and distances obtained from our user experiment. This suggests that the computational model that we propose in this paper can be used for judging the perceived compliance and distances between objects with different mechanical properties. Additionally, we show that our model outperforms a simple L2 norm operating directly on force-displacement data. When compared to the results obtained for simple shapes (Section 5.3, Table 5.3), the correlation coefficients for complex geometries (Table 5.5) are lower. We believe that this reduced performance can be attributed to: (a) accuracy of the interaction performed by participants (i.e., contact position and direction of interaction), and (b) the contact area of the finger and the object, which depends on the neighborhood of the point of interaction.

Using our data, we can also evaluate whether or not the perceived compliance depends on the geometry or solely on the material properties of an object. To test the hypothesis that the geometry influences perception, we computed two-way analysis of variance (ANOVA) on our data from the experiment. The two variables considered in the test were the location of the interaction and the silicone used for casting the shapes. The ANOVA revealed a significant effect for

both variables with $p < 0.05$. This shows that people are not able to factor out the geometry while judging the mechanical properties of the object, and that our approach that takes into account this factor is justified.

5.5 Application to 3D Printing

In this section we propose three possible usage scenarios for our model. First, we show that it can help in the process of choosing the right material for fabricating objects with prescribed compliance properties. The second use case demonstrates that our compliance prediction can speed up prototyping processes by providing suggestions to designers. Finally, we discuss how our model can affect current state-of-the-art fabrication algorithms.

5.5.1 Accuracy

Reproducing desired deformation properties is an important problem. Recent work tackles this problem by fabricating metamaterials made of micro-structures [Bickel et al., 2010; Schumacher et al., 2015; Panetta et al., 2015]. However, due to the limitations imposed by designers or 3D printers, complex internal structures are not always desirable or fabricable. In practice, the designer might be able to use only a limited number of materials to reproduce a target object. As shown in previous sections, our model has a higher correlation with user data than the L2 norm, which suggests that one could use it to improve the accuracy of optimization for material assignment.

Figure 5.14. From left to right: a rendering of the original seahorse model. Seahorse models used for the study: Dragon Skin 30, Ecoflex 00-30, TangoBlack+, Flexible resin, TPU 92A-1.

Experiment To examine this claim, we conducted the following experiment. We fabricated a seahorse model using five different materials: TangoBlack+

(Objet500 Connex), TPU 92A-1 (Laser printer), Flexible resin (Ember printer), Smooth-On Dragon Skin 30 silicone rubber (casting), and Smooth-On Ecoflex 00-30 silicone rubber (casting) shown in Figure 5.14. The experiment consisted of generating all triplets of materials. From each triplet the object with medium stiffness was removed and the task was to decide which of the two remaining materials should be used to replicate the removed object. We assumed that during interaction each seahorse is held by its body and participants examine the tip of its mouth. In total we had 10 trials. In 4 of these trials the L2 norm and our model did not predict the same material. We conducted a user study with 16 participants. All of them were presented with each triplet and asked to choose the best replica.

Results Detailed results of this study are presented in Table 5.6. The L2 norm correctly predicted 85 out of a total of 160 selections (53%), which is on the level of a random guess. In contrast, our model was able to predict 125 answers (78%). The prediction was always consistent with the majority vote. This suggests that during the design process our model can provide meaningful suggestions to artists regarding the material choice. The L2 norm failed to predict the correct answers in some cases. All of them involved situations when the harder material was the correct answer.

People's choice		Reference material	Less similar sample	
Material	People		Material	People
Dragon Skin 30	100%	TangoBlack+	Ecoflex 00-30	0%
Ecoflex 00-30	69%	TangoBlack+	TPU 92A-1	31%
Ecoflex 00-30	81%	Dragon Skin 30	TPU 92A-1	19%
TPU 92A-1	69%	Flexible resin	Dragon Skin 30	31%
Ecoflex 00-30	56%	Dragon Skin 30	Flexible resin	44%
TPU 92A-1	94%	Flexible resin	Ecoflex 00-30	6%
TangoBlack+	88%	Dragon Skin 30	Flexible resin	12%
Ecoflex 00-30	69%	TangoBlack+	Flexible resin	31%
TangoBlack+	94%	Dragon Skin 30	TPU 92A-1	6%
TPU 92A-1	63%	Flexible resin	TangoBlack+	37%

Table 5.6. Results of the perceptual study. From each trial the triplet is shown as well as the corresponding participants' picks. Our predictions always correspond to the participants' choices. Cases when the L2 norm did not agree with our prediction are highlighted in red.

5.5.2 Material Selection for Prototyping

Recent advances in 3D printing enable designers to rapidly fabricate and test prototype versions of their designs. However, printing a single prototype can still take hours and consume a considerable amount of material. Therefore, it is desirable to optimize for the cost and the printing time by minimizing the number of printed models. In the context of compliance, the material properties are not very meaningful in the perceptual sense, and the variation of material parameters does not always result in the expected changes in the object behavior due to the geometry of printed models. Our model can help in this process by providing suggestions about the material parameters to designers.

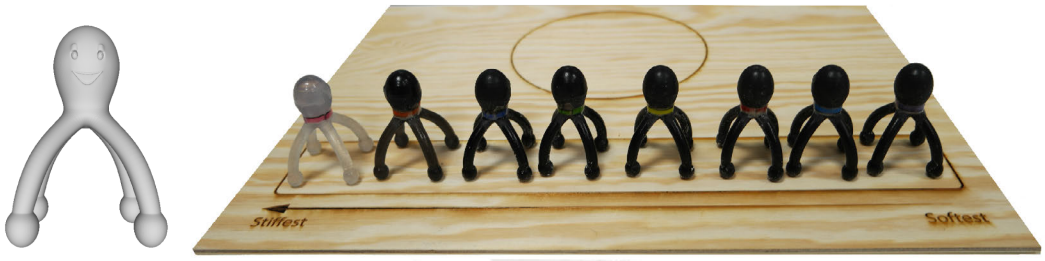

Figure 5.15. Left: a rendering of the original octopus model. Right: all printed octopuses ordered by stiffness.

Experiment We consider a situation of fabricating an object with unspecified compliance properties. A designer does not know the exact material parameters, and therefore, they wish to print a few versions of the object from which they can choose. To aid the process, it is desired to print objects that will cover the range of the possible outcomes in a uniform way. For the purpose of demonstration we considered an octopus shaped toy (Figure 5.15, left) that is to be printed on the Objet500 Connex printer. Using digital materials composed from TangoBlack+ and VeroClear one can fabricate 8 different models. To simulate the design process, we conducted an experiment in which 20 participants were asked to put themselves in the role of the designer. They had already tried models with the two base materials and now would like to print some in-between samples. Due to time and cost limitations they can, however, print only a limited subset of the samples. Their task is to pick these samples such that they maximize the difference in compliance between each other. In the first trial users were asked to pick 2 out of 6 possible samples and in the second trial they were allowed to pick 3. To check whether our model can help them in this task, we predicted the user

selection by taking sets of samples that maximize the minimal distance measured with our model. For the comparison, we also used the L2 norm.

Results The outcome is presented in Table 5.7. For the trial of picking two representative samples, the L2 norm predicted the pair 4,2 while our model predicted the pair 6,4. For the trial of picking three samples, our model predicted the triplet 8,6,4 while the L2 norm predicted the triplet 6,4,2. The prediction of our model was always the set with the highest occurrence count; meanwhile, the L2 norm matched people's choice only partially. This demonstrates that our model can be used for designing pallets of materials/objects that exhibit sufficient variation in perceived compliance. The main difference between our model and the L2 norm is that in both trials the L2 norm is biased towards stiffer samples. This is similar to the previous experiment and may suggest that the L2 norm overestimates human sensitivity to hard materials.

Picking 2 samples						
People's choice	6,4	7,4	7,6	8,6	...	4,2
Probability of choice	30%	20%	20%	15%	...	0%
Picking 3 samples						
People's choice	8,6,4	7,6,4	9,6,4	8,7,4	...	6,4,2
Probability of choice	35%	20%	10%	10%	...	0%

Our prediction — The L2 norm prediction

Table 5.7. Results of the perceptual study. For both trials of picking two and three samples we present the people's choice and its corresponding probability. Sets predicted by our model and the L2 norm are underlined in green and red, respectively.

5.5.3 Improving Fabrication Algorithms

Optimizing material/metamaterial assignment in order to obtain desired mechanical properties for a given object is a computationally expensive operation. We have already demonstrated that our model provides a more accurate prediction of differences between objects with respect to the L2 norm. This suggests that we can replace the L2 model in fabrication methods, such as Spec2Fab [Chen

et al., 2013], and obtain a closer match (in the perceptual sense) with the desired mechanical properties. However, our model can also shorten the time that such an optimization requires. Given desired properties of an object, we can run and terminate such optimizations when the error drops below one just-noticeable difference.

Experiment To test this hypothesis we ran two Spec2Fab [Chen et al., 2013] optimizations that minimized the L2 norm. One of them was performed until convergence, while the other was stopped when the error measured using our model reached one JND. The resulting curves can be seen in Figure 5.16. Sixty-four simulation steps were required for the optimization to converge, while only nine simulations were needed for the solution obtained using the perceptual cutoff. Interestingly, when we compute the perceptual difference between target and optimized curves, the converged curve has an error above one JND. Meanwhile, the curve for the terminated optimization has an error under half JND. This indicates that the curve obtained with the perceptual cutoff is a much better reproduction.

Discussion This behavior can be explained by two facts. First, our model takes into account user interaction. Second, it investigates the slope of the force-displacement relation instead of the direct difference between curves. In Figure 5.16 (left), we can see the plot of force-displacement data for the target object as well as the optimized ones. On the right, we present plots of the corresponding slope values (derivatives). One can clearly see that in the region of main interaction between 18 and 35 N (high values of force histogram), the curve obtained with the perceptual cutoff provides almost a perfect match in terms of the slope. This gives us an important insight about how our model could be used to design new fabrication techniques that optimize the perceived differences directly.

5.6 Limitations and Future Work

In our experiments, we demonstrated that the compliance model and metric presented in this paper can be successfully used in many fabrication-oriented applications. However, there are some limitations to our technique.

Our model was constructed using only 12 samples. However, our experiments demonstrated that they cover the compliance range well (Figure 5.6) and the model achieves good correlation with the measured data. Furthermore, our experiments showed that our model extends beyond the original stimuli database

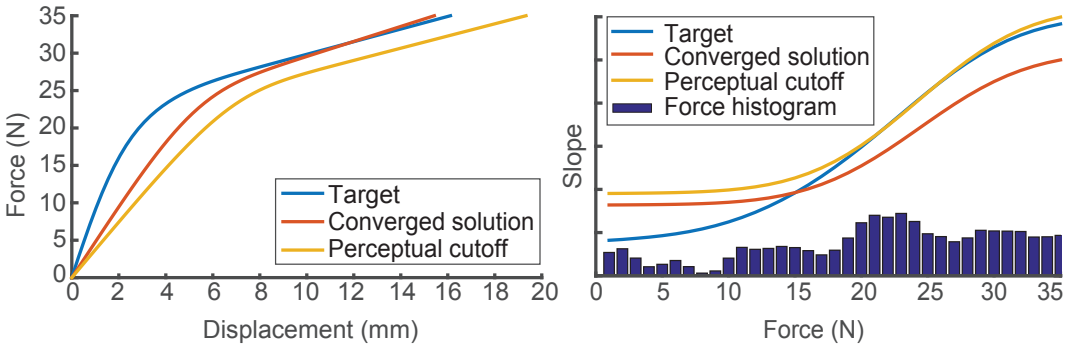

Figure 5.16. Left: force-displacement curves obtained from Spec2Fab optimization. Right: Slopes of the force-displacement curves along with corresponding histogram of forces.

for both softer (e.g., Ecoflex 00-30) and harder (e.g., VeroClear) materials.. The limited number of materials was also a design choice. More samples would make our user study infeasible. We believe that it is straightforward and beneficial to extend the force histograms database to improve the performance of our technique for complex geometries.

More limiting than the above, our model only supports a single mode of interaction (poking). We limit ourselves to elastic material models which assume no velocity dependence on resulting forces. Although for many of the materials used in our applications (i.e., silicones, Flexible resin, TPU 92A-1) this is acceptable, for some viscoelastic 3D printable materials, e.g., TangoBlack+, this is not an ideal assumption. Visual feedback was avoided for the purpose of modeling compliance perception in Sections 5.2 and 5.3. For practical reasons we did not avoid it during evaluation of the applications for complex geometries. We restricted our study to nonlinearly elastic materials that were produced by 3D printing metamaterials. Although they cannot be considered as natural materials, the benefit of using them for such experiments is the fact that people are unfamiliar with them. Therefore, additional cues such as textures are eliminated. An obvious direction of future work would be to relax one or more of these assumptions. Exploring free-form interactions or more complex material models could be particularly interesting, as they may result in a higher-dimensional perceptual space. Incorporating modeling of visual and texture feedback may further improve the performance of our model in fabrication applications.

The other set of limitations comes from how the perceptual model was applied to complex geometries. In particular, for each point on the surface, we examine only a single poking direction perpendicular to the object's surface. While

this is a natural direction for poking, in practice it is nearly impossible to poke objects exactly along the normal of the surface. This can negatively affect the compliance prediction of very complex objects where even a slight deviation from the normal direction results in perceivably different compliance properties. To apply our model in such a case, one could simulate the interaction for multiple poking directions and integrate the individual predictions to get the resulting compliance. How to integrate those predictions is an interesting problem for future work. Such integration of compliance would require prior knowledge of the interaction and the ability to predict which poking directions are more probable than others.

5.7 Conclusions

We presented a psychophysical experiment for evaluating compliance of different materials. Utilizing the collected data, and the numerical modeling of NMDS experiments, we found a perceptual embedding of 12 blocks of metamaterials. Our results suggest that despite the high dimensionality of possible nonlinear materials, the perceptual space can be estimated well using a one-dimensional embedding. This may suggest that people try to relate nonlinear materials to corresponding linear materials. Since the numerical model does not provide us with an explanation of the recovered space we evaluate different computational models for predicting perceived compliance. Based on the above experiments and modeling, we proposed an interaction-aware compliance model which is capable of estimating perceived differences between compliant 3D objects. We demonstrated several applications what showcase the capabilities of our model at creating more intuitive user interfaces, and improving numerical optimization to achieve more faithful reproductions. Both the model and the applications are evaluated in a series of user experiments in which we show that our perceptual model outperforms state-of-the-art approaches. The proposed model can be easily integrated in many existing computational fabrication algorithms, such as [Bickel et al., 2010; Chen et al., 2013; Schumacher et al., 2015; Panetta et al., 2015].

Chapter 6

Perception-Aware Modeling and Fabrication of Digital Drawing Tools

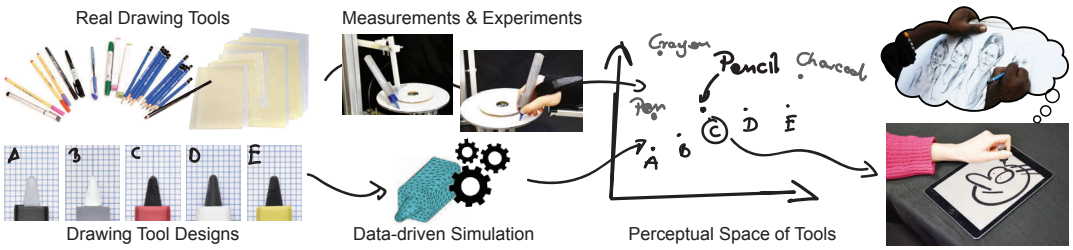

Figure 6.1. We propose a system for fabricating digital drawing tools that mimic the feel of real tools. To this end, we measure properties of different real drawing tools, study their perception, and design a perception-aware space of drawing tools. We later develop a simulation technique which allows us to embed new designs into the space, evaluate the pairwise similarity between them and the tools we want to replicate. This drives the design process of different digital tools.

In the previous chapter, we demonstrate the capabilities of perception-aware fabrication at reproducing the compliance of non-linear elastic objects. Unfortunately, we also observed limitations of the methodology. A significant experimental effort was required to find the physical phenomena governing even a relatively small one-dimensional space of compliance. Moreover, to project new designs into the perceptual space we relied on an off-the-shelf numerical simulation that has a significant computational cost. In this chapter, we lift both of these constraints by proposing novel algorithms for NMDS analysis and for numerical

modeling of haptic feedback. Equipped with the new tools we demonstrate an application of our methodology to the challenging problem of designing digital drawing tools.

Drawing and writing are among the oldest methods of recording and communicating information. While the culture and technology have evolved, traditional drawing and writing instruments, such as charcoal, crayons, pencil, ballpoint or fountain pens, remain unchanged and are still widely used. They are made of different materials and often used with a different drawing substrate producing a drastically different style and serving a different application. Each of them also has specific haptic feedback easily recognizable by professional and casual users. The relation between the haptic feedback and the tool is not only appreciated by artists but also used to their advantage [Danna and Velay, 2015; Annett et al., 2014]. The specific haptic feedback determines the degree to which the artist can control the tool. While in some cases, less precise control is desired to achieve less regular strokes; in other cases, excellent control is critical. The learned by artists relation between the aggregated haptic feedback of the traditional drawing tools and the stroke shape also allows them to refrain from relying on precise hand-eye coordination while drawing.

Recently, the digital drawing tools (styli and tablets) became an attractive alternative for many users as they offer comfort, portability, and integration with the abundance of new software solutions that significantly aid the creative process. The most recent software tools[1] enable taking advantage of this digitalization while trying to maintain the appearance of strokes produced by the traditional instruments. Unfortunately, as we take advantage of the digital drawing tools, we also lose the distinct haptic feedback associated with traditional instruments together with all its advantages. As a consequence, many artists return to traditional tools when possible.

A number of recent commercial and research efforts have focused on constructing new stylus devices and tablet surfaces to provide better feels-like haptic response behavior for the stylus-to-tablet interaction. Recent commercial products, including the Apple Pencil[2], the Microsoft Surface Pen[3], the Wacom stylus[4], the PaperLike cover[5], and the reMarkable system[6] focus on leveraging *passive* feedback with their main goal the faithful reproduction of haptic sensations from

[1] https://www.adobe.com/products/fresco.html
[2] https://www.apple.com/apple-pencil/
[3] https://www.microsoft.com/en-us/surface/accessories/surface-pen
[4] http://www.wacom.com/
[5] https://paperlike.com/
[6] https://remarkable.com/

drawing and writing with real materials [Williams, 2015]. These styli and surfaces have, in almost all cases, been painstakingly designed by trial and error: swapping in and hand testing differing stylus and surface pairs and iterating over shapes and materials[7]. Even so, as we will show in Section 6.5, current tools only cover a small subset of the range spanned by even standard drawing materials. Embedded actuation devices have also been explored [Romano and Kuchenbecker, 2012; Cho et al., 2016; Wang et al., 2016]. However, current latencies in available actuation puts these methods well out of reach for realistic haptic-rate feedback response to stroke gestures [Annett et al., 2014].

Our goal is to enable the automated design of styli that will passively deliver the haptic cues of all preferred drawing and writing implements; or even to blend between multiple tools to create styli that respond to mark making with novel feedback. To investigate perceptually-relevant properties of drawing tools we start with physical measurements of a representative set of traditional drawing tools (Section 6.2). Next, we embed these tools into a perceptual space (Section 6.3). To this end, we conduct an extensive psychophysical experiment. The results of this study are analyzed using our novel method that jointly optimizes for the perceptual space and its correlation to physical parameters of the drawing tools, Section 6.3.1. Critically this correlation with the physical properties enables a direct application to stylus design and fabrication. To design and evaluate new drawing tools without the need for fabrication and measurement, we propose a new method for simulating the interaction between a probe and a substrate (Section 6.4). Since many of the physical phenomena governing the response behavior are complex and expensive to model [Romano and Kuchenbecker, 2012; Otaduy et al., 2004], we propose a data-driven forcing term governed by our previous measurements (Section 6.4). Thus the primary contributions of this work include:

- measurement of the interactions between drawing tools and different substrates commonly used for traditional writing and drawing,
- perceptual experiments evaluating similarities between differing drawing tools,
- a perceptual space optimization that builds a space of tools whose dimensions are correlated with the physical properties of the drawing materials,
- a new data-driven method for simulating drawing tools, and

[7] `https://paperlike.com/`

- application of the above methodology to the design and evaluation of the haptic sensation of digital drawing tools.

6.1 Overview

We formulate the problem of designing digital drawing tools as a design optimization process which seeks a synthetic drawing tool that generates similar haptic sensation to the one generated by a given real tool. The process requires two major components. First, we need to define essential properties of drawing tools which enable performing a meaningful comparison between them. Second, to allow for optimizing design parameters, we need a technique that simulates the surface-tool interaction and obtains the properties that are necessary for the comparison.

To address the first point, we follow the observation from the literature (Chapter 2) and assume that the haptic sensation is induced mainly by the resistance and vibration resulting from the tool-surface interaction. Using our custom measuring device, we obtain the characteristic of a representative sample of drawing tools (Section 6.2) and find a set of physical properties that can describe them. To acquire information on how these properties influence perception and how they contribute to the perceived differences between them, we perform an extensive psychophysical experiment with a set of commonly used drawing tools and paper samples (Section 6.3.3). The results of the study allow us to embed the tools in a so-called perceptual space where the perceived distances between different tool are defined by the Euclidean distance between them. From the design point of view, it is essential that the tools can be easily embedded in the space based on their physical properties. Therefore, the derivation of the spaces is defined as a likelihood optimization which tries to match the experimental data and correlate the axes of the space with physical properties of the tools.

To embed new tools into the perceptual space and validate the haptic sensation they produce, we propose a new physical simulation based on perception-aware numerical coarsening which characterizes the vibrations caused by the tool-surface interaction (Section 6.4). Since accurate simulation of all the contact effects at the tip of the tool is challenging, we propose to model some of the effects using a data-driven approach. More precisely, we model the propagation of the vibration initiated at the tip of the tool using our exponential Euler integrator, while the forces acting at the tip of the tool are generated using our data-driven approach. The data-driven forces encode the complex contact characteristic between the tool and the surface.

Both the derivation of the perceptual space and the custom simulation enable applications to the digital drawing tool design which we demonstrate and validate in Section 6.5.

6.2 Physical Measurements

To study the feedback transferred from a drawing tool to a users hand, we built a custom measurement device and used it to characterize several traditional drawing tools.

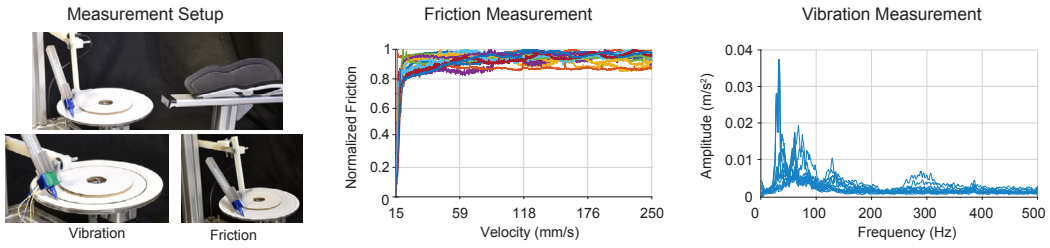

Figure 6.2. Measurements and perceptual user study setup. Turntable on which a mechanical arm is used to capture friction and vibration of drawing tools. There is a holder for a human arm (during user studies) to provide consistent grip and orientation. In the middle, we can see normalized friction measurements recovered using our setup. On the right is the accelerometer response for a fixed velocity.

6.2.1 Measuring setup

In our work, we follow observations from the literature (Chapter 2.2.2) that haptic sensation is induced primarily via resistance and vibration generated by the tool-surface interaction. Consequently, our device (Figure 6.2) measures resistance and vibration transferred to fingertips while drawing. The base of the device is a turntable operated by a DC motor with controllable speed to which different types of drawing substrate can be attached. To simulate an artist drawing on a surface, we design an arm with an enclosure for various drawing tools. By adding extra weight, the device can simulate different pressures applied by a user. The vibration of a tool is measured by an accelerometer attached to the enclosure, while the resistance is captured by a force sensor placed on the arm.

6.2.2 Defining measurement parameters

The vibration and the resistance produced by a drawing tool depend on its speed and the pressure applied. To reflect standard drawing scenarios, we first performed a pilot study to determine these parameters. We invited four amateur artists to a short drawing session and asked them to draw basic shapes [Garcia, 2003] (Figure 6.3) using an iSKN[8] tablet. The participants used five different drawing materials: ballpoint pen, soft (8B) and hard (2H) pencil, fine-liner, and charcoal, on three different drawing surfaces: standard 80-gram office paper, rough paper for pencil drawing, and smooth stone paper (Figure 6.4). During each session, the velocity of the tool was recorded using the positioning system of the tablet, while the pressure applied by the participants was recorded by a force sensor mounted below the tablet. To eliminate the effect of hand pressure, we placed a hand rest next to the tablet. Figure 6.5 shows the histograms of recorded velocities and forces. Based on the results, we decided to perform further measurements of drawing tools using the mean force applied by the artists, i.e., 2.2 N, and velocities between the 5th and 95th percentile, i.e., 17-250 mm/s.

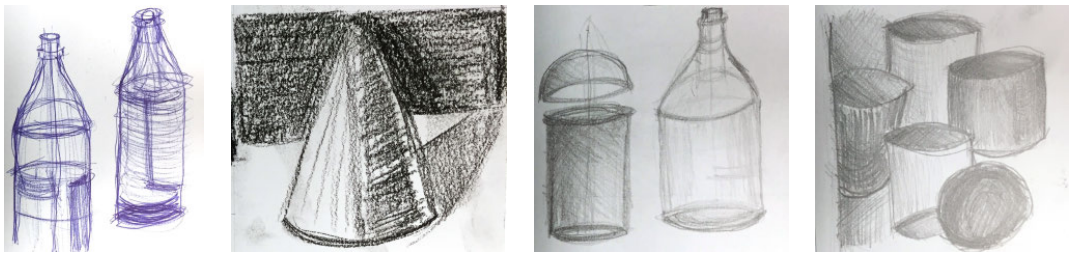

Figure 6.3. Sample images drawn during our preliminary user study.

6.2.3 Measurements

Vibration measurements were executed on our setup in the vibration measurement configuration (Figure 6.6). In this configuration, we use a shorter arm to reduce the vibrations coming from our setup. The arm has attached weight which exerts 2.2 N of force on the turntable. Additionally an ADXL335 accelerometer is attached onto the drawing tool that provides us with real-time vibration capture at 1000 Hz. The vibration test aims to capture the velocity dependent spectrogram of response. To this end, we turn the turntable from 0 to 350 mm/s in 1% increments. Each increment is maintained for 1.5 second to gather enough data

[8]https://www.iskn.co/

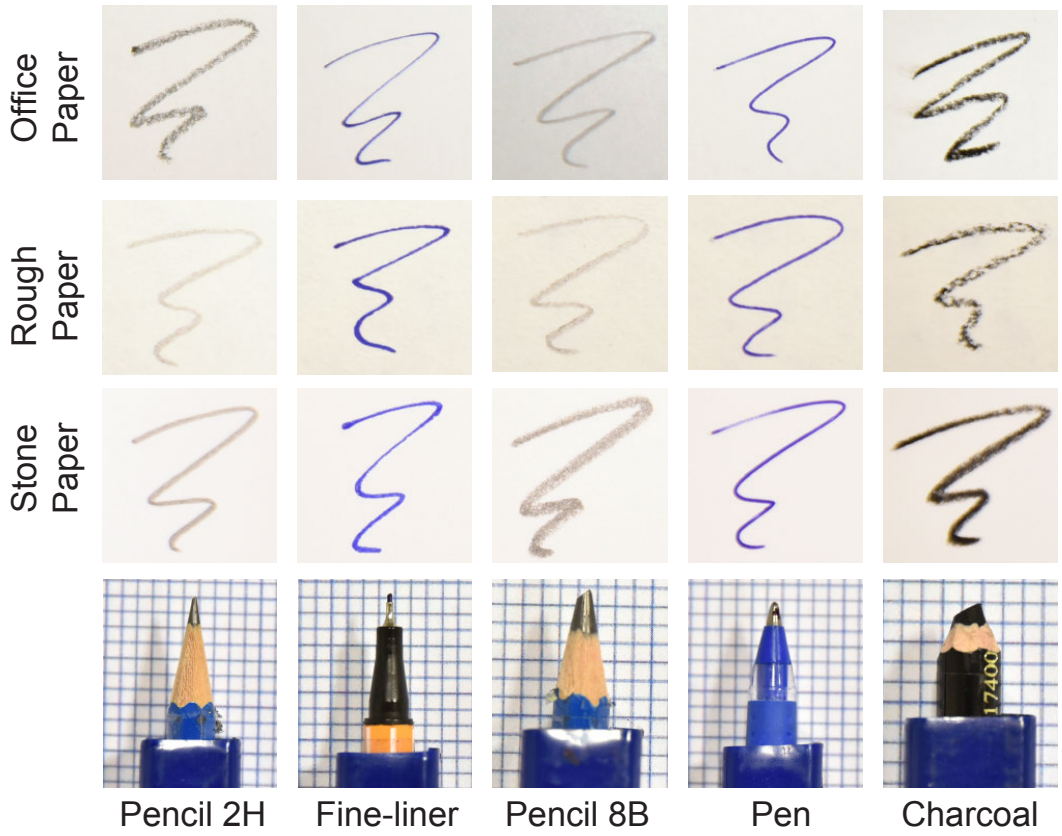

Figure 6.4. Our user study considers 5 drawing tools: 2H and 8B pencils, fine-liner, ballpoint pen, and charcoal and 3 drawing substrates: 80-gram office paper, rough artist paper, and stone paper.

for Fourier analysis. One experiment trial takes 2.5 minutes. Since our measurements show bands with no frequency shifts to accelerate the data collection for 3D printed tools, we modify the experiment and use 10% increments, which shortens one trial to 15 seconds. Finally, to reduce the impact of measurement noise, we repeat each measurement three times. To get the final results, we average the amplitudes of each signal in frequency domain.

For friction measurements, we reconfigure our setup. Instead of an accelerometer we attach a force sensor OptoForce OMD-10-SE-10N (Figure 6.7). To increase the sensitivity of the sensor, we attach the drawing tool on a 3D printed lever which provides us ten-fold amplification of frictional force. Similarly to the previous experiment, we attach an additional weight to the so it applies 2.2 N of force on the turntable. The turntable is set to move from 17 to 250 mm/s in smooth increments. The frictional force is captured at 1000 Hz during this

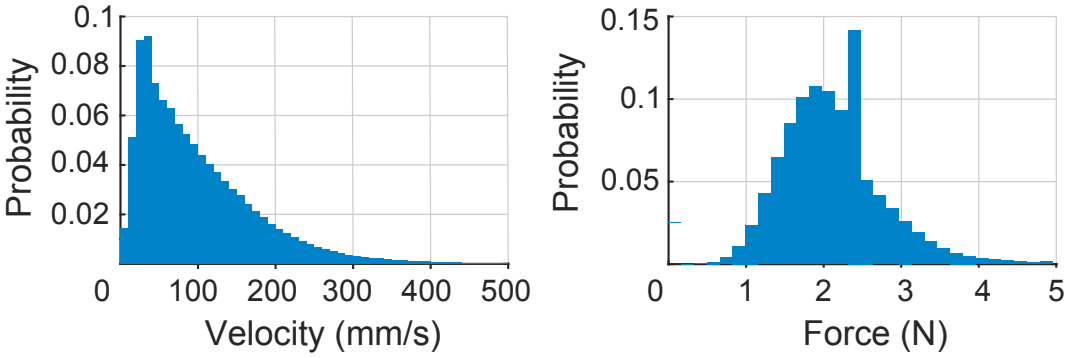

Figure 6.5. Velocity and pressure histograms recorded from free drawings during our pilot study.

motion.

Using our device, we measured each drawing tool and surface combination from our preliminary study (Figure 6.4). The normalized frictional force measurements (Figure 6.2, middle), revealed a small interaction between velocity and the friction coefficient. Consequently, we decided to use Coulomb friction to model measurements, and denote a single frictional coefficient for drawing tool and particular surface pairs. The recovered frictional coefficients range from 0.13 to 0.33. The vibration measurements (Figure 6.2, right) revealed a broad-band characteristic with significant dependence on velocity. We captured the complex vibration characteristic for each combination of tool and surface using a velocity dependent spectrogram of vibratory response. For more detailed measured data please refer to Appendix B.

Since the goal of the experiment is to measure vibration caused by tool-surface interaction, it is critical to assure that the measurement device itself does not produce significant vibration. To measure the vibration produced by our device, we experimented with a pen on an oiled acrylic sheet. Figure 6.8 demonstrates vibration spectrograms for the oiled surface and standard drawing paper. It can be observed that, although very weak, there is a vibration produced by the system on the oiled surface (due to the DC motor) which manifests as a diagonal line on the spectrogram. The weak horizontally structured signals, since not velocity dependent, are most likely due to the remaining interaction of the pen with the oiled surface. This is possible since our base surface is not perfectly smooth, and some asperities exist which actuate a drawing tool. We verified that the vibration caused by the system does not affect the perception of the drawing tools in an informal test during which the acrylic surface was perceived as smooth and vibration-free. For more comparisons of measurements on

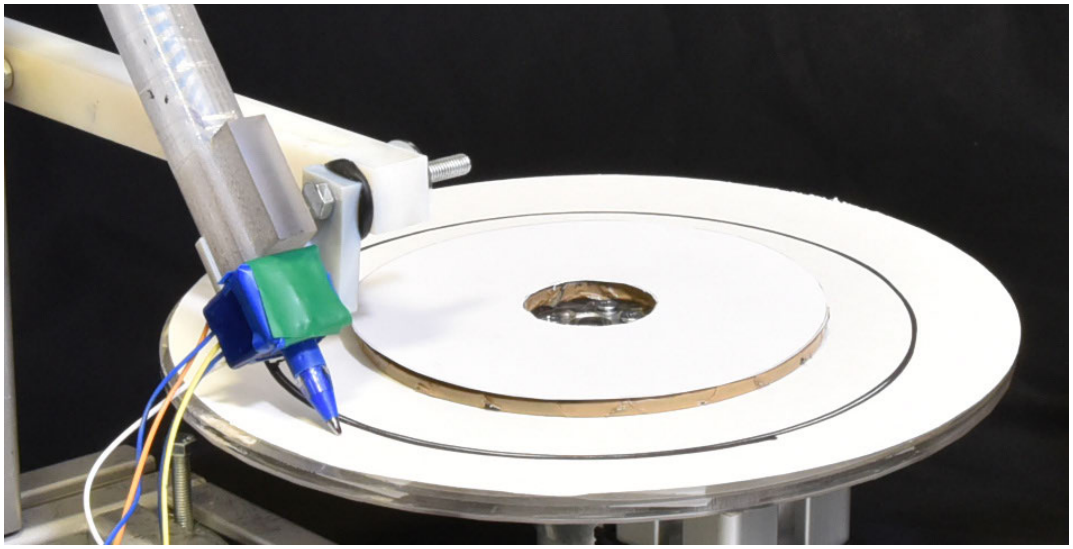

Figure 6.6. Acceleration measurement setup. The arm holds a drawing instrument, and the drawing substrate is attached via a magnetic plate. The vibrational response of the drawing tool is captured using an accelerometer.

paper and oiled acrylic please see Appendix B.

6.3 Perceptual Space Optimization

In this section, we describe the method and experimental setup we use to derive a perceptual space of drawing tools. We begin with non-metric MDS [Wills et al., 2009] and extend the formulation to automatically correlate the dimensions of the space with physical properties.

6.3.1 Recovering The Perceptual Space

We adapt a Bayesian method developed by Silverstein and Farrell [2001] which translates the result of pairwise comparison into scalar data. The advantage of this approach is that it is robust to noise, missing data, and cases when some of the comparisons are not performed the same number of times.

The experimental data required for computing the space consists of triplets of stimuli i, j, and k, with human judgments about which of i and j is more similar to k. To find the perceptual space, we want to find positions of our stimuli in a M-dimensional space such that the probability of our experiment occurring is

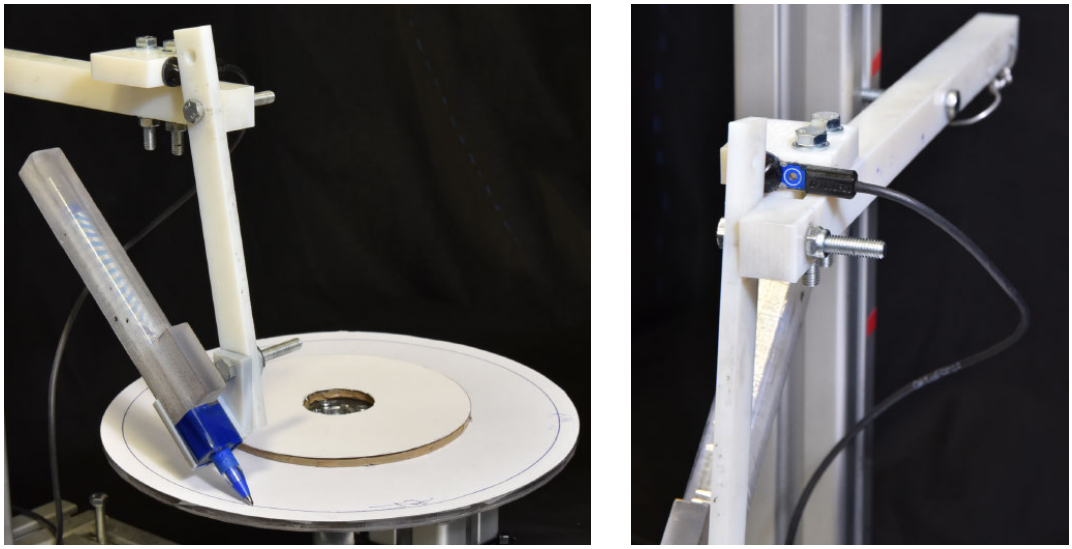

Figure 6.7. Friction measurement setup. The arm is enhanced with a lever that amplifies the frictional force excreted on the force sensor.

maximized. Under assumptions of Thurstone Case V law of comparative judgment [Thurstone, 1927], the probability can be computed as:

$$P_{exp} = \prod_{ijk} \binom{C_{ijk} + C_{jik}}{C_{ijk}} P_{ijk}^{C_{ijk}} (1 - P_{ijk})^{C_{jik}}, \tag{6.1}$$

where P_{ijk} is the expected percentage of subjects reporting that the sample i is more similar to a reference k than sample j, and C_{ijk} is the actual number of subjects that preferred sample i in our experiment. The binomial represents the total number of ways to pick C_{ijk} stimuli from a population of $(C_{ijk} + C_{jik})$, while the rest of the equation represents the probability of observing the particular sequence of decisions in the experiment. To express the probability P_{ijk} as the distances in the space we follow [Silverstein and Farrell, 2001] and use a cumulative normal function. To speed up the evaluation of the function, we approximate it as [Vazquez-Leal et al., 2012]:

$$\begin{gathered} P_{ijk} = 1 - \exp\left(-10.38\, d_{ijk} + 111 \arctan\left(0.09\, d_{ijk}\right) + 1\right)^{-1} \\ d_{ijk} = ||\mathbf{x}_i - \mathbf{x}_k|| - ||\mathbf{x}_j - \mathbf{x}_k||, \end{gathered} \tag{6.2}$$

where $\mathbf{x}_i$, $\mathbf{x}_j$, and $\mathbf{x}_k$ are locations of stimuli i, j, and k in the perceptual space.

Typically, perceptual spaces are computed for only a fixed set of stimuli, and positions of new stimuli are unknown. However, from the application point of

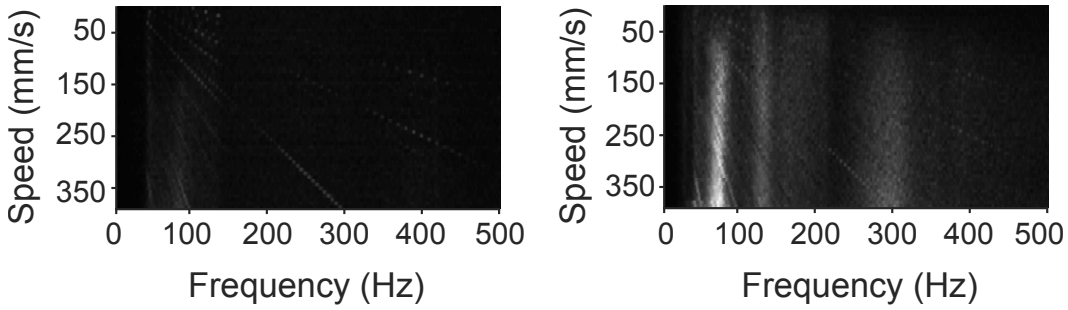

Figure 6.8. Measurements of a ballpoint pen on oiled acrylic (left), and 80-gram office paper (right).

view, it is critical that new stimuli can be easily embedded into the space without performing additional user experiments. To solve the problem it is desirable that the axes of the space are correlated with the physical properties of the stimuli. Finding a suitable definition of dimensions is usually a challenging task and often done as an additional step after deriving the space. Instead, we propose to jointly optimize for the embedding and correlation with the physical properties of the stimuli. The additional correlation requirement can be expressed as maximizing:

$$P_{corr} = \prod_{d} (Corr[X^d, D^d]), \tag{6.3}$$

where X^d is a vector of d-th coordinates of points X in the perceptual space and D^d are physical properties we wish to correlate with Q^d . To compute a perceptual space, we combine Equations 6.2 and 6.3 to formulate an optimization problem:

$$\arg\max_{X} P_{exp} \cdot P_{corr}^{\lambda}, \tag{6.4}$$

where λ is driving the trade-off between good agreement with experimental data and good correlation with measurable properties of the stimuli. In practice, to avoid very small values of the above function we minimize its negative logarithm,

$$\arg\min_{X} -\log(P_{exp}) - \lambda \log(P_{corr}). \tag{6.5}$$

To resolve translational ambiguity, we constrain a single element to the origin of our space. In contrast to standard, non-metric MDS methods, our formulation does not have a scale and rotational ambiguity. Scale is constrained by Equation 6.2 relating probabilities P_{ijk} with distances in the space, while rotation is handled by our correlation term (Equation 6.3). Our full optimization problem

is then nonconvex [Silverstein and Farrell, 2001]. To find a perceptual space minimizing Equation 6.5, we apply a quasi-Newton solver [Avriel, 2003] with multiple random starting points.

6.3.2 Experiment Design

To solve the above optimization, we must acquire an approximation of probabilities of sample preferences from a user study. Obtaining good estimates for all probabilities requires an unfeasible amount of experimental data in terms of number of experiments to perform. We therefore need to carefully decide which stimuli triplets should be evaluated in order to obtain a good estimation of the space. To this end, we make use of two observations: (1) evaluating triplets for which the outcome is expected brings little information; and (2) due to the sigmoid characteristic of the function relating distance in the space and probability (Equation 6.2), less obvious comparisons bring more reliable information to the optimization. In a similar fashion to the "exploration-exploitation" strategy from machine learning, this leads us to a two-stage procedure. In the first exploration stage, subjects evaluate all triplets but only with a few comparisons. Next, we identify the non-obvious triplets and conduct a study with a larger number of comparison. The final space is computed using P_{ijk} estimated from the data obtained in both stages.

6.3.3 Experiment

Our experimental apparatus was an extended version of our measurement device presented in Section 6.2. Besides the turntable on which we can place different samples of drawing surfaces, we added a support for the participants hands. To prevent visual feedback, which could affect the similarity judgments, we installed a black curtain that separated the participant from the turntable. Our experimental setup is demonstrated in Figure 6.9.

Stimuli The stimuli consisted of 15 combinations of five drawing tools and three different kinds of paper that were previously measured (Section 6.2), which resulted in 1365 triplets. Half were presented to the participants twice during the first stage of our experiments. From this set of triplets, we selected 60 stimuli which resulted in a tie and used them in the second stage where each of them was evaluated ten times. For the optimization of the perceptual space we used all the triplets from the second stage and those that did not result in a tie in the first stage.

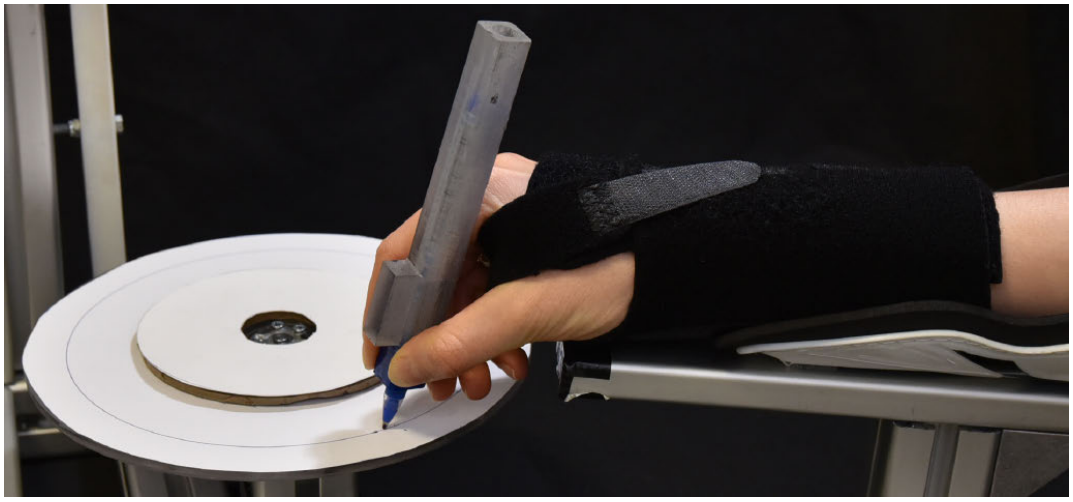

Figure 6.9. Turntable setup for user studies. Participants' hands rest in a holder and an arm-wrap is used to limit wrist motion.

Task Each trial involved investigating and comparing three combinations of drawing tools and papers. During each trial, the participants were asked to sit in front of our apparatus and to rest one of their hands on the support. Then, the instructor demonstrated each of the three stimuli. This required changing the surface on the turntable, sharpening the tools which undergo wear, and handing the tools to the participants. Next, the subjects were asked to lower the drawing tool until it reaches the surface. At this point, the instructor activated the turntable which was rotating with a previously determined speed range from 17 up to 250 mm/s (Section D.5). The full period of the velocity change was 3 seconds which was sufficient to appreciate the velocity-dependent effects while keeping the length of the user-study short. The subjects were instructed which of the stimuli were the reference and the tests, and then asked to identify which of the two test stimuli was more similar to the reference. Before answering the question, the participants could investigate each of the stimuli an unlimited number of times. To prevent any visual and auditory feedback, the participants were asked to sit behind the curtain and to wear noise-canceling headphones.

Participants For the first stage of our experiment, we invited 34 participants (20-30 years old, M/F ratio 20/14). Each of them performed an equal number of different randomly chosen comparisons. For the second stage, we asked 10 new participants (20-30 years old, M/F ratio 5/5) and each of them evaluated all 60 curated samples that in the first stage resulted in a tie. Evaluation of one

triplet took on average 30 seconds, and the whole study took approximately 60 minutes due to the additional setup performed by the instructor. Due to the length of the study, the participants were free to take a break at any time or even split the study into multiple sessions. All the participants received financial compensation.

6.3.4 Perceptual Space of Drawing Tools

When recovering a perceptual space from experimental data, it is crucial to determine its dimensionality. On the one hand, higher dimensionality allows for explaining experimental data better, but on the other hand, it may lead to overfitting and more challenging correlation with measurable properties of stimuli. Our formulation explicitly tries to find a space that correlates well with physical attributes, which is controlled by parameter λ. Here, we analyze the trade-offs between different number of dimensions and different choices of λ.

To design the perceptual descriptor of drawing tools, we use our gathered measurements. As our first dimension we opted for the Coulomb friction coefficient. Next, as a second descriptor we use overall force of vibrations. Since human perception of vibrations is a U-shaped function peaking at around 200-300 Hz. We account for this nonlinearity by equalizing the vibrational force by human sensitivity thresholds [Israr, Choi and Tan, 2006]. For the 1D space we optimize two times and separately correlate for friction and overall vibration amplitude. For the 2D space we use friction and overall vibrational power of the signal. Finally, for higher dimensional spaces we split the vibration data into uniform bands based on dimensionality of the space.

To determine the dimensionality of the space and a good value of the parameter λ, we used the accurate estimations of probabilities from the second stage of the experiment and compared them to the predictions given by perceptual spaces computed using different λ values and dimensionality. Figure 6.10 visualizes the mean match error defined as the average error in prediction of pairwise distances. We can observe an apparent gain in performance when using a 2D space which quickly tapers and does not significantly improve with higher dimensions. Based on this analysis, we decided to use a 2-dimensional space. The best one correlates with frictional coefficient and a mean value of vibration spectrogram, with a linear correlation of 0.98 and 0.95 respectively.

We present our recovered perceptual space of drawing tools in Figure 6.11. As expected, it forms clusters of the same kind of drawing tools, but also captures differences related to using different papers.

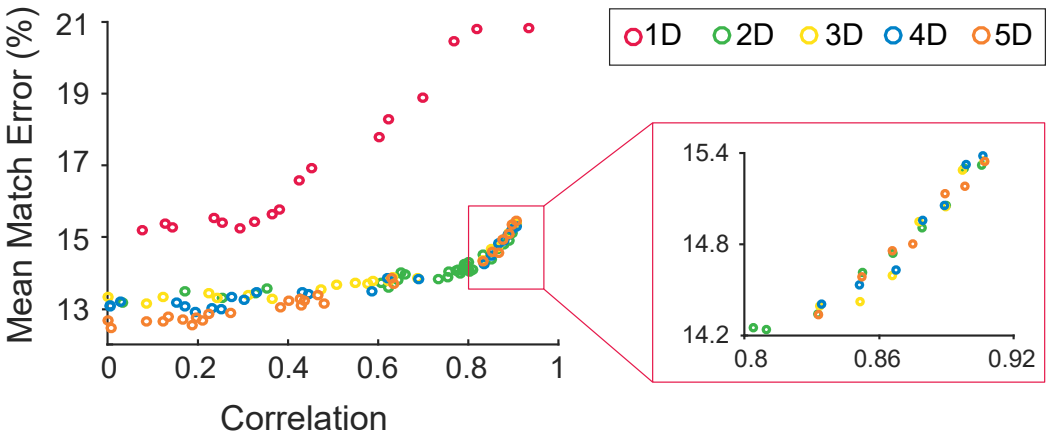

Figure 6.10. The plot shows the match of our optimized perceptual spaces to the experimental data. Different colors correspond to spaces of different dimensionality, while points with the same color correspond to different values of λ.

6.3.5 Accuracy

We performed two tests to further validate the accuracy of our space. First, we performed 5-fold validation. We splited the data from the second stage of our experiment into training and testing. Next, we optimized for 2-dimensional spaces using only training sets and use the testing for validation. An average error was 17 % which is similar to the error obtained in our dimensionality test (Figure 6.10).

To test the reliability of our perceptual space, we computed confidence intervals of tool placement using bootstrapping. We performed random sampling of the experimental data and used the re-sampled data to generate a new space which was then aligned with our original perceptual space (Figure 6.11, points). We repeated this procedure 1000 times and generated confidence intervals by drawing ellipsoids that enclose 95% and 68% of points corresponding to the same tool (Figure 6.11, ellipses). In most cases, the estimated confidence intervals were smaller than the distances between the individual tools. This suggests that the placement of the tools is reliable. The differences in confidence intervals can be explained by wear characteristics of different tools. For example, ballpoint pen and multiliner are resistant to wear and do not change over time. On the other hand pencils and charcoals get dulled by wear. This results in change of tip size affecting the physical properties of the tool, which introduces noise to human judgments.

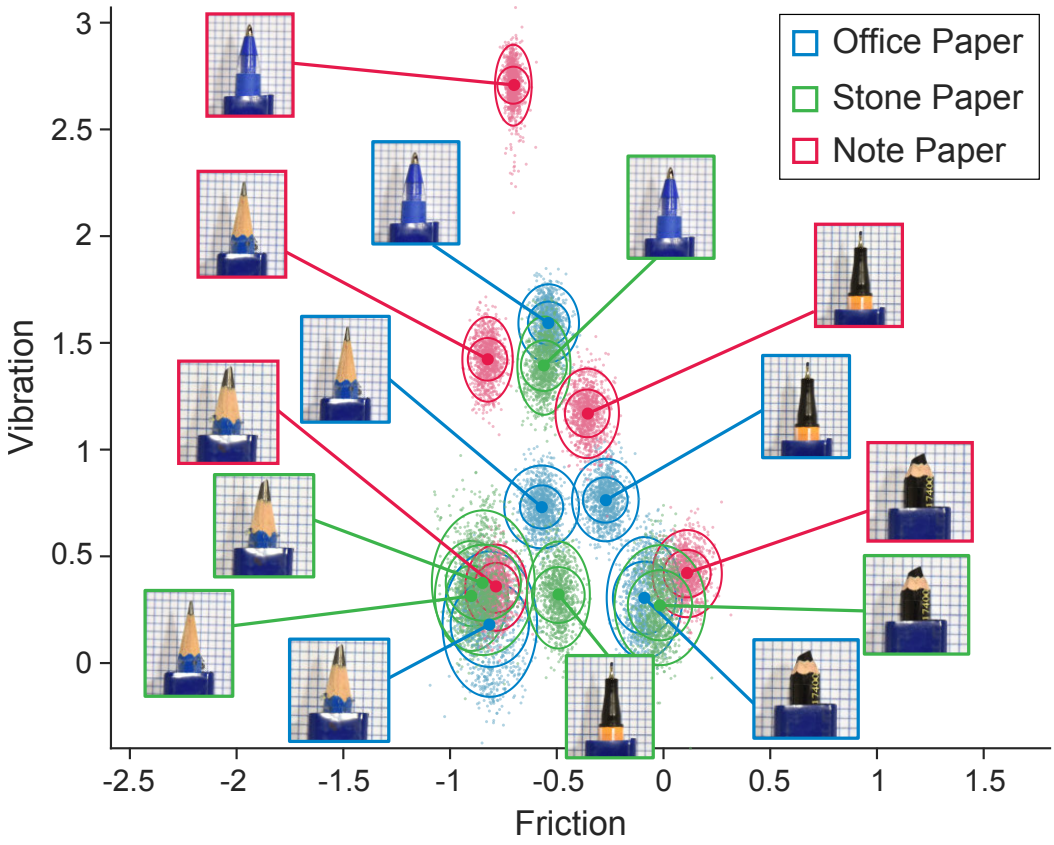

Figure 6.11. The perceptual space obtained using our data and optimization. The axis are correlated with vibration and friction measurements. Confidence intervals obtained by bootstrapping visualize 95% and 68% regions.

We also validated our optimization and data selection approach on synthetic data sets. To this end, we designed 100 synthetic tests. Each of them consisted of a random perceptual space with a fixed dimensionality, which we used to emulate our triplet selection process and a user study in a Monte-Carlo fashion. Finally, we used the synthetic data to recover the perceptual space. Figure 6.12 shows the mean and standard deviation of the match between optimized and ground-truth spaces for the various dimensionality of the initial data. The results demonstrate the the error introduced by our optimization technique is smaller than the errors reported for our space of drawing materials. In Section 6.5, we further demonstrate that despite all the inaccuracies reported in this section, our space can facilitate the process of designing drawing tools.

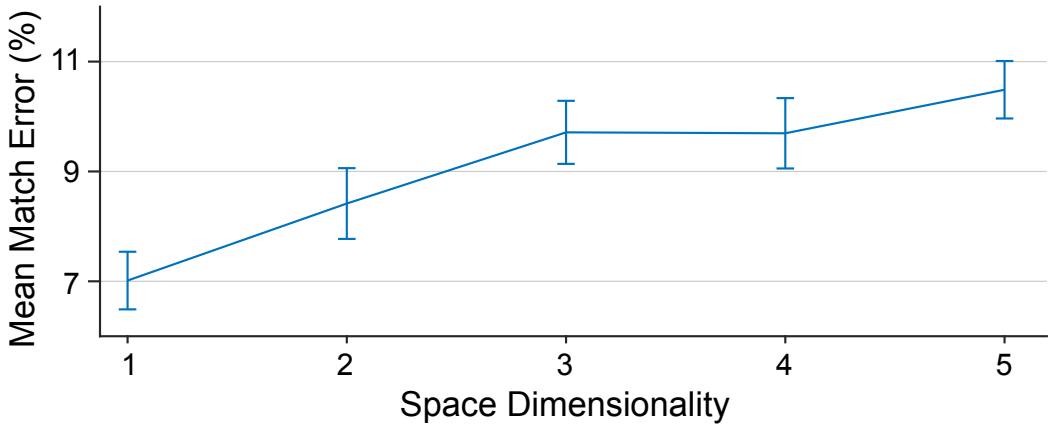

Figure 6.12. Synthetic results of optimizing perceptual spaces using our experiment design as a function of space dimensionality. The bars show mean error and standard deviation.

6.4 Physical Simulation

We now have a procedure to estimate the placement of a tool in our perceptual space given a set of physical measurements. Generating these measurements from a physical setup is inconvenient so instead we turn to physical simulation. A full contact simulation of a drawing tool with paper micro-asperities is overly expensive and many effects are still open problems in computational modeling, (Chapter 3.2.1). These difficulties inspire us to take a partially data-driven approach, wherein we model as much of the problem as we can mathematically, using fitted models to fill in the gaps. Our method is motivated by approaches in musical instrument simulation [Blood, 2009] in which the main resonator of an instrument is simulated but the complex driving force (the mouthpiece) is tackled in a data-driven fashion [Li et al., 2016]. More precisely, we model the propagation of the vibration initiated at the tip of the drawing tool using our exponential Euler integrator, while the forces acting at the tip of the stylus are generated using our data-driven approach. The data-driven forces encode the complex contact characteristic between the stylus and the surface.

6.4.1 Perception-Aware Coarsening for Exponential Euler Integrator

Our goal is to simulate the vibrational behaviors of elastic geometries in driven contact with rough surfaces. The tools and materials we work with exhibit high

stiffness, as reflected in their Young's moduli, and so are effectively modeled as linearly elastic using a small deformation assumption. The equations of motion of such the physical system is then:

$$\mathbf{M}\ddot{\mathbf{x}}(t) + \mathbf{D}\dot{\mathbf{x}}(t) + \mathbf{K}\mathbf{x}(t) + \Lambda(\mathbf{x}(t)) = \mathbf{0}, \tag{6.6}$$

where $\mathbf{x}(t)$ is the displacements at time t. $\mathbf{M}$, $\mathbf{D}$, and $\mathbf{K}$, are the mass, damping and stiffness matrices, respectively, and $\Lambda(\mathbf{x}(t))$ is the vector of external forces. In this regime, we find using only the stiffness term of Rayleigh damping model works well for our application, $\mathbf{D} = \lambda\mathbf{K}$. We can rewrite Equation 6.6 to solve for accelerations:

$$\ddot{\mathbf{x}}(t) = -\mathbf{M}^{-1}\mathbf{D}\dot{\mathbf{x}}(t) - \mathbf{M}^{-1}\mathbf{K}(\mathbf{x}(t) - \mathbf{x}_0) - \mathbf{M}^{-1}\Lambda(\mathbf{x}(t)).$$

Now we reduce the system to a set of first order equation using the substitutions:

$$\mathbf{y}_1(t) = \mathbf{x}(t), \quad \mathbf{y}_2(t) = \dot{\mathbf{x}}(t),$$

$$\begin{aligned}
\mathbf{y}_1'(t) &= \dot{\mathbf{x}}(t) = \mathbf{y}_2, \\
\mathbf{y}_2'(t) &= \ddot{\mathbf{x}}(t) = -\mathbf{M}^{-1}\mathbf{D}\dot{\mathbf{x}}(t) - \mathbf{M}^{-1}\mathbf{K}(\mathbf{x}(t) - \mathbf{x}_0) - \mathbf{M}^{-1}\Lambda(\mathbf{x}(t)),
\end{aligned}$$

or more compactly in matrix form:

$$\begin{bmatrix} \dot{\mathbf{x}}(t) \\ \ddot{\mathbf{x}}(t) \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{M}^{-1}\mathbf{K} & -\mathbf{M}^{-1}\mathbf{D} \end{bmatrix} \begin{bmatrix} (\mathbf{x}(t) - \mathbf{x}_0) \\ \dot{\mathbf{x}}(t) \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ -\mathbf{M}^{-1}\Lambda(\mathbf{x}(t)) \end{bmatrix}.$$

Assuming Rayleigh damping we can express the damping matrix as $D = \mu\mathbf{M} + \lambda\mathbf{K}$:

$$\begin{bmatrix} \dot{\mathbf{x}}(t) \\ \ddot{\mathbf{x}}(t) \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{M}^{-1}\mathbf{K} & -\mathbf{M}^{-1}(\mu\mathbf{M} + \lambda\mathbf{K}) \end{bmatrix} \begin{bmatrix} (\mathbf{x}(t) - \mathbf{x}_0) \\ \dot{\mathbf{x}}(t) \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ -\mathbf{M}^{-1}\Lambda(\mathbf{x}(t)) \end{bmatrix}.$$

Using the substitution $\mathbf{A} = \mathbf{M}^{-1}\mathbf{K}$ we arrive to the solution

$$\begin{bmatrix} \dot{\mathbf{x}}(t) \\ \ddot{\mathbf{x}}(t) \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{A} & -(\mu + \lambda\mathbf{A}) \end{bmatrix} \begin{bmatrix} (\mathbf{x}(t) - \mathbf{x}_0) \\ \dot{\mathbf{x}}(t) \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ -\mathbf{M}^{-1}\Lambda(\mathbf{x}(t)) \end{bmatrix},$$

or in a more compact form:

$$\dot{\mathbf{X}} = \mathfrak{U}\mathbf{X}(t) + \Gamma(\mathbf{X}(t)).$$

This equation has a known analytical solution given as:

$$\mathbf{X}(t) = e^{(t-t_0)\mathfrak{U}}\mathbf{X}(t_0) + e^{t\mathfrak{U}} \int_{t_0}^{t} e^{-\tau\mathfrak{U}}\Gamma(\mathbf{X}(\tau))\,\mathrm{d}\tau, \tag{6.7}$$

which is equivalent to its recursive form:

$$\mathbf{X}(t+\Delta t) = e^{\Delta t \mathfrak{U}}\mathbf{X}(t) + \int_{t}^{t+\Delta t} e^{(t+\Delta t-\tau)\mathfrak{U}}\Gamma(\mathbf{X}(\tau))\,\mathrm{d}\tau. \tag{6.8}$$

where $\mathbf{X}$ is the stacked vector of system displacements and velocities, $\mathfrak{U}$ is the system matrix and Γ collects our external forcing terms. This system can be used as the basis to build an exponential integrator [Hochbruck and Ostermann, 2010] as long as a method for evaluating the forcing integral can be devised.

Our ability to perform predictive simulation of real-world styli hinges on the representation of the Γ forcing term in Equation 6.8 and our capacity to integrate it efficiently. In our setting this term bundles a wide range of complex phenomena that drive the probe interaction behavior. The first of which are the oscillations induced by micro-scale surface contact. The fine scale at which the contact occurs quickly makes the computation intractable. We address this issue by applying perception-aware numerical coarsening. Humans do not perceive vibratory feedback uniformly. We are most sensitive around 240 Hz and regarded as insensitive bellow 50 Hz and above 500 Hz [Israr, Choi and Tan, 2006], (Figure 6.13). As a result, we can significantly speed up the computation by limiting the numerical model only to the range that is appreciable by the observer. Given the structure of our problem we adopt the spectral perspective and represent the oscillatory driving force as a linear combination of sinusoidal terms of varying phase. This enables efficiency and flexibility in simulation and modeling.

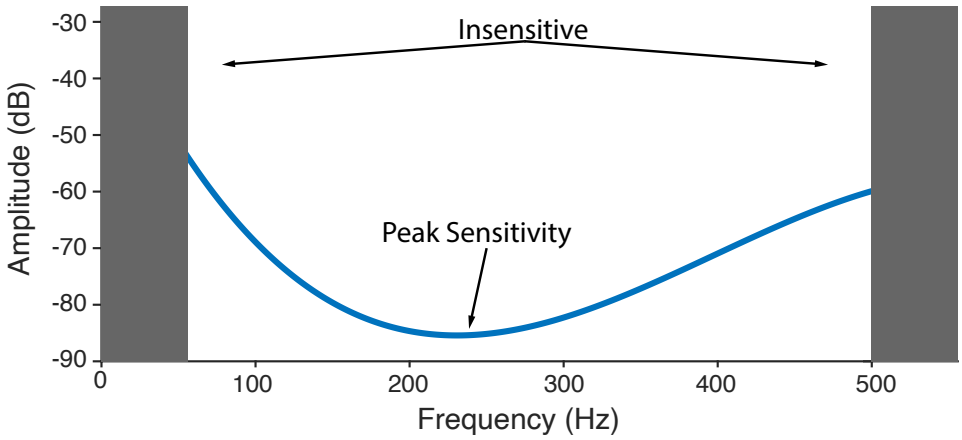

Figure 6.13. Human sensitivity to vibration as a function of frequency. We are most sensitive at around 240 Hz and regarded as insensitive bellow 50 and above 500 Hz.

First, for a single sinusoidal forcing term we can construct an exact integrator.

For clarity we will use the following substitutions:

$$\mathbf{A} = \text{diag}(\alpha) \qquad\qquad \Omega = \text{diag}(\omega),$$

where Ω designates an identity matrix scaled by phase ω, and $\mathbf{A}$ a diagonal matrix of per degree of freedom amplitudes α. We start with the base recursive version of the equation:

$$\mathbf{X}(t+\Delta t) = e^{\Delta t \mathfrak{U}}\mathbf{X}(t) + \int_t^{t+\Delta t} e^{(t+\Delta t-\tau)\mathfrak{U}}\Gamma(\mathbf{X}(\tau))\,\mathrm{d}\tau, \tag{6.9}$$

where $\Gamma(X(t)) = \mathbf{A}\sin(\omega t + \phi)$, and ϕ is the frequency shift. We are interested in finding a solution of the integral:

$$e^{(t+\Delta t)\mathfrak{U}} \int_t^{t+\Delta t} e^{-\tau\mathfrak{U}}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau.$$

First we find a solution to the integral in the base form:

$$\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau.$$

Using per partes substitution:

$$\begin{aligned} u &= e^{\tau M}\mathbf{A} & v' &= \sin(\omega\tau+\phi) \\ u' &= Me^{\tau M}\mathbf{A} & v &= -\Omega^{-1}\cos(\omega\tau+\phi) \end{aligned}$$

$$\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau = a + \int_t^{t+\Delta t} Me^{\tau M}\mathbf{A}\Omega^{-1}\cos(\omega\tau+\phi)\,\mathrm{d}\tau,$$

where $a = [-e^{\tau M}\mathbf{A}\Omega^{-1}\cos(\omega\tau+\phi)]_t^{t+\Delta t}$. Which we can reorder to:

$$\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau = a + M\Omega^{-1}\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\cos(\omega\tau+\phi)\,\mathrm{d}\tau.$$

Now we can again apply per partes:

$$\begin{aligned} u &= e^{\tau M}\mathbf{A} & v' &= \cos(\omega\tau+\phi) \\ u' &= Me^{\tau M}\mathbf{A} & v &= \Omega^{-1}\sin(\omega\tau+\phi) \end{aligned}$$

$$\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau = a + b - M\Omega^{-1}\int_t^{t+\Delta t} Me^{\tau M}\mathbf{A}\Omega^{-1}\sin(\omega\tau+\phi)\,\mathrm{d}\tau,$$

where $b = M\Omega^{-1}[e^{\tau M}\mathbf{A}\Omega^{-1}\sin(\omega\tau+\phi)]_t^{t+\Delta t}$. Now we can reorder again to get:

$$\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau = a + b - M^2\Omega^{-2}\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau,$$

To remove the inverse of Ω we multiply both sides by Ω^2 from the left. Note since this is a scaled identity matrix it holds the commutative law and we can move it around as needed.

$$\Omega^2\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau = \Omega^2 a + \Omega^2 b - M^2\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau$$

Now we subtract $-M^2\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau$, take the integral in front of parenthesis:

$$(\Omega^2+M^2)\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau = \Omega^2 a + \Omega^2 b.$$

And finally the solution is:

$$\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau = (\Omega^2+M^2)^{-1}(\Omega^2 a + \Omega^2 b),$$

which expands to:

$$\begin{aligned}\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau =& (\Omega^2+M^2)^{-1}(\\ &\Omega^2[-e^{\tau M}\mathbf{A}\Omega^{-1}\cos(\omega\tau+\phi)]_t^{t+\Delta t}\\ &+\Omega^2 M\Omega^{-1}[e^{\tau M}\mathbf{A}\Omega^{-1}\sin(\omega\tau+\phi)]_t^{t+\Delta t}).\end{aligned}$$

Which finally expands to:

$$\begin{aligned}\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau =& (\Omega^2+M^2)^{-1}(\\ &\Omega e^{tM}\mathbf{A}\cos(\omega t+\phi)\\ &-\Omega e^{(t+\Delta t)M}\mathbf{A}\cos(\omega(t+\Delta t)+\phi)\\ &+M e^{(t+\Delta t)M}\mathbf{A}\sin(\omega(t+\Delta t)+\phi)\\ &-M e^{tM}\mathbf{A}\sin(\omega t+\phi)).\end{aligned}$$

We can take the common term e^{tM} in front of parenthesis:

$$\begin{aligned}\int_t^{t+\Delta t} e^{\tau M}\mathbf{A}\sin(\omega\tau+\phi)\,\mathrm{d}\tau =&(\Omega^2+M^2)^{-1}(e^{tM})(\\ &\Omega\mathbf{A}\cos(\omega t+\phi)\\ &-\Omega e^{\Delta t M}\mathbf{A}\cos(\omega(t+\Delta t)+\phi)\\ &+Me^{\Delta t M}\mathbf{A}\sin(\omega(t+\Delta t)+\phi)\\ &-M\mathbf{A}\sin(\omega t+\phi)).\end{aligned}$$

Now we can plug thevsolution back to Equation 6.9 with the substitution of $M = -\mathfrak{U}$.

$$\begin{aligned}X(t+\Delta t) =&e^{\Delta t\mathfrak{U}}\mathbf{X}(t)+e^{(t+\Delta t)\mathfrak{U}}(\Omega^2+\mathfrak{U}^2)^{-1}(e^{-t\mathfrak{U}})(\\ &\Omega\mathbf{A}\cos(\omega t+\phi)\\ &-\Omega e^{-\Delta t\mathfrak{U}}\mathbf{A}\cos(\omega(t+\Delta t)+\phi)\\ &\mathfrak{U}\mathbf{A}\sin(\omega t+\phi)\\ &-\mathfrak{U}e^{-\Delta t\mathfrak{U}}\mathbf{A}\sin(\omega(t+\Delta t)+\phi)).\end{aligned}$$

The factor $e^{(t+\Delta t)\mathfrak{U}}$ could have been stored in the integral and therefore it can be moved inside the parenthesis:

$$\begin{aligned}X(t+\Delta t) =&e^{\Delta t\mathfrak{U}}\mathbf{X}(t)+(\Omega^2+\mathfrak{U}^2)^{-1}(e^{(t+\Delta t)\mathfrak{U}}e^{-t\mathfrak{U}})(\\ &\Omega\mathbf{A}\cos(\omega t+\phi)\\ &-\Omega e^{-\Delta t\mathfrak{U}}\mathbf{A}\cos(\omega(t+\Delta t)+\phi)\\ &\mathfrak{U}\mathbf{A}\sin(\omega t+\phi)\\ &-\mathfrak{U}e^{-\Delta t\mathfrak{U}}\mathbf{A}\sin(\omega(t+\Delta t)+\phi)).\end{aligned}$$

After simplifying we get to the final solution:

$$\begin{aligned}X(t+\Delta t) =&e^{\Delta t\mathfrak{U}}\mathbf{X}(t)+(\Omega^2+\mathfrak{U}^2)^{-1}(\\ &e^{\Delta t\mathfrak{U}}(\Omega\mathbf{A}\cos(\omega t+\phi)+\mathfrak{U}\mathbf{A}\sin(\omega t+\phi))\\ &-\Omega\mathbf{A}\cos(\omega(t+\Delta t)+\phi)-\mathfrak{U}\mathbf{A}\sin(\omega(t+\Delta t)+\phi)).\end{aligned}$$

The exponential integrator solves for displacements and velocities. Since we would like to get accelerations we need to calculate the first derivative of the solution:

$$X(t) = e^{t\mathfrak{U}}\int_0^t e^{-\tau\mathfrak{U}}\Gamma(X(\tau))\,\mathrm{d}\tau \tag{6.10}$$

We start by applying the derivative product rule:

$$X'(t) = (e^{t\mathfrak{U}})' \int_0^t e^{-\tau\mathfrak{U}}\Gamma(X(\tau))\,\mathrm{d}\tau + e^{t\mathfrak{U}}\left(\int_0^t e^{-\tau\mathfrak{U}}\Gamma(X(\tau))\,\mathrm{d}\tau\right)',$$
$$X'(t) = \mathfrak{U}e^{t\mathfrak{U}} \int_0^t e^{-\tau\mathfrak{U}}\Gamma(X(\tau))\,\mathrm{d}\tau + e^{t\mathfrak{U}}\left(\int_0^t e^{-\tau\mathfrak{U}}\Gamma(X(\tau))\,\mathrm{d}\tau\right)'.$$

Now we use the property of definitive integrals:

$$\left(\int_0^b f(x)\,\mathrm{d}x\right)' = (F(b)-F(0))' = F'(b)-0 = (b)'f(b) = f(b),$$

and we get that:

$$X'(t) = \mathfrak{U}e^{t\mathfrak{U}} \int_0^t e^{-\tau\mathfrak{U}}\Gamma(X(\tau))\,\mathrm{d}\tau + e^{t\mathfrak{U}}e^{-t\mathfrak{U}}\Gamma(X(t))$$
$$X'(t) = \mathfrak{U}e^{t\mathfrak{U}} \int_0^t e^{-\tau\mathfrak{U}}\Gamma(X(\tau))\,\mathrm{d}\tau + \Gamma(X(t)).$$

In our case of sinusoidal forces the integral has already known solution from Equation 6.10:

$$X'(t) = \mathfrak{U}X(t) + \mathbf{A}\sin(\omega t + \phi)$$

An interesting property of exponential integrators is that the contribution of each individual forcing term can be evaluated separately and then combined to form the final result. This property can be shown as follows.

$$X(t) = e^{t\mathfrak{U}} \int_0^t e{-}\tau\mathfrak{U} \sum_i \Gamma_i(X(\tau))\,\mathrm{d}\tau$$
$$X(t) = e^{t\mathfrak{U}} \sum_i \int_0^t e{-}\tau\mathfrak{U}\Gamma_i(X(\tau))\,\mathrm{d}\tau$$
$$X(t) = \sum_i e^{t\mathfrak{U}} \int_0^t e{-}\tau\mathfrak{U}\Gamma_i(X(\tau))\,\mathrm{d}\tau$$

Another interesting property of the exponential integrator is that the effect of a forcing term can be arbitrarily scaled which allows to compute the effect of each forcing term for a unitary force and then scale the result appropriately. The

proof of this property comes properties of the integral:

$$X_k(t) = e^{t\mathfrak{U}} \int_0^t e^{-\tau\mathfrak{U}} k\Gamma(X(\tau))\,\mathrm{d}\tau$$
$$X_k(t) = ke^{t\mathfrak{U}} \int_0^t e^{-\tau\mathfrak{U}} \Gamma(X(\tau))\,\mathrm{d}\tau$$
$$X_k(t) = kX(t)$$

These two properties allow a set of solutions to the single sinusoidal term problem to form a linear basis that can be applied to represent any signal made up of a linear combination of these sinusoids. Practically, we exploit this during design tasks by precomputing a simulation basis for a probe and then synthesizing its interaction with new surfaces by summation of these pre-simulated results.

Special care should be also taken to achieve optimal efficiency with the recursive formulation. There are two options how to tackle the computation of matrix exponential. In previous work [Michels et al., 2014] tools such as expokit [Sidje, 1998] were used to directly calculate the effect of matrix exponential on a vector and bypass the expensive computation of a dense matrix. The speed of such solution depends on the stiffness of the simulated material. Very good performance can be achieved for Young's moduli in the range of kilo to mega Pascals. However, for stiff materials, such as many 3D printed polymers the Young's modulus is in range of giga Pascals. Such extreme stiffness slows down the computation and makes it impractical. In our experiments we have found out that the most efficient method is to precompute the matrix exponential. This can be done once at the start of the simulation and the most expensive operation becomes dense matrix times vector multiplication that we can offload to the GPU.

To validate our simulation and gather parameters for our method, we run a test simulation. We 3D printed a bar on Objet260 printer using VeroClear material. The bar was $210 \times 20 \times 2$ mm and was rounded at the end (radius 10 mm) with a circular cutout of 5 mm radius located 10 mm from the end. The bar was clamped to a table on a 30 mm section, preloaded to 5 cm and released. First, we captured the bar's oscillations using an attached accelerometer. Next, material parameters (Young's modulus and Rayleigh damping) were optimized to fit our measurements. The recovered numerical Youngs modulus was 0.45 GPa and the stiffness term of Rayleigh damping was 1.25e−3. Figure 6.14 demonstrates the excellent agreement between simulation and experiment.

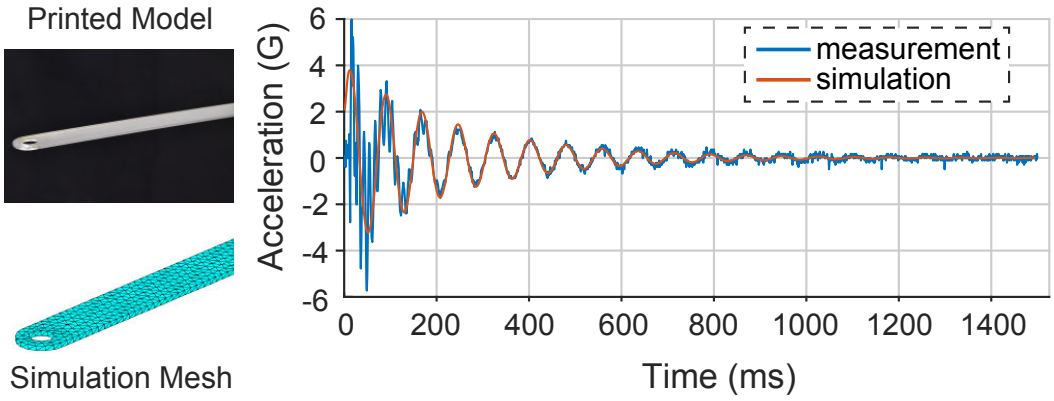

Figure 6.14. Oscilating 3D printed bar captured with accelerometer (blue) and physical simulation of the bar (red).

6.4.2 Recovering the Forcing Signal

To predict the behavior of drawing tools we need to recover the forces that are acting on the tool due to surface interaction. This could be traditionally done by simulating the range of diverse physical behaviors acting on the tool, e.g., frictional stiction, stick-slip behavior, viscous damping, wear, deposition and so forth. However, this simulation strategy poses a number of modeling and performance obstacles in our setting. Many of these effects are expensive to simulate, while for others it is not at all clear that suitable models exist to be simulated, irrespective of cost. We thus adopt a hybrid data-driven approach to model the forces acting on the drawing tool mediated by our exponential integrator derived above. We first construct our data-driven forces below and then apply them to drive our simulation.

Inspired by our integrator we choose to model the surface forcing term as a linear combination of sinusoidal components that are scaled as a function of both the compliance of the stylus tip material and the tip size itself. Concretely we represent the surface in the frequency domain as amplitudes of sinusoids from the range relevant to haptic feedback (1-500 Hz) Israr, Choi and Tan [2006], and scaling factors as piecewise cubic Hermite polynomials with ten control points uniformly spaced at 50 Hz intervals.

We compute parameters for our model by solving the optimization:

$$\underset{F_S * t_\Psi * m_f}{\arg\min} \sum_{\Psi} \sum_{S} \left|\left| Integrator(\Psi, F_S * t_\Psi * m_f) - M_{\Psi,S} \right|\right|, \tag{6.11}$$

where Ψ is a shape and material combination, F_S is forcing term of surface S, t_Ψ, and m_Ψ are tip and material scaling parameters. Here integrator is the Ex-

ponential Euler integrator we derive in Chapter 6.4. The integrator takes as an input the assembled forcing term and returns a simulated measurement. We minimize the L2 norm between our simulation and the measured ground truth $M_{\Psi,S}$. The required simulations can be quickly evaluated using our pre-simulated basis results.

We printed, measured, and simulated a total of 27 tools with a unitary forcing signal in the range of frequencies relevant to haptic feedback (1-500 Hz). We then optimized this data with Equation 6.11 to acquire our final material forcing terms.

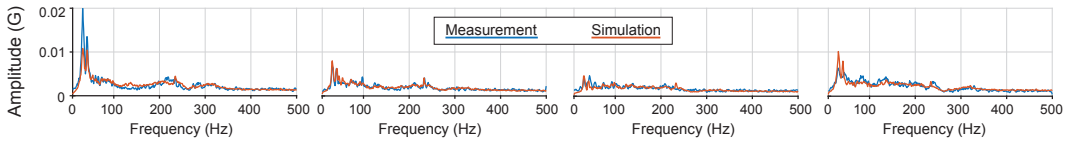

Figure 6.15. Fit of testing dataset at a fixed velocity. We compare physical measurements (blue) to our data-driven simulation (red). We can see a good match of estimated vibration for the testing dataset.

To validate the data-driven model we performed a leave one out test of the recovered forcing term. One exemplar was removed from the data set, our model was trained on the remaining samples and then evaluated by fitting the test exemplar, (Figure 6.15). We can see that we achieved good match with testing data. The main discrepancies come from sharp peaks that are very challenging to fully reproduce. For full evaluation please refer to Appendix C. Next, we also evaluated the effects of the scaling terms, (Figure 6.16). The material scaling factor (Figure 6.16 left) is inversely proportional to material softness. A soft material is predicted to damp the vibrations of the tool. However, increased softness leads to more pronounced stick-slip behavior which is then compensated for by increasing the material scaling term. The effect of tip size (Figure 6.16 right) is more subtle. In general large tip size leads to more vibration damping. The scaling factors for tip size of 1 and 2 mm are very similar. This is caused by the wear experienced by the material which sands down the finer tip quicker, therefore, making them equivalent.

Our simulations exhibit a minor mismatch in lower frequencies around 25 Hz. This is caused by our choice of eleven evenly spaced control points at 50 Hz intervals, which fails to reconstruct this sharp jump. Our control point setup was chosen to avoid standard issues with over-fitting of polynomial curves. For general simulation, we would be interested in a precise match of the full spectra. Here, however, the slight mismatches for low frequencies are not reflected in user

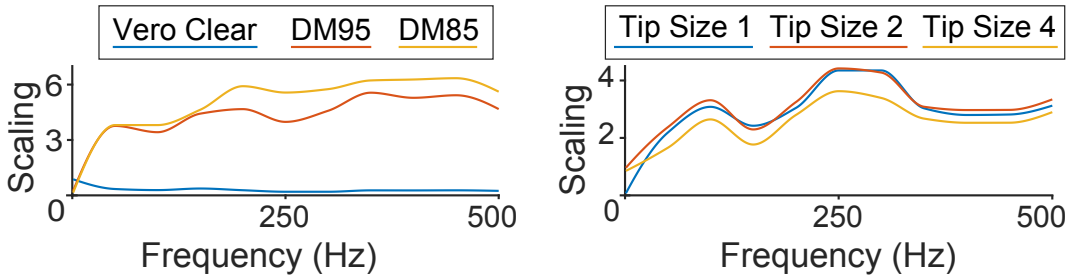

Figure 6.16. Material (left) and tip (right) scaling factor. Both factors were optimized jointly.

data as humans are less sensitive to vibrations in the low range than to higher frequency vibrations where our simulator matches experimental data well.

6.4.3 Results

The purpose of the simulation is to predict the behavior of 3D printed tools without physically manufacturing them. To evaluate the precision of our simulation we conducted a test in which we predicted the vibratory response of an interpolated tool. We interpolated between two designs. The first design has a half-sphere tip of 4 mm in diameter printed in DM85. The second design has a half-sphere tip of 1 mm diameter and is printed in VeroClear. Both designs were measured on standard office paper at various velocities (Figure 6.17 blue). The new tool has linearly interpolated material and tip parameters. To predict its behavior we used our full simulation pipeline. First, we measured the response of the original designs to recover tip and material scaling factors. Next, we used our exponential Euler integrator to simulate the response of the interpolated tool. We used forcing term we recovered for office paper modulated by linearly interpolated material and tip scaling factors (Figure 6.17 red). We compare our full simulation pipeline to a simulation where we would not interpolate the material and tip parameters (Figure 6.17 gray). We can see that our prediction matches reality well. The deviations can be attributed to the general noisiness of vibrational measurements, as well as, our printing process since we used dithering to generate the interpolated material.

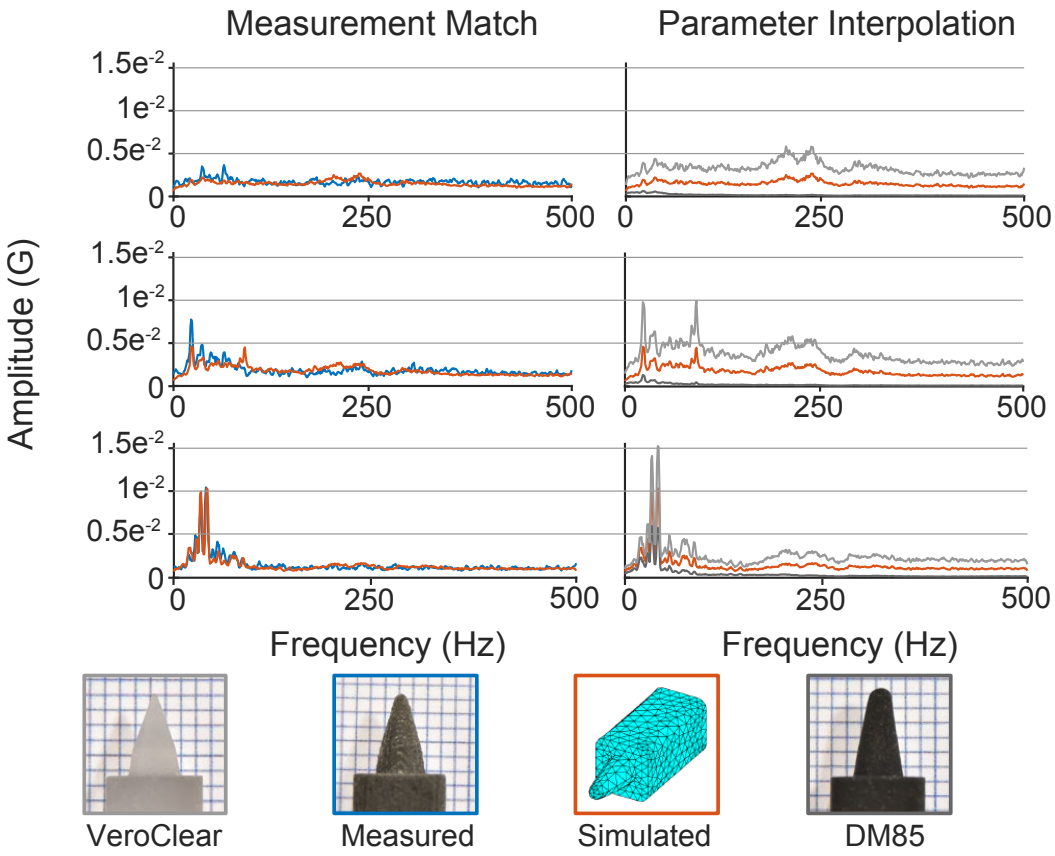

Figure 6.17. An interpolated design was measured (blue line) and simulated using our pipeline (red line). We also compare to results of our simulation if we would not interpolate tip and material scaling parameters (gray lines).

6.5 Application For Drawing Tool Design

Finally, we demonstrate how our perceptual space and simulator can be used to aid the process of designing digital drawing tools.

6.5.1 Perceptual Space Exploration

Because we correlate the dimensions of the perceptual space with measurable properties of drawing tools, we can easily embed new tools into the space and evaluate their pairwise similarities based on their measurements or simulation. To demonstrate this, we extended our previously derived space with several other drawing tools. We considered four different categories. The first group consists of our initial set of traditional drawing tools extended by two new materials:

a crayon and a felt-tip marker (Figure 6.18, red). The second category is 3D-printed tools (Figure 6.18 green). They were fabricated on a Stratasys Objet 260 Connex printer using different materials ranging from VeroClear (low friction) to DM85 (high friction). The tip size of the tools varied from 1 to 4 mm. This category also includes tools that were covered with a Teflon tape to lower the friction further. All tools in the first and the second category were used on our three different kinds of paper (Figure 6.4). The third group consisted of the same 3D printed tools as in the previous category, but this time they were used on two different artificial surfaces: glass and glass with a screen protector (Figure 6.18 blue). This category of drawing tools demonstrates what one can achieve with today's tablets and multi-material printing technologies. We also included commercial solutions such as a Wacom Tablet Pen, Microsoft Surface Pen, and Apple Pencil (Figure 6.18 yellow).

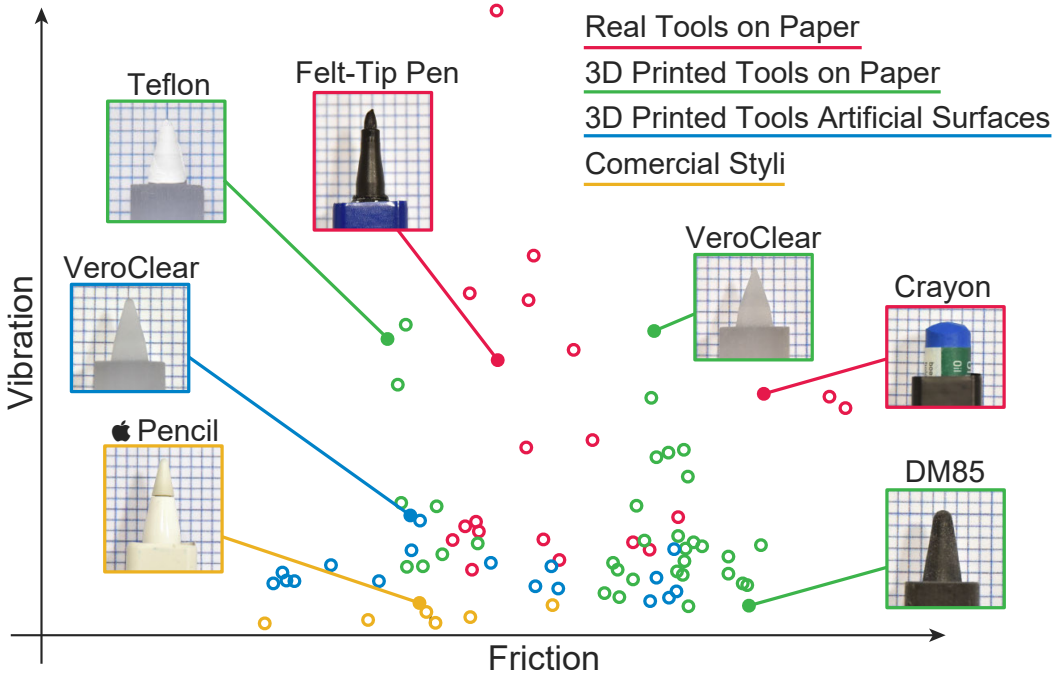

Figure 6.18. Perceptual space grouped by means of fabrication: real tools (red), 3D printed tools on paper (green), and artificial substrates (blue), and commercial styli (yellow).

To validate the accuracy of placing novel drawing tools into our perceptual space, we conducted a user experiment.

Stimuli We randomly chose seven new surface-tool combinations that were not used for for the computation of the perceptual space. They included four traditional drawing tools (a crayon on stone and office paper, and felt-tip on note and stone paper), two 3D printed tools (Teflon covered VeroClear on note paper and DM95 on office paper), as well as two digital styli (Apple Pencil, and Wacom rubber nib on glass). For each of the tools, we formed three triplets by adding random tools as tests. In total, we obtained 24 stimuli.

Participants 12 participants (20-30 years old, 7/5 M/F ratio) took part in this experiment. Each of them was financially compensated.

Procedure Participants evaluated each triplet by identifying the test tool that was more similar to the reference. Similarly to our original experiment, we avoided visual and auditory feedback but allowed participants to draw with the tools freely. The participants could also swap the tools as many times as they wanted and they were given unlimited time for experimentation.

Results We compared the results of the experiment with predictions provided using our space. The prediction matched the popular opinion in 19 out of 24 cases. Z-test revealed that our model predicts preference significantly better than random chance (p-value < 0.0022). On average our prediction had an error of 19% and correlated to the experimental results with a Pearson correlation of 0.81 (p-value < 0.0001). This suggests a strong linear relationship between our predictions and ground truth. The results support our data validation (Section 6.3.5) and suggest that our methodology translates well to freehand drawing. Figure 6.19 presents the detailed results of the experiment.

Our perceptual space can be used to draw several conclusions. First, the group of commercial products overlaps to only a small extent with traditional tools. Interestingly, our 3D-printed tools, on artificial surfaces, extended the overlap, which suggests that there is room to improve digital styli. The plot also reveals that digital styli on the glass surface and the screen protector can achieve a large variation in friction, which is almost sufficient to represent all traditional drawing tools, but they lack correct vibratory feedback which is currently limited to drawing implements such as an 8B pencil on smooth stone paper. This suggests that the community should explore the design of tablet covers which can provide this missing cue.

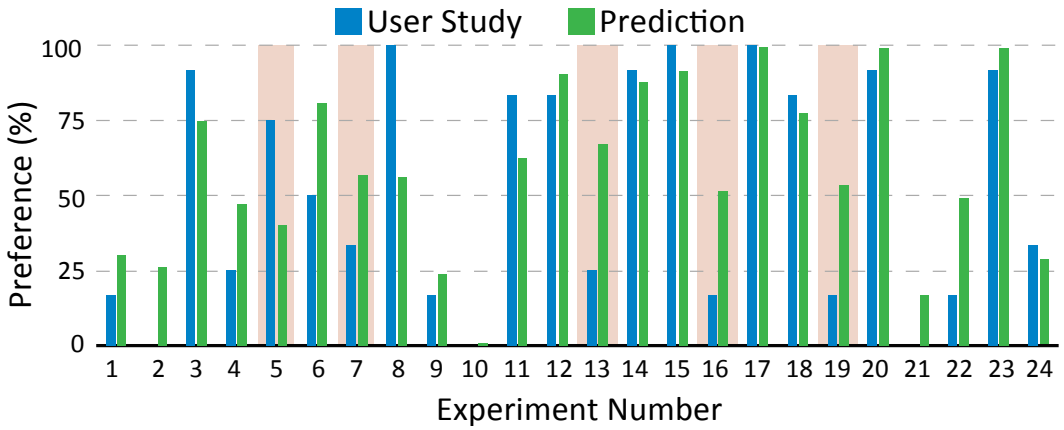

Figure 6.19. The results of the free drawing user study. For each out 24 triplets, we plot preference obtained in our validation experiment next to the prediction coming from the perceptual space. The red background indicates the failure cases where the majority vote did not agree with our prediction.

6.5.2 Optimizing Digital Styluses

Finally, we leverage the ability of our perceptual space to estimate similarities in haptic feedback and demonstrate the application of our work to digital stylus fabrication. For this purpose, we consider optimizing a stylus for two different surfaces: glass and a screen protector, and validate our results in a user experiment.

Stimuli In the first case, we considered a screen protector as our artificial surface. On this protector, we measured two 3D-printed styli, one in VeroClear with a tip diameter of 1 mm, and one in DM85 with a tip diameter of 4 mm. Figure 6.20 (left) shows their relation to the real drawing tool in our space. One can see that neither of them matches the real tool, and a better design can be achieved. We used a line search powered by our simulation pipeline (Section 6.4) to generate a design in between the two synthetic styli. For the frictional coefficient, we used linear interpolation. The second case consider an improvement of digital styli on a glass surface typical for tablets. We measured the two 3D-printed designs from the previous experiment. Figure 6.20 (right) shows the relative placement of the real tool, commercial stylus, and our 3D printed tips. Once again we use our simulation pipeline to design an interpolated stylus.

Participants We invited 16 participants to investigate the case with a screen protector and 12 for the glass surface. The participants were between 20-35 years old, and M/F ratio was 9/7. Each participant was compensated for taking part in the study.

Procedure Each participant was presented with all the digital tools at once and asked to select one of them that is most similar to the real tool that served as a reference. We decided not to limit visual and auditory feedback, as the study consisted solely of artificial materials so the bias from these cues should be minimal.

Results For the case with the screen protector, 11 participants selected our interpolated design, 4 the design printed in VeroClear material, and 1 a design made of DM85 material. For the second case, 10 out of 12 participants preferred our interpolated design and 2 preferred the VeroClear design. Chi-square goodness of fit showed that there exists a significant difference in preference among the groups of stimuli (p-values of < 0.018 and < 0.001). In a post-hoc analysis a pairwise comparison using Z-tests with Holm-Bonferroni correction revealed that the preference towards our designs is also statistically significant (p-values of < 0.036 and < 0.008).

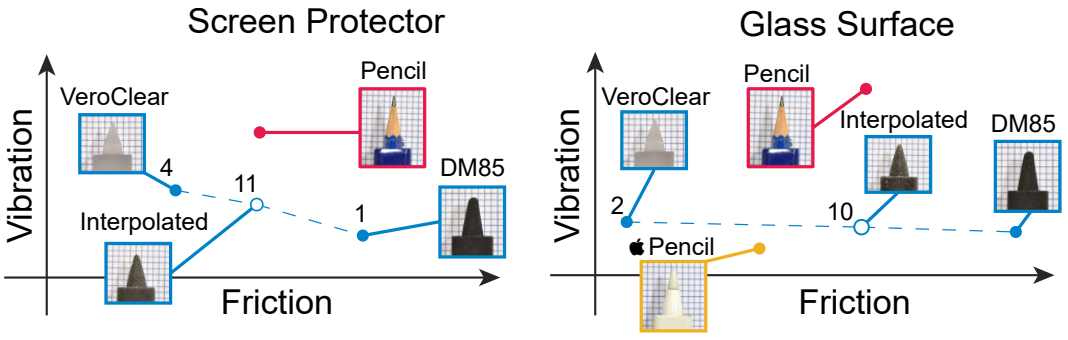

Figure 6.20. Details of our perceptual space for optimization of digital styli. We show two cases: interpolation of 3D-printed styli on a plastic screen protector (left), and interpolation of 3D-printed styli on the glass surface in comparison to the Apple Pencil (right). Number indicate subject preference when compared to real pencil.

6.6 Limitations and Future Work

Our perceptual model for drawing tools was designed under the assumption that such tools exhibit isotropic behavior. While this is a good assumption for a large group of commonly used tools (pens, pencils, charcoals) there are cases when this assumption is violated. For example, ink pens are designed to deposit material only in specific drawing directions and only ride smoothly on the surface in that direction. When used in the other direction, they oppose the motion, creating a drastically different response. Another limitation of our model, inherited from the tools themselves, is that they manifest little variation in compliance. However, tools such as brushes and brush pens are often used based on how compliant they are. We also do not account for the wear of drawing materials. If a tool is not always resharpened, the wear will eventually result in a measurable and perceivable change in haptic behavior. In the future, it would be interesting to account for anisotropies, compliance, and wear, as well as to investigate what their impact on perception is.

Usually, perceptual spaces are scaled in just-noticeable-difference units. This proved challenging in our case. The identification of tools varies with experience. While a novice artist might have a hard time distinguishing between various pencils, an expert could readily tell the difference. Our perceptual space was created with average users. We believe it still applies to more experienced users regarding relative rather than absolute similarity predictions.

Our simulation has the limitations of a typical data-driven method. We cannot simulate tools that are vastly different from what we measured and our approach only handles interaction with surfaces for which the forcing term is known. Including new surfaces requires additional measurements. The same is true for material scaling factors and tip geometries, they have to be recovered for new tools that can not be interpolated. Our applications also requires friction information. We use simple linear interpolation to estimate it for the newly interpolated tools. While such a simple solution proved useful in our application, a more involved approach could be useful to more precisely pinpoint the location of an interpolated material in the perceptual space.

In our design, we did not consider durability constraints. Therefore, some of our current designs exhibit wear and cannot be directly used as commercial products. In the future, it would be interesting to augment our perceptual design problem with additional constraints such as durability or cost. Another interesting possibility is to design tools which represent a certain group of real materials.

6.7 Conclusion

In this work, we have examined the problem of modeling and fabricating digital styli that have the *feel* of real drawing tools. To this end, we measured, analyzed, and characterized properties of a representative set of traditional tools. We formulated a probabilistic likelihood-based NMDS optimization to identify features that influence perception and to derive a perceptual space of drawing materials. This enabled the evaluation of perceived differences between designs. Furthermore, we applied perception-aware coarsening to design a Exponential Euler integrator that uses a data-driven forcing term to estimate the vibratory feedback of a digital stylus. The resulting simulation method allows us to embed new tool designs into the perceptual space without the need for fabricating them or conducting expensive and time-consuming experiments. Finally, we demonstrated that our approach can aid in the process of designing new styli, and validated our results in user experiments.

Chapter 7

Fabrication-in-the-Loop Co-Optimization of Surfaces and Styli for Drawing Haptics

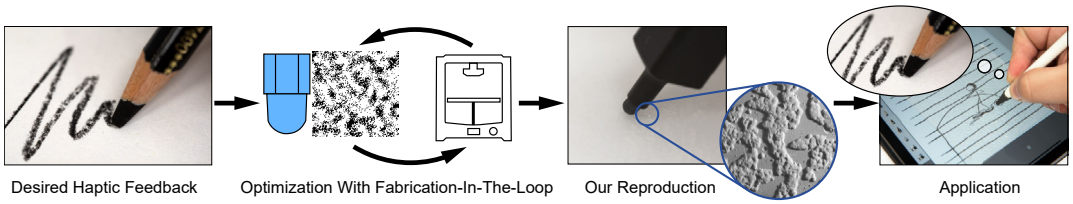

Figure 7.1. We propose a data-driven method for mimicking haptic feedback of drawing tools. Our method uses fabrication-in-the-loop design enabled by our data-driven surrogate model which automatically handles exploration-exploitation trade-offs and minimizes the amount of printed samples. The final stylus-surface combinations are manufacturable on commonly available hardware and can be directly integrated into current digital drawing solutions.

In Chapter 6 we demonstrated the benefits of exploiting perception in the design of digital styli with desired haptic response. However, optimizing the stylus alone is not sufficient to mimic the traditional instruments. The haptic feedback of drawing tools is a result of a complex interaction of the drawing tool with the substrate. The results of such interaction are transferred to the finger as a resistance to the movement and vibration. This interaction is governed by multiple coupled phenomena [Blau and Gardner, 1996]. The contact of a drawing tool with the micro-geometry of the substrate produces a specific frictional response.

The friction is further modulated by the worn material acting as a lubricant. Additionally, the impacts of cellulose layers of drawing substrate coupled with frictional stick-slip evoke a vibratory response (Figure 7.2). In this chapter, we build upon the perceptual space presented in Chapter 6 and tackle the problem of fabricating digital tools and substrates that jointly closely reproduce the haptic feedback of their traditional counterparts. Unfortunately, as in many fabrication problems, a direct reproduction of the physics governing the haptic feedback is not currently a viable solution. First, the reproduction of drawing surfaces and tools has to respect the limitations of fabrication techniques (e.g., printer resolution, material availability). Second, some of the phenomena driving the haptics of real tools are not desiredable for digital tools. For example, it is impractical to make digital tools wear as pencils do or to introduce a rolling ball lubricated by ink for replicating a ballpoint pen. As a result, the limitations of the fabrication techniques, as well as the absence of some physical processes, must be compensated by modifying other aspects of the digital designs. This leads to a challenging high-dimensional co-optimization between the stylus (shape and material) and the surface of the drawing tablet (microgeometry and material).

Finding a set of digital drawing tools replicating a particular feel can be formulated as an optimization similar to a typical specification-to-fabrication process [Chen et al., 2013]. Such an approach, however, usually requires efficient and accurate numerical simulation, which in this case, would simulate the intimate contact between a drawing instrument and a substrate. However, the complexity of the coupled phenomenons, the scale at which they occur, as well as imperfections in fabrication processes pose significant challenges to simulating all required effects. Accurate simulation of some of the effects, e.g., the coupling between friction and viscous damping, is still an open research problem [Chen et al., 2017]. Data-driven simulation techniques [Chen et al., 2015] could potentially aid modeling the complex and coupled phenomena. However, building a general data-driven simulation capable of handling a wide range of digital drawing tools requires sampling the high-dimensional design space of stylus-surface interaction.

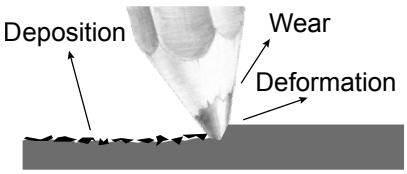

Figure 7.2. Frictional contact

To address the above challenges, we propose a novel fabrication-in-the-loop method for co-optimizing the desired haptic feel of a stylus-surface combination. We base our method on two key components. First, to address the challenging co-optimization, we refrain from directly matching the properties of the traditional tools (e.g., geometry and material). Instead, we optimize the tools based on a

characterization of the haptic feedback given by our proposed perceptual space of drawing tools, which enables focusing on perceptually-relevant tool characteristics. Second, we employ a fabrication-in-the-loop approach that systematically explores the space of possible designs in a search for the optimal one. We minimize the number of fabricated tools by incrementally building a data-driven surrogate model of haptic feedback. The model is based on Gaussian Processes [Rasmussen and Williams, 2005] and provides us with confidence bounds on predicted haptic behavior. Our method uses these bounds to formulate an efficient sampling strategy that automatically balances the exploration-exploitation trade-off. Additionally, the surrogate model allows us to transfer the knowledge between fabrication processes to accelerate optimization. Thanks to the above approach, our optimization is feasible despite the time-consuming fabrication process included in the optimization loop.

We evaluate the effectiveness of our solution in a series of free-hand drawing experiments performed with casual and professional users. The results demonstrate that our method enables the fabrication and customization of digital drawing tools such that they closely match traditional drawing instruments and users' expectations regarding haptic feedback. In some cases, our reproductions have proven to be hard to distinguish even from traditional counterparts, which is not the case for the state-of-the-art solutions. When compared to the industrial and research reproduction, our tools outperform them by a significant margin in all test-cases. The wide range of fabrication techniques used in our experiments makes our tools ready for integration with off-the-shelf digital drawing solutions. To summarize, our main contributions are:

- formulation of perception-driven optimization of target haptic feedback;
- parametrization of a design space for co-optimization of haptic feedback produced by a stylus-surface pair;
- a data-driven surrogate model for predicting haptic feedback of stylus-surface pairs;
- an algorithm for practical fabrication-in-the-loop optimization with an efficient sampling strategy of the design space that maximizes expected improvement;
- an application of the algorithm to the design of drawing tools;
- validation of manufactured replicas of drawing tools in blind, free-hand experiments with casual and professional artists.

7.1 Overview

The goal of our work is to formulate an optimization procedure for digital drawing tools (i.e., drawing surface and stylus) such that they provide desired haptic feedback. We start by formulating the low-dimensional design space of the tools, which describe both a wide range of styli (Section 7.2.1) and drawing surfaces (Section 7.2.2). To measure the error between different designs, we employ a perceptual metric for drawing tools from the previous chapter. To evaluate the error, we propose a fabrication-in-the-loop approach where we use a data-driven surrogate model (Section 7.3.2). We build the model to handle the stylus-surface co-optimization (Section 7.3.2) and include transfer learning to accelerate optimization of unobserved fabrication processes. To minimize the fabrication effort, we propose an acquisition function that efficiently explores the design space by automatically handling the exploration-exploitation trade-off (Section 7.3.3). We demonstrate how our model can be used to optimize for the desired interaction of a stylus-surface combination (Section 7.4). We validate our method by comparing with naive approaches and observing the quality of our manufactured tools (Section 7.5). Additionally, we compare our method in a series of free drawing user studies. In blind experiments, we demonstrate that our tools can model target haptic feedback, (Section 7.6.2), are preferred to state of the art approaches (Section 7.6.3), and produce realistic haptic sensations (Section 7.6.4). Finally, we evaluate our reproductions in a survey with professional artists (Section 7.7).

7.2 Problem Modeling

The parametrization of our drawing tools has to provide durable and manufacturable tools. It also has to allow for achieving a wide range of haptic feedback. We reach these goals by taking inspirations from traditional drawing tools and numerical simulation.

7.2.1 Stylus Parametrization

A drawing stylus design can have many degrees of freedom, which define the shape of the stylus, drawing nib, and the material from which these parts are manufactured. To investigate the importance of different parameters, we used numerical simulation and performed modal analysis on differently shaped styli made of different materials. The analysis revealed a significant influence of ma-

terial on vibrational modes but little influence of the overall shape of the stylus. To decrease the fabrication time and remain consistent with commercial styli we treat the body of the stylus as made from a rigid material. Consequently, we do not consider the entire stylus for our optimization and focus on the nib, which directly interacts with a drawing surface. We assume that the shape of the nib is a cone with a hemispherical tip of radius r. Also, it is made of one homogenous printing material m (Figure 7.3).

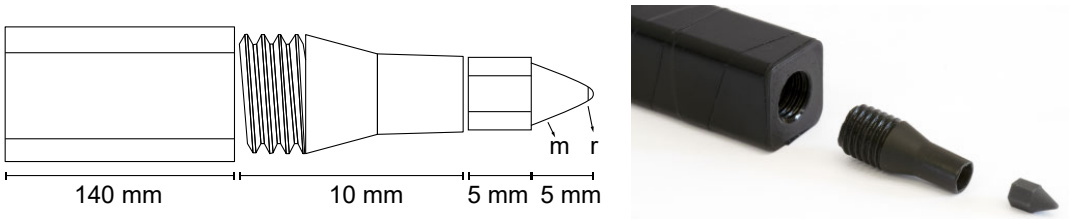

Figure 7.3. The modular design used for our styli (left). The replaceable tip can be made of different material m and have a different radius r. A picture of a manufactured stylus (right).

Fabrication We fabricate our styli using Formlabs Form 2 printer with Black and Flexible resin. To produce durable styli, we limit the radius r of the tip to be within 0.5 mm and 2 mm. We vary the material by mixing the two resins with a continuous ratio of m. Pre-mixing the materials ensured a smooth blend between them and the durability of the print. To minimize the printing time, we also designed the styli as nibs that fit into a universal holder (Figure 7.3).

7.2.2 Surface Parametrization

We inspire the parametrization of the drawing surfaces by first performing imaging of several drawing substrates (office paper, rough paper, and smooth stone paper) using the Gelsight system Yuan et al. [2017]. The scans (Figure 7.4) reveal that the geometry of each substrate is governed by the distribution and thickness of the cellulose fibers. Despite the locally anisotropic structure, fibers create an isotropic substrate on a global scale. This is important at low drawing speeds when the local anisotropy becomes more apparent. Additionally, the surfaces manifest very small variation in height (approx. five microns).

To capture the character of traditional drawing surfaces, we model our surfaces as a heightfield defined by a noise generated using isotropic Gabor kernels Lagae et al. [2009]. Isotropic Gabor kernel behaves similarly to the anisotropic variant but has random orientation at each spatial location, which creates a glob-

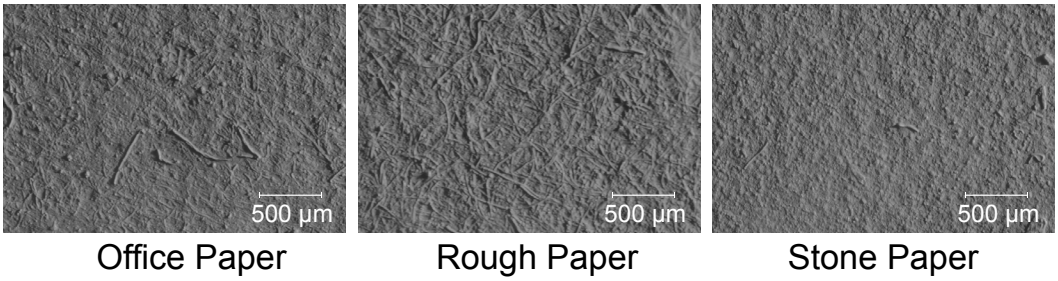

Figure 7.4. Gelsight scans of drawing substrates used in our experiments.

ally isotropic but locally anisotropic surface. We parametrize our surfaces in the power spectra domain where a single isotropic Gabor kernel is represented by a radial Gaussian, thus requires only three parameters: frequency, amplitude, and standard deviation. To model our surfaces, we consider two Gabor filters [Lagae et al., 2011] resulting in six parameters. Since no single fabrication process can cover the entire gamut of drawing tools, we add a categorical parameter to model the manufacturing procedure. To produce our substrates we consider three options: 2D inkjet printing on a transparency sheet, the same process ink-jet where now the printed structure is additionally covered with a fixative spray, and a FormLabs printer with transparent rigid material. Each surface is also scaled and quantized according to the thickness of the printing layer. The scaling is another parameter we included in our optimization.

Fabrication We fabricate our surfaces using a 600 dpi Canon i-SENSYS LBP6780x and Formlabs Form 2 printers. The laser-jet printer uses a plain transparency sheet as the substrate and prints a single ink layer. As an optional post-processing step, we spray the transparency sheet with a fixative spray to increase the friction of the surfaces. We quantize the Gabor noise using the layer height of each manufacturing process. We measured the ink layer to be approximately 8.5 microns thick. We use the Formlabs printer to manufacture surfaces with 25-micron layers of rigid, transparent material. Figure 7.5 shows surfaces generated with our method.

7.3 Efficient Fabrication-In-The-Loop Optimization

Given the parameters of our digital drawing tools, i.e., the radius and the material of the stylus tip, six surface geometry parameters, the manufacturing method of the surface, and the scaling parameter, we seek an efficient optimization which

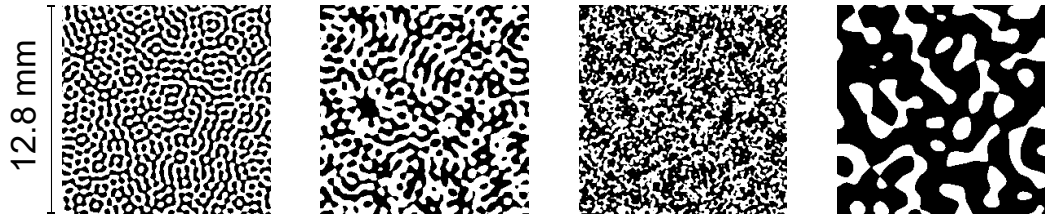

Figure 7.5. Sample surfaces from our parametrization quantized to one layer.

finds the combination of digital stylus and surface replicating desired haptic feedback.

A traditional approach to such a problem is to formulate an optimization procedure guided by numerical simulation. We have experimented with FEM-based frictional contact models by incrementally building an accurate numerical simulator. Unfortunately, as we were adding support for different phenomena to achieve the required accuracy, the simulation became prohibitively expensive to use in an optimization loop. Based on this investigation, we concluded that due to the complexity of the phenomena we need to model as well as imperfections introduced in the fabrication process, the investigated simulations do not provide sufficient efficiency and accuracy. Interestingly, the time required for simulating our designs exceeds the time needed for fabricating and measuring them. Inspired by this observation, we propose to incorporate fabrication proces directly into the optimization loop. This design decision leads to a gradient-free black-box function optimization that cannot be efficiently solved using methods that densely sample the design space, e.g., stochastic optimizations. Instead, we propose a data-driven model based on Gaussian Processes, which allows for efficient sampling of the design space, and therefore, minimizing the number of samples that have to be fabricated and measured before reaching the optimal solution.

7.3.1 Haptic Feedback Similarity

Before formulating the optimization, we need to define what the goal of the optimization is. Haptic feedback is a complex phenomenon, and it is unclear how to evaluate the similarity between different tools. Here, we leverage our work from Chapter 6, where we investigate the perception of haptic feedback induced by drawing tools and discover two primary cues used to distinguish between tool-surface combinations: Coulomb friction between the drawing tool and substrate, and a velocity-dependent spectrogram of vibratory feedback. The two perceived quantities define a perceptual space where the similarity between tool corre-

sponds with Euclidean distance. In order to objectively verify the quality of our reproductions, we extend this perceptual space by recovering the Just Noticeable Difference (JND) units Fechner [1860]. We define 1 JND as a difference between two stimuli which can be detected with a 75% chance. To compute the JND units we analyze the free drawing study Chapter 6.5.1. During the study, participants were presented with a mixture of traditional and 3D printed drawing tools. The experiment design was a two-alternative forced-choice experiment, where one tool serves as a reference and two tools as possible reproductions. The task was to pick the reproduction that is more similar to the reference. The outcome of the study is the probability of preferring one tool over another. We utilize these probabilities and optimize for two scaling factors for friction and vibration respectively such that the resulting JND conforms to the conducted psychophysical studies.

7.3.2 Gaussian Process Surrogate Model

The input to our optimization is a characterization of haptic feedback based on a perceptual space. The output is a design from our parametrization which can be directly manufactured. To evaluate the perceived difference between our designs and the target haptic behavior we rely on minimizing the Euclidean distance in a perceptual space of haptic feedback which requires physically manufacturing and measuring the sample. Since the manufacturing process is expensive (both time and cost-wise) and derivative-free we propose to utilize a data-driven surrogate model. The input to the model is a design from our parametrization and previously observed designs. The output is predicted perceived haptic feedback based on the data evidence. The model is based on two building blocks. The first block is an efficient approximation of the objective function which can be evaluated quickly and provides confidence bounds on the prediction. The confidence bounds are a key element to build the second block which is an acquisition function that searches through the input design space. By incorporating the certainty of prediction we can automatically handle the exploration-exploitation trade-off during optimization.

To incorporate confidence in the prediction we first have to assume the uncertainty of the data. To this end, we model the mapping from design space to measurements with uncertainty explained by a Gaussian distribution. Assuming normality of the distribution is key in formulating an efficient analytical solution to predict and sample new designs. Under such an assumption, each predictor is considered to be a random variable with multi-variate Gaussian distribution. The Gaussian Process is then the joint distribution of the observed variables on

an infinite continuous domain. A Gaussian Process is itself a normal distribution and can be parametrized with a mean and a covariance matrix. The infinite domain of the Gaussian Process allows us to express predictions of unobserved designs. To achieve this we formulate the covariance matrix as a kernel on a continuous domain. The kernel function defines how strongly is a design correlated with our previous observations. The correlation typically smoothly varies across the domain. By aggregating these correlations for all observed data we can predict the location of a design point as a normal distribution giving us both the predicted value and confidence intervals on the prediction.

In our setting the problem is further complicated by introducing a categorical variable that models the fabrication process. The smoothness assumptions imposed on the kernel function require special treatment of categorical variables. Therefore, we seek two predictors on perceived haptic feedback: one based on surface parametrization κ_1 and one based on the manufacturing process κ_2. By multiplying these two predictors we create a predictor of haptic feedback κ that smoothly varies across all dimensions. The predictor κ is calculated as:

$$\kappa(\mathbf{x}, \mathbf{x}') = \kappa_1(\mathbf{x}, \mathbf{x}') * \kappa_2(\mathbf{x}, \mathbf{x}'), \tag{7.1}$$

where $\mathbf{x}$ is a design point we wish to predict, and $\mathbf{x}'$ is an observed data point. The data points $\mathbf{x} = [\mathbf{t}, \mathbf{s}, k]$, where $\mathbf{t}$ is a vector of stylus parameters, $\mathbf{s}$ is a vector of surface parameters, and k is the manufacturing process.

Predicting Stylus-Surface Haptic Feedback

The Gaussian Process takes as an input our stylus-surface parametrizations and predicts the perceived haptic feedback. This prediction is guided by the covariance kernel which imposes smoothness assumptions on our predictor. There is a range of kernel functions that were developed based on different applications (Exponential kernel, Matern 3/2 kernel, Matern 5/3 kernel, and their Automatic Relevance Determination (ARD) variants [Rasmussen and Williams, 2005]). To identify an appropriate kernel function we designed a test case. We used the measurements collected in Chapter 6.2 for 3D printed tools on paper substrates. Then we projected the substrates into our parameterization. We use cross-validation to optimize for a model that best explains the observed data. Based on the cross-validation results we select the ARD Matern 3/2 kernel:

$$\kappa_1(\mathbf{x}, \mathbf{x}') = \sigma_f^2(1 + \sqrt{3}r_\kappa)e^{-\sqrt{3}r_\kappa}, \tag{7.2}$$

$$r_\kappa = \sqrt{\sum_{i=1}^{D} \frac{(x_j - x_j')^2}{\sigma_i^2}}, \tag{7.3}$$

where, $D = 9$ is the dimensionality of the predictor. σ_f is the signal standard deviation of our data which defines how certain we are about observed predictions, and σ_i is the characteristic length scale different for each dimension. The characteristic length scale defines at which range values are correlated with observed measurements and is set by fitting the Gaussian Process into observed data using a L-BFGS optimization [Nocedal and Wright, 2006]. Larger values lead to smoother functions and smaller to functions with faster variation.

Transfer Learning Between Fabrication Processes

The employed parametrization is not bound to a particular manufacturing process. Introducing new manufacturing processes will have an effect on the perceived haptic feedback. On the one hand, each process will produce different haptic feedback for the same parameters. On the other hand, it is reasonable to believe that trends in feedback change would remain similar across fabrication processes. Therefore, we would like to encode categorical variables in a way that exploits observed data when few observations of the category were made but once available favors predictions made using the data gathered from the category.

To model this behavior we first encode the categorical variables using one-hot encoding [Garrido-Merchán and Hernández-Lobato, 2020]. Next, we use an ARD kernel [Neal, 1996]:

$$\kappa_2(x, x') = \exp\left(-\frac{1}{2}\sum_{j=1}^{K}\frac{1}{\sigma_j^2}(x_j - x_j')^2\right), \tag{7.4}$$

where κ_2 is the kernel function, K is the number of categories, and σ_j is the characteristic length parameter. By optimizing for the length parameter σ_j we can tweak the effect of individual categories. We can see the effect of multiple categories on an example (Figure 7.6). For cases where data about a category are sufficient, we can use $\sigma_j \to \infty$ leading to effectively no transfer of knowledge. On the other hand, if we do not have sufficient data about a category a lower sigma will appropriately scale the predictions of other categories. One disadvantage of our approach is that for unobserved categorical parameters we can not predict their expected behavior. Instead, to include a new category i.e., a new fabrication process we initialize the category with a single measurement which will be used to find a first estimate of the mixing ratios.

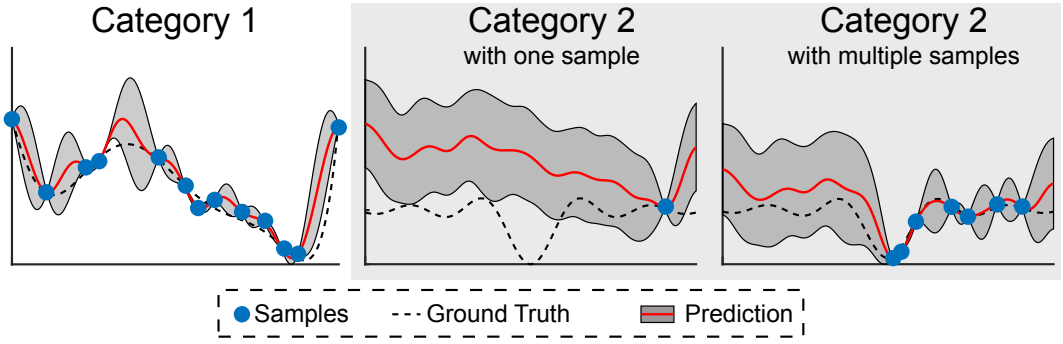

Figure 7.6. A one-dimensional example of information transfer between multiple categories. The estimate of category 1 (left) is used to improve prediction in uncertain regions of category 2 (middle). With more samples the prediction is weighted towards observations of the new category (right).

Friction and Vibration Prediction

In our setting, we are interested in jointly predicting the frictional and vibrational response of a tool and surface combination. In the context of Gaussian Processes the joint prediction can be formulated by stacking the Gaussian Processes explaining individual responses into a single multivariate-normal distribution $\mathcal{N}$ which can be written as:

$$\begin{bmatrix} F \\ V \end{bmatrix} = \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{K}_F & a \\ a & \mathbf{K}_V \end{bmatrix} + \begin{bmatrix} \sigma_F^2\mathbf{I} & 0 \\ 0 & \sigma_V^2\mathbf{I} \end{bmatrix}\right), \tag{7.5}$$

where F, V are predicted friction and vibration respectively, $\mathbf{K}_F$, $\mathbf{K}_V$ are their corresponding correlation matrices, and σ_F, σ_V is the noise of prediction in each Gaussian Process, and a defines the correlation matrix between the two predicted values. The kernel function for a has to be carefully selected. Poor selection imposes unwanted structure on the problem which leads to poor performance [Bonilla et al., 2008]. To select the appropriate kernel a we can use an interesting property of the formulation. If at each parameter value we have noise-free observations for both predicted values then the shared information between the predicted values vanishes. The intuition behind this observation is that by first decorrelating the response variables by a we can estimate the relationship as the covariance of the decorrelated data. For detailed derivation please see [Bonilla et al., 2008].

We assume that the bulk of the cost is in fabrication (stylus and surface). Measurements for friction and vibration are both fast and it is feasible to recover both functional values for each predictor. Therefore, to gain useful information

from cross-correlation we would need significant noise in the measurements. To evaluate the experimental noise we repeated measurements for a ballpoint pen. We found a 0.3 JND standard deviation on the placement of the tool in the perceptual space which is too small to be noticeable. Based on the findings we conclude that the modeling can be done separately for friction and vibration. This means we set $a = 0$ and we have two separate Gaussian Processes: one predicts perceived friction, the other perceived vibration.

7.3.3 Acquisition Function

Efficient sampling of the design space is the most crucial factor in quickly optimizing an expensive function. To this end, at each manufacturing step, we would like to produce the sample that either brings us closest to target or provides the most information to improve our surrogate model. The Gaussian Process surrogate model enables us to predict the expected behavior of a tool-surface combination with confidence bounds. We can exploit these confidence bounds to formulate an acquisition function of new measurements that maximize the expected improvement of our model towards a target behavior.

The input of the acquisition function is the desired haptic behavior. The output of the function is a new stylus-surface design which maximizes the expected improvement towards the target behavior. Unlike traditional minimization problems, we are interested in finding a specific value of a two-dimensional function, (Equation 7.5), rather than its minima. One option would be to directly estimate the distance to our target with the surrogate model. However, the probability as a function of distance to the target can not be modeled by Gaussian Processes. To address this issue we reformulate the improvement as a two-dimensional distance minimizing function. At each sample point we seek to estimate the function:

$$u(\mathbf{x}) = \max(0, d^* - (\mathcal{N}_F(\mathbf{x}) - F)^2 - (\mathcal{N}_V(\mathbf{x}) - V)^2), \tag{7.6}$$

where d^* is the square distance between target and current best estimate, $\mathcal{N}_F(\mathbf{x})$ and $\mathcal{N}_V(\mathbf{x})$ are the estimated friction and vibrational behavior respectively and, F and V is the target friction and vibration respectively.

The problem is visualized in Figure 7.7 where we wish to optimize for target friction and vibration. The current best estimate P^* defines a circle around the target T with radius d^*. Designs within this region have expected improvement larger than zero. Since our predictor variables are defined as Gaussian distributions we can not simply evaluate the value of the improvement function. Instead, we have to estimate its expected value which can be formulated as a double

integral over the design space weighting functional values by their probability of occurring. As an illustrative example to evaluate the expected improvement of design with fixed vibration v we predict the friction as a Gaussian distribution (Figure 7.7 gray). The expected improvement then includes the frictional values (Figure 7.7 orange) located within the integration bounds multiplied by their probability (Figure 7.7 blue). In order to formulate an analytical version of the integral, we need to set tight integration bounds which allow us to implicitly handle the max operator. A definite double integral integrates over a rectangular subspace (Figure 7.7 dashed). We set this subspace to be the inscribed square of the circle which defines the area of possible improvement by setting $d = \sqrt{\frac{d^*}{2}}$. The integral then has the following form:

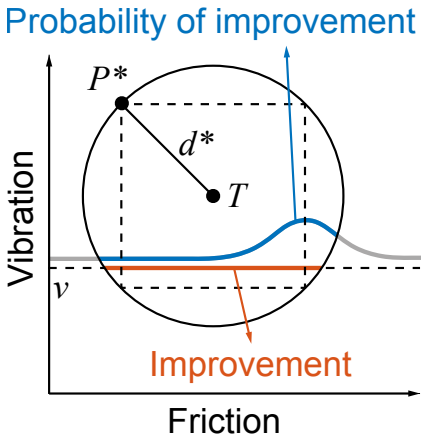

Figure 7.7. Visualization of acquisition function.

$$\int_{V-d}^{V+d} \int_{F-d}^{F+d} (d^* - (x-F)^2 - (y-V)^2) p_x p_y \, \mathrm{d}x \, \mathrm{d}y, \tag{7.7}$$

where p_x and p_y is probability of predicting frictional and vibrational value respectively. To find an analytical solution we first express the probability as a Gaussian distribution:

$$\frac{1}{\sqrt{2\pi\sigma_F^2}\sqrt{2\pi\sigma_V^2}} \int_{V-d}^{V+d} \int_{F-d}^{F+d} (d^* - (x-F)^2 - (y-V)^2) e^{-\frac{(x-\mu_F)^2}{2\sigma_F^2}} e^{-\frac{(y-\mu_V)^2}{2\sigma_V^2}} \, \mathrm{d}x \, \mathrm{d}y.$$

We start by grouping the constant terms while treating y as a constant:

$$s \int_{V-d}^{V+d} e^{-\frac{(y-\mu_V)^2}{2\sigma_V^2}} \int_{F-d}^{F+d} (D - (x-F)^2) e^{-\frac{(x-\mu_F)^2}{2\sigma_F^2}} \, \mathrm{d}x \, \mathrm{d}y.$$

Next, we solve the inner integral:

$$s \int_{V-d}^{V+d} \frac{\sigma_F}{2} \left[-\sqrt{2\pi} \left((y-V)^2 - d^* + (F-\mu_F)^2 + \sigma_F \right) \left(\operatorname{erf}\left(\frac{d+F-\mu_F}{\sqrt{2}\sigma_F} \right) + \operatorname{erf}\left(\frac{d-F+\mu_F}{\sqrt{2}\sigma_F} \right) \right) + 2\sigma_F \left((d+F-\mu_F) e^{-\frac{(d-F+\mu_F)^2}{2\sigma_F^2}} + (d-F+\mu_F) e^{-\frac{(d+F-\mu_F)^2}{2\sigma_F^2}} \right) \right] e^{-\frac{(y-\mu_V)^2}{2\sigma_V^2}} \, \mathrm{d}y.$$

We again substitute constant terms:

$$s\frac{\sigma_F}{2}\int_{V-d}^{V+d}\left(-\sqrt{2\pi}l\left((y-V)^2+m\right)+k\right)e^{-\frac{(y-\mu_V)^2}{2\sigma_V^2}}\,\mathrm{d}y,$$

where $m=-d^*+(F-\mu_F)^2+\sigma_F^2$, $l=-\sqrt{2\pi}\left(\operatorname{erf}\left(\frac{d+F-\mu_F}{\sqrt{2}\sigma_F}\right)+\operatorname{erf}\left(\frac{d-F+\mu_F}{\sqrt{2}\sigma_F}\right)\right)$, and $k=2\sigma_F\left((d+F-\mu_F)e^{-\frac{(d-F+\mu_F)^2}{2\sigma_F^2}}+(d-F+\mu_F)e^{-\frac{(d+F-\mu_F)^2}{2\sigma_F^2}}\right)$.
And finally solve for the second integral:

$$\frac{\sigma_F\sigma_V}{4\sqrt{2\pi\sigma_F^2}\sqrt{2\pi\sigma_V^2}}\left[\sqrt{2\pi}\left(k+l\left(m+\sigma_V^2+(V-\mu_V)^2\right)\left(\operatorname{erf}\left(\frac{d-\mu_V+V}{\sqrt{2}\sigma_V}\right)+\operatorname{erf}\left(\frac{d+\mu_V-V}{\sqrt{2}\sigma_V}\right)\right)\right)-\right.$$
$$\left.2l\sigma_V\left((d-\mu_V+V)e^{-\frac{(d+\mu_V-V)^2}{2\sigma_V^2}}+(d+\mu_V-V)e^{-\frac{(d-\mu_V+V)^2}{\sqrt{2}\sigma_V^2}}\right)\right].$$

Practical Consideration of Varying Fabrication Time

An interesting property of manufacturing processes is that modification of different parameters can lead to a non-constant change in manufacturing time, e.g., time difference between modifying the geometry of a single layer and printing at a higher resolution. In such cases, a small change in parameters can lead to a significant increase in time while providing minimal gain. It is possible to include the effects of different time costs for parameter change in the acquisition function by weighting the expected improvement by the manufacturing time. However, in case of extreme time differences, the weight would get unbalanced and the acquisition function would over-explore the parameters with lower time costs. The unbalanced time factors could be solved by adding a constant to each manufacturing process. However, the constant depends on the specific setup and has to be manually tuned.

Instead, we propose to utilize another property of the fabrication: the ability to fabricate designs in parallel. Every time one of our manufacturing devices is free we generate a new design with the current state of the surrogate model. When one fabrication process is significantly faster it is naturally explored more. The fast parallel updates to the surrogate model can be utilized to more efficiently use the slower manufacturing process. It is possible that after a number of fast iterations we find a new design with better performance than the one being manufactured. In such a case, we can use the expected improvement over time to evaluate the payoff of stopping the current print and starting a new one. The ability to stop prints that no longer provide any value before they finish enables us to more efficiently incorporate different fabrication procedures into our design.

7.4 Application To Stylus-Surface Design

We use our surrogate model to explore the design space of stylus-surface interaction by manufacturing. Due to differences in our fabrication processes a stylus design takes significantly longer to manufacture than a new surface. We incorporate this knowledge into the design of our optimization loop by optimizing for new styli and new surfaces in parallel, (Pseudocode 1). To perform one iteration we start by fitting two Gaussian Processes: one for frictional response and one for the vibrational response of our designs. Next, we search for a new tool to print given our set of available substrates $\mathscr{S}$ by maximizing the expected improvement defined in Equation 7.6:

$$\max_{\mathbf{t},i} u(\mathbf{t}, \mathbf{s}_i), \mathbf{s}_i \in \mathscr{S}, \tag{7.8}$$

where $\mathbf{t}$ is the parametrization of the stylus, and $\mathbf{s}_i$ is a surface from the surface set $\mathscr{S}$. If the time-weighted expected improvement is larger than the currently printing design we start to print the new design.

Next, we seek a new surface given our set of available styli $\mathscr{T}$. We search for the surface by maximizing:

$$\max_{\mathbf{s},i} u(\mathbf{t}_i, \mathbf{s}), \mathbf{t}_i \in \mathscr{T}, \tag{7.9}$$

where $\mathbf{s}$ is the parametrization of a the surface, and $\mathbf{t}_i$ is a tool from tool set $\mathscr{T}$. To maximize the expected improvement we use a genetic algorithm [Conn et al., 1991]. These two optimization steps are repeated until we reach an acceptable reproduction with a JND bellow 0.3. We can see the behavior of our optimization on an example in Figure 7.8.

7.4.1 Reliability Test

Our data-driven surrogate model needs a warm start with a set of measurements. To generate the initial sampling we considered six initial styli and surface designs resulting in a total of 36 initial measurements. As manufacturing process we use the laser-jet printer with a transparency substrate. We pick styli parameters that uniformly sample the design space with three radii (0.5, 1, 2 mm) and two materials (black, and flexible), Figure 7.9 left. The initial surface designs were generated automatically using random sampling, Figure 7.9 right. We measured each initial design and found their appropriate placement in the perceptual space of haptic feedback, (Figure 7.10 left). We can observe that our initial sampling provides good coverage of tools with lower perceived friction.

Algorithm 1: Perceptual Optimization Pseudocode

```
Data: T = target behavior
Data: o = observations
while true do
    M_F(x) ← gaussian_process(friction(o));
    M_V(x) ← gaussian_process(vibration(o));
    t = pick new tool by solving Equation 7.8;
    if not printing or tool has better improvement then
        print new tool;
    s = pick new surface by solving Equation 7.9;
    m = measure new tool;
    o add [m,s] pair to observations dataset;
    if ||T − Perceptual(m)|| < 0.3 then
        done;
```

To verify the effect of initial sampling on our optimization we removed the initial surface which provided haptic response close in the perceptual space to a ballpoint pen. We selected this design because it lies within the gamut of traditional drawing tools and is relatively isolated. This isolation requires the optimization process to explore the design space and not rely on already observed samples. The optimization process required 9 iterations to reach a satisfactory close reproduction (Figure 7.10 right). Please note that the optimization includes exploration of the design space and does not behave like gradient descent. An interesting observation is that the new optimized surface is very similar to the removed design.

7.4.2 Optimizing Haptic Feedback

To demonstrate the capability of our method to efficiently model the behavior of drawing tools we optimize stylus-surface replicas of four traditional tools from the perceptual space: 2H pencil on stone paper, 8B pencil on stone paper, ballpoint pen on rough paper, and charcoal on rough paper. We opted for these tools due to their variety of haptic responses which covers well the gamut of haptic feedback provided by traditional tools.

Our initial sampling already provides a faithful reproduction of the 2H pencil, (Figure 7.11 2H pencil). We use our surrogate model to optimize for the haptic feedback of an 8B pencil on stone paper, (Figure 7.11 8B pencil). The optimiza-

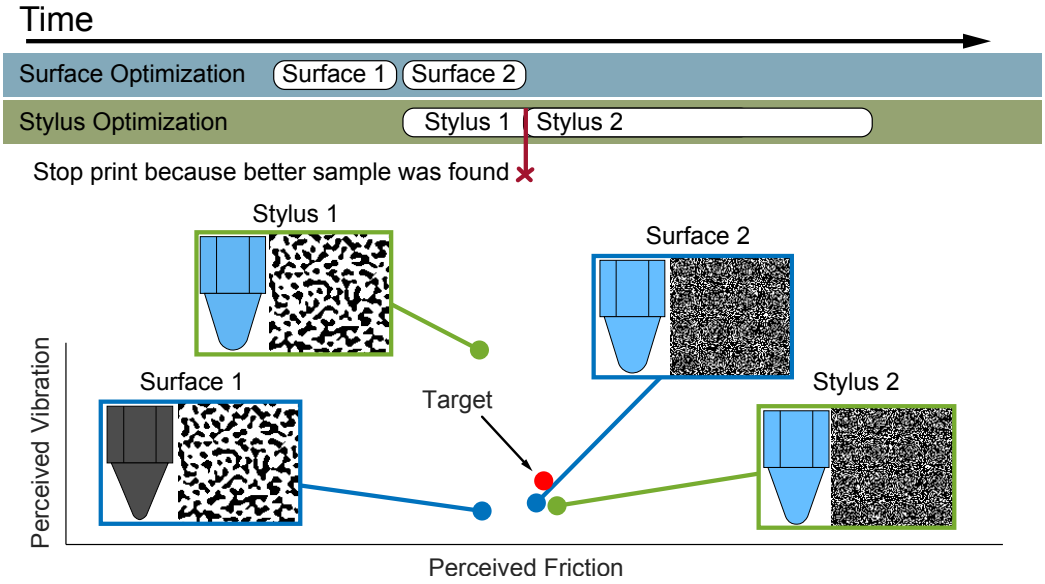

Figure 7.8. Example of parallel execution of stylus and surface manufacturing. In the first iteration, we manufacture a surface and enhance our data-driven model. The model then predicts which tool we should print to best match the target. We prepare the print for the tool and while the tool is printing we optimize for a new surface. The new surface causes an update to the surrogate model which results in a new suggestion for a stylus to print. Since the printer is currently occupied we compare the tools based on their expected improvement over time. We find that the expected improvement of the new tool outweighs the difference in printing time. Therefore, we stop the current print and prepare the new tool. In order to verify that the model prediction was correct, we manufactured and measured both styli. We can observe that the newly predicted stylus indeed achieved better improvement.

tion converges within one iteration to a satisfactory solution. Next, we optimize for the ballpoint pen on a rough surface. The surrogate suggests printing a new tool with a radius of 0.5 millimeters composed of 14% black material and 86% flexible material. During printing time the optimizer was suggesting five surface designs that could not achieve satisfactory vibrational feedback. Once the new tool was printed we measured it using the original surface for which it was suggested. Afterward, it took the optimizer 4 steps to converge to a satisfactory solution, (Figure 7.11 ballpoint pen). Finally, we optimize for the haptic feedback produced by charcoal on rough paper. Since the friction of charcoal is significantly higher than what is achievable by the ink-jet printer we change

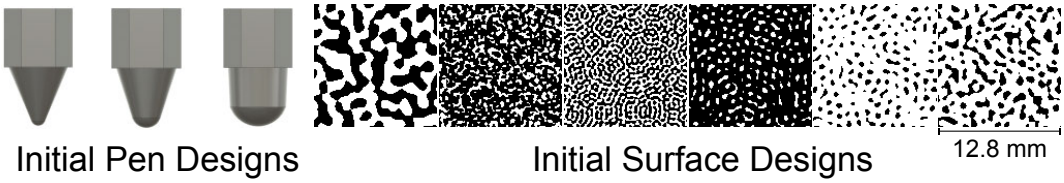

Figure 7.9. Pen and surface designs used to initialize the surrogate model.

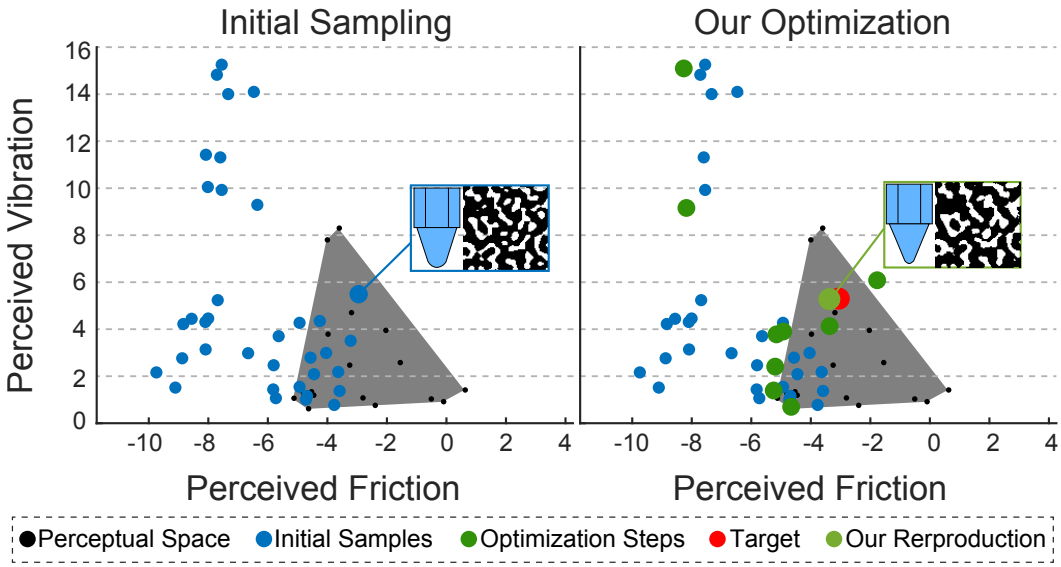

Figure 7.10. Initial samples measurements are shown on the left. We can observe that initial sampling can capture a ballpoint pen on office paper. We verify our algorithm by removing the surface from initial sampling and optimize for the same tool, right.

the fabrication method to a 3D printer. We initialize the new fabrication process by manufacturing and measuring a design which was closest to charcoal from original sampling. Afterward, we use the optimizer and find a satisfactory reproduction in 3 steps, (Figure 7.11 charcoal). The resulting optimized stylus-surface combinations are shown in Figure 7.12.

The timing of one iteration loop is a combination of predicting a new sample, manufacturing the design, and measuring. The prediction of new designs takes an average of 2 minutes. Manufacturing depends on the selected method fabrication method. For surface fabrication, we use an inkjet printer, fixative, and a 3D printer. The inkjet printer is the fastest and takes about 60 seconds. The fixative agent is applied on the inkjet-printed surface and takes about 15 minutes to dry, and lastly 3D printing a new surface takes about 3 hours. A similar fab-

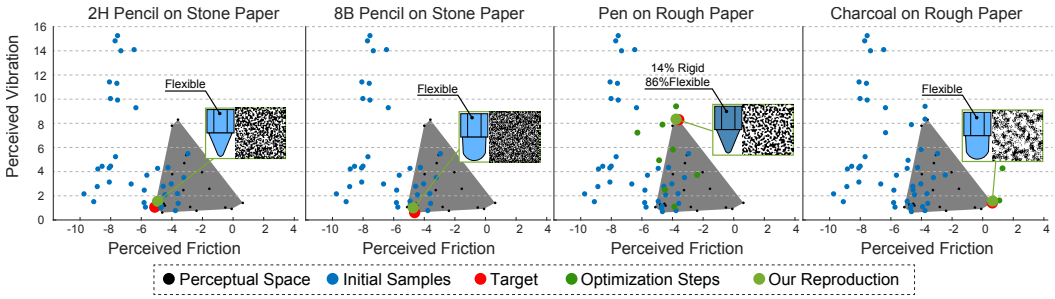

Figure 7.11. Results of optimizing haptic feedback using our surrogate model. Black samples mark original perceptual space. Blue samples are sample surfaces used to initialize the model. The red color is the position of a tool we would like to reproduce. Dark green dots mark the path our optimizer took in the perceptual space to find the final reproduction (light green).

rication time is required to 3D print a new pen design. Lastly, the measurement itself takes about 15 minutes.

7.4.3 Gamut of Haptic Feedback

Estimating the gamut of haptic feedback achievable through different fabrication processes is advantageous during design and allows us to quantify the expressiveness of our design space. To quantify the gamut we cannot rely on a parameter sweep. The non-linearity and high dimensionality of our space would require a far too great number of samples to be practical. Instead, we use our surrogate model to acquire an estimate of the gamut. To contrast our digital styli with traditional instruments we use the gamut recovered by [Piovarči et al., 2018] (Figure 7.13 black) that is composed of various drawing tools and substrates.

By searching for the convex hull of haptic feedback of traditional drawing instruments we recovere a good estimate of haptic feedback achievable by our laser-jet printer with a transparency substrate. For visualization we marked the gamut as a convex hull, (Figure 7.13 blue). To estimate the haptic feedback achievable by a new process we took the surface-stylus combinations forming the convex hull of our laser-jet process and manufactured them using the fixative and 3D printer. Due to frictional and vibrational coupling, the mapping of haptic feedback between surfaces is not a simple transformation of the space. To further explore the gamut of haptic feedback of the new processes we, therefore, use our optimization. We expand the gamut of the new processes by optimizing for traditional drawing tools outside of the approximate convex hull. We can see the optimized gamuts of haptic feedback in Figure 7.13. Each manufacturing

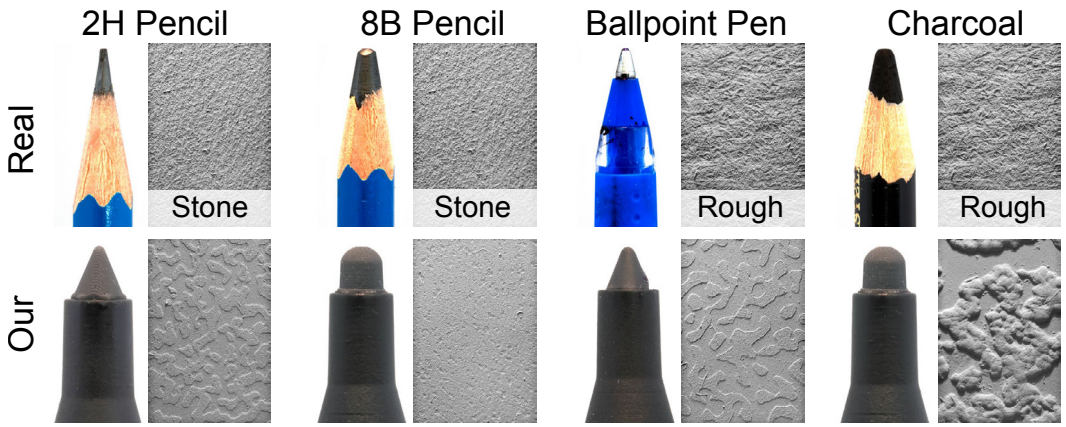

Figure 7.12. Four traditional drawing tools (top) and their optimized digital counterparts (bottom). We show both the optimized stylus and the Gelsight scan of the optimized surface. Please note that our surfaces do not match the original substrates yet manifest the same haptic response.

process is capable of achieving satisfactory vibrational feedback. In terms of frictional feedback, the best performing process is a transparency sheet covered with a fixative agent that covers almost the entire gamut of traditional drawing tools.

7.4.4 Universal Drawing Surface

We demonstrated that our algorithm can generate stylus-surface combinations designed to reproduce a specific drawing instrument. For practical purposes, it is typically easier to manufacture a device with a single surface and use multiple nibs for the stylus to customize the haptic feedback. We use our surrogate model to optimize for a surface pattern that can capture the largest coverage of the perceptual space of drawing tools. Our fabrication method of choice is the transparency sprayed with fixative as its gamut encompasses the entire perceptual space.

The input to our function is an area of haptic feedback we would wish to reproduce. The output is a surface design that can capture the largest cross-section of the input area given our toolset $\mathcal{T}$. Since we seek to find a surface that manifests different haptic feedback for many styli we can not rely on the acquisition function of our surrogate model. Instead, we use the surrogate model to estimate the expected haptic feedback achievable on the surface as:

$$E[\mathbf{s}] = ConvexHull(\{\mathcal{N}_{F_i}(\mathbf{s}), \mathcal{N}_{V_i}(\mathbf{s})\}), \tag{7.10}$$

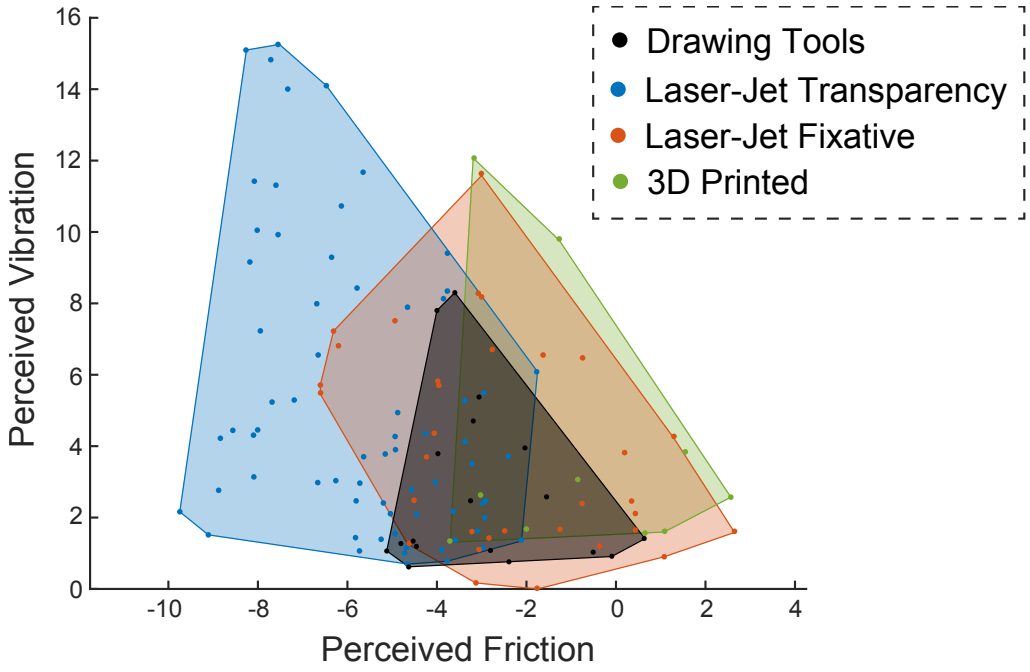

Figure 7.13. The gamut of haptic feedback achievable by our different manufacturing processes: laser-jet on transparency sheet (blue), laser-jet on transparency sheet with fixative (red), 3D printer (green) compared with traditional drawing tools (black).

where $\mathbf{s}$ is the parametrization of the surface, $\mathcal{N}_{F_i}(\mathbf{s})$ and $\mathcal{N}_{V_i}(\mathbf{s})$ are the Gaussian processes predicting friction and vibration of tool i from the toolset $\mathcal{T}$ on the surface $\mathbf{s}$ respectively. $ConvexHull$ estimates the convex hull of haptic feedback achievable by the surface. Based on experiments performed by [Piovarči et al., 2018] we assume that the haptic feedback between individual styli can be interpolated linearly. Since we would like to produce realistic haptic sensations we set the input area to the convex hull of haptic provided by traditional drawing tools. To find a universal drawing surface we maximize:

$$\max_{\mathbf{s}} Intersect(input, ConvexHull(\{\mathcal{N}_{F_i}(\mathbf{s}), \mathcal{N}_{V_i}(\mathbf{s})\})). \tag{7.11}$$

The function contains an estimate of the expected intersection area that we calculate using Monte-Carlo evaluation. The functional is then maximized using LBFGS method Nocedal and Wright [2006].

We can see the results of the optimization in Figure 7.14. We stop the optimization when the next sample has both predicted and measured coverage within 0.1% of the previous iteration. The algorithm converged within 4 iterations. The

main difference between iterations 3 and 4 is in the type of haptic feedback provided. The surface from iteration 3 provides less vibratory feedback than the surface from iteration 4. Additionally, there are two interesting observations about iteration 4. The surfaces optimized by our method are always slanted. This is a result of using various materials. On one hand, hard plastic has the tendency to produce more vibration but also glides smoother on the surface. On the other hand, rubbery material damps the vibratory feedback and creates more friction with the surface. The second interesting observation is the final surface which contains larger spacing between features. The large spacing directly interacts with tip size. A small tip takes longer to hit a new feature and therefore has less vibratory feedback. Conversely, a large tip is hitting the surface features more often and create an illusion of a rougher surface.

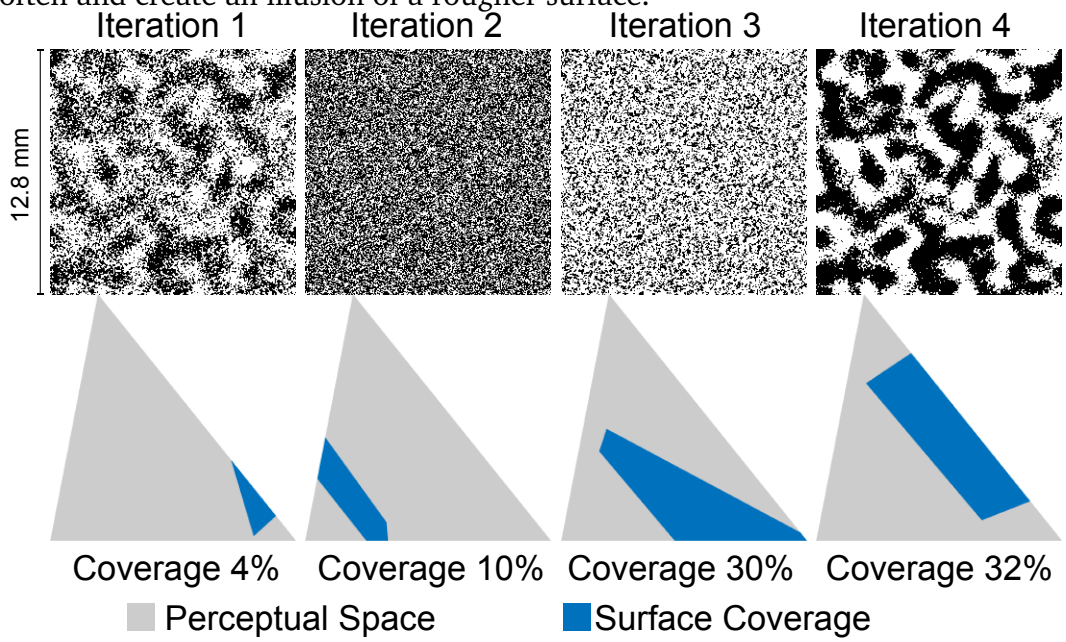

Figure 7.14. Four iteration of optimal surface optimization.

7.5 Comparison and Evaluation

To evaluate the quality of our optimization we perform a series of tests. First, we compare our method with naive surface reproduction and show that special care has to be taken to optimize the stylus-surface haptic feedback. Next, we evaluate the quality of our reproductions by comparing raw measurements. Finally, we investigate the optical properties of our surfaces and their suitability for direct application on tablet screens.

7.5.1 Comparison With Geometry-Inspired Baseline

We designed an algorithm that can generate stylus-surface combinations with a desired haptic response. It is possible that similar haptic feedback of traditional tools could be also achieved by reproducing the geometry of the paper substrate. To verify this approach we approximate a manufacturing pipeline by scanning the rough paper sample with the Gelsight system. The sample showed a peak to peak variation of 10 microns, (Figure 7.15 left). Such extreme resolution is well beyond the capabilities of our manufacturing hardware [Sitthi-Amorn et al., 2015]. However, it is still possible to modify the surface to make it printable by scaling the height-field. To not alter the frequency of the surface features we scale only along the Z-axis. Next, from the measured height-field, we generated a tillable surface [Embark Studios, 2019]. To validate the quality of the reproduction we manufactured the surface on transparency and 3D printing substrate at different scaling factors, (Figure 7.15 right, red and blue respectively). We then measured the perceived haptic feedback on the substrate by using a ballpoint pen. We plot the different measurements as paths signifying the effects of scaling on haptic feedback. We can observe that both reproductions generate different paths in the perceptual space. As we scale up the transparency film we first observe an increase in vibratory feedback which is later followed by a decline. This is a result of the transparency sheet saturating with the ink. On the other hand, the 3D printed surface can scale arbitrarily high so the haptic feedback keeps increasing. The minor dip observed is caused by the first layer saturating.

7.5.2 Measurements of Tool-Surface Combinations

The perceptual space uses an aggregate value to represent the measurement of vibration. This step was motivated by Piovarči et al. due to the overall velocity-independent broadband response of traditional drawing instruments. Our 3D printed designs could potentially violate this assumption. If so this could result in a perceivable difference due to a clear frequency shift when compared to a traditional instrument. We show the original and reproduced spectrograms in Figure 7.16. We can observe that the 3D printed surfaces achieve very similar vibratory response as the original instruments.

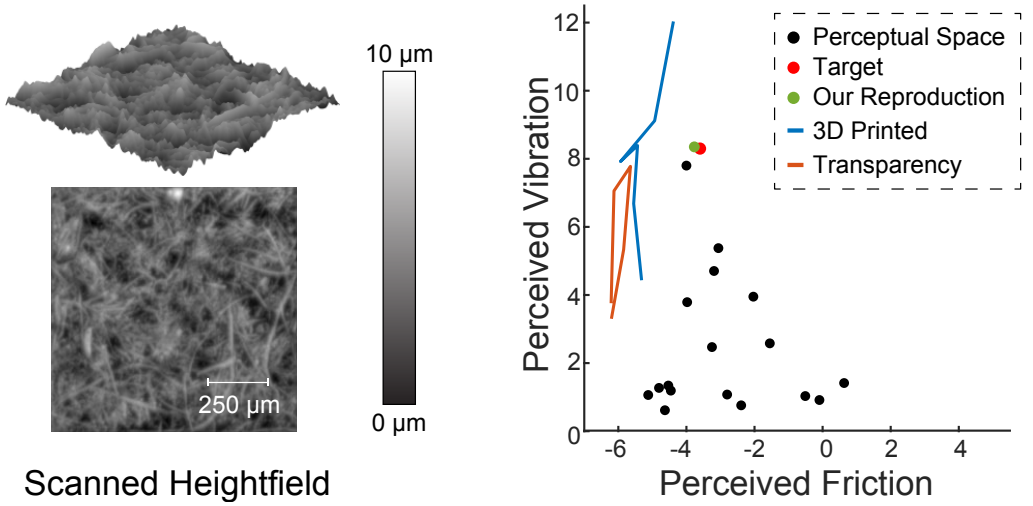

Figure 7.15. Results of surface reproduction of a ballpoint pen. The scanned height-field (left) and measured haptic feedback (right) form paths in perceptual space parametrized by the scaling factor.

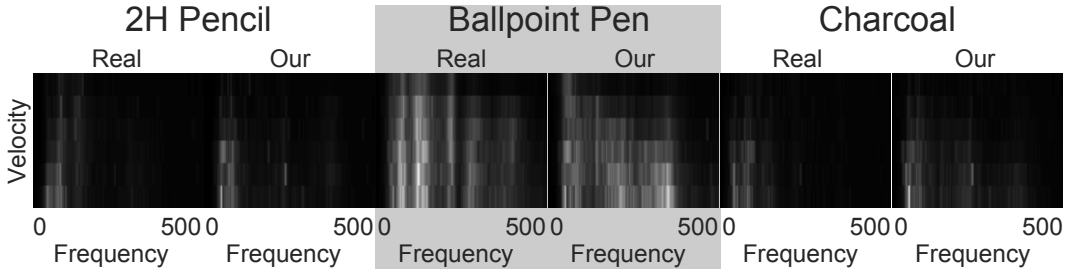

Figure 7.16. Vibration measurements of original drawing instruments and their 3D printed counterparts.

7.5.3 Optical Properties of Generated Patterns

The optical properties of the generated pattern are important for applications on graphical tablets with displays. To verify the image quality of our screens we manufactured our optimized charcoal reproduction design using a clear resin. We then placed the surface on an LCD screen with a checkerboard pattern (Figure 7.17). We can see that the pattern has similar properties to a diffuse screen protector. The surface geometry creates more significant distortion which is a trade-off for achieving desired haptic properties.

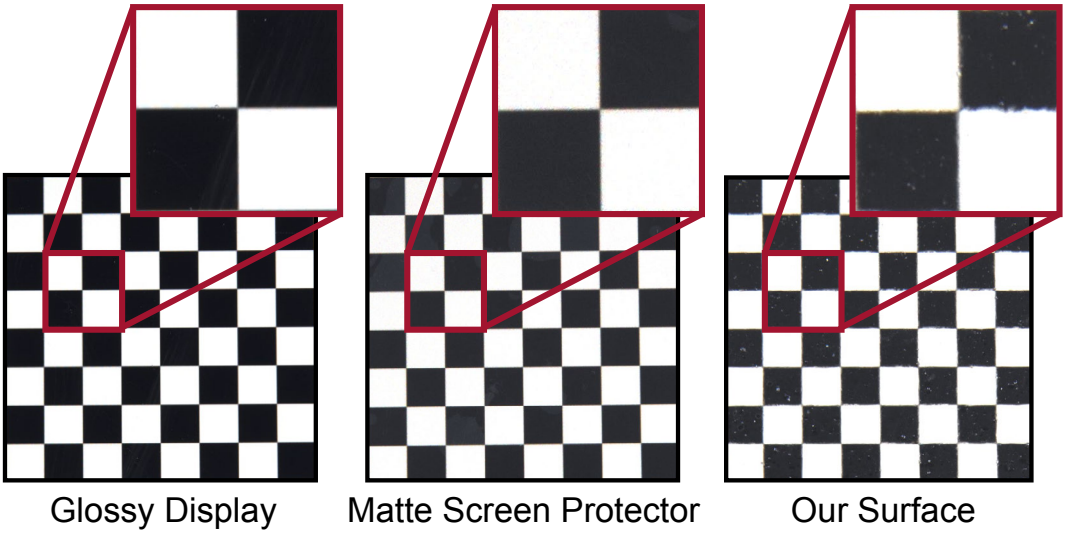

Figure 7.17. Visual properties of our surface evaluated by taking photos of an LCD screen displaying a checkerboard pattern.

7.6 Experimental Validation

Our method heavily relies on the applicability of the perceptual space to solve a non-trivial optimization problem. Such an approach requires careful verification in user experiments. To this end, we start by evaluating peoples' ability to recognize different traditional tools. Next, we compare our tools to current commercial and research solution as well as a baseline solution which attempts replicated material and geometry for drawing tools. We also compare our designs to their traditional counterparts to verify the accuracy and realism of our haptic feedback. Finally, we describe a survey with professional artists.

7.6.1 Study 1: Acquiring Vocabulary

The goal of the first experiment is to make sure all subjects are sufficiently familiar with traditional drawing tools, can describe differences between the tools, and that they are comfortable with our experimental setup. This experiment serves as a base of all our studies which are performed under the same conditions.

Task and Stimuli The participants were sat in front of a cloth, which eliminates visual cues and asked to wear noise-canceling headphones to eliminate auditory cues. To eliminate the tool-shape bias we place each drawing tool in a unified holder based on Copic markers. At each trial, the participants were given one

drawing tool (material and a surface) and asked to draw with it using a simple back and forth motion resembling shape filling (inset Figure) with the task of describing the perceived haptic feedback verbally. After drawing with each tool, they were presented with pairs of tools and asked to describe perceived differences. The study was concluded when the participants felt confident in describing both the feedback and the differences. For this experiment, we selected three tools providing substantially different haptic feedback: ballpoint pen on rough paper, 8B pencil on stone paper, and charcoal on rough paper. To not bias the participants, we referred to the tools as letters A, B, and C. During the preliminary experiment, we realized that without visual and auditory cues, people often use an unusual amount of force. Therefore, the task was first performed on a scale, and the participants were informed in case the force they used was outside the range of commonly used pressure values for drawing [Piovarči et al., 2018].


8B Penci

Participants A total of 22 participants (15 male and 7 female) aged between 22-33 participated in the experiment.

Observations The limited sensory conditions initially affected the pressure applied by the participants. After a couple of trials, they all got accustomed to the task and naturally maintained pressure used for drawing. Most participants used common descriptive terms, such as vibration, drag, friction, smoothness. The tools also were often compared to known drawing instruments. The ability to describe the tools was facilitated by the comparison task when users were able to contrast the feedback of different drawing tools.

7.6.2 Study 2: Distinguishing Drawing Tools

In the second experiment, we verify whether participants can distinguish between the three drawing instruments from Study 1. Additionally, we introduce our replicas to test whether the participants can associate them with the correct traditional counterparts.

Task and Stimuli In each trial, the participants were presented with one tool (A, B, or C) at random and asked to identify it. After three trials with each traditional tool, we introduced our reproductions and performed two repetitions

with each 3D printed tool and one repetition with each traditional tool again. The order of trials was randomized. We assume the study to be successfully completed if a participant correctly identified the type of all traditional tools with an 80% success rate.

Participants All participants of Study 1.

Results All but one participant completed the experiment successfully. Having completed successfully was a prerequisite to participate in further studies. In the vast majority of incorrect answers, participants confused the pencil with the charcoal. This can be explained by the orientation-dependent haptic feedback produced by the charcoal. Using the tool with a short edge can momentarily lower the drag making it similar to the pencil. In 98% percent of cases, our tools were associated with correct traditional counterparts. Interestingly, there was less confusion between our replicas of the pencil and the charcoal than between the original tools. We attribute the increase to more consistent feedback provided by our tools, i.e., absence of wear. Chi-square goodness of fit revealed that there is a significant effect (p-value < 0.001) of the type of the presented tool on the associated traditional counterpart. In a post hoc analysis, a pairwise comparison using binomial tests with Holm-Bonferroni correction revealed that the preference of assigning our replicas to their corresponding drawing tool is statistically significant (p-values < 0.001 for all drawing tools). The above results demonstrate that not only can the participants correctly identify traditional drawing tools using only the haptic feedback, but also, our replicas are most of the time (98% cases) associated with the correct traditional counterparts.

7.6.3 Study 3: Comparison with State of the Art

Our third study builds upon Study 2 and compares our replicas with industrial and state-of-the-art approaches in academia to verify whether users show a stronger preference towards our solution.

Task and Stimuli In each trial, participants were presented with a triplet of drawing instruments, i.e., reference and two tests. The reference tool was always taken from the following set of traditional drawing tools: 8B pencil on stone paper, ballpoint pen on rough paper, and charcoal on rough paper. One of the tests was our replica of the reference tool, while the second test was an

alternative reproduction. The alternative solutions consisted of (a) the reproduction produced by Piovarči et al. [2018] for a glass substrate of a standard tablet, (b) surface geometry reproduction using a 3D printer with ballpoint pen which measured closest to the real counterpart (Section 7.5.1), and (c) a range of commercial styli with drawing surfaces: Apple pencil, Wacom tablet with different drawing nibs, Paperlike surface. The task in each trial was to identify one of the test tools whose haptic feedback is most similar to the reference.

Participants A total of 11 participants (8 male and 3 female) aged between 21-33 participated in the experiment.

Results The results of the experiment (Figure 7.18) demonstrate that our reproduction provides a better match of haptic reproduction when compared to any other drawing tools considered in the experiment. Our tools were chosen in 98% of cases. The remaining 2% are the replicas produced using the recent work of Piovarči et al. [2018]. An analysis of pairwise differences using binomial tests with Holm-Bonferroni correction revealed that the preference of our tools in each trial is significant (p-values of < 0.006).

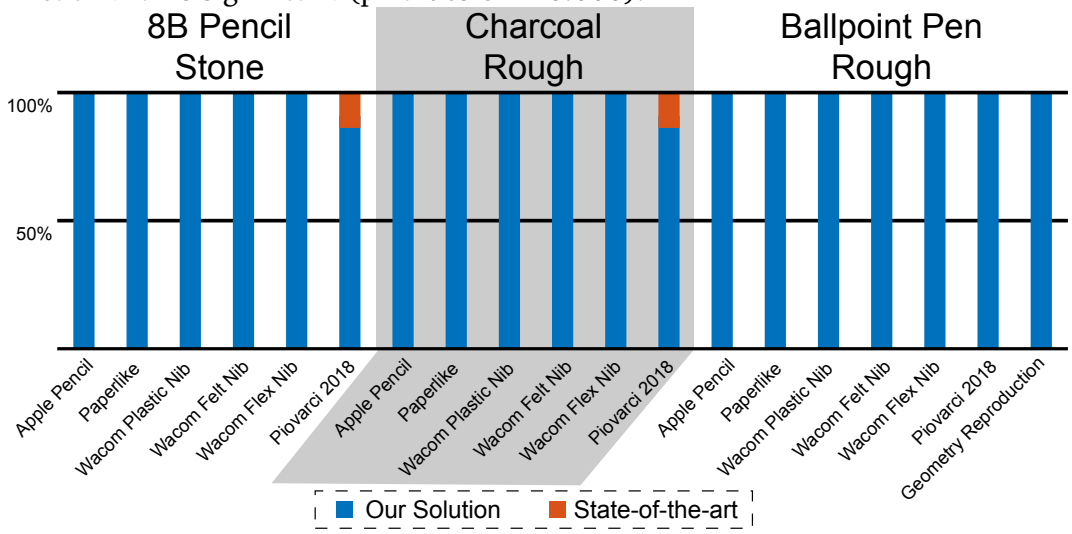

Figure 7.18. Bar-plot showing the similarity of our tools when compared to other digital styli. For each reference, indicated on the top, and an alternative solution, indicated below the bars, the plot shows the percentage of participants finding our replica more similar than the alternative commercial or research solution.

7.6.4 Study 4: Realism of Reproductions

In the last experiment, we want to assess how close our optimized replicas are to their traditional counterparts. Since showing that two tools are indistinguishable is difficult from the statistic point of view, we instead analyze the perceived distances between our replicas and the original tools.

Task and Stimuli Similarly to the previous experiment, in each trial, participants were presented with one reference and two test sets of drawing instruments. The reference tool was one from the following set: 2H pencil on stone paper, 8B pencil on stone paper, ballpoint pen on rough paper, and charcoal on rough paper. One of the tests was our replica of the reference tool. The second test was a traditional drawing tool with a paper substrate. The task was to identify which of the two test tools were closer to the reference. The task was performed 6 times for each reference tool in progressively harder trials. We first compared our replica with the two other tools from the convex hull of haptic feedback. Next, we compared our replica to a novel sensation produced by either varying the drawing paper or the drawing tool. Finally, we compare our replicas with the original tools they were meant to reproduce. All of the trials were randomized to avoid ordering effects.

Participants A total of 10 participants (4 females and 6 males) aged between 22-28 participated in the experiment.

Results The raw data from the experiment (Figure 7.19) demonstrates that even though not in all cases our replicas are often confused or preferred over traditional drawing tools. We argue that this is already a good result since it is often reported that commercial solutions feel nothing like traditional tools. Additionally, our replicas of the ballpoint pen, 8B pencil, and even 2H pencil, were often confused with their traditional counterparts. To further analyze the results, we apply the algorithm of [Piovarči et al., 2018], to recover a small two-dimensional perceptual space for the set of the tools considered in this experiment (Figure 7.20). We can observe that our reproductions are close to the target tools and do not form a separate cluster. This suggests that our stylus-surface combinations produce realistic haptic feedback similar to traditional drawing instruments. The biggest discrepancy is between the reproduction of the 2H pencil on stone paper. This can be likely explained by the wear of the hard pencil and its reliance on orientation to produce a consistent haptic response. Nevertheless, even in such a challenging scenario, our reproduction is closer than alternatives.

It is interesting to note that the recovered perceptual space approximates the original space of [Piovarči et al., 2018], which suggests that our studies are consistent with the previous work.

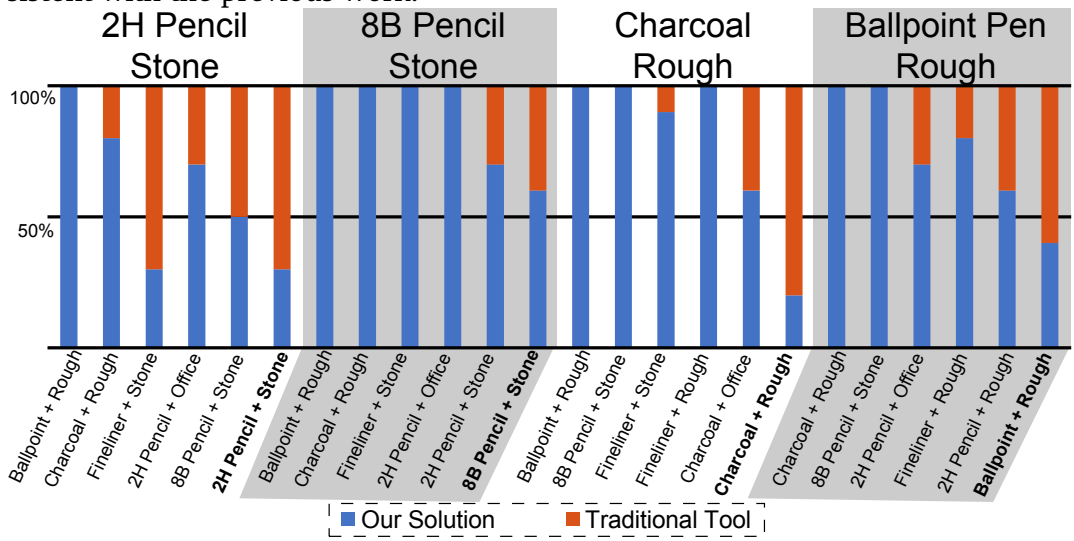

Figure 7.19. Bar-plot showing the similarity of our tools when compared to traditional drawing tools. For each reference, indicated on the top, and an alternative traditional tool, indicated below the bars, the plot shows the percentage of participants finding our replica more similar than the alternative tool.

7.6.5 Summary

The experiments above allow us to draw two main conclusions. First, the replicas of the traditional tools provided by our technique outperform all the investigated research and commercial solutions. Second, our replicas provide haptic feedback that is close to the one generated by their traditional counterparts. While the reproduction of the feedback is not exact, it is essential to note that digital drawing tools should not replicate all phenomena governing the haptic feedback of the traditional tools. For example, the wear of the tool changes how the tool feels over time, but it is not desired for the digital tools to replicate the effect.

7.7 Survey With Professional Artists

One of the main goals of reproducing haptic feedback on digital tablets is to improve the drawing experience of professionals as well as boost their produc-

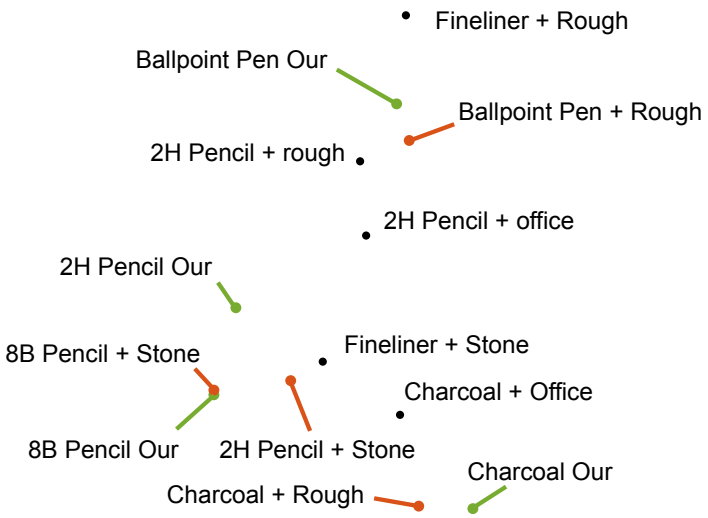

Figure 7.20. Two-dimensional perceptual space computed for tools in our experiments. The Euclidean distance between the samples can be interpreted as a perceived distance between different drawing tools combinations.

tivity and creativity. We invited several professional artists to draw with our 3D printed stylus-surface combinations and asked them how they compare to traditional drawing tools and commercial alternatives.

Questions and Stimuli The survey started with a short interview to acquire information about the background and skill set of each artist. We then followed with questions about traditional tools and their strengths/weaknesses when compared to digital styli. Afterward, we presented them with a selection of traditional drawing tools. They were able to compare them to commercial tablets from Wacom, Apple, as well as our 3D printed replicas. To eliminate the visual bias we use a passive setup with no stroke rendering for the digital styli. The artists were motivated with questions regarding the advantages and disadvantages of our tools. Finally, we asked them if they would use the proposed designs in their daily work.

Participants 5 professional artists (3 female, 2 male) aged 21-35 were invited to participate in the study.

Results Most artists reported that missing haptic feedback is an issue that affects the concentration and ability to swap between traditional and digital tools quickly. However, one can get used to it when drawing on a daily basis. When

comparing with traditional tools, the currently available digital tools lack the feeling of surface texture. Paperlike was closest to the feeling of a soft pencil on paper, but the structure was still weaker.

Our ballpoint pen reproduction was received very well, and it was reported as feeling very close to the real pen, with the main difference being the lack of the feel of the rolling ball. Our charcoal reproduction was also received well and reported to be very similar to the real charcoal. The main difference was the wear that our tool does not reproduce. One participant noticed the stickiness of our materials and claimed it to be unpleasant. The 2H pencil reproduction was also considered close to the real pencil. The main disadvantage was the wrong auditory feedback. Nevertheless, the tool felt real as it reproduced the feel of the paper texture realistically. Our reproduction of 8B pencil was ranked as the best replica with multiple participants rating it to be the same as the reference.

All of the participants expressed interest in trying to use our tools for a more extended time. They also expressed their interest in testing whether the enhanced haptic feedback leads to strokes more similar to traditional tools. The survey demonstrates that our digital tools were well-received, and the professional artists appreciated the haptic feedback. Despite some limitations, the feedback provided by our tools makes them an exciting alternative for currently available digital tools. For a detailed transcript of the study, please refer to Appendix D.

7.8 Limitations And Future Work

Our results are limited by the available materials and resolution of fabrication devices. To reproduce materials with a large friction coefficient, we are limited to use rubbery resins, which have an unintentional side effect of adding stickiness to our tools. Unfortunately, this is not captured by the perceptual space we use, and while most artists did not report the problem, stickiness should be minimized at least through material selection. An additional fabrication constraint lies in the simplicity of our designs. Our styli are composed of a rigid holder and a swappable cone-shaped nib. This design limitation is most notable for the ballpoint pen, which was reportedly close to the target but perceived by some participants as a hard pencil, which indeed is very close in the perceptual space to the pen. The reported difference was the perceivable lack of a rolling ball in our design. In future work, it would be interesting to see whether addressing the limitations in hardware, design space, and perceptual modeling can make the replicas of the tools feel even more similar to the traditional tools. Our fabrication-in-the-loop optimization procedure is a natural approach to try incorporating such exten-

sions.

Besides manufacturing constraints, there are two modalities of drawing tools that we did not consider during optimization: sound and wear. Even though the sound of our tools is already similar to the one produced by real tools, the overall drawing experience would likely improve with accurately synthesized audio. The wear also does not match that of traditional instruments. This is mostly noticeable on tools such as charcoal, which behave differently based on wear. However, not mimicking the wear of traditional tools is a practical consideration since introducing significant wear would significantly lower the lifespan of our tools and add unacceptable debris in electronics.

During formative experiments, we attempted to manually optimize for a surface that obtains similar haptic feedback to a ballpoint pen. After printing more than 60 designs, we were unable to achieve a satisfactory solution. This initial experiment motivated us to design an automatic method that provides significant time savings. In the future, it would be interesting to conduct detailed studies with expert designers to compare quantitative speed-up and improvement provided by methods like ours to direct and manual human-driven search.

Exciting avenues for future work also include improvements to our fabrication-in-the-loop methodology. Currently, our algorithm explores new designs via greedy sampling based on expected improvement for which we provide an efficient, closed-form solution. Alternatively, one can exploit the capabilities of 3D printers to produce multiple designs in parallel. Such an approach requires finding a set of best candidates from a continuous domain. Unfortunately, we are not aware of an analytical solution to this problem; therefore, the key here lies in formulating an efficient numerical algorithm. Our surrogate model can be potentially enhanced by reusing previously printed designs to perform more measurements in combination with newly printed ones. However, since not all combinations provide the same information gain, focusing on efficient ways of selecting the combinations that maximize the improvement given a time budget for measurements is a promising extension.

Further improvements and future work were also suggested in our survey by artists. Some of them reported that our tools have the quick-and-dirty feeling of traditional instruments, which is associated with the stochastic nature of the drawn strokes. It would be interesting to consider a co-optimization of the drawing tools and the synthesized stroke to match not only the haptic but also the visual feedback. Another direction of future work includes the optimization of a limited set of distinct and representative drawing tools to provide a user with a small set of tools fulfilling their needs.

7.9 Conclusion

Despite the success of digital drawing tools, fabricating tools which can closely replicate the haptic feedback of the traditional drawing instruments is a challenging problem. The problem requires a joint optimization of the drawing surface and a stylus, which accounts for limitations of fabrication techniques. Unfortunately, due to the complexity of the phenomena which govern the feel of the tools and the scale at which they occur, standard optimization techniques that rely on numerical simulation or direct reproduction of material properties and geometry do not lead to successful solutions. In this work, we demonstrate that for the class of problems, where simulating physical phenomena is more expensive than fabricating and measuring exemplar solutions, it becomes beneficial to replace simulation with fabrication. Consequently, we propose a fabrication-in-the-loop optimization procedure for replicating traditional drawing tools. A key ingredient for making such a procedure successful is an efficient sampling of the design parameter space, which, in our case, is realized using Gaussian Processes. Such an approach not only enables efficient sampling of the design space but also allows for performing the optimization directly in the perceptual space of drawing tools, which focuses the search on perceptually-relevant features. We applied our technique to fabricate a wide range of tools using several fabrication techniques. The user experiments with casual users, as well as a survey with professional artists, confirmed that our optimization strategy produces tools with realistic haptic feedback, which closely resembles the behavior of the traditional tools. When compared to existing solutions, our tools are preferred over all investigated alternatives.

Chapter 8

Towards Spatially Varying Gloss Reproduction for 3D Printing

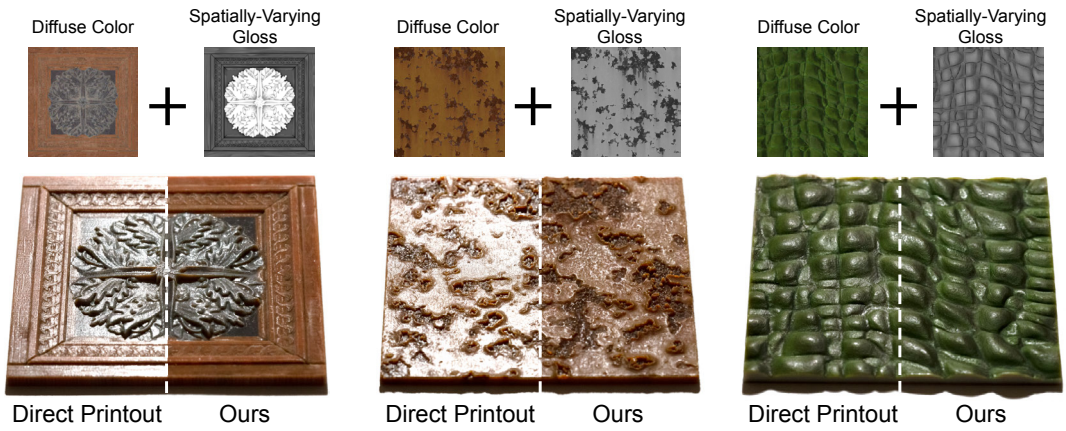

Figure 8.1. The input to our system is a diffuse color and spatially-varying gloss. We first reproduce the color using commercial ink-jet printers (left halves). Next, as a post-processing step we use our custom printer to jet varnishes that match the input reflectance (right halves).

In previous chapters we applied perception-aware fabrication to haptic reproduction. We demonstrated that it is possible to go beyond the hardware gamut of a manufacturing process. By carefully optimizing the digital designs we can mimic the haptics of materials with an order of magnitude finer geometry than available on a commercial 3D printer, as well as, the haptic feedback of materials that lie outside of the printing gamut like metals and graphite. However, perception-aware fabrication is not limited to haptic reproduction. The proposed methodology can be applied to other problems where human observers judge the

quality of manufactured goods. In this context a prominent direction of future work lies in applying perception-aware fabrication in the field of appearance reproduction.

Reproducing objects' appearance is essential for their visual appeal. Recent advances in multi-material 3D printing allow us to create intricate geometrical shapes with faithful color [Sumin et al., 2019] and translucency [Urban et al., 2019] reproduction. The method of choice for fabrication is ink-jet 3D printing [Sitthi-Amorn et al., 2015; Stratasys, 2016], which allows for deposition of colored materials with high spatial resolution. Going beyond color and translucency, however, another crucial component of an object's appearance is gloss. When carefully chosen, it can convey objects' functional properties, enhance perceived value, or create intricate patterns that can improve the visual appeal of objects.

Despite the significant visual influence and potential of carefully edited gloss, we rarely observe 3D-printed parts with spatially-varying gloss properties. Current state-of-the-art 3D printers, such as [Stratasys, 2016], and their software offer minimal choice. Naturally, printed objects have a glossy finish. If a matte finish is desired, it is created by covering the object with support material that, upon removal, introduces roughness. Irrespectively of the chosen option, the final gloss of the object further depends on the printed geometry.

One of the main challenges in fabricating spatially-varying gloss properties lies in the printing process itself. This limitation is directly linked to the hardware design. The ink-jet printheads used in 3D printing have tiny nozzles, typically around 20 microns in diameter, to enable high-resolution printing. Consequently, the deposition of highly viscous materials with larger particle size is challenging and hampers the robustness of the printing process. On the other hand, a wide range of reflectance properties is usually achieved by different concentrations of reflective or absorbant particles, which increase the viscosity of the materials significantly [Tipsotnaiyana et al., 2013]. As a result, the same nozzles that are required for achieving high-resolution prints are not suitable for printing a wide range of gloss. Various alternative manufacturing techniques with a wider gamut of printable gloss were proposed [Matusik et al., 2009; Pereira et al., 2017] to alleviate this issue. However, due to the difficulty of integration with the high-resolution color 3D printing process, there is still no well-established workflow for full appearance fabrication. The only available solution currently allowing for modifying gloss using commercial ink-jet printers is to change the surface roughness by introducing a surface geometry [Rouiller et al., 2013; Elkhuizen et al., 2019]. However, the range of gloss that can be produced by these techniques is limited not only by the properties of 3D printable materials but also by the printing resolution.

To address the current limitation in gloss reproduction, we propose a novel printing system based on PICO Pμlse® [1] technology. The printer is capable of depositing materials with very high viscosity and relatively large particle sizes. This gives us the opportunity to control the glossiness of a surface by using off-the-shelf varnishes. We develop a pipeline to support the hardware with a full workflow for manufacturing objects with fine-grained spatially-varying gloss. For optimal printability of varnishes with such a wide range of properties, we develop a calibration procedure that fine-tunes varnish-specific jetting parameters. We use our setup to jet and characterize a large set of varnishes including varnishes based on mineral-spirits, water, and oil. We characterize our varnishes by performing dense measurements of reflectance and fitting them to analytical BRDF models. Based on this characterization, we perform varnish selection to maximize the gloss gamut. We analyze the quality of the samples and show that mixing varnishes affects not only the reflectance but also the quality of the resulting halftone pattern due to on-surface spatial mixing before fully curing. Consequently, we propose a data-driven, simplex-based prediction model for computing the aggregate gloss of a halftoned mixture. Its additional feature is the capability of accounting for the visual quality of the generated halftone pattern. The model allows us to create surfaces with controllable and spatially-varying gloss properties. We demonstrate the capabilities of the entire pipeline by manufacturing several 2.5 D objects with spatially-varying gloss.

The main contribution of this paper is the complete system for gloss reproduction using varnishes with a wide range of viscosities and particle sizes. We discuss all the building blocks required for such a system and show their practical implementations. Our system and the produced results demonstrate the feasibility of using varnishes for fine control of gloss properties of objects produced using 3D printing technology.

8.1 Overview

The rest of the paper describes our system, both hardware, and software, for modifying the gloss properties of 3D objects by applying different mixtures of varnishes. More specifically, we first describe our deposition setup, which is capable of placing different kinds of varnishes atop a 3D surface (Section 8.2). Next, we consider a set of varnishes for our system and demonstrate a procedure for exploring the parameter space of our deposition system in order to find material-specific parameters that guarantee high-quality varnish surfaces (Sec-

[1] https://www.nordson.com/en/divisions/efd/products/jet-dispensers/pico-pulse-valve

tion 8.3). We then present measurements of the reflectance properties of the surfaces produced using our system, and choose a set of varnishes as base materials for our gloss editing printer (Section 8.4). Given the measurements, we also propose a simple yet powerful prediction model that enables the creation of surfaces with predefined gloss properties by halftoning our base materials (Section 8.5). Finally, we demonstrate the results obtained using our system for 3D printed samples (Section 8.6).

8.2 Hardware Apparatus

To modify the reflectance properties of the surface we seek to deposit varnishes with high spatial accuracy. Conventional inkjet printers are not capable of jetting varnish materials due to their high viscosity and large particle sizes. In order to print such challenging materials, we designed a hardware setup based on piezo-actuated needle valves, (Figure 8.2 left). The jetting head consists of a pressurized varnish reservoir, the valve body, a piezo actuator, a nozzle of variable diameter (50-300 microns), and a spring-loaded needle valve (Figure 8.2 middle).

One activation cycle of the printing head is depicted in Figure 8.2 bottom right. To dispense the varnish, the needle valve is opened by the piezoelectric actuator. The valve remains open for a set duration, during which the varnish can freely flow from the nozzle. Afterward, the valve quickly closes. The back pressure on the printing material and force of the closing needle sever a portion of printing material from the print head and project it onto the printing surface. The shape and size of the formed droplets depend on the material properties of the printing material and the amount of energy transferred to the droplet during the jetting process. The energy transferred to the droplet is a non-linear combination of the pressure in the varnish chamber, the time to open and close the valve, and the stroke power of shutting down the needle valve. This relatively simple jetting process coupled with a larger nozzle diameter offers a powerful combination capable of processing a much wider variety of materials than inkjet-based printers at the cost of producing larger droplet size.

Our apparatus combines three jet-valve dispensers (PICO Pμlse®
from Nordson EFD, Providence, RI, USA) with a Cartesian robot (Hiwin Mikrosystems, Taichung, Taiwan), a gantry, and a controller. The Cartesian robot is used to move and locate the dispensers in Z and the sample in X and Y. The dispensers are used to deposit different varnishes. The movement of the Cartesian robot and the timing of the dispensers is coordinated by the controller.

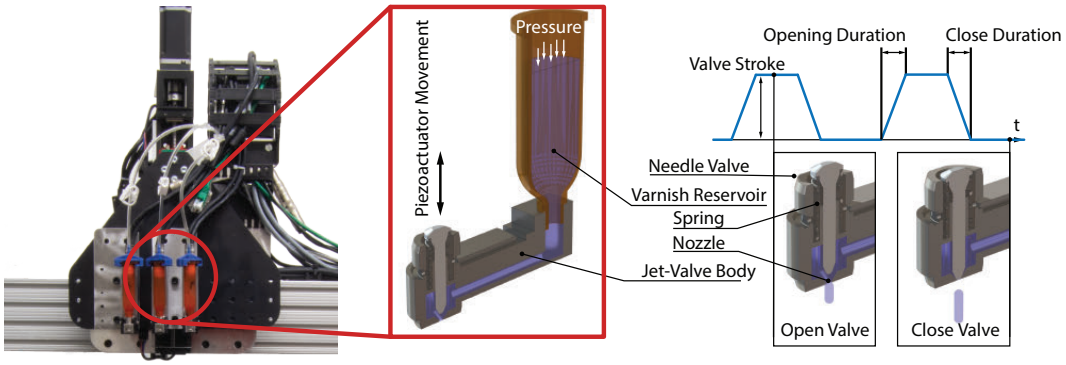

Figure 8.2. Varnish printing apparatus (left) consists of needle jetting valves (middle). To jet with the valve, the needle rises allow pressurized material to flow. By quickly shutting the valve the material is jetted onto the substrate.

8.3 Varnish Jetting

In our experiments, we use of-the-shelf varnishes available in art stores. The varnishes are based on various solvents (water, mineral spirits, oil) and contain various amounts and sizes of particles to produce gloss, matte, or satin finish. The different mechanical and chemical composition of each varnish affect the behavior of droplets during jetting. At the resolution of our jetting processes, even small changes between material batches lead to a decrease in printing quality. As a result, the parameters have to be fine-tuned for each varnish at printing time to achieve consistent droplet quality. In this section, we describe the effect of printing parameters on the resulting droplet shape. Based on this, we propose a varnish calibration scheme. Next, we investigate the effect of droplet spacing on the uniformity of the surface coverage. Finally, we present a set of varnishes that can be consistently applied using our hardware setup.

8.3.1 Effects of Printing Parameters on Droplet Shape

The printing parameters have a direct but non-linear influence on the shape of the printed droplet. To better understand this relationship we performed an experiment. We manually optimized parameters for a varnish to be able to jet the material. Next, we modified different parameters in isolation to investigate their effect on droplet shape. To compare the shape and geometry of the droplets we use Gelsight scans [Yuan et al., 2017] (Figure 8.3).

By visual analysis of the results, we can observe that the valve stroke has the highest impact on energy introduced into the droplet. Excessively high pressure

results in a forceful surface impact, causes the droplet to splash into the surrounding area. The valve open and close time affect how much of the varnish is allowed to escape the valve and accumulate at the nozzle. The accumulation has a direct influence on the dot size. Pushing the intervals too far results in the spreading of the varnish on the printing nozzle due to surface tension; this results in non-circular droplets. The pressure affects both the amount of varnish that escapes the nozzle during the opening and how much force is introduced into the droplet. This results in a combined effect, of changing both droplet size and shape.

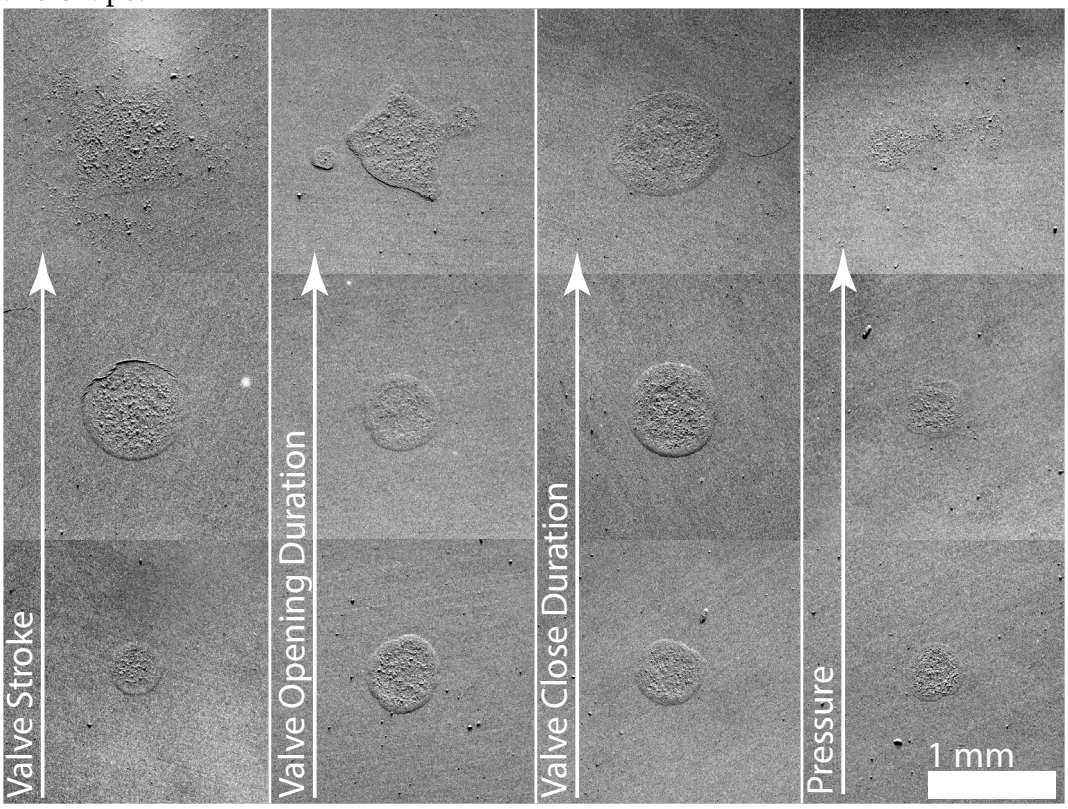

Figure 8.3. Gelsight captures of droplets created by varying the jetting parameters: valve stroke, valve open duration, valve close duration, and air pressure.

To achieve consistent droplets, we perform a parameter sweep before printing to calibrate each of our varnishes. Each varnish is prepared based on manufacturers data-sheets. In particular, we mix the golden varnishes in a volume ratio of $1:1$ with their corresponding solvents. We start from a known working state which allows us to jet the material through the nozzle. Next, we perform a parameter sweep on each of our printing parameters: time the valve is opened,

the time it takes the valve to close, how much the valve opens, and the air pressure of the varnish. Using this procedure, we calibrated a set of varnishes. We present the results of the calibration in Figure 8.4. The varnishes' droplets are captured using both Gelsight (top) and an optical microscope (bottom). The optical microscope was set up with a blue LED to increase the visual contrast. Overall shape of our droplets can be well approximated with circular patterns. Most of our varnishes could print reliably at 450-micron resolution. The exceptions are Golden matte heavy gel, Golden gloss extender mixture; where due to the extreme viscosity the printing resolution dropped to >1000 microns.

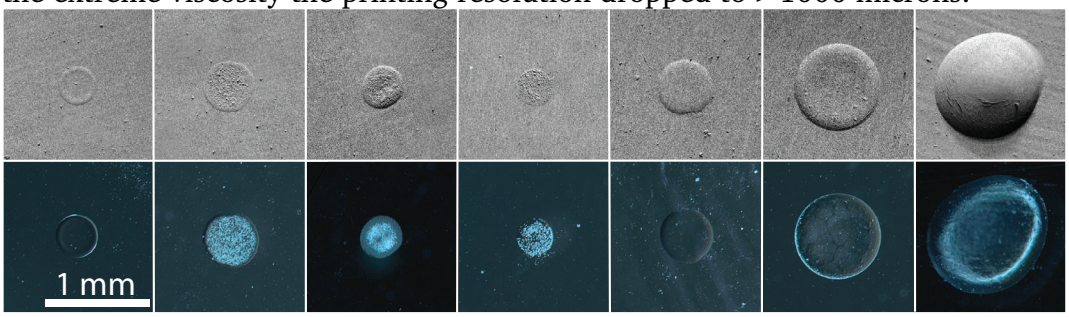

Figure 8.4. Gelsight (top) and optical (bottom) scans of varnish droplets for a selection of off-the-shelf varnishes. From left to right, 1:1 ratio of Golden MSA gloss varnish and Golden MSA solvent, 1:1 ratio of Golden MSA satin varnish and Golden MSA solvent, 1:1 ratio of Golden MSA matte varnish and Golden MSA solvent, Schmincke 611 matte varnish, Schmincke 610 gloss varnish, Amsterdam 115 Matte Varnish, 1:1 ratio of Golden matte heavy gel and Golden gloss extender.

8.3.2 Effect of Varnish Spacing

The reflectance of an object depends not only on the material but also on the microstructure of the geometry. In our printing apparatus, the microgeometry is affected not only by the jetting parameters but also by the spacing between printed varnish droplets. In our setting, we do not wish to affect the glossiness through geometry variation. We rather rely on our varnishes which have reported specular gloss[2] ranging from very matte (0.9 gloss units measured at 60-degrees) to high gloss (87 gloss units measured at 60-degrees).

To investigate the effect of droplet spacing, we conduct the following experiment. We deposit matte varnish with optimized parameters using progressively smaller spacing between the dots and observe the resulting surface roughness

[2] https://www.goldenpaints.com/technicalinfo/technicalinfo_polvar

(Figure 8.5). We can note that in the selected range the surface roughness does not deviate significantly due to spacing. On the other hand, the thickness of the applied film is correlated with the change in spacing. We attribute this behavior to the relatively long time (∼30 minutes) it takes for our varnishes to cure after printing. Varnishes in the liquid state are still allowed to flow and form a uniform film on the surface. In order to maintain a uniform coverage, we opted for the spacing of 320 microns which corresponds to tiling the inscribed squares of our varnishes.

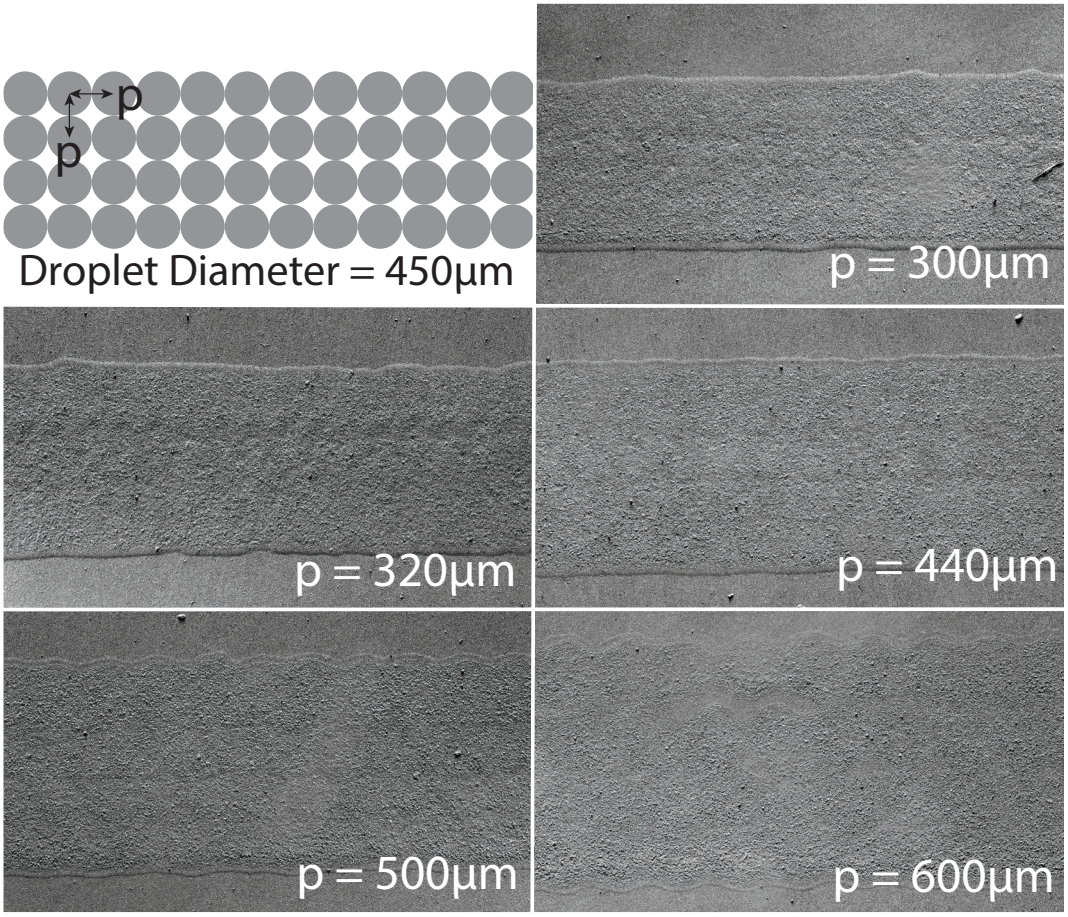

Figure 8.5. Effect of varying spacing. In the observed range increasing the spacing leads to thinning of the film created by the deposited varnish.

8.3.3 Optimized Jetable Varnishes

Using our hardware apparatus we optimized a set of printable varnishes capable of consistently jetting at 450-micron resolution (Figure 8.6). The varnishes were

applied on a transparency sheet. The translucency and high degree of smoothness of the transparency sheet guarantee that it does not have an effect on the varnish appearance. With our system, we can successfully jet varnishes with a wide range of viscosities, and more interestingly, handle matte varnishes with particles. This capability allows us to achieve broad coverage of reflectance from very glossy to matte.

Figure 8.6. Photos of our printed varnishes on a cylindrical setup with a line light source.

8.4 Varnish Selection

The set of varnishes printable on our hardware covers a large range of visual gloss. To quantify the exact reflectance properties we measure the varnishes on a custom-built setup. Next, we approximate the BRDF of varnishes with a low-parameter analytical model, which allows us to represent them in a low-dimensional space. We use this space to visualize the gamut covered by our varnishes and perform varnish selection.

8.4.1 Measurement Setup

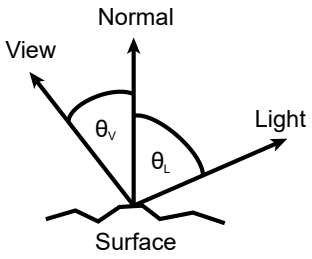

Off-the shelf varnishes are designed to provide optimal surface properties from all viewing directions. Our setup does not introduce significant anisotropy thanks to the spatial mixing that occurs at the boundary of deposited varnishes. Consequently, we assume that our varnishes have isotropic reflectance properties. Such an assumption significantly simplifies the hardware required to capture the reflectance since the four-dimensional space reduces to three dimensions. We base our setup on a design proposed by Ngan et al. [2005]. Our setup (Figure 8.7) consists of a black cylinder to which the measurement sample is attached and a series of light sources located at 15, 30, 46, and 60 degrees. We use a matte black cylinder to eliminate the influence of the measuring setup on the

recovered reflectance. Under the assumption of isotropic uniform reflectance we can capture a representative dense reflectance measurement from a single shot of each light source (Figure 8.7 right). The reflectance is plotted as a function of the signed difference between the angle of the view direction with the surface normal and the angle of light direction with the surface normal: $\Delta\theta = \theta_V - \theta_L$.

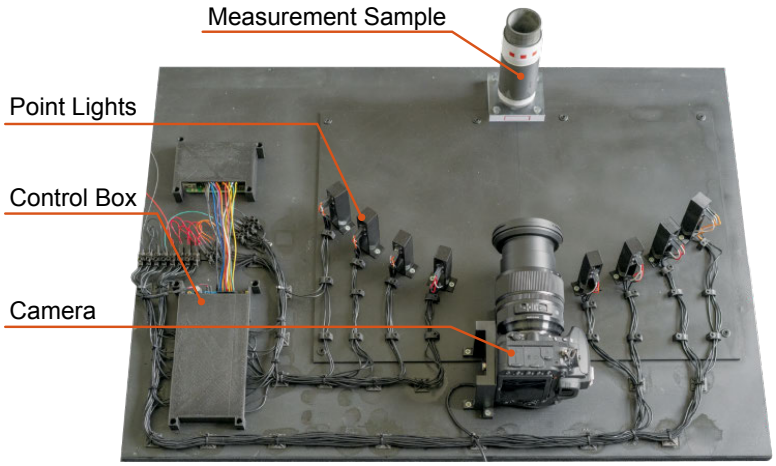

Figure 8.7. Our measurement setup and sample capture of matte and satin varnish.

8.4.2 BRDF Model Fitting

To perform a selection of our varnishes we seek a low-dimensional embedding in which we can compare the reflectance. Good candidates for such embedding are analytical reflectance models. We consider fitting multiple models into our data: Ward isotropic model [Ward, 1992], Cook-Torrance model [Cook and Torrance, 1982] with Backmann distribution [Cook and Torrance, 1982], and Cook-Torrance model with GGX distribution [Trowbridge and Reitz, 1975; Walter et al., 2007] and their multi-lobe variants [Lafortune et al., 1997]. We fit the model to our data by minimizing:

$$\min_{m,F_0} |D - M(d, m, F_0)|, \tag{8.1}$$

where M is the analytical model; the analytical model has input roughness $m \in \mathbb{R}^l$, and Fresnel factor $F_0 \in \mathbb{R}^l$ (or intensity in case of Ward model); l is the number of lobes that we are fitting, d is the diffuse component estimated as the 5th percentile of intensities captured in our measurement, and D are our measured values. We minimize Equation 8.1 using L-BFGS optimization [Nocedal

and Wright, 2006]. Our fitting attempts to approximate the data as closely as possible and does not impose constraints on physically correct reflectance.

An interesting feature of our fitting is that, unlike previous work [Ngan et al., 2005], it does not need additional weights to properly approximate the measured lobe. This is thanks to our hardware setup. By measuring samples on a cylinder we achieve denser sampling in the regions of mirror reflection and, therefore, additional reweighting is unnecessary.

We analyzed the data by comparing fitting errors and found that the Cook-Torrance model with GGX distribution best approximates our measurements. We also verified that a single-lobe model is sufficient to capture the appearance of our samples and a multi-lobe model does not significantly improve our results. Figure 8.8 shows the final fits to three of our varnishes and their corresponding renderings using the Mitsuba renderer [Jakob, 2010].

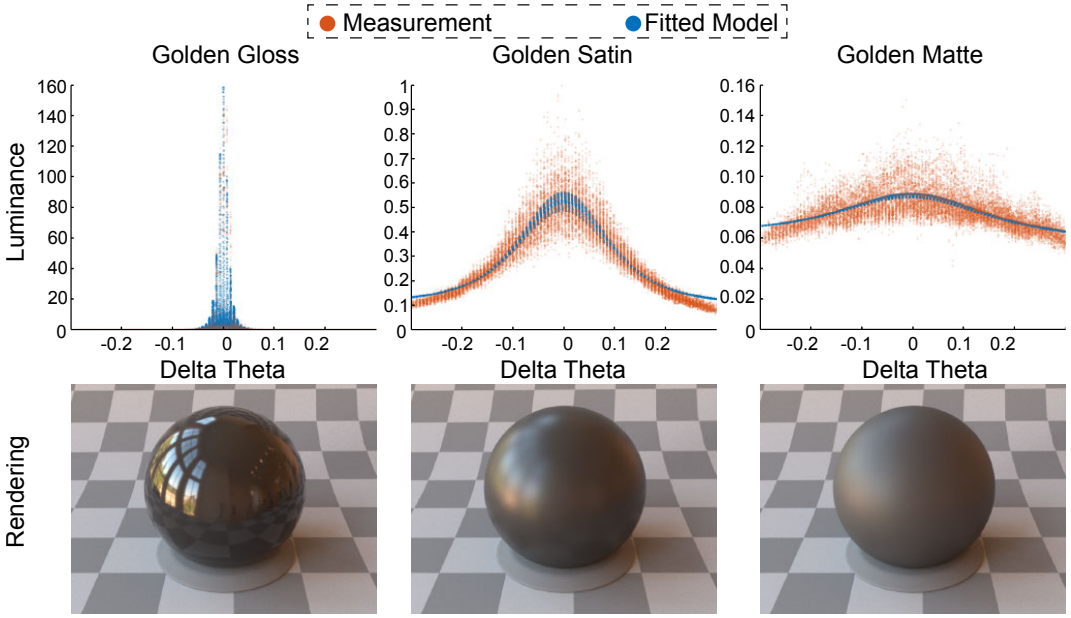

Figure 8.8. Fitted raw measurements of our varnishes (top) and corresponding rendering with the fitted model parameters (bottom).

8.4.3 Reflectance Gamut

Fitting a Cook-Torrance model allows us to plot our varnishes in a two-dimensional space formed by roughness and Fresnel factor. Due to the non-physical nature of our fitting process roughness and Fresnel factor are not orthogonal parameters in the analytical model. Therefore, we cannot simply quantify our gamut by a convex hull of our varnishes in the parameter space. Instead, we visually inspect

the varnishes and select those that maximize the perceived gloss.

Our hardware apparatus can print simultaneously with three varnishes. We wish to select these varnishes such that they span the largest possible gamut of reflectance. We opt for Golden Matte and Golden Gloss varnishes as they represent our most diffuse and most specular sample respectively. For the last varnish, we decided to use Golden Satin. We motivate this choice by three factors: an intermediate varnish allows us to generate halftone patterns with lower visibility of dithering artifacts, the varnish is visually similar to Schmincke Matte, and the varnish is spirit-based as the other Golden varnishes.

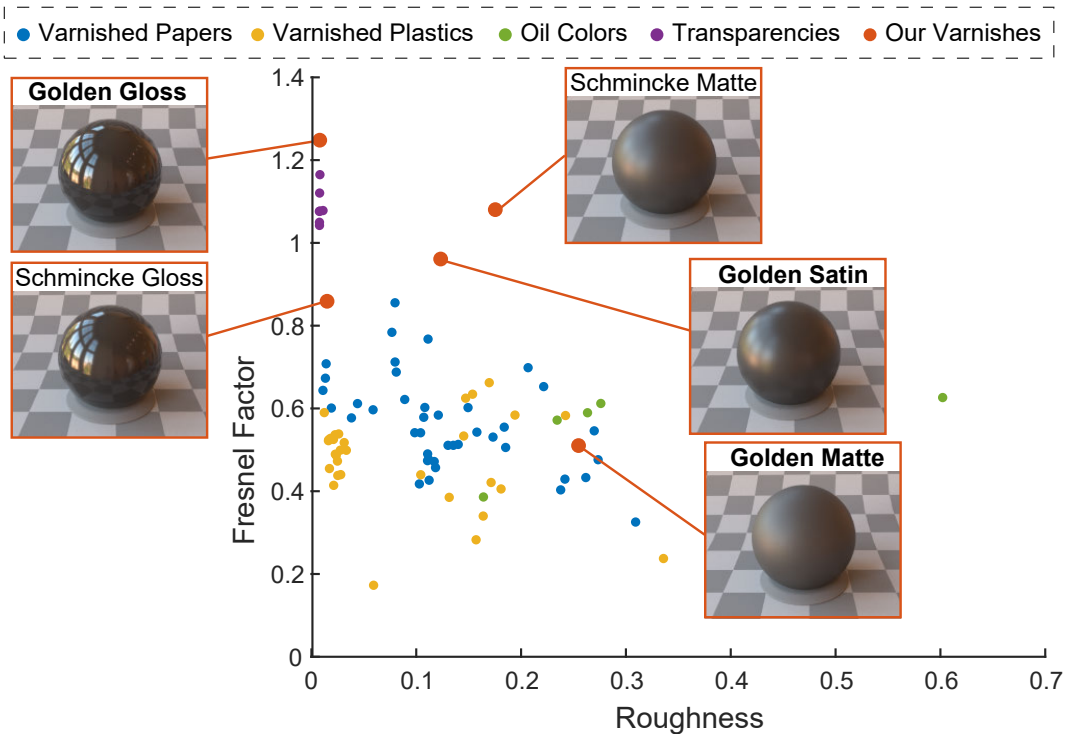

Figure 8.9. The reflectance gamut achievable by our varnishes (orange). The gamut contains various hand-made samples for comparison: varnished paper (blue), varnished plastics (yellow), oil colors (green), and transparencies (purple). The varnishes we selected for printing are in bold.

To estimate the gamut of achievable gloss we assume a single lobe model as it holds well for the base varnishes. However, during halftoning the surface is covered with up to three varnish primaries. These varnishes form individual droplets that manifest in the measurements as a mixture of three separate lobes. As a result, the phenomenological model is only a rough predictor of the gloss achievable by our setup and we rely on raw measurements for reflectance predictions.

8.5 Gloss Reproduction

To reproduce the desired reflectance with a limited set of primary varnishes we can use a halftoning algorithm. For each spatial location, the algorithm selects the primary closest to the target. This produces an error which is compensated for by propagating to neighboring locations. The algorithm operates under the assumption that the deposited materials combine linearly. However, since our varnishes are not cured at deposition time they can mix in a liquid state and manifest non-linear mixing behavior. Therefore, we need to verify the linearity assumption of varnish halftones. We prepared three halftoning patterns linearly interpolating from matte to gloss in uniform steps, and printed them on our setup (Figure 8.10 top). We measured the applied varnishes and compare the physical measurement with prediction based on linear mixing (Figure 8.10 bottom).

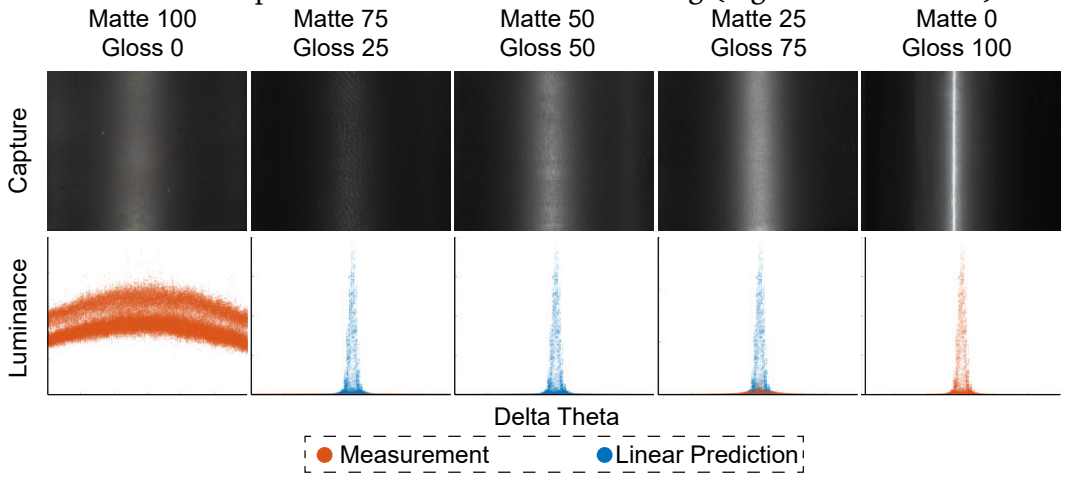

Figure 8.10. Test of linearity of varnish halftoning. We start by generating halftone patterns and applying them on a substrate. The samples are then measured and compared with linear prediction.

We can observe that the patterns of halftoned gloss are mixing during printing. As a result, the 50/50 pattern looks almost as a single material. This mixing affects the visual properties and violates the linearity assumption which means we cannot rely on the linear mixing to produce correct results.

8.5.1 Simplex-Interpolation for Halftone Reflectance Prediction

When observing the interpolated samples we can see that while the linearity assumption does not hold the change in varnish properties smoothly varies from matte to glossy. We leverage this observation to formulate a simplex-interpolation

model for the prediction of the reflectance of halftone samples. Due to varnish dithering the final appearance can be a combination of multiple phenomenological effects. To not impose any assumption on the predicted data we directly use the measurements. The input to the model is a mixing ratio and a data set of previously observed mixing ratios with associated measurements. The output is a prediction of reflectance measurement.

A halftone screen can be treated as a spatial approximation of a given mixture of varnish primaries $\mathscr{P}$. The mixture is parameterized by the mixing coefficients α_i where $i \in \mathscr{P}$ and $\sum_i \alpha_i = 1$. Using this definition a set of N varnishes defines a standard $N-1$ simplex. With a linear halftoning model we could predict the reflectance of various mixtures by using the barycentric coordinates of varnishes in our base set $\mathscr{P}$. However, due to nonlinear mixing such interpolation would not hold. To address this issue we propose a more intensive piecewise linear data-driven model. We manufacture and print a set of interpolated varnishes to better approximate the non-linear mixing (Figure 8.11 left). We used the following mixing ratios: 75/25/0, 50/50/0, 50/25/25, 33/33/33 and their permutations. Due to the non-linear mixing of multiple varnishes, we can not rely on reflectance models. Instead, we represent the reflectance as quantized raw measurements. We parameterize the measurements in the angular domain as the difference between the view and light direction. The difference $\Delta\theta$ is then quantized with a resolution of 0.1-degree change. To predict new samples we perform a hyper-tessellation of the measured data points (Figure 8.11 left). Then to predict a new unobserved mixture we first find its enclosing sub-simplex and then use the barycentric coordinates within the sub-simplex (Figure 8.11 left) to predict a reflectance measurement (Figure 8.11 right).

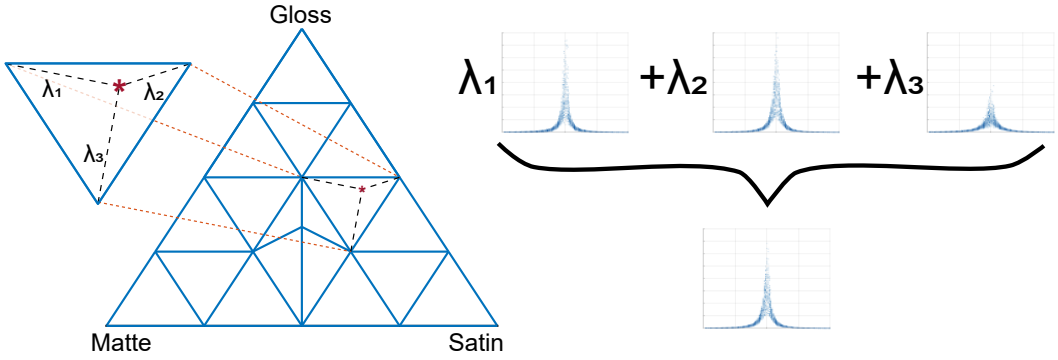

Figure 8.11. A two-dimensional simplex created to predict the reflectance of gloss halftones. To predict a new sample we locate the enclosing simplex and use barycentric coordinates to interpolated our measurements.

We evaluate our model using a cross-validation scheme. We repeatedly leave

out three varnishes and rebuild our model (Figure 8.12 left). Next, we predict the reflectance of the missing varnishes (Figure 8.12 right). We can observe the model has an overall good performance. The largest discrepancy from the prediction occurs when dealing with a glossy material. This is likely due to the highest non-linearity of our function when mixing the glossy varnish since a small amount of different varnish already produces an observable difference in gloss.

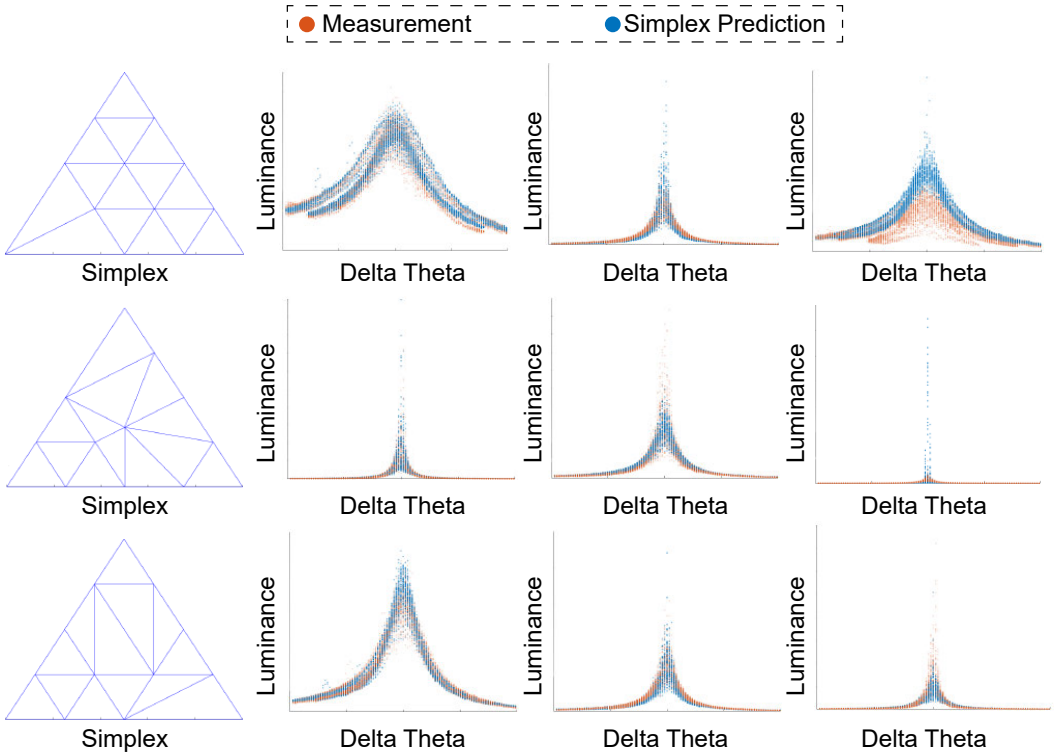

Figure 8.12. Cross-validation of our simplex-based reflectance prediction. Leaving out multiple samples, creating a new tesselation, and predicting the missing samples.

8.5.2 Predicting Halftone Pattern Quality

Due to the large size of our varnish dots, the halftoning pattern can be observed by naked eye from a sufficiently close distance. However, the physical mixing of varnishes masks the halftone pattern and improves the quality of the observed gloss. To take advantage of this property and incorporate it into our halftoning algorithm we propose a heuristic model based on observation of the varnish behavior.

The input to our heuristic model is a desired mixing ratio. The output is a

single number representing the quality of the halftone pattern where a lower number represents a surface with higher quality. To define the heuristic we observed the spatial mixing of our varnishes. The spatial mixing causes a smooth material transition between two varnish droplets. Additionally, a droplet always mixes only in its one-ring neighborhood. To model the spatial blurring of our varnishes we use Gaussian filters with a standard deviation of 0.5 and kernel support of 3×3 which represents the droplets smoothly mixing with the surrounding droplets. The blurring is applied onto halftone screens corresponding to 4×8 cm patches used for our measurements. This is equivalent to 158×316 droplets. Next, we compute the standard deviation of the pattern. The standard deviation estimates the visibility of the blurred halftone pattern. The visibility predictor σ can be evaluated as:

$$\sigma(\alpha) = \sum_{i}^{N} \sqrt{\frac{1}{M-1} \sum (G * P(\alpha_i) - \alpha_i)^2}, \tag{8.2}$$

where M is the number of droplets in the dithering pattern, G is a Gaussian kernel with 3×3 support and standard deviation of 0.5, and P is a function that generates halftone pattern for mixture α. We can see the results of predicted and observed dithering artifacts in Figure 8.13.

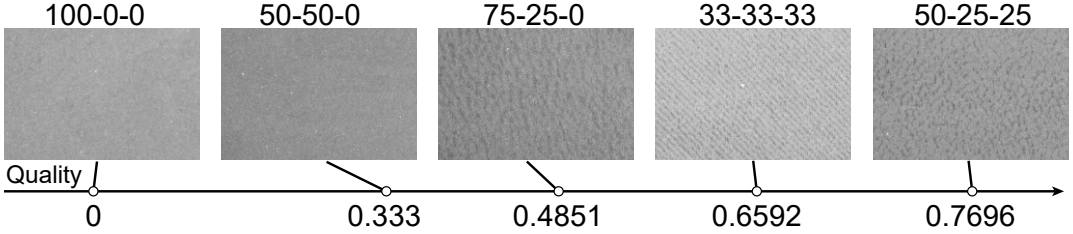

Figure 8.13. Predicting dithering pattern visibility for different varnish mixtures of matte-glossy-satin.

8.5.3 Prescribed Reflectance Reproduction

We use our simplex model to reproduce prescribed reflectance. The input to our method is a per-pixel assignment of the desired reflectance. The output is per-pixel mixing ratios of our base varnishes that best match the input. At each spatial location, we solve a minimization problem searching for the optimal mixing ratio that reproduces the desired reflectance. Since our varnish combinations can result in similar reflectance we encourage halftone pattern uniformity by adding pattern quality as a regularizer. The final minimization problem is then:

$$\min_{\alpha} \sum (R_T - S(\alpha))^2 + \sigma(\alpha), \tag{8.3}$$

where $\alpha = \alpha_1, \alpha_2, ...\alpha_N$ is the set of the mixing ratio of our base varnishes, R_T is target reflectance, S is our simplex predictor which transforms mixing ratios into reflectance measurements and $\sigma(\alpha)$ is the predictor of dithering pattern visibility for mixture α computed using Equation 8.2. We minimize Equation 8.3 with a L-BFGS method [Nocedal and Wright, 2006]. After the mixing ratios are generated we use vector error diffusion with Stucki weights to dither the desired mixing ratios into our varnish primaries [Stucki, 1982; Lau and Arce, 2007].

The regularization term is ensuring to prefer patterns with less visible halftone screens. As an indirect consequence, this leads to a preference for mixing more similar varnishes as lower mixing ratios and mixing ratios close to 50:50 create more uniform patterns. We can observe the effect in Figure 8.14. The roughness map directly corresponds to the Cook-Torrance model. Lower values represent glossier materials and higher values matte materials (Figure 8.14 left). Our dithering pattern (Figure 8.14 middle) is encoded into RGB such that R corresponds to matte material, G corresponds to gloss and B corresponds to satin. For comparison we also show the dithering pattern without our visibility predictor (Figure 8.14 right). We can observe that the visibility predictor successfully reduces the dithering noise by removing isolated varnish droplets.

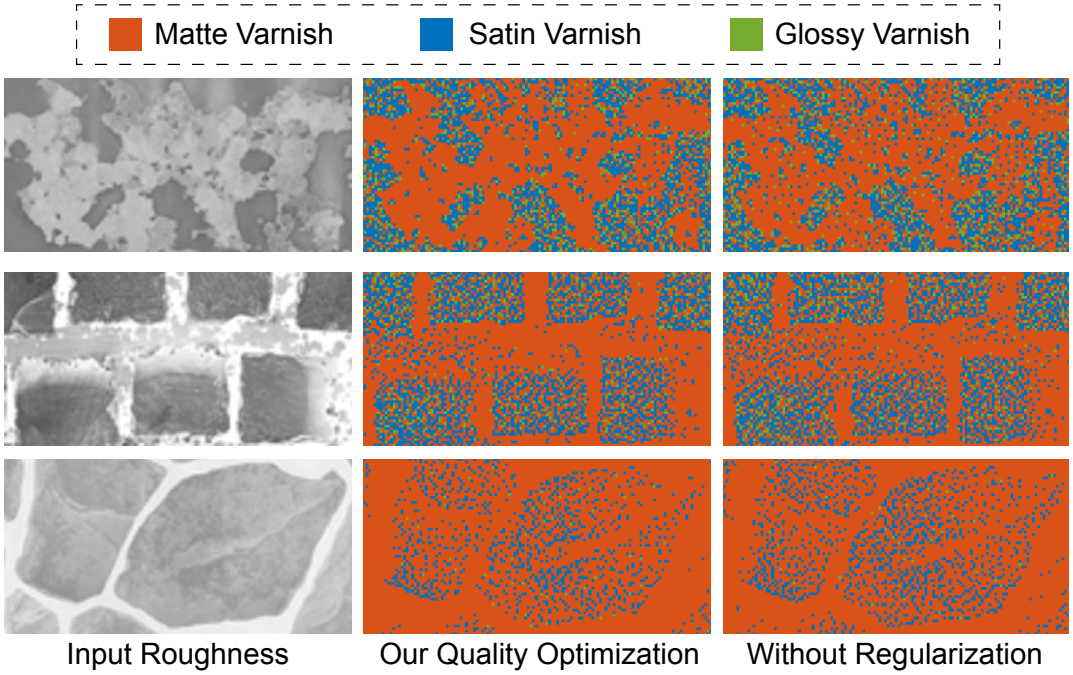

Figure 8.14. Prescribed roughness for a target reflectance (left) is dithered using our simplex model with (middle) and without (right) dithering pattern visibility optimization.

8.6 Results

We demonstrate the capabilities of our system by creating several samples with spatially-varying gloss. We fabricate flat samples of varnishes applied on transparency sheets that showcase the fine-grained spatially-varying gloss achievable using our hardware. Additionally, we chose three height-fields with associated spatially-varying BRDF that we reproduce using our system. We compare the final manufactured pieces against direct printouts from a commercial printer.

8.6.1 Flat Samples

We visualize the range of achievable reflectance by manufacturing a gradation of gloss ranging from highly matte to highly glossy in Figure 8.15. The transition is composed of uniform patches of gloss manufactured by dithering a predefined mixture of our base varnishes. To capture the samples, we place them perpendicularly to a display showing a checkerboard with a grid size of approximately 3 millimeters. We position the camera in the middle of each sample at a 30-degree angle to capture the reflection of the checkerboard pattern. The resulting transition demonstrates the range of gloss we can achieve with our system.

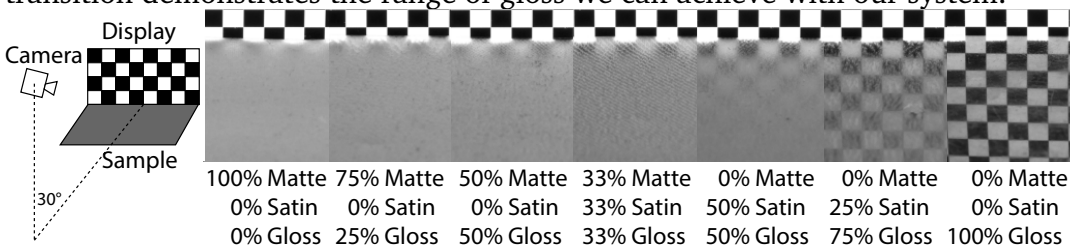

Figure 8.15. A gradation of gloss formed by jetting uniform patches of various varnish mixtures. The images are captured by placing the samples against a display showing a checkerboard pattern.

To validate the capabilities of our setup in creating fine-grained spatially-varying gloss, we prepare two samples: *Waterfall*, and *Knight* (Figure 8.16). The input gloss on *Waterfall* is created by treating its luminance image as a gloss map. For *Knight*, we edited the gloss manually. Both gloss maps are halftoned using vector error diffusion to generate the input to our printer. To capture the manufactured pieces we illuminate them from a 60-degree angle with an area light source and capture the photos with a camera at the specular direction (Figure 8.16 left). To predict the final appearance, we reconstruct the setup in the Mitsuba renderer. For reflectance parameters of the halftoned varnishes we use the fitted BRDFs explained in Section 8.4.2. Similarly to our dithering pattern

visibility predictor we simulate the spatial mixing of our varnishes by blurring the halftoned maps with a Gaussian filter of 3×3 with a standard deviation of 0.5. We can observe that this simple phenomenological model is effective at predicting the appearance of our manufactured pieces (Figure 8.16). The main discrepancy in the photos comes from imprecisions in the light setup used to capture the physical samples. More precisely, the lamp used for capture is not a perfect spotlight which is used for the rendering. In particular, the real lamp provides a bright spotlight with a gradual falloff towards the boundaries.

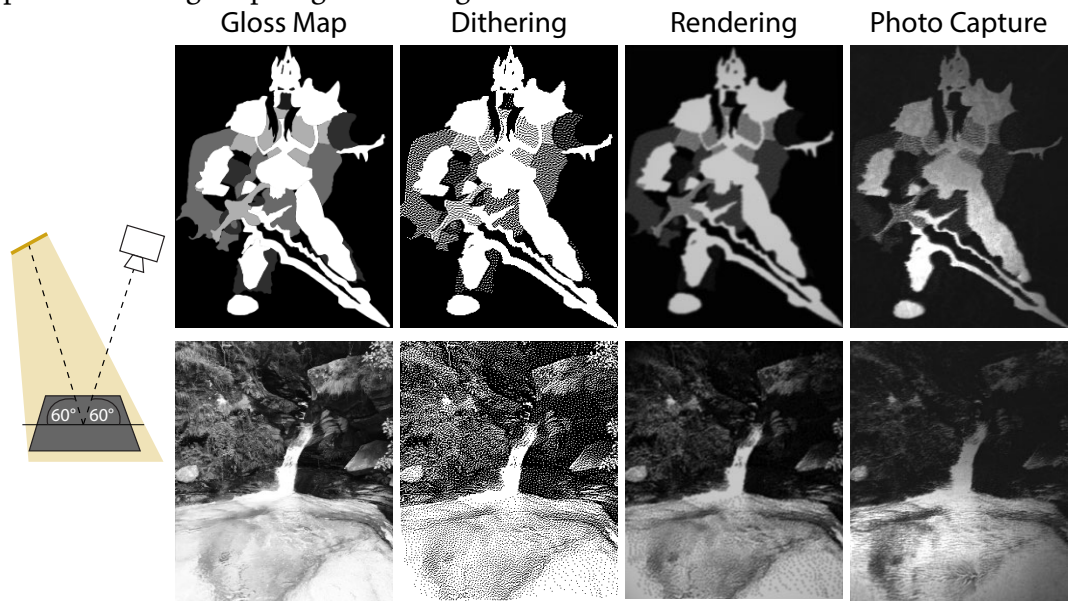

Figure 8.16. Spatially-varying gloss is defined with two varnishes: matte (black) and glossy (white); and dithered using our device. To validate the fabrication we capture the fabricated samples with an area light source and compare with rendered predictions.

Finally, we highlight the capability of our setup to separately manufacture color and gloss. To this end, we printed two colored photos of *Knight* and *Waterfall* on a commercial inkjet printer. To demonstrate the separability we modify the reflectance by placing transparency sheets with spatially-varying gloss created using our setup on top of the printed photos. We capture the samples with a camera at an elevation of 40-degrees and with a light source moving azimuthally from 0 to 180 with a 40-degree elevation, as shown in Figure 8.19. We can observe that the final fabricated samples manifest spatially-varying gloss and high-resolution color.

8.6.2 Height-field Samples

To demonstrate potential integration with current 3D printers, we use our system to apply gloss on previously-3D-printed 2.5D samples. To this end, we prepare three height-fields: a door ornament, a rusted paint, and a polished leather patch. The models have associated albedo, roughness, and metallic maps that encode parameters of a Cook-Torrance model with GGX distribution. The roughness map modifies the microfacet distribution, and the metallic map determines the Fresnel term. We use these maps to obtain parameters of spatially-varying BRDF, which is later inputted to our model for generating a halftoning pattern. The inputs and the corresponding halftoning patterns are presented in Figure 8.17.

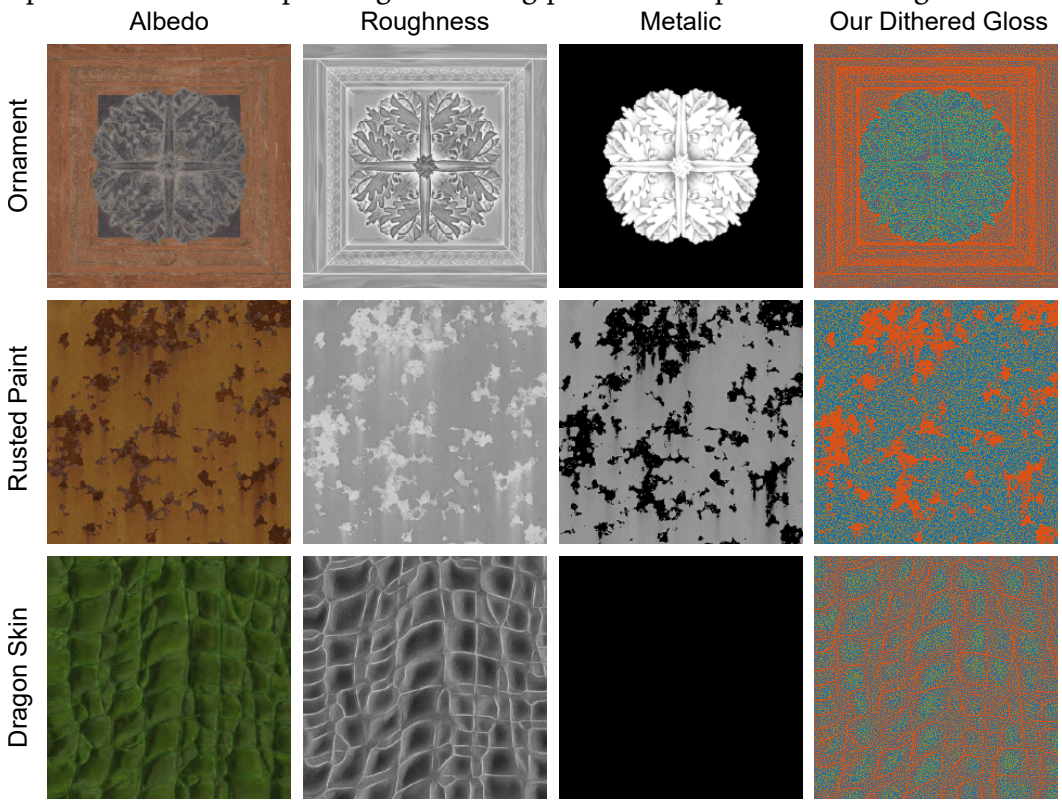

Figure 8.17. Albedo, roughness, and metallic textures for our objects with their corresponding generated halftone patterns.

We print the colored models using Stratasys ObjetJ750 with a matte finish (Figure 8.20, top rows). Then, we use our device to modify the gloss of the printouts (Figure 8.20, bottom rows). We align the printed samples with the jetting head by placing them into the corner of the print bed which is calibrated with respect to the nozzles. To capture the samples, we placed a camera at an elevation of 40-degrees and rotated the light source around the sample at a con-

stant elevation angle (Figure 8.20 top). Figure 8.20 shows captures from three light source positions at azimuth directions of 0, 45, and 90 degrees. To visualize the spatially-varying gloss, we select two regions and plot the 95th percentile of luminance captured by the camera (Figure 8.20 right). It can be observed that even though the initial samples were printed in the printer's matte mode, they exhibit significant specularity seen as peaks in the plots. The luminance profiles for our samples significantly differ. More importantly, in contrast to direct printouts, our system produces samples with a measurable variation between different locations on the samples. Even though the perceived effect is affected by the underlying geometry [Ho et al., 2008], we can observe visible variation in spatially-varying gloss properties. Another interesting outcome is that thanks to our capability of depositing highly viscous matte varnishes, we can achieve significantly more matte finish than the one available on the color printer. This is visualized by almost flat profiles for some locations on our samples.

8.7 Limitations and Future Work

While this paper presents a complete system for applying varnishes on 3D printed objects, there are several areas which should be further investigated before the capabilities of such a system are fully exploited.

In our experiments, we use varnishes with isotropic reflectance. A potential direction of future work is to enhance the varnishes with particles that can produce the anisotropic appearance, though this would also require a method for aligning these particles [Pereira et al., 2017]. Such alteration opens up questions on how to efficiently capture, model, and predict the appearance of anisotropic varnishes. Another limitation of off-the-shelf varnishes is that while they strive to be color neutral, often yellowing or bluing of the underlying substrate can happen. To this end, an interesting direction of future work is to estimate how the varnishes affect the spectral colors and compensate for any undesired color shifts. Additionally, our system assumes that the varnishes are deposited on top of a colored substrate. As a result, the highlights created by our printer are always white. A potential avenue for future work is to investigate how to alter the color of the glossy highlight.

Our varnishes were measured on transparency sheets. While this procedure is useful for characterizing varnish properties, the final reflectance of the surface is a combination of the varnish and the underlying structure, in particular, microgeometry [de la Rie et al., 2010]. Any imperfections in the surface finish have an effect on the final appearance. It is possible to minimize such effects by using a

self-correcting hardware setup [Sitthi-Amorn et al., 2015]. However, a more viable solution is to include the hardware imprecision into the gloss modeling and halftoning algorithm to enable high-precision gloss editing on current hardware solutions.

As we demonstrated in Figure 8.18, when applying our varnishes on tilted surfaces, they stick to the substrate. Unfortunately, we observed that the appearance of the surface changes slowly as the slope increases. We attribute this effect to the staircase artifacts introduced by the 3D printer as well as a different mixing behavior of varnishes on inclined surfaces. While the former can be overcome by sandpapering the surface [Elek et al., 2017; Sumin et al., 2019], handling both requires extending our model such that it compensate for the influence of the slant on the final appearance. This leads to an interesting observation that, in order to print a complex 3D object with uniform gloss, the system should vary the varnish mixtures according to the underlying geometry.

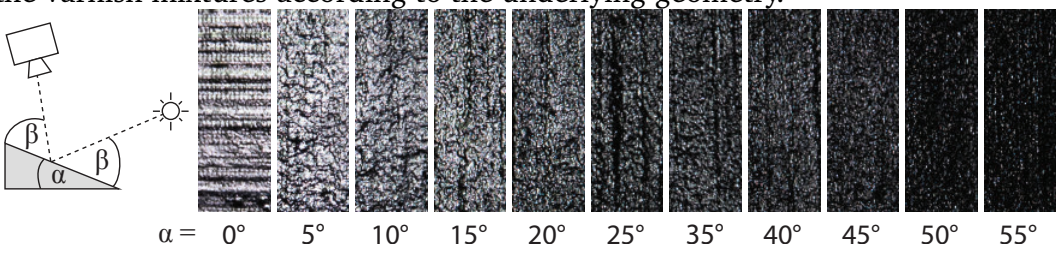

Figure 8.18. Differently slanted 3D printed surfaces covered with the same coverage of Golden gloss varnish, captured at specular configuration. The angles below the pictures indicate the deviation of the surface from the 3D printer's tray. The appearance start to show significant deviations from the flat surface at approximately 20°.

The proposed printing hardware enables separation between color and gloss manufacturing. This is advantageous from the fabrication point of view as fewer materials are necessary to achieve faithful reproductions. However, the joint color and gloss appearance management remains a coupled and challenging problem [Wills et al., 2009]. In this context, an interesting question is how to leverage gloss metamerism i.e., perceiving lobes of different shapes as being equivalent, for appearance reproduction. Here, the main challenge is that gloss lobes with a similar width and height can be difficult to distinguish at a quick glance, but it is possible to see the difference under carefully controlled conditions. Hence, true metamerism of gloss is a function of how carefully one looks. Quantifying this has been attempted by works developing a BRDF similarity metric [Fores et al., 2012; Pereira and Rusinkiewicz, 2012; Sun et al., 2017], but to our knowledge, no standardized solution exists yet. As a result, future work for

understanding the physical and perceptual coupling between color and gloss is critical in achieving high-quality reproductions.

Our model for predicting the appearance of different mixtures of varnishes inherits the general limitations of data-driven approaches. While we demonstrate that a relatively sparse sampling of different mixing ratios leads to a good prediction, due to the non-linearity of the modeled function, better prediction can be achieved by careful selection of the sampling nodes [Warburton, 2006]. Besides improving the model accuracy, the sampling strategy can also accelerate the interpolation with larger number of primary varnishes [Babaei and Hersch, 2016].

The dithering pattern visibility predictor assigns a general score to a pattern irrespective of the appearance of its materials. The relatively simple predictor helps improving the pattern quality by removing isolated droplets and regularizing the dithered patterns. An interesting direction of future work is incorporating the reflectance of the dithered varnishes into the prediction. This way the generated halftones could prefer more similarly looking materials that can lead to smoother transitions.

Since our varnish deposition system uses only a single nozzle it is slower than commercial inkjet printers at depositing a single layer of material. More specifically, covering a 100×50 millimeters rectangle takes approximately 30 minutes on our setup. This process can be significantly optimized by enhancing the hardware setup. However, printing a medium-sized object on a 3D printer requires thousands of layers. In contrast, our prototype varnish application is only a single layer and as such it does not introduce a significant bottleneck into the manufacturing process.

Finally, a single deposition pass of varnishes can handle only 3D objects that are height-fields. Therefore, modifying the gloss of an arbitrary 3D printed object requires multiple passes during which the object is placed at a different orientation. The process requires an accurate alignment [Sitthi-Amorn et al., 2015], or ideally, an automated system. Alternatively, the object can be printed in multiple parts, each of them being a height-field [Herholz et al., 2015]. While this may be considered a significant limitation, in fact, there is no other viable solution to this problem. Current color 3D printing devices generate appearance, which differs depending on the supporting material placement and the local slope of the surface. With printers such as ObjetJ750 [Stratasys, 2016], it is not possible to affect the appearance at the bottom of the object which is attached to the print bed. Considering these limitations, we believe that our deposition system is a viable solution for physical gloss modification with a clear, albeit challenging, path for fabrication of arbitrary 3D printed objects.

8.8 Conclusion

Despite steady progress in 3D printing technology, full appearance reproduction is still a challenging task due to the difficulties in depositing a wide range of materials, lack of established appearance fabrication processes, and limitations on surface finish imposed by hardware designs. In this work, we took a step towards addressing these limitations, enabling further development of appearance fabrication using 3D printing technology. To this end, we presented a novel hardware apparatus capable of jetting highly viscous materials, which enables reliable deposition of a wide range of varnishes. We demonstrated the required steps for building a system for achieving high-quality, spatially-varying gloss of 3D printed samples. We presented a method for selecting printing materials and calibrating their deposition. To address complex spatial on-surface material mixing, we proposed a data-driven model for predicting the appearance. Furthermore, we take into account the influence of the halftoning pattern to improve the achieved finish quality. Finally, we demonstrated the system's performance and expressiveness by manufacturing several examples with spatially-varying gloss. While each of these steps can be further improved, we believe that the presented workflow serves as a complete set of basic building blocks of a varnish-based appearance fabrication process that can guide future improvements and integration with existing 3D printing systems.

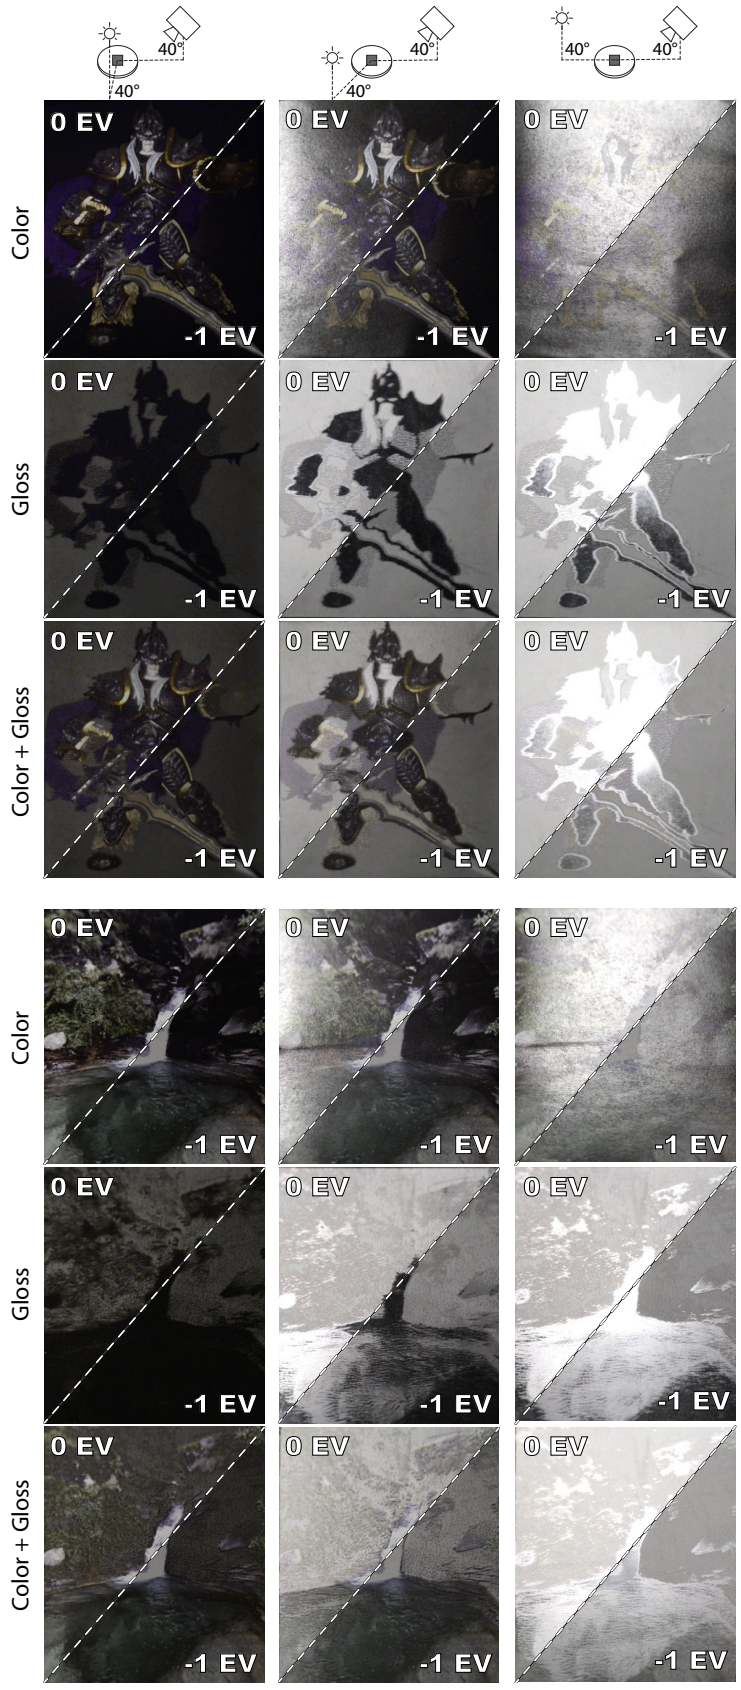

Figure 8.19. Appearance manufacturing with separate fabrication of color (ink-jet printer) and gloss (our device). The final combined appearance manifests both high-resolution color and spatially-varying gloss. Photos are captured at two exposure values (+0, -1) with a still camera and a moving light source. For a full visualization please see the supplementary video.

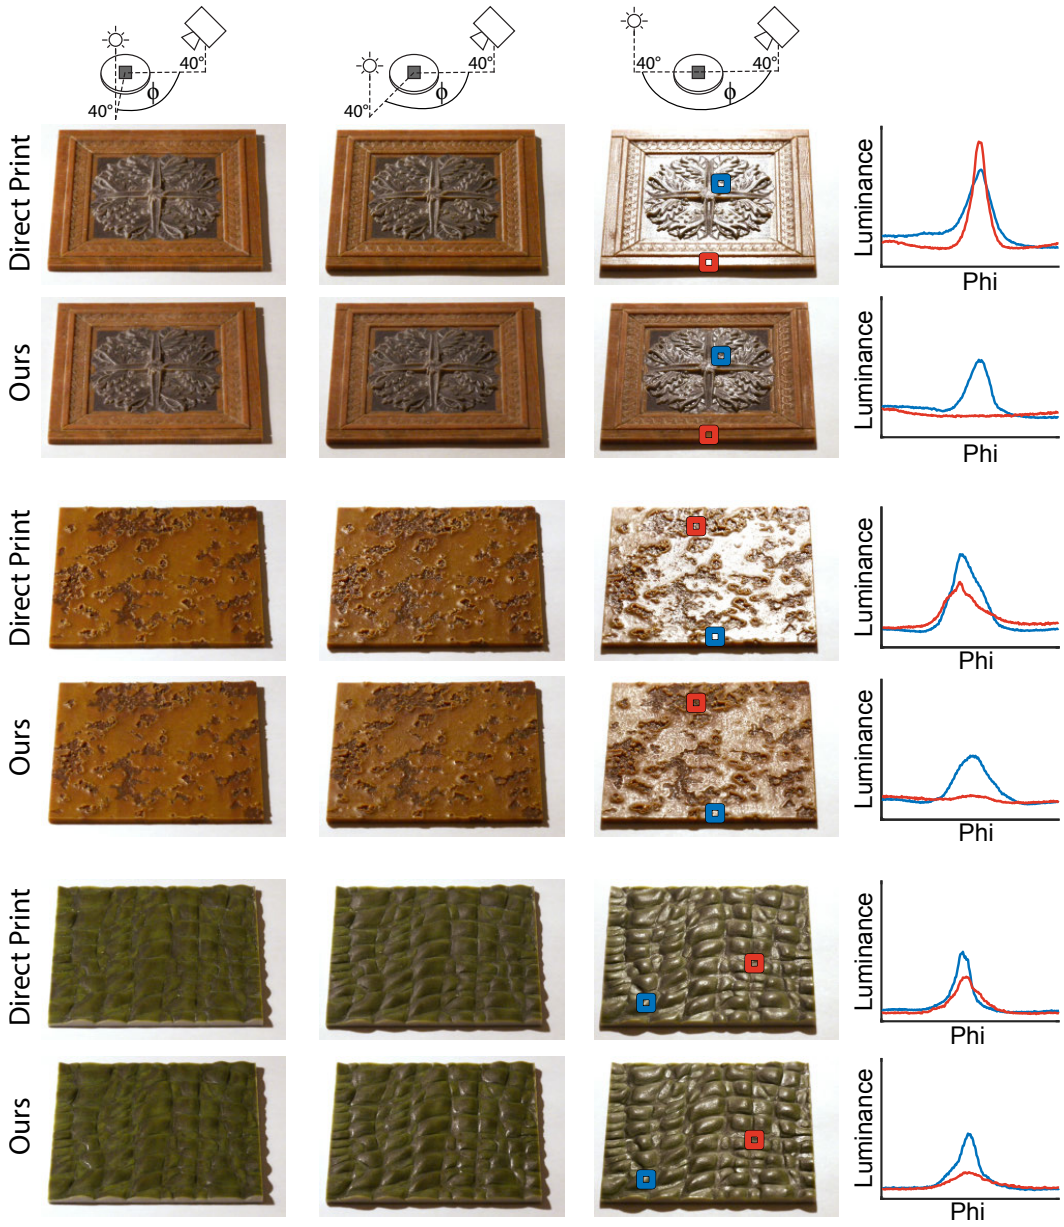

Figure 8.20. Manufactured height-fields without varnish (top) and with halftoned varnish using our system (bottom). Photos are captured with a still camera and moving light source. We plot the luminance of two locations as the light rotates around to showcase the gloss variation achieved by our system. For the full capture please see the supplementary video.

Chapter 9

Conclusions And Future Work

In this dissertation, we propose to leverage the limitations of the human sensorial system to overcome the computational and hardware limits of fabrication, effectively increasing the *apparent* gamut of a 3D printer. We show the feasibility of the idea on the example of fabricating objects with non-linear elastic behavior. The design of elastic objects is constrained both computationally, by having too large design space to find the global optimum, and physically, by the limited amount of materials available on 3D printers. We show that these limits can be bypassed by combining computation with perception. The core of our approach is the construction of a so-called perceptual space in which the Euclidean distance between stimuli corresponds with a perceived difference. To recover such a space we propose a design of psychophysical studies suitable for exploration of haptic properties. We formulate a numerical model by analyzing the outcome of the studies and find that a one-dimensional space is sufficient to explain the perception of compliance. We further analyze several computational models that use the physical properties of the materials to explain the recovered perceptual space and find a good correlation with a model based on the local stiffness of the sample. To generalize the model to arbitrary geometries we rely on numerical simulation of contact between a finger and an elastic object. We demonstrate direct benefits of the model to manufacturing in the ability to perform better material selection, design more intuitive user interfaces, and most importantly compute higher quality reproductions than state-of-the-art methods.

While the case-study of compliance reproduction provides encouraging results it also reveals the weaknesses of the approach that limit the application to more complex problems: expensive perceptual studies, difficulty in finding a computational model that explains a recovered perceptual space, and infeasibly slow numerical simulation for complex physical phenomena. We address these

limitations in the second project to enable the design of digital styli that more closely resemble traditional drawing tools. We propose novel experiment designs that aim to reduce the required number of participants. We support the designs with a new likelihood-based computational model. The proposed model is more robust to missing samples which further improves its performance for studies with sparse sampling. Additionally, the model automatically generates a perceptual space explained by physical measurements, thereby eliminating the tedious manual process of identifying governing physical phenomena of each axis. By using the new study designs and optimization, we construct a perceptual space of drawing tools explained by perceived friction and perceived vibration. To apply our model to optimization of digital designs, we use numerical simulation. Due to the complexity of the governing physical phenomena and the scale at which they occur, we cannot rely on off-the-shelf numerical models. Instead, to accelerate the computation, we propose to use perception-aware coarsening of the numerical model. We exploit the fact that humans can perceive vibratory feedback only up to 500 Hz and design an efficient numerical simulator that estimates the behavior of a digital stylus only in the perceptually relevant range. Additionally, we use an efficient data-driven term to model the complex coupled interaction between a stylus and a drawing substrate. We show applications of these improvements to the design of digital styli that provide haptic feedback more similar to traditional instruments than commercial solutions.

In a followup work, we validate the possibility of increasing the *apparent* gamut of a 3D printer by manufacturing perceptually equivalent solutions instead of direct replicas. We stay in the realm of drawing tools and set our goal to mimic the behavior of traditional drawing instruments with digital styli. This is a challenging problem because the interaction between a drawing instrument and a substrate happens on an order of magnitude finer scale than reproducible on a commercial 3D printer. To bypass this limitation, we leverage our perceptual space of drawing tools. We start by formulating a parametrization of the problem that is composed of the joint interaction between a stylus and a surface. Since we explicitly require to estimate the coupled behavior, we cannot rely on our data-driven simulator. Interestingly, we observed that when building a more exact numerical model the speed of computation approached the manufacturing time. This motivated us to employ a fabrication-in-the-loop optimization. The key challenge in applying such optimization in practice lies in efficient sampling of the design space that takes into account the time complexity that comes with evaluating the objective function by physically manufacturing and measuring samples. To this end, we propose to use a probabilistic surrogate model that gives us confidence bounds on the predicted perceived friction and perceived vi-

bration. We use these confidence estimates to formulate an acquisition function that samples the design space by maximizing the improvement of each sample towards our goal. We demonstrate that this process leads to designs that closely match the haptic feedback of traditional tools while manifesting different geometries and different materials. We verify the quality of our reproduction in blind studies with casual users and a survey with professional artists.

In the last scenario, we apply the methodology to the rapidly growing field of appearance reproduction. Since appearance reproduction is such a young field no standardized method exists for reproducing the full appearance of an object. The main challenge is that to modify the surface reflectance (gloss) the current manufacturing techniques rely on modifying the surface microgeometry or on depositing materials with varying appearance properties. This leads to a coupling of objects' gloss with diffuse (sub-surface) color. To decouple this influence, we propose to create objects where the color information is printed separately at a high resolution, and the gloss is modified in a post-processing step by spatial deposition of varnishes. The varnishes are formulated to be translucent and only modify the surface gloss through diffusive particles. Unfortunately, these particles increase the viscosity of the varnishes to the point that they can not be jetted using commercial inkjet printers. To enable production with such challenging materials, we propose a custom printing device capable of jetting highly viscous materials. We characterize the device from hardware setup, through identifying effects of printing parameters on droplet shape, to gloss gamut estimation and selection of varnish primaries for printing. We support the printing device with a predictive model that estimates the reflectance of varnish mixtures. To minimize the visibility of dithering artifacts, we propose a quality metric inspired by the properties of the human visual system and physical mixing of varnishes on the objects' surface. We show the capabilities of our setup at reproducing objects with spatially-varying color and gloss.

In conclusion, this thesis presents methods for incorporating human perception into the design of objects with desired haptic and visual properties. Our proposed design and evaluation of perceptual studies provide a robust method for estimating the physical phenomena that govern the perceived interactions. The application of perceptual spaces for fabrication hinges on the ability to efficiently estimate the physical properties of digital designs. To this end, we propose to use perception-aware coarsening of the numerical model and a fully data-driven fabrication-in-the-loop optimization. These two more general approaches can find applicability in increasing simulation speed beyond the examples presented in this thesis. More specifically, thanks to its data-driven nature the fabrication-in-the-loop optimization can be employed even for problems that do not rely on

human perception. A particularly interesting area of application is the design of metamaterials where numerical simulation is accelerated by assuming homogenization. However, this assumption often does not hold in practice and, as a result, a significantly more complex numerical model is required. Our fabrication-in-the-loop optimization could be used here to replace the expensive numerical simulation.

Another exciting avenue for future work is incorporating our compliance model with our work on the reproduction of drawing tools. Such a system would be capable of replicating the three primary cues that govern the haptic perception of objects, i.e., friction, vibration, and compliance (Chapter 2). Applications of such a method developed for styli range from reproducing the haptic feedback of softer drawing instruments, e.g., brushes, to mimicking soft tissues investigated by a surgeon during laparoscopic surgery. The system could be further extended to support full touch interaction and enable the reproduction of 3D models with prescribed haptic feedback. The main challenge here lies in designing sufficiently expressive parametrization that manifests a large range of haptic properties, and effectively navigating this high-dimensional space towards the desired haptic response.

Finally, we believe that our methodology can be also used to increase the *apparent* gamut of appearance available on modern manufacturing devices. While many works treat color and gloss reproduction as separate steps these two properties are jointly interpreted by the human visual system. As a result it is possible to affect the perceived color contrast by modifying the gloss of an object [Dalal and Natale-Hoffman, 1999] or to affect the perceived gloss level by tweaking the objects color [Pellacini et al., 2000]. As future work, we plan to utilize our spatially-varying gloss printing setup to generate samples for psychophysical experiments and to study the joint perception of color and gloss under natural light conditions. This will allow us to extend the predictive model with insights from human perception to maximize the *apparent* gamut of appearance reproduction.

Bibliography

3D Systems [2020]. https://www.3dsystems.com/haptics-devices/touch-x. Accessed: 2020-07-11.

Akrour, R., Schoenauer, M. and Sebag, M. [2011]. Preference-based policy learning, *in* D. Gunopulos, T. Hofmann, D. Malerba and M. Vazirgiannis (eds), *Machine Learning and Knowledge Discovery in Databases*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 12–27.

Allaire, G., Bonnetier, E., Francfort, G. and Jouve, F. [1997]. Shape optimization by the homogenization method, *Numerische Mathematik* **76**: 27–68.

Ambrosi, G., Bicchi, A., Rossi, D. D. and Scilingo, E. P. [1999]. The role of contact area spread rate in haptic discrimination of softness, *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, Vol. 1, pp. 305–310 vol.1.

Annett, M., Anderson, F., Bischof, W. F. and Gupta, A. [2014]. The pen is mightier: understanding stylus behaviour while inking on tablets, *Proceedings of Graphics Interface 2014*, Canadian Information Processing Society, pp. 193–200.

Arasan, A., Basdogan, C. and Sezgin, T. M. [2013]. Haptic stylus with inertial and vibro-tactile feedback, *2013 World Haptics Conference (WHC)*, pp. 425–430.

Archard, J. F. [1953]. Contact and rubbing of flat surfaces, *Journal of Applied Physics* **24**(8): 981–988.

Armstrong-Hélouvry, B. [1991]. *Control of Machines with Friction*, Kluwer Academic Publishers, Boston, Massachusetts.

AST [2010]. Standard test method for rubber property - durometer hardness., *Standard*, ASTM International.

Avriel, M. [2003]. *Nonlinear Programming: Analysis and Methods*, Dover Publications.

Baar, T., Samadzadegan, S., Brettel, H., Urban, P. and Segovia, M. V. O. [2014]. Printing gloss effects in a 2.5 d system, *Measuring, Modeling, and Reproducing Material Appearance*, Vol. 9018, International Society for Optics and Photonics, p. 90180M.

Babaei, V. and Hersch, R. D. [2016]. *n*-ink printer characterization with barycentric subdivision, *IEEE Transactions on Image Processing* **25**(7): 3023–3031.

Babaei, V., Vidimče, K., Foshey, M., Kaspar, A., Didyk, P. and Matusik, W. [2017]. Color contoning for 3d printing, *ACM Transactions on Graphics* **36**(4): 124.

Bau, O., Poupyrev, I., Israr, A. and Harrison, C. [2010]. Teslatouch: Electrovibration for touch surfaces, *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology*, UIST '10, ACM, New York, NY, USA, pp. 283–292.

Beck, J. and Prazdny, S. [1981]. Highlights and the perception of glossiness., *Perception & Psychophysics* .

Bensmaïa, S., Hollins, M. and Yau, J. [2005]. Vibrotactile intensity and frequency information in the pacinian system: A psychophysical model, *Attention, Perception, & Psychophysics* **67**(5): 828–841.

Bergmann Tiest, W. M. and Kappers, A. M. [2006]. Analysis of haptic perception of materials by multidimensional scaling and physical measurements of roughness and compressibility, *Acta psychologica* **121**(1): 1–20.

Bertails-Descoubes, F., Cadoux, F., Daviet, G. and Acary, V. [2011]. A nonsmooth Newton solver for capturing exact Coulomb friction in fiber assemblies, *ACM Transactions on Graphics* **30**(1): 6:1–6:14.

Bharaj, G., Levin, D. I. W., Tompkin, J., Fei, Y., Pfister, H., Matusik, W. and Zheng, C. [2015]. Computational design of metallophone contact sounds, *ACM Transactions on Graphics* **34**(6).

Bickel, B., Bächer, M., Otaduy, M. A., Lee, H. R., Pfister, H., Gross, M. and Matusik, W. [2010]. Design and fabrication of materials with desired deformation behavior, *ACM Transactions on Graphics* **29**(3).

Blau, P. J. and Gardner, J. K. [1996]. Tribological characteristics of graded pencil cores on paper, *Wear* **197**(1): 233 – 241.

Blood, B. [2009]. Physics of musical instruments.

Bonilla, E. V., Chai, K. M. and Williams, C. [2008]. Multi-task gaussian process prediction, *in* J. C. Platt, D. Koller, Y. Singer and S. T. Roweis (eds), *Advances in Neural Information Processing Systems 20*, Curran Associates, Inc., pp. 153–160.

Bonneel, N., van de Panne, M., Paris, S. and Heidrich, W. [2011]. Displacement interpolation using lagrangian mass transport, *ACM Transactions on Graphics* **30**(6): 158:1–158:12.

Boundless [2013]. Boundless biology: Somatosensation. `https://courses.lumenlearning.com/boundless-biology/chapter/somatosensation/`.

Britannica, I. E. et al. [1957]. *Encyclopædia britannica*, Encyclopaedia Britannica, Incorporated.

Brunton, A., Arikan, C. A., Tanksale, T. M. and Urban, P. [2018]. 3d printing spatially varying color and translucency, *ACM Transactions on Graphics* **37**(4): 157.

Brunton, A., Arikan, C. A. and Urban, P. [2015]. Pushing the limits of 3d color printing: Error diffusion with translucent materials, *ACM Transactions on Graphics* **35**(1): 4.

Burdea, G. C. [2000]. Haptics issues in virtual environments, *Proceedings Computer Graphics International 2000*, pp. 295–302.

Canudas de Wit, C., Olsson, H., Astrom, K. J. and Lischinsky, P. [1995]. A new model for control of systems with friction, *IEEE Transactions on Automatic Control* **40**(3): 419–425.

Carter, T., Seah, S. A., Long, B., Drinkwater, B. and Subramanian, S. [2013]. Ultrahaptics: Multi-point mid-air haptic feedback for touch surfaces, *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, UIST '13, Association for Computing Machinery, New York, NY, USA, pp. 505–514.

Chadwick, A. C. and Kentridge, R. [2015]. The perception of gloss: A review, *Vision research* **109**: 221–235.

Chae, M. P., Lin, F., Spychal, R. T., Hunter-Smith, D. J. and Rozen, W. M. [2015]. 3D-printed haptic "reverse" models for preoperative planning in soft tissue reconstruction: A case report, *Microsurgery* **35**(2): 148–153.

Chen, D., Levin, D. I. W., Didyk, P., Sitthi-Amorn, P. and Matusik, W. [2013]. Spec2Fab: A reducer-tuner model for translating specifications to 3D prints, *ACM Transactions on Graphics* **32**(4): 135:1–135:10.

Chen, D., Levin, D. I. W., Matusik, W. and Kaufman, D. M. [2017]. Dynamics-aware numerical coarsening for fabrication design, *ACM Transactions on Graphics* **36**(4): 84:1–84:15.

Chen, D., Levin, D. I. W., Sueda, S. and Matusik, W. [2015]. Data-driven finite elements for geometry and material design, *ACM Transactions on Graphics* **34**(4).

Chen, X., Shao, F., Barnes, C., Childs, T. and Henson, B. [2009]. Exploring relationships between touch perception and surface physical properties, *International Journal of Design* **3**: 67–76.

Cho, Y., Bianchi, A., Marquardt, N. and Bianchi-Berthouze, N. [2016]. Realpen: Providing realism in handwriting tasks on touch surfaces using auditory-tactile feedback, *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, UIST '16, ACM, New York, NY, USA, pp. 195–205.

Choi, S. and Tan, H. Z. [2005]. Toward realistic haptic rendering of surface textures, *ACM SIGGRAPH 2005 Courses*, SIGGRAPH '05, ACM, New York, NY, USA.

Chouvardas, V. G., Miliou, A. N. and Hatalis, M. K. [2008]. Tactile displays: Overview and recent advances, *Displays* **29**(3): 185–194.

Cirac, J. I. and Zoller, P. [2012]. Goals and opportunities in quantum simulation, *Nature Physics* **8**(4): 264–266.

Conn, A. R., Gould, N. I. and Toint, P. [1991]. A globally convergent augmented lagrangian algorithm for optimization with general constraints and simple bounds, *SIAM Journal on Numerical Analysis* **28**(2): 545–572.

Connor, C. E., Hsiao, S. S., Phillips, J. R. and Johnson, K. O. [1990]. Tactile roughness: neural codes that account for psychophysical magnitude estimates, *Journal of Neuroscience* **10**(12): 3823–3836.

Cook, R. L. [1986]. Stochastic sampling in computer graphics, *ACM Transactions on Graphics* **5**(1): 51–72.

Cook, R. L. and Torrance, K. E. [1982]. A reflectance model for computer graphics, *ACM Transactions on Graphics* pp. 7–24.

Cooke, T., Kannengiesser, S., Wallraven, C. and Bülthoff, H. H. [2006]. Object feature validation using visual and haptic similarity ratings, *ACM Transactions on Applied Perception (TAP)* **3**(3): 239–261.

Cooke, T., Wallraven, C. and Bülthoff, H. H. [2010]. Multidimensional scaling analysis of haptic exploratory procedures, *ACM Transactions on Applied Perception (TAP)* **7**(1): 7.

Culbertson, H., Unwin, J. and Kuchenbecker, K. J. [2014]. Modeling and rendering realistic textures from unconstrained tool-surface interactions, *IEEE Transactions on Haptics* **7**(3): 381–393.

Curcio, C. A., Sloan, K. R., Kalina, R. E. and Hendrickson, A. E. [1990]. Human photoreceptor topography, *Journal of comparative neurology* **292**(4): 497–523.

Dahl, P. R. [1976]. Solid friction damping of mechanical vibrations, *AIAA Journal* **14**(12): 1675–1682.

Dalal, E. N. and Natale-Hoffman, K. M. [1999]. The effect of gloss on color, *Color Research & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur* **24**(5): 369–376.

Danna, J. and Velay, J.-L. [2015]. Basic and supplementary sensory feedback in handwriting, *Frontiers in Psychology* **6**: 169.

Davison, A. C. and Hinkley, D. V. [1997]. *Bootstrap Methods and their Application*, Cambridge University Press.

de la Rie, E. R., Delaney, J. K., Morales, K. M., Maines, C. A. and Sung, L.-P. [2010]. Modification of surface roughness by various varnishes and effect on light reflection, *Studies in Conservation* **55**(2): 134–143.

Degraen, D., Zenner, A. and Krüger, A. [2019]. Enhancing texture perception in virtual reality using 3d-printed hair structures, *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, Association for Computing Machinery, New York, NY, USA.

Didyk, P. [2012]. *Perceptual Display: Exceeding Display Limitations by Exploiting the Human Visual System*, Doctoral dissertation, Universität des Saarlandes.

Ding, S. and Bhushan, B. [2016]. Tactile perception of skin and skin cream by friction induced vibrations, *Journal of Colloid and Interface Science* **481**: 131 – 143.

Diolaiti, N., Niemeyer, G., Barbagli, F. and Salisbury, J. K. [2006]. Stability of haptic rendering: Discretization, quantization, time delay, and coulomb effects, *IEEE Transactions on Robotics* **22**(2): 256–268.

Dong, Y., Wang, J., Pellacini, F., Tong, X. and Guo, B. [2010]. Fabricating spatially-varying subsurface scattering, *ACM Transactions on Graphics* **29**(4): 62.

Dorfmann, A. and Muhr, A. [1999]. *Constitutive models for rubber*, Vol. 1, CRC Press.

Dudley, J. J., Jacques, J. T. and Kristensson, P. O. [2019]. Crowdsourcing interface feature design with bayesian optimization, *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, Association for Computing Machinery, New York, NY, USA.

Elek, O., Sumin, D., Zhang, R., Weyrich, T., Myszkowski, K., Bickel, B., Wilkie, A. and Křivánek, J. [2017]. Scattering-aware texture reproduction for 3d printing, *ACM Transactions on Graphics* **36**(6): 241.

Elkhuizen, W., Essers, T., Song, Y., Geraedts, J., Weijkamp, C., Dik, J. and Pont, S. [2019]. Gloss, color, and topography scanning for reproducing a painting's appearance using 3d printing, *Journal on Computing and Cultural Heritage (JOCCH)* **12**(4): 1–22.

Embark Studios [2019]. Ttexture synthesis, `https://github.com/EmbarkStudios/texture-synthesis`.

Fechner, G. T. [1860]. *Elemente der Psychophysik*, number v. 1 in *Elemente der Psychophysik*, Breitkopf und Härtel.

Ferwerda, J. A., Pattanaik, S. N., Shirley, P. and Greenberg, D. P. [1996]. A model of visual adaptation for realistic image synthesis, *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pp. 249–258.

Fischer, W. J., Mayr, M., Spirk, S., Reishofer, D., Jagiello, L. A., Schmiedt, R., Colson, J., Zankel, A. and Bauer, W. [2017]. Pulp fines-characterization, sheet formation, and comparison to microfibrillated cellulose, *Polymers* **9**(8).

Fleming, R. W., Dror, R. O. and Adelson, E. H. [2003]. Real-world illumination and the perception of surface reflectance properties, *Journal of vision* **3**(5): 3–3.

Fleming, R. W., Torralba, A. and Adelson, E. H. [2004]. Specular reflections and the perception of shape, *Journal of vision* **4**(9): 10–10.

Fores, A., Ferwerda, J. and Gu, J. [2012]. Toward a perceptually based metric for brdf modeling, *Final Program and Proceedings - IS and T/SID Color Imaging Conference* **CIC 12**: 142–148.

Frazier, P., Powell, W. and Dayanik, S. [2009]. The knowledge-gradient policy for correlated normal beliefs, *INFORMS Journal on Computing* **21**(4): 599–613.

Friedman, R. M., Hester, K. D., Green, B. G. and LaMotte, R. H. [2008]. Magnitude estimation of softness, *Experimental brain research* **191**(2): 133–142.

Fujii, G., Takahashi, M. and Akimoto, Y. [2018]. Cma-es-based structural topology optimization using a level set boundary expressionâĂŤapplication to optical and carpet cloaks, *Computer Methods in Applied Mechanics and Engineering* **332**: 624 – 643.

Garcia, C. W. [2003]. *Drawing for the Absolute and Utter Beginner*, Watson-Guptill.

Garrido-Merchán, E. C. and Hernández-Lobato, D. [2020]. Dealing with categorical and integer-valued variables in bayesian optimization with gaussian processes, *Neurocomputing* **380**: 20–35.

Genecov, A. M., Stanley, A. A. and Okamura, A. M. [2014]. Perception of a haptic jamming display: Just noticeable differences in stiffness and geometry, *Haptics Symposium (HAPTICS), 2014 IEEE*, IEEE, pp. 333–338.

Gescheider, G. A. [2013]. *Psychophysics: the fundamentals*, Psychology Press.

Gkioulekas, I., Xiao, B., Zhao, S., Adelson, E. H., Zickler, T. and Bala, K. [2013]. Understanding the role of phase function in translucent appearance, *ACM Transactions on Graphics* **32**(5): 147.

Goodman, J. M. and Bensmaia, S. J. [2017]. A variation code accounts for the perceived roughness of coarsely textured surfaces, *Scientific reports* **7**: 46699.

Guruswamy, V. L., Lang, J. and Lee, W.-S. [2009]. Modelling of haptic vibration textures with infinite-impulse-response filters, *2009 IEEE International Workshop on Haptic Audio visual Environments and Games*, pp. 105–110.

Hale, K. S. and Stanney, K. M. [2014]. *Handbook of Virtual Environments: Design, Implementation, and Applications*, 2nd edn, CRC Press, Inc., USA.

Hansen, N., Müller, S. D. and Koumoutsakos, P. [2003]. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es), *Evolutionary Computation* **11**(1): 1–18.

Harnoy, A., Friedland, B. and Cohn, S. [2008]. Modeling and measuring friction effects, *IEEE Control Systems* **28**(6): 82–91.

Harper, R. and Stevens, S. [1964]. Subjective hardness of compliant materials, *Quarterly Journal of Experimental Psychology* **16**(3): 204–215.

Hartung, B. and Kersten, D. [2002]. Distinguishing shiny from matte, *Journal of Vision* **2**(7): 551–551.

Hašan, M., Fuchs, M., Matusik, W., Pfister, H. and Rusinkiewicz, S. [2010]. Physical reproduction of materials with specified subsurface scattering, *ACM Transactions on Graphics*, Vol. 29, ACM, p. 61.

Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J. and Tibshirani, R. [2009]. *The elements of statistical learning*, Vol. 2, Springer.

Havran, V., Filip, J. and Myszkowski, K. [2016]. Perceptually motivated brdf comparison using single image, *Computer graphics forum*, Vol. 35, Wiley Online Library, pp. 1–12.

Hayward, V. and Cruz-Hernandez, M. [2000]. Tactile display device using distributed lateral skin stretch, *Proceedings of the haptic interfaces for virtual environment and teleoperator systems symposium*, Vol. 69, ASME, pp. 1309–1314.

Helps, R. and Helps, C. [2016]. Measuring stylus and tablet performance for usability in sketching, *Proceedings of the 5th Annual Conference on Research in Information Technology*, ACM, pp. 19–24.

Hennig, P. and Schuler, C. J. [2012]. Entropy search for information-efficient global optimization, *Journal of Machine Learning Research* **13**(null): 1809–1837.

Herholz, P., Matusik, W. and Alexa, M. [2015]. Approximating free-form geometry with height fields for manufacturing, *Computer Graphics Forum* **34**(2): 239–251.

Ho, Y.-X., Landy, M. S. and Maloney, L. T. [2008]. Conjoint measurement of gloss and surface texture, *Psychological Science* **19**(2): 196–204.

Hochbruck, M. and Ostermann, A. [2010]. Exponential integrators, *Acta Numerica* **19**: 209–286.

Holliins, M., Faldowski, R., Rao, S. and Young, F. [1993]. Perceptual dimensions of tactile surface texture: A multidimensional scaling analysis, *Perception & psychophysics* **54**(6): 697–705.

Hollins, M., Bensmaïa, S., Karlof, K. and Young, F. [2000]. Individual differences in perceptual space for tactile textures: Evidence from multidimensional scaling, *Perception & Psychophysics* **62**(8): 1534–1544.

Hollins, M., Lorenz, F. and Harper, D. [2006]. Somatosensory coding of roughness: The effect of texture adaptation in direct and indirect touch, *Journal of Neuroscience* **26**(20): 5582–5588.

Hoshi, T., Takahashi, M., Iwamoto, T. and Shinoda, H. [2010]. Noncontact tactile display based on radiation pressure of airborne ultrasound, *IEEE Transactions on Haptics* **3**(3): 155–165.

Hullin, M. B., Ihrke, I., Heidrich, W., Weyrich, T., Damberg, G. and Fuchs, M. [2013]. State of the Art in Computational Fabrication and Display of Material Appearance, *Eurographics Annual Conference (STAR)*, Girona, Spain.

Hullin, M. B., Lensch, H. P. A., Raskar, R., Seidel, H.-P. and Ihrke, I. [2011]. Dynamic display of BRDFs, *in* O. Deussen and M. Chen (eds), *Computer Graphics Forum*, Eurographics, Blackwell, Llandudno, UK, pp. 475–483.

Hunt, R. W. G. and Wright, W. D. [1995]. *The reproduction of colour*, Vol. 4, Fountain press England.

Hunter, R. S. [1958]. Photoelectric color difference meter, *J. Opt. Soc. Am.* **48**(12): 985–995.

Hunter, R. S. et al. [1937]. Methods of determining gloss, *NBS Research paper RP* **958**.

Hutmacher, F. [2019]. Why is there so much more research on vision than on any other sensory modality?, *Frontiers in psychology* **10**: 2246.

ICC 1:2004-10 [2003]. Image technology colour management - architecture, profile format, and data structure, *Technical report*, ICC.

Ion, A., Kovacs, R., Schneider, O. S., Lopes, P. and Baudisch, P. [2018]. Metamaterial textures, *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, Association for Computing Machinery, New York, NY, USA.

Israr, A., Choi, S. and Tan, H. Z. [2006]. Detection threshold and mechanical impedance of the hand in a pen-hold posture, *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 472–477.

Israr, A., Tan, H. Z. and Reed, C. M. [2006]. Frequency and amplitude discrimination along the kinesthetic-cutaneous continuum in the presence of masking stimuli, *The Journal of the Acoustical Society of America* **120**(5): 2789–2800.

Jakob, W. [2010]. Mitsuba renderer. http://www.mitsuba-renderer.org.

Johnson, K. O. [2001]. The roles and functions of cutaneous mechanoreceptors, *Current Opinion in Neurobiology* **11**(4): 455 – 461.

Jones, D. R., Schonlau, M. and Welch, W. J. [1998]. Efficient global optimization of expensive black-box functions, *Journal of Global Optimization* **13**(4): 455–492.

Jones, L. A. and Hunter, I. W. [1990]. A perceptual analysis of stiffness, *Experimental Brain Research* **79**(1): 150–156.

Joshi, A., Kale, S., Chandel, S. and Pal, D. K. [2015]. Likert scale: Explored and explained, *Current Journal of Applied Science and Technology* pp. 396–403.

Karnopp, D. [1985]. Computer simulation of stick-slip friction in mechanical dynamic systems, *ASME Journal of dynamic Systems, Measurement and Control* **107**(1): 100–103.

Kato, K., Ishizuka, H., Kajimoto, H. and Miyashita, H. [2018]. Double-sided printed tactile display with electro stimuli and electrostatic forces and its assessment, *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, Association for Computing Machinery, New York, NY, USA.

Kaufman, D. M., Edmunds, T. and Pai, D. K. [2005]. Fast frictional dynamics for rigid bodies, *ACM Transactions on Graphics* **24**(3): 946–956.

Kaufman, D. M., Sueda, S., James, D. L. and Pai, D. K. [2008]. Staggered projections for frictional contact in multibody systems, *ACM Transactions on Graphics* **27**(5): 164:1–164:11.

Kim, S.-C., Israr, A. and Poupyrev, I. [2013]. Tactile rendering of 3d features on touch surfaces, *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, UIST '13, ACM, New York, NY, USA, pp. 531–538.

Klare, S. and Peer, A. [2015]. The formable object: A 24-degree-of-freedom shape-rendering interface, *IEEE/ASME Transactions on Mechatronics* **20**(3): 1360–1371.

Klatzky, R. L. and Lederman, S. J. [2002]. Perceiving texture through a probe, *Touch in virtual environments* pp. 180–193.

Klatzky, R. L. and Lederman, S. J. [2008]. Perceiving object properties through a rigid link, *Haptic Rendering: Algorithms and Applications* **1**: 7–19.

Klatzky, R. L., Lederman, S. J., Hamilton, C., Grindley, M. and Swendsen, R. H. [2003]. Feeling textures through a probe: Effects of probe and surface geometry and exploratory factors, *Attention, Perception, & Psychophysics* **65**(4): 613–631.

Koçak, U., Palmerius, K. L., Forsell, C., Ynnerman, A. and Cooper, M. [2011]. Analysis of the JND of stiffness in three modes of comparison, *Haptic and Audio Interaction Design*, Springer, pp. 22–31.

Koyama, Y., Sato, I., Sakamoto, D. and Igarashi, T. [2017]. Sequential line search for efficient visual design optimization by crowds, *ACM Transactions on Graphics* **36**(4).

Kry, P. G. and Pai, D. K. [2006]. Interaction capture and synthesis, *ACM Transactions on Graphics* **25**(3): 872–880.

Kuschel, M., Di Luca, M., Buss, M. and Klatzky, R. L. [2010]. Combination and integration in the perception of visual-haptic compliance information, *Haptics, IEEE Transactions on* **3**(4): 234–244.

Lafortune, E. P. F., Foo, S.-C., Torrance, K. E. and Greenberg, D. P. [1997]. Non-linear approximation of reflectance functions, *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, ACM Press/Addison-Wesley Publishing Co., USA, pp. 117–126.

Lagae, A., Lefebvre, S., Drettakis, G. and Dutré, Ph. [2009]. Procedural noise using sparse gabor convolution, *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2009)* **28**(3): 54–64.

Lagae, A., Lefebvre, S. and Dutré, P. [2011]. Improving gabor noise, *IEEE Transactions on Visualization and Computer Graphics* .

Lagunas, M., Malpica, S., Serrano, A., Garces, E., Gutierrez, D. and Masia, B. [2019]. A similarity measure for material appearance, *ACM Transactions on Graphics* **38**(4).

Laming, D. and Laming, J. [1992]. F. hegelmaier: On memory for the length of a line, *Psychological Research* **54**(4): 233–239.

Lan, Y., Dong, Y., Pellacini, F. and Tong, X. [2013]. Bi-scale appearance fabrication, *ACM Transactions on Graphics* **32**(4): 145:1–145:12.

Larson, G. W., Rushmeier, H. and Piatko, C. [1997]. A visibility matching tone reproduction operator for high dynamic range scenes, *IEEE Transactions on Visualization and Computer Graphics* **3**(4): 291–306.

Lau, D. L. and Arce, G. R. [2007]. *Modern Digital Halftoning, Second Edition*, CRC Press, Inc., USA.

Lee, J. C., Dietz, P. H., Leigh, D., Yerazunis, W. S. and Hudson, S. E. [2004]. Haptic pen: A tactile feedback stylus for touch screens, *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology*, UIST '04, ACM, New York, NY, USA, pp. 291–294.

Leib, R., Nisky, I. and Karniel, A. [2010]. Perception of stiffness during interaction with delay-like nonlinear force field, *Haptics: Generating and Perceiving Tangible Sensations*, Springer, pp. 87–92.

Leonardis, D., Solazzi, M., Bortone, I. and Frisoli, A. [2015]. A wearable fingertip haptic device with 3 dof asymmetric 3-rsr kinematics, *2015 IEEE World Haptics Conference (WHC)*, pp. 388–393.

Levin, A., Glasner, D., Xiong, Y., Durand, F., Freeman, W., Matusik, W. and Zickler, T. [2013]. Fabricating BRDFs at high spatial resolution using wave optics, *ACM Transactions on Graphics* **32**(4): 144:1–144:14.

Leškowský, P., Cooke, T., Ernst, M. O. and Harders, M. [2006]. Using multidimensional scaling to quantify the fidelity of haptic rendering of deformable objects, *Proceedings of the EuroHaptics 2006 International Conference (EH 2006)*.

Li, D., Levin, D. I., Matusik, W. and Zheng, C. [2016]. Acoustic voxels: Computational optimization of modular acoustic filters, *ACM Transactions on Graphics* **35**(4).

Lichtenauer, M. S., Schuetz, P. and Zolliker, P. [2013]. Interaction improves perception of gloss, *Journal of Vision* **13**(14): 14–14.

Long, B., Seah, S. A., Carter, T. and Subramanian, S. [2014]. Rendering volumetric haptic shapes in mid-air using ultrasound, *ACM Transactions on Graphics* **33**(6).

Lukac, R. [2010]. *Computational Photography: Methods and Applications*, 1st edn, Harry N Abrams, USA.

Luongo, A., Falster, V., Doest, M., Ribo, M., Eiriksson, E., Pedersen, D. and Frisvad, J. [2019]. Microstructure control in 3d printing with digital light processing, *Computer Graphics Forum*, Wiley Online Library.

Malzbender, T., Samadani, R., Scher, S., Crume, A., Dunn, D. and Davis, J. [2012]. Printing reflectance functions, *ACM Transactions on Graphics* **31**(3): 20:1–20:11.

Mannos, J. and Sakrison, D. [1974]. The effects of a visual fidelity criterion of the encoding of images, *IEEE transactions on Information Theory* **20**(4): 525–536.

Marlow, P. J., Kim, J. and Anderson, B. L. [2012]. The perception and misperception of specular surface reflectance, *Current Biology* **22**(20): 1909–1913.

Martínez, J., Dumas, J. and Lefebvre, S. [2016]. Procedural voronoi foams for additive manufacturing, *ACM Transactions on Graphics* **35**(4).

Martínez, J., Skouras, M., Schumacher, C., Hornus, S., Lefebvre, S. and Thomaszewski, B. [2019]. Star-Shaped Metrics for Mechanical Metamaterial Design, *ACM Transactions on Graphics* **38**(4): Article No. 82 :1–13. Special issue, SIGGRAPH 2019.

Martinez, M. O., Morimoto, T. K., Taylor, A. T., Barron, A. C., Pultorak, J. D. A., Wang, J., Calasanz-Kaiser, A., Davis, R. L., Blikstein, P. and Okamura, A. M. [2016]. 3-D printed haptic devices for educational applications, *2016 IEEE Haptics Symposium (HAPTICS)*, pp. 126–133.

Masia, B., Wetzstein, G., Didyk, P. and Gutierrez, D. [2013]. A survey on computational displays: Pushing the boundaries of optics, computation, and perception, *Computers & Graphics* **37**(8): 1012 – 1038.

Massie, T. H., Salisbury, J. K. et al. [1994]. The phantom haptic interface: A device for probing virtual objects, *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems*, Vol. 55, Chicago, IL, pp. 295–300.

Matusik, W., Ajdin, B., Gu, J., Lawrence, J., Lensch, H. P. A., Pellacini, F. and Rusinkiewicz, S. [2009]. Printing spatially-varying reflectance, *ACM Transactions on Graphics* **28**(5): 128:1–128:9.

Matusik, W., Pfister, H., Brand, M. and McMillan, L. [2003]. A data-driven reflectance model, *ACM Transactions on Graphics* **22**(3): 759–769.

Mazzone, A. and Kunz, A. [2005]. Sketching the future of the smartmesh wide area haptic feedback device by introducing the controlling concept for such a deformable multi-loop mechanism, *First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics Conference*, pp. 308–315.

Mead, A. [1992]. Review of the development of multidimensional scaling methods, *Journal of the Royal Statistical Society: Series D (The Statistician)* **41**(1): 27–39.

Meyer, D. J., Peshkin, M. A. and Colgate, J. E. [2016]. Tactile paintbrush: A procedural method for generating spatial haptic texture, *2016 IEEE Haptics Symposium (HAPTICS)*, pp. 259–264.

Michels, D. L., Sobottka, G. A. and Weber, A. G. [2014]. Exponential integrators for stiff elastodynamic problems, *ACM Transactions on Graphics* **33**(1): 7:1–7:20.

Mirtich, B. and Canny, J. [1995]. Impulse-based simulation of rigid bodies, *Proceedings of the 1995 Symposium on Interactive 3D Graphics*, I3D '95, ACM, New York, NY, USA, pp. 181–ff.

Misra, S., Fuernstahl, P., Ramesh, K., Okamura, A. M. and Harders, M. [2009]. Quantifying perception of nonlinear elastic tissue models using multidimensional scaling, *EuroHaptics conference, 2009 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009. Third Joint*, IEEE, pp. 570–575.

Mockus, J. [1989]. *Bayesian Methods for Global Optimization in the Gaussian Case*, Springer Netherlands, Dordrecht, pp. 79–116.

Morovic, J. and Luo, M. R. [2001]. The fundamentals of gamut mapping: A survey, *Journal of Imaging Science and Technology* **45**(3): 283–290.

Motoyoshi, I. and Matoba, H. [2012]. Variability in constancy of the perceived surface reflectance across different illumination statistics, *Vision Research* **53**(1): 30–39.

Myshkin, N. K., Petrokovets, M. I. and Chizhik, S. A. [1998]. Simulation of real contact in tribology, *Tribology International* **31**(1): 79 – 86.

Neal, R. M. [1996]. *Bayesian Learning for Neural Networks*, Springer-Verlag, Berlin, Heidelberg.

Ngan, A., Durand, F. and Matusik, W. [2005]. Experimental analysis of brdf models, *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques*, EGSR '05, Eurographics Association, Goslar, DEU, pp. 117–126.

Nisky, I., Pressman, A., Pugh, C. M., Mussa-Ivaldi, F. A. and Karniel, A. [2011]. Perception and action in teleoperated needle insertion, *IEEE Transactions on Haptics* **4**(3): 155–166.

Nocedal, J. and Wright, S. [2006]. *Numerical optimization*, Springer New York.

Otaduy, M., Jain, N., Sud, A. and Lin, M. [2004]. Haptic rendering of interaction between textured objects.

Pacchierotti, C., Prattichizzo, D. and Kuchenbecker, K. J. [2016]. Cutaneous feedback of fingertip deformation and vibration for palpation in robotic surgery, *IEEE Transactions on Biomedical Engineering* **63**(2): 278–287.

Pacchierotti, C., Sinclair, S., Solazzi, M., Frisoli, A., Hayward, V. and Prattichizzo, D. [2017]. Wearable haptic systems for the fingertip and the hand: Taxonomy, review, and perspectives, *IEEE Transactions on Haptics* **10**(4): 580–600.

Panetta, J., Rahimian, A. and Zorin, D. [2017]. Worst-case stress relief for microstructures, *ACM Transactions on Graphics* **36**(4).

Panetta, J., Zhou, Q., Malomo, L., Pietroni, N., Cignoni, P. and Zorin, D. [2015]. Elastic textures for additive fabrication, *ACM Transactions on Graphics* **34**(4): 135:1–135:12.

Panozzo, D., Diamanti, O., Paris, S., Tarini, M., Sorkine, E. and Sorkine-Hornung, O. [2015]. Texture mapping real-world objects with hydrographics, *Computer Graphics Forum* **34**(5): 65–75.

Pantone [2018]. `https://www.xrite.com/de/blog/tolerancing-part-3`. Accessed: 2020-07-11.

Papas, M., Regg, C., Jarosz, W., Bickel, B., Jackson, P., Matusik, W., Marschner, S. and Gross, M. [2013]. Fabricating translucent materials using continuous pigment mixtures, *ACM Transactions on Graphics* **32**(4): 146.

Paperlike [2020]. `https://paperlike.com/`. Accessed: 2020-07-11.

Pellacini, F., Ferwerda, J. A. and Greenberg, D. P. [2000]. Toward a psychophysically-based light reflection model for image synthesis, *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., pp. 55–64.

Pereira, T., Leme, C. L. A. P., Marschner, S. and Rusinkiewicz, S. [2017]. Printing anisotropic appearance with magnetic flakes, *ACM Transactions on Graphics* **36**(4).

Pereira, T. and Rusinkiewicz, S. [2012]. Gamut mapping spatially varying reflectance with an improved brdf similarity metric, *Computer Graphics Forum* **31**: 1557–1566.

Perez, A. G., Lobo, D., Chinello, F., Cirio, G., Malvezzi, M., San Martín, J., Prattichizzo, D. and Otaduy, M. A. [2015]. Soft finger tactile rendering for wearable haptics, *Proceedings of World Haptics Conference*.
URL: *http://www.gmrv.es/Publications/2015/PLCCMSPO15*

Perez, A. G., Lobo, D., Chinello, F., Cirio, G., Malvezzi, M., San Martin, J., Prattichizzo, D. and Otaduy, M. A. [2017]. Optimization-based wearable tactile rendering, *IEEE transactions on haptics* .

Perret, J. and Vander Poorten, E. [2018]. Touching virtual reality: A review of haptic gloves, *ACTUATOR 2018; 16th International Conference on New Actuators*, pp. 1–5.

Peshkin, M. A., Colgate, J. E., Wannasuphoprasit, W., Moore, C. A., Gillespie, R. B. and Akella, P. [2001]. Cobot architecture, *IEEE Transactions Robotics and Automation* **17**: 377–390.

Pfund, A. [1930]. The measurement of gloss, *JOSA* **20**(1): 23–25.

Piovarči, M., Foshey, M., Babaei, V., Rusinkiewicz, S., Matusik, W. and Didyk, P. [2020]. Towards spatially varying gloss reproduction for 3d printing, *ACM Transactions on Graphics* **39**(66).
URL: *https://doi.org/10.1145/3414685.3417850*

Piovarči, M., Kaufman, D. M., Levin, D. I. W. and Didyk, P. [2020]. Fabrication-in-the-loop co-optimization of surfaces and styli for drawing haptics, *ACM Transactions on Graphics (To Appear)* .

Piovarči, M., Levin, D. I., Rebello, J., Chen, D., Ďurikovič, R., Pfister, H., Matusik, W. and Didyk, P. [2016]. An interaction-aware, perceptual model for non-linear elastic objects, *ACM Transactions on Graphics* **35**(4).

Piovarči, M., Levin, D. I. W., Kaufman, D. M. and Didyk, P. [2018]. Perception-aware modeling and fabrication of digital drawing tools, *ACM Transactions on Graphics* **37**(4).

Piovarči, M., Wessely, M., Jagielski, M., Alexa, M., Matusik, W. and Didyk, P. [2017]. Directional screens, *Proceedings of Symposium on Computational Fabrication '17*, Vol. 35 of *SCF '17*.

Polygerinos, P., Wang, Z., Galloway, K. C., Wood, R. J. and Walsh, C. J. [2015]. Soft robotic glove for combined assistance and at-home rehabilitation, *Robotics and Autonomous Systems* **73**: 135 – 143. Wearable Robotics.

Pongrac, H. [2008]. Vibrotactile perception: examining the coding of vibrations and the just noticeable difference under various conditions, *Multimedia Systems* **13**(4): 297–307.

Poupyrev, I., Okabe, M. and Maruyama, S. [2004]. Haptic feedback for pen computing: Directions and strategies, *CHI '04 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '04, ACM, New York, NY, USA, pp. 1309–1312.

Pressman, A., Karniel, A. and Mussa-Ivaldi, F. A. [2011]. How soft is that pillow? The perceptual localization of the hand and the haptic assessment of contact rigidity, *The Journal of Neuroscience* **31**(17): 6595–6604.

Pressman, A., Welty, L. J., Karniel, A. and Mussa-Ivaldi, F. A. [2007]. Perception of delayed stiffness, *The International Journal of Robotics Research* **26**(11-12): 1191–1203.

Rasmussen, C. E. and Williams, C. K. I. [2005]. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press.

Reinhard, E., Heidrich, W., Debevec, P., Pattanaik, S., Ward, G. and Myszkowski, K. [2010]. *High dynamic range imaging: acquisition, display, and image-based lighting*, Morgan Kaufmann.

Ren, Z., Yeh, H., Klatzky, R. and Lin, M. C. [2013]. Auditory perception of geometry-invariant material properties, *IEEE Transactions on Visualization and Computer Graphics* **19**(4): 557–566.

Repaper [2020]. `https://www.iskn.co/eu/`. Accessed: 2020-07-11.

Robles-De-La-Torre, G. [2006]. The importance of the sense of touch in virtual and real environments, *IEEE MultiMedia* **13**(3): 24–30.

Romanelli, J. R. and Earle, D. B. [2009]. Single-port laparoscopic surgery: an overview, *Surgical endoscopy* **23**(7): 1419–1427.

Romano, J. M. and Kuchenbecker, K. J. [2012]. Creating realistic virtual textures from contact acceleration data, *IEEE Transactions on Haptics* **5**(2): 109–119.

Rouiller, O., Bickel, B., Matusik, W., Alexa, M. and Kautz, J. [2013]. 3D printing spatially varying BRDFs, *IEEE Computer Graphics and Applications* **33**: 48–57.

Saga, S. and Raskar, R. [2012]. Feel through window: Simultaneous geometry and texture display based on lateral force, *SIGGRAPH Asia 2012 Emerging Technologies*, SA '12, Association for Computing Machinery, New York, NY, USA, pp. 1–3.

Salisbury, K., Conti, F. and Barbagli, F. [2004]. Haptic rendering: introductory concepts, *IEEE Computer Graphics and Applications* **24**(2): 24–32.

Samadzadegan, S., Baar, T., Urban, P., Segovia, M. V. O. and Blahová, J. [2015]. Controlling colour-printed gloss by varnish-halftones, *Measuring, Modeling, and Reproducing Material Appearance 2015*, Vol. 9398, International Society for Optics and Photonics, p. 93980V.

Schmidt, R. and Ratto, M. [2013]. Design-to-fabricate: Maker hardware requires maker software, *IEEE Computer Graphics and Applications* **33**(6): 26–34.

Schorr, S. B. and Okamura, A. M. [2017a]. Fingertip tactile devices for virtual object manipulation and exploration, *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, Association for Computing Machinery, New York, NY, USA, pp. 3115–3119.

Schorr, S. B. and Okamura, A. M. [2017b]. Three-dimensional skin deformation as force substitution: Wearable device design and performance during haptic exploration of virtual environments, *IEEE Transactions on Haptics* **10**(3): 418–430.

Schüller, C., Panozzo, D., Grundhöfer, A., Zimmer, H., Sorkine, E. and Sorkine-Hornung, O. [2016]. Computational thermoforming, *ACM Transactions on Graphics* **35**(4).

Schumacher, C., Bickel, B., Rys, J., Marschner, S., Daraio, C. and Gross, M. [2015]. Microstructures to control elasticity in 3d printing, *ACM Transactions on Graphics* **34**(4).

Schwartzburg, Y., Testuz, R., Tagliasacchi, A. and Pauly, M. [2014]. High-contrast computational caustic design, *ACM Transactions on Graphics* **33**(4): 74:1–74:11. Proc. SIGGRAPH 2014.

Serrano, A., Gutierrez, D., Myszkowski, K., Seidel, H.-P. and Masia, B. [2016]. An intuitive control space for material appearance, *ACM Transactions on Graphics* **35**(6).

Shi, L., Babaei, V., Kim, C., Foshey, M., Hu, Y., Sitthi-Amorn, P., Rusinkiewicz, S. and Matusik, W. [2019]. Deep multispectral painting reproduction via multi-layer, custom-ink printing, *ACM Transactions on Graphics* **37**(6): 271.

Sidje, R. B. [1998]. Expokit: A software package for computing matrix exponentials, *ACM Transactions on Mathematical Software (TOMS)* **24**(1): 130–156.

Sigal, L., Mahler, M., Diaz, S., McIntosh, K., Carter, E., Richards, T. and Hodgins, J. [2015]. A perceptual control space for garment simulation, *ACM Transactions on Graphics* **34**(4): 117:1–117:10.

Silverstein, D. A. and Farrell, J. E. [2001]. Efficient method for paired comparison, *Journal of Electronic Imaging* **10**(2): 394–398.

Sitthi-Amorn, P., Ramos, J. E., Wangy, Y., Kwan, J., Lan, J., Wang, W. and Matusik, W. [2015]. Multifab: A machine vision assisted platform for multi-material 3d printing, *ACM Transactions on Graphics* **34**(4).

Skedung, L., Arvidsson, M., Chung, J., Stafford, C., Berglund, B. and Rutland, M. [2013]. Feeling small: Exploring the tactile perception limits, *Scientific reports* **3**: 2617.

Skedung, L., Danerlöv, K., Olofsson, U., Johannesson, C. M., Aikala, M., Kettle, J., Arvidsson, M., Berglund, B. and Rutland, M. W. [2011]. Tactile perception: Finger friction, surface roughness and perceived coarseness, *Tribology International* **44**(5): 505 – 512. Special Issue: ECOTRIB 2009.

Stachowiak, G. W. [2006]. *Wear: materials, mechanisms and practice*, John Wiley & Sons.

Stetten, G., Wu, B., Klatzky, R., Galeotti, J., Siegel, M., Lee, R., Mah, F., Eller, A., Schuman, J. and Hollis, R. [2011]. Hand-held force magnifier for surgical instruments, *in* R. H. Taylor and G.-Z. Yang (eds), *Information Processing in Computer-Assisted Interventions*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 90–100.

Stevens, S. S. [2017]. *Psychophysics: Introduction to its perceptual, neural and social prospects*, Routledge.

Stoll, A. and Cavalcante, M. [2011]. Stick-slip characteristics of leather/artificial leather, *Automotive Buzz, Squeak and Rattle: Mechanisms, Analysis, Evaluation and Prevention* p. 63.

Stratasys [2016]. Stratasys J750 the ultimate full-color multi-material 3D printer, `http://www.stratasys.com/3d-printers/production-series/stratasys-j750`. [Online; Accessed 15-08-2017].

Stucki, P. [1982]. *A Multiple-error Correction Computation Algorithm for Bi-level Image Hardcopy Reproduction*, R. Oldenbourg.
URL: *https://books.google.ch/books?id=iaT_GwAACAAJ*

Sturm, J. F. [1999]. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones, *Optimization Methods and Software* **11–12**: 625–653. Version 1.05 available from `http://fewcal.kub.nl/sturm`.

Sumin, D., Rittig, T., Babaei, V., Nindel, T., Wilkie, A., Didyk, P., Bickel, B., KŘivánek, J., Myszkowski, K. and Weyrich, T. [2019]. Geometry-aware scattering compensation for 3d printing, *ACM Transactions on Graphics* **38**.

Sun, T., Serrano, A., Gutierrez, D. and Masia, B. [2017]. Attribute-preserving gamut mapping of measured brdfs, *Computer graphics forum*, Vol. 36, Wiley Online Library, pp. 47–54.

Tada, M. and Pai, D. K. [2008]. Finger shell: Predicting finger pad deformation under line loading, *2008 Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 107–112.

Takeuchi, Y., Kamuro, S., Minamizawa, K. and Tachi, S. [2012]. Haptic duplicator, *Proceedings of the 2012 Virtual Reality International Conference*, VRIC '12, Association for Computing Machinery, New York, NY, USA.

Tan, H. Z., Durlach, N. I., Shao, Y. and Wei, M. [1993]. Manual resolution of compliance when work and force cues are minimized, *ASME Dyn. Syst. Control Div. Publ. DSC, ASME, NEW YORK, NY,(USA), 1993,* **49**: 99–104.

Tan, H. Z., Pang, X. D. and Durlach, N. I. [1992]. Manual resolution of length, force, and compliance, *Advances in Robotics* **42**: 13–18.

Thurstone, L. L. [1927]. A law of comparative judgment, *Psychological Review* **34**(4).

Tiest, W. M. B. and Kappers, A. M. L. [2009]. Cues for haptic perception of compliance, *Haptics, IEEE Transactions on* **2**(4): 189–199.

Tipsotnaiyana, N., Jarupan, L. and Pechyen, C. [2013]. Printing qualities on inkjet-printed paper from varnish coating agent with rice husk silica particles, *Advanced Materials Engineering and Technology*, Vol. 626 of *Advanced Materials Research*, Trans Tech Publications Ltd, pp. 691–695.

Torres, C., Campbell, T., Kumar, N. and Paulos, E. [2015]. Hapticprint: Designing feel aesthetics for digital fabrication, *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, UIST '15, Association for Computing Machinery, New York, NY, USA, pp. 5830–591.

Treutwein, B. [1995]. Adaptive psychophysical procedures, *Vision research* **35**(17): 2503–2522.

Trowbridge, T. and Reitz, K. P. [1975]. Average irregularity representation of a rough surface for ray reflection, *JOSA* **65**(5): 531–536.

Tymms, C., Gardner, E. P. and Zorin, D. [2018]. A quantitative perceptual model for tactile roughness, *ACM Transactions on Graphics* **37**(5).

Urban, P., Tanksale, T. M., Brunton, A., Vu, B. M. and Nakauchi, S. [2019]. Redefining a in rgba: Towards a standard for graphical 3d printing, *ACM Transactions on Graphics* **38**(3): 1–14.

van Laarhoven, P. J. M. and Aarts, E. H. L. [1987]. *Simulated annealing*, Springer Netherlands, Dordrecht, pp. 7–15.

Vazquez-Leal, H., Castaneda-Sheissa, R., Filobello-Nino, U., Sarmiento-Reyes, A. and Orea, J. S. [2012]. High accurate simple approximation of normal distribution integral, *Mathematical Problems in Engineering* **2012**.

Verrillo, R. and Gescheider, G. [1975]. Enhancement and summation in the perception of two successive vibrotactile stimuli, *Perception and Psychophysics* **18**: 128–136.

Wacom [2020]. `https://estore.wacom.com/en-US/accessories/accessories-nibsc.html`. Accessed: 2020-07-11.

Walker, J. M., Colonnese, N. and Okamura, A. M. [2016]. Noise, but not uncoupled stability, reduces realism and likeability of bilateral teleoperation, *IEEE Robotics and Automation Letters* **1**(1): 562–569.

Walter, B., Marschner, S. R., Li, H. and Torrance, K. E. [2007]. Microfacet models for refraction through rough surfaces, *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, EGSR'07, Eurographics Association, Goslar, DEU, pp. 195–206.

Wandell, B. and Thomas, S. [1997]. Foundations of vision, *Psyccritiques* **42**(7).

Wang, Q. and Hayward, V. [2008]. Tactile synthesis and perceptual inverse problems seen from the viewpoint of contact mechanics, *ACM Transactions on Applied Perception* **5**(2).

Wang, Q., Ren, X., Sarcar, S. and Sun, X. [2016]. Ev-pen: Leveraging electro-vibration haptic feedback in pen interaction, *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces*, ISS '16, ACM, New York, NY, USA, pp. 57–66.

Warburton, T. [2006]. An explicit construction of interpolation nodes on the simplex, *Journal of engineering mathematics* **56**(3): 247–262.

Ward, G. J. [1992]. Measuring and modeling anisotropic reflection, *SIGGRAPH Computational Graphics* **26**(2): 265–272.

Wehner, M., Quinlivan, B., Aubin, P. M., Martinez-Villalpando, E., Baumann, M., Stirling, L., Holt, K., Wood, R. and Walsh, C. [2013]. A lightweight soft exosuit for gait assistance, *2013 IEEE International Conference on Robotics and Automation*, pp. 3362–3369.

Wendt, G., Faul, F. and Mausfeld, R. [2008]. Highlight disparity contributes to the authenticity and strength of perceived glossiness, *Journal of Vision* **8**(1): 14–14.

Weyrich, T., Peers, P., Matusik, W. and Rusinkiewicz, S. [2009]. Fabricating microgeometry for custom surface reflectance, *ACM Transactions on Graphics* **28**(3): 32:1–32:6.

Williams, R. [2015]. Jony Ive interview: The story of the Apple Pencil, `http://www.telegraph.co.uk/technology/apple/11988396/Jony-Ive-interview-The-story-of-the-Apple-Pencil.html`. accessed on: 8 January 2018.

Wills, J., Agarwal, S., Kriegman, D. and Belongie, S. [2009]. Toward a perceptual space for gloss, *ACM Transactions on Graphics* **28**(4): 1–15.

Winfield, L., Glassmire, J., Colgate, J. E. and Peshkin, M. [2007]. T-PaD: Tactile pattern display through variable friction reduction, *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)*, pp. 421–426.

Withana, A., Groeger, D. and Steimle, J. [2018]. Tacttoo: A thin and feel-through tattoo for on-skin tactile output, *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, UIST '18, Association for Computing Machinery, New York, NY, USA, pp. 365–378.

Wu, W.-C., Basdogan, C. and Srinivasan, M. A. [1999]. Visual, haptic, and bimodal perception of size and stiffness in virtual environments, *ASME Dyn. Syst. Control Div. Publ. DSC* **67**: 19–26.

Wyszecki, G. and Stiles, W. S. [1982]. *Color science*, Vol. 8, Wiley New York.

Yoshioka, T., Bensmaïa, S., Craig, J. C. and Hsiao, S.-F. [2007]. Texture perception through direct and indirect touch: An analysis of perceptual space for tactile textures in two modes of exploration, *Somatosensory & Motor Research* **24**(1-2): 53–70.

Yoshioka, T., Gibb, B., Dorsch, A. K., Hsiao, S. S. and Johnson, K. O. [2001]. Neural coding mechanisms underlying perceived roughness of finely textured surfaces, *Journal of neuroscience* **21**(17): 6905–6916.

Yoshioka, T. and Zhou, J. [2009]. Factors involved in tactile texture perception through probes, *Advanced Robotics* **23**(6): 747–766.

Yuan, W., Dong, S. and Adelson, E. H. [2017]. Gelsight: High-resolution robot tactile sensors for estimating geometry and force, *Sensors* **17**(12).

Zhang, J., Fiers, P., Witte, K. A., Jackson, R. W., Poggensee, K. L., Atkeson, C. G. and Collins, S. H. [2017]. Human-in-the-loop optimization of exoskeleton assistance during walking, *Science* **356**(6344): 1280–1284.

Zhang, Y., Yin, C., Zheng, C. and Zhou, K. [2015]. Computational hydrographic printing, *ACM Transactions on Graphics* **34**(4).

Zhu, B., Skouras, M., Chen, D. and Matusik, W. [2017]. Two-scale topology optimization with microstructures, *ACM Transactions on Graphics* **36**(5).

Appendix A

Compliance Measurements

Figure A.1 shows on the left shows the force applied by all participants on a block during the psychophysical study. On the right the corresponding peak-force histogram stored in our database is shown. Figure A.2 contains the simulated force-displacement data for all of our testing objects.

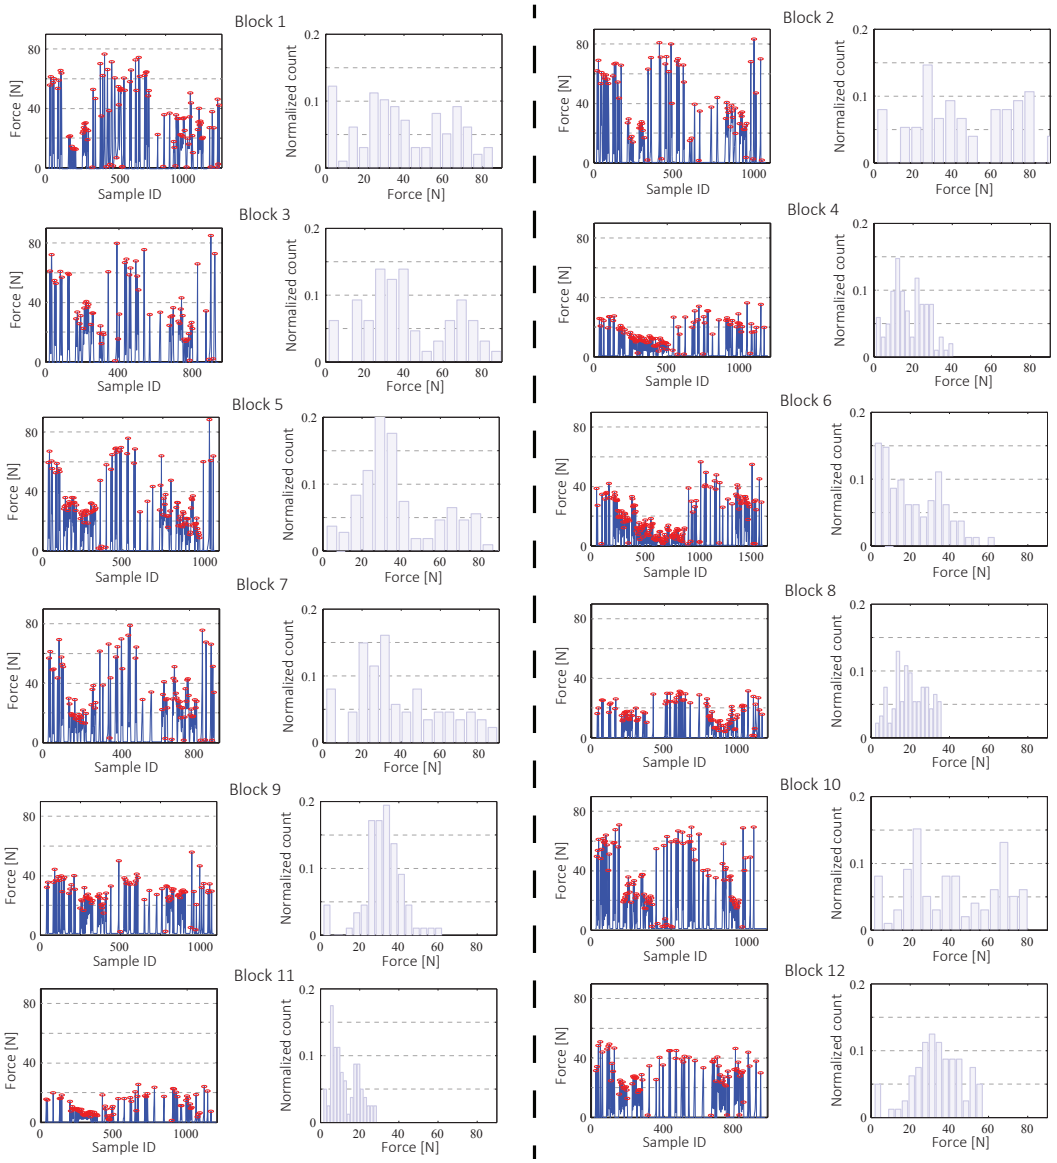

Figure A.1. Left: peak-force data recovered for each block during psychophysical experiment. Right: histograms stored in our database.

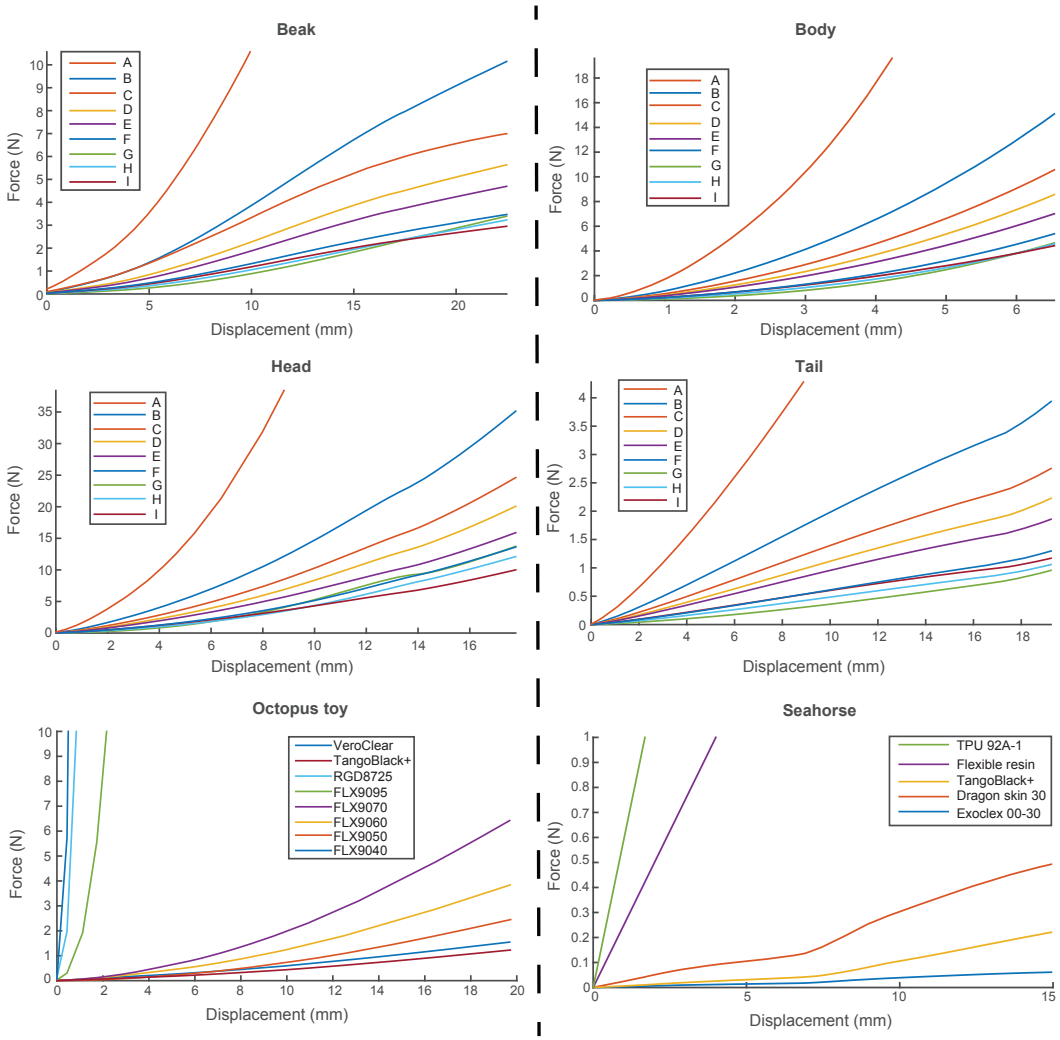

Figure A.2. Force-displacement data recovered for interaction with all fabricated objects.

Appendix B

Drawing Tools Measurements

We present detailed captures of vibration spectograms for traditional drawing tools (Figure B.1), our 3D printed styli (Figure B.2), and commercial styli (Figure B.3), as well as, frictional measurements of traditional (Figure B.4) and digital tools (Figure B.5). In the following text and figures, we use these abbreviations: 2H pencil (p2h), 8B pencil (p8b), ballpoint pen (pen), Stabilo point 88 fine-liner (fine), Stabilo pen 68 (stab), charcoal (charc), 80 g office paper (office), rough artistic paper (rough), smooth stone paper (stone).

For artificial tools we 3D printed nine different material-shape combinations using the Objet260 printer to train our data-driven forcing term. We used three base materials: VeroClear, DM95, and DM85. The shape variation consisted of dome-shaped tips with varying diameter: 1, 2, and 4 mm. Each of these tools was measured on our surfaces: glass, screen protector, office paper, rough paper, and stone paper.

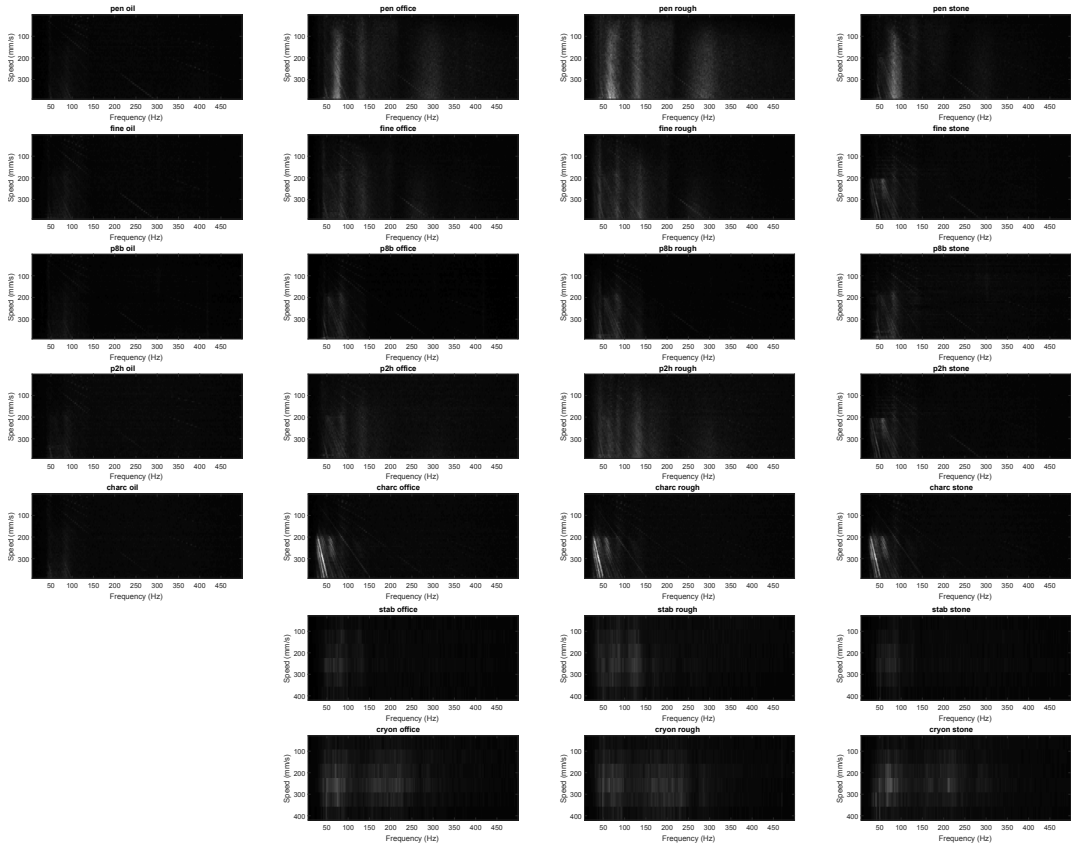

Figure B.1. Vibration spectrograms of real tools recorded using our setup. Each title of the plot consist of name of the tool and the surface. The first column provides measurements of the tools on oiled surface in order to demonstrate the lack of vibrations coming from the measurement setup.

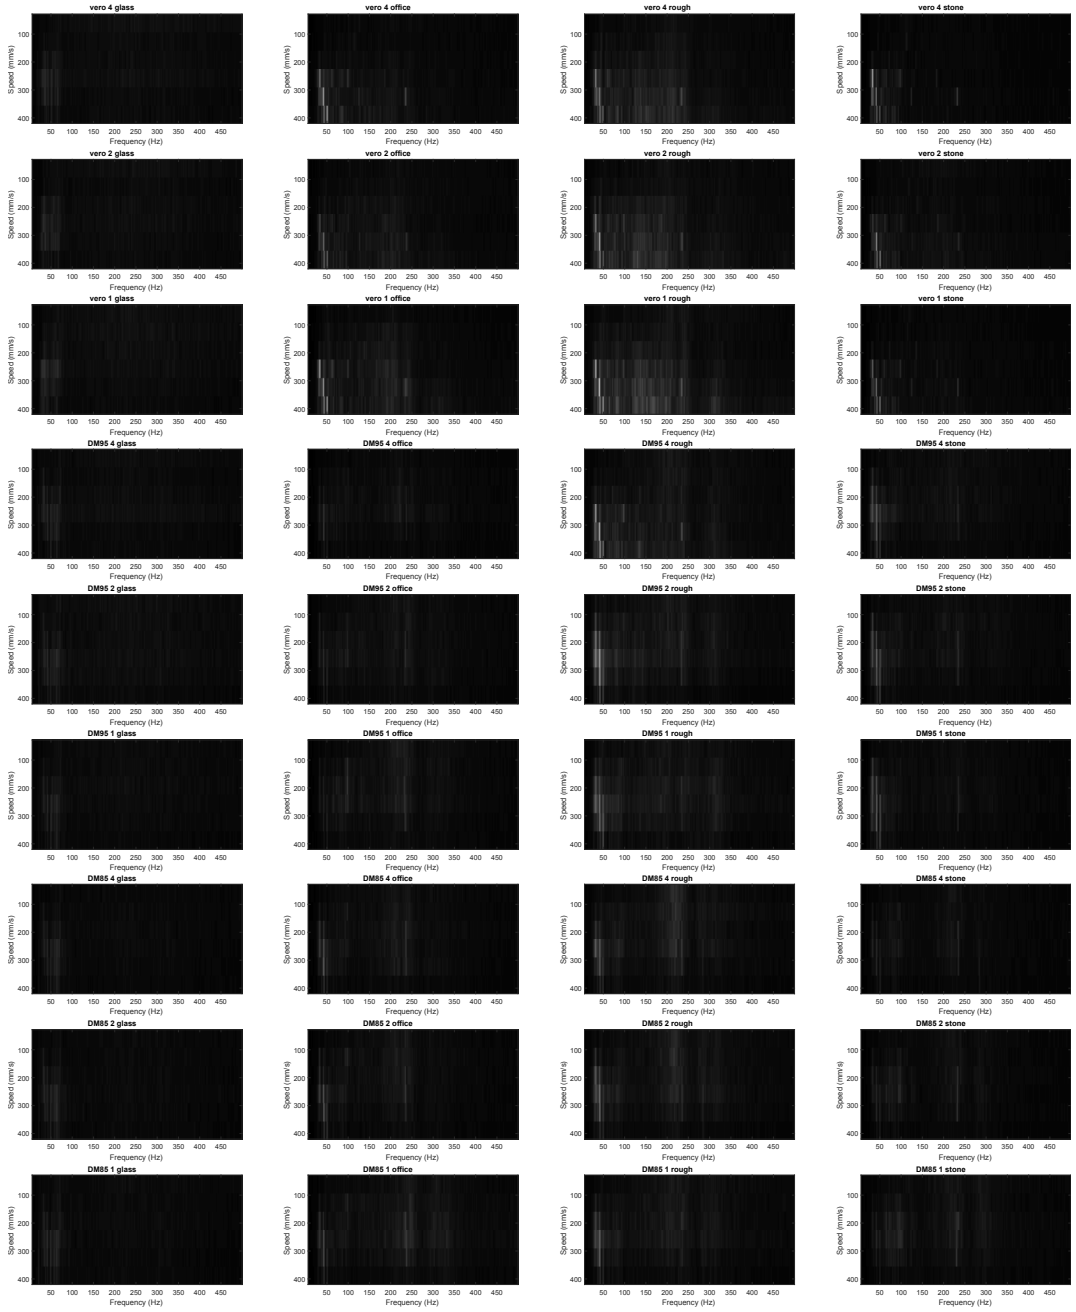

Figure B.2. Vibration spectrograms of 3D printed tools recorded using our setup. The titles of the plots consist of the name of the material, size of the tip, and the name of the surface.

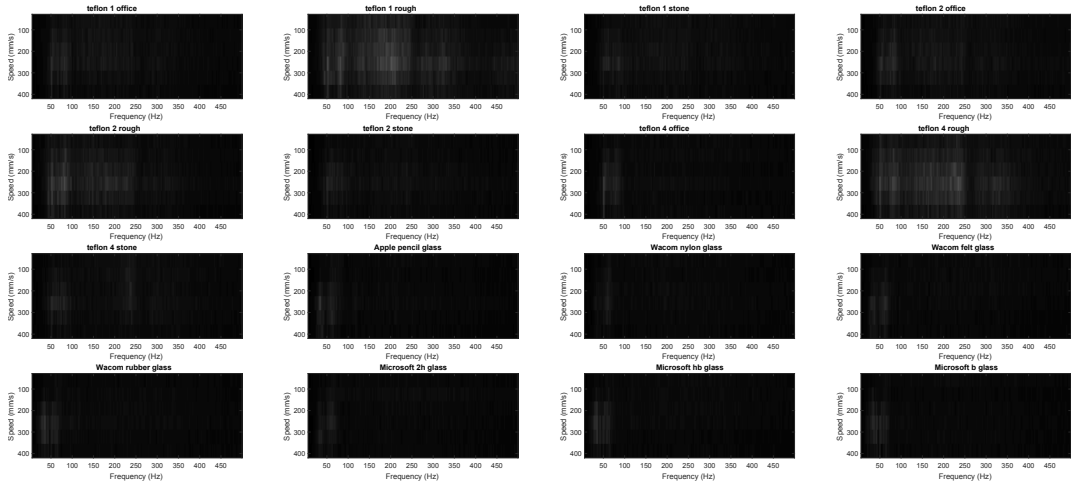

Figure B.3. Vibration spectrograms of 3D printed and commercial tools recorded using our setup. The titles of the plots consist of the name of the material/tool (for commercial products), size/type of the tip, and the name of the surface.

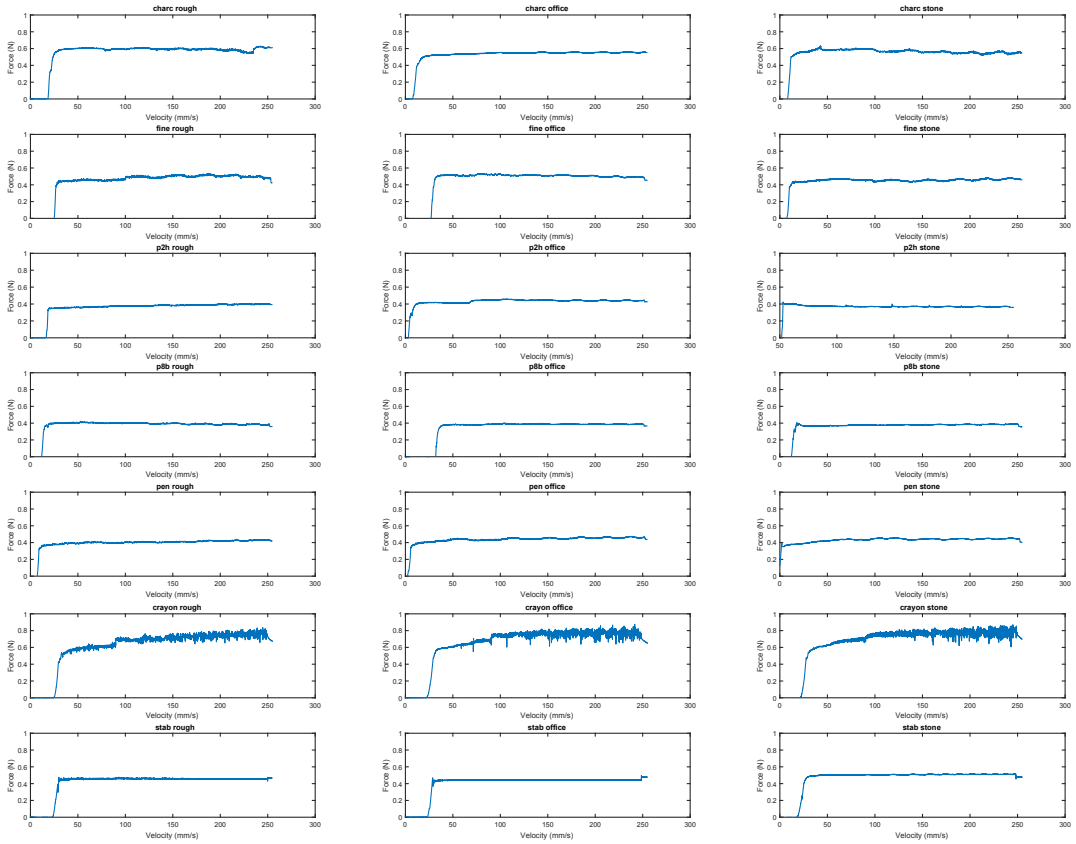

Figure B.4. Friction measurement of drawing tools on various kinds of paper recorded using our setup.

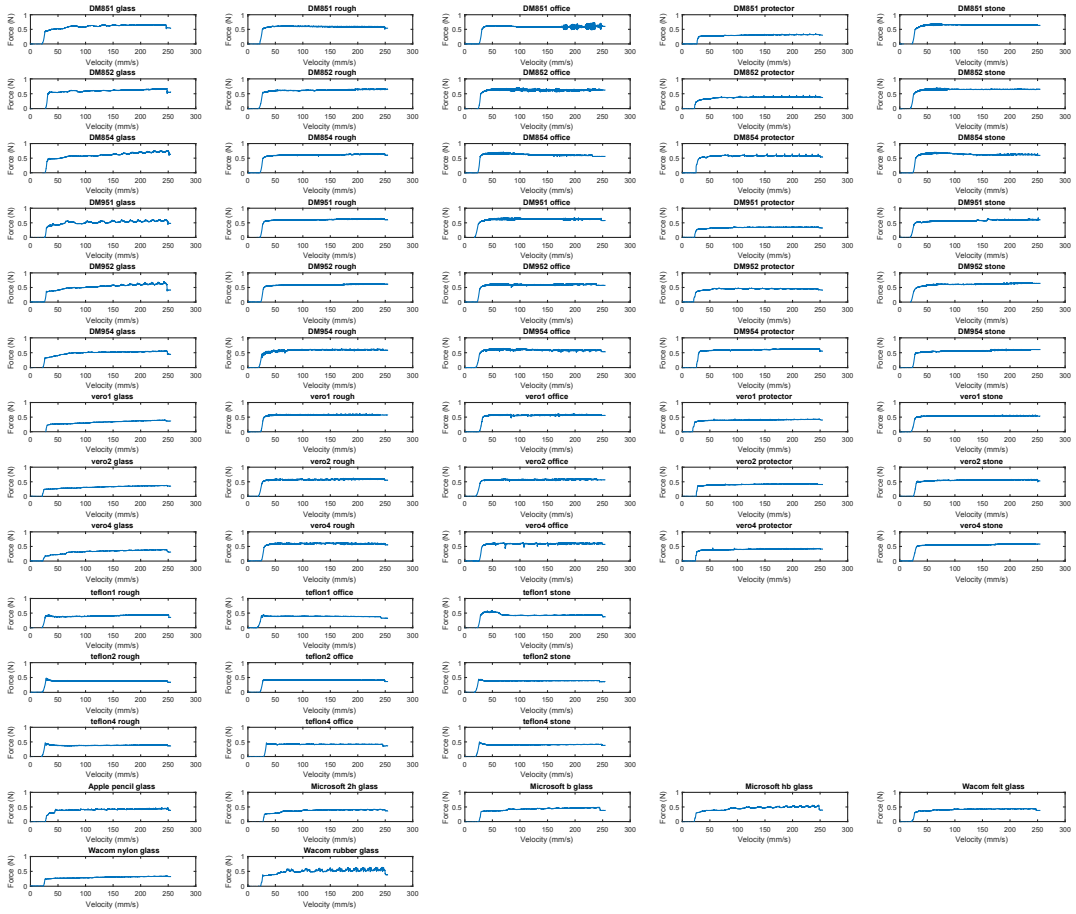

Figure B.5. Friction measurement of 3D printed and commercial styli recorded using our setup.

Appendix C

Stylus Simulation Validation

We run a leave-one-out cross-validation of our simulation. We consider our 3D printed designs on our three paper substrates and a fixed velocity giving us a total of 27 optimizations. In each optimization, one sample was left out and then reconstructed using data optimized for other samples. These left out samples are collected in Figure C.1.

We validated our spectrogram interpolation by 3D printing two surfaces. Each surface was measured at various velocities and then reconstructed using our full data-driven exponential Euler pipeline. We consider two cases. In the first case, we interpolate a mixture of 25% VeroClear and 75% DM85 on office paper with a tip size of 1 mm (Figure C.2 top). In the second case, we simulate a mixture of 50% VeroClear and 50% DM85 with a tip size of 4 mm (Figure C.2 bottom). Both mixtures were measured on our turntable on an office paper (Figure C.2 blue lines) and are compared to our simulator (Figure C.2 red lines). We also compare to not interpolated forcing terms (Figure C.2 gray lines).

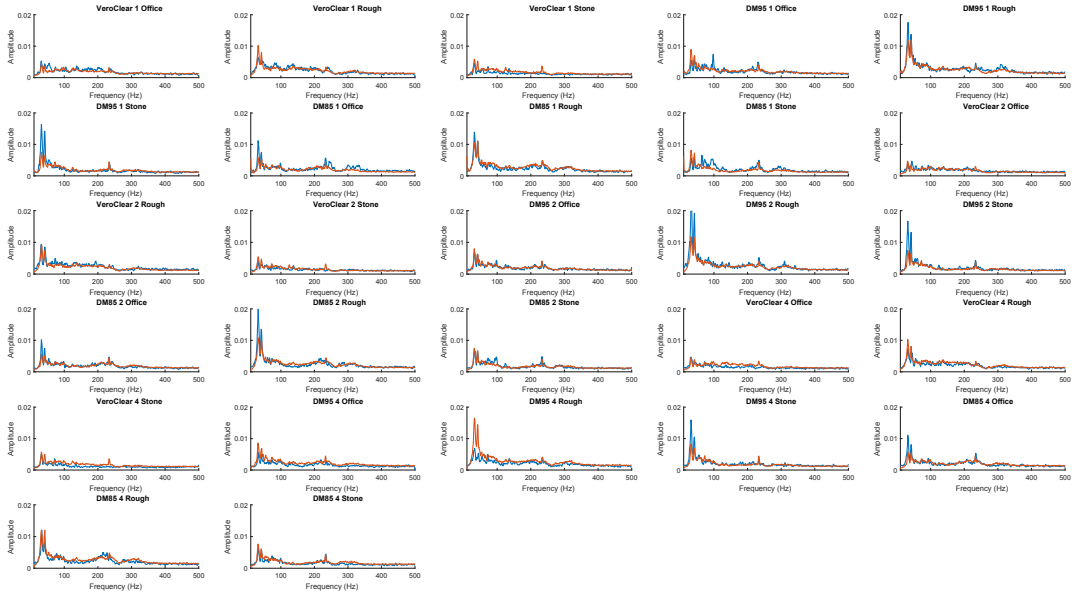

Figure C.1. Full leave-one-out cross-validation of our data-driven forcing term.

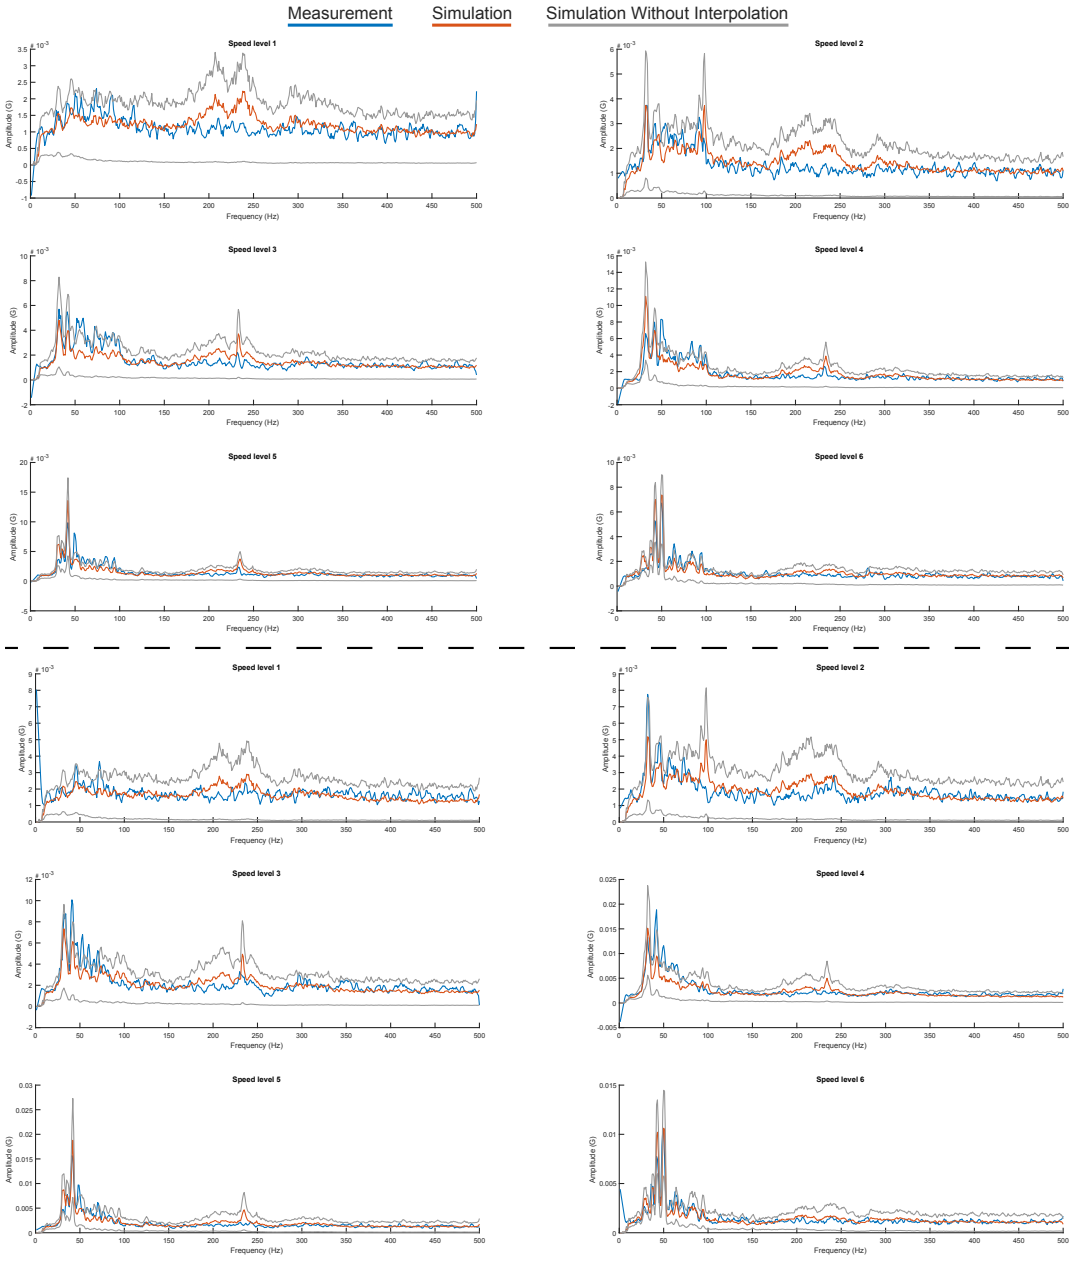

Figure C.2. Interpolated forcing term evaluation. We predicted the vibration response of an interpolated forcing term using our simulator (red line) and compare it to measurements on our turn-table (blue line). We also compare it to the simulations without interpolation of forcing terms (gray lines).

Appendix D

Survey With Professional Artists

We present a detailed transcript of the survey conduct with professional artists.

D.1 Artist 1

What is your current occupation?

2D/3D artist for a gaming studio.

What kind of tools did you use in the past e.g. charcoals, pencils, digital styli?

Pencils, ballpoint pens, charcoal, chalk, pastels, oil colors, markers, different thickness of papers, watercolor papers, soft and hard papers. Then Genius tablet and afterwards Huyon tablet with a screen, Wacom Intuous and iPad Pro with a paperlike.

How often do you use traditional drawing tools? What kind of traditional drawing tools do you use?

Every other day. Usually a pencil or a marker and a piece of paper.

How often do you use digital drawing tools? What kind of digital drawing tools do you use? (manufacturer/display/tips/any extra modifications, e.g., custom surface)

Every other day. Usually start the work on paper and then mirror it in digital. I use a Wacom Intuous tablet.

When/for what tasks do you use traditional tools?

Sketching of ideas. It is faster, more portable. Once it is done we copy it in digital. I have an iPad with paperlike which I use often but there is something which makes me gravitate towards traditional tools. Paperlike is better but does not provide the feeling of drawing with an instrument. I do not notice this when writing but when drawing the pen is less slippery and provides more control and has more pleasant surface for drawing.

When/for what tasks do you use digital tools?

For producing art in the company which needs to be modifiable and repeatable. Wanted to switch to all digital but keep returning to traditional tools.

Why you do not use only traditional tools?

Need to create art for the job.

Why you do not use only digital tools?

Personal preference of traditional tools when possible.

What, in your opinion, are the advantages and disadvantages of traditional tools?

Cheaper, you can take it with you, but does not have undo. Eraser just destroys the paper. You can feel the instrument and control the strokes by using different pressure or orientation.

What, in your opinion, are the advantages and disadvantages of digital tools?

Better higher quality output, better color management, easier to mix. Can use layers and opacity to modify the same piece later.

How important is the feel of the materials when drawing? Why?

I think it is important because you have different feeling in hand based on the strokes. Depends how you can use the materials.

Did you ever notice the difference in haptic feedback between traditional and digital drawing tools? If yes did it ever impact the way you use the drawing tools?

Yes and as a result I keep returning to traditional drawing tools.

Are you using any means of customizing the haptic feedback of the digital tools? How do you customize your tools? Do you think you have sufficient flexibility in the customization? Do you consider the feel of styli in your choice of tool?

I am using paperlike but otherwise no. I did not search specifically for means of adjusting the feedback but I do not think they exists otherwise one of my colleagues would mention it. I would definitely consider the feel of the stylus when picking a new tablet.

Do you think it would be beneficial if the digital drawing tools could provide the same feel as the traditional tools? Why?

It would be awesome. Because digital drawing has way more advantages than disadvantages. But the missing feedback which also makes some people stay with traditional media could transfer to digital. Which would make the workflow easier and eliminate the scanning, sending, and redrawing part.

A session with traditional tools, digital, and our reproductions

Uses 8B pencil on stone paper The paperlike gives you a bit of scratching which is missing on glass surface. The Wacom is not very similar to the traditional tools. I liked the surface. It does not feel exactly the same but you can definitely feel the paper structure. It resembles a marker more than like a pencil due to the sound. Also the tip is less sharp than the pencil. It makes me want to fill in shapes with color I think it would be great for that. I would really have a feeling like using an actual tool.

Uses ballpoint pen on rough paper This is really good, I have a feeling this is like the ballpoint pen. When I compare it with a tablet the difference is really obvious and I can feel the structure that is missing on my digital devices. I would try it out on my tablet. I think it could actually make me use only my tablet.

Uses charcoal on rough paper The sample is very similar to charcoal. You can feel the charcoal from it. It has a lot of resistance from motion. I do not even have to try the digital tools to know they do not feel like that.

Uses 2H pencil on stone paper This has more structure than the reference pencil. You can feel it is a pencil. It is more similar to the reference when using the sharp edge. I think with this I would mind less the use of classical tablets. I would be able to feel the tool and estimate where I draw without looking directly which is a problem for many people. Here you would feel the structure and be able to draw better. Now I would like to draw only on such surfaces.

What are the main differences between the three groups?

The digital tools always feel the same. I like how these mimic the traditional tools.

What do you think about the new reproductions of the traditional drawing tools?

I like that I can feel the structure and have the feeling that I know what I draw without looking.

What do you think about the haptic feedback provided by our tools comparing to the traditional and the digital tools?

You can feel the pencil reproductions are not 100% but are approaching the real tools. They are much more natural.

What is still missing?

I am unsure, maybe mostly the mismatch in sounds.

Would you prefer using a tablet with our enhanced haptic response? Why?

I would prefer to use your drawing tools. They are much more natural and you have a feeling you can draw with them. Drawing and writing would get more precise with tools that feel more similar to traditional instruments.

What are the drawbacks of our tools?

You still have to change tools. I would not like to constantly swap tips and surfaces. Would be good to have an efficient method to change. Alternatively it would be possible to have multiple pens with preinstalled tip.

How could you imagine using these tools?

I would use them in a full pipeline. The pencil replica I would do a sketch, and with the pen I would highlight the edges. Finally, with the 8B replica I would fill in the shape due to the larger tip area. I would likely not use the charcoal reproductions since I also do not like to use charcoal.

D.2 Artist 2

What is your current occupation?

Studying game art at a university.

What kind of tools did you use in the past e.g. charcoals, pencils, digital styli?

Pastels, pencils, pens, oil/acrylic colors, scratching into clay. Genius and later Wacom tablet.

How often do you use traditional drawing tools? What kind of traditional drawing tools do you use? 5 days of a week. Most often I use a pencil How often do you use digital drawing tools? What kind of digital drawing tools do you use? (manufacturer/display/tips/any extra modifications, e.g., custom surface)

Also 5 times a week most often using Wacom Intuous tablet.

When/for what tasks do you use traditional tools?

For longer art pieces, ideation. Also at school for studies of human body drawing.

When/for what tasks do you use digital tools?

For school and work. Particularly, when the finished drawing needs to be shared or produced quickly.

Why you do not use only traditional tools?

Digital tools have many advantages and I would like to get more accustomed to them.

Why you do not use only digital tools?

The traditional tools feel more natural which helps me concentrate on the task. When using digital tools I get distracted with feeling hungry, cold, etc. With traditional tools I do not think about such issues as much.

What, in your opinion, are the advantages and disadvantages of traditional tools?

They are cheaper and easier to use. When you press with a pencil you know what kind of stroke you will get.

What, in your opinion, are the advantages and disadvantages of digital tools?

Digital tools provide undo, layers, color mixing which are more convenient to get the desired color. Also digital pieces are often wanted in the industry.

How important is the feel of the materials when drawing? Why?

Not very important for drawing itself but important for concentration. With traditional tools I can concentrate on the piece but with digital I switch focus from feeling they are just not right and eventually get distracted.

Did you ever notice the difference in haptic feedback between traditional and digital drawing tools? If yes did it ever impact the way you use the drawing tools?

Yes I noticed the difference. Traditional tools behave in different ways based on tool which then impacts how I use the tool. But digital feel the same.

Are you using any means of customizing the haptic feedback of the digital tools? How do you customize your tools? Do you think you have sufficient flexibility in the customization? Do you consider the feel of styli in your choice of tool?

At the moment I do not use anything to customize my drawing tools. I use the software sensitivity to adjust the stylus.

Do you think it would be beneficial if the digital drawing tools could provide the same feel as the traditional tools? Why?

Yes and no. I do like the haptic feedback of traditional tools and it helps me concentrate. But digital tools can be more than traditional so why limit ourselves to realistic feelings? Maybe there is some way they can be different from traditional but better suited for drawing.

A session with traditional tools, digital, and our reproductions

Uses 8B pencil on stone paper The pencil feels very smooth and light to draw with. Your reproduction is very similar. It also glides but has a feeling of a bit of structure. The apple pencil feels like rubbing on glass similarly to Wacom rubber nib. The nylon nib has way too little resistance. The paperlike is quite similar but you can feel the pencil has more structure. The Wacom felt nib has similar resistance but even less structure than paperlike.

Uses ballpoint pen on rough paper The pen has a very clear surface structure which I can feel with your reproduction. I can feel the pen jumping up and down. Also it is very close in terms of resistance. When I close my eyes I can feel a pen. The digital counterparts do not have such pronounced surface. Paperlike would be closer to a pencil than the pen and other alternatives feel very smooth.

Uses 2H pencil on stone paper The reproduction feels very similar to the pencil. The sound is much closer to a marker and the surface structure is more pronounced. Similar to before the closest to pencil in terms of digital tools is paperlike. But here the difference in texture is even more pronounced that your reproduction feels more like it has texture than paperlike.

Uses charcoal on rough paper The charcoal has a lot of drag. It has much more resistance than other digital tools. The reproduction has the same feeling

of making real lines. But the tool feels more rigid and less hollow than a charcoal (hollow as in charcoal would wear down very quickly when pressing but our tool would not). In terms of resistance your tool is much more similar than alternatives.

What are the main differences between the three groups?

The feeling of texture. You can feel the texture in traditional tools and in your tools. Whereas the digital tools feel flat and affect mostly the resistance.

What do you think about the new reproductions of the traditional drawing tools?

I like that I can draw the lines in a similar way that with the traditional tools.

What do you think about the haptic feedback provided by our tools comparing to the traditional and the digital tools?

Digital tools feel very different from traditional. Your tools are coming close.

What is still missing?

The sound is different from traditional tools. The pencil reproductions sounded like markers.

Would you prefer using a tablet with our enhanced haptic response? Why?

I would try it. Perhaps it would help with my concentration thanks to being closer to traditional tools.

What are the drawbacks of our tools?

There might be too many options. I would try all of them but likely end up with one or two.

How could you imagine using these tools?

—

D.3 Artist 3

What is your current occupation?

3D artist in game development company.

What kind of tools did you use in the past e.g. charcoals, pencils, digital styli?

I used charcoal, chalk, pencil, marker and primarily a ballpoint pen. From digital tools I used Wacom or a mouse. At home I have an iPad with paperlike.

How often do you use traditional drawing tools? What kind of traditional drawing tools do you use? About once a week. I use a small notebook and a pen. How often do you use digital drawing tools? What kind of digital drawing tools do you use? (manufacturer/display/tips/any extra modifications, e.g., custom surface)

I use them daily for work. Mainly Wacom Cintique.

When/for what tasks do you use traditional tools?

Mainly for free time drawing. I do not use it for work so I am not investing my time to improve my skillset.

When/for what tasks do you use digital tools?

I use primarily a mouse for 3D modelling but use stylus in Maya. It depends on the software tool which one feels better to use. I am using mainly the Cintique tablet and the eye-hand coordination is much better.

Why you do not use only traditional tools?

I need digital tools for work.

Why you do not use only digital tools?

I use traditional tools because they are cheaper. I am less worried they would break, got stolen, or run out of battery.

What, in your opinion, are the advantages and disadvantages of traditional tools?

They are small, compact and always with you. If you do not have a display tablet then correlation of hand and eye is for free. It is also taken as a learning step that one should start with traditional tools. Though I think the digital tools evolved enough the one could start drawing with them. They also provide more range in terms of working area.

What, in your opinion, are the advantages and disadvantages of digital tools?

Charging battery, fragility, portability. From advantages you can set brushes, colors, undo and they work will with some 3D software.

How important is the feel of the materials when drawing? Why?

Very important. I need a specific pen type for drawing which is very thin. I do not draw or write very well with a larger size tool. I would specifically seek out a pen that works well for me.

Did you ever notice the difference in haptic feedback between traditional and digital drawing tools? If yes did it ever impact the way you use the drawing tools?

Yes, there is a difference when using the pencil with different pressure I have different wear and different types of strokes. You can also use it on a side to have varying contact area. The digital tablets have this but one needs to get used to it. But it does not affect me much since I mainly work with 3D.

Are you using any means of customizing the haptic feedback of the digital tools? How do you customize your tools? Do you think you have sufficient flexibility in the customization? Do you consider the feel of styli in your choice of tool?

I was not changing the nibs but for my iPad I have paperlike and it is very noticable. Makes the surface feel much closer to paper. It more familiar and natural to me but I do not use anything for my work tablet just the base configuration. I would rather compare the ergonomics of the tablet, the quality of the tracking, and available software. I think I could adjust to a different stylus.

Do you think it would be beneficial if the digital drawing tools could provide the same feel as the traditional tools? Why?

For me personally it does not matter. It might be good that it is not the same. Because there is a large variety of traditional tools and it would be unclear which is best to use.

A session with traditional tools, digital, and our reproductions

Uses 2H pencil on stone paper It is quite similar to the pencil. The sound reminds me more of a marker. When using the noise cancelling headphones the feeling is very similar to the pencil. It has more structure than paperlike. The wacom styli are flat in comparison. It is a shame you can not see the stroke. I would be interested to see how the structure affects the generated lines and if they look more pencil like.

Uses charcoal on rough paper Replicating charcoal is very challenging. The orientation affects the properties a lot. When I use it on the sharp side it glides very easily unlike your reproduction. But when I use it on the larger cross section it has significant drag and is similar to your reproduction. It does not capture the wear but is more similar than other digital tools.

Uses 8B pencil on stone paper Very smooth feeling the pencil almost glides on the surface. It is very similar to the paperlike but you can feel a bit more texture. Overall very close the pencil, paperlike, and your solution. The sound again reminds a bit more of a marker than a pencil.

Uses ballpoint pen on rough paper This feels very much like a sharp ballpoint pen. Feels very close to the one I use at home for drawing. You can feel the substrate when drawing and the small tip. I would be interested in trying this one long term.

What are the main differences between the three groups?

Your tools have stronger feeling of surface compared to digital tools.

What do you think about the new reproductions of the traditional drawing tools?

There are still something missing to make them perfect. But they are close much closer than digital tablets.

What do you think about the haptic feedback provided by our tools comparing to the traditional and the digital tools?

Feels like a step in the right direction but for a complete reproduction would need some work.

What is still missing?

Mainly the sound difference and the missing wear.

Would you prefer using a tablet with our enhanced haptic response? Why?

For 3D I prefer smooth surfaces for higher production speed. But for free time drawing I would test them to see how they perform long term.

What are the drawbacks of our tools?

—

How could you imagine using these tools?

You could use the large tip size for filling large areas of the piece.

D.4 Artist 4

What is your current occupation?

Student and now hobby artist.

What kind of tools did you use in the past e.g. charcoals, pencils, digital styli?

Pencils, charcoal, pens for learning and now mostly moved to digital using Wacom Intuos.

How often do you use traditional drawing tools? What kind of traditional drawing tools do you use? Less than I would like about once a week. Mainly a pencil and a pen. How often do you use digital drawing tools? What kind of digital drawing tools do you use? (manufacturer/display/tips/any extra modifications, e.g., custom surface)

Couple of times per week I use the Wacom tablet to draw something.

When/for what tasks do you use traditional tools?

Quick sketches and concepts.

When/for what tasks do you use digital tools?

For full drawings with coloring. Lately also more for sketching to get more practice.

Why you do not use only traditional tools?

The advantages of digital tools coupled with the ability to not have to buy materials is and be able to draw with many colors.

Why you do not use only digital tools?

I do not like the smoothness of digital tools. Specially the glass ones where I have the feeling of the glass. This makes me want to use the tools less and keeps me returning to traditional media.

What, in your opinion, are the advantages and disadvantages of traditional tools?

I can add character to the piece. Each stroke with a tool is unique and can not be reproduced exactly. Also has synchronization with eye that is not there with my tablet.

What, in your opinion, are the advantages and disadvantages of digital tools?

Digital tools make much easier coloring, no need to own all of the individual colors and have undo functionality.

How important is the feel of the materials when drawing? Why?

The feel of the materials is very important. the different textures of drawing tools lead to different patterns of each stroke. This adds character to the piece.

Did you ever notice the difference in haptic feedback between traditional and digital drawing tools? If yes did it ever impact the way you use the drawing tools?

Yes I noticed it. The main impact for me is that I avoided glass surfaces even though hand eye synchronization is nice to have.

Are you using any means of customizing the haptic feedback of the digital tools? How do you customize your tools? Do you think you have sufficient flexibility in the customization? Do you consider the feel of styli in your choice of tool?

No at the moment I do not use any customization.

Do you think it would be beneficial if the digital drawing tools could provide the same feel as the traditional tools? Why?

Yes. If the glass tablets feel better I could use them. It would be also interesting to have strokes more similar to traditional media.

A session with traditional tools, digital, and our reproductions

Uses ballpoint pen on rough paper You can feel the vibrations when drawing. The rough surface and pen vibrating up and down as you go over the surface. The feeling is very close but there is something different. I think the ball in the pen is missing which makes the pen roll more than your tool. The digital tablets are nowhere near in terms of vibration.

Uses 2H pencil on stone paper It reminds me of a marker due to the sound. The sound is quite affecting my judgment of the tool. With noise cancelling headphones the tool feels much more like the pencil. Now I would say they are quite similar. Your tool has a bit more vibration than the real pencil. The closest comparison from digital would be paperlike but the vibration is too weak when compared to the pencil. Your reproduction is closer.

Uses charcoal on rough paper The reproduction is very similar it has the feeling of drag similar to charcoal which is opposing my motion and making it difficult to move the tool. Also a bit of texture similar to charcoal is present. The tablets feel too smooth to be like a charcoal.

Uses 8B pencil on stone paper I would say this one is the closest to the real tool. It feels like the real pencil. The digital tablets are not similar. The felt nib is closest in terms of resistance but your tools has more surface structure. Paperlike has similar texture but the tool glides too easily when compared to the pencil.

What are the main differences between the three groups?

igital tools feel very smooth and very perfect with no external influence. Your and digital tools have a feeling of texture which gives the tools randomness.

What do you think about the new reproductions of the traditional drawing tools?

Your reproductions feel quite similar to the real tools. I did not draw with visual feedback but they have potential. I would like to test a working prototype and see how my drawing style would change with your surfaces. They feel like they can add the randomness and uniqueness of drawing with a traditional tool which would be interesting to try.

What do you think about the haptic feedback provided by our tools comparing to the traditional and the digital tools?

They feel good. The extra texture feels better than standard tablets.

What is still missing?

The shape should be improved to be more ergonomic. The tip itself is nice.

Would you prefer using a tablet with our enhanced haptic response? Why?

I would definitely prefer a tablet with texture, I do not like too smooth surfaces specially glass tablets. Having haptic feedback similar to real tools and hand eye coordination would be good.

What are the drawbacks of our tools?

The sound affects my perception of the tool. It would be good to match the sound as well.

How could you imagine using these tools?

I would pick them based on the traditional tool I like most. Or when I would like to transfer the drawing style of the traditional media into the virtual world.

D.5 Artist 5

What is your current occupation?

3D art lead in a gaming studio.

What kind of tools did you use in the past e.g. charcoals, pencils, digital styli?

Almost everything. Oil/acryl colors, pencils, pens, charcoals. Wacom tablets, Apple iPad, even have a pocket Wacom with me.

How often do you use traditional drawing tools? What kind of traditional drawing tools do you use? Quite often when I need to relax but its not a part of my job anymore. How often do you use digital drawing tools? What kind of digital drawing tools do you use? (manufacturer/display/tips/any extra modifications, e.g., custom surface)

On a daily basis for work mostly for 3D modelling.

When/for what tasks do you use traditional tools?

I use pencils for concept art because it is much faster. Also for portraits in free time.

When/for what tasks do you use digital tools?

The same tasks as the traditional tools

Why you do not use only traditional tools?

Why you do not use only digital tools?

I use both media to not fall into stereotypes and push myself. Not having undo trains me to pay more attention to the drawing.

What, in your opinion, are the advantages and disadvantages of traditional tools?

One big advantage of traditional tools is that you can make multiple pieces, place them on the ground and pick the direction. This is more complicated in digital format due to limited screen size. Using a pencil and paper is also faster for sketches.

What, in your opinion, are the advantages and disadvantages of digital tools?

Undo is a huge advantage. After getting the idea with a pencil quick and dirty the tablet is much faster at making a production piece. Easier to adjust colors, reproducibility are demanded in the industry.

How important is the feel of the materials when drawing? Why?

It is very important for the immersion. It helps to get into the drawing and not think about other things. Maybe this will change if people stop drawing with traditional tools.

Did you ever notice the difference in haptic feedback between traditional and digital drawing tools? If yes did it ever impact the way you use the drawing tools?

Yes I noticed the difference. I am a fan of Wacom tablets but I also criticize the feeling of plastic on plastic or glass surfaces. It is important to get closer to real tools. For me it is very binary. If I do not like the feeling I do not use the device.

Are you using any means of customizing the haptic feedback of the digital tools? How do you customize your tools? Do you think you have sufficient flexibility in the customization? Do you consider the feel of styli in your choice of tool?

I was using a piece of paper strapped to a Wacom stylus. It was better but I could worse see the piece. I have a friend that uses toothpicks and I heard a broken spaghetti can work too. In general I found a tip that works for me (Wacom felt tip) and bought a lot of them. The shape in hand for a stylus is very important for the drawing. The length also needs to be just right.

Do you think it would be beneficial if the digital drawing tools could provide the same feel as the traditional tools? Why?

I think it would be useful since it would be easier for young people to switch between the two kinds.

A session with traditional tools, digital, and our reproductions

Uses ballpoint pen on paper This is not exactly the same feeling. I was not thinking about it before but I think with the pen I can feel moving the ball which is missing with your tool. I think having the sound more closer would help. I have an app in which when I draw the app makes a sound of the pen which helps me to focus on the work. When I focus on the haptic feedback itself it feels good. Maybe what I am missing most is that your prototype does not draw. In general, this combination might be too rough for a tablet. It would be interesting to see the difference in drawing without software smoothing on your surface if it would get similar to real tool. Paperlike is even smoother than the glass tablet which is counter intuitive. The feeling is not close to a pen, not even a pencil since its has so little resistance and is so smooth. The feeling of surface is stronger in your tool than the digital ones. But perhaps the surface is a bit too rough for the paper which makes the surface feeling more coarse.

Uses 2H pencil on stone paper The feeling of the pencil and the paper is nice. It is very close to the pencil. The feeling of drawing is much better and the sound is also closer. It feels quick and dirty and I would like to use it for concept art. All of the digital tools are too soft to compare. The wacom felt nib is close in terms of resistance but the structure feeling is 3 times more intensive in real tools.

Uses charcoal on rough paper I can feel the charcoal has a lot of resistance. That is similar to the Wacom rubber nib which also drags but does not have the texture. Your tool drags similar to the charcoal too. I liked it, it is not bad you can feel the texture similar to charcoal. One disadvantage is the stickiness. If you could remove that it would be perfect.

Uses 8B on stone paper The pencil feels awesome, this is pencil porn. The reproduction looks really good feels very close to the pencil. It has real feeling of paper and proper contrast of substrate. I think this is really awesome if it was transparent I would use it. There is still a little bit of stick. It is not bothering but a real pencil does not have it. From all the tools this is something I would like the most to use. But it is very subjective and other people might prefer other surfaces.

What are the main differences between the three groups?

The feeling of texture which is absent in digital tools.

What do you think about the new reproductions of the traditional drawing tools?

The reproductions are very interesting and I like that they challenge the status quo.

What do you think about the haptic feedback provided by our tools comparing to the traditional and the digital tools?

I like that your tools are rougher than digital tools. The rougher feeling is closer to a paper.

What is still missing?

The stickiness at the end is bothering me. It would be good to reduce it.

Would you prefer using a tablet with our enhanced haptic response? Why?

I would like to have your surfaces as an option. Ideally I would pick 3 which provide nice range with clear distinction. Like pen, like charcoal, like pencil. This would make it easy for people to choose but if you give people too much choice they could struggle.

What are the drawbacks of our tools?

Personal bias. The smooth tool worked well for me because that is what I would also prefer to draw. For other people they might prefer the rougher tools and then also respond better to them.

How could you imagine using these tools?

—