

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

MICHAL PIOVARČI, Università della Svizzera italiana, Switzerland

DANNY M. KAUFMAN, Adobe Research

DAVID I. W. LEVIN, University of Toronto, Canada

PIOTR DIDYK, Università della Svizzera italiana, Switzerland




Fig. 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.

Digital drawing tools are now standard in art and design workflows. These tools offer comfort, portability, and precision as well as native integration with digital-art workflows, software, and tools. At the same time, artists continue to work with long-standing, traditional drawing tools. One feature of traditional tools, well-appreciated by many artists and lacking in digital tools, is the specific and diverse range of haptic responses provided by them. Haptic feedback in traditional drawing tools provides unique, per-tool responses that help determine the precision and character of individual strokes. In this work, we address the problem of fabricating digital drawing tools that closely match the haptic feedback of their traditional counterparts. This requires the formulation and solution of a complex, co-optimization of both digital styli and the drawing surfaces they move upon. Here, a potentially direct formulation of this optimization with numerical simulation-in-the-loop is not yet viable. As in many complex design tasks, state-of-the-art methods do not currently offer predictive modeling at rates and scales that can account for the numerous, coupled, physical behaviors governing the haptics of styli and surfaces, nor for the limitations and uncertainties inherent in their fabrication processes. To address these challenges, we propose fabrication-in-the-loop optimization. Critical to making this strategy practical we construct our objective via a Gaussian Process that does not require computing derivatives with respect to design parameters. Our Gaussian Process surrogate model then provides both function estimates and confidence intervals that guide the efficient sampling of our design space. In turn, this sampling critically reduces the numbers of fabricated examples during exploration and automatically handles exploration-exploitation trade-offs. We apply our method to fabricate drawing tools that provide a wide range

of haptic feedback, and demonstrate that they are often hard for users to distinguish from their traditional drawing-tool analogs.

CCS Concepts: • **Computing methodologies** → **Graphics input devices; Perception**; • **Human-centered computing** → **Haptic devices**;

Additional Key Words and Phrases: perception, fabrication, drawing tools, haptics, fabrication-in-the-loop, co-optimization

ACM Reference Format:

Michal Piovarči, Danny M. Kaufman, David I. W. Levin, and Piotr Didyk. 2020. Fabrication-in-the-Loop Co-Optimization of Surfaces and Styli for Drawing Haptics. *ACM Trans. Graph.* 39, 4, Article 1 (July 2020), 16 pages. <https://doi.org/10.1145/3386569.3392467>

1 INTRODUCTION

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 [Annett et al. 2014; Danna and Velay 2015]. 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.

Authors' addresses: Michal Piovarči, Università della Svizzera italiana, Via Giuseppe Buffi 13, Lugano, 6900, Switzerland, michal.piovarci@usi.ch; Danny M. Kaufman, Adobe Research, Seattle; David I. W. Levin, University of Toronto, Toronto, Canada; Piotr Didyk, Università della Svizzera italiana, Lugano, Switzerland.

© 2020 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3386569.3392467>.

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 tools¹ 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.

The haptic feedback of drawing tools is a result of a complex interaction of the drawing tool with the substrate. The results of this 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 2). In this paper, we tackle the problem of fabricating digital tools that 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 desired 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 (micro-geometry and material). In this paper, we reduce the design space to nine dimensions. This still makes straightforward solutions, such as random sampling, impractical due to the massive amounts of designs that would need to be fabricated and evaluated.

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 phenomena, 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




Fig. 2. Frictional contact

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.

To address the above challenges, we present a new 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 a recently proposed perceptual space of drawing tools [Piovarči et al. 2018], 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 tools 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.

¹<https://www.adobe.com/products/fresco.html>

2 RELATED WORK

Our work draws inspiration from several research areas, including haptic feedback reproduction, contact modeling, and fabrication-oriented optimization. Below, we provide an overview of the most relevant work from these areas.

2.1 Haptic Feedback Reproduction

Our sense of touch is vital for efficiently navigating both real and virtual environments [Robles De La Torre 2006]. The vibrations and forces perceived on our skin enable precise interaction with objects even without directly looking at them. Unfortunately, replicating a specific sense of touch using digital devices is still an open and challenging problem. Most of the haptic devices are active. They range from Phantoms [Massie et al. 1994] and tactile displays [Hayward and Cruz-Hernandez 2000; Perez et al. 2017], which apply direct force using motors and actuators, to micro and electro vibrators [Bau et al. 2010; Kim et al. 2013], which stimulate the perception of friction and shape using local vibrations. For a more in-depth discussion of such devices, we refer to the survey by Chouvardas et al.'s [2008]. In this work, we do not tackle the general problem of reproducing haptic feedback but focus on replicating the feedback of drawing tools using passive devices.

The combination of a drawing tool and a paper creates a unique, intimate coupling that produces a characteristic haptic response. The manufacturers of digital devices acknowledge the problem [Annett et al. 2014] and address it by providing users a selection of stylus nibs (Wacom, Microsoft), aftermarket patterned screen protectors (PaperLike), or using micro-scratched glass surface (reMarkable). However, realistic reproduction of haptic feedback remains a challenging open problem [Choi and Tan 2005], which motivates companies to develop solutions that can automatically digitize drawings (Wacom Bamboo, iSkin) at the cost of the digital convenience.

Researches took an alternative approach and reintroduced the haptic feedback of traditional drawing tools using vibration motors [Arasan et al. 2013; Lee et al. 2004; Poupyrev et al. 2004]. To mimic textures investigated with a stylus-like probe Romano and Kuchenbecker [2012] proposed a data-driven algorithm where collected accelerometer data drives a vibration motor connected to a stylus. They demonstrated the ability to reproduce the haptic response of a wide range of materials. Building upon this idea Cho et al. [2016] constructed a stylus with a vibratory and auditory response of drawing tools. Wang et al. [2016] proposed to use electrovibration to replicate the frictional properties of a stylus. The main limitation of active modulation is the significant latency introduced by the mechanical components [Annett et al. 2014; Helps and Helps 2016]. Our passive tools avoid this problem and benefit from providing an immediate haptic response to the users. Additionally, our tools can be easily manufactured using available 2D and 3D printers, which avoids any mechanical and electrical components.

Closest to our approach is the work of Piovarči et al. [2018]. The authors propose a perceptual space of drawing tools which combined with a data-driven simulation, leads to an optimization of digital styli. The method can only interpolate between previously manufactured styli, and therefore, cannot expand the range of achievable haptic feedback. In contrast, we directly optimize for

both stylus and surface given target haptic feedback. This approach allows us to match the haptics of traditional drawing tools closely.

2.2 Stylus-Surface Contact Modelling

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. One possibility to mimic the drawing substrates is to scan the surface. Either using a 3D scanner, shape from shading [Prados and Faugeras 2006], or a capturing system like Gelsight [Yuan et al. 2017]. Afterward, one can use an appropriate manufacturing method to reproduce the scanned mesh. However, 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.

In computer graphics, the frictional contact is typically limited to the Coulomb friction model [Harnoy et al. 2008], where frictional force is directly proportional to the applied load. However, in the context of drawing tools effects like lubrication [Harnoy et al. 2008], wear [Stachowiak 2006], and material deposition [Archard 1953] become relevant. The problem is further complicated by the coupling between friction and contact with elasticity and viscous damping [Chen et al. 2017]. Proper frictional modeling under these conditions is an open area of research [Bertails-Descoubes et al. 2011; Kaufman et al. 2005, 2008], and many of the underlying factors are still not well understood.

2.3 Goal-based Computational Fabrication

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]. 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 tractable, the choice of suitable parametrization of the design space becomes critical. 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 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 produced by stylus-surface interaction.

Apart from the numerical simulations, which make the optimization procedure challenging, goal-based computational fabrication

techniques must also account for the existing limitations of fabrication devices. For example, when 3D printing objects with prescribed elastic behavior, a method can only use the materials available on the 3D printer. A possible solution is to design so-called metamaterials that approximate the desired behavior not only by varying material distribution within the printing volume but also the geometry [Bickel et al. 2010; Martínez et al. 2016; Panetta et al. 2015; Schumacher et al. 2015]. Here we draw inspiration from these works: instead of creating a carbon copy of drawing tools, we focus on designing tools that obtain similar behavior.

2.4 Gaussian Processes and Active Learning

Minimizing costly and difficult-to-evaluate error functions, which are common in computational fabrication, is an open and challenging problem [Nocedal and Wright 2006]. One 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 2004], 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 data-points iteratively. Also here, surrogate models like Gaussian Processes are often used to generate the datasets presented to the participants [Akrouy 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 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 4.1) and drawing surfaces (Section 4.2). To measure the error between different designs, we employ a perceptual metric for drawing tools (Section 5.1). To evaluate the error, we propose a fabrication-in-the-loop approach where we use a data-driven surrogate model (Section 5.2). We build the model to handle the stylus-surface co-optimization (Section 5.2.3) and include transfer learning to accelerate optimization of unobserved fabrication processes (Section 5.2.2). 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 5.3). We demonstrate how our model can be used to optimize for the desired interaction of a stylus-surface combination (Section 6). We validate our method by comparing with naive approaches and observing the quality of our manufactured tools

(Section 7). 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 8.2), are preferred to state of the art approaches (Section 8.3), and produce realistic haptic sensations (Section 8.4). Finally, we evaluate our reproductions in a survey with professional artists (Section 9).

4 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.

4.1 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 material 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 3).




Fig. 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 3).

4.2 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 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).




Fig. 4. Gelsight scans of drawing substrates used in our experiments.

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 globally 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 5 shows surfaces generated with our method.




Fig. 5. Sample surfaces from our parametrization quantized to one layer.

5 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 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 process 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.

5.1 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 recent work by Piovarči et al. [2018], where the authors 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. To recover these properties the authors propose a custom measurement setup² where a drawing tool is placed in a unified holder and the interaction with a substrate is automatically measured. The authors proposed a likelihood-based optimization that correlates perceived friction with Coulomb friction and perceived vibration with the mean value of the spectrogram across frequency and velocity. The two perceived quantities define a perceptual space where the similarity between tool corresponds 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 study conducted by Piovarči et al. [2018]. 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 psychophysical studies conducted by Piovarči et al.

²[misop.github.io/Perception-Aware-Modeling-and-Fabrication-of-Digital-Drawing-Tools/](https://github.com/misop/Perception-Aware-Modeling-and-Fabrication-of-Digital-Drawing-Tools/)

5.2 Gaussian Process Surrogate Model

The input to our optimization is a characterization of haptic feedback based on the perceptual space: Coulomb friction and average vibratory feedback. The output is a stylus and surface 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. 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:

$$\kappa(\mathbf{x}, \mathbf{x}') = \kappa_1(\mathbf{x}, \mathbf{x}') * \kappa_2(\mathbf{x}, \mathbf{x}'), \quad (1)$$

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.

5.2.1 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 by [Piovarči et al. 2018] 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:

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

$$r_\kappa = \sqrt{\sum_{i=1}^D \frac{(x_i - x'_i)^2}{\sigma_i^2}}, \quad (3)$$

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.

5.2.2 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 frictional and vibratory feedback for the same surface. 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]:

$$\kappa_2(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{1}{2} \sum_{j=1}^K \frac{1}{\sigma_j^2} (x_j - x'_j)^2\right), \quad (4)$$

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 6). For cases where data about a category are sufficient, we can use $\sigma_j \rightarrow \infty$ 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.




Fig. 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).

5.2.3 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:

$$\begin{bmatrix} F \\ V \end{bmatrix} = \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{K}_F & a \\ a & \mathbf{K}_V \end{bmatrix} + \begin{bmatrix} \sigma_F^2 \mathbf{I} & 0 \\ 0 & \sigma_V^2 \mathbf{I} \end{bmatrix} \right), \quad (5)$$

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.

5.3 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 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:

$$u(\mathbf{x}) = \max(0, d^* - (\mathcal{N}_F(\mathbf{x}) - F)^2 - (\mathcal{N}_V(\mathbf{x}) - V)^2), \quad (6)$$

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. 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




Fig. 7. Acquisition function

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 gray). The expected improvement then includes the frictional values (Figure 7 orange) located within the integration bounds multiplied by their probability (Figure 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 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:

$$\int_{V-d}^{V+d} \int_{F-d}^{F+d} (d^* - (x - F)^2 - (y - V)^2) p_x p_y dx dy, \quad (7)$$

where p_x and p_y is probability of predicting frictional and vibrational value respectively. The analytical solution of this integral is described in the Supplementary Material.

Constraining the integration region encourages exploration and prevents the over-exploitation of observed data. Additionally, it is

possible to modify the trade-off by adjusting the means of each predicted distributions. By moving the means closer towards the target the acquisition function prefers to exploit the data and search for local minima of the function. Conversely, by moving the means further away from the target we encourage exploration of uncertain regions in the surrogate model.

5.3.1 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.

6 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 \mathcal{S} by maximizing the expected improvement defined in Equation 6:

$$\max_{\mathbf{t}, i} u(\mathbf{t}, \mathbf{s}_i), \mathbf{s}_i \in \mathcal{S}, \quad (8)$$

where \mathbf{t} is the parametrization of the stylus, and \mathbf{s}_i is a surface from the surface set \mathcal{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 \mathcal{T} . We search for the surface by maximizing:

$$\max_{\mathbf{s}, i} u(\mathbf{t}_i, \mathbf{s}), \mathbf{t}_i \in \mathcal{T}, \quad (9)$$

where \mathbf{s} is the parametrization of the surface, and \mathbf{t}_i is a tool from tool set \mathcal{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 below 0.3. We can see the behavior of our optimization on an example in Figure 8.

ALGORITHM 1: Perceptual Optimization Pseudocode

```

Data:  $T$  = target behavior
Data:  $o$  = observations
while true do
   $\mathcal{N}_F(x) \leftarrow$  gaussian_process(friction( $o$ ));
   $\mathcal{N}_V(x) \leftarrow$  gaussian_process(vibration( $o$ ));
   $t =$  pick new tool by solving Equation 8;
  if not printing or tool has better improvement then
    | print new tool;
  end
   $s =$  pick new surface by solving Equation 9;
   $m =$  measure new tool;
   $o$  add  $[m, s]$  pair to observations dataset;
  if  $||T - \text{Perceptual}(m)|| < 0.3$  then
    | done;
  end
end

```

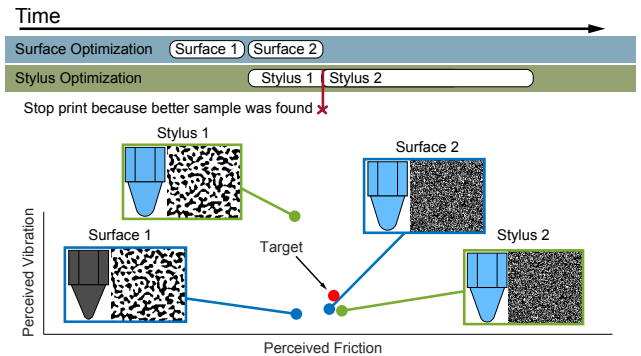


Fig. 8. Example of parallel execution of stylus and surface manufacturing. In the first iteration, we manufacture a surface and enhance our data-driven model. The model then predicts which tool we should print to best match the target. We prepare the print for the tool and while the tool is printing 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.

6.1 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 9 left. The initial surface designs were generated automatically using random sampling, Figure 9 right. We measured each initial design and found their appropriate placement in the perceptual space of haptic feedback, (Figure 10 left). We can observe that our initial sampling provides good coverage of tools with lower perceived friction.




Fig. 9. Pen and surface designs used to initialize the surrogate model.

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 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.

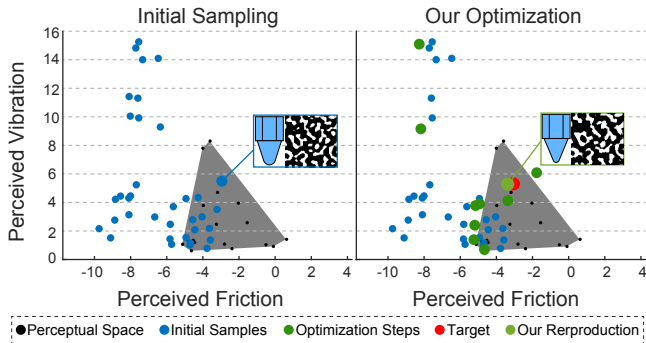


Fig. 10. Initial samples measurements are shown on the left. We can observe that initial sampling can capture a ballpoint pen on office paper. We verify our algorithm by removing the surface from initial sampling and optimize for the same tool, right.

6.2 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 11 2H pencil). We use our surrogate model to optimize for the haptic feedback of an 8B pencil on stone paper, (Figure 11 8B pencil). The optimization 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 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 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 11 charcoal). The resulting optimized stylus-surface combinations are shown in Figure 12.




Fig. 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.

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

6.3 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

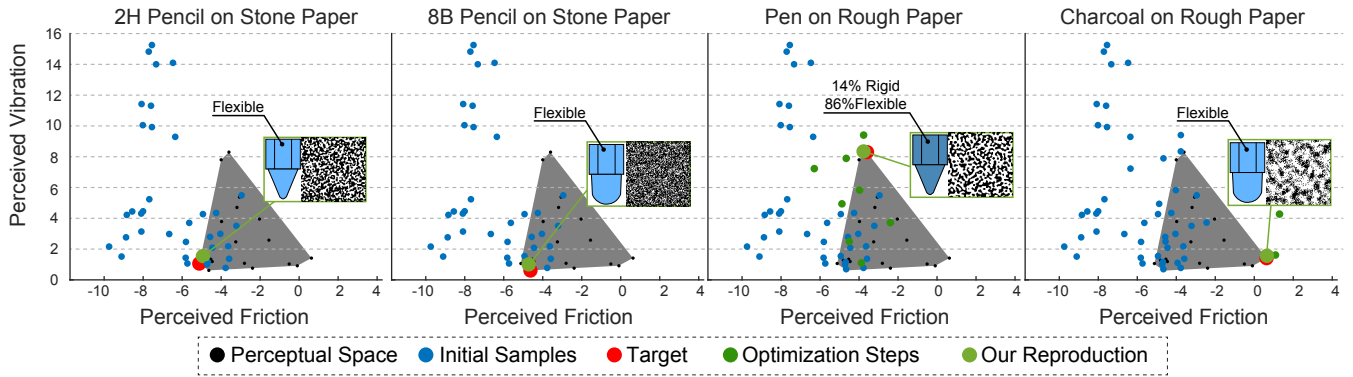


Fig. 11. Results of optimizing haptic feedback using our surrogate model. Black samples mark original perceptual space. Blue samples are sample surfaces used to initialize the model. The red color is the position of a tool we would like to reproduce. Dark green dots mark the path our optimizer took in the perceptual space to find the final reproduction (light green).

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 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 recover 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 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 13. Each manufacturing 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.

6.4 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

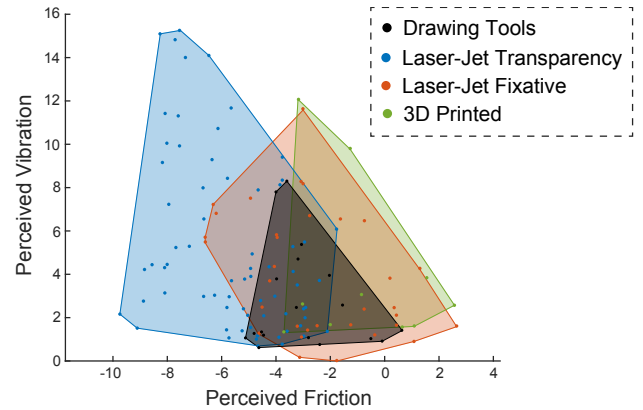


Fig. 13. The gamut of haptic feedback achievable by our different manufacturing processes: laser-jet on transparency sheet (blue), laser-jet on transparency sheet with fixative (red), 3D printer (green) compared with traditional drawing tools (black).

surrogate model. Instead, we use the surrogate model to estimate the expected haptic feedback achievable on the surface as:

$$E[s] = \text{ConvexHull}(\{N_{F_i}(s), N_{V_i}(s)\}), \quad (10)$$

where s is the parametrization of the surface, $N_{F_i}(s)$ and $N_{V_i}(s)$ are the Gaussian processes predicting friction and vibration of tool i from the toolset \mathcal{T} on the surface 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:

$$\max_s \text{Intersect}(\text{input}, \text{ConvexHull}(\{N_{F_i}(s), N_{V_i}(s)\})). \quad (11)$$

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 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.

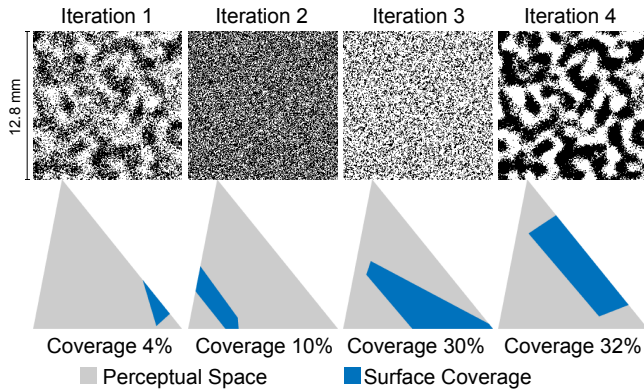


Fig. 14. Four iteration of optimal surface optimization.

7 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.

7.1 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 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 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.




Fig. 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.

7.2 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 16. We can observe that the 3D printed surfaces achieve very similar vibratory response as the original instruments.




Fig. 16. Vibration measurements of original drawing instruments and their 3D printed counterparts.

7.3 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 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.




Fig. 17. Visual properties of our surface evaluated by taking photos of an LCD screen displaying a checkerboard pattern.

8 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.

8.1 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].

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.

8.2 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.

8.3 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 7.1), 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 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).

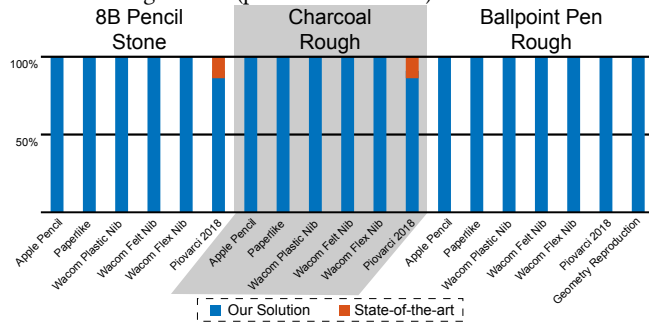


Fig. 18. Bar-plot showing the similarity of our tools when compared to other digital styli. For each reference, indicated on the top, and an alternative solution, indicated below the bars, the plot shows the percentage of participants finding our replica more similar than the alternative commercial or research solution.

8.4 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 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 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.

8.5 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

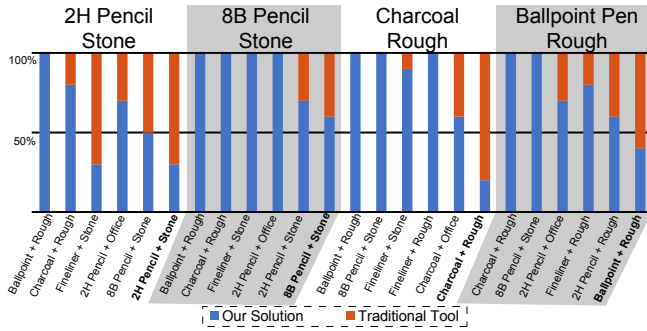


Fig. 19. Bar-plot showing the similarity of our tools when compared to traditional drawing tools. For each reference, indicated on the top, and an alternative traditional tool, indicated below the bars, the plot shows the percentage of participants finding our replica more similar than the alternative tool.

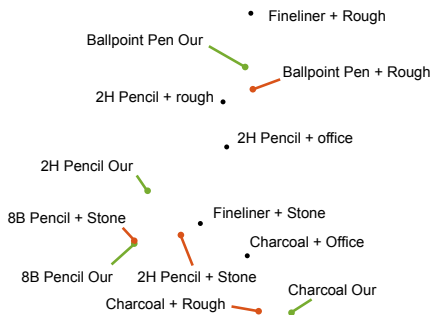


Fig. 20. Two-dimensional perceptual space computed for tools in our experiments. The Euclidean distance between the samples can be interpreted as a perceived distance between different drawing tools combinations.

changes how the tool feels over time, but it is not desired for the digital tools to replicate the effect.

9 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 productivity 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 the supplementary materials.

10 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 extensions.

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.

11 CONCLUSION

Despite the success of digital drawing tools, fabricating tools which can closely replicate the haptic feedback of the traditional drawing tools 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.

12 ACKNOWLEDGEMENTS

We would like to thank Okan Tarhan Tursun for his help with the video. This work is graciously supported by the following grant agencies: ERC Starting Grant (PERDY-804226), NSERC Discovery (RGPIN-2017-05524), NSERC Accelerator (RGPAS-2017-507909), Connaught Fund (503114), Canadian Foundation for Innovation, Canada Research Chairs Programme.

REFERENCES

- Riad Akrouf, Marc Schoenauer, and Michele Sebag. 2011. Preference-Based Policy Learning. In *Machine Learning and Knowledge Discovery in Databases*, Dimitrios Gunopulos, Thomas Hofmann, Donato Malerba, and Michalis Vazirgiannis (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 12–27.
- Michelle Annett, Fraser Anderson, Walter F Bischof, and Anoop Gupta. 2014. The pen is mightier: understanding stylus behaviour while inking on tablets. In *Proceedings of Graphics Interface 2014*. Canadian Information Processing Society, 193–200.
- Atakan Arasan, Cagatay Basdogan, and Tefik Metin Sezgin. 2013. Haptic stylus with inertial and vibro-tactile feedback. In *2013 World Haptics Conference (WHC)*. 425–430. <https://doi.org/10.1109/WHC.2013.6548446>
- John F. Archard. 1953. Contact and Rubbing of Flat Surfaces. *Journal of Applied Physics* 24, 8 (1953), 981–988. <https://doi.org/10.1063/1.1721448> arXiv:<https://doi.org/10.1063/1.1721448>
- Olivier Bau, Ivan Poupyrev, Ali Israr, and Chris Harrison. 2010. TeslaTouch: Electro-vibration for Touch Surfaces. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, New York, NY, USA, 283–292. <https://doi.org/10.1145/1866029.1866074>
- Florence Bertails-Descoubes, Florent Cadoux, Gilles Daviet, and Vincent Acary. 2011. A nonsmooth Newton solver for capturing exact Coulomb friction in fiber assemblies. *ACM Transactions on Graphics* 30, 1 (Feb. 2011), 6:1–6:14.
- Gaurav Bharaj, David I. W. Levin, James Tompkin, Yun Fei, Hanspeter Pfister, Wojciech Matusik, and Changxi Zheng. 2015. Computational Design of Metallophone Contact Sounds. *ACM Transactions on Graphics* 34, 6 (Oct. 2015), 13. <https://doi.org/10.1145/2816795.2818108>
- Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. 2010. Design and Fabrication of Materials with Desired Deformation Behavior. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 3 (2010).
- Peter J. Blau and Joseph K. Gardner. 1996. Tribological characteristics of graded pencil cores on paper. *Wear* 197, 1 (1996), 233–241. [https://doi.org/10.1016/0043-1648\(96\)06952-9](https://doi.org/10.1016/0043-1648(96)06952-9)
- Edwin V Bonilla, Kian M. Chai, and Christopher Williams. 2008. Multi-task Gaussian Process Prediction. In *Advances in Neural Information Processing Systems 20*, J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis (Eds.). Curran Associates, Inc., 153–160. <http://papers.nips.cc/paper/3189-multi-task-gaussian-process-prediction.pdf>
- Desai Chen, David I. W. Levin, Piotr Didyk, Pitchaya Sitthi-Amorn, and Wojciech Matusik. 2013. Spec2Fab: A reducer-tuner model for translating specifications to 3D prints. *ACM Transactions on Graphics (Proceedings SIGGRAPH 2013, Anaheim, CA)* 32, 4 (2013), 135:1–135:10.
- Desai Chen, David I. W. Levin, Wojciech Matusik, and Danny M. Kaufman. 2017. Dynamics-aware Numerical Coarsening for Fabrication Design. *ACM Transactions*

- on *Graphics* 36, 4 (July 2017), 84:1–84:15. <https://doi.org/10.1145/3072959.3073669>
- Desai Chen, David I. W. Levin, Shinjiro Sueda, and Wojciech Matusik. 2015. Data-Driven Finite Elements for Geometry and Material Design. *ACM Transactions on Graphics* 34, 4 (July 2015), 10. <https://doi.org/10.1145/2766889>
- Youngjun Cho, Andrea Bianchi, Nicolai Marquardt, and Nadia Bianchi-Berthouze. 2016. RealPen: Providing Realism in Handwriting Tasks on Touch Surfaces Using Auditory-Tactile Feedback. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 195–205. <https://doi.org/10.1145/2984511.2984550>
- Seungmoon Choi and Hong Z. Tan. 2005. Toward Realistic Haptic Rendering of Surface Textures. In *ACM SIGGRAPH 2005 Courses (SIGGRAPH '05)*. ACM, New York, NY, USA, Article 125. <https://doi.org/10.1145/1198555.1198612>
- Vasilios G Chouvardas, Amalia N Miliou, and Miltiadis K Hatalis. 2008. Tactile displays: Overview and recent advances. *Displays* 29, 3 (2008), 185–194.
- Andrew R Conn, Nicholas IM Gould, and Philippe Toint. 1991. A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds. *SIAM J. Numer. Anal.* 28, 2 (1991), 545–572.
- Jérémy Danna and Jean-Luc Velay. 2015. Basic and supplementary sensory feedback in handwriting. *Frontiers in Psychology* 6 (2015), 169. <https://doi.org/10.3389/fpsyg.2015.00169>
- John J. Dudley, Jason T. Jacques, and Per Ola Kristensson. 2019. Crowdsourcing Interface Feature Design with Bayesian Optimization. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, Article Paper 252, 12 pages. <https://doi.org/10.1145/3290605.3300482>
- Oskar Elek, Denis Sumin, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Krivánek. 2017. Scattering-aware Texture Reproduction for 3D Printing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36, 6 (2017), 241:1–241:15.
- Embark Studios. 2019. Ttexture Synthesis. <https://github.com/EmbarkStudios/ttexture-synthesis>.
- Gustav Theodor Fechner. 1860. *Elemente der Psychophysik*. Number v. 1 in *Elemente der Psychophysik*. Breitkopf und Härtel. <https://books.google.ch/books?id=6rINAAAAAAAJ>
- Wolfgang Johann Fischer, Melanie Mayr, Stefan Spirk, David Reishofer, Lukas Andreas Jagiello, Romana Schmiedt, Jerome Colson, Armin Zankel, and Wolfgang Bauer. 2017. Pulp Fines-Characterization, Sheet Formation, and Comparison to Microfibrillated Cellulose. *Polymers* 9, 8 (2017). <https://doi.org/10.3390/polym9080366>
- Peter Frazier, Warren Powell, and Savas Dayanik. 2009. The Knowledge-Gradient Policy for Correlated Normal Beliefs. *INFORMS Journal on Computing* 21, 4 (2009), 599–613. <https://doi.org/10.1287/ijoc.1080.0314> arXiv:https://doi.org/10.1287/ijoc.1080.0314
- Eduardo C. Garrido-Merchán and Daniel Hernández-Lobato. 2020. Dealing with categorical and integer-valued variables in Bayesian Optimization with Gaussian processes. *Neurocomputing* 380 (Mar 2020), 20–35. <https://doi.org/10.1016/j.neucom.2019.11.004>
- Avraham Harnoy, Bernard Friedland, and Simon Cohn. 2008. Modeling and measuring friction effects. *IEEE Control Systems* 28, 6 (Dec 2008), 82–91. <https://doi.org/10.1109/MCS.2008.929546>
- Vincent Hayward and M Cruz-Hernandez. 2000. Tactile display device using distributed lateral skin stretch. In *Proceedings of the haptic interfaces for virtual environment and teleoperator systems symposium*, Vol. 69. ASME, 1309–1314.
- Richard Helps and Clarissa Helps. 2016. Measuring Stylus and Tablet Performance for Usability in Sketching. In *Proceedings of the 5th Annual Conference on Research in Information Technology*. ACM, 19–24.
- Philipp Hennig and Christian J. Schuler. 2012. Entropy Search for Information-Efficient Global Optimization. *J. Mach. Learn. Res.* 13 (June 2012), 1809–1837.
- Donald R. Jones, Matthias Schonlau, and William J. Welch. 1998. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization* 13, 4 (01 Dec 1998), 455–492. <https://doi.org/10.1023/A:1008306431147>
- Danny M. Kaufman, Timothy Edmunds, and Dinesh K. Pai. 2005. Fast Frictional Dynamics for Rigid Bodies. *ACM Transactions on Graphics* 24, 3 (July 2005), 946–956. <https://doi.org/10.1145/1073204.1073295>
- Danny M. Kaufman, Shinjiro Sueda, Doug L. James, and Dinesh K. Pai. 2008. Staggered Projections for Frictional Contact in Multibody Systems. *ACM Transactions on Graphics* 27, 5 (Dec. 2008), 164:1–164:11. <https://doi.org/10.1145/1409060.1409117>
- Seung-Chan Kim, Ali Israr, and Ivan Poupyrev. 2013. Tactile Rendering of 3D Features on Touch Surfaces. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. ACM, New York, NY, USA, 531–538. <https://doi.org/10.1145/2501988.2502020>
- Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. 2017. Sequential Line Search for Efficient Visual Design Optimization by Crowds. *ACM Transactions on Graphics* 36, 4 (July 2017), 11. <https://doi.org/10.1145/3072959.3073598>
- Ares Lagae, Sylvain Lefebvre, George Drettakis, and Philip Dutré. 2009. Procedural Noise using Sparse Gabor Convolution. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2009)* 28, 3 (July 2009), 54–64. <https://doi.org/10.1145/1531326.1531360>
- Ares Lagae, Sylvain Lefebvre, and Philip Dutré. 2011. Improving Gabor noise. *IEEE Transactions on Visualization and Computer Graphics* (2011). <http://www-sop.inria.fr/reves/Basilic/2011/LLD11>
- Johnny C. Lee, Paul H. Dietz, Darren Leigh, William S. Yerazunis, and Scott E. Hudson. 2004. Haptic Pen: A Tactile Feedback Stylus for Touch Screens. In *Proceedings of the 17th Annual ACM Symposium on User Interface Software and Technology (UIST '04)*. ACM, New York, NY, USA, 291–294. <https://doi.org/10.1145/1029632.1029682>
- Dingzeyu Li, David I.W. Levin, Wojciech Matusik, and Changxi Zheng. 2016. Acoustic Voxels: Computational Optimization of Modular Acoustic Filters. *ACM Transactions on Graphics* 35, 4 (2016). <https://doi.org/10.1145/2897824.2925960>
- Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. 2016. Procedural Voronoi Foams for Additive Manufacturing. *ACM Transactions on Graphics* 35, 4 (July 2016), 12. <https://doi.org/10.1145/2897824.2925922>
- Jonàs Martínez, Méline Skouras, Christian Schumacher, Samuel Hornus, Sylvain Lefebvre, and Bernhard Thomaszewski. 2019. Star-Shaped Metrics for Mechanical Metamaterial Design. *ACM Transactions on Graphics* 38, 4 (July 2019), 82:1–13. <https://doi.org/10.1145/3306346.3322989> Special issue, SIGGRAPH 2019.
- Thomas H. Massie, J. Kenneth Salisbury, et al. 1994. The phantom haptic interface: A device for probing virtual objects. In *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems*, Vol. 55. Chicago, IL, 295–300.
- Jonas Mockus. 1989. *Bayesian Methods for Global Optimization in the Gaussian Case*. Springer Netherlands, Dordrecht, 79–116. https://doi.org/10.1007/978-94-009-0909-0_5
- Radford M. Neal. 1996. *Bayesian Learning for Neural Networks*. Springer-Verlag, Berlin, Heidelberg.
- Jorge Nocedal and Stephen Wright. 2006. *Numerical optimization*. Springer New York.
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic Textures for Additive Fabrication. *ACM Transactions on Graphics* 34, 4 (July 2015), 135:1–135:12.
- Alvaro G. Perez, Daniel Lobo, Francesco Chinello, Gabriel Cirio, Monica Malvezzi, Jose San Martin, Domenico Prattichizzo, and Miguel A. Otaduy. 2017. Optimization-based wearable tactile rendering. *IEEE transactions on haptics* (2017).
- Michal Piovarči, David I. W. Levin, Danny M. Kaufman, and Piotr Didyk. 2018. Perception-Aware Modeling and Fabrication of Digital Drawing Tools. *ACM Transactions on Graphics* 37, 4 (July 2018), 15. <https://doi.org/10.1145/3197517.3201322>
- Michal Piovarči, Michael Wessely, Michal Jagielski, Marc Alexa, Wojciech Matusik, and Piotr Didyk. 2017. Directional Screens. In *Proceedings of SCF '17 (SCF '17)*, Vol. 35. Ivan Poupyrev, Makoto Okabe, and Shigeaki Maruyama. 2004. Haptic Feedback for Pen Computing: Directions and Strategies. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems (CHI EA '04)*. ACM, New York, NY, USA, 1309–1312. <https://doi.org/10.1145/985921.986051>
- Emmanuel Prados and Olivier Faugeras. 2006. *Shape From Shading*. Springer US, Boston, MA, 375–388. https://doi.org/10.1007/0-387-28831-7_23
- Carl Edward Rasmussen. 2004. *Gaussian Processes in Machine Learning*. Springer Berlin Heidelberg, Berlin, Heidelberg, 63–71. https://doi.org/10.1007/978-3-540-28650-9_4
- Carl Edward Rasmussen and Christopher K. I. Williams. 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- Gabriel Robles De La Torre. 2006. The importance of the sense of touch in virtual and real environments. *IEEE MultiMedia* 13, 3 (July 2006), 24–30. <https://doi.org/10.1109/MMUL.2006.69>
- Joseph M. Romano and Katherine J. Kuchenbecker. 2012. Creating Realistic Virtual Textures from Contact Acceleration Data. *IEEE Transactions on Haptics* 5, 2 (April 2012), 109–119. <https://doi.org/10.1109/TOH.2011.38>
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to Control Elasticity in 3D Printing. *ACM Transactions on Graphics* 34, 4 (July 2015), 13. <https://doi.org/10.1145/2766926>
- Yuliy Schwartzburg, Romain Testuz, Andrea Tagliasacchi, and Mark Pauly. 2014. High-contrast Computational Caustic Design. *ACM Transactions on Graphics* 33, 4 (July 2014), 74:1–74:11. <https://doi.org/10.1145/2601097.2601200>
- Pitchaya Sitthi-Amorn, Javier E. Ramos, Yuwang Wang, Joyce Kwan, Justin Lan, Wenshou Wang, and Wojciech Matusik. 2015. MultiFab: A Machine Vision Assisted Platform for Multi-Material 3D Printing. *ACM Transactions on Graphics* 34, 4 (July 2015), 11. <https://doi.org/10.1145/2766962>
- Gwidon W Stachowiak. 2006. *Wear: materials, mechanisms and practice*. John Wiley & Sons.
- Qinglong Wang, Xiangshi Ren, Sayan Sarcar, and Xiaoying Sun. 2016. EV-Pen: Leveraging Electro-vibration Haptic Feedback in Pen Interaction. In *Proceedings of the 2016 ACM on Interactive Surfaces and Spaces (ISS '16)*. ACM, New York, NY, USA, 57–66. <https://doi.org/10.1145/2992154.2992161>
- Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. 2017. GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force. *Sensors* 17, 12 (2017). <https://doi.org/10.3390/s17122762>