Perception-Aware Modeling and Fabrication of Digital Drawing Tools

MICHAL PIOVARČI, MPI Informatik, Saarland University, MMCI, and Università della Svizzera italiana, Switzerland
DAVID I. W. LEVIN, University of Toronto, Canada
DANNY M. KAUFMAN, Adobe Research, United States of America
PIOTR DIDYK, MPI Informatik, Saarland University, MMCI, and Università della Svizzera italiana, Switzerland

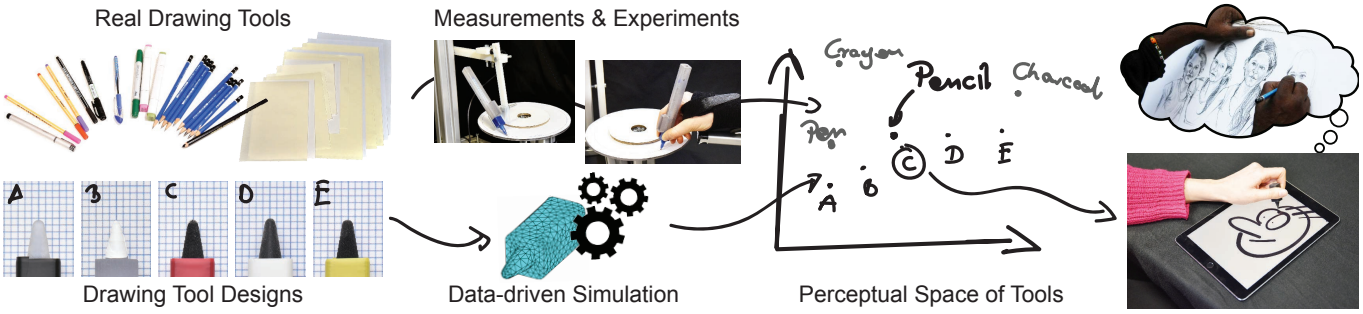

Fig. 1. We propose a system for fabricating digital drawing tools that mimic the feel of real tools. To this end, we measure properties of different real drawing tools, study their perception, and design a perception-aware space of drawing tools. We later develop a simulation technique which allows us to embed new designs into the space, evaluate the pairwise similarity between them and the tools we want to replicate. This drives the design process of different digital tools.

Digital drawing is becoming a favorite technique for many artists. It allows for quick swaps between different materials, reverting changes, and applying selective modifications to finished artwork. These features enable artists to be more efficient and creative. A significant disadvantage of digital drawing is poor haptic feedback. Artists are usually limited to one surface and a few different stylus nibs, and while they try to find a combination that suits their needs, this is typically challenging. In this work, we address this problem and propose a method for designing, evaluating, and optimizing different stylus designs. We begin with collecting a representative set of traditional drawing tools. We measure their physical properties and conduct a user experiment to build a perceptual space that encodes perceptually-relevant attributes of drawing materials. The space is optimized to both explain our experimental data and correlate it with measurable physical properties. To embed new drawing tool designs into the space without conducting additional experiments and measurements, we propose a new, data-driven simulation technique for characterizing stylus-surface interaction. We finally leverage the perceptual space, our simulation, and recent advancements in multi-material 3D printing to demonstrate the application of our system in the design of new digital drawing tools that mimic traditional drawing materials.

Authors' addresses: Michal Piovarči, MPI Informatik, Saarland University, MMCI, Università della Svizzera italiana, Switzerland; David I. W. Levin, University of Toronto, Canada; Danny M. Kaufman, Adobe Research, United States of America; Piotr Didyk, MPI Informatik, Saarland University, MMCI, Università della Svizzera italiana, Switzerland.

0730-0301/2018/8-ART123 $15.00
https://doi.org/10.1145/3197517.3201322

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

Additional Key Words and Phrases: digital drawing, haptic, perception, stylus, contact simulation

ACM Reference Format:
Michal Piovarči, David I. W. Levin, Danny M. Kaufman, and Piotr Didyk. 2018. Perception-Aware Modeling and Fabrication of Digital Drawing Tools. *ACM Trans. Graph.* 37, 4, Article 123 (August 2018), 15 pages. https://doi.org/10.1145/3197517.3201322

1 INTRODUCTION

Everyone draws and everyone writes: some with skill and precision, while many of us simply aim to communicate and record. All of us begin by writing and drawing with simple yet highly effective hand tools, e.g. crayons, pencils and pens, technologies that have largely remained unchanged for well over a century [Petroski 1992; Van Dulken 2002]. As we turn to digital devices, styli and tablets offer the promise of integrating our ease and comfort in the affordances of hand implements with the advantages of software tools.

Many previous obstacles to stylus adoption are gone. Stylus response time, pressure sensitivity, stroke capture and palm rejection have all been rapidly improving. Concurrently, drawing apps like *Paper by FiftyThree*[1] have successfully focused on replicating the wide variety of physical marks made by differing implements and materials. However, the complementary goal of emulating the physical *feel* of real-world drawing and writing materials with styli remains an important and ongoing challenge.

[1]https://www.fiftythree.com/

A number of recent commercial and research efforts have focused on constructing new stylus devices and tablet surfaces to provide better feels-like haptic response behavior [Annett et al. 2014; Cho et al. 2016; Wang et al. 2016] for the stylus-to-tablet interaction. The starting point in all cases is the simple-to-state but challenging goal of replicating the "hand feel" of pens and pencils. These are by far the most used writing and drawing implements and users consistently report the lack of similar feedback in digital styli to be a hindrance to their adoption of styli [Annett et al. 2014].

To address these needs we take initial steps towards enabling the consistent and systematic design of *passive feedback* styli that feel like targeted drawing and writing tools. We propose a suite of computational models and a method for evaluating, designing, and optimizing stylus components that emulate the physical feel of drawing tools.

Passive feedback is the natural response given by the combined physical interactions of the stylus in frictional contact with the tablet surface as experienced by the hand gripping the stylus. In turn, this is mediated by a large number of complex and coupled physical phenomena. These include the mechanical vibrations the tablet surface excites in the stylus, combined with frictional stiction, material wear and deposition, stick-slip behavior and viscous damping at the tip that changes (e.g., Coulomb's Law) with the amount of pressure applied by the hand.

Each implement thus has distinct characteristics that lend themselves to different applications, e.g., broad hatch and fill marks vs precise lining, and they each have distinctly different haptic behaviors that identify them to the user. Thus there is an intimate connection between the way a drawing tool feels and how we use it. Without this haptic feedback we lose much of the finely tuned ability we have long integrated to make controlled marks [Danna and Velay 2015]. However, fidelity rendering of haptic feedback remains a challenging and open research problem [Choi and Tan 2005].

It is equally important to note that mark-making styles vary with how *controllable* an implement is - the smoother the response the more gestural and loose a marking we make, e.g., for artwork; while in the midground reasonably frictional response enables fine control, e.g., for note taking and mechanical drawing; finally, going all the way to extremely rough feedback creates another range of more loosely controlled artistic gestures associated with noisy pattern generation for textures; e.g., with charcol and chalks [Mayer 1991].

Recent commercial products, including the Apple Pencil[2], the Microsoft Surface Pen[3], the Wacom stylus[4], the PaperLike cover[5], and the reMarkable system[6] similarly focus on leveraging *passive* feedback with their main goal the faithful reproduction of these haptic sensations from drawing and writing with real materials [Williams 2015]. These styli and surfaces have, in almost all cases, been painstakingly designed by trial and error: swapping in and hand testing differing stylus and surface pairs and iterating over shapes and materials[7]. Even so, as we will show in Section 6, current tools only cover a small subset of the range spanned by even standard drawing materials. Embedded actuation devices have also been explored [Cho et al. 2016; Romano and Kuchenbecker 2012; Wang et al. 2016]. However, current latencies in available actuation puts these methods well out of reach for realistic haptic-rate feedback response to stroke gestures.

Our goal is to enable the automated design of styli that will passively deliver the haptic cues of all preferred drawing and writing implements; or even to blend between multiple tools to create styli that respond to mark making with novel feedback.

To investigate perceptually-relevant properties of drawing tools we start with physical measurements of a representative set of traditional drawing tools (Section 3). Next, we embed these tools into a perceptual space where the perceived distances between different tools are defined by the Euclidean distance between them (Section 4). To this end, we conduct an extensive psychophysical experiment. The results of this study are analyzed using our novel method that jointly optimizes for the perceptual space and its correlation to physical parameters of the drawing tools. Critically this correlation with the physical properties enables a direct application to stylus design and fabrication. To design and evaluate new drawing tools without need for fabrication and measurement, we propose a new method for simulating the interaction between drawing tools and substrates (Section 5). Since many of the physical phenomena governing the response behavior are complex and expensive to model [Otaduy and Lin 2004; Romano and Kuchenbecker 2012], we propose a data-driven simulation method governed by our previous measurements. Thus the primary contributions of this work include:

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

2 PREVIOUS WORK

Our work combines insights from perception, haptics, and physical simulation. In this section, we provide an overview of previous work and its relation to our method.

2.1 Human Perception

While the perception of drawing tools is a largely unexplored research area, many works have considered texture exploration using a rigid probe [Klatzky and Lederman 2002]. In this context, the most dominant attribute is roughness – the perception of which is primarily governed by vibrations of the tool [Klatzky and Lederman 2008]. Vibrations are sensed by Pacinian corpuscles [Hollins et al. 2006; Klatzky et al. 2003; Yoshioka and Zhou 2009] – mechanoreceptors

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

[7]https://paperlike.com/

sensitive to high-frequency vibrations with a peak sensitivity at 160-320 Hz [Israr et al. 2006a]. The vibrations sensed through the tool depend on the properties of the tool and the surfaces. Klatzky and Lederman [2002] studied surfaces composed of raised pins. They demonstrated that along with surface properties, i.e., spacing of pins, tool size, exploratory speed, and applied force influence roughness judgments. Furthermore, their results suggest that probe size and pin spacing are primary factors influencing the vibrations and so perceived roughness. Bensmaina et al. [2005] proposed to quantify the aggregated effect of the above factors on roughness perception using the spectral power of the vibrations produced by the probe.

Yoshioka et al. [2007] conducted user experiments in which subjects explored a wide range of real materials using a probe. They found that primary driving attributes for texture perception are roughness, stickiness, and compliance. They also proposed to model these attributes using logarithms of vibratory power, the coefficient of friction, and relative compliance, respectively. This supports the earlier findings that perception of vibration, force, and compliance follow Weber's law [Israr et al. 2006b; Jones and Hunter 1990; Nisky et al. 2011; Pongrac 2008].

In contrast to the above studies, we investigate the perception of drawing tools. Consequently, we account for both differing surfaces and probes (styli). Similarly to Yoshioka et al. [2007], we explain the perception using measurable physical properties; however, we also focus on critical additional properties including geometry and material of styli. We demonstrate that these attributes are crucial in the context of styli design.

2.2 Perceptual Spaces

Perception of varying physical properties can be analyzed via a perceptual space that embeds a set of stimuli in a space where the distances between the samples correspond directly to the perceived similarity between them. Perceptual spaces are commonly constructed using non-metric MDS [Wills et al. 2009] which computes an embedding based on relative similarity comparisons between stimuli pairs. The method is usually considered to be more stable and reliable than standard MDS [Piovarči et al. 2016; Wills et al. 2009] since it is based on a forced-choice experiment rather than magnitude estimation. Usually, the dimensions of a perceptual space can be correlated with physical properties which helps place new samples in the space without additional experiments.

Perceptual spaces have been widely applied to analyze haptic perception, e.g., roughness [Bergmann Tiest and Kappers 2006; Hollins et al. 2000], compliance [Piovarči et al. 2016; Tiest and Kappers 2009], and shape and texture [Cooke et al. 2006, 2010]. Beyond haptics, the methodology has been applied to understand the perception of reflectance properties, such as gloss [Pellacini et al. 2000; Wills et al. 2009], translucency [Gkioulekas et al. 2013] and more intuitive animation controls for artists [Sigal et al. 2015]. Our work computes the perceptual space of digital drawing materials using the non-metric MDS approach. To find the space, we adapt a likelihood-maximization method [Silverstein and Farrell 2001] which was proposed for scaling one-dimensional data. In contrast to previous work which treats correlating the perceptual axes with physical properties as separate step, we propose to modify multi-dimensional scaling to jointly optimize for a perceptual space which both explains the experimental data and correlates with a set of physical properties of the drawing tools.

2.3 Digital Drawing Solutions

Over the years a variety of haptic devices have been proposed ranging from Phantoms to tactile displays, each with different limitations. Phantom devices [Massie et al. 1994] are capable of reproducing multiple forces acting upon the human hand or finger at once. Since a device handles all interactions, its limited force output leads to a poor reproduction of solid surfaces. Tactile displays [Chouvardas et al. 2005] aim at controlling properties of a surface when inspected with fingertips. The approaches vary from stationary and wearable pin-based displays [Hayward and Cruz-Hernandez 2000; Perez et al. 2017] to those that rely on micro-vibrations introduced to the surface [Bau et al. 2010; Kim et al. 2013]. For a more in-depth overview of these devices, please refer to Chouvardas et al.'s [2008] excellent survey. In this work, we are concerned with reproduction of sensations generated by a surface in combination with a probe that, in our case, mimics a drawing tool.

Commercial Solutions. Several commercial solutions directly address the challenge of creating digital drawing tools which replicate the sensation and experience provided by the traditional materials. Products such as Microsoft's Surface Pen, Wacom's Intuos, and reMarkable, offer nibs that are designed to replicate different drawing materials. Despite these efforts, users notoriously report lack of proper feel and dissimilarity to the traditional materials [Annett et al. 2014]. Some products try to improve these technologies. The PaperLike cover is a screen protector that significantly enhances the interaction of a stylus with the tablet surface. Other products, such as the Wacom Bamboo, focus on digitization by allowing users to draw on paper atop a force sensitive tablet. iSkn provides a similar solution, but applies an additional ring, attached to drawing tools, to recover orientation. This enables simulation of more realistic strokes. Although these solutions often provide good haptic experience, they do not offer the full range of the advantages of digital drawing.

Research Prototypes. Researchers have sought to improve haptic feedback of styli via active pens that are equipped with motors for creating artificial vibrations [Arasan et al. 2013; Lee et al. 2004; Poupyrev et al. 2004]. Researchers have also investigated additional modalities of digital styli, e.g., pen bending [Fellion et al. 2017] and specific hand grips [Song et al. 2011], to provide richer input capabilities. In the context of drawing, Romano and Kuchenbecker [2012] used the concept of active styli and proposed a fully data-driven system where interactions are recorded via accelerometer, and then, reproduced using a pair of vibration motors based on speed and pressure data provided by the stylus. This system could recreate the haptic sensation of a wide range of different surfaces. Similarly, Cho et al. [2016] recorded several tools and designed a system for replaying the feedback. Wang et al. [2016] proposed to replace an active stylus with an active surface and use electrovibrations of the drawing surface to control its friction properties. However, active solutions require additional electronics which complicate the design, restrict the usage of these devices, and introduce significant latency

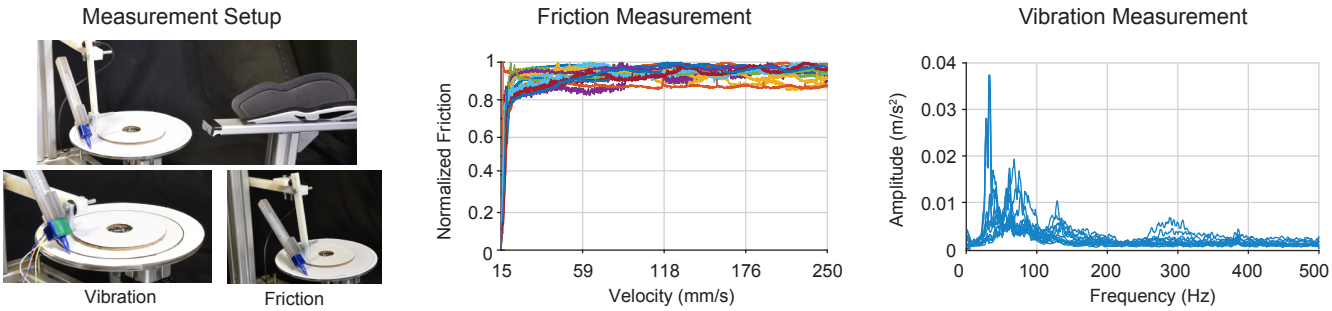

Fig. 2. Measurements and perceptual user study setup. Turntable on which a mechanical arm is used to capture friction and vibration of drawing tools. There is a holder for a human arm (during user studies) to provide consistent grip and orientation. In the middle, we can see normalized friction measurements recovered using our setup. On the right is the accelerometer response for a fixed velocity.

[Annett et al. 2014; Helps and Helps 2016]. Our system for designing a passive stylus mitigates the problem of additional electronics and provides instantaneous adaptation of feedback to changes in stylus orientation and applied pressure.

2.4 Simulating Surface Interaction

Standard friction models generally begin with Coulomb's model [Harnoy et al. 2008] – an assumption of frictional resistance proportional to normal load and in opposition to velocity. These assumptions are often reasonable for solid-to-solid frictional contact. In the presence of lubrication, however, viscosity effects become important. The Stribeck effect [Harnoy et al. 2008] describes frictional behavior in lubricated contact. Initially, with increased speed, friction forces drop. As speed is further increased the Stribeck model assumes an increase in friction linearly proportional to speed. Frictional forces can generate nonsmooth stick-slip behaviors. Karnopp et al. [1985] model a regularized stick-slip behavior that removes nonsmoothness from the Coulomb model, simplifying numerical integration. A wide range of more complex frictional models [Armstrong-Hélouvry 1991; Canudas de Wit et al. 1995; Dahl 1976] are available and range in suitability depending on modeling needs. Our particular problem, that of stylus-surface interaction is further complicated by the effects of material wear [Stachowiak 2006], material deposition [Archard 1953] on the surface and the coupling between friction and contact with elasticity and viscous damping [Chen et al. 2017]. Applying a standard frictional contact modeling is not sufficient to obtain accurate numerical estimates of stylus behavior. Our simulator builds on impulse-based frictional contact modeling [Mirtich and Canny 1995]. This remains an active research area in simulation [Bertails-Descoubes et al. 2011; Kaufman et al. 2005, 2008] and many of the additional complicating factors described above are not yet well understood in this context. To overcome this difficulty, we opt for an exponential integrator [Hochbruck and Ostermann 2010; Michels et al. 2014] with a data-driven surface-interaction model. In combination, our exponential integration gives us an analytical solution to the linear elastodynamic equations of motion while achieving excellent match with experimental data captured during stylus-surface interaction.

3 PHYSICAL MEASUREMENTS

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

3.1 Measuring setup

In our work, we follow observations from the literature (Section 2) that haptic sensation is induced primarily via resistance and vibration generated by the tool-surface interaction. Consequently, our device (Figure 2) measures resistance and vibration transferred to fingertips while drawing. The base of the device is a turntable operated by a DC motor with controllable speed to which different types of drawing substrate can be attached. To simulate an artist drawing on a surface, we design an arm with an enclosure for various drawing tools. By adding extra weight, the device can simulate different pressures applied by a user. The vibration of a tool is measured by an accelerometer attached to the enclosure, while the resistance is captured by a force sensor placed on the arm.

3.2 Defining measurement parameters

The vibration and the resistance produced by a drawing tool depend on its speed and the pressure applied. To reflect standard drawing scenarios, we first performed a pilot study to determine these parameters. We invited four amateur artists to a short drawing session and asked them to draw basic shapes [Garcia 2003] (Figure 3) using an iSKN[8] tablet. The participants used five different drawing materials: ballpoint pen, soft (8B) and hard (2H) pencil, fine-liner, and charcoal, on three different drawing surfaces: standard 80-gram office paper, rough paper for pencil drawing, and smooth stone paper (Figure 4). During each session, the velocity of the tool was recorded using the positioning system of the tablet, while the pressure applied by the participants was recorded by a force sensor mounted below the tablet. To eliminate the effect of hand pressure, we placed a hand rest next to the tablet. Figure 5 shows the histograms of recorded velocities and forces. Based on the results, we decided to perform further measurements of drawing tools using the mean force applied by the artists, i.e., 2.2 N, and velocities between the 5th and 95th percentile, i.e., 17-250 mm/s.

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

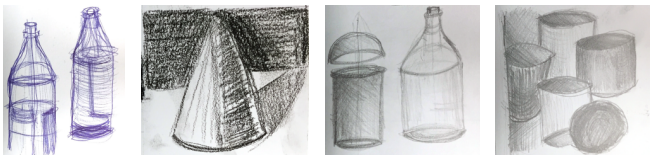

Fig. 3. Sample images drawn during our preliminary user study.

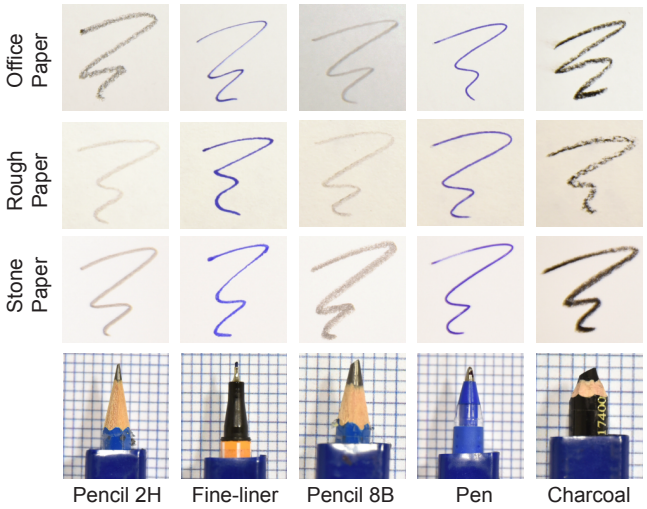

Fig. 4. Our user study considers 5 drawing tools: 2H and 8B pencils, fine-liner, ballpoint pen, and charcoal and 3 drawing substrates: 80-gram office paper, rough artist paper, and stone paper.

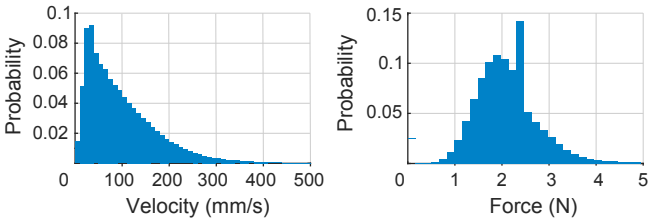

Fig. 5. Velocity and pressure histograms recorded from free drawings during our pilot study.

3.3 Measurements

Using our device, we measured each drawing tool and surface combination from our preliminary study (Figure 4). The normalized frictional force measurements (Figure 2, middle), revealed a small interaction between velocity and the friction coefficient. Consequently, we decided to use Coulomb friction to model measurements, and denote a single frictional coefficient for drawing tool and particular surface pairs. The recovered frictional coefficients range from 0.13 to 0.33. The vibration measurements (Figure 2, right) revealed a broad-band characteristic with significant dependence on velocity. We captured the complex vibration characteristic for each combination of tool and surface using a velocity dependent spectrogram of vibratory response. For more details about the measurements as well as all captured data please refer to supplementary material.

Since the goal of the experiment is to measure vibration caused by tool-surface interaction, it is critical to assure that the measurement device itself does not produce significant vibration. To measure the vibration produced by our device, we experimented with a pen on an oiled acrylic sheet. Figure 6 demonstrates vibration spectrograms for the oiled surface and standard drawing paper. It can be observed that, although very weak, there is a vibration produced by the system on the oiled surface (due to the DC motor) which manifests as a diagonal line on the spectrogram. The weak horizontally structured signals, since not velocity dependent, are most likely due to the remaining interaction of the pen with the oiled surface. This is possible since our base surface is not perfectly smooth, and some asperities exist which actuate a drawing tool. We verified that the vibration caused by the system does not affect the perception of the drawing tools in an informal test during which the acrylic surface was perceived as smooth and vibration-free. For more comparisons of measurements on paper and oiled acrylic please see the supplementary material.

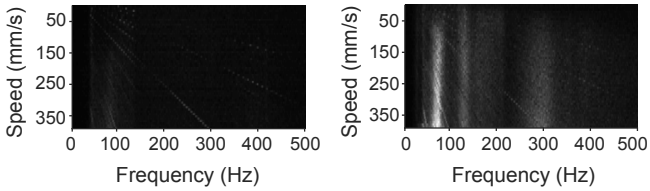

Fig. 6. Measurements of a ballpoint pen on oiled acrylic (left), and 80-gram office paper (right).

4 PERCEPTUAL SPACE OPTIMIZATION

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

4.1 Recovering the Perceptual Space

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

The experimental data required for computing the space consists of triplets of stimuli A, B, and C, with human judgments about which of A and C is more similar to B. To find the space, we want to find positions of our stimuli in an n-dimensional space such that the probability of our experiment occurring is maximized. Under assumptions of Thurstone Case V law of comparative judgment [Thurstone 1927], the probability can be computed as:

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

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

$$P_{ijk} = 1 - \exp\left(-10.38\, d_{ijk} + 111 \arctan\left(0.09\, d_{ijk}\right) + 1\right)^{-1} \quad (2)$$

$$d_{ijk} = ||Q_i - Q_k|| - ||Q_j - Q_k||, \quad (3)$$

where Q_i, Q_j, and Q_k are locations of stimuli i, j, and k in the perceptual space.

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

$$P_{corr} = \prod_d (Corr[Q^d, D^d]), \quad (4)$$

where Q^d is a vector of d-th coordinates of points Q in the perceptual space and D^d are physical properties we wish to correlate with Q^d . To compute a perceptual space, we combine Equations 2 and 4 to formulate an optimization problem:

$$\arg\max_Q P_{exp} \cdot P_{corr}^{\lambda}, \quad (5)$$

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

$$\arg\min_Q -\log(P_{exp}) - \lambda \log(P_{corr}). \quad (6)$$

To resolve translational ambiguity, we constrain a single element to the origin of our space. In contrast to standard, non-metric MDS methods, our formulation does not have scale and rotational ambiguity. Scale is constrained by Equation 2 relating probabilities P_{ijk} with distances in the space, while rotation is handled by our correlation term (Equation 4). Our full optimization problem is then nonconvex [Silverstein and Farrell 2001]. To find a perceptual space minimizing Equation 6, we apply a quasi-Newton solver [Avriel 2003] with multiple random starting points.

4.2 Experiment Design

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

4.3 Experiment

Our experimental apparatus was an extended version of our measurement device presented in Section 3. Besides the turntable on which we can place different samples of drawing surfaces, we added a support for the participants hands. To prevent visual feedback, which could affect the similarity judgments, we installed a black curtain that separated the participant from the turntable. Our experimental setup is demonstrated in Figure 7.

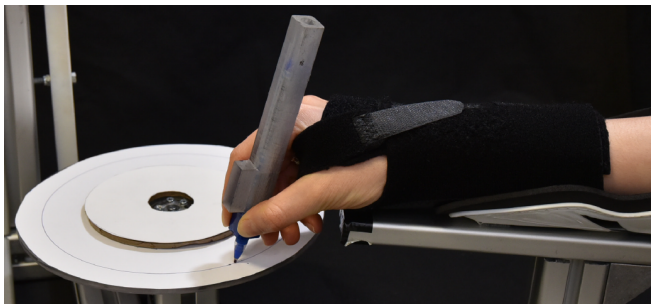

Fig. 7. Turntable setup for user studies. Participants' hands rest in a holder and an arm-wrap is used to limit wrist motion.

Stimuli. The stimuli consisted of 15 combinations of five drawing tools and three different kinds of paper that were previously measured (Section 3), which resulted in 1365 triplets. Half were presented to the participants twice during the first stage of our experiments. From this set of triplets, we selected 60 stimuli which resulted in a tie and used them in the second stage where each of them was evaluated ten times. For the optimization of the perceptual space we used all the triplets from the second stage and those that did not result in a tie in the first stage.

Task. Each trial involved investigating and comparing three combinations of drawing tools and papers. During each trial, the participants were asked to sit in front of our apparatus and to rest one

of their hands on the support. Then, the instructor demonstrated each of the three stimuli. This required changing the surface on the turntable, sharpening the tools which undergo wear, and handing the tools to the participants. Next, the subjects were asked to lower the drawing tool until it reaches the surface. At this point, the instructor activated the turntable which was rotating with a previously determined speed range from 17 up to 250 mm/s (Section 8). The full period of the velocity change was 3 seconds which was sufficient to appreciate the velocity-dependent effects while keeping the length of the user-study short. The subjects were instructed which of the stimuli were the reference and the tests, and then asked to identify which of the two test stimuli was more similar to the reference. Before answering the question, the participants could investigate each of the stimuli an unlimited number of times. To prevent any visual and auditory feedback, the participants were asked to sit behind the curtain and to wear noise-canceling headphones.

Participants. For the first stage of our experiment, we invited 34 participants (20-30 years old, M/F ratio 20/14). Each of them performed an equal number of different randomly chosen comparisons. For the second stage, we asked 10 new participants (20-30 years old, M/F ratio 5/5) and each of them evaluated all 60 curated samples that in the first stage resulted in a tie. Evaluation of one triplet took on average 30 seconds, and the whole study took approximately 60 minutes due to the additional setup performed by the instructor. Due to the length of the study, the participants were free to take a break at any time or even split the study into multiple sessions. All the participants received financial compensation.

4.4 Perceptual Space of Drawing Tools

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

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

To determine the dimensionality of the space and a good value of the parameter λ, we used the accurate estimations of probabilities from the second stage of the experiment and compared them to the predictions given by perceptual spaces computed using different λ values and dimensionality. Figure 8 visualizes the mean match error defined as the average error in prediction of pairwise distances. We can observe an apparent gain in performance when using a 2D space which quickly tapers and does not significantly improve with higher dimensions. Based on this analysis, we decided to use a 2-dimensional space. The best one correlates with frictional coefficient and a mean value of vibration spectrogram, with a linear correlation of 0.98 and 0.95 respectively.

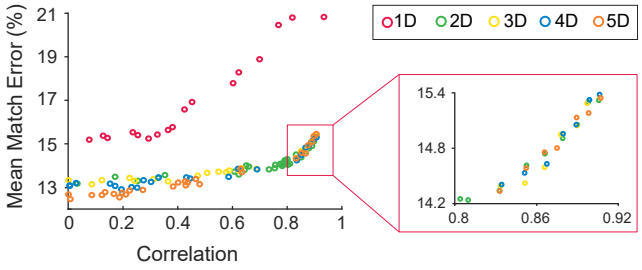

Fig. 8. The plot shows the match of our optimized perceptual spaces to the experimental data. Different colors correspond to spaces of different dimensionality, while points with the same color correspond to different values of λ.

We present our recovered perceptual space of drawing tools in Figure 9. As expected, it forms clusters of the same kind of drawing tools, but also captures differences related to using different papers.

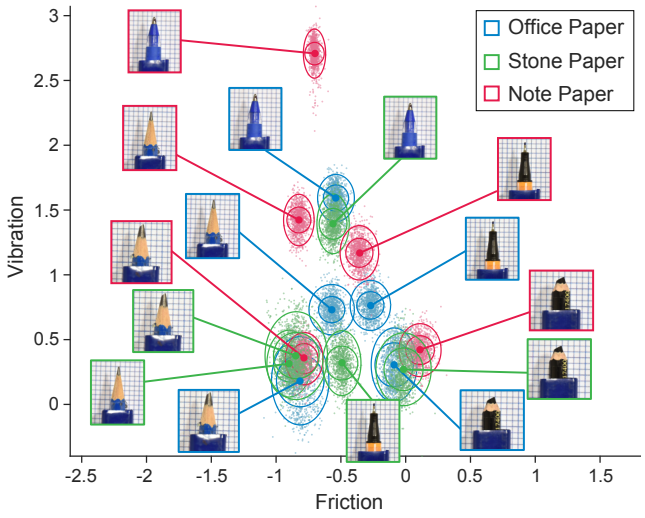

Fig. 9. The perceptual space obtained using our data and optimization. The axis are correlated with vibration and friction measurements. Confidence intervals obtained by bootstrapping visualize 95% and 68% regions.

4.5 Accuracy

We performed two tests to further validate the accuracy of our space. First, we performed 5-fold validation. We splited the data from the second stage of our experiment into training and testing. Next, we optimized for 2-dimensional spaces using only training sets and

use the testing for validation. An average error was 17 % which is similar to the error obtained in our dimensionality test (Figure 8).

To test the reliability of our perceptual space, we computed confidence intervals of tool placement using bootstrapping. We performed random sampling of the experimental data and used the re-sampled data to generate a new space which was then aligned with our original perceptual space (Figure 9, points). We repeated this procedure 1000 times and generated confidence intervals by drawing ellipsoids that enclose 95% and 68% of points corresponding to the same tool (Figure 9, ellipses). In most cases, the estimated confidence intervals were smaller than the distances between the individual tools. This suggests that the placement of the tools is reliable. The differences in confidence intervals can be explained by wear characteristics of different tools. For example, ballpoint pen and multiliner are resistant to wear and do not change over time. On the other hand pencils and charcoals get dulled by wear. This results in change of tip size affecting the physical properties of the tool, which introduces noise to human judgments.

We also validated our optimization and data selection approach on synthetic data sets. To this end, we designed 100 synthetic tests. Each of them consisted of a random perceptual space with a fixed dimensionality, which we used to emulate our triplet selection process and a user study in a Monte-Carlo fashion. Finally, we used the synthetic data to recover the perceptual space. Figure 10 shows the mean and standard deviation of the match between optimized and ground-truth spaces for the various dimensionality of the initial data. The results demonstrate the the error introduced by our optimization technique is smaller than the errors reported for our space of drawing materials. In Section 6, we further demonstrate that despite all the inaccuracies reported in this section, our space can facilitate the process of designing drawing tools.

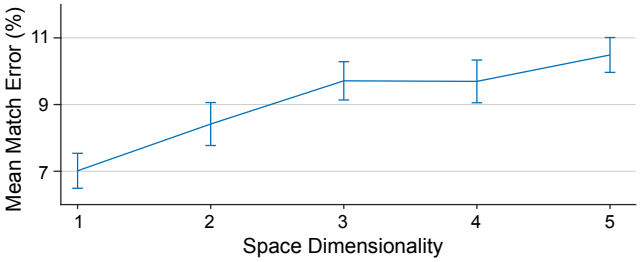

Fig. 10. Synthetic results of optimizing perceptual spaces using our experiment design as a function of space dimensionality. The bars show mean error and standard deviation.

5 PHYSICAL SIMULATION

We now have a procedure to estimate the placement of a tool in our perceptual space given a set of physical measurements. Generating these measurements from a physical setup is inconvenient so instead we turn to physical simulation. A full contact simulation of a drawing tool with paper micro-asperities is overly expensive and many effects are still open problems in computational modeling, (Section 2.4). These difficulties inspire us to take a partially data-driven approach, wherein we model as much of the problem as we can mathematically, using fitted models to fill in the gaps. Our method is motivated by approaches in musical instrument simulation [Blood 2009] in which the main resonator of an instrument is simulated but the complex driving force (the mouthpiece) is tackled in a data-driven fashion [Li et al. 2016]. More precisely, we model the propagation of the vibration initiated at the tip of the tool using our exponential Euler integrator, while the forces acting at the tip of the tool are generated using our data-driven approach. The data-driven forces encode the complex contact characteristic between the tool and the surface.

5.1 Exponential Euler Integrator

Our goal is to simulate the vibrational behaviors of elastic geometries in driven contact with rough surfaces. The tools and materials we work with exhibit high stiffness, as reflected in their Young's moduli, and so are effectively modeled as linearly elastic using a small deformation assumption. The equations of motion of such the physical system is then:

$$\mathbf{M}\ddot{\mathbf{x}}(t) + \mathbf{D}\dot{\mathbf{x}}(t) + \mathbf{K}\mathbf{x}(t) + \Lambda(\mathbf{x}(t)) = \mathbf{0}, \tag{7}$$

where $\mathbf{x}(t)$ is the displacements at time t. $\mathbf{M}$, $\mathbf{D}$, and $\mathbf{K}$, are the mass, damping and stiffness matrices, respectively, and $\Lambda(\mathbf{x}(t))$ is the vector of external forces.

In this regime, we find using only the stiffness term of Rayleigh damping model works well for our application, $\mathbf{D} = \lambda\mathbf{K}$. In this case, these equations admit an analytical solution given as

$$\mathbf{X}(t) = e^{(t-t_0)\mathfrak{U}}\mathbf{X}(t_0) + e^{t\mathfrak{U}}\int_{t_0}^{t} e^{-\tau\mathfrak{U}}\Gamma(\mathbf{X}(\tau))\,\mathrm{d}\tau, \tag{8}$$

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

Our ability to perform predictive simulation of real-world drawing tools hinges on the representation of the Γ forcing term in Equation 8 and our capacity to integrate it efficiently. In our setting this term bundles a wide range of complex phenomena that drive the drawing tool interaction behavior. The first of which are the oscillations induced by micro-scale surface contact. Given the structure of our problem we adopt the spectral perspective and represent the oscillatory driving force as a linear combination of sinusoidal terms of varying phase. This, in turn, enables efficiency and flexibility in simulation and modeling.

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

$$\begin{aligned}\mathbf{X}(t+\Delta t) =& e^{\Delta t\mathfrak{U}}\mathbf{X}(t) + (\Omega^2 + \mathfrak{U}^2)^{-1}(\\ & e^{\Delta t\mathfrak{U}}(\Omega\mathbf{A}\cos(\omega t+\phi) + \mathfrak{U}\mathbf{A}\sin(\omega t)+\phi)\\ & - \Omega\mathbf{A}\cos(\omega(t+\Delta t)+\phi)) - \mathfrak{U}\mathbf{A}\sin(\omega(t+\Delta t)+\phi).\end{aligned} \tag{9}$$

Here and in the following Ω designates an identity matrix scaled by phase ω, $\mathbf{A}$ a diagonal matrix of per degree of freedom amplitudes, and ϕ is the frequency shift.

Additionally, this representation also allows a set of solutions to the single sinusoidal term problem to forms a linear basis that can

be applied to represent any signal made up of a linear combination of these sinusoids. Practically, we exploit this during design tasks by precomputing a simulation basis for a drawing tool and then synthesizing its interaction with new surfaces by summation of these pre-simulated results. For more details about the derivation please refer to Appendix A.

To validate our simulation and gather parameters for our method, we run a test simulation. We 3D printed a bar on Objet260 printer using VeroClear material. The bar was 210 × 20 × 2 mm and was rounded at the end (radius 10 mm) with a circular cutout of 5 mm radius located 10 mm from the end. The bar was clamped to a table on a 30 mm section, preloaded to 5 cm and released. First, we captured the bar's oscillations using an attached accelerometer. Next, material parameters (Young's modulus and Rayleigh damping) were optimized to fit our measurements. The recovered numerical Youngs modulus was 0.45 GPa and the stiffness term of Rayleigh damping was 1.25e−3. Figure 11 demonstrates the excellent agreement between simulation and experiment.

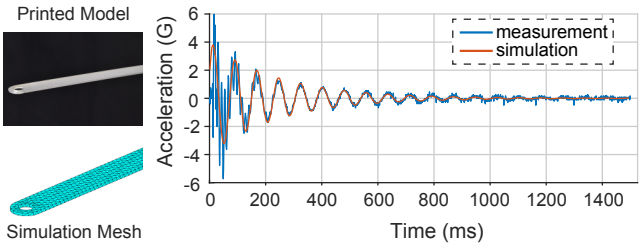

Fig. 11. Oscilating 3D printed bar captured with accelerometer (blue) and physical simulation of the bar (red).

5.2 Recovering the Forcing Signal

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

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

We compute parameters for our model by solving the optimization:

$$\underset{F_S * t_\Psi * m_f}{\arg\min} \sum_{\Psi} \sum_{S} \left\| Integrator(\Psi, F_S * t_\Psi * m_f) - M_{\Psi,S} \right\|, \quad (10)$$

where Ψ is a shape and material combination, F_S is forcing term of surface S, t_Ψ, and m_Ψ are tip and material scaling parameters. Here integrator is the Exponential Euler integrator we derive above. The integrator takes as an input the assembled forcing term and returns a simulated measurement. We minimize the L2 norm between our simulation and the measured ground truth $M_{\Psi,S}$. The required simulations can be quickly evaluated using our pre-simulated basis results.

We printed, measured, and simulated a total of 27 tools with a unitary forcing signal in the range of frequencies relevant to haptic feedback (1-500 Hz). We then optimized this data with Equation 10 to acquire our final material forcing terms.

To validate the data-driven model we performed a leave one out test of the recovered forcing term. One exemplar was removed from the data set, our model was trained on the remaining samples and then evaluated by fitting the test exemplar, (Figure 12). We can see that we achieved good match with testing data. The main discrepancies come from sharp peaks that are very challenging to fully reproduce. For full evaluation please refer to the supplementary material. Next, we also evaluated the effects of the scaling terms, (Figure 13). The material scaling factor (Figure 13 left) is inversely proportional to material softness. A soft material is predicted to damp the vibrations of the tool. However, increased softness leads to more pronounced stick-slip behavior which is then compensated for by increasing the material scaling term. The effect of tip size (Figure 13 right) is more subtle. In general large tip size leads to more vibration damping. The scaling factors for tip size of 1 and 2 mm are very similar. This is caused by the wear experienced by the material which sands down the finer tip quicker, therefore, making them equivalent.

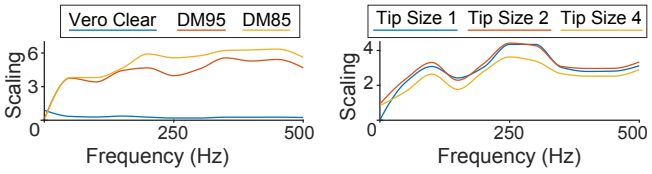

Fig. 13. Material (left) and tip (right) scaling factor. Both factors were optimized jointly.

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

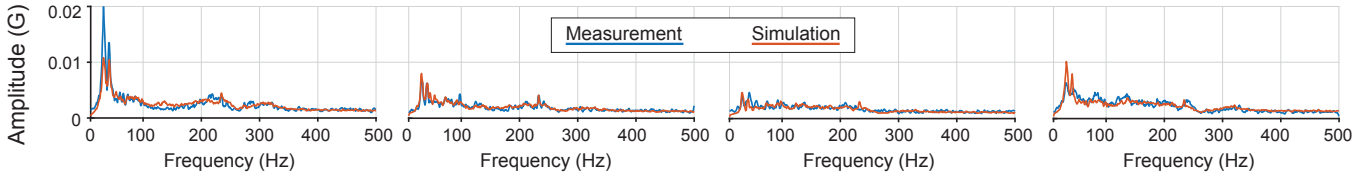

Fig. 12. Fit of testing dataset at a fixed velocity. We compare physical measurements (blue) to our data-driven simulation (red). We can see a good match of estimated vibration for the testing dataset.

5.3 Results

The purpose of the simulation is to predict the behavior of 3D printed tools without physically manufacturing them. To evaluate the precision of our simulation we conducted a test in which we predicted the vibratory response of an interpolated tool. We interpolated between two designs. The first design has a half-sphere tip of 4 mm in diameter printed in DM85. The second design has a half-sphere tip of 1 mm diameter and is printed in VeroClear. Both designs were measured on standard office paper at various velocities (Figure 14 blue). The new tool has linearly interpolated material and tip parameters. To predict its behavior we used our full simulation pipeline. First, we measured the response of the original designs to recover tip and material scaling factors. Next, we used our exponential Euler integrator to simulate the response of the interpolated tool. We used forcing term we recovered for office paper modulated by linearly interpolated material and tip scaling factors (Figure 14 red). We compare our full simulation pipeline to a simulation where we would not interpolate the material and tip parameters (Figure 14 gray). We can see that our prediction matches reality well. The deviations can be attributed to the general noisiness of vibrational measurements, as well as, our printing process since we used dithering to generate the interpolated material.

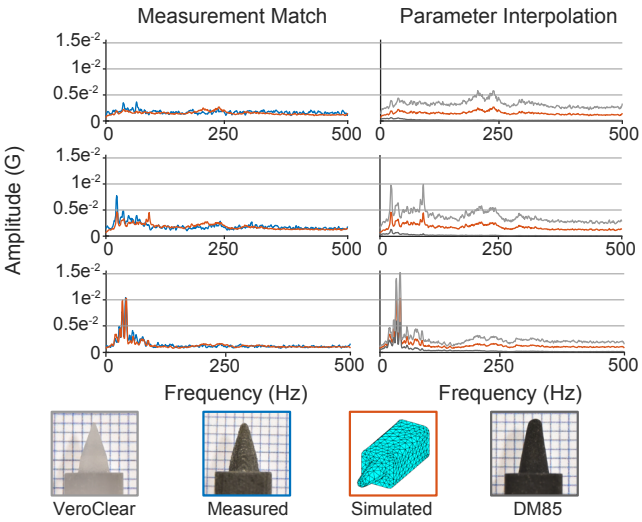

Fig. 14. An interpolated design was measured (blue line) and simulated using our pipeline (red line). We also compare to results of our simulation if we would not interpolate tip and material scaling parameters (gray lines).

6 APPLICATION FOR DRAWING TOOL DESIGN

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

6.1 Perceptual Space Exploration

Because we correlate the dimensions of the perceptual space with measurable properties of drawing tools, we can easily embed new tools into the space and evaluate their pairwise similarities based on their measurements or simulation. To demonstrate this, we extended our previously derived space with several other drawing tools. We considered four different categories. The first group consists of our initial set of traditional drawing tools extended by two new materials: a crayon and a felt-tip marker (Figure 15, red). The second category is 3D-printed tools (Figure 15 green). They were fabricated on a Stratasys Objet 260 Connex printer using different materials ranging from VeroClear (low friction) to DM85 (high friction). The tip size of the tools varied from 1 to 4 mm. This category also includes tools that were covered with a Teflon tape to lower the friction further. All tools in the first and the second category were used on our three different kinds of paper (Figure 4). The third group consisted of the same 3D printed tools as in the previous category, but this time they were used on two different artificial surfaces: glass and glass with a screen protector (Figure 15 blue). This category of drawing tools demonstrates what one can achieve with today's tablets and multi-material printing technologies. We also included commercial solutions such as a Wacom Tablet Pen, Microsoft Surface Pen, and Apple Pencil (Figure 15 yellow).

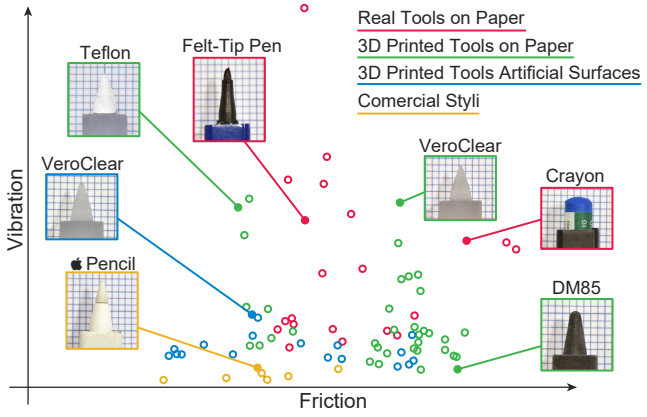

Fig. 15. Perceptual space grouped by means of fabrication: real tools (red), 3D printed tools on paper (green), and artificial substrates (blue), and commercial styli (yellow).

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

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

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

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

Results. We compared the results of the experiment with predictions provided using our space. The prediction matched the popular opinion in 19 out of 24 cases. Z-test revealed that our model predicts preference significantly better than random chance (p-value < 0.0022). On average our prediction had an error of 19% and correlated to the experimental results with a Pearson correlation of 0.81 (p-value < 0.0001). This suggests a strong linear relationship between our predictions and ground truth. The results support our data validation (Section 4.5) and suggest that our methodology translates well to freehand drawing. Figure 16 presents the detailed results of the experiment.

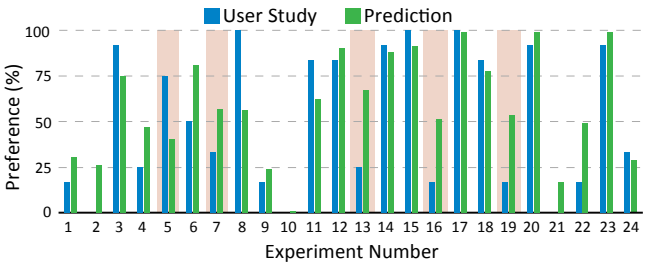

Fig. 16. The results of the free drawing user study. For each out 24 triplets, we plot preference obtained in our validation experiment next to the prediction coming from the perceptual space. The red background indicates the failure cases where the majority vote did not agree with our prediction.

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

6.2 Optimizing Digital Styluses

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

Stimuli. In the first case, we considered a screen protector as our artificial surface. On this protector, we measured two 3D-printed styli, one in VeroClear with a tip diameter of 1 mm, and one in DM85 with a tip diameter of 4 mm. Figure 17 (left) shows their relation to the real drawing tool in our space. One can see that neither of them matches the real tool, and a better design can be achieved. We used a line search powered by our simulation pipeline (Section 5) to generate a design in between the two synthetic styli. For the frictional coefficient, we used linear interpolation. The second case consider an improvement of digital styli on a glass surface typical for tablets. We measured the two 3D-printed designs from the previous experiment. Figure 17 (right) shows the relative placement of the real tool, commercial stylus, and our 3D printed tips. Once again we use our simulation pipeline to design an interpolated stylus.

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

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

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

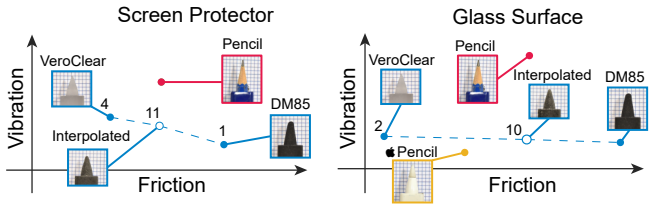

Fig. 17. Details of our perceptual space for optimization of digital styli. We show two cases: interpolation of 3D-printed styli on a plastic screen protector (left), and interpolation of 3D-printed styli on the glass surface in comparison to the Apple Pencil (right). Number indicate subject preference when compared to real pencil.

7 LIMITATIONS AND FUTURE WORK

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

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

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

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

8 CONCLUSION

In this work, we have examined the problem of modeling and fabricating digital styli that have the *feel* of real drawing tools. To this end, we measured, analyzed, and characterized properties of a representative set of traditional tools. We identified features that influence perception and used them to derive a perceptual space of drawing materials. This enabled the evaluation of perceived differences between designs. Furthermore, we proposed a data-driven simulation method that allows us to embed new tool designs into the perceptual space without the need for fabricating them or conducting expensive and time-consuming experiments. Finally, we demonstrated that our approach can aid in the process of designing new styli, and validated our results in user experiments.

ACKNOWLEDGMENTS

We would like to thank Elena Arabadzhiyska and Celso Gomes for their help with the video, Itrat Rubab and Chitra Singh for their help in collecting data during the user studies, and all the anonymous participants of our experiments.

REFERENCES

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 Tevfik 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

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

Mordecai Avriel. 2003. *Nonlinear Programming: Analysis and Methods*. Dover Publications.

Olivier Bau, Ivan Poupyrev, Ali Israr, and Chris Harrison. 2010. TeslaTouch: Electrovibration 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

Sliman Bensmaïa, Mark Hollins, and Jeffrey Yau. 2005. Vibrotactile intensity and frequency information in the Pacinian system: A psychophysical model. *Attention, Perception, & Psychophysics* 67, 5 (2005), 828–841.

Wouter M Bergmann Tiest and Astrid ML Kappers. 2006. Analysis of haptic perception of materials by multidimensional scaling and physical measurements of roughness and compressibility. *Acta psychologica* 121, 1 (2006), 1–20.

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.

Brian Blood. 2009. Physics of musical instruments. (2009).

Carlos Canudas de Wit, Henrik Olsson, Karl J. Astrom, and Pablo Lischinsky. 1995. A new model for control of systems with friction. *IEEE Trans. Automat. Control* 40, 3 (Mar 1995), 419–425. https://doi.org/10.1109/9.376053

Desai Chen, David I. W. Levin, Wojciech Matusik, and Danny M. Kaufman. 2017. Dynamics-aware Numerical Coarsening for Fabrication Design. *ACM Trans. Graph.* 36, 4, Article 84 (July 2017), 15 pages. https://doi.org/10.1145/3072959.3073669

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.

Vasilios G Chouvardas, Amalia N Miliou, Miltiadis K Hatalis, et al. 2005. Tactile displays: a short overview and recent developments. In *Proceedings of the 5th International Conference on Technology and Automation (ICTA'05)*. 246–251.

Theresa Cooke, Sebastian Kannengiesser, Christian Wallraven, and Heinrich H Bülthoff. 2006. Object feature validation using visual and haptic similarity ratings. *ACM Transactions on Applied Perception (TAP)* 3, 3 (2006), 239–261.

Theresa Cooke, Christian Wallraven, and Heinrich H Bülthoff. 2010. Multidimensional scaling analysis of haptic exploratory procedures. *ACM Transactions on Applied Perception (TAP)* 7, 1 (2010), 7.

Philip R. Dahl. 1976. Solid Friction Damping of Mechanical Vibrations. *AIAA Journal* 14, 12 (1976), 1675–1682. https://doi.org/10.2514/3.61511

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

Nicholas Fellion, Thomas Pietrzak, and Audrey Girouard. 2017. FlexStylus: leveraging flexion input for pen interaction. In *UIST 2017 (UIST 2017)*. Québec, Canada, 2482 – 2489. https://doi.org/10.1088/0957-0233/13/10/303

Claire Watson Garcia. 2003. *Drawing for the Absolute and Utter Beginner*. Watson-Guptill.

Ioannis Gkioulekas, Bei Xiao, Shuang Zhao, Edward H Adelson, Todd Zickler, and Kavita Bala. 2013. Understanding the role of phase function in translucent appearance. *ACM Transactions on Graphics (TOG)* 32, 5 (2013), 147.

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.

Marlis Hochbruck and Alexander Ostermann. 2010. Exponential integrators. *Acta Numerica* 19 (2010), 209–286.

Mark Hollins, Sliman Bensmaïa, Kristie Karlof, and Forrest Young. 2000. Individual differences in perceptual space for tactile textures: Evidence from multidimensional scaling. *Perception & Psychophysics* 62, 8 (2000), 1534–1544.

Mark Hollins, Florian Lorenz, and Daniel Harper. 2006. Somatosensory Coding of Roughness: The Effect of Texture Adaptation in Direct and Indirect Touch. *Journal of Neuroscience* 26, 20 (2006), 5582–5588. https://doi.org/10.1523/JNEUROSCI.0028-06.2006

Ali Israr, Seungmoon Choi, and Hong Z. Tan. 2006a. Detection Threshold and Mechanical Impedance of the Hand in a Pen-Hold Posture. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 472–477. https://doi.org/10.1109/IROS.2006.282353

Ali Israr, Hong Z. Tan, and Charlotte M. Reed. 2006b. Frequency and amplitude discrimination along the kinesthetic-cutaneous continuum in the presence of masking stimuli. *The Journal of the Acoustical Society of America* 120, 5 (2006), 2789–2800. https://doi.org/10.1121/1.2354022

Lynette A Jones and Ian W Hunter. 1990. A perceptual analysis of stiffness. *Experimental Brain Research* 79, 1 (1990), 150–156.

Dean Karnopp. 1985. Computer Simulation of Stick-Slip Friction in Mechanical Dynamic Systems. *ASME Journal of dynamic Systems, Measurement and Control* 107, 1 (Mar 1985), 100–103. https://doi.org/10.1115/1.3140698

Danny M. Kaufman, Timothy Edmunds, and Dinesh K. Pai. 2005. Fast Frictional Dynamics for Rigid Bodies. *ACM Trans. Graph.* 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 Trans. Graph.* 27, 5, Article 164 (Dec. 2008), 11 pages. 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

Roberta L. Klatzky and Susan J. Lederman. 2002. Perceiving texture through a probe. *Touch in virtual environments* (2002), 180–193.

Roberta L. Klatzky and Susan J. Lederman. 2008. Perceiving object properties through a rigid link. *Haptic Rendering: Algorithms and Applications* 1 (2008), 7–19.

Roberta L. Klatzky, Susan J. Lederman, Cheryl Hamilton, Molly Grindley, and Robert H. Swendsen. 2003. Feeling textures through a probe: Effects of probe and surface geometry and exploratory factors. *Attention, Perception, & Psychophysics* 65, 4 (2003), 613–631.

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 Trans. Graph.* 35, 4, Article 88 (July 2016), 12 pages. https://doi.org/10.1145/2897824.2925960

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.

Ralph Mayer. 1991. *Artist's Handbook of Materials and Techniques: Revised and Updated*. Viking Books.

Dominik L. Michels, Gerrit A. Sobottka, and Andreas G. Weber. 2014. Exponential Integrators for Stiff Elastodynamic Problems. *ACM Trans. Graph.* 33, 1, Article 7 (Feb. 2014), 20 pages. https://doi.org/10.1145/2508462

Brian Mirtich and John Canny. 1995. Impulse-based Simulation of Rigid Bodies. In *Proceedings of the 1995 Symposium on Interactive 3D Graphics (I3D '95)*. ACM, New York, NY, USA, 181–ff. https://doi.org/10.1145/199404.199436

Ilana Nisky, Assaf Pressman, Carla M. Pugh, Ferdinando A. Mussa-Ivaldi, and Amir Karniel. 2011. Perception and Action in Teleoperated Needle Insertion. *EEE Trans. Haptics* 4, 3 (July 2011), 155–166. https://doi.org/10.1109/TOH.2011.30

Miguel A. Otaduy and Ming C. Lin. 2004. A Perceptually-inspired Force Model for Haptic Texture Rendering. In *Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization (APGV '04)*. ACM, New York, NY, USA, 123–126. https://doi.org/10.1145/1012551.1012574

Fabio Pellacini, James A. Ferwerda, and Donald P. Greenberg. 2000. Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 55–64.

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).

Henry Petroski. 1992. *The pencil: A history of design and circumstance*. Alfred a Knopf Incorporated.

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

Helena Pongrac. 2008. Vibrotactile perception: examining the coding of vibrations and the just noticeable difference under various conditions. *Multimedia Systems* 13, 4 (2008), 297–307. https://doi.org/10.1007/s00530-007-0105-x

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

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

Leonid Sigal, Moshe Mahler, Spencer Diaz, Kyna McIntosh, Elizabeth Carter, Timothy Richards, and Jessica Hodgins. 2015. A Perceptual Control Space for Garment Simulation. *ACM Trans. Graph.* 34, 4, Article 117 (July 2015), 10 pages. https://doi.org/10.1145/2766971

D. Amnon Silverstein and Joyce E. Farrell. 2001. Efficient method for paired comparison. *Journal of Electronic Imaging* 10, 2 (2001), 394–398.

Hyunyoung Song, Hrvoje Benko, Francois Guimbretiere, Shahram Izadi, Xiang Cao, and Ken Hinckley. 2011. Grips and Gestures on a Multi-Touch Pen. 1323–1332. https://www.microsoft.com/en-us/research/publication/grips-and-gestures-on-a-multi-touch-pen/

Gwidon W Stachowiak. 2006. *Wear: materials, mechanisms and practice*. John Wiley & Sons.

Luis L. Thurstone. 1927. A law of comparative judgment. *Psychological Review* 34, 4 (1927).

Wouter M Bergmann Tiest and Astrid ML Kappers. 2009. Cues for haptic perception of compliance. *Haptics, IEEE Transactions on* 2, 4 (2009), 189–199.

Stephen Van Dulken. 2002. *Inventing the 20th Century: 100 Inventions that Shaped the World from the Airplane to the Zipper*. NYU Press.

Hector Vazquez-Leal, Roberto Castaneda-Sheissa, Uriel Filobello-Nino, Arturo Sarmiento-Reyes, and Jesus Sanchez Orea. 2012. High Accurate Simple Approximation of Normal Distribution Integral. *Mathematical Problems in Engineering* 2012 (2012). https://doi.org/10.1155/2012/124029

Qinglong Wang, Xiangshi Ren, Sayan Sarcar, and Xiaoying Sun. 2016. EV-Pen: Leveraging Electrovibration 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

Rhiannon Williams. 2015. Jony Ive interview: The story of the Apple Pencil. http://www.telegraph.co.uk/technology/apple/11988396/Jony-Ive-interview-The-story-of-the-Apple-Pencil.html. (2015). accessed on: 8 January 2018.

Josh Wills, Sameer Agarwal, David Kriegman, and Serge Belongie. 2009. Toward a perceptual space for gloss. *ACM Transactions on Graphics* 28, 4 (2009), 1–15.

Takashi Yoshioka, Sliman Bensmaïa, James C. Craig, and Shen-Fu. Hsiao. 2007. Texture perception through direct and indirect touch: An analysis of perceptual space for tactile textures in two modes of exploration. *Somatosensory & Motor Research* 24, 1-2 (2007), 53–70. https://doi.org/10.1080/08990220701318163

Takashi Yoshioka and Julia Zhou. 2009. Factors Involved in Tactile Texture Perception Through Probes. *Advanced Robotics* 23, 6 (2009), 747–766. https://doi.org/10.1163/156855309X431703

A EXPONENTIAL INTEGRATOR DERIVATIONS

In this section we will introduce derivation of our Exponential integrator as well as proof of some important properties. We start by deriving the general integrator from general equation of motion. Next, we express the integrator for a sinusoidal forcing term. Finally, we show that the results computed by our integrator can be arbitrary rescaled and added together to create novel simulations.

A.1 Exponential Integrator Derivation

The equation of motion of a stiff elastic system under small deformation assumption can be expressed as:

$$\mathbf{M}\ddot{\mathbf{x}}(t) + \mathbf{D}\dot{\mathbf{x}}(t) + \mathbf{K}(\mathbf{x}(t) - \mathbf{x}_0) + \mathbf{\Lambda}(\mathbf{x}(t)) = \mathbf{0},$$

which we can rewrite to solve for accelerations:

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

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

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

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

or more compactly in matrix form:

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

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

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

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

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

or in a more compact form:

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

This equation has a known analytical solution

$$\mathbf{X}(t) = e^{(t-t_0)\mathfrak{U}}\mathbf{X}(t_0) + e^{t\mathfrak{U}}\int_{t_0}^{t} e^{-\tau\mathfrak{U}}\Gamma(\mathbf{X}(\tau))\,\mathrm{d}\tau, \tag{11}$$

which is equivalent to its recursive form

$$\mathbf{X}(t+\Delta t) = e^{\Delta t\mathfrak{U}}\mathbf{X}(t) + \int_{t}^{t+\Delta t} e^{(t+\Delta t-\tau)\mathfrak{U}}\Gamma(\mathbf{X}(\tau))\,\mathrm{d}\tau. \tag{12}$$

A.2 Sinusoidal Forcing Term

Here we present the derivation of sinusoidal forcing term for recursive exponential integrators (Equation 12) starting from rest state that were used in this work. Analogous derivations can be done for the non-recursive version of the integrator, as well as, for integrators starting at a prestressed state. For clarity we will use the following substitutions:

$$A = \mathrm{diag}(\alpha) \qquad \Omega = \mathrm{diag}(\omega)$$

We start with the base recursive version of the equation:

$$\mathbf{X}(t+\Delta t) = e^{\Delta t\mathfrak{U}}\mathbf{X}(t) + \int_{t}^{t+\Delta t} e^{(t+\Delta t-\tau)\mathfrak{U}}\Gamma(\mathbf{X}(\tau))\,\mathrm{d}\tau, \tag{13}$$

where $\Gamma(X(t)) = A\sin(\omega t+\phi)$. We are interested in finding a solution of the integral:

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

First we need to find a solution to the integral:

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

Using per partes substitution:

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

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

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

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

Now we can again apply per partes:

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

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

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

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

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

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

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

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

And finally the solution is:

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

which expands to:

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

Which finally expands to:

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

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

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

Now we can plug it back to Equation 13 with the substitution of $M=-\mathfrak{U}$.

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

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

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

After simplifying we get to the final solution:

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

A.3 Analytical Acceleration

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

$$X(t)=e^{t\mathfrak{U}}\int_0^t e^{-\tau\mathfrak{U}}\Gamma(X(\tau))\,\mathrm{d}\tau \tag{14}$$

We start by applying the derivative product rule:

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

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

Now we use the property of definitive integrals:

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

and we get that:

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

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

In our case of sinusoidal forces the integral has already known solution from Equation 14:

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

A.4 Summation of Exponential Integrators

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

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

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

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

A.5 Multiplying Forces by a Constant

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

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

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

$$X_k(t)=kX(t)$$