DISS. ETH NO. 30040

COMPUTATIONAL DESIGN AND FABRICATION OF WEARABLE HAPTICS

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES
(Dr. sc. ETH Zurich)

presented by

VELKO VECHEV
M.Sc in Information Technology
Chalmers University of Technology

born on 20.04.1984

accepted on the recommendation of

Prof. Dr. Stelian Coros, examiner
Dr. Bernhard Thomaszewski, examiner
Prof. Dr. Alexandra Ion, examiner

2024

Velko Vechev: *Computational Design and Fabrication of Wearable Haptics* , © 2024

DOI: 10.3929/ethz-b-000668226

ABSTRACT

Wearable technologies are increasingly embedded in our lives, from smartwatches on our wrists and augmented reality displays on our heads to more advanced sensing garments that can estimate our pose in real-time. While wearable sensors aim to collect and analyze information from our bodies, wearable haptics aim to physically provide information back to our bodies. The two technologies are potentially complementary, for example, allowing users to physically interact with objects in Virtual Reality (VR). However, while wearable sensing has advanced significantly, wearable haptics remains a challenging problem.

The main challenge in designing wearable haptics are the opposing objectives of performance and wearability. Higher performance devices are likely to be bulkier and heavier, thereby reducing comfort. In addition, haptic devices need to be in close proximity to the body, and unlike wearable sensors, must account for the physical interaction with the soft outer layers of the body. Because of these considerations, wearable haptics has largely been confined to basic vibrotactile actuation on the body (e.g. buzzing), which falls far short of fully harnessing our proprioceptive and tactile abilities.

We address this challenge in two parts. First, we develop actuators that are designed from the ground-up for compliance - the property of flexible and semi-stretchable soft components that makes them ideal for on-body use. Counter intuitively, the property of compliance also makes actuators incredibly inefficient (as they are poor at force transmission) if not properly designed, placed, and connected to the body. Thus our second idea is to design such wearable haptic systems via computational methods.

Specifically, we use physical modeling and simulation in combination with topology optimization to create inverse design methods. This allows for the automatic creation of a variety of garments designed to resist specific motions (kinesthetic garments). We further extend our framework to leverage stateful components (haptic devices that can become stiff on command), and optimize for multiple objectives depending on these states. The resulting wearable haptics are a novel type of soft exoskeleton which are comfortable to wear, and yet can resist a variety of motions on-demand.

Going forward, our results based on force characterization and in various VR tasks, show that automatically designed haptic systems have the potential to perform significantly better than manually designed counterparts.

ZUSAMMENFASSUNG

Tragbare Technologien, von Smartwatches an unseren Handgelenken bis zu Augmented-Reality-Brillen und Sensorkleidungsstücken die unsere Körperhaltung in Echtzeit analysieren können, finden zunehmend Einzug in unseren Alltag. Während tragbare Sensoren darauf ausgelegt sind, Daten von unserem Körper zu erfassen und zu verarbeiten, dienen tragbare Haptik-Systeme dazu, physische Rückmeldungen direkt an unseren Körper zu übermitteln. Diese beiden Technologien können sich ergänzen, um Nutzern die Möglichkeit zu geben, mit Objekten in der virtuellen Realität physisch zu interagieren. Die Entwicklung tragbarer Haptik-Systeme stellt jedoch, trortz des signifikanten Fortschritts bei tragbaren Sensoren, weiterhin ein schwieriges Problem dar.

Die größte Herausforderung bei der Entwicklung tragbarer Haptik-Systeme sind die gegensätzlichen Ziele von Leistung und Tragbarkeit. Leistungsstärkere Geräte neigen dazu, größer und schwerer zu sein, was den Tragekomfort mindert. Zudem erfordern haptische Geräte, im Gegensatz zu tragbaren Sensoren, einen direkten Kontakt zu der weichen äußeren Schicht des Körpers und müssen daher hautnah getragen werden. Aus diesem Grund beschränken sich tragbare Haptik-Systeme meist auf einfache vibrotaktile Stimuli, die unsere propriozeptiven und taktilen Fähigkeiten bei weitem nicht voll ausschöpfen können. Wir gehen diese Herausforderung in zwei Teilen an. Zunächst entwickeln wir Aktuatoren, die speziell auf Flexibilität und Dehnbarkeit ausgelegt sind, um eine optimale Anpassung an den Körper zu gewährleisten. Entgegen der Intuition sind diese flexiblen Aktuatoren jedoch ohne sorgfältige Gestaltung, Platzierung und Verbindung mit dem Körper in der Kraftübertragung ineffizient. Daher entwickeln wir, in einem zweiten Schritt, solche tragbaren haptischen Systeme mithilfe von rechnergestützten Methoden.

Konkret nutzen wir physikalische Modellierung und Simulation in Kombination mit Topologieoptimierung, um inverse Designmethoden zu entwickeln, die es ermöglichen, automatisch Kleidungsstücke zu entwerfen, die spezifischen Bewegungen Widerstand leisten können (kinästhetische Kleidungsstücke). Wir erweitern unser Framework in dem wir zustandsabhängige Komponenten integrieren (haptische Geräte, die auf Befehl steif werden), die es uns erlauben Systeme in Abhängigkeit von diesen Zuständen zu optimieren.

Die resultierenden tragbaren Haptik-Systeme stellen eine innovative Form weicher Exoskelette dar, die nicht nur angenehm zu tragen sind, sondern auch eine Vielzahl von Bewegungen effektiv unterstützen können. Unsere Ergebnisse aus Experimenten zur Kraftcharakterisierung und in unterschiedlichen VR-Anwendungsszenarien belegen, dass automatisch entwickelte haptische Systeme ein signifikant höheres Leistungspotenzial besitzen als manuell entworfene Systeme.

ACKNOWLEDGEMENTS

I would first like to recognize and thank Prof. Otmar Hilliges, who believed in me and game me the opportunity to go on this academic journey. But giving the opportunity is only the first step, and Otmar was continually present to push me and encourage me to aim high week after week. I would also like to thank Dr. Bernhard Thomaszewski, who joined as co-supervisor about half way through the PhD. In addition to being a walking encyclopedia, Bernhard helped to guide me towards more scientific and principled approaches to solving problems, and was always there to get me back on the path when I veered off. I was told before I began the PhD that it is more like a marathon than a sprint, and for this I am grateful to have had such supporting and hands-on advisors along the way.

I want to acknowledge my committee, which included Prof. Dr. Stelian Coros - who was also a collaborator, and stepped up as my main supervisor when it became necessary, and also Prof. Alex Ion from CMU, whose opinion I highly value - having had the brief chance to work with her when she was still at ETH.

It would not have been possible to cross the finish-line without my close collaborators, Dr. Ronan Hinchet and Dr. Juan Zarate, who both started at the LMTS lab at EPFL (working under Prof. Dr. Herbert Shea), but through our multi-year collaborations, somehow both migrated to join ETH. They both dedicated so much of their time and energy to help me turn my raw ideas into something tangible and scientifically grounded. I want to specifically thank Juan for having his usual Maté drink one particular June morning, where his insights during our 5 minute conversation turned around a struggling project into several publications. Don't underestimate the power of the water-cooler talk!

One of the biggest positives of doing a PhD was doing it along the members of the AIT lab. I was placed temporarily in the office that would become my permanent home with mates Manuel and Stefan (sorry and thanks Tobi for your desk!). Both are great listeners, both are great at linear algebra, and at least one is great at hitting high pitch notes during late night submissions. I watched and learned from the originals - Jie the 'groupfather', the always positive but deceivingly productive Wookie, Adrian - possibly the best min-maxer of all time, and Emre who had high standards for both work and cuisine. I want to thank Christoph and Thomas for their help

during tight deadlines, TA sessions during the HCI course, and related forbidden HCI discussions. I enjoyed very much the afternoon sugar runs with Sammy and Thomas, and vibrant lunch discussion with Xu, Marcel and the other ever-present members of the lab. I want to also thank David and Anna for creating a great atmosphere while they were with us, and David in particular for jumping into a project quickly and fearlessly with a get-things-done approach.

There were several people who were instrumental in getting me on this track. First my brother Martin, both as an example that sets a high bar, and also as a mentor throughout the PhD. Second, Dr. Alexandru Dancu, for inspiring me to pursue academics during my Masters and showing me the ropes. Finally, I want to thank my mother Anna, the most important person in my life who doesn't stop pushing me relentlessly to reach and exceed my potential.

Throughout the PhD I was incredibly fortunate to also start a family with 2 kids with my partner in crime - Lilia, who took on the massive responsibility to raise them while I studied and worked. So this PhD is dedicated to them - for Lilia, Rayna, and Luke.

CONTENTS

1

INTRODUCTION

Wearable technology is on the rise - from smartwatches on our wrists that keep track of our bio signs, headphones inside and over our ears, and densely packed head-worn displays for great immersion in AR and VR. Even engineered compression garments are a type of passive wearable technology, specifically designed for high-performance athletic use. The close proximity of wearable systems to the body have the potential to efficiently detect inputs (via sensing), and to provide outputs (via haptics). While wearable sensing has recently made substantial progress (via body-worn IMUs [1] or capacitive touch and strain sensors [2]), wearable haptic output remains either limited to tactile vibrations or trends towards higher-performance, but bulky devices that are not desirable for all-day wear. In this thesis, we will focus on exploring the output side of wearable technology (i.e. haptics) - although developments in methods and fabrication there could also have an impact on sensing technologies as well.

Wearable haptics are body-worn systems that use actuators to provide physical stimuli either to the muscles and joints in our body (kinesthetic feedback), or to the range of mechanoreceptors embedded in our skin (tactile feedback). In contrast to handheld haptic devices such as VR controllers and tabletop feedback devices [3] that the user must approach before using, wearable haptics are *always available*. This crucial property allows them to be activated *on-demand*, and in-turn, this enables a crucial control loop - the ability for a user to perform some input (i.e. touching an object in VR), and then *immediately* receive the appropriate output, allowing them to react accordingly. We leverage this control-loop on a daily basis - easily grasping objects by applying just enough force to prevent slipping [4] and performing object recognition through touch alone [5]. Thus, in the same way that displays have significantly increased in resolution and fidelity to saturate the receptors in our eyes, it makes sense to push wearable haptics to leverage the potential of our bodies haptic senses.

A particular setting where haptic output is highly valued is in Augmented Reality (AR) and Virtual Reality (VR) systems, where the *lack* of haptic output is easily noticed by users. Haptics plays an especially important role in such a setting, as it can increase the realism experienced by the user [6, 7] and afford them the ability to interact with the virtual environment in a similar fashion as we do in real-life [8]. Beyond the virtual setting, wearable haptics have the potential to assist us in everyday tasks by providing physical feedback directly to our bodies. The types of feedback can vary greatly, from the simple vibrotactile information of a smartwatch, to the kinesthetic feedback felt when wearing a posture-correcting or compression garments, and finally to a full-body powered exoskeleton that can physically assist users during locomotion [9, 10].

In contrast to wearable sensors, wearable haptics have one additional requirement that makes their design and deployment even more challenging - *they must exert a force on the user*. Fundamentally, they cannot be mounted far away from the body, and they must be grounded in such a way as to efficiently apply a force. The output of wearable haptics should also be high-fidelity, actuating at different frequencies in order to stimulate the various type of mechanoreceptors in our skin. We can categorize these as central *performance considerations*. Orthogonally to this, mounting such devices onto the body also introduces a new set of *wearability considerations* [11]) that include: *comfort*, *form-factor* (i.e. size, shape), *weight*, and *impairment* (the effect on the range-of-motion of the user). The balance of performance with wearability considerations is the core challenge undertaken in this thesis.

A particular problem of current-generation wearable haptics is the *rigidity* of their components. Rigid components can have high performance, but they are fundamentally incompatible with the body's dynamic movements and soft outer layers. Rather than fitting rigid components onto a soft body, we look to recent advancements in the fields of soft materials and robotics [12, 13], where a *compliance-first* approach results in materials that are inherently wearable and bio-compatible, regardless of where and how they are instrumented on the body. The compliance of a wearable system determines to what degree it is able to bend, conform, and stretch with the body's movements. While compliance is a great starting point for wearable haptics, it introduces a number of technical challenges - how to *fabricate* stretchable and bendable actuators that retain their high performance, and how to *design* and place them on the body without breaking the aforementioned wearability considerations. Furthermore, any bending or stretching in the overall system will create slack and *absorb* forces rather than transmit them to the user, degrading overall performance. In other words, the very same property that makes haptic devices comfortable to wear, may also make them ineffective at doing their job. This is a significant problem, and one of the major obstacles to larger adoption of compliant on-body interfaces.

1.1 APPROACH

We propose a two-fold approach to overcome these challenges: 1) we first develop actuators and connecting structures that ground them to the body such that they are compliant and yet remain performant, 2) we then consider the design problem holistically in terms of how to place and connect haptic components on the body by modeling their material properties, behaviors, and physical interactions with the body.

1.1.1 *Design and fabrication of Compliant Wearable Haptics*

We begin by examining potential actuator technologies as candidates for compliant on-body haptic interfaces. In particular, we focus on providing feedback to the hand, where both kinesthetic and tactile feedback are highly desirable. We developed two types of haptic interfaces - one that provides kinesthetic feedback via electrostatic clutches (ES Clutches), and another that provides tactile feedback using latching electromagnetic elements. Both haptic interfaces rely on a flexible textile substrate, allowing the hand to flex to almost any position. High fidelity output is also targeted, in the sense of providing either high holding forces or a high range of renderable frequencies.

We further propose the concept of *kinesthetic garments* - garments that can resist particular motions by inducing a high-energy state in their deformed pose. We show that it is possible to create such garments automatically. A particular property of these garments is that they use only passive resistive material, which has high applicability in many situations, but cannot be programmed to react to on-demand inputs. In turn, we combine a textile-packaged version of our electrostatic clutches with kinesthetic garments to create *active kinesthetic garments* - soft exoskeletons that can resist motions on-demand.

1.1.2 *Computational Methods for Wearable Haptics*

In the wearable setting, wearable haptics based on compliant materials are not enough, we need to also consider how they impact freedom-of-movement and comfort of their wearer, both in their active and non-active states. There is also an interactive component - interfaces stretch and deform depending on whether they are active or not. Thus, a significant amount of time is devoted to balancing these objectives through real-world iteration and testing, significantly increasing the overall effort required of designers. Instead, we consider recent approaches in graphics literature that develop physics-based models of textiles mounted on human bodies [14, 15]. We similarly model the interface layer (the garment) as a thin-sheet sliding on the body and we build on these methods by further modeling the actuators and sensors attached to them. This enables a more fundamental and scalable approach to designing wearable haptics - through computational design methods that leverage algorithms to support the designer.

Crucially, our pipeline provides an automatic way to perform *inverse design* in order to satisfy the high-level objectives of the designer. This not only reduces design effort, but excels at producing highly efficient designs, especially as we scale up our haptic interfaces to the whole body. Our pipeline is built by combining two disparate disciplines - topology optimization and parametric human-body models - we call this on-body topology optimization. Finally, we extend our formulation to include active strain-based components, factoring in states directly in the formulation, thus allowing for the optimization of the opposing objectives of wearability *and* performance.

1.2 CONTRIBUTIONS

1.2.1 *Compliant Hand-Worn Haptics for VR*

Our first contribution is based on novel hardware. We introduce two sets of haptic interfaces that both use a textile glove substrate. TacTiles is designed to allow texture and object recognition in VR using a flexible haptic array to provide localized tactile feedback. Each element in the array uses an electromagnetic actuator that weight just 1.8g, consumes 130 mW, and measures 1cm^3. The hardware contributions are a custom 8-layer PCB, asymmetric latching, and materials to dampen secondary vibrations during actuation. Based on this hardware, we propose two distinct modes of actuation: contact and pulse. We found that users were able to localize feedback, discriminate between modes with high accuracy, and differentiate objects from 2d and 3d haptic surfaces, and localize virtual objects without looking at them.

Our second wearable haptic device (DextrES) is a flexible haptic glove that integrates both kinesthetic and cutaneous feedback in a thin and light form factor (weighing less than 8g). DextrES not only vibrates when making contact with objects, it also resists finger flexion to prevent penetration into virtual objects. It does so in a thin, lightweight, and compliant package. This is accomplished using an electrostatic clutch generating up to 20 N of holding force on each finger by modulating the electrostatic attraction between flexible elastic metal strips to generate an electrically-controlled friction force. For end-users, this allows them to feel the feedback required for fine-grained grasps, enabling dexterous articulation in VR. A controlled experiment indicates that this interface can improve the grasping precision for different types of virtual objects.

1.2.2 *Automatic Design and Modeling of Kinesthetic Garments*

Our second contribution is a method that can automatically create passive (unpowered) motion-resisting garments (Kinesthetic Garments) through tailored distributions of a stiff material that is used for reinforcement. Finding designs that distribute a given amount of stiff material to maximally stiffen the response to specified motions is a challenging problem. In this work, we propose an optimization-driven approach for automated design of layout patterns for kinesthetic garments. Our main contribution is to cast this design task as an on-body topology optimization problem. Our method allows designers to explore a continuous range of designs corresponding to various amounts of stiff material coverage. Our model captures both tight contact and lift-off separation between cloth and body. We demonstrate our method on a variety of design-layout problems for different body sites and motions. Automatically generated designs lead to a two- to threefold improvement in performance in terms of energy density. A set of automatically generated designs that were fabricated and compared with baselines in a blind user-study were consistently rated as providing more resistance than baselines in a comparative user study.

1.2.3 *Dual-Objective Optimization of Kinesthetic Garments for Performance and Wearability*

Our third and final contribution links our two previous contributions by extending our on-body topology optimization framework to support active components. It is able to generate efficient *connecting structures* between active components—one of the central challenges in this context. We focus on electrostatic (ES) clutches that are compliant in their passive state while strongly resisting elongation when activated. Our method automatically computes optimized connecting structures that efficiently resist a range of pre-defined body motions on demand. We propose a novel dual-objective optimization approach to simultaneously maximize the resistance to motion when clutches are active, while minimizing resistance when inactive. We demonstrate our method on a set of problems involving different body sites and a range of motions. We further fabricate and evaluate a subset of our automatically created designs against manually created baselines using mechanical testing and in a VR pointing study. We found that the automatically designed garments significantly outperform manually designed ones across a wide variety of tasks in simulation, force characterization studies, and user evaluations.

1.3 THESIS OUTLINE

The thesis structure is as follows: we first establish state-of-the-art in wearable haptics and devices, as well as theoretical concepts behind structural optimization on the body. We then present 4 chapters related to our contributions in hardware and methods.

CHAPTER 1 presents state-of-the-art in wearable haptic devices. Theory is also presented on capacitive-based devices related to electrostatic clutches.

CHAPTER 2 presents the theory required for simulation of haptic interfaces on the body. Additional theory is also provided on structural optimization, specifically bi-directional evolutionary structural optimization methods.

CHAPTER 3 proposes a novel tactile hand-worn device.

CHAPTER 4 proposes a novel kinesthetic feedback glove-based device.

CHAPTER 5 presents a novel method for automatically designing motion-resisting garments based on topology optimization.

CHAPTER 6 extends our on-body topology optimization method with active components (i.e. clutches) and adds a dual-objective to optimize for both performance and wearability.

CHAPTER 7 summarizes our contributions and makes suggestions about future work.

1.4 PUBLICATIONS

This thesis is based on the following peer-reviewed publications, in the order of the chapters they are presented:

Vechev, V. *et al. Tactiles: Dual-mode low-power electromagnetic actuators for rendering continuous contact and spatial haptic patterns in vr* in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2019), 312

Hinchet, R. *et al. Dextres: Wearable haptic feedback for grasping in vr via a thin form-factor electrostatic brake* in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (2018), 901. DextrES was also presented as a demo at UIST 2018.

Vechev, V. *et al. Computational Design of Kinesthetic Garments* in *Computer Graphics Forum* **41** (2022), 535

Vechev, V. *et al. Computational Design of Active Kinesthetic Garments* in *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (2022), 1

2

BACKGROUND

2.1 HAPTICS THEORY AND DEFINITIONS

2.1.1 *Haptic Perception*

Human haptic perception is a primary sense that is distributed throughout the body in our skin, muscles, and joints. It provides us the capability to sense our natural or simulated environment through two distinct modalities: *tactile* and *kinesthetic* [20]. Our tactile sense provides feedback on pressure, shear, and vibration through mechanoreceptors embedded in our skin. Kinesthetic feedback is closely related to proprioception - our ability to localize our limbs and posture in space and to sense large scale forces exerted on our hands and bodies via muscle and joint activations. The haptic sense distinguishes itself from other senses in that it is mechanically *active* - we act on the environment by manipulating it with touch and forces, and use the information of this interaction to continually modulate further actions.

2.1.2 *Haptic Interfaces*

A *haptic interface* enables a similar type of bi-directional information flow - allowing for the simultaneous exchange of information between the user and the machine. Haptic-interfaces are unlike haptic displays (and related graphical displays), which only offer unidirectional information transfer. Rather, they provide (haptic) information in-accordance with the input, establishing the so-called haptic-loop. In Fig. 2.1, a normal mouse (left) can be enhanced to provide haptic feedback in accordance with input (right), for example to resist movement when a user is nearing a disabled action on-screen. Hayward *et al.* [21] further characterizes the exchange of information as an exchange of *energy*, a useful notion that we can leverage when using physics-based optimization to design haptic-interfaces.

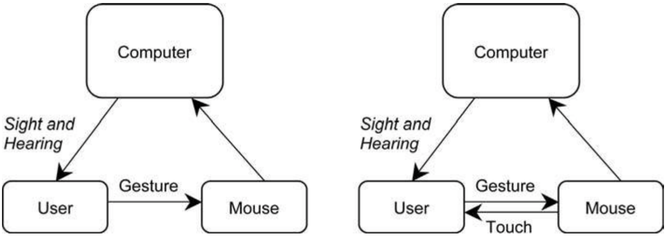

FIGURE 2.1: The Haptic-Loop allows for bi-directional exchange of information between machine and user (Hayward *et al.* [21]).

2.1.3 *Haptic components of grasping*

The perceptual mechanisms behind the experience of holding an object or exploring the shape and texture of its surface is composed of *kinesthetic* and *cutaneous* components [21]. During object manipulation, the typical cycle starts when type 1 fast receptors in the fingertips are excited for about 1 second, indicating the contact boundary of an object [22]. After initial contact, kinesthetic forces are transmitted through the joints and muscles, informing us of relative limb and finger positions through the sense of proprioception. Kinesthetic and cutaneous channels work in tandem to provide an accurate sensation of touch [23] that also acts as a feedback loop to accurately control the grasping force exerted on an object [4].

2.1.4 *Tactile perception of object properties*

When exploring object properties such as shape, size, and texture, we perform corresponding movement patterns known as exploratory procedures (EPs) [24]. For example, to recover texture, we slide our fingers laterally over an object, and to recover shape, we mold our hand to the surface of an object. Gibson refers to these information seeking movements as Active Touch, which differs from passive touch, or being touched [5], in that our sensory perception of material properties is significantly enhanced [25]. Studies in object identification through touch found that using five fingers instead one can decrease recognition time, and that we are overall more accurate and faster at identifying 3D objects than 2D raised surfaces [26]. When recovering the shape of objects, increasing the distance between touch points allows for finer discrimination of curvature [27].

2.1.5 *Wearable body-grounded kinesthetic and tactile haptic devices*

A wide range of haptic devices exist - from tabletop devices that allow touching [3, 28] and grasping [29], handheld devices that can be picked up to provide feedback to the hand [8], to wearable devices that provide either kinesthetic and tactile feedback. In this thesis we focus exclusively on wearable, body-grounded haptic devices (Fig. 2.2 a.Kinesthetic and b.Tactile), which exert forces onto the user through the use of a counteracting force anchoring to some location on the user's body. While they enable full mobility, the effective anchoring and distribution of such counteracting forces is precisely what makes them so difficult to design.

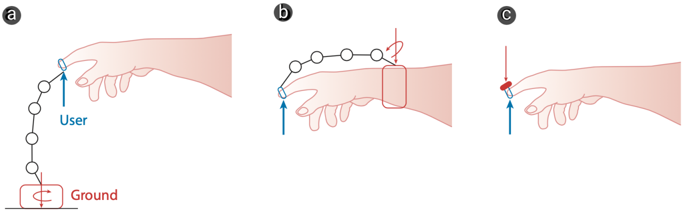

FIGURE 2.2: Haptic device grounding showing a) World-grounding, b) Body-grounding (kinesthetic feedback), and c) Local-grounding (tactile feedback). Blue and red arrows indicate forces and counter-forces respectively. Reproduced from Culbertson, Schorr & Okamura [20].

2.1.6 *Wearability performance and design factors*

In regards to the wearability of haptic devices, Pacchierotti *et al.* [11] define a number of performance factors and design considerations. We can characterize performance in terms of degrees-of-freedom (the dimensionality of the feedback), workspace (shape and size), peak-force, (perceived) inertia and friction - which can be affected by weight, precision and resolution (repeatability of stimuli and saturation of mechanoreceptor ability), and bandwidth (frequency response). Such performance factors are often in contrast to wearability factors - comfort, weight, impairment (impediment to motion), and form-factor, where for example, increasing peak-force requires a larger device that could interfere with motion and increase weight. Designing wearable haptic interfaces requires at the minimum to consider, and ideally, to optimize the balance of performance and wearability considerations.

2.1.7 *Soft-wearable and textile-integrated haptic interfaces*

A fundamentally different approach to designing haptic interfaces involves using textiles and soft materials (i.e. silicone-based) as substrates to carry active components, including actuators and sensors [12, 13]. As an already familiar body-worn interface, garments provide an ideal substrate that can be both instrumented and modified at the fabric level with custom properties. For example, higher tensile stiffnesses can be used as anchoring and connecting structures to transfer forces onto the body efficiently [30]. Garment-based haptic interfaces are highly dependant on a precise fit and placement of components, necessitating the need for customization or personalization. Attempting to reduce comfort factors such as pressure can also lead to increased slack and reduced stiffness, resulting in lower overall performance. Despite these challenges, garment-based haptic interfaces can be comfortably used in everyday scenarios, leading to many novel applications including locomotion assistance, grasping assistance, thermoregulation, and general load holding and carrying [12, 31]. In an AR/VR setting, Yin *et al.* [12] provide a large categorization of kinesthetic and tactile haptic devices for the body and the hand based on soft technologies.

2.2 ON-BODY PHYSICS-BASED OPTIMIZATION

In this thesis, we perform both *forward* simulation and *inverse* design of structures in contact with the human body. Here we outline theory related to classical finite element modeling (FEM) that we use in our simulation of garments, topology optimization methods and theory typically used for engineering stiff structures with reduced material usage, and the parametric models of the human body that are typically used as input.

2.2.1 *Garment Modeling and Simulation*

In order to measure deformation in a given domain, we use the deformation gradient $\mathbf{F} = \frac{\partial \hat{x}}{\partial \bar{x}}$, where $\hat{x}$ describes the mapping from the continuous rest domain $\Omega_{\bar{x}}$ to the deformed domain Ω_x. Typically x and $\bar{x}$ are discretized as vertices $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_m) \in \mathbb{R}^{3n}$ and $\bar{\mathbf{x}} = (\bar{\mathbf{x}}_1, \ldots, \bar{\mathbf{x}}_m) \in \mathbb{R}^{3n}$. In $\mathbb{R}^3$ space, $\mathbf{F}$ is a 3x3 matrix, allowing us to map the rest vertices to the deformed vertices by simple matrix multiplication: $x = \mathbf{F}\bar{x}$.

Crucially, while $\mathbf{F}$ is translation invariant, it still contains rotation, scaling, and reflection. Therefore, $\mathbf{F}$ is typically processed into a set of quantities that help us build various constitutive models. These are described in Tbl. 2.1.

Symbol	Equation	Definition
$\mathbf{C}$	$\mathbf{F}^T\mathbf{F}$	Right Cauchy Green Tensor
$\mathbf{E}$	$\frac{1}{2}(\mathbf{C} - I)$	Green Strain Tensor
I_C	$tr(\mathbf{C})$	First Invariant of Right Cauchy Green Tensor
J	$det(\mathbf{F})$	Relative Volume Change

Table 2.1: Quantities derived from F.

Here, both $\mathbf{C}$ and $\mathbf{E}$ are non-linear measures of deformation, and both are rotation invariant. While I_C captures deformations in *all directions*, including shearing, J captures only volume changes.

Next, we can use these quantities to build the aforementioned constitutive material models. In particular, we are interested in hyper-elastic materials that exhibit large deformations. First, we specify the elastic behavior of the model through the energy density function Ψ, in this case based on the classical compressible neo-Hookean model [32]:

$$\Psi(\mathbf{F}) = \frac{\mu}{2}(I_C - 3) - \mu \log J + \frac{\lambda}{2}(\log J)^2 , \tag{2.1}$$

where λ and μ correspond to Lame's first parameter and shear modulus, allowing us to penalize shearing and volume changes separately. This energy density is then integrated over the domain, giving rise to an energy potential W:

$$W = \int_{\Omega_{\bar{x}}} \Psi(\mathbf{F})\, d\Omega_{\bar{x}} , \tag{2.2}$$

Under static equilibrium state $x*$, the internal forces of $f_i^{internal} = \frac{\partial W}{\partial x_i}$, plainly the derivative of the energy potential with respect to its nodal positions, are balanced with external forces applied to the system:

$$\mathbf{f}(x^*(p)) = f_i^{internal} + f^{external} = \mathbf{0} . \tag{2.3}$$

where such a state is typically achieved through an optimization scheme such as the Newton-Raphson procedure described in Nocedal & Wright

[33]. The variable p in this case represents high-level design parameters such as the patterns of cloth that make up a garment.

Garments, in particular skin-tight garments, consist of thin layers of fabric. We can therefore model them as thin shell - triangular surfaces whose rest state is in the domain in $\bar{x} \subset \mathbb{R}^2$, but it's deformed state is defined by vertices $x \subset \mathbb{R}^3$. The procedure to derive a deformation gradient F thus requires constructing one additional degree-of-freedom that is orthnonormal to the existing plane defined by the edges of $\bar{x}$ in $\mathbb{R}^2$. The procedure can be found in [34]. It is important to note that since the rest state $\bar{x}$ does not change, the rest configuration as projected into the orthno-normal basis only needs to be computed once.

The resulting deformation gradient $\mathbf{F}_C$ is a 3×2-matrix and $\mathbf{C}_C$ is a 2×2-matrix describing the deformation of the element with respect to rest state coordinates. The energy density Ψ must be modified to account for the 2x2 matrix by replacing $(I_C - 3)$ with $(I_C - 2)$. While thin shells typically also include a bending energy - skin tight garments in particular exhibit much higher energy in-plane rather than as a result of bending. As such, the bending term is often omitted.

2.2.2 *Inverse Design*

Forward simulation allows designers of wearable haptic interfaces to *evaluate* existing designs, however, it does not allow them to drive the creation of such designs via high-level objectives. In particular, we are interested in satisfying objectives related to performance, such as maximizing forces felt by the user, while minimizing wearability factors such as impediment and comfort by reducing the overall size and form-factor of the interface. Inverse design aims to allow optimization of such high-level objectives, while abstracting away the physical simulation step in between the setting of design objectives and their resulting equilibrium configuration. A number of approaches to inverse design have been explored:

SENSITIVITY ANALYSIS leverages the fact that for every set of design parameters p, there exists a local unique equilibrium configuration (See Eq. 5.10), therefore creating an implicit map between p and the equilibrium state $x*$. The problem can then be framed as an unconstrained minimization that depends only on the high-level parameters p. Sensitivity analysis has been employed to assist in designing garments automatically, for example, by allowing designers to edit patterns directly [35], or to automatically create patterns that minimize friction and pressure [14].

TOPOLOGY OPTIMIZATION aims to find an optimal material distribution that minimizes a particular cost function (e.g. compliance) while satisfying a set of given constraints, for example, a target material budget [36]. The design variable in this case is a per-element density value $d \in [0, 1]$, that represents whether a particular element is part of the domain. Using this approach it is possible to create connecting structures that efficiently carry loads and route forces to boundary areas. In order to solve this problem, a number of algorithms have been proposed. A popular approach is SIMP (Solid Isotropic Materials with Penalization), which is a continuous method, that is, d is continuous between 0 and 1. As the optimization progresses, it penalizes intermediate material densities with the power-law pushing the values of d to either 0 or 1. However, there is no guarantee that the design will reach a state with fully binary distribution. The presence of intermediate material densities can also be problematic as the stress constraints of the material at those particular densities is not known *a priori* [37]. Related to this is the *singularity* problem, where the elements with assigned low-density exhibit higher than expected strain values, making optimization and convergence difficult [37, 38].

BI-DIRECTIONAL EVOLUTIONARY TOPOLOGY OPTIMIZATION (BESO) is an alternative formulation of the optimal material distribution problem in which the design variable d are constrained to binary values $\{0, 1\}$ [39]. As only fully solid or void materials are permitted, this allows BESO to avoid the singularity issue. The algorithm consists of an iterative procedure where material is added and removed simultaneously based on sensitivity of elements with respect to the objective function, determined through finite element simulation. A target material budget constraint is typically imposed at each successive iteration. In terms of performance, previous work has shown that BESO can achieve similar designs and results as SIMP [40, 41], however, intermediate BESO designs with monotonically decreasing material budgets allow for a degree of design exploration not possible with SIMP.

Two general approaches exist: hard-kill, where elements are fully removed from the structure, and soft-kill, where elements are replaced by a weaker material, but still present in the structure. The hard-kill approach has previously been found to lead to non-optimal solutions, in particular when using a coarse mesh structure [42], whereas the soft-kill approach is more stable and lends itself to a multi-material setting [43]. The multi-material setting is more appropriate for our problem as garments and haptic interfaces include substrates and stiff structures that require multiple materials to model, thus we focus on this version of the BESO algorithm.

multi-material beso begins with the following problem statement:

$$\begin{aligned} \mathbf{d}^* = \arg\max_{\mathbf{d}} \quad & f(x^*, \mathbf{d}) \\ \text{s.t.} \quad & \sum_e V^e d^e = V^* \ , \end{aligned} \tag{2.4}$$

where f is the objective function, d is the design vector $d = [d_1, \ldots, d_n]$, $x*$ is the configuration of the shape as a result of a simulation, V_e is the elemental volume and $V*$ is the maximum allowable volume. A typical minimization objective function could be the mean compliance: $C = \frac{1}{2}\mathbf{f}^T\mathbf{u}$, where $\mathbf{f}$ and $\mathbf{u}$ are the load and displacement vectors [39]. In an elastic setting such as garment simulation, an objective function can instead minimize the elastic energy potential W (See E.2.2). Under a multi-material interpolation scheme, the Young's modulus of a particular element can be computed as $E(x_e) = E_1 d_e^p$, where E_1 is a fully solid material, and p is the penalty exponent.

The algorithm takes as input the sensitivity of each element with respect to the objective function. In the mean compliance setting, the sensitivity is equal to the integrated elemental strain energy:

$$\alpha_e = \frac{\partial C}{\partial x_e} = \begin{cases} \frac{1}{2}\left[1 - \frac{E_2}{E_1}\right] \mathbf{u}_e^T \mathbf{K}_e^1 \mathbf{u}_e & \text{material 1} \\ \frac{1}{2} \frac{d_{\min}^{p-1}(E_1 - E_2)}{d_{\min}^p E_1 + \left(1 - d_{\min}^p\right) E_2} \mathbf{u}_e^T \mathbf{K}_e^2 \mathbf{u}_e & \text{material 2} \ . \end{cases} \tag{2.5}$$

where the penalty exponent p here is used to control the sensitivity of the less dense material (material 2), and by setting p to infinity, the sensitivity tends to 0, turning it into the hard-kill version of BESO. Additionally, a number of other issues such as checkerboarding and mesh-dependency are resolved through spatial and temporal filtering of the sensitivity values [39].

The BESO structural refinement procedure consists of the following steps. The design vector d is typically initialized to a fully dense structure resulting in a starting volume V_0. We first define the target volume $V*$ and evolutionary ratio ER. Then, finite element analysis is performed

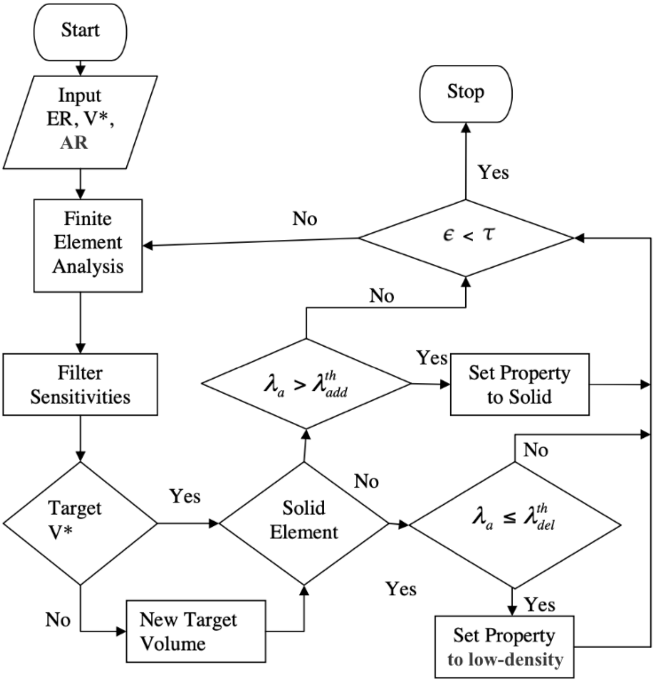

FIGURE 2.3: The BESO structural refinement procedure (Based on Aremu *et al.* [40]).

and the sensitivity numbers are extracted based on Eq. 2.6. Sensitivity numbers α_e are filtered to produce λ_e which are then sorted and ranked. The threshold λ^{th} is then calculated such that the area of all elements with higher sensitivities satisfies the current iteration's target area V_i. If the sensitivity of a low-density element exceeds this threshold, add it to the domain, and vice-versa, set solid elements to low-density if they are below. The add and remove thresholds, λ_{add}^{th} and λ_{del}^{th} may differ as the volume of material to be added may be upper bounded by an addition ratio AR in order stabilize the procedure. The target volume is updated for the next iteration $V_{i+1} = V_i(1 \pm ER)$, and if $V_i \equiv V^*$, then set ER to 0. The procedure is visualized in Fig. 2.3.

The procedure is considered converged under the following criteria:

$$\epsilon = \frac{|\sum_{i=1}^{N}(f_{[k-1+1]}) - f_{[k-1]}|}{\sum_{i=1}^{s}(f_{[k-1+1]})} \leq \tau \ , \tag{2.6}$$

where $f_{[n]}$ is the value of the objective function at the n-th iteration, k is the current iteration, and N is the window size over which the *relative change* ϵ of the objective value must remain below the allowable threshold τ.

2.2.3 *Modeling the Surface of the Human Body*

Many approaches exist to modelling the surface of the human body, however here we examine approaches which aim to construct anatomically correct representations, namely i) data-driven methods and ii) physical simulations of it's surface properties.

SMPL [44] is a parametric model of the surface of the body based on a mesh template T with N = 6890 vertices. The vertices are offset from the statistically mean shape based on the training scans T_μ. These offsets are determined by 10 shape $\boldsymbol{\beta}$ and 72 pose $\boldsymbol{\theta}$ parameters, representing a kinematic chain of 24 joints with 3 degrees of freedom each.

W is a linear blend-skinning function that takes skin vertices in the rest pose T, joint locations J, a pose θ, and the blend weights B_w, and returns the posed skin vertices:

$$\begin{aligned} T(\beta, \theta) &= T_\mu + B_s(\beta) + B_p(\theta) \\ M(\beta, \theta) &= W(T(\beta, \theta), J(\beta), \theta, B_w) \end{aligned}$$

The rest pose $T(\beta, \theta)$ is based on the skeleton joints $J(\beta)$ after pose-dependent deformations $B_p(\theta)$ and shape dependent deformations $B_s(\theta)$. Shape-dependent deformations $B_s(\beta)$ model subject identity.

SMPL-H [45] extends SMPL with additional pose and shape parameters for the joints of the hand, and more recently SMPL-X [46] has added parameters for the face. SMPL-X also proposes the use of a human body prior VPoser, which is encoded from training data using a Variational autoencoder (VAE) [47].

STAR [48] further extends SMPL by using an auto-encoder architecture that encourages local joint to surface dependency, and training on a larger subset of human scans that SMPL. STAR includes 300 β components, to give it higher expressivity. Due to this larger sample size and expressivity, STAR has higher accuracy specifically when it comes to representing muscle bulging and other local surface deformations.

SMPL has also been extended to model the soft tissue dynamics of the human body [49]. Casas et al. has similarly applied the use of a VAE to add nonlinear soft-tissue dynamic effects to skinned meshes [50], allowing the model to run in real-time.

A concurrent but equally important research path is through modeling *physical properties* of the human body, typically using FEM simulations. Here we again focus only on models of the surface of the body.

Kim et al. [51] create a volumetric version of SMPL which is overlayed with an FEM driven physics-based skin layer. They use a relative distance field based on a biharmonic function to smoothly interpolate between the root bones and the skin surface. This allows them to successfully register subject-specific parameters for each 4D scan.

One weakness of such a method is it's non-contact measurement approach, limiting the type of mechanical properties of the surface of the body (i.e. the skin) that we can infer. Pai et al. [52] aim to build a more accurate model of the skin by employing a custom-designed measurement probe. The probe is equipped with 3 endothermic cameras, a force sensor, and it's position is continually measured with an external optical tracking system. They place optical markers on human subjects and measure the deformation of the skin after applying a force in the direction of the skin normal, and laterally. Using this data, they model a sliding thick skin surface layer over an internal layer represented by an implicit function. FEM simulations of their model show that it can accurately model contact with large objects such as a belt tightening around the waist.

3

TACTILE ON-HAND HAPTIC-INTERFACES

In this chapter, we focus on designing a tactile-feedback interface for the hand. The Palmer surface of the hand is a region with very high density of mechanoreceptors. However, stimulating the different receptive frequencies of these receptors is challenging, particularly with a single actuator. Another challenge is to densely mount many such actuators in close proximity to cover the surface of the hand. Our technical contributions are as follows. We developed and fabricated a lightweight and compact actuator that is able to span a wide range of frequencies - from 1Hz to 200Hz. We then created a dense array of 15 such actuators and mounted them to a textile glove. Device characterization showed that an individual actuator can sustain its peak frequency of 200Hz for 50ms, or we it can actuate at 2Hz for an indefinite period of time. User evaluations revealed high localisation ability for the majority of the array, and new use-cases such as object recognition through haptic feedback alone. We summarize these findings and the key takeaways at the end of this chapter.

3.1 INTRODUCTION

When interacting with objects in real life, we make use of our sense of touch to recover properties such as objects' shape, size, and texture [5, 24]. Allowing for similar haptic-based scene exploration in Virtual Reality (VR) is desirable as it can increase realism [53] and enhance our interactive capabilities [54, 55]. In particular, we seek to leverage the Palmar surface of the hand which contains the highest density of touch receptors [22] and is the most natural place to provide tactile feedback.

Designing a hand-based cutaneous feedback system for use in VR poses significant challenges as it must be i) *wearable* – to allow the hand to move freely and to ensure skin conformity, (ii) *high-fidelity* in terms of density and the ability to render different modes of actuation (e.g. contact events and continuous touch) and yet (iii) *low-power* – to prevent overheating during frequent actuation typical in VR. Even at the single actuator level, these design challenges are at odds with each other, and become even more complex to integrate in a system designed for the whole surface of the hand.

The most common actuator design is based on vibrotactile feedback such as those built into VR controllers (e.g. Oculus Touch) and in a glove form factor such as the Cybertouch [56] and the Hi5 VR Glove. However, as vibrotactile feedback is only able to stimulate high frequency mechanoreceptors [22], it can not convincingly render discrete touch events or the continuous sensation of touch required for high-fidelity and localized feedback.

Increasing fidelity is possible by using mechanically-driven tilting platforms and high resolution pin arrays [8, 11, 57]. More complex actuators, however, make it difficult to integrate and place them flexibly around the hand. Another strategy is to use simple actuators with external driving mechanisms, for example using servo motors [58], or pneumatic inflation [59]. This allows the integration of multiple actuators on the hand, however, at the cost of additional bulky hardware which limits freedom of movement in VR. More recently, actuators have been proposed based on electromagnetic latching [60] and direct electrotactile stimulation [61]. While these devices present promising developments, they must carefully balance power-draw, force-generation capabilities, and skin conformity in order to avoid overheating during demanding VR usage.

To address the diverging requirements of form factor, power envelope, and rendering fidelity we propose a novel type of multi-mode actuator, called TacTiles, which can be placed anywhere on the hand, and comes in a form factor and power envelope which is closer to current vibrotactile actuators. Each TacTile is self-contained and includes the actuation mechanism capable of rendering discrete contact events, continuous contact, and short high-frequency bursts. This allows an array of TacTiles to seamlessly transition between object surface exploration and object enclosure in a single interaction.

Similar to our earlier work on MagTics [60], TacTiles utilizes a bi-stable electromagnetic latching mechanism, but takes up half the physical space and reduces power consumption by a factor of four, while maintaining a high holding force of 200 mN. This is enabled by a novel hardware design consisting of 4-layer latching plates and custom-designed 8-layer printed circuit boards (PCBs). The small form factor (1 cm^3) and low weight (1.8 g) of the actuators allows for denser and more flexible arrangements on the hand. The low power consumption (130 mW on average, pulsed activation at 5 V and 3.6 A) allows for continuous activation without overheating, which in turn enables the actuator to render a high dynamic range of frequencies. To avoid undesired haptic sensations from actuator recoil when disengaging (e. g. releasing contact with a virtual object), we furthermore propose adding dampening material into the magnet chamber which drastically reduces the perceived recoil. We experimentally show that the dampening is effective across a wide frequency band.

These results are used to build a 15 actuator array of TacTiles, which is mounted on a glove and tested under various scenarios including discriminating haptic pattern frequencies and haptic surface and object exploration in VR. We experimentally confirm that users can discriminate the location and mode of individual TacTiles with an average accuracy of 78.7%. Furthermore, we show that participants can discriminate the frequency of different haptic patterns with a just-noticeable-difference of 6% on average. TacTiles enable a wide range of applications, which we show in a first experience test in VR. When users were shown multiple different haptic patterns and objects for exploration, they were able to discriminate objects from haptic patterns, different patterns, and could detect objects on virtual surfaces even without looking at the objects.

3.2 RELATED WORK

We focus on wearable tactile feedback devices most relevant to TacTiles. For a full review of haptic devices designed for AR and VR, we refer readers to the work of Bermejo & Hui [62].

3.2.1 *Wearable vibrotactile displays.*

Wearable haptic devices are a natural fit for VR as they are always-available and at the user's fingertips. The most commonly used haptic technology in such devices is based on vibrotactile actuators which can render contact on different parts of the hand [56, 63] and are common in commercial gloves designed for VR [64, 65]. Vibrotactile feedback harnesses the piezoelectric effect to efficiently convert electrical current into vibrating mechanical energy. However, such feedback is primarily capable of stimulating Pacinian mechanoreceptors responsible for high-frequency sensations [22]. While vibrotactile actuators lack the rendering fidelity to create a sensation of localized touch, they are well suited to rendering textures which have natural vibratory patterns [**culbertson2014one**]. Recent advances in asymmetric vibrotactile actuation allow rendering directional cues on the fingertips [66–

68]. An important design goal of TacTiles is to retain the small form-factor and power-envelope of vibrotactile actuators, while expanding on the types of modes it can render.

3.2.2 *High-fidelity haptic feedback.*

Several types of haptic devices for VR have been proposed to increase fidelity. The high receptor density of the fingertips makes them an ideal place for high resolution feedback [69]. Benko et al. proposed two devices, *NormalTouch* and *TextureTouch* based on a handheld mechanically-tilting platform and a 4×4 articulated pin array respectively [8]. They found that both platforms increased tracing and pointing task accuracy, however, the high resolution pin-array did not perform better than the tilting platform. Whitmire et al. [70] proposed an articulated and interchangeable haptic wheel that could accurately render 1D shear forces. Surprisingly, they found that wheel spin direction had little impact on realism. Multi-DOF devices have also been proposed in a more wearable form-factor where the articulated platform is driven by small motors mounted on top of the finger [11, 57, 71, 72] and are capable of rendering contact angles and varying degrees of pressure to the fingertips. These devices are typically characterized by mechanical actuation and because of their form-factor, are not designed to cover arbitrary parts of the hand.

Integrating tactile feedback on to the whole hand can also be accomplished by offloading the driving mechanisms. Son and Park [58] proposed a tactile device to provide feedback at 10 locations on the user's palm, however, it requires attaching the relatively large motor-pack directly onto the user's hand. Gloves based on pneumatics use valves to activate tactile pixels on the hand [59]. However, these require complex routing and pumps to activate which may limit wearability. Since each TacTile is designed to be a self-contained unit and does not require surface area beyond its immediate footprint, it can be placed anywhere on the hand without restricting range-of-motion. When configured in such an array, TacTiles can provide higher fidelity localized feedback to render interactions such as object enclosure and surface exploration.

3.2.3 *Low-power actuation mechanisms.*

Power draw is crucial in the context of VR, where activations are frequent and sessions can last a long time. Therefore, we consider two actuation technologies with low-power draw, but with the potential to render multiple modes of actuation: electrotactile and electromagnetic.

Electrotactile actuation directly stimulates mechanoreceptors in the hand to render touch in VR [55, 71]. To be most effective, such devices require high conformity to the skin and typically use clipping systems that limit placement to the fingertips. Recent work has addressed this problem by embedding electrotactile tactors in a flexible substrate [61] that can adhere to the skin. Rendering continuous contact with electrotactile actuation is still problematic as it requires a constant power-draw. Yem et al. also found mechanical skin displacement to be a better approximator of properties such as hardness and macro roughness [71].

Electromagnetic actuation was used by Yang et al. in a 3×3 tactile array driven by embedded solenoids and returned to resting states by springs [73]. The measured force output was rated at 5.6 mN which is an order of magnitude less than TacTiles. It also requires continuous power output to provide feedback. Improving on this design, Pece et al. proposed MagTics, a flexible tactile device using a bi-stable latching mechanism that provides continuous force output without the need for further power [60]. Despite many appealing properties, the device cannot sustain continuous usage in VR due to its power draw of 140W per activation. In addition, when disengaging from the skin, the bi-stable actuator causes an unwanted haptic vibration as the pin hits the wall of the chamber. More recently,Duvernoy *et al.* [74] showed a stationary conformal haptic interface interface using a braking mechanism based on magnetic force repulsion [74].

We base the core design of a single TacTile on the power-efficient MagTics actuator, but make key improvements by significantly reducing power consumption, adding damping to reduce unwanted feedback, and reducing the size, thus enabling a multi-mode actuator that is suitable for continuous and demanding use in VR.

3.3 SYSTEM OVERVIEW

The aim of our work is to provide a tactile feedback system for VR that is capable of rendering realistic tactile information when touching virtual shapes and exploring object surfaces. Given the design considerations of *wearability*, *high-fidelity*, and *low-power*, we select the bi-stable electromagnetic latching actuator used in MagTics [60] and make several key improvements that significantly decrease it's form-factor, harness the full mode capabilities of the actuator, and lower its power usage by a factor of four. We further

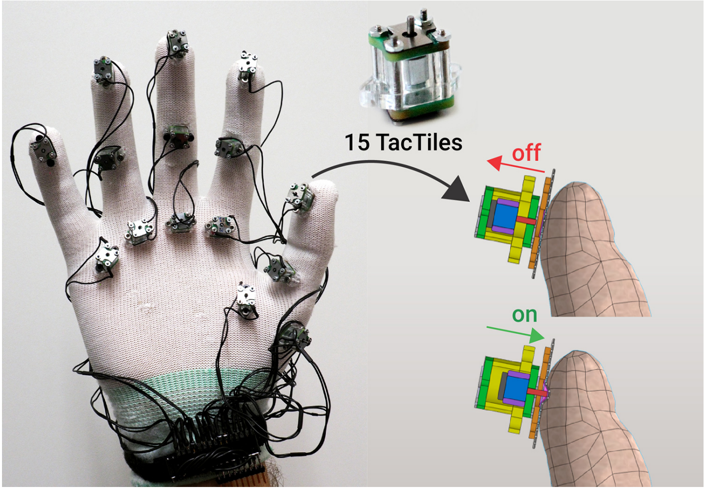

FIGURE 3.1: The final array of 15 TacTiles mounted on a glove. Each actuator can be controlled individually, rendering a haptic sensation by making contact to the skin with a retractable pin (red).

integrate 15 of these actuators in a dense arrangement on a textile glove, exploiting a large part of the Palmar (inner) surface of the hand. Together, these improvements enable continuous usage under high-stress scenarios in VR and open up new types of interaction. The proposed system called TacTiles, is shown in 3.1.

3.3.1 *Principle of Operation*

We employ the same principle of operation as MagTics, which is based on a bi-stable electromagnetic latching mechanism. This consists of a moving permanent magnet switching between two stable positions via application of short pulses of electrical current (see Figure 3.2.b). The magnet moves inside a cylinder, covered by coils on a PCB to deliver the switching impulse. The magnet stays in any of the stable positions due to attraction to the latching plates mounted in parallel to the PCBs. Once latched, the mechanism consumes *no further power* until switching again. The magnets are laterally shielded to avoid the magnetic cross-talk between neighbouring actuators. A pin attached to to the magnetic side facing the skin can transmit the movement and deliver contact/no-contact sensation to the user. The original design achieved a holding force of up to 200 mN with a stroke of 2 mm, values that were shown to provide a convincing and distinguishable tactile feedback [60].

3.3.2 *Design for VR*

The design in MagTics has several drawbacks that limit its applicability to fine-grained VR interaction, which we overcome in this thesis. First, 140 *W* of pulsed power were required over a time span of 20 *ms* for switching. This large amount of Joule heating limits the switching frequency, as the heat is accumulated faster than the device can dissipate it. This limitation renders active surface and texture exploration infeasible, which requires high-frequency and continuous activation. Second, the symmetric by-stable design leads to strong sensation during both engagement and disengagement with the skin. In our preliminary tests, we found that this is perceived as highly unnatural. It is therefore desirable to have strong impact force but no sensation during disengagement. Finally, MagTics require a pitch of 1.7 *cm* between actuators, making it difficult to arrange them densely.

We alleviate these main limitations of MagTics through 1) asymmetric latching, 2) smaller form factor, 3) new electronic PCBs consuming far less power, resulting in less heat, 4) faster fabrication and much easier assembly and re-configuration, 5) a damping mechanism to reduce feedback during disengagement.

The resulting improvements allow TacTiles to perform an average of 120 switches per minute without overheating, enough for demanding VR interactions such as exploring a haptic surface. TacTiles can also be activated in short bursts of 50 *ms* at up to 200 *Hz*, giving it a wide dynamic range. Asymmetric latching and damping address the unwanted sensations when disengaging from the skin. The reduced form-factor enables denser arrays allowing for glove integration.

3.4 HARDWARE DESIGN

3.4.1 *Materials and Assembly*

TacTiles are haptic feedback cells designed to be created with a simple and reliable fabrication process. They only require parts that are readily available (e. g. in modern FabLabs), shown in Fig. 3.2.a. Latching plates are made from low carbon laser-cut steel and plastic parts are made of laser-cut acrylic (PMMA). Screws, nuts, pins, magnets and shields are standard elements and the PCBs can be ordered online. The PCBs are symmetrical and electrical contacts are established by exerting pressure only. This allows for quick assembly, reconfiguration and tuning of parameters (e. g. adding PCB layers to increase latching force).

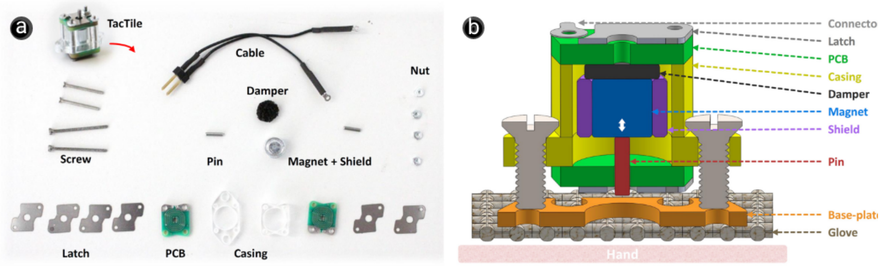

FIGURE 3.2: TacTiles parts and structure. A) Pre-assembly view of the TacTiles parts and B) schematic of TacTiles assembly.

3.4.2 *Form Factor*

While actuators need to be compact, the force provided by electromagnetic actuators scales with the volume of the magnet. This leads to a trade off between actuator size and perceived force.

To achieve a higher actuator density while maintaining a high force profile, we reduced the footprint by 41.6%, compared to MagTics. By employing a PCB technology with thicker and larger number of layers, we were able to improve the power consumption and make the coil more compact. A single TacTile measures $8 \times 9.5\ mm$ (i. e. footprint of only $71.5\ mm^2$). Using mostly plastics reduced its weight by 45.5% to $1.8\ g$. This weight reduction is crucial for comfort when the actuators are worn on the fingertips with no perceivable sagging of the glove's fabric even when the user's hand is facing down.

3.4.3 *Power Envelope*

While bistability allows persistent haptic contact/no-contact sensation without drawing power, power consumption during switching must be low enough to allow switching at higher frequencies without failure due to excessive heating. Our design only requires a square voltage pulse of maximum 5 *V* and 10 *ms* to switch the actuators, depending on travel time and travel distance. *Once latched, the actuator does not consume any power*. The power consumption is further improved by leveraging a denser PCB design, composed of six 105 *µm* thick copper layers inside and two 75 *µm* thick copper layers above the PCB. Each layer contains five turns of 220 *µm* wide tracks. Compared to MagTics, we thus increase the amount of copper per PCB by 390%. However, the PCBs' total thickness increases only from 1.2 *mm* to 1.4 *mm*. As a result, the new PCBs have twice as many turns while having the same total electrical resistance of $1.38 \pm 0.1\Omega$. Our design is 7 times more power efficient and requires pulses of only 18 *W* to switch reliably, compared to 140 *W* reported for MagTics. During our user studies, we recorded a mean rate of 44 switches per minute in each actuator. This represents an average power of only 130 *mW* on each TacTile under usual operation and thus a large reduction in heat generation (see Section 3.5.2).

3.4.4 *Asymmetric Latching and Damping*

Touch interaction in VR can be broken down into three separate events that the user feels: 1) the initial contact, 2) the continuous sensation of pressure on the finger, and 3) disengaging from the object. In the real world, we naturally integrate these three events. At each stage there is a particular expectations as to how they feel. For example, when sliding the finger over an object we expect to feel pressure, whereas on object release, we expect minimal output (unless the object surface is sticky). When designing our actuators for VR gloves we experimentally found that during the actuator recoil, the vibration created by the magnet latching to the 'off' state may be perceived as an additional impact, leading to an unnatural sensation when releasing contact with virtual objects. We therefore introduce an asymmetric latching mechanism, that can maintain force generation on the skin, while reducing the recoil sensation. We found that an asymmetric latching plate of 400 *µm* thickness on the side facing the skin, and of 200 *µm* thickness respectively, gave the best results in terms of force generation, i. e. a strong contact force while maintaining a soft landing when retracting the actuator. When the magnets hits the PCB it produces an acceleration peak of 10 *g*, followed by a phase of magnet bounce, which is perceived as high frequency vibration (Fig. 3.3). Adding dampening materials between the magnet and the 'off' latching plate drastically reduces this effect (Fig. 3.3, dotted green and orange line).

3.4.5 *Modes of Actuation*

Due to its ability to provide actuation with high frequency, we propose two different modes of actuation. In *contact mode*, the pin of an actuator moves towards the skin until it makes contact, i. e. the skin gets stimulated directly by the pin. From our damping tests, we know that this requires a movement time of 6 *ms* (see Figure 3.3). The pin then rests on the skin (using the bistability of the device) until disengagement is triggered (e. g. a user ending a collision with a virtual object). In *pulse mode*, the pin moves towards the skin for only 3 *ms*, and then retracts immediately. Therefore, while users feel a pulse, the skin is not directly stimulated by the pin but the indirect movement of the whole actuator. This is comparable to the haptic feedback of piezo actuators, but only a single vibration. Pulse mode can also be fired in successive fashion using a 2 *ms* pulse towards the skin, and a 3 *ms* pulse in the reverse direction to render vibration sensations. In our experiments in VR, we use the *contact mode* for rendering collisions with objects and the *pulse mode* for rendering spatial haptic patterns.

3.4.6 *Integration of the Actuators into a VR Glove*

Compact and lightweight actuators are the first step towards a comfortable haptic glove. Electromagnetic actuators must also be able to be placed arbitrarily in order to coincide with areas of high mechanoreceptor density on the user's hand. We hence chose to forgo a monolithic approach in favor of a distributed, reconfigurable array of actuators. Based on the work by Murakami et al. [**murakami2010decision**], we choose the most relevant locations to provide tactile information when manipulating objects, as shown in Figure 3.1. These locations cover the five fingertips, four intermediate phalanges (little finger excluded), five palm locations close to each finger (reverse of each knuckle), and the exterior of the palm below the thumb. We use a commercial textile glove as starting point for our system and glue a 2 *mm* thick base-plate with 2 threads to the textile in the desired locations. This way, the actuators can be mounted quickly by simply screwing them onto the base-plates (see Figure 3.2).

3.5 SYSTEM EVALUATION

We evaluate key parameters of the hardware design: the effect of additional dampening, and the thermal behaviour of the actuators.

3.5.1 *Effect of Damping*

During switching, the magnet is accelerated from being in contact with one latching plate to the other (2 *mm* travel distance). Due to the high acceleration, the magnet bounces and generates high-frequency vibrations which are perceived as being unnatural. We evaluated different damping materials to alleviate this effect. We compared discs of 5 *mm* in diameter and 1 *mm* thickness made of two different materials: PDMS silicone and foam. We expect that silicone will present an elastic bounce combined with damping, while for the foam we expect the damping to be the dominant.

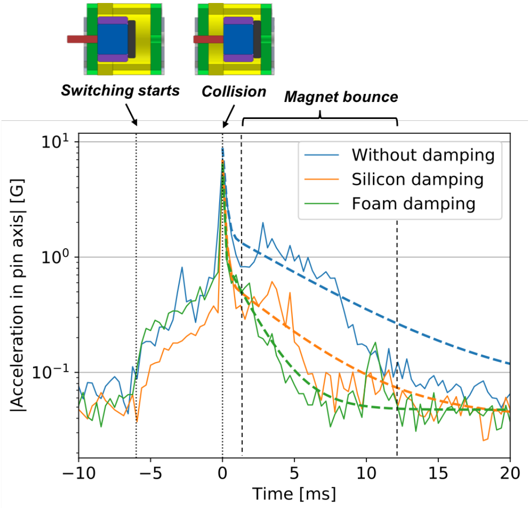

FIGURE 3.3: Acceleration switching profiles with and without damping (log-scale). Each curve is the average over 20 switches measured with an accelerometer *ADXL345*, previously aligned in the time domain. At time = 0 the magnet hits the 'off' state (max acceleration due to the collision $10g$). Dashed lines show exponential decay fit (see text).

To experimentally verify the accelerations during switching, we attached differently damped TacTiles to a one meter long pendulum of 40 *g* mass with an accelerometer *ADXL345* attached to it. We sampled the acceleration at a rate of 3200 *Hz* over 20 switches. Fig. 3.3 shows the acceleration profile for the three different cases over time: no damping, silicone damping and foam damping. The time axis (t) was offset so the collision with the latching plate corresponds to 0 *ms* (maximum acceleration). The acceleration axis is in log scale. Two regimes can be distinguished: a switching period from $t = -6\ ms$ up to the collision, and a second period of magnet bounce from $t > 0$ until it reaches the resting position. We observe that the vibrations during the bouncing period decay much faster in the case of using damping compared to the baseline. Furthermore, foam appears to be more promising since the accelerations decay fastest. An exponential decay fit (dashed lines) confirms this: $\tau_{no-damped} = 5.9\ ms$; $\tau_{silicone} = 4.1\ ms$; $\tau_{foam} = 1.8\ ms$. Based on this result, we built our VR glove with 15 TacTiles using foam as damping material.

3.5.2 *Experimental Test of Heat*

The Joule heat generated on electromagnetic actuators is a well known limitation when using them intensively. TacTiles benefits from the latching mechanism that requires only power for switching, added to the improved coil design described in Section 3.4.3. We measured the actuator heating under an intense switching rate to test the thermal behaviour of the device. Figure 3.4 shows the temperature as a function of time when read from a thermocouple placed between the latching plates and the PCB of the device. For this experiment we switch the actuator continuously for 6 minutes at 2 Hz (average power $P_{av} = 0.36\ W$). Then we let the device cool down under natural convection with no power applied. The curve of temperature vs. time is well described by Newton's law of cooling and its heating equivalent,

$$\begin{aligned} T_{heating}(t) &= T_{room} + \Delta T_{max} \left(1 - e^{-(t-t_0)/\tau_{therm}}\right) \\ T_{cooling}(t) &= T_{room} + \Delta T_0\ e^{-(t-t_0)/\tau_{therm}}\ , \end{aligned}$$

where T represents temperature and t time. These equations describe how temperature rises and decreases always with a characteristic time τ_{therm}. The asymptotic value of the temperature reached in 3 to 5 τ_{therm} is proportional to the average power applied, that means $\Delta T_{max} = c * P_{av}$. According to the measured values $\tau_{therm} = 102 \pm 5$ seconds and $c = 69 \pm 10$ °C/W for our device.

Thermally speaking, we can consider the actuator as a mass able to accumulate heat and then dissipated it into the environment. Only the mean power matters, i. e. the average number of switches occurred in the past 3 to 5 τ_{therm}. As shown in Figure 3.4, we can ensure TacTiles are kept in a safe temperature range if the average switching rate is below the 120 switches

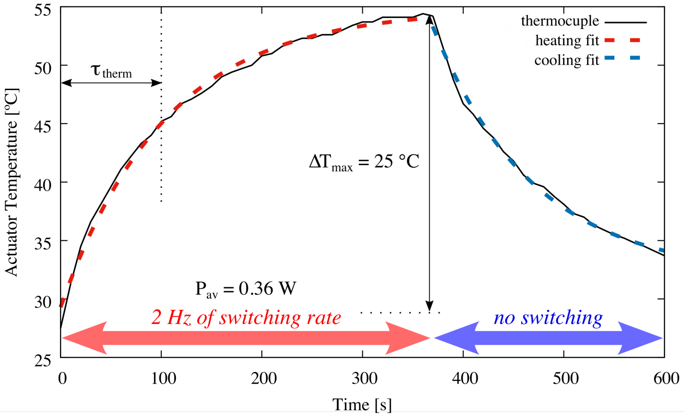

FIGURE 3.4: Heating (6 min at 2 Hz, red arrow region) and cooling temperature curve (blue arrow region) of a single TacTiles actuator. Dashed lines indicate the exponential fit according to the Newton's law of cooling/heating. We obtain the values $\tau_{therm} = 102 \pm 5$ seconds and $\Delta T_{max} = 25$ °C/W, when actuating with an average of 120 switches per minute ($P_{av} = 0.36$ W).

per minute (2 *Hz*). This does not impact the maximum speed of activation, which is constrained only by the amount of time required to move a single pin. For example, it is possible to render repeated *pulses* at up to 200 *Hz* in short 50 *ms* bursts on average every 5 seconds without overheating. For real-use testing, we recorded activations for each actuator on the hand during study 3, which constitutes a demanding haptic environment, and found that the actuators were on average switched between 38 and 49 times per minute. Using the upper bound, this can be translated in a temperature increase over room temperature of only $\Delta T = 10.1°$C. The increase in temperature was also not noticed by participants, likely due to thermal insulation from the glove. This shows that TacTiles are able to be go under intense use without any overheating problems.

3.6 USER EVALUATION

To better understand the efficacy of TacTiles and its possible applications, we conducted 3 studies with studies 1 and 2 focusing on quantitative aspects and study 3 exploring qualitative aspects. The studies were designed to answer the following research questions:

- **RQ1:** How well can participants localize individual actuators on their hand when using TacTiles?
- **RQ2:** How well can participants discriminate the two modes (i. e. contact mode and pulse mode) of TacTiles?
- **RQ3:** What is the *just noticeable difference* (JND) of frequency for spatial haptic patterns with TacTiles?
- **RQ4:** How well can participants perceive differences in *direction* and *frequency* between visual and corresponding haptic spatial patterns with TacTiles, and what is the effect this has on realism?

3.6.1 *Study 1: Mode Discrimination and Localization*

The aim of this study was to understand how well participants can 1) localize individual actuators on their hand and 2) differentiate between different modes of tactile feedback. TacTiles supports two modes of tactile feedback: direct skin contact (*contact mode*) and pulse without skin contact (*pulse mode*). For each actuation, participants were tasked to answer at which locations they felt the feedback, and which mode they felt.

3.6.1.1 *Participants*

We recruited 10 unpaid participants (age M = 28.3 years; SD = 3.1; 3 female) from the local university campus. 8 participants had used VR equipment before, 7 had experienced haptic feedback devices, based on self-reports. Each participant signed an informed consent form prior to the study informing them about the data we recorded and the possibility to interrupt or end the study at any time.

3.6.1.2 *Procedure and Tasks*

The study was conducted in a quiet experimental room. Participants were equipped with the TacTiles glove on their right hand, shown in Figure 3.5 (*left*). They were introduced to the device and completed a short training. During the training, five individual TacTiles were actuated in both modes (i. e. 10 training trials). Participants were informed about the mode and were allowed to look at the device during actuation. During the study, participants wore noise-canceling headphones through which white noise was played. This was done to mask any auditory sensation of the magnetic actuators.

FIGURE 3.5: Apparatus for the study 1 and 2. *Left*: a participant in part 1 (localization and mode). They were asked to indicate the position and mode of the stimulus and not look at the device during actuation. *Right* shows a participants in part 2 (JND) wearing a VR headset.

Participants were asked *not to look at the device* but an instruction sheet showing the locations of the individual TacTiles (Figure 3.5, *left*). Based on a predefined random order, individual TacTiles were actuated with a random mode. Participants were asked to state the location of the actuated TacTile and its mode. A trial consisted of rendering a single mode of feedback at one of 15 locations on the hand. Each mode was played back once at each location for a total of 30 trials per participant. There was no time limit set for answering, and participants could ask for an actuation to be repeated once. This experiment lasted approximately 15 minutes per participant.

3.6.1.3 *Results*

Participants were able to correctly identify the location and mode of individual TacTiles with an accuracy of 78.7% (SD = 11.2%). Looking at localization only revealed an accuracy of M = 87.7% (SD = 8.6%). Participants were able to discriminate the mode in M = 87.0% (SD = 7.6%) of trials. Figure 3.6 shows the results in more detail.

3.6.2 *Study 2: JND of Spatial Haptic Patterns*

While several studies have been conducted to understand the influence of a single actuator on the perception of parameters such as softness (e. g. Perez et al. [**perez2015soft**]) or edge sharpness (e. g. Park et al. [**park2017sharpness**]), we were interested in perceptual effects of our multi-actuator device. Specif-

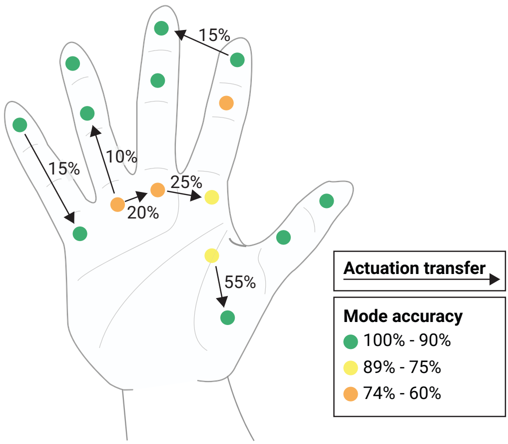

FIGURE 3.6: Breakdown of mode errors and localization errors. Colored dots indicate actuator positions and their mode accuracy. Arrows and number indicate how often an actuator (arrow start) was mistaken for another actuator (end arrow). Localization errors below 5% (i. e. 1 out of 20) were omitted from this figure.

ically, we were interested in participants' ability to discriminate *frequencies* of different spatial haptic patterns.

3.6.2.1 *Stimuli*

We chose 4 different spatial patterns from literature [75] (see Figure 3.7) that were replayed to participants while they moved their hand over a virtual object. The speed of the patterns was not connected with the speed of their hand movement. Each pattern is encoded in the time domain by setting the mode of each actuator to render a *pulse* at discrete time points. Although the participant moves their hand to trigger the patterns, each pattern is played back at a fixed speed to avoid the confounding variable of hand speed. This allows us to render haptic patterns on the hand in a repeatable and controllable manner (i. e. separate frequency and movement velocity). The presentation order of the patterns was counterbalanced using a Latin square.

3.6.2.2 *Procedure and Tasks*

In this second study, participants were wearing a VR headset and the TacTiles glove on their right hand, shown in Figure 3.5 (*right*). Since this

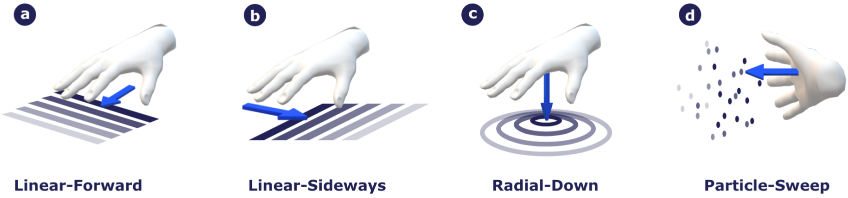

FIGURE 3.7: Spatial haptic patterns coupled to their respective hand movements: a) Linear pattern as the hand moves forward, b: Linear pattern as the hand moves sideways, c: Radial pattern as the hand sweeps down, c: Particle pattern as the hand sweeps sideways. Each line, ring, or set of dots is rendered once as it propagates through the hand. Opacity represents, i. e. lines rendered later are have less opacity.

study was conducted immediately after study 1, the same set of participants stayed in the same room and were given a new set of instructions and tasks. During the study, participants wore noise-canceling headphones through which white noise was played. Participants' hand posture and position was tracked using a Leap Motion controller that was mounted to the front of the headset. The experimental software was programmed in Unity 2017. The virtual environment consisted of two cubes with a side length of 30 cm, with a horizontal distance of 15 cm. When touching an object, a spatial haptic pattern with a specific frequency was played back. Participants were asked to decide which of the two objects exhibited a higher frequency. Based on pilot studies, we select the reference stimuli of the four patterns to be one second long in total. The line patterns (a and b in Figure 3.7) refresh five times during a stimuli, giving 5 Hz reference patterns. The radial pattern (Figure 3.7.c) refreshes 4 times, while the particle pattern (Figure 3.7.d) refreshes 10 times within a stimuli. This gives 4 Hz and 10 Hz of reference frequency for the radial and the particle patterns, respectively. In all cases, we chose a variable step size with an initial value of 5%. In each trial, participants touched both objects (i. e. experienced both frequencies). They were asked to move their hand across the objects to reflect the direction of the pattern, e. g. sweep from the back to the front of the object for the first linear pattern in Figure 3.7. For one object, the reference frequency was played back. For the other object, the approaching frequency was played back. Participants were unaware which frequency was used and were allowed to touch each object an unlimited number of times. The assignment of frequency to object was randomized every time participants provided an answer. The initial value for the reference frequency was always +15%

of the reference frequency. We only used a positive values to decrease the total time needed for the experiment given that positive and negative JND approaches are typically symmetric [76]. After each trial, we asked participants to identify which of the two frequencies was perceived higher. A correct response brought the frequency in the next trial a step size towards the reference frequency, and vice versa [77]. The step size was decreased to 1% once participants hit an absolute delta of 5%. We chose this procedure in order to get higher accuracy after the initial approach. The procedure was repeated until the direction was reversed 3 times and the reversal points were averaged to get the JND of each pattern. The experiment lasted approximately 30 minutes per participant.

3.6.2.3 *Results*

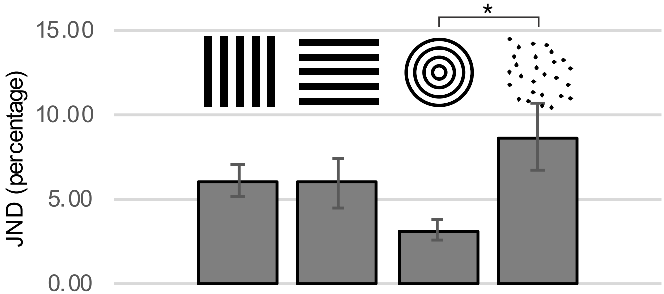

FIGURE 3.8: Results of the JND experiment. Error bars indicate standard error.

Participants were able to adequately sense a difference in frequency of *M* = 6% (*SD* = 4.5%). Figure 3.8 illustrates the results. Note that lower JND is better. For the vertical and horizontal linear pattern, the JND was *M* = 6.2% (*SD* = 2.9%) and *M* = 6.0% (*SD* = 4.6%), respectively. The radial pattern exhibited the lowest JND, *M* = 3.2% (*SD* = 2.1%). The particle pattern, exhibiting no distinct structure, had the highest JND of *M* = 8.7% (*SD* = 6.2%).

A Kruskal-Wallis test revealed a statistical main effect between the four patterns, $\chi^2(3) = 9.498$, $p = 0.023$. A series of post-hoc t-tests revealed that only the difference between the radial and the particle pattern was statistically significant, t=2.879, p = 0.04.

This shows that for distinct patterns (e. g. linear or radial), participant can distinguish very small differences in frequencies, and even for patterns that lack this structure, these noticeable differences are well below 10%.

3.6.3 *Study 3: Experience Test*

To gather more insights into the usage of TacTiles in VR, we performed an experiential study where participants were seated in front of a virtual desk (see Figure 3.9) containing multiple objects with different visual and spatial haptic patterns. We were interested in the realism of interactions when augmented with TacTiles, as well as the influence of spatial haptic pattern frequency and direction on the participant's sensation. We furthermore wanted to know participants reaction to the two modes and their combination in a single scene.

3.6.3.1 *Participants and Apparatus*

We re-invited 6 unpaid participants from the first study to participate (age M = 29 years; SD = 2.7; 2 female) in the experiential study. The study took place one day after the first. We decided against performing all 3 studies in one session to avoid fatigue.

In this study, we replaced the Leap Motion controller used for hand tracking with an OptiTrack setup (Motive:Tracker 2.1, 10 Prime3 cameras) to increase the hand pose tracking stability during 3D scene exploration. The TacTiles glove was augmented with 4 passive infrared markers and 5 active LED tags at the fingers to enable hand pose tracking. The coordinate systems were aligned via the built-in calibration procedure in the Motive:Tracker software.

3.6.3.2 *Environment and Procedure*

We designed a simple virtual environment (Figure 3.9) with 5 different types of objects: a table exhibiting a particle type haptic pattern; 3 bamboo segments and 3 wood blocks with a linear pattern that a) matched the visual texture direction and frequency b) matched the direction but had twice the frequency, or c) matched the frequency but not the direction; a fan that rendered a radial haptic pattern in mid-air, the frequency of which increased with proximity to the object; a cup and a computer mouse, that would render continuous (permanent) contact rather than pulses (i. e. contact mode triggered on collision). Spatial haptic patterns were rendered during collisions between the hand mesh and the object, while continuous contact mode activated a predefined set of actuators for that object.

Participants were familiarized with the setup and guided through the environment. The experimenter asked them to explore the wood blocks

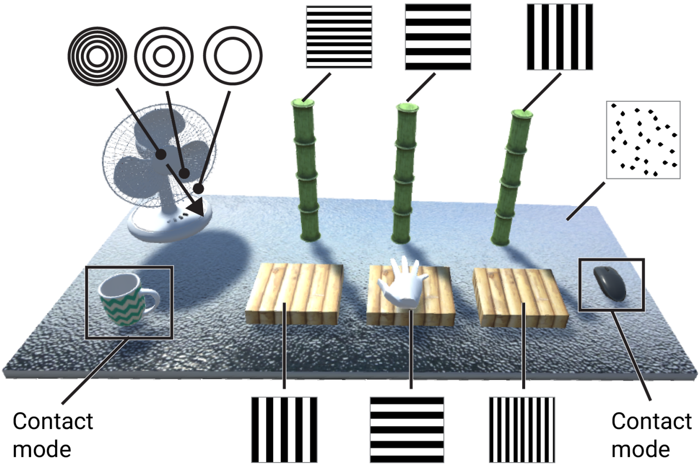

FIGURE 3.9: Environment with objects and haptic surfaces used in our experience test.

by gliding their hands horizontally over their surfaces, and similarly, by exploring the bamboo segments by moving from the top to the bottom of the segment. Participants were asked to pay attention to how well the spatial patterns they felt on their hands matched the object's visual texture and based on this, select the object which they perceived had the closest match. Next, participants were asked to localize the mouse by touch alone, first by placing their hand on the table and then moving towards the mouse until they felt the contact mode engage on the actuators. After completing the tasks, participants could interact with any object in the scene and comment on the realism of the interaction. The study lasted approximately 25 minutes per participant.

3.6.3.3 *Results*

Overall, the virtual objects with matching visual and haptic direction and frequency were the rated as the best matched, with three out of six participants choosing the first wood block and again, three out of six choosing the second bamboo segment. All six participants could easily tell the difference between the two modes of interaction.

SENSITIVITY TO FREQUENCY AND DIRECTION OF PATTERNS. All participants were able to discriminate the object in the set of three with higher frequency than the others. However, for directionality, only two participants noticed that the second wood block texture was not matching, while four out of six participants noticed that the third bamboo segment texture did not match. Participants often commented when they perceived a mismatch between visual and haptic patterns, for example, when referring to the visual bump frequency on the middle bamboo: *"The middle one is more confusing. There is a mismatch."* (P3) and *"The middle one is random bumps. They don't correspond to what I'm seeing visually."* (P2). Referring to the in the middle wood block, one participant correctly perceived that it was direction that was mismatched: *"It's rotated. If I imagine it being rotated it would align nicely."* (P4). Some participants noted that the visual sense could override their haptic sensation: *"I had a feeling they all matched the visuals. Maybe because I am so focused on the visuals. I had the feeling that they all go from bottom to top."* (P1).

CONTACT AND PULSE MODE. Since in the previous study, contact mode was only rendered for brief periods at a single location, participants were surprised that the actuators had the capability to continuously push against their skin, which was magnified by the combined effect of the array. Participants considered the effect of touching an object both distinct and convincing: *"With the cup I feel that I've made contact with the object."* (P2) and *" The feeling is very penetrating, but very convincing. You cannot miss the cup."* (P1). While not explicitly asked to do so, many stated a preference for the contact mode over the pulse mode used for spatial patterns: *"Purely just in terms of touch, it more closely resembles what I would expect."* (P6) and *"At least with the feedback I can feel when I've made the contact with the cup, but I don't with the block. So I like the cup better."* (P2).

REALISM OF FEEDBACK. In regards to realism, the effect of 3D geometry played a heightened role in the expectations of users: *"The bamboo has the extrusions - which I should feel. The wood blocks are flat, so I should feel a shear, not a travelling bump."* (P3). However, for the wooden blocks, the tactile feedback did not match the flat visual appearance: *"I see wood, and I am expecting it to be a coarse surface. Not smooth."* (P6) On the other hand, participants felt the spatial particle pattern on the table was a good match for the visually bumpy surface: *"I like the table though. Its meant to be textured. The table is great."* (P3).

EYES-FREE LOCALIZATION. Four of the six participants could localize the mouse on their first attempt without explicitly looking at it by gliding their hand over the desk. While having tactile feedback was an effective aid, it was not as efficient as it could be due to the lack of grounded feedback: *"If you're on the right plane, then it's very easy. But you can go through the table."* (P4).

3.7 DISCUSSION

The main contribution of our work is a novel design of an electromagnetic tactile actuator that exhibits a small form factor, low power consumption, low heat generation, and a strong holding force. By introducing a dampening material directly in the actuators, we were able to mitigate the transfer of unwanted vibrations onto the hand. We further characterized the thermal behavior of a single actuator and found that even in a demanding haptic environment the heat generation was at most 10.1°C above room temperature. To cover a large set of mechanoreceptors on the hand, we carefully chose the placement of 15 actuators and integrated them into a wearable form factor, i. e. a haptic glove that can be used in VR. By introducing two different modes of actuation, TacTiles can render both haptic surface information (see Figure 3.7) and the distinct sensation of touching an object. Taken together, these contributions allow TacTiles to satisfy the proposed design guidelines of *wearability*, *low-power*, and *high-fidelity*.

In our first experiment on mode discrimination and localization, participants achieved 87.0% and 87.7% accuracy respectively. One actuator on the palm exhibited comparably poor localization of 55%, indicating that at this specific location, a single actuator instead of two might have been sufficient. This goes in line with findings from neurobiology on the sensitivity of different areas of the human hand due to differently distributed mechanoreceptors [69, 78]. Actuators along the base of the fingers exhibited higher localization, although they were closer together (between 1.5 *cm* and 2 *cm*). This needs to be considered for rendering haptic sensations, for example, by increasing the density of haptic actuators near the fingertips, while spreading them out near the palm. However, it should also be taken into account that it's even possible to confuse localization between two fingertips [79]. Increased density at the fingertips may exacerbate this effect.

We also found that participants were able to easily distinguish differences in the frequencies of spatial haptic patterns of 3% to 6%. Even for quasi-random patterns, participants were able to distinguish frequencies with a difference of only 8.7%. Pairing this with visual feedback from our second experience test, we believe that TacTiles enables a very rich set of haptic rendering capabilities.

In our experience test, participants confirmed the importance of a *correct correspondence* between a visual and haptic representations of patterns. Previous work on rendering tactile feedback to the fingertip in VR has shown a high tolerance for directional mismatch between visual texture and haptic patterns (e. g. Haptic Revolver [70]). In our experiments, however, when rendering feedback to the whole hand, we found that participants were able to perceive both orientation differences and frequency differences. Thus, while full-hand tactile rendering affords richer experiences in VR, the directional and frequency compatibility between visual and haptic patterns should be considered carefully.

In terms of realism, in pulse mode, the haptic sensations produced by TacTiles would be more appropriate to render 3D features rather than 2D visual textures. Participants were also more sensitive to directional mismatches when presented with 3D features, such as in the bamboo segments. The bump-mapped 2D table surface felt realistic, while the flat appearance of the wood blocks elicited expectations of shearing. Enclosing or grasping objects triggered contact mode. Participants commented that this more closely resembles their expectations in terms of haptic feeling. This mode could additionally provide information about the location and state (i. e. grasped) of an object which could be used in more interactive scenarios.

Providing both contact and haptic surface rendering capabilities in a single device is desirable as we use these modes routinely in everyday life, for example, during object identification tasks [80]. With TacTiles, we show that by combining both modes in a single device, we can enable new forms of interaction in VR such as the localization of objects that are out of sight by touch alone.

3.7.1 *Limitations and Future Work*

While the 15 actuators of TacTiles are able to render haptic feedback to key regions on the fingers and palm, the full haptic surface of the hand still has much room before its fully saturated. In addition, areas on the hand such as the sides of the fingers are difficult to exploit, as the size of the actuator can physically impede finger flexion and abduction. Future work can address this by further reducing the size of TacTiles and placing them in even denser arrangements.

Another area which could improve VR realism is providing different levels of contact pressure which could be based on the levels of penetration into an object.

Going beyond purely tactile feedback, it would be a fruitful direction of future research to combine tactile arrays with kinesthetic haptic devices such as our recently introduced device DextrES [54] in order to increase realism of touching *and* holding objects (cf. user comment on muscle tension).

We believe that TacTiles has the potential to serve as haptic feedback device beyond the current form factor as a glove, for example, in an forearm sleeve. TacTiles enable the design of novel, larger scale devices due to its low power consumption and low heat generation. This may enable so far unexplored rendering of tactile sensations in VR.

3.8 SUMMARY AND OUTLOOK

This chapter presented a novel on-hand haptic tactile device that integrates an array of electromagnetic actuators.

The main technical contributions are the actuator and the array integration on the hand. The actuator has a small form factor (1 cm^3, 1.8 g), minimal power consumption (130 mW) and low heat output, enabling long use without overheating. The array packs 15 actuators in close proximity, and enables novel VR scenarios.

We leverage these technical contributions to enable new use-cases. First, we developed two different actuation modes, i. e. *pulse* and *contact* mode to allow the rendering of different spatial patterns and object contact. Users successfully spatially localized actuations on the hand, and discriminated between the two modes.

An experiential study showed that the array can convincingly render continuous touch with an object, convey haptic surface information, and even allow users to localize objects in a scene without looking at them. This is an important and novel finding that has not been previously demonstrated. It shows that users are able transfer their already existing tactile skills to the VR domain.

The main takeaway is that a combination of form-factor, efficiency, and performance are all required to enable more advanced use cases in VR. Without a small enough actuator, achieving a minimum saturation (targeting a minimum of 1 actuator for each section of each finger) is not possible. For *interactive* uses cases that are omni-present in VR, we also require continued use without overheating, which requires power efficiency. TacTiles also represents a significant step towards employing an expanded range of our mechanoreceptors ability to sense different frequencies, that is beyond the typical high-frequency vibration of vibrotactile arrays.

Although the tactile glove can provide feedback across a range of frequencies, it is missing a key component of haptics - kinesthetic feedback. The actuators developed in this chapter, although miniaturized and designed to reduce encumbrance, cannot stop the hand from moving through objects. In the next chapter, we explore the direction of kinesthetic feedback on the hand based on compliant elements that can flex and deform with the complex movements of the hand.

4

KINESTHETIC ON-HAND HAPTIC-INTERFACES

In the previous chapter, we focused on providing on-hand *tactile* feedback through a miniaturized tactile array. In this chapter, we continue this line of work, but we address an even more challenging problem - providing *kinesthetic* feedback to the hand. The challenge stems from the limited space available on the hand, and in particular the high dexterity of the hand, that can both be difficult to resist due to the high number of degrees-of-freedom, and difficult to not impede due to that same flexibility. As our hand is a primary means of input and feedback gathering, The creation of a haptic interface that both provides enough resistance while not limiting the *dexterity* of the hand remains an important open problem for VR immersion and utility.

In this chapter, we develop a novel flexible electrostatic-based clutch mechanism that can provide physical feedback in a range of grasps performed in a VR environment. The two key developments are the flexible strips that provide extremely high force-density at low weight (8 grams per strip), and the customizable 3D-printed mounting system to enable both sliding and contraction, crucial for interactive tasks. The approach of textile substrate combined with compliant active components is inline with the vision of *robotic garments* [13]. It is worth to emphasize that there are no motors or any kind of cable-pulley systems employed, the flexible clutches are *both*. During testing, users showed significant improvements using the novel interface in the majority of supported grasps over Piezeo/Vibrotactile-only baselines.

4.1 INTRODUCTION

The dexterity of the human hand enables us to perform a number of useful everyday tasks such as actively exploring surfaces and grasping and moving objects [5, 81]. In Virtual Reality (VR), dexterous manipulation using the hand is a popular means of interaction. It allows us to leverage learned motor skills and vice versa, to train for real-world scenarios in VR [76]. While rapid progress has been made on the input side (display and sensing technologies), haptic interfaces providing physical feedback to the hand lag behind in their fidelity. In particular, the lack of appropriate kinesthetic feedback limit our ability to precisely steer and place grasped objects in 3D space [82].

The ability to grasp objects is amongst the most useful skills we can perform in VR [29]. One challenging aspect is the wide array of possible grasps which require the fingers to be free to move into different configurations [81]. Traditionally, grasping feedback in VR has been supported via glove-based exoskeletons which create braking forces on the fingers [83, 84], render localized tactile feedback on the fingertips [56, 63], or combine aspects of both [59, 85]. These devices often employ complex mechanisms placed around the hand which may either add weight, constrain the movement of the fingers, or both. As a result, the full range of interaction capabilities of the human hand are under-utilized.

To address this challenge, we introduce DextrES, a finger-mounted haptic mechanism capable of achieving up to 20N of holding force on each finger when flexing inward. Our novel approach is based on electrostatic attraction to create a rapidly controlled braking force between two electrically charged strips of metal. We harness the resulting braking force to rapidly render on-demand kinesthetic feedback which blocks the motion of the fingers. Crucially, this allows for the design of a very thin and flexible form factor haptic interface for grasping objects in VR - a long standing goal which has thus far relied on space-inefficient and bulky mechanisms. Such an interface may also be generalized to function beyond VR, for example in Augmented Reality (AR), robotic tele-operation, and rehabilitation applications.

In contrast to a one-size-fits-all mounting solution, we integrate DextrES onto the index finger and thumb using modular fittings with different strip lengths inserted into 3D printed articulated guides to keep them moving smoothly. The strips are anchored onto the fingertip and wrist resulting in controlled frictional forces due to sliding when the finger is flexed. This mounting strategy allows for easy adaptation to different hand sizes. We couple our kinesthetic brake with miniature vibration motors mounted at each fingertip to signal initial contact events, mimicking a typical object manipulation cycle [22]. The resulting integration into VR allows freedom of movement for both the fingers and hand. The volume of the control electronics can be reduced to a few cm^3 with off-the shelf components, and the very low power consumption (less than 100 mW) allows for battery powered operation, providing a straightforward path to widespread real-world implementation.

We test the capabilities of DextrES in two experiments. First, we establish the just noticeable difference (JND) at different voltage levels and associate this to equivalent holding forces and perceived stiffness values. Second, we explore the impact of our feedback mechanism on the precision of four different grasps (see 4.5) in a VR environment. Results indicate that DextrES provides effective feedback and improves precision. Finally, we report qualitative results and user feedback on the perceived user experience when interacting in a free-form VR environment.

4.2 RELATED WORK

GRASPING IN VIRTUAL REALITY Researchers have replicated various types of forces in VR which are rendered when grasping an object, including gravity [86, 87], contact [63], shearing [57], rendering hard surfaces [84], and spring-back [29]. While many different types of grasps are possible [81], most grasping devices focus on finger-opposition power grasps. Grounded devices can create high fidelity feedback [88, 89], but are fixed in position. Hand-held VR controllers such as the Oculus Rift and HTC Vive allow the user to move his or her arms freely, but occupy the grasp thus prevent most hand movements, as well as only render coarse vibro-tactile feedback. Our approach consists of a thin form-factor electrostatic brake which can render kinesthetic and cutaneous haptic feedback in a wide range of grasps, affording a rich set of interactive capabilities.

KINESTHETIC HAPTIC FEEDBACK GLOVES Haptic feedback gloves have a long history in HCI and VR research [11]. A number of exoskeletal devices have been proposed to provide kinesthetic haptic feedback by blocking fingers' movement. We can distinguish between gloves based on pneumatic or hydraulic systems, and those based on electromechanical systems. Gloves based on fluids generally use pumps and valves [59] to displace pistons [90, 91] or activate jamming layers [92, 93] on the glove. These technologies are well-known, but difficult to miniaturize and can result in complex and bulky systems. Gloves using magnetorheological fluids have also been reported [94, 95]. Gloves using electromechanical principles mostly make use of motors or brakes directly on the glove (early versions used very long cables [83]). They use servo motors linked to bars/rods over the top of the hand [72, 96] or cables (tendon-based) across the top of the hand [83, 85, 97] to control finger position. They can actively steer mechanical linkages for finer control. Others have used motors to drive clamp braking mechanisms in order to block the finger [29, 84].

Motors or pumps can offer significant forces, but their performance decreases quickly when scaled down. It is very difficult to maintain sufficient force if scaled to volumes of a few cubic centimeters. Larger motors or pumps may be acceptable for VR but it would be disturbing for AR [11]. It would be ideal to directly be able to lock the finger position without using motors, just by blocking or clutching directly jointless flexible links (cables or strips) connected to the finger. We report here such a solid-state device, a jointless exoskeleton where the only moving parts are actuated by the user.

ELECTROSTATIC BRAKING MECHANISMS Numerous types of brakes, and more generally clutches, have been developed. Electromagnetic clutches are the most common type of electrically driven clutch, but are bulky and have high power consumption [98]. Mechanical latches can significantly decrease power consumption at the cost of increased complexity and reduced speed. Magnetorheological (MR fluid) clutches are simpler but heavier and consume more energy [99–101]. Such clutches or brakes are not well suited to haptic gloves for VR and AR applications, where the wearable haptic systems should be as comfortable and discreet as possible.

Early development of electrostatic (ES) clutches started in 1923 [102] and brakes have been used since 1957 [103]. Similar principles were used for lateral sliding ES actuators [104, 105]. ES clutches have been developed more recently for applications in robotics [106–109]. ES clutches are an interesting alternative to electromagnetic clutches at mm scale [106]. Their design is compact and simple, they can be low profile, flexible, and lightweight. Furthermore, they can generate high forces with ms time. Once the electrodes are charged, power consumption becomes extremely low. In light of these advantages, we leverage ES clutches as the core element of thin-form factor electrically driven haptic gloves able to provide kinesthetic haptic feedback.

4.3 SYSTEM OVERVIEW

The aim of our work is to provide haptic feedback for dexterous manipulation of virtual objects in VR and AR. The ideal feedback mechanism would be able to provide both *kinesthetic* and *cutaneous* feedback [21] while not encumbering the natural motion of the users fingers and requiring minimal user instrumentation. This is challenging to achieve since most

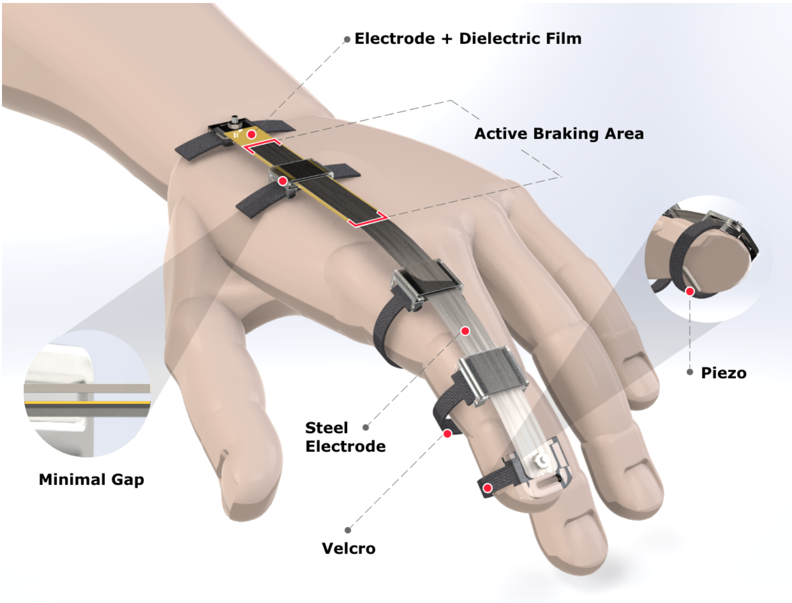

FIGURE 4.1: Schematic representation of the components of our haptic feedback mechanism. ES brakes along the fingers provide kinesthetic feedback and piezo actuators provide cutaneous feedback. The total weight is less than 8 grams. The ES brake is flexible, allowing for natural hand motion and hence a variety of grasps.

mechanisms that can provide sufficient force to block finger motion also require significant user instrumentation and are bulky.

Recent work has therefore opted to only support haptic feedback for a limited number of hand poses via actuation mechanisms that are build into a VR controller [8, 110].

In contrast, our work explores the use of an ES brake as thin-form factor kinesthetic feedback mechanism that can be integrated into a VR glove without hindering the natural motion of the fingers. Together with piezo actuators, mounted at the fingertips, DextrES, illustrated in Fig. 4.1, provides both kinesthetic and cutaneous feedback and can thus support a variety of grasps and enables precise VR manipulations.

4.3.1 *Challenges*

While the basic concept of ES brakes is straightforward and has been leveraged for non-haptic applications, designing an effective feedback mechanism for VR is not. To allow for unencumbered motion, the device needs to be thin-form factor and lightweight. Yet to provide effective haptic sensations, it must produce sufficient forces and must be easily mounted on users' hands of varying sizes. Throughout this paper we discuss our solutions to the following challenges:

(1) Fabricating an ES Brake. The brake must have sufficient force, speed, be robust, have a low-form factor and sub-Watt power consumption. These requirements impact material and thickness choices for the conductor and insulator layers. Since the brake must conform to the finger shape, the metal strips must be made from a strong and flexible material able to repeatedly bend, yet also provide a small restoring force, thus excluding ductile metals like copper or soft materials like conductive fabrics. The dielectric layer also impacts mechanical and electrostatic aspects and must hence be chosen to be thin enough to attain useful forces without requiring tens of kV, must have a high breakdown field and very low leakage current. Furthermore the dielectric must be flexible and smooth enough to allow accurate control of frictional forces to allow or block the sliding of the strips.

(2) Haptic Glove Integration The human hand moves in complex ways, typically modeled by a total of 27 degrees-of-freedom [111]. This dexterity poses significant challenges for the design of haptic feedback mechanisms. First, the braking mechanism must be securely mounted onto the users' hand such that it can effectively brake the motion of fingers in arbitrary poses. To effectively deal with metacarpal abduction (particularly challenging for the thumb) the force needs to be anchored at the back of the hand. Furthermore, to allow for natural motion the brake should not create friction when disengaged. Finally, to accommodate varying hand sizes the mounting mechanism must be flexible and modular.

(3) VR Integration. Since VR affords a very immediate form of interaction, a haptic feedback mechanism should be able to function efficiently under arbitrary hand poses. In particular, humans use a variety of grasps [81] and as many as possible of these should be supported. To support dexterous object manipulation in VR, the braking mechanism must be able to engage and dis-engage almost instantaneously to allow for rapid, natural hand motion corresponding to a realistic sensation of grasping and releasing virtual objects.

4.4 ELECTROSTATIC BRAKING MECHANISM

In this section we present the working principle, fabrication process and performance of the electrostatic kinesthetic haptic feedback brakes.

4.4.1 *Operation principle*

At the heart of our approach is a laminar electrostatic (ES) brake. Our ES brake consists of 18 cm long thin flexible metal strips that slide freely when no control voltage is applied, but generate up to 20 N of holding force per pair of strips when a suitable control voltage is applied. One of the key features of the ES brake is its thin form-factor, ideal for wearable applications. The active part of the brake is conformable to fingers and can be directly mounted or inserted on a glove. The brake mass on the glove is 8 g, and it is 6 mm high (including attachments).

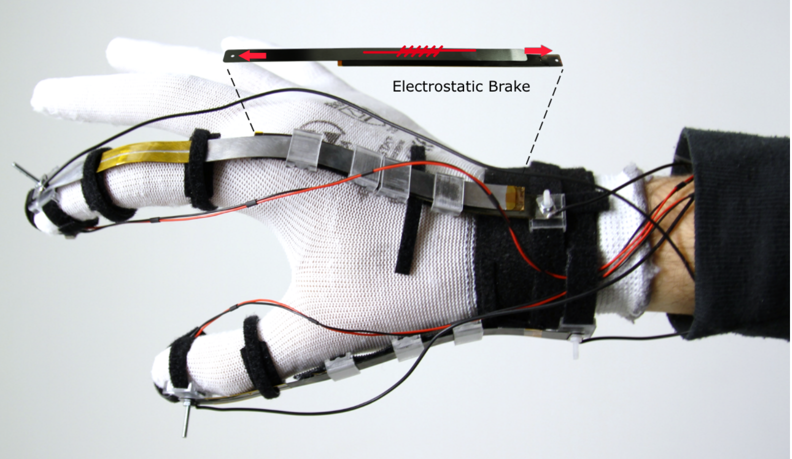

FIGURE 4.2: DextrES hardware and integration:The electrostatic brake is mounted on the back of the hand, and piezoelectric actuators on the fingertips for cutaneous feedback.

As shown in 4.2, the ES brake is attached to the glove, covering the back of the hand and the back of the finger. The high degree of flexibility allows excellent conformity to any hand shape. Fig. 4.3 shows the structure of a single ES brake strip (the strips can be stacked to increase force). Each brake element consist of two 100 μm micron thick steel strips, separated by a thin insulation layer bonded to one strip, thus forming a variable capacitor C_{strip}:

$$C_{strip} = \frac{\epsilon_r \epsilon_0 A}{d}, \tag{4.1}$$

where ϵ_r is the relative permittivity of the insulator between the electrodes, ϵ_0 is the permittivity of vacuum, A is the overlap area between the electrodes, and d is the thin dielectric gap between the electrodes. One strip (the "hand strip") is attached via the glove to a fixed point on the back of the hand, while the other strip (the "finger strip" is attached via the glove to a fingertip. When the voltage difference between the strips is zero, the strips freely slide with a very low friction, enabling full and unimpeded finger movements (Fig. 4.3b). In the simplest model of the device, when a voltage V is applied between the strips, an attractive electrostatic force $F_{compression}$ is generated between the strips, pulling them together (Figure 4.3c) :

$$F_{compression} = \frac{\epsilon_r \epsilon_0 A V^2}{2d^2}, \tag{4.2}$$

This electrically-controlled normal force leads to frictional forces between the strips, partially or fully blocking the movement of the finger. The friction force is less than or equal to the friction coefficient μ times $F_{compression}$:

$$F_{friction} \leq \mu F_{compression}. \tag{4.3}$$

The higher the applied voltage, the higher the friction force. Using this ES brake, we can thus apply a high blocking force to the fingers, providing kinesthetic haptic feedback.

The power consumption $P_{ESbrake}$ of the brake is determined by the energy to charge the capacitor multiplied by the switching frequency f:

$$P_{ESbrake} = \frac{E}{t} = \frac{1}{2} C V^2 f \tag{4.4}$$

Operating at 20Hz and 1.5 kV, the device power consumption is less than 60 mW.

4.4.2 *Fabrication of the ES brake*

After introducing the general working principle we now detail our solutions to challenge Nr. 1 as outlined in the System Overview Section.

We chose stainless steel as conductor since it is a reliable spring material. The bending stiffness of a strip scales approximately with the cube of

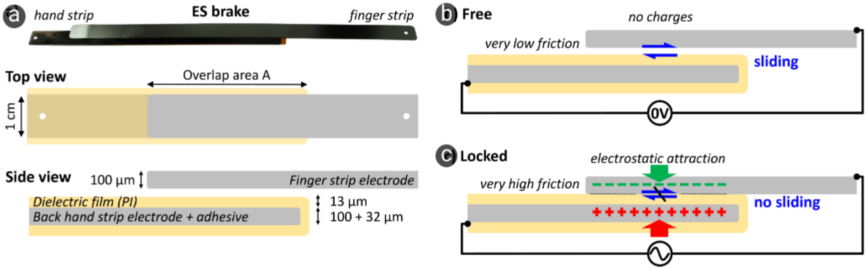

FIGURE 4.3: Schematic of the ES brake. a) Picture of the ES brake's hand strip and finger strip. Top view of the 1 cm wide hand strip and finger strip overlapping one another. Side view of the strips describing the materials and thickness of each layers. b) Free sliding between the strips when the voltage difference is 0 V. c) When a voltage is applied between strips, the electrostatic attraction pulls the strips together, drastically increasing the sliding friction.

the shim thickness. One must find a suitable compromise between being thick enough for the shim to slide easily without buckling or plastically deforming, yet thin enough so that the force to bend the strip is nearly imperceptible.

The fabrication of the ES brake strips consists of 3 steps: first, two strips 18 cm long and 1 cm wide were laser cut from 100 μm thick stainless steel sheets. Strips are shortened at a later time to fit the user's hand and fingers. Second, after polishing the edges, we deposited by hand onto the top surface of the "hand strip" a 32 μm thick conductive double-side adhesive and a 13 μm thin polyimide film, slightly wider than the steel shim to avoid short circuits. Polyimide has a high breakdown field of over 300 V/μm and a dielectric constant of 3.4. Using a conductive adhesive to attach the polyimide to the steel was a key step in reducing the driving voltage, as the adhesive is thus part of the electrode rather than being part of the dielectric. The fabrication process is straightforward, low-cost and readily industrialized.

4.4.3 *Control Electronics for ES brake*

To actuate and control the ES brake, we assembled a custom high voltage (HV) power source [112], based on a 2000 V DC-DC converter (XP Power, EMCO) with a maximum output power of 1 W and a maximum current of 500 μA for safety. It is coupled with an H-bridge using opto couplers (MPI

Technologies) to generate a square-wave AC signal at a frequency of up to 1 kHz. The HVS was controlled by an Arduino micro controller via a USB connection to a laptop. For a fully portable application, the electronics can be scaled down to a few cm^3. The use of HV in wearable devices is not a problem as long as the current is very small. Hence, we limited the current to 500 μA. In addition, strips and connections can easily be enameled or insulated. In our case, the users' hand was insulated using a nylon glove and Polyimide tape.

We used bipolar square waves at 10 or 20 Hz. AC operation eliminates charge injection in the dielectric layers, a problem we had observed after continuous DC actuation. AC actuation thus allows the electrostatic force to be turned off as soon as the voltage is reduced, even after hours of continuous operation. It comes at the cost of marginally higher power consumption (tens of mW).

4.5 HAPTIC GLOVE INTEGRATION

The ES braking mechanism needs to be integrated into a glove form-factor so as to effectively provide haptic feedback (Challenge Nr. 2). We explored multiple configurations for clutch placement (e.g., finger vs. back of the hand), attachment means to finger tip, wrist, or forearm, investigated different positions to account for the multiple degrees of freedom of the thumb, etc. We report here only the configuration that gave the best results.

4.5.1 *Glove assembly*

We mount the ES brakes on a glove covering the index finger and the thumb via velcro fabric hook and loop fasteners, 3D printed wrist and finger tip anchors (4.5 mm high and 16 mm wide) and 3D printed guides (6 mm high and 14 mm wide), see Figure 4.4. Assembly is straightforward and can account for variations of hand size, geometry (static) and flexibility (dynamic) across users. To deliver effective haptic feedback, the ES Brake must conform as much as possible to the hand and be firmly attached to it.

4.5.1.1 *Finger flexion and abduction*

The fingers consist of 3 phalanges (2 for the thumb), with joints able to bend up to 90 degrees and with radius of curvature of just a few *mm*. This range of motion can cause problems when bending a stack of sliding

Figure 4.4: Assembly of the ES brake, 3D printed guides and attachments fastened to thin nylon glove using velcro.

strips. We therefore designed the “hand strip” to be slightly shorter than the “finger strip”. The “finger strip” is attached to the finger tip and covers the phalanges and the metacarpus while the “hand strip” is attached to the wrist and only covers the carpus and metacarpus. The overlap region covers the metacarpus on the back of the hand, anchoring the force so that it can counteract finger flexion.

Another challenge arises from the metacarpo-phalangeal (MCP) joint which can both flex and abduct. If not counteracted, under abduction the free end of the “finger strip” will laterally slide on the back of the hand while rotating the finger, causing misaligned strips and reduced braking force. It is important to maintain a constant distance between the strips, even under deformation. Finally, a misaligned strip can damage the insulating layer. To avoid this, we polished the strips’ edges and covered the strips’ free ends with insulating tape. We use 3D printed guides to keep the strips aligned on the back of the hand, requiring the “finger strip” to be flexible enough to accomodate finger rotation.

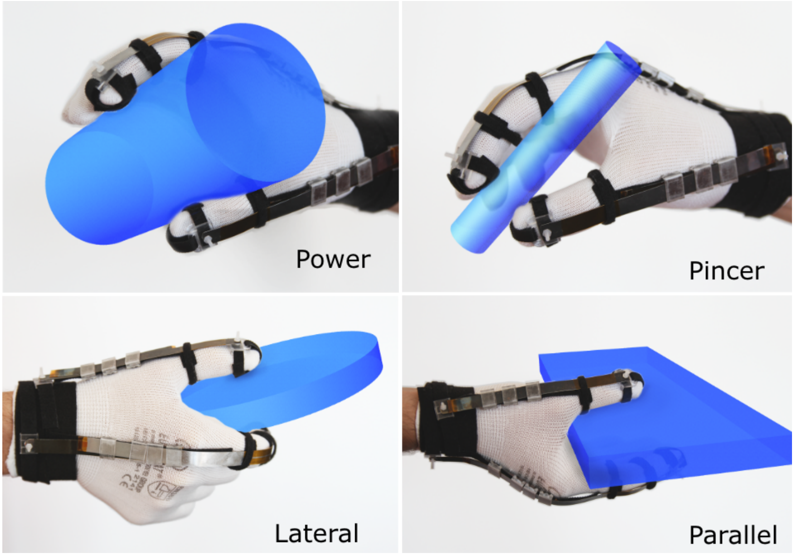

FIGURE 4.5: Supported Grasps: Power, Pincer, Lateral, and Parallel

4.5.1.2 *Thumb*

The thumb is composed of only two phalanges and a flexible metacarpus. Designing an ES Brake for the thumb proved to be more difficult than for other fingers. The ES Brake device is the same for all fingers (only its length changes) but its integration on the glove is different for the thumb. We empirically found that to be effective and to support the Power, Pincer, Lateral and Parallel grasps (Figure 4.5) we had to tilt the thumb strips anchors 30 degree outward. Moreover, the hand strip had to be attached further back on the wrist compared to the index finger (Figure 4.2).

4.5.2 *Glove activation and deactivation*

In the simplest model, the electrostatic force scales as $1/\text{gap}^2$, thus having a small initial distance between strips is critical. We carefully designed our 3D printed guides to ensure that a) the strips are as close together as possible, and b) leave just enough room for smooth gliding to allow fast retraction. To activate the brakes, each ES strip is set to 1 kV at 20 Hz. Once a small region of the strip pairs come into contact, the adhesion propagates in a zipper-like effect. To deactivate the brakes, the difference of electrical potential between the strips is set back to 0 V.

4.6 INTEGRATION INTO VR

4.6.0.1 *Tracking and Haptic Device Control*

Creating a convincing method of grasping objects in VR requires precise tracking of the fingers in order to determine when contact has been made. For tracking, we use an OptiTrack tracking system with 10 Prime 13 W cameras running at 240 Hz and custom designed rigid bodies that screw into the tips of the fingers. The centroids of the rigid bodies are calibrated to sit in the center of the finger such that finger collisions in real life match finger collisions in VR. The mean tracking error after calibration of the whole system was < 1 mm. An Oculus CV1 headset is used to display the virtual scene. The coordinate systems are aligned via a calibration procedure built into the Motive:Tracker software. We use Unity to render the VR scenes. The position of the fingers are displayed as small spheres. Each haptic controller (index, thumb, piezo) has a separate physical connection (USB) and are controlled individually over different serial ports. Fig. 4.6 shows the complete setup.

4.6.0.2 *Grasping Method*

We implement a custom grasping algorithm, similar to Choi et al. [110] using a kinematic approach. A grasp begins when the position of each finger (index, thumb) are within 5 mm of a virtual object and the object to be grasped is between the fingers. Once the object is grasped, the resulting ray between the two fingers is used to kinematically rotate and re-position the object in real time, and to calculate the amount of object penetration for analysis. The grasp ends when the ray between the fingers exceeds its original starting (euclidean) distance. This approach ensures a steady and natural feeling grasp and supports more types of grasps than off-the-shelf solutions such as the Leap Motion Interaction Engine.

4.6.1 *Haptic Rendering*

When a user grasps an object, we activate index and thumb brakes simultaneously. Any slack will initially be perceived as zero blocking force, but as soon as the slack is taken up, it will be perceived as a sudden locking of

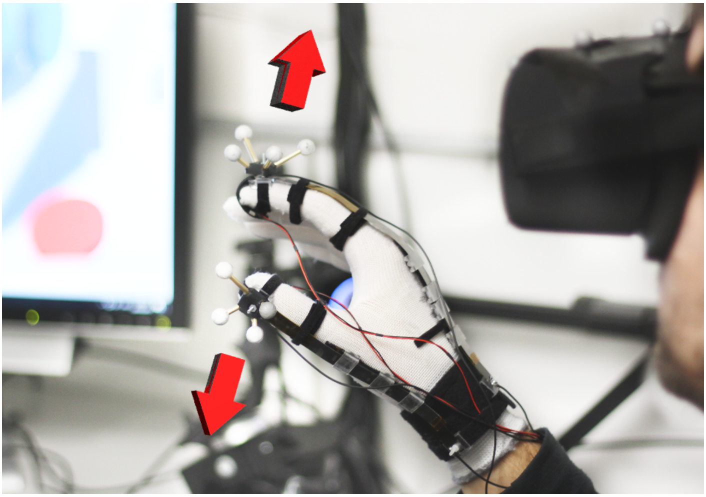

Figure 4.6: VR Integration: The user's index finger and thumb are tracked provided via OptiTrack rigid bodies and feedback is provided when grasps are triggered on virtual objects in Unity.

the finger. Mechanically, the strips counteract torque in the DIP and hence directly brake the downward motion of the fingertip. The perceived effect is that of grasping an object in real life, despite not directly generating a normal force. It is possible that a user could squeeze hard enough to break the adhesion, however, this would exceed normal grip forces during grasps which are shown to be two times the load force [4].

In accordance with perceptual theory on initial contact during object manipulation [22], the feeling of grasping can be improved by adding tactile feedback at the fingertips. For this task, we use tiny vibration motors (PiezoVibe from Murata) which measure 3.8 x 10.5 x 2 mm and vibrate at 240 Hz. They generate an acceleration of 1.2 G for a mass of 20 g and consume 6 mW. When a grasp begins, we briefly activate the piezos for 0.3s to indicate the start of a touch event. They are not re-activated during release.

4.7 SYSTEM EVALUATION

Before reporting on our usability experiments we briefly characterize the ES brakes in terms of blocking force vs. applied voltage and for response speed.

4.7.1 *ES Brake predicted force and speed*

Based on Eq. 4.2 and Eq. 4.3, and considering an ES brake having an overlap of A=11 cm^2, a 13 μm Polyimide insulator film with relative dielectric constant 3.4 and a friction coefficient of 0.2 between kapton and steel [113], a voltage of 1500 V should generate a compression force of 220 N, resulting in a friction force of 44 N. Using our HV supply with a maximum current of 500 μA, it should take 50 ms to fully charge the strips, thus enabling up to 20 Hz operation. At this frequency, the ES brake consumes 57 mW (Eq. 4.4).

4.7.2 *Measurement method*

To measure the braking force of our ES brake, we placed it in a pull tester (Instron 3344L) equipped with a 50 N load cell (Instron 2519). This allows us to pull on the ES brake over 10 mm while measuring the braking force (Figure 4.7a).

4.7.3 *Experimental results*

Figure 4.7b plots the braking force generated by the ES brake vs. time for a pull speed of 1 mm/s for a 10 Hz AC actuation voltages ranging from 0 V to 1500 V (Figure 4.7b). The force starts at zero and quickly increases as any slack is taken up. The force then reaches a plateau corresponding to the braking force. At maximum load, we noted repeated slipping and slip-stick behavior as a result of the AC actuation.

Results are summarized in (Figure 4.7c) for 16 measurements on several devices. Our ES brakes can block up to 20 N at 1500 V and 10 Hz with variations of 10 %. This corresponds to a force of 2 N/cm^2. It is possible to stack ES brakes to achieve higher forces or to reduce the operating voltage at constant force. When switching between 0 V (free) and 1500 V (locked), we observed a response time of less than 100 ms. This correspond to a force slew rate higher than 200 N/s.

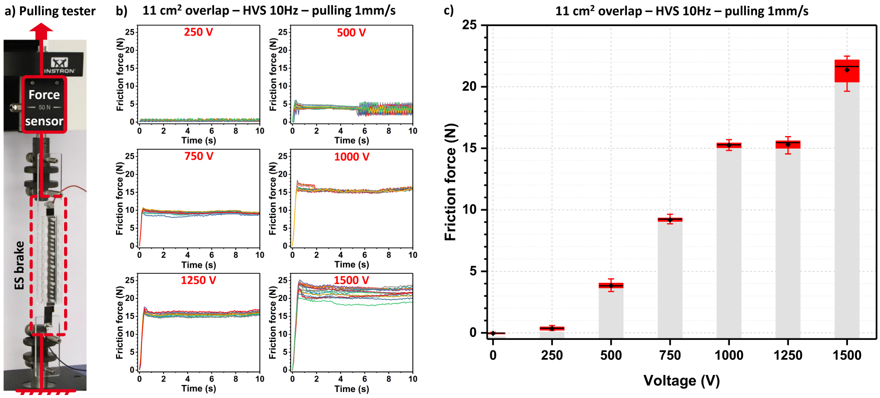

FIGURE 4.7: Friction force measurements. a) Pull tester setup to measure the ES brake's friction force. b) Measured braking force for voltages from 250 V to 1500 V vs time as the strips are pulled apart at 1 mm/s. The initial overlap area was 11 cm^2. The frequency of the AC actuation was 10 Hz.c) Friction force vs. applied voltage. The black dot is the average, the black line is the median, the box correspond to the IRQ and the bars to the min-max. Higher is better.

4.8 USER EVALUATION

To better understand the efficacy of DextrES and its possible applications, we conduct a quantitative and a qualitative experiment. First, a psychophysical evaluation measures the *just noticeable difference* (JND) of stiffness which can be felt on each finger. Second, we explore the grasping precision afforded by DextrES and its effect on the immersion of the user. Each study has been designed to answer the following research questions respectively:

- **RQ1**: What is the *just noticeable difference* of blocking force at different voltages at each finger (index, thumb) and what is their associated perceived stiffness?
- **RQ2**: What effect do the kinesthetic, cutaneous, and combined modes of DextrES have on the *precision* and *immersion* while grasping and manipulating objects in VR?

4.8.1 *Study 1: Force Discrimination*

While grasping devices have been created that can exert up to 100N of force per finger [29], it is not clear that such high forces are actually needed for dexterous manipulation of objects in VR. In this study, we are interested in the perceived stiffness rendered on each finger (thumb, index) and its JND. In order to create the feeling of different levels of perceived stiffnesses, we vary a reference input voltage, and use an adaptive staircase method [77] to determine the JND at each reference voltage. Based on pilot studies, we select three reference voltages (200 V, 400 V, 800 V) and a variable step size with an initial value of ΔV to 7.5 %. Before the study, we measured the force output of the strips at each of the reference voltages and noted this for later analysis (see Fig. 4.7.

4.8.1.1 *Participants*

We recruited six healthy adult unpaid participants (*M*=30.8; *SD*=3.1; 1 female) from the ETHZ university campus. Participants had an average hand span of 21.7 cm (*SD*=1.5) as measured from the end of the pinky finger to the thumb. Each participant signed an informed consent form prior to the study.

4.8.1.2 *Procedure and Task*

The braking mechanism was mounted on the index finger and thumb of each participant. Since we only test one finger at a time, the mounting of the 3D printed guides and their velcro holders can be placed in straight lines extending from the tip of both the index finger and thumb, with the thumb configured in the abducted position. Participants are then given some practice time to get accustomed to the device, after which they put on noise-canceling headphones and a blindfold in order to eliminate interference from external visual and auditory senses. The JND for each finger is then determined using the adaptive staircase procedure [77] described above.

Each trial consists of two runs with randomized presentation order. In one run, the fixed reference voltage is activated, and in the other run, the approaching voltage is activated. On each run, the participant flexes their finger inwards until they sense the blocking force. The participant is unaware of which voltage is used. The initial value for the approaching voltage is +25 % of the reference value. In the case of 800 V, it was set at -25 % of the reference voltage, however positive and negative JND approaches are typically symmetric [76]. After each trial, we ask participants to identify which of the two voltages was perceived to be blocking their finger more. A correct response brings the voltage in the next trial a step size towards the reference voltage, and vice versa [77]. The step size was halved to 3.75 % after the first direction reversal in order to get more accuracy after the initial approach. The procedure is repeated until the direction is reversed 3 times and the reversal points are averaged to get the JND of each starting condition. At the end of the procedure, participants answered how stiff they perceived the blocking force to be.

Location	JND			Measured Force			Perceived Stiffness		
	200V	400V	800V	200V	400V	800V	200V	400V	800V
Index Finger	20.6 %	5.6 %	7.8 %	0.10 N	1.34 N	6.04 N	1.2	2.3	3.8
Thumb	15.2 %	4.6 %	10.1 %	0.16 N	1.05 N	3.91 N	1.5	1.8	4.2

TABLE 4.2: JND Study results at three different reference voltages. Blocking force equivalent was pre-measured for each set of strip combinations. Perceived stiffness is rated on a 5-points Likert scale where 1 is easily deformable and 5 is a rigid object that cannot be deformed.

4.8.1.3 *Results*

Table 4.2 summarizes the JND for each finger in different positions and reference voltages. The smallest JND for both fingers is in the middle range (400 V), where participants could adequately sense about 5% differences in blocking force.

Based on the perceived stiffness at this voltage, it is possible for DextrES to render objects with different levels of deformable stiffness. At the low end (200 V), the JND rises significantly as we approach a perceptibility threshold. At the high end (800 V), participants were considerably more perceptible than expected and perceived stiffness is still not close to maximum, meaning there is still some room to render objects with very hard but still deformable stiffnesses if the voltage is increased even further. Comparing between the middle and high reference, the results are non-linear, which suggests that there is an upper bound on rendering an object of maximum stiffness.

4.8.2 *Study 2: Grasping Precision and Realism in VR*

To answer RQ2, we conduct a VR study measuring both quantitative aspects of precision during object manipulation, and qualitative aspects of realism during the grasping of objects. We use the same definition of grasp as Feix et al. [81] where grasping stipulates that objects are held firmly in the hand (rigid) and not rotated by moving the fingers (static).

4.8.2.1 *Participants*

Ten healthy adult subjects (*M*=27.6; *SD*=4.14; 2 female) were recruited for our study. Two participants had no previous experience with VR and 1 participant was left-handed. Each participant signed an informed consent form prior to the study.

4.8.2.2 *Procedure*

The procedure is described to the participant alongside a brief introduction to the device and its function. The ES brake strips are then mounted to the back of the hand and adjusted as described earlier. Participants could sit in a 1x1 meter tracking area wearing an Oculus headset. Participants wore noise-canceling headphones to remove external audio cues when grasping objects. The experiment consisted of two scenarios, the first evaluating the quantitative aspect measuring grasping precision, the second, qualitative aspects comparing the realism of grasping objects between different haptic feedback conditions. The experiment took 1.5 hours to complete.

Participants could practice freely in order to learn four different grasps (*Lateral, Parallel, Pincer, Power* see 4.5) and to get used to the different types of feedback. After training, 4 blocks were completed of 16 unique grasp/condition combinations for 64 trials in total. The quantitative part was followed by approximately 1 min of each condition in a physics playground which was designed to resemble a typical desk with various sized items which could be grasped and interacted with. The experiment concluded by participants indicating their subjective preferences and a short interview on the overall impression of the experience in using DextrES. We also recorded suggestions for possible applications of our device and informal comments.

4.8.2.3 *Design*

For the quantitative study, a within-subject design was used with two independent variables: *Grasp {Lateral, Parallel, Pincer, Power}* and *Feedback*

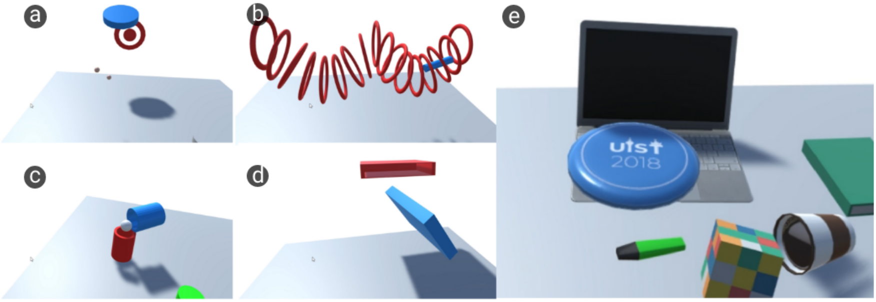

Figure 4.8: Experiment tasks. Manipulated objects in red, target in blue: a) hitting target with frisbee (*lateral* grasp). b) Steering task (*pincer*). c) Ball-in-cup (*power*). d) Bookshelf (*parallel*). e) Physics playground.

{Visual, Piezo, Brake, Both}. The order of each stimuli was randomized such that each combination of *Grasp* and *Feedback* was presented once per block. As dependent variables, we measured *time* and *precision* for each trial. Precision was measured by the percentage of finger-object penetration averaged over the whole trial.

At the end of the physics playground scenario, we asked participants to rate how realistic the sensation of holding an object is in each feedback condition on a 7-point Likert-scale (1: Extremely unrealistic, 7: Extremely realistic). We also ask about the comfort of the device while turned off (1: Very uncomfortable, 7: Very comfortable) and the freedom of movement (1: Fully blocking, 7: Full range of motion).

4.8.2.4 *Task, Stimuli and Apparatus*

Trials were initialized and terminated via pressing a virtual button, or after a 20 second timeout. Each *Grasp* has an associated task (see Fig. 4.8 for visuals and explanation) which is derived from real-life tasks. Participants were instructed to complete the task in timely manner as accurately as possible. The instructions given in regards to grasping an object were to perform it as naturally as possible. After each block participants could take a break before proceeding.

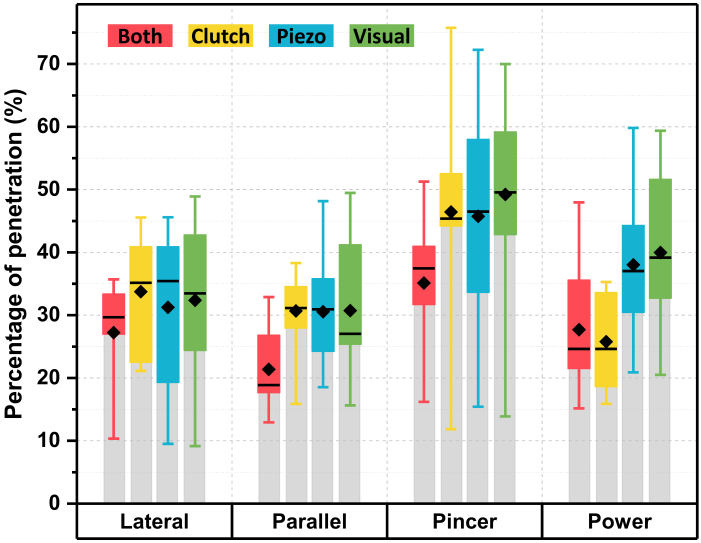

FIGURE 4.9: Effect of feedback mechanism on precision of 4 grasp types. Conditions are *both* (red), *brake* (blue), *piezo* (yellow) and *visual* (green). The black dot is the average, the black line is the median, the box correspond to the IRQ and the bars to the min-max. Lower is better.

4.8.3 *Quantitative Results*

To assess the effect of the different feedback mechanisms on grasp precision we ran a repeated measures ANOVA for each grasp. There was no difference in terms of task completion times. Since the four grasps are significantly different it does not make sense to compare feedback mechanisms across the four grasps. We now report main effects and pairwise post-hoc comparisons for all four grasps (*lateral*, *parallel*, *pincer*, *power*) respectively. The sphericity assumption was not violated for any of the grasps. All p-values of pairwise comparisons are Bonferroni corrected.

For the *lateral* condition (e.g., turning a key) a main effect for the feedback mechanism ($F_{3,27} = 5.17, p < 0.01$) was detected. A post-hoc analysis reveals that *both* is significantly more precise than *brake* ($p = 0.02$) but differences to other conditions are not statistically significant.

The *parallel* grasp (e.g., lifting a book) yields similar results. The analysis again shows a main effect ($F_{3,27} = 4.86, p < 0.01$) and post-hoc analysis reveals that there is a significant difference between *both* and *brake* ($p = 0.04$), albeit inspecting the plot in (Figure 4.9) shows that the differences are very small and this result should be interpreted carefully.

The remaining two grasps show a more marked effect of the feedback mechanism on the precision of the grasp, perhaps because both *pincer* and *power* admit much more finger motion (cf. Discussion section). There is a main effect for *pincer* ($F_{3,27} = 12.24, p < 0.01$) and pairwise comparisons indicate that *both* is significantly more precise than *brake* ($p = 0.01$), *piezo* ($p = 0.02$), and *visual* ($p < 0.01$). Finally, the *power* grasp also yields a main effect ($F_{3,27} = 21.32, p < 0.01$). The pairwise comparisons indicate that for this grasp *brake* is the most precise. However, compared with *both* ($p > 0.05$) the difference is not statistically significant. Both (*brake, both*) feedback mechanisms are however statistically more significant than the (*piezo, visual*) baselines (*both* vs *piezo*: ($p = 0.01$), *both* vs *visual*: ($p < 0.01$)).

4.8.4 *Qualitative Results*

4.8.4.1 *Subjective Rankings*

The physics playground gave participants a chance to interact freely with virtual objects. In terms of the realism (see Fig. 4.10) of the sensation of holding an object participants consistently rated the *both* feedback condition the highest (M=5.3; SD=0.5), followed by the *brake* (M=4.4; SD=0.8) and *piezo* (M=3.5; SD=1.1), and finally the *visual* only feedback (M=2.2; SD=0.8). Participants rated the device as mounted on the hand as fairly comfortable (M=5.1; SD=1.4), and the freedom of movement as neutral in terms of limiting finger motion (M=4.6; SD=0.8).

4.8.4.2 *Participant Comments*

Participants strongly favored the combined feedback, which produced the highest sense of VR immersion (e.g., *"Haptic is missing from VR, and I'm really impressed that it can block."* and *"It's pretty cool. It adds a lot of immersion."*), and also helped to identify when grasping an object begins *"It helped me understand when I should stop applying force"*. The brake-only condition was slightly less preferred due to missing collision information *"When you go and touch something you expect your skin to be bump into it"*. When the brake was missing, participants could sense its absence *"As soon as you go to the next trial and it's off, then you miss the feedback"*. Two participants preferred the piezo over the brake on its own, however, it was still rated as less realistic as just the brake (e.g. *"Its very useful to know when you touch it, but its not* *a realistic feeling"*. Furthermore, the brake adds realism in the context of limiting range of motion (e.g. *"If you hold a bigger object, then if you rotate you should have more limited range of motion"*) (because the forearm muscles engage). Some participants found DextrES to provide physical support, specifically in the lateral grasp *"When I got tired, I was using the brake to rest my thumb"*). The main issue w.r.t. to comfort was the feeling of velcro on the hand, but on the whole, participants felt that DextrES was comfortable to wear. In terms of applications, participants wanted to use the device to play games, for virtual typing, and also for creative tools such as CAD tools and painting (e.g. *"You can 3D paint, but maybe different sized brushes have different weight"*).

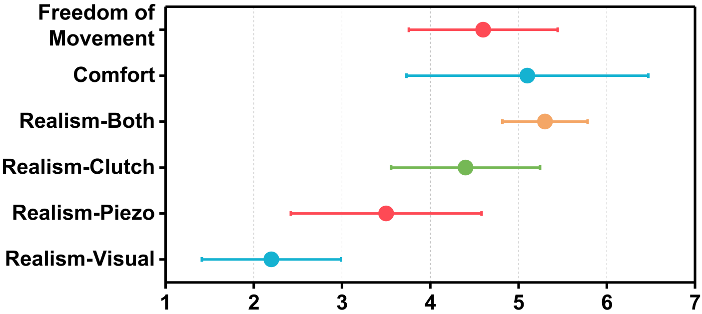

FIGURE 4.10: Subjective feedback on 7-point Likert scale. Dot is the mean and bars indicate confidence interval. Higher is better.

4.9 DISCUSSION

A major contribution of our work is to **fabricate an ES brake** for VR. Through careful material selection, specifically the use of a conductive adhesive to minimize effective dielectric thickness and a dielectric with high breakdown fields, in combination with a flexible mechanical design that allows for dimension tailoring and achieves reliable sliding over curled fingers, and by using AC switching of the 1.5 kV power supply to avoid charge injection, we were able to develop an ES brake with suitable force generation capabilities while allowing natural hand motion. Including mounting hardware, DextrES weighs under 8 g, yet can block 20 N when on, with only a few mN of residual force when off. We thus have a blocking force density over 2500 N/kg, while delivering a device so flexible it is barely perceivable on the hand. The materials are all readily commercially available, and can be machined in minutes with a laser cutter.

The ES brake was **integrated into a glove** using close-fitting 3D printed attachments on the fingertips and wrist, by precise positioning of 3D printed guides on the glove, by aligning the index brake with the index's long extensor tendon, and by placing most of the active part of the brake on the back of the hand where there is much less deformation. The thumb was most challenging, principally due to its more complex motion (not only can it flex like the fingers, but it can also pivot on 2 axis), making it harder to mount the ES brake in a way that effectively blocks flexion. There is also little room on the back of the hand to mount the brake and a medium curvature close to the wrist attachment. We found that aligning the thumb ES Brake 30 degrees outward of the thumb flexion axis gave the most blocking. While more direct forms of tactile feedback are available [8, 60], integrating such devices in their current form may interfere with the mounting of the strips and the natural motion of the hand.

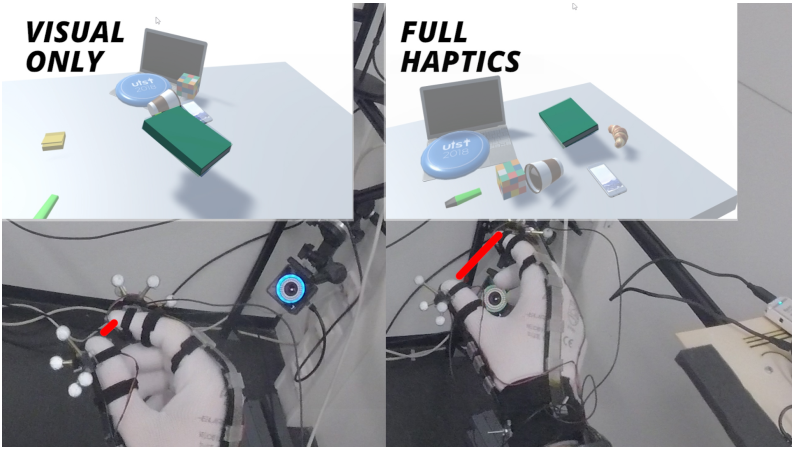

FIGURE 4.11: Examples of differences in grasps across conditions. Left: without haptic feedback, participants penetrate virtual objects. Right: with haptic feedback, fingers conform to the object's shape (green book).

The final challenge of **VR integration** was also met. The results from our VR grasping study indicate that DextrES is able to support three of the four grasps (Parallel, Pincer, Power), in particular when both the brake and piezo work in tandem. Participants were able to both pick up and drop objects in a natural fashion and experience a sensation of holding an object.

The main differentiating factor between grasps is the distance between the tip of the index finger and the thumb. In the Power grasp, this distance is wide, and thus gives some space for the brake to engage and to exhaust any mechanical slack. As a result, the *brake* and *both* conditions perform similarly. In the Pincer and Parallel grasps, this distance is small, thus they greatly benefit from the additional collision signal from the Piezo. With regards to the Lateral grasp, this distance is also small. Furthermore, the inward flexion of the thumb does not necessarily induce any sliding motion on the strips of the brake, and thus neither the piezo or the brake have much stopping effect. While the brake cannot constrain certain degrees of freedom of the thumb in the Lateral grasp, participants enjoyed that they could rest their thumb after the brake has engaged.

The above results show that DextrES is able to increase precision during specific VR manipulations. In same cases and for some participants the differences were small in the controlled experiment. However, in the more natural setting of the physics playground, participants exhibited very different behavior. They were less careful when picking up and handling objects and thus tended to penetrate through them completely in the visual and piezo conditions (see Fig. 4.11, left), whereas with haptic feedback they conformed to the object's shape (see Fig. 4.11, right). In part, this explains the large differences in the perceived realism in holding an object compared to the smaller differences we see when looking at the percentage of penetration in the controlled experiment. While we do not directly test grasping objects of variable stiffness in VR, results from the force discrimination study suggests that this is a possibility.

4.10 SUMMARY AND OUTLOOK

In this chapter, we presented a novel haptic glove integrating electrostatic braking using flexible components. With its low mass (under 8 g) and high force (over 20 N) it overcomes limitations of more traditional motors and pumps. Our experimental results indicate that DextrES is a very promising step towards soft, flexible, high-speed wearable haptics conveying the sense of grasping with high realism. We tested the device for 4 grasps and found improved grasping precision for different virtual objects. By including small piezoactuators at the fingertips, we further increased the grasping precision.

Naturally, there is much room for future work. The operating voltage could be reduced significantly by printing thinner dielectric layers or layers with higher permittivity. Lower voltage operation will: i) make the control electronics more compact and much cheaper since all components can be sourced in surface mount format (SMD), ii) reassure users who may be concerned about high voltages, iii) ease regulatory processes for wearables. Users currently are aware of the 20Hz switching, which can be distracting. Lower voltage operation would allow the device to be run with a sine wave rather than a square wave, greatly reducing the audible vibration. Further, the force generation capabilities may be increased by stacking several ES brakes.

Prior work in the space of on-hand kinesthetic feedback employs rigid components either on or near the hand. The development of a compliant mechanism that can clutch on-demand paves the way for radically different designs, allowing for integration with textiles as shown in this chapter, and instrumenting the hand *directly*, rather than offloading heavy pumps or motors elsewhere on the body. This addresses an important challenge in VR haptics, where there is demand for richer feedback without excessive user instrumentation.

While the work presented in this chapter targets on-hand kinesthetic feedback in a compliant package, it remains unclear how to scale this approach up to the whole body. In particular, the grasps under which the device failed to increase user precision were those with the least amount of movement in the hand resulting in a lack of *sliding* in strips, thereby producing no feedback. An important deduction is to be made here, that no matter how efficient or energy dense a clutch-type device is, it is fundamentally limited by it's design and integration onto the body. The design must be such that the desired motions induce sliding in the device, and also efficiently route those forces onto the body (See Fig. 2.2). Scaling up such a device to the whole body and with the ability to resist even more motions, would only multiply the design effort required. In our follow-up work, we redirect our focus to addressing this design problem, applying a computational approach based on a physical simulation and reducing the design effort required through inverse design methods.

5

COMPUTATIONAL DESIGN OF KINESTHETIC GARMENTS

This chapters proposed the idea of using computational methods for the design of wearable haptics. Beyond designing wearable haptics for the hand, designing for larger parts of the body introduces difficult design challenges: more degrees of freedom to account, more complex motions, and more active components to place and connect. As this complexity increases so does the effort on the part of the designer (See Fig. 5.1). In order to reduce this effort, we propose the idea of a computational pipeline that leverages physical modeling and simulation. Our pipeline is able to automatically create efficient designs for *kinesthetic garments*, garments that are able to passively resist a particular motion. Although such garments can only resist motion passively, the aim of this chapter is to lay the foundation for later work by first introducing a suitable garment-on-body model and the idea of using topology optimization to address the inverse-design problem.

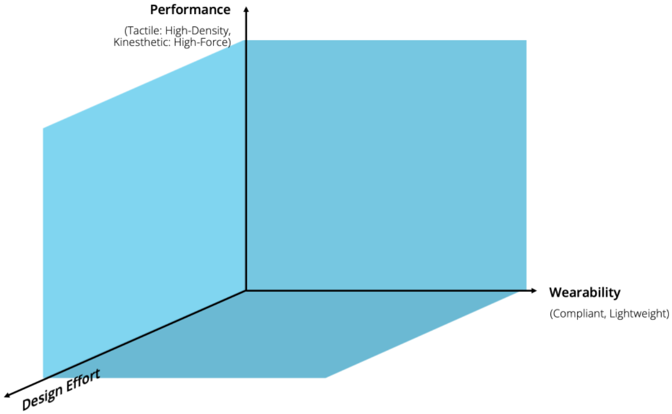

FIGURE 5.1: In addition to Performance and Wearability factors, Design Effort is a key variable when scaling up haptic-interfaces to the whole body.

5.1 INTRODUCTION

Kinesthetic feedback during body motion has wide-ranging applications in posture-correction, locomotion assistance [9, 10], and enhanced immersion in mixed reality [6, 7]. In many recent works, *wearable compliant interfaces* have emerged as the preferred means of transmitting forces to the human body, owing to their lightweight and conforming properties. An important consideration in these interfaces is to provide feedback or assistance for specific motions, while not overly instrumenting the user. However, unlike their stiff and rigid counterparts, wearable compliant interfaces behave in a non-linear and difficult to predict manner, creating a challenging problem for designers.

The goal of this work is to facilitate the design of *Kinesthetic garments*—lightweight and compliant apparel that, when deformed during body motion, deliver kinesthetic feedback to the user via forces felt in the muscles. This is accomplished by reinforcing the garment with stiffer material in order to resist specific motions. In this task, designers must balance the conflicting goals of resisting specific motions and retaining as much of the garment's flexibility as possible. We propose to combine these objectives into the task of *maximizing design efficiency*: only the minimal amount of reinforcement material required to achieve a desired stiffening effect should be used. Conversely, a given budget of material should be distributed such as to maximally resist the specified motion.

The challenge of designing garments that consider body motion is an emerging topic in computer graphics [14, 114] and wearable robotics [13, 115]. A number of technical challenges arise: both the cloth and the body behave in a complex non-linear manner, and their coupling under tight contact is difficult to model. In addition, the human body itself deforms significantly under kinematic motion (i.e. muscle bulging). Optimizing for particular material budgets in an on-body and compliant setting has yet to be considered. Thus, the task of designing garments that provide kinesthetic feedback under desired motions remains difficult and time-consuming for non-expert designers.

Our idea is to cast this design task as an *on-body topology optimization* problem in which we seek to compute optimal layouts for garment reinforcements. Based on this idea, we propose an automatic design algorithm that, given a material budget and specific motion will maximize the energy needed to deform the garment, and thus the mechanical work users generate during the motion and the kinesthetic feedback they experience. As the desired material budget may not be known a-priori, our method supports designers by generating a range of distinct designs, enabling the exploration of the trade-off between material budget and resulting performance.

As with most topology optimization methods, our approach relies on finite element analysis during optimization. For this purpose, we create a customized garment-on-body model of skin-tight cloth that is elastically stretched over a body. This model allows cloth to smoothly slide on the body, resulting in garments that fit tightly to convex parts of the body. It also allows for lift-off separation between cloth and body, and exhibits close-to-zero stiffness under compression to emulate wrinkling.

We leverage our model and automatic design method to produce a number of kinesthetic garment designs for different body sites and motions. Optimal designs are 2 to 3 times more efficient (in terms of energy density) than fully reinforced designs, and are characterized by complex branching structures that route around and across the body to exploit pose-induced body deformations. Physical validation tests in 2D show our simulation model is in agreement with experimental results. We fabricated physical prototypes for the arm and the knee using a simple but easy to deploy heat-transfer vinyl (HTV) process. We found that users consistently rated our optimal designs as more resistive when compared to baselines. Additionally, we explored the performance of our method in more complex use cases, such as back posture support.

5.2 RELATED WORK

DESIGN OF GARMENTS Designing garments with specific properties has been extensively investigated in graphics literature. Typical approaches consist of adjusting 2D pattern designs with regards to optimization criteria [116], or by directly controlling textile properties through knit and stitch patterns [114]. Physics-based pattern generation has been employed in multiple works, including for friction and pressure distribution [14, 117] in skin-tight clothing, friction minimization [14], and for automatic pattern generation from high-level user input [118]. Wang et al. develop a method for creating woven 2d patterns of elastic braces based on tunable springs that induce the desired normal pressure in the garment's worn state [119]. Aesthetic criteria are considered by Kwok et al. who use an evolutionary strategy to generate a diverse set of garment designs [120]. An emergent topic is designing garments that consider motion [114]. These so called 4D-garments can be designed to minimize pressure from hard materials and

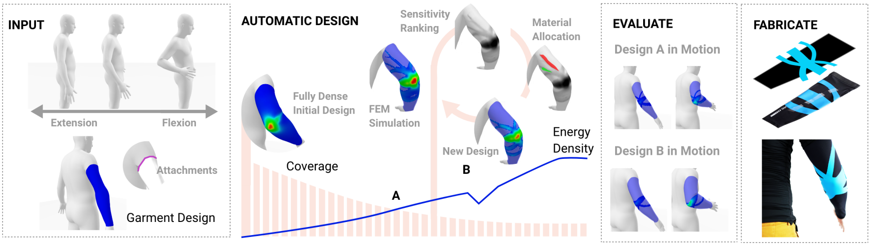

FIGURE 5.2: System overview description using the example of the arm flexion as an input motion, and the arm sleeve region as the garment to be reinforced. From left to right. *Input:* A designer imports a sequence of body poses and labels the rest pose (zero energy) and final pose (optimization objective). The garments are defined on the surface of the body. *Design Exploration:* The on-body topology optimization produces a sequence of optimal designs as a function of the reinforced area (budget). The background bars represent the area coverage of reinforcements. *Evaluate:* Designer picks between two candidates and evaluates them in simulation over the full motion sequence. *Fabricate:* the 3D on-body surface design of the reinforce regions is then flatted, cut and heat-transferred into the soft sleeve to produce the kinesthetic garment.

sliding from soft-material by creating a multi-material integrated knitting map. In our work, we extend the notion of considering motion in computational garment design by providing kinesthetic feedback in response to particular motions. Similar to Montes et al. [14], our work employs a physically based model, however, allowing for additional degrees of freedom in terms of cloth lift-off and leveraging a structural optimization approach instead of 2d pattern optimization.

STRUCTURAL OPTIMIZATION The question of how to best distribute a given amount of material such as to obtain optimal performance is a central problem in engineering [36]. Topology Optimization (TO) has been employed in graphics to explore the intersection of structural objectives with aesthetic guidance by the user [121, 122]. Recent work from Liu et al. [123] pushed the limits in terms of grid resolution, demonstrating that, when given sufficient resolution, TO is able to recover structures similar in complexity to those found in Nature. In an elastic material setting, Skouras et al. leverage a material interpolation and penalization approach (SIMP) to optimize the material distribution of actuated characters in order to achieve a target deformation behaviour [124]. Bruns et al. apply the SIMP method for elastic structures undergoing large displacements [125], while Huang et al. demonstrated that gradient-free, evolutionary strategies such as BESO can also handle such cases [126]. Closer to our setting, structural optimization is a common approach used in cast design. Zhang and Kwok [127] adapt SIMP into a two-manifold surface and enable efficient personalized cast designs. Personalized casts have also been optimized for thermal comfort [128] using an FEM in the loop strategy to selectively thicken material to increase structural stability. Topology optimization on manifolds can also be achieved by establishing a conformal map from 2D to 3D surfaces and evolving the boundary using the level-set method [129]. While previous work has focused on topology optimization in 2D or 3D settings, we explore an on-body TO approach for automated synthesis of reinforcement structures embedded in complex-shaped 3D surfaces.

AUGMENTING TEXTILES Kinesthetic garments belong to a wider class of textile-based compliant interfaces that are augmented with sensing and actuating systems [13]. Passive and quasi-passive systems have been developed to support locomotion [130] and provide lifting support [131] using shape memory alloys. Ortiz et al. [115] take an optimization approach to optimally combine and place passive elastic cords, clutches, and dampers to reduce the force and power required by a person to generate lower body motion. Fully active systems embedded onto soft garments have been designed for flexion/extension of the elbow joint [132], grasping [17, 132], and for locomotion assistance [9, 10]. In these works, load paths and anchoring body sites need to be carefully considered for effective transmission of forces. Sensor placement on the lower body was investigated by Gholami et al. [133], who use an evolutionary method to find optimal placement of fabric sensors to recover lower body kinematics. Whole-body sensing and vibrotactile actuation was used in providing feedback on back posture information [134]. While kinesthetic garments are purely passive, they still allow for an information feedback loop that is specified at design time. They also similarly depend on the same efficient routing and anchoring to effectively transmit their forces.

5.3 OVERVIEW

Our computational pipeline supports designers throughout the complex task of designing kinesthetic garments. See Fig. 6.1 for a visual summary.

INPUT AND PARAMETRIZATION As high level input to our pipeline, designers specify the motion they wish to provide feedback to as a single rest and deformed state of the body.

In our tool, motions are specified based on the STAR/SMPL parametric human body model [44, 48]. We use STAR for its large space of body shapes and its ability to represent the effects of pose-induced deformations such as muscle bulging. Realistic motions can be sampled from the AMASS dataset [135] or, alternatively, using existing methods that can recover user motion and personalized shape from RGBD sensors [136].

Next, designers select the dimensions and placement of the soft garment onto which reinforcements will be placed (black sleeve in Fig. 6.1). We parametrize the rest configuration of garments as 2D surfaces embedded on the body mesh, while during simulation, the garment is free to lift-off from the surface. Reinforcements are simulated on the garment mesh as a per-element material property: each element of garment material is assigned one of two materials, i.e., *Cloth* or *Reinforced Cloth*.

REINFORCEMENT OPTIMIZATION We first consider the performance criteria a designer might employ, and therefore, the choice of objective function of our method. Since the overall goal is to resist a specific pose or motion, the work that the user puts in to reach that state should be maximized. In other words, the more energy stored in the reinforcements, the more energy the user has to spend getting into the specified pose. Thus, the overall energy of the kinesthetic garment in the deformed pose is a useful metric for determining how well it would resist the given motion. Comfort and weight is another factor in the design. In general, more coverage by reinforcement material can make the garment stiffer, possibly heavier, and more difficult to put on. Thus, reducing the effective coverage of reinforcements is also a useful objective for a designer. Our automatic design method combines these two objectives, by maximizing the energy density of the garment reinforcements, while progressively reducing the amount of garment covered by reinforcements.

DESIGN SELECTION AND EVALUATION The output of our optimization stage is a continuous sequence of garment reinforcement designs, from none to full coverage, that designers can explore and analyze. Each design is the result of an FEM simulation that is based on a tailor built garment-on-body model, constrained to a particular material budget. In the evaluation stage, designs can be compared by simulating the mounted garment over the full motion sequence (that includes in-between frames). This allows designers to understand how the garment deforms, how the energy storage changes as a function of time, and to make an informed selection.

EXPORT AND FABRICATION In the last step of the process, one final design is exported and fabricated using conventional multi-layer fabric process. In this paper, we demonstrate it via the a standard heat-transfer vinyl (HTV) application process, adding as a reinforcement material a stiffer and stretchable thin film onto the full stretchable garments (see blue lines on top of the black sleeve in output garment of Fig 6.1). While developing the meshes (i.e. unrolling onto a flat surface) is currently a manual task, it could be automated in a straightforward manner.

5.4 GARMENT-ON-BODY MODEL

In order to evaluate the performance of a given design for a specific motion, we must compute the deformations induced in the reinforced garment by given body poses. To this end, we build a computational model for kinesthetic garments that supports on-body stretching and sliding, softens under compression, and allows for lift-off separation. This model allows us to evaluate the performance of a given design in terms of energy density, and this metric can then be exploited by our automatic design method. In the following section, we describe each part of the model.

5.4.1 *Body*

As a key requirement, our computational model must ensure that the garments do not penetrate into the body and that they are able to smoothly slide over its surface. While the STAR/SMPL parametric model is able to produce realistic body shapes, the corresponding piece-wise linear surface meshes induce gradient discontinuities for sliding motion that are problematic for continuous optimization methods. For this reason, we convert the discrete surface meshes from STAR/SMPL into a smooth representation based on implicit moving least squares (IMLS) [137]. The resulting smooth surfaces are implicitly defined as the zero level set of a signed distance field

$$\Phi(x) = \frac{\sum_k \mathbf{n}_k \cdot (\mathbf{x} - \mathbf{v}_k)\phi_k(x))}{\sum_k \phi_k(x)}\,, \tag{5.1}$$

where $\mathbf{v}_k$ and $\mathbf{n}_k$ are vertex positions and normals of the body, respectively, $\mathbf{x}$ is a garment vertex, and ϕ is a locally supported kernel function,

$$\phi_k(x) = \left(1 - \frac{\|\mathbf{x} - \mathbf{v}_k\|^2}{h^2}\right)^4 , \tag{5.2}$$

that vanishes beyond the cut-off distance h. In practice, h is pre-computed as a function of the local neighbourhood vertex distance defined by each $\mathbf{v}_k$—we use twice the average edge length of its one-ring. Based on this distance field, we formulate a penalty energy that attracts garment vertices inside the body to its surface, where $\mathbf{x}_i$ are the garment vertices:

$$E_{\text{body}}(\Phi) = \begin{cases} \sum_i \Phi(x_i)^2 & \Phi(x_i) \leq 0 \\ 0 & \text{otherwise} . \end{cases} \tag{5.3}$$

Note that the above expression is a uni-lateral penalty function, i.e., it prevents the garment from penetrating into the body but allows for lift-off separation as required. An example where this uni-lateral formulation is essential can be in Fig 5.3, where taut material over the inner elbow lifts off under tension during arm extension.

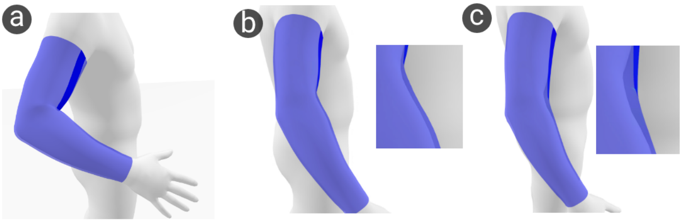

FIGURE 5.3: Starting from a neutral position (a) and moving to an extended position, we compare the behaviour between bi-lateral (b) and our uni-lateral penalty contact model (c) that prevents intersections between cloth and body while allowing lift-off. The total energy of reinforced material in (b) is 30% higher than in (a).

5.4.2 *Garment and Reinforcements*

We model reinforced garments as bi-material distributions where each triangle element e is assigned a specific material property. The Young's modulus and Poisson ratio of the garment material are set to either *cloth* or *reinforced cloth* through the design variable $d^e \in [0,1]$, resulting in a per element strain energy density.

The garments are discretized as a triangle mesh with sets of nodal positions $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_m) \in \mathbb{R}^{3n}$ and $\bar{\mathbf{x}} = (\bar{\mathbf{x}}_1, \ldots, \bar{\mathbf{x}}_m) \in \mathbb{R}^{3n}$ for deformed and undeformed poses, respectively.

To quantify the deformation for a given triangle, we compute the deformation gradient $\mathbf{F}_C$ that maps edge vectors $\bar{\mathbf{e}}_{i,j}$ from the undeformed triangle to their deformed counterparts $\mathbf{e}_{i,j}$, i.e.,

$$\mathbf{e}_{i,j} = \mathbf{F}_C \bar{\mathbf{e}}_{i,j} \ . \tag{5.4}$$

We note that computing $\mathbf{F}_C$ in this way is equivalent to a standard finite element discretization with linear triangle elements.

Based on the per-triangle deformation gradient $\mathbf{F}_C$, we compute the corresponding right Cauchy-Green tensor $\mathbf{C}_C$,

$$\mathbf{F}_C = \frac{\partial \mathbf{x}}{\partial \bar{\mathbf{x}}} \quad \text{and} \quad \mathbf{C}_C = \mathbf{F}_C^T \mathbf{F}_C \ . \tag{5.5}$$

Since $x \in \mathbb{R}^3$ and $\bar{x} \in \mathbb{R}^2$, the deformation gradient $\mathbf{F}_C$ is a 3×2-matrix and $\mathbf{C}_C$ is a 2×2-matrix describing the deformation of the element with respect to rest state coordinates.

Based on $\mathbf{C}$, we define an elastic energy density based on the classical compressible neo-Hookean model [32]:

$$W^e_{\text{garment}}(\mathbf{C}, d^e) = \frac{\mu}{2}(I_C - 2) - \mu \log J + \frac{\lambda}{2}(\log J)^2 \ , \tag{5.6}$$

where $I_C = \operatorname{tr}(\mathbf{C})$ is the first invariant of the right Cauchy-Green tensor that captures all deformations, and $J = \det(\mathbf{F}_C)$ is the relative area change. The material parameters λ and μ correspond to Lame's first parameter and shear modulus, which can be converted to Young's modulus and Poisson's ratio.

In a setting of fully constraining the CST elements to a manifold (i.e. mesh surface), this energy would result in garments that resist compression, a property that cloth garments do not exhibit. Prior work [14, 138] deals with this problem by defining a relaxed strain energy [139] for compressive modes of the respective materials. Our formulation allows garments to lift off the surface of the body (see 5.1) and thus produce wrinkles in areas that experience compressive forces. However, without sufficiently high mesh resolution and proper handling of self contacts, these wrinkles tend to create spurious deformation patterns (see 5.4) that can mislead the optimization method. We therefore adopt a mixed approach that does not penalize compression but allows cloth to lift off the body surface, e.g., in response to stretching over concave regions (see Fig. 5.7).

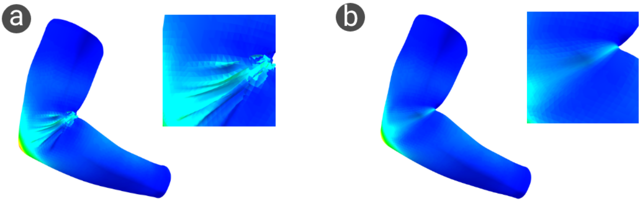

FIGURE 5.4: Wrinkles that emerge in response to compression can generate spurious deformation patterns and geometric artefacts (a). With a relaxed energy formulation that does not penalize compression (b), deformations are smooth and free from compression artefacts.

Since the energy density is constant over each triangle element as per its definition, we can integrate this quantity into a final per element energy and sum over these to arrive at the final garment energy:

$$E_{\text{garment}}(C,d) = \sum_e t^e A^e W^e_{\text{garment}}(C,d^e) \;, \tag{5.7}$$

where A^e is the per-triangle area in the undeformed configuration and t^e is the per element thickness of the fabric.

5.4.3 *Attachments and Contact*

Kinesthetic garments typically have straps or elastics at their boundaries that prevent excessive sliding. Therefore, we specify Dirichlet conditions at the attachment points that we enforce through soft constraints. We allow the designer to specify areas of the garment that would be attached to the body, for example, via belts, velcro loops, or using sticky materials such as silicone. Typically, these might be placed at the boundaries of the garment.

Similar to d^e, for each triangle we specify whether it is attached or not. To avoid creating overly stiff boundary elements, we use the triangle center as the origin rather than its vertices. We formulate this coupling potential as a zero-length spring:

$$E_{\text{attachments}} = \frac{1}{2}k(x^c - x^c_0)^T(x^c - x^c_0) \;, \tag{5.8}$$

where k is the spring stiffness, and x^c and x_0^c denote the deformed and undeformed positions corresponding to the centers of attached elements. We choose a relatively low value of $k = 0.002$ (compared to the base fabric) to discourage designs that do rely on 'pulling on the boundary'.

Attachments ensure contact between the garment and the body. Contact between body and garment may create an additional stimulus, but it is not required for leveraging strain energy for feedback. While our formulation allows lift-off, some part of the garment is always in contact with the body, either through a loop structure or an attachment point.

5.4.4 *Simulation*

With the model and energies defined above, we perform a quasi-static simulation by solving the following unconstrained optimization problem:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \quad E_{\text{body}}(v, \mathbf{x}) + E_{\text{garment}}(\mathbf{x}, d) + E_{\text{attachments}}(\mathbf{x}) \ . \tag{5.9}$$

To ensure that $\mathbf{x}^*$ is in a state of equilibrium, the minimization must ensure that derivative of the above system vanishes with respect it's degrees of freedom $\mathbf{x}$ in the deformed configuration:

$$\mathbf{f}(x^*) = \frac{\partial E_{\text{body}}}{\partial x} + \frac{\partial E_{\text{garment}}}{\partial x} + \frac{\partial E_{\text{attachments}}}{\partial x} = \mathbf{0} \ . \tag{5.10}$$

To solve this minimization problem, we use the L-BFGS [140] optimizer provided by PyTorch [141] which affords us with super-linear convergence, while also allowing us to simulate fairly dense topologies. We consider simulations converged once their gradient norm reaches 1e-7. The simulation is performed on a PC running a quad core 3.6Ghz I9 CPU and an NVidia 2080TI GPU. All calculations are done on the GPU.

5.4.4.1 *Initialization*

Since the garment is parametrized as a subdivided subset of the body mesh, we can initialize the garment positions in rest and deformed states in a straightforward manner by simply copying the position of the body vertices to the corresponding garment vertices. This initialization is important for our local IMLS field to produce meaningful gradients, as any garment vertices that are initialized too far inside the body would have no energy to push them to the surface.

5.5 ON-BODY TOPOLOGY OPTIMIZATION

In order to automatically generate energy-dense designs for a given motion and garment, we look towards topology optimization (TPO). The aim of such methods is to find an optimal *material distribution* within a prescribed domain under load, minimizing for a particular cost function (i.e. compliance) and staying within a material budget [36]. We surveyed two widely used TPO continuous and discrete approaches. SIMP is a continuous method that penalizes intermediate material densities with the power-law to encourage near-binary material distributions [142]. BESO [39] is a bidirectional evolutionary method that iteratively refines a binary material density $\{0, 1\}$ distribution towards a given material budget. In penalization based approaches like SIMP, intermediate results contain invalid non-binary material parameters and are thus not usable. Intermediate results in BESO on the other hand are valid as they are always in the $\{0, 1\}$ set. Comparisons between SIMP and BESO have been undertaken and show that they can produce similar results [41], however, only BESO can create a continuous space of optimal designs in one shot. The BESO algorithm is a better fit for our design scenario, allowing a continuous and coherent *exploration* of the design space of possible reinforcement layouts.

5.5.1 *Design Objective*

As a basis of our algorithm, we use the multi-material version of BESO proposed by [43]. In the standard setting, elements are entirely added or removed from the domain, while in the multi-material setting, removed elements are set to a weaker material. These two approaches are known as hard-kill and soft-kill BESO respectively [143].

As noted in Section 5.3, our goal is to increase the energy density of the garment while gradually decreasing the surface area instrumented with reinforced material. We define the following optimization problem:

$$\begin{aligned} \mathbf{d}^* = \arg\max_{\mathbf{d}} \quad & E_{\text{garment}}(x^*, \mathbf{d}) \\ \text{s.t.} \quad \sum_e A^e d^e = A^* \ , \quad & \mathbf{f}(x^*) = \mathbf{0} \end{aligned} \tag{5.11}$$

where A^* is the prescribed area of the domain, and $\mathbf{d}$ is the design vector determining the material assignment which, for each element, can take on the values $d^e \in 0, 1$. Note that when d^e is 0, its area does not count toward the target material area. In effect, for a given target area A^*, we are looking for the optimal material assignment $\mathbf{d}^*$ that maximizes the energy of the garment in its equilibrium state $\mathbf{x}^*$.

5.5.2 *Sensitivity to Objective*

In the BESO method, the optimal topology d^* is determined according to a relative ranking of sensitivity numbers, where the sensitivity of each element in the domain is evaluated w.r.t. to the objective function. The sensitivity number of linear materials has been derived simply as the per-element strain energy E^e [39], that is, the *variation* in the total strain energy is the same as the strain energy element if it were to be added or deleted. To take into account our non-regular mesh topology, we modify this by dividing by the element area. Thus, in our case, the per-element sensitivity α^e is equal to the elemental strain energy density W^e_{garment}.

In the multi-material setting, sensitivity numbers need to be modified to take into account the elastic moduli of each material. We follow a similar approach as Huang et al. [43], integrating the concept of penalization and material interpolation from SIMP, resulting in the following formulation for two materials:

$$\alpha^e = \begin{cases} \frac{1}{2}\left[1 - \frac{E_2}{E_1}\right] W^e_{\text{garment}} & \text{material 1 (Reinforced Cloth)} \\ \frac{1}{2}\frac{d_{\min}^{p-1}(E_1 - E_2)}{d^p_{\min}E_1 + \left(1 - d^p_{\min}\right)E_2} W^e_{\text{garment}} & \text{material 2 (Cloth)}\ . \end{cases} \tag{5.12}$$

where E_1 and E_2 correspond to the Young's modulus of reinforced cloth and cloth respectively, and $d_{\min}$ is the minimum value for the design variable d. While in our case d is in the set $[d \in 0, 1]$, setting $d_{\min}$ to 0 would simply switch the method to the hard-kill version as the secondary material would have 0 sensitivity. The power exponent p interpolates between the influence of the first and second material. For all experiments, we set $d_{\min} = 0.001$ and $p = 1.6$.

5.5.3 *Smoothing and Filtering*

Performing topology optimization with constant strain elements can result in the well known problem of *checkerboarding*, where the sensitivity of elements can become discontinuous across boundaries. To overcome this problem, a filter is applied to both smooth the sensitivity across elements, and to interpolate sensitivity numbers between the boundaries of multiple materials. The filter works by transferring elemental sensitivities into their connected nodes:

$$\tilde{\alpha}_i = \frac{\sum_{e \in \epsilon} A^e \alpha^e}{\sum_{e \in \epsilon} A^e} \quad . \tag{5.13}$$

where ϵ is the set of elements connected to node i, A^e and α^e are the element area and sensitivity respectively. It is important to note that nodal sensitivities carry no physical meaning, but can be interpreted as nodal sensitivity *density*. We use simple averaging $\widetilde{\alpha^e} = \sum_{i \in e}^{N} \tilde{\alpha}_i / N$ to redistribute sensitivity values back onto elements. The sensitivity numbers are further temporally smoothed with their historical information:

$$\tilde{\alpha}_i = \frac{1}{2}\left(\widetilde{\alpha_{i,k}} + \alpha_{i,\widetilde{k-1}}\right) \quad . \tag{5.14}$$

where k corresponds to the current iteration of the algorithm.

5.5.4 *BESO Procedure*

With the elemental sensitivities $\widetilde{\alpha^e}$ in hand, we can apply the BESO structural refinement procedure. First, The design vector d is initialized with a fully dense structure resulting in a starting area A_0. The following steps are repeated until convergence:

1. Finite element analysis is performed bringing the garment mesh into an equilibrium state x^*, resulting in sensitivity numbers α^e.

2. Apply spatial and temporal filtering to obtain $\widetilde{\alpha^e}$

3. Rank sensitivity numbers $\widetilde{\alpha^e}$ of all elements in the domain, and find the threshold α_{th} such that the area of all elements with higher sensitivities satisfies the current iteration's target area A_i.

4. Update the design variable d^e for all elements with sensitivities below α_{th} to 0, and set all those above to 1

5. To stabilize the procedure , a ratio $AR_{\max}$ controls the upper limit on elements that can be switched from 0 to 1.
6. Update the target area for the next iteration $A_{i+1} = A_i(1 \pm ER)$, where ER is the evolutionary ratio. If $A_i \equiv A^*$, then set ER to 0.

We consider the optimization converged when the target A^* is reached and the change in energy in Eq.A.1 is less than a predefined threshold τ over N previous iterations. A detailed procedure of two material BESO can be found in [43].

5.6 RESULTS

We conduct a three-part evaluation of the proposed computational pipeline. We start by validating our simulation model against a set of physical cloth samples in 2D. Next, we explore the capabilities of our method in a simulated on-body setting, generating designs for a diverse set of motion and garment pairs. We then fabricate and physically test 2 of these designs in a blind comparative user study.

5.6.1 *Material Parameters and 2D Validation*

We first conducted an experimental characterization of the material parameters of *cloth* and *reinforced cloth*. For this purpose, we custom built a pull-testing apparatus (See fig. 5.5) and measured the uni-axial forces resulting from a controlled in-plane displacement. For each material, we fabricated a 10×14 cm of fabric and clamp it at its two ends. This gives a 10×10 cm patch of active deformable area. Tensile forces were measured on a Pesola spring-scale. Fabric samples were taken directly from the arm sleeves used in the fabrication section. Based on force-displacement measurements, we determined a Young's modulus of 0.5 MPa for *cloth* 5.7 MPa for *reinforced cloth*. The thicknesses of the base cloth was 0.27 mm and 0.35 with reinforcements added. A Poisson's ratio of 0.33 was used based on prior work on elastic knitted fabrics [144].

Using the values derived above, we constructed an identical virtual testing apparatus and perform simulation studies in 2D. We compare the force output of two virtual samples against their physical counterparts, each with 40% area covered by reinforcements (See Fig 5.5). Our results show that the simulated and physical behaviour are in agreement for this validation setup. For example, at 10% induced deformation in both samples,

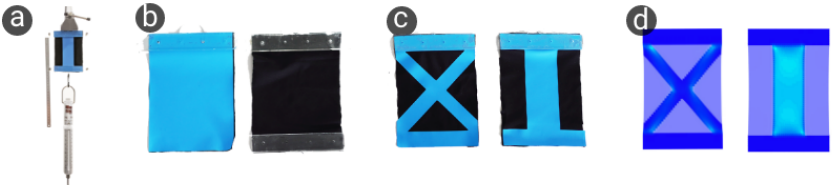

FIGURE 5.5: Material stiffens and 2D model validation. a) Built pull-tester. b) Samples of *cloth* and *reinforced cloth* used for Young's modulus to recover (0.5 MPa and 5.7 MPa, respectively). c) A separate set of samples was fabricated and tested against d) simulated designs within an agreement of 10%.

for a) the *Line*: we measured 5.5 N with the experiment and obtain 5.4 in simulation (2% difference); while for b) the *X* design we measured 10.8 N against the 9.9 N of simulations (8% difference).

5.6.2 *Simulated On-body Designs*

The goal of this study is to demonstrate both the diversity of designs as well as their performance in terms of energy per unit area. We begin by sampling the required poses and motions from the AMASS dataset [135]. We use only the pose parameters and zero out the β (shape) parameters. Four garments are designed using our tool to cover a variety of body sites (See Fig. 5.6). We set the following BESO parameters for all experiments: ER = 1.5%, AR_{max} = 1.5%, A^* = 15%, except in the case of knee-flexion and slouching, where A^* = 10%. As a primary measure, we use the objective function defined in Eq. A.1, that is, the energy density of just the reinforced part of the garment. We normalize this value to the fully dense design (equal to the first iteration of the algorithm) to allow comparison of designs in terms of their relative energy density. All designs are evaluated in simulation, and a subset of these designs are fabricated and evaluated by users. Parameters and performance for all experiments are listed in Table 5.1.

SIMPLE MOTIONS We begin by studying low degree-of-freedom joints at the elbow and knee. Results for flexion of these joints is shown in Fig 5.7 and 5.8. In the *Arm Flexion* case, we see a smooth and monotonically increasing objective function, with a smoothly decreasing area profile. When target area is reached, the designs and objective function are stable. A number of designs can be sampled from this space, depending on user needs.

model	vertices	BESO steps	time / iter.	energy density increase
Arm Flexion*	7.1k	150	10.65s	1.97x
Arm Extension	7.1k	180	10.88s	2.26x
Knee Flexion*	3.2k	180	10.47s	2.16x
Knee Extension	3.2k	150	10.24s	2.91x
Slouch*	28.3k	200	15.80s	2.15x
Crouch	9.5k	150	11.70s	2.32x
Lunge	9.5k	180	11.55s	3.02x

TABLE 5.1: Summary of parameters and performance for simulated results. * Indicates designs that were fabricated.

FIGURE 5.6: A range of garments designed with our tool. Pink areas represent attached areas. All garments start in a fully reinforced setting, which is represented with opaque blue color.

In the *Knee Flexion* case, we intentionally specify a smaller area budget (10%), resulting in the sub-optimal design shown in (e). This is due to insufficient material to sustain the more energy dense X design found in (d). In this case, our tool allows designers to simply revert to the more energy dense design with slightly more area coverage. In both optimization routines, the energy density of designs increases by over two times. In both designs, we observe the formation of *loops*—key structural features that emerge when attachments cannot be relied upon to hold the structure in place. One difference between the two most optimal designs in (d), is the lack of a loop directly over the joint. As the optimization is exploiting muscle bulging, there is less of this effect on the knee than in the arm.

Next, we look at designs for the same joints and garments in the case of arm extension and knee hyper-extension. Results for these motions are shown in Fig 5.10. Knee-Hyper extension, where the knee joints becomes

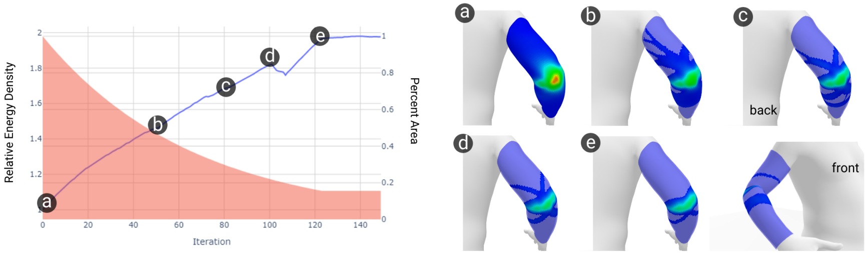

FIGURE 5.7: Optimization of Arm Sleeve targeting Arm Flexion. Starting with a fully dense design (a), the target area is reduced and reinforcements are placed in such a way as to smoothly increase the energy density of the garment. Many efficient designs can be sampled from this space (b-e).

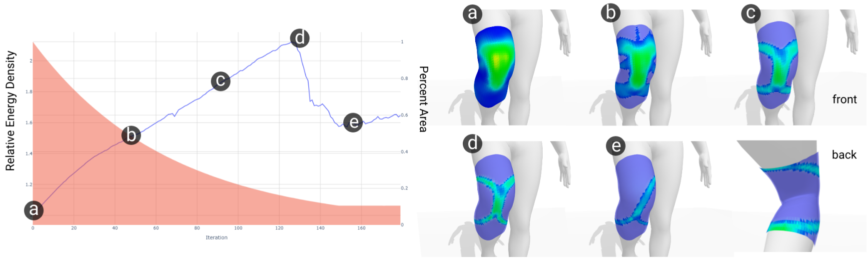

FIGURE 5.8: Optimization of Knee Sleeve targeting Knee Flexion. The energy density increases up to a point (a-d), but falls off drastically if too much material is trimmed. While the design in (e) is still more energy dense than (a), there is not enough material budget to sustain the more energy dense design in (d).

locked is not a desirable pose and can negatively affect posture. In this case, our optimization discovers two interesting features: for knee-hyper extension, a web-like structure emerges with a loop on the front of the knee. In the arm extension case, a two-sided loop wraps around to the back of the arm.

COMPLEX MOTIONS We evaluated our method on motions that involve multiple joints in both the upper and lower body. For the upper body, correcting bad posture induced by slouching is one possible application of kinesthetic garments. Fig 5.9 shows the optimization result for counteracting such a motion. The optimized design exhibits a harness-like topology with

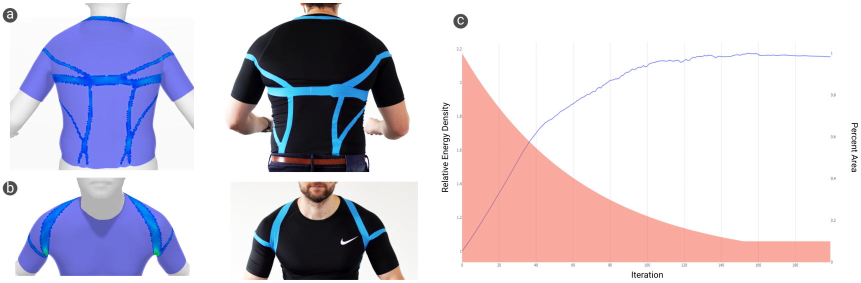

FIGURE 5.9: Optimized posture correcting shirt viewed from the back (a) and front (b). Area coverage is 10%. Note the smooth increase in energy density (c) over the optimization process.

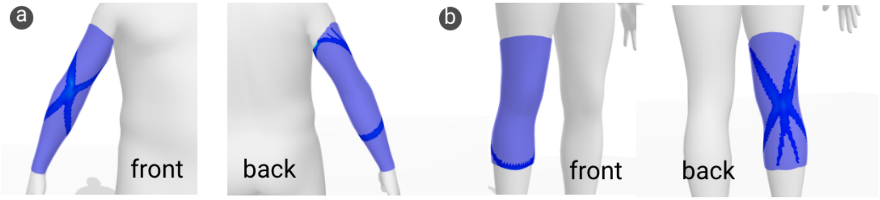

FIGURE 5.10: Knee (a) and Arm Sleeves (b) for Knee and Arm Extension. Loops are present in both designs and a spider-like web is apparent on the knee-design.

complex network of branches that connect to the base. Multiple loops are present, extending over the arms, around the front of the shoulders, and under the armpits. It is worth noting that the slight asymmetry observed in this design is due to the person-specific slouching pose.

For the lower body, we show the effect that different motions can have on a garment design. In this case, a potential application are cycling shorts designed for resisting motions during exercise. The resulting designs shown in Fig 5.11 are both appropriate and unique to the motion under which they were optimized.

WRINKLING Thin sheet materials generally wrinkle at the onset of compression. Resistance to compression is therefore orders of magnitude weaker than for stretching. Our relaxed energy formulation captures this behavior at the macro level for both garment and reinforcement materials without the need for modeling fine-scale wrinkles. Importantly, it tolerates compression in a given direction while accurately capturing the more vital stretching

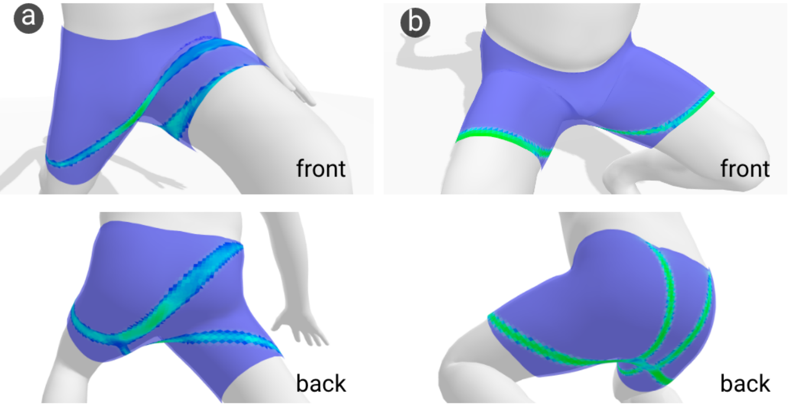

Figure 5.11: Optimized exercise cycling shorts for lunging (a) and crouching (b). Area coverage is 15%.

in the transverse direction—a combination that arises frequently in our physical samples. Figure 5.12 demonstrates a close match between our simulated compression field and the wrinkles in the fabricated garment.

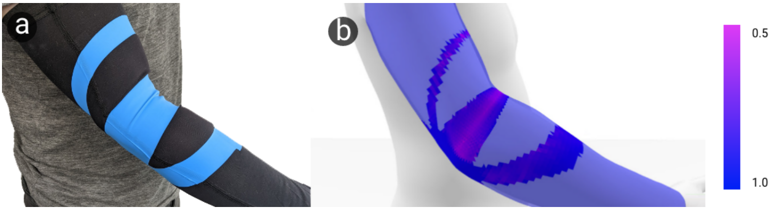

Figure 5.12: Wrinkling in simulation (a), compared to wrinkling in fabricated garment (b). Compression field is visualized from (blue) no compression, to (purple) compression in the transverse direction of strain.

5.7 USER EVALUATION

In order to validate the effectiveness of automatically produced designs, we select a subset of motions *Arm Flexion and Knee Flexion* and fabricate one optimal design for each. These designs are compared to 2 baselines: an X shape and a Line in a simple head-to-head user study. Each baseline was designed to be exactly the same area as the optimal design. We chose a line as it forms a geodesic between attachment points, while the X forms a loop on the arm.

FABRICATION We fabricated 3 designs for each joint, as shown in Fig 5.14. Optimal designs are exported as meshes, developed onto a flat surface, and processed in 2D to smooth edges. We maintain the area and shape of each design throughout the process. As our material, we used the Siser EasyWeed Stretch as it is one of few vinyls that is stretchable and can recover its form after deformation. For the stretchable garments for the arm and knee we used GripGrab UV sleeves. Designs are cut with a Siluhette Cameo 4 Pro plotter and applied with a Transmatic TMH 28 heat press for 16 seconds at 160 C. Designs are weeded in reverse fashion (peeling off the material to be applied), and all material is placed on the garment simultaneously, allowing for more accurate seam placement. The fabrication process is shown in Figure 5.13.

PROCEDURE Five healthy adult subjects (*M*=27.6; *SD*=4.14;) were recruited. Since we only fabricated one size of our designs, participants were all male and similar in size to the template SMPL mesh. All participants wore a blindfold so they could not identify which sample they were wearing. For each trial, two different designs were mounted on each arm. The center of the garment designs were carefully aligned to the center of the elbow (See fig 5.15. Participants were then asked to rotate their elbow to between 45-60 degrees slowly, to get a sensation of how much the resistance each garment on their arms produces. No time limit was set, and multiple attempts were allowed. After this, participants gave their preference for which garment produced more resistance to their motion. Each comparison was repeated with the arms switched, so the effect of arm strength could be balanced. The order in which garments was worn was randomized. A total of 12 trials were completed for the Arm Sleeve comparisons, and the procedure was repeated in identical fashion for the Knee Sleeve.

RESULTS For *Arm Flexion*, participants ranked the optimal design as providing more resistance in all but one case. The line design was never selected once. It's worth noting that both the X design and the optimal design contain loops in the upper and lower part of the arm, however, only the optimal design contains the loop in the center of the elbow. This could be one reason for performing better in this comparison.

For *Knee Flexion*, the X design was chosen over the optimal design in two out of ten comparison, and the Line design was selected three times over the

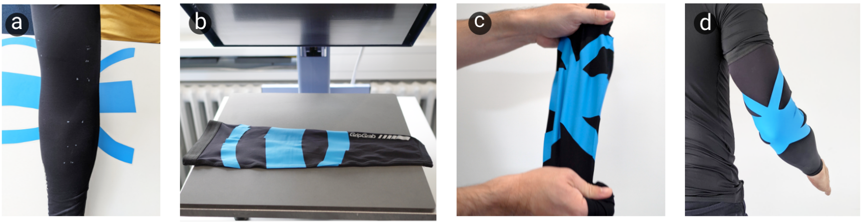

FIGURE 5.13: Garments reinforcement material (HTV) is cut (a) and transfered onto the base garment using a heat press (b). The resulting garment has high compliance (c) and can be worn as a typical sleeve (d).

FIGURE 5.14: Fabricated Garments for user study. Line, X, and Optimal designs for elbow sleeves (a-c), and knee sleeves (d-e) respectively.

X. Participants noted several times that the Knee Flexion comparison was more difficult to make. The less consistent ranking for the knee example could be due to the comparatively strong leg muscles, which might reduce sensitivity to changes in garment stiffness. One participant noted that he felt pressure on the front of his leg with the optimal design.

5.8 LIMITATIONS

A limitation of our method is that it does not account for the effect on nearby motions that should not be resisted. In particular, garments that cover areas with high mobility (i.e. around the torso) could impact freedom of movement. In these cases, the objective function could be modified with additional terms penalizing energy density in such nearby motions. To generalize the method further, multiple possible motions can be taken into account where each motion can be resisted or kept free.

A limitation of our model is the simplified attachment and contact potential, in particular, it cannot account for stick-slip behaviour that occurs when the strain energy of the garment causes the attachments to slip from

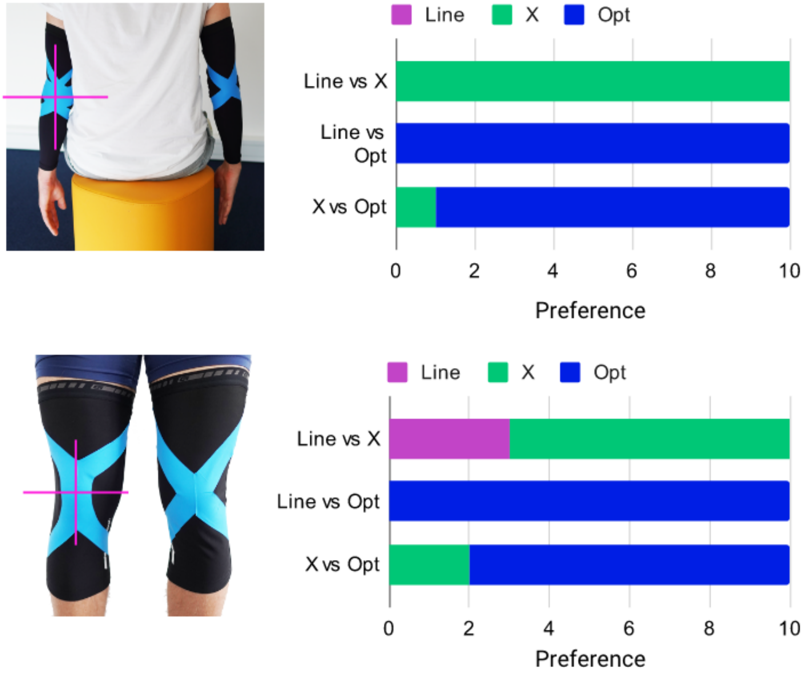

FIGURE 5.15: User study setup and preferences for Arm Flexion and Knee Flexion. The center of garments (represented with pink cross) was aligned to the mid point of the elbow/knee. Users prefer the BESO optimized design a 90% of the times in the arm (top) and 80% for the knee (bottom).

their original mounting positions. Attachments must also be specified manually by the designer which is appropriate only for pre-existing garment substrates (e.g. sleeves). Automatic placement of attachments could be integrated into the method via rope caging [145] to enable design of custom garments.

With regards to comfort, we assume that minimizing coverage will reduce the overall impact on the garments' original behavior. While coverage is not a direct measure of comfort, even more sophisticated functions of comfort will often involve some degree of coverage reduction [128]. A quantitative measure of comfort such as dense pressure can be integrated into the method to produce comfort aware designs.

5.9 SUMMARY AND OUTLOOK

In this chapter, we presented a topology optimization approach to the problem of designing garments that provide kinesthetic feedback. Our method was able to automatically generate efficient designs that maximally resist specified motions for given material budgets. Our in-simulation results showed that our approach generalizes well to a large space of human motion, ranging from simple flexion and extension examples for single joints to complex motions spanning multiple joints, such as posture changes. In our user studies, blind forced-choice comparisons of perceived resistance, user chose had a near-universal preference for our automatic designs.

Potential follow-up directions include adding a degree of user-input, for example to be able to influence the design of structures in desired directions (similar to [122]), and leveraging the computational pipeline for designing strain-based *sensors*. The basic principles of maximizing energy density while considering body deformation also apply in active systems, such as the ones found in wearable robotic garments [13].

This chapter provided a foundation for bridging two disparate disciplines: classical engineering structural optimization approaches with parametric human-body models. This is what enables the idea of *on-body topology optimization*. While in this work we focused on designing garments with kinesthetic properties on the body, there are many potential applications outside of garment design. One aspect that is particularly lacking in the current formulation is the idea of on-demand feedback, or more generally, the integration of active components. In the following chapter, we show that on-body TPO is a sound foundation by extending it to support such use cases, opening up additional real-time applications such as feedback in VR.

6

COMPUTATIONAL DESIGN OF ACTIVE KINESTHETIC GARMENTS

The focus of this chapter is to build up on the foundation of on-body topology optimization (TPO) by introducing *active components* into the formulation. Active components are an essential part of wearable systems that provide feedback on-demand. The main challenge in integrating such components into the on-body TPO formulation is to take their state into account, where depending on the state, the wearable may have different performance and wearability characteristics. From this perspective, this chapter makes several technical contributions: formalizing the notion of impediment to the user as strain energy when active components are *Off* and the performance of the garment as strain energy when the active components are *On*, and a dual-objective that combines both. We further extend the formulation from the previous chapter to support the resistance of multiple motions. The resulting physical haptic-interfaces are termed Active Kinesthetic Garments, soft-exoskeletal like garments that can resist multiple motions *on-demand*. Results showed that when comparing to manually design equivalent garments, our method is able to yield drastically higher physical resistance in simulation and in force characterization experiments. In a VR user study task, it performed significantly better for some garments than manually designed counterparts.

6.1 INTRODUCTION

Kinesthetic garments are an efficient and non-obtrusive way of providing force feedback for human body motion. By augmenting stretchable fabric with strategically designed reinforcements, they offer targeted resistance to motions along specific directions [146]. They are part of an emerging trend of soft robotic garments [13] that have the potential to assist human wearers in various ways such as during locomotion [9, 10], rehabilitation [147], and increasing immersion in mixed reality [6, 7, 148]. However, relying only on *passive* mechanical structure for feedback prevents their use in such applications because they require *active* feedback.

In this work, we propose a computational approach for designing *active kinesthetic garments* that can resist user-defined motions *on demand*. To implement such adaptive resistance, we rely on electrostatic clutches [149], i.e., pre-fabricated components that provide extremely high stiffness contrast between their active and inactive states. Designing active kinesthetic garments then amounts to determining clutch placements, typically placed over high-strain areas, and finding a passive structure that connects and anchors the active components. Crucially, this layout should result in the garment providing maximal resistance when clutches are active, but minimally interfere with motion otherwise. Designing effective connecting structures requires the understanding of the interaction between stretchable garments in multiple states sliding over a deforming body in multiple poses, a very difficult and unintuitive task.

To address this challenge, we formalize the design of active kinesthetic garments as an on-body topology optimization problem whose objective function explicitly balances the opposing goals for active and inactive states. By maximizing the difference in elastic energy between active and inactive states, our formulation encourages layouts in which clutches link disconnected parts of the passive structure. In this way, clutches leverage the passive structure to establish strong, load-carrying paths when active while maintaining freedom of movement otherwise.

We implement our formulation within a standard evolutionary optimization algorithm, and produce a set of active kinesthetic garment designs that each target multiple motions spanning different body sites. Our results indicate that designs produced with our approach are highly effective and outperform manually-designed alternatives by significant margins. To further substantiate this analysis, we manufacture a subset of our designs for experimental evaluation. Both mechanical testing and a VR pointing task indicate clear advantages for the designs created with our method. To summarize, we make the following contributions:

- A *computational design pipeline* for the automatic creation of active kinesthetic garments that includes a novel objective function that considers active components and multiple motions.

- A set of *fabricated active kinesthetic garments* built on compliant material integrating ES clutches as kinesthetic feedback components.

- A comprehensive *evaluation* showing the effectiveness of our method in simulation, in a physical validation, and in a VR user study against manually-designed and visual baselines.

6.2 RELATED WORK

We summarize works in the areas of computational methods in garment modeling and augmentation, intersecting with hardware and devices capable of providing body-scale kinesthetic feedback.

BODY-SCALE KINESTHETIC HAPTIC FEEDBACK SYSTEMS Early work to provide kinesthetic feedback to the body used motors and hydraulic pistons to actuate heavy bulky haptic platforms. More recently, several wearable body kinesthetic feedback systems have been developed, mostly based on electromagnetic motors [150, 151] with rods [152, 153] or cables [154, 155] transmission, and based on pneumatic actuators [6, 156] which are soft and more comfortable at the detriment of a bulkier equipment (pumps, compressors, valves). An alternative way to provide body kinesthetic feedback are passive blocking mechanisms like vacuum jamming [93] (still requiring pumps) and ES clutches [17, 106, 157, 158]. In particular, ES clutches offer the advantages of being ultra-thin, light, and soft enabling the design of compliant kinesthetic garment designs. Such kinesthetic systems are typically manually designed to specifically fit a limb/joint and block a certain motion. In contrast, we leverage an automatic design method that models and simulates clutches, allowing us to accommodate any set of motions and body areas.

TOPOLOGY OPTIMIZATION Topology optimization is a powerful method used in engineering disciplines to most efficiently distribute a finite amount of material, typically to minimize compliance [36, 159]. The graphics community has also combined compliance minimization with user guided input [121, 122]. It has also been demonstrated on elastic materials [124] as well as structures undergoing large displacements [125]. Closer to our work, topology optimization has moved into the on-body domain where it has been used for personalized cast design [127] and casts designed for thermal comfort [128]. Most recently, Vechev et al. demonstrated the design of *kinesthetic garments*, which are passively reinforced garments designed to resist a single motion [146]. However, this work only formulates a single compliance minimization objective, and thus cannot be used in a setting that leverages active components. We extend this approach in two important ways, first by the addition of a dual objective that considers the active and inactive states of our components. A second important contribution is a formulation that enables optimization for multiple motions.

INTELLIGENT GARMENT AUGMENTATION The intelligent design of garments is an emerging discipline with important applications for the general population. Computational design approaches to garment design have recently started to consider motion as a fundamental design quantity in so-called 4D garments [114] that minimize friction and pressure via integrated knitting maps. In addition to minimizing friction during motion, Montes et al. also optimize for pressure distributions and body fit by employing a physically based model of skin-tight garments on the body [14]. Vechev et al. augment existing skin-tight clothing with passive reinforced materials to resist a single given motion, employing a more flexible model of the garment that allows cloth to slide and lift-off from the body [146]. Optimization of component placement has also been used in soft-robotic garments, in combining elastic cords, clutches, and dampers to reduce the force and power required by a person to generate lower body motion [115]. Evolutionary optimization techniques were employed by Gholami et al. for designing garments with optimally placed fabric sensors [133]. Muthukumarana et al. integrated combinations of active shape-memory based components into garments allowing for actuation on clothing [160]. In our work, we augment garments with active components that generate kinesthetic feedback and design supporting optimization objectives to create efficient structures connecting them.

6.3 COMPUTATIONAL DESIGN PIPELINE

Our method supports designers in the task of creating *active* kinesthetic garments that can resist any motion from a predefined set of movements. The design goals of our pipeline are to enable kinesthetic garments that maximize the feedback felt by users when ES Clutches are active, while minimizing interference with their motion when inactive. Our pipeline consists of three main phases: 1) input — where designers specify motions, garment designs, and clutch placements; 2) automatic design — where our method automatically links clutches with stiff material to satisfy the above design goals; and 3) a fabrication phase. The full computational pipeline is illustrated in Fig 6.1.

6.3.1 *Input*

Our pipeline requires three components as input: a set of body motion, a base garment, and a predefined number of ES clutches.

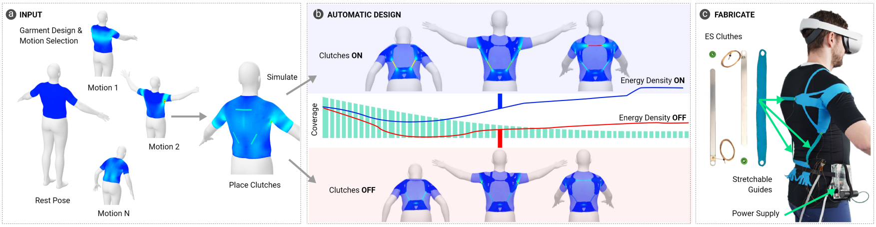

FIGURE 6.1: Pipeline overview, from left to right (a) *Input:* Designers specify N motions with a single rest pose, and a garment (shirt). The simulated behaviour of the garment is shown, with light blue indicating high strain. Clutches are then placed on the garment over high strain areas. (b) *Automatic Design:* We simulate the garment under each motion with all clutches ON (top), or OFF (bottom), noting that the energy difference between these states increases until the target coverage is reached. Designs can be sampled at any point in the evolutionary progress. (c) *Fabricate:* We assemble the ES clutches made of flexible strips sliding in a stretchable textile guide and fabricate the connecting structures attaching everything together to make the final active kinesthetic garment (right).

MOTIONS are specified using the STAR/SMPL parametric human body model [44, 48] which produces a surface mesh v with N = 6890 vertices $\in \mathbb{R}^{3n}$ based on a 72 pose ` parameters. To create target poses, we sample from the AMASS dataset [135] and make individual adjustments to the ` as needed. We define a motion as a single rest pose $\bar{\mathbf{v}}$ and an accompanying deformed pose $\mathbf{v}$. A *set* of motions is defined with $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_i)$ deformed poses, and a common rest pose $\bar{\mathbf{v}}$.

GARMENTS AND CONNECTING STRUCTURES are modeled as 2D mesh surfaces embedded on the body, initialized with the same rest and deformed nodal positions as the underlying body mesh. A garment in its rest state is defined through nodal positions $\bar{\mathbf{x}} = (\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n) \in \mathbb{R}^3$ and $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n) \in \mathbb{R}^{3n}$ when deformed. The connecting structure of the garment is modeled using a bi-material distribution where each triangle element e of the garment mesh is assigned a specific material property. This property is set through the design variable $d^e \in [0, 1]$ for each element e, where 0 and 1 correspond to *cloth* and *reinforced cloth* respectively.

ACTIVE COMPONENTS In our formulation, ES clutches are modeled as rectangular surface meshes that are attached to the garment at a predefined set of vertices. A key requirement for optimal ES clutch operation is that they are initialized in a *taut* state, that is, all slack must be removed from the system before forces are felt at the endpoints. We create a low-dimensional parametrization of ES clutches that is defined by the following variables: a *starting point*, a *surface direction*, and a *length and width*. From this, we procedurally generate a spline, and extrude a mesh (see Appendix A.1) with rest vertices $\mathbf{q} = (\mathbf{q}_1, \ldots, \mathbf{q}_m) \in \mathbb{R}^{3n}$ and $\bar{\mathbf{q}} = (\bar{\mathbf{q}}_1, \ldots, \bar{\mathbf{q}}_m) \in \mathbb{R}^{3n}$ when deformed. The endpoints q^c of the clutch mesh (three at each end) are connected to the garment mesh using simple quadratic penalty functions, which allows for firm attachment.

6.3.2 *Automatic Design*

Finding a passive mechanical structure that optimally connects electrostatic clutches placed by the user is a key challenge in the design of active kinesthetic garments. In the previous chapter we demonstrated a method for on-body topology optimization using a single compliance-minimization objective (summarized in Appendix A.3). Such an objective cannot be applied in our setting, as it has no notion of component states, and the single objective does not sufficiently capture the high-level goal of minimizing motion interference when components are inactive. Therefore, we propose to extend this formulation from passive reinforcements to our setting of active kinesthetic garments by *(1)* distinguishing between active and inactive clutch states by extending the simulation model with stateful components, *(2)* reconciling the different design goals for active and inactive states through a new state-dependant dual-objective, and *(3)* accounting for multiple motions.

ACTIVE COMPONENT MODEL AND SIMULATION ES Clutch stiffness varies according to their state, thus, we model their behaviour using a bimodal material. We implement this as a neo-Hookean material that resists compression and changes modes depending on the activation vector $\gamma = [\gamma_0, \gamma_1, ..., \gamma_n], \quad \gamma_n \in [0, 1]$. Each γ_i determines the state of clutch i, with 0 and 1 corresponding to inactive and active states, respectively. The Young's modulus of the clutch material is then set to $Y^i_{\text{clutch}} = \gamma_i Y_{\text{stiff}} + (1 - \gamma_i) Y_{\text{cloth}}$. The elastic energy stored in the clutches during deformation is defined as $E_{\text{clutches}}(q, \gamma)$. We define penalty terms $E_{\text{body}}(v, \mathbf{q})$ preventing clutches from entering the body, and an additional term $E_{\text{attach}} = \frac{1}{2}k(q^c - x^x_c)^T(q^c - x^x_c)$ that attaches the six endpoint vertices to their respective locations x_c on the garment. Throughout all examples, we set constant values for Young's Modulus to $Y_{\text{cloth}} = 0.5\text{MPa}$, $Y_{\text{reinforced\_cloth}} = 0.5\text{GPa}$, and $Y_{\text{stiff}} = 3.0\text{GPa}$. We use a Poisson's ratio of 0.33 for all materials.

We combine our active component model, with the garment-on-body model described in [146]. The terms $E_{\text{garment}}(\mathbf{x}, d)$, $E_{\text{body}}(v, \mathbf{x})$, and E_{attach} are summarized in Appendix A.2. With the model and energies defined above, we perform a quasi-static simulation by solving an unconstrained optimization problem,

$$\begin{aligned} \mathbf{x}^*, \mathbf{q}^* = \arg\min_{\mathbf{x},\mathbf{q}} \quad & E_{\text{garment}}(\mathbf{x}, d) + E_{\text{body}}(v, \mathbf{x}) + E_{\text{attach}}(\mathbf{x}) + \\ & E_{\text{clutches}}(\mathbf{q}, \gamma) + E_{\text{body}}(v, \mathbf{q}) + E_{\text{attach}}(\mathbf{q}) \ , \end{aligned} \tag{6.1}$$

using the GPU-based L-BFGS [140] optimizer provided by PyTorch [141]. We take advantage of GPU parallelism by simulating all states (poses) simultaneously. We consider simulations converged once the gradient norm of (6.1) reaches 1e-7.

STATE-DEPENDANT DUAL-OBJECTIVE A central goal for the structural optimization step is to find a material layout such that the garment resists the specified motions as strongly as possible when clutches are active, while showing minimal resistance otherwise. Assuming all-elastic materials, we translate this goal into the requirement that the stored energy of the garment should be maximized when clutches are active, and minimized when they are inactive. Our key insight is to introduce an *energy differential objective* that combines these opposing goals as

$$\begin{aligned} \mathbf{d}^* = \arg\max_{\mathbf{d}} \quad & E_{\text{garment}}(x^*_{ON}(\mathbf{d}, \mathbf{q}, \mathbf{fl}), \mathbf{d}) \\ & - E_{\text{garment}}(x^*_{OFF}(\mathbf{d}, \mathbf{q}, \mathbf{fl}), \mathbf{d}) \ , \\ \text{s.t.} \quad \sum_e A_e d_e = A^* \ , \quad & \mathbf{f}(x^*_{ON}) = \mathbf{0}, \quad \mathbf{f}(x^*_{OFF}) = \mathbf{0} \end{aligned} \tag{6.2}$$

where $\mathbf{q}$ holds the variables of all clutches, and x^*_{ON}, x^*_{OFF} are distinct equilibrium states corresponding to all clutches being active ($\gamma_i = 1 \forall i$) and inactive ($\gamma_i = 0 \forall i$), respectively.

To solve this optimization problem with the BESO algorithm, we must compute the per-element sensitivities, i.e., the partial derivatives of the objective function with respect to per-element material assignment variables d^e. Following (6.2), we simply have to sum the sensitivity values for the active and inactive states to obtain a single value that is used in the BESO ranking procedure. Everything else follows the procedure described in [146] and is summarized in Appendix A.3.

MULTIPLE MOTIONS Whereas the method described in [146] computes static reinforcements for a single target motion, we ultimately want to move towards *programmable garments* that can resist many motions by use of their active components. To this end, we extend (6.2) to the multi-motion setting by summing contributions for all poses as

$$\begin{aligned} \mathbf{d}^* = \arg\max_{\mathbf{d}} \quad & \sum_k \hat{E}^k_{\text{garment}}(x^*_{k,ON}(\mathbf{d},\mathbf{q},\mathbf{fl}),\mathbf{d}) \\ & - \sum_k \hat{E}^k_{\text{garment}}(x^*_{k,OFF}(\mathbf{d},\mathbf{q},\mathbf{fl}),\mathbf{d}) \\ \text{s.t.} \quad \sum_e A^e d^e = A^* \ , \quad & \mathbf{f}(x^*_{k,ON}) = \mathbf{0}, \quad \mathbf{f}(x^*_{k,OFF}) = \mathbf{0} \ \forall k \ , \end{aligned} \tag{6.3}$$

where k runs over all input poses. A problem with this simple approach is that the optimization may receive larger rewards for increasing an already good performance for a given pose instead of improving the worst-performing case. We address this problem by normalizing the strain energy density for each pose in a pre-processing step

$$\hat{E}^k_{\text{garment}} = \sum_e t^e A^e \hat{W}^{k,e}_{\text{garment}}(x^*, d^e), \ \hat{W}^{k,e} = \frac{W^{k,e}}{\max_e(W^{k,e})} \ . \tag{6.4}$$

In this way, each pose is given the same importance, irrespective of its initial strain energy, thus encouraging material layouts that more evenly distribute the garment's performance across all input motions.

6.3.3 *Hardware Details and Fabrication*

In the last step of the pipeline, designs are fabricated.

ES CLUTCHES provide resistance to elongation when active [149], while remaining stretchable with low resistance when inactive. They are thin, light and flexible which make them highly compliant and consume very low power when engaged (e.g. one 14cm by 1cm clutch consumes 12 mW at 350V). The ES clutches from [149] were modified for better integration by making them stiffer to reduce bending, packaging them in elastic guides to keep them fully self-retractable and safer for on-body use. Each ES clutch is composed of 3 parts: an electrode strip, an insulating strip, and a stretchable textile guide. Strips are made of 125 μm metalized polyester films from McMaster-Carr. Films are laser cut into long 1cm wide strips of various lengths. Additionally, insulating strips are covered with a 25μm thick layer of poly (vinylidene fluoride-trifluoroethylene-chlorotrifluoroethylene) from Piezotech-Arkema [149].

GARMENTS AND ATTACHMENTS All designs are exported as meshes and manually processed in Blender. We simplify geometry, and unroll the designs onto flat surfaces using the Paper model plugin (without changing area). As our connecting material, we attached a layer of polyurethane (Siser EasyWeed Stretch) onto 100% cotton fabric. This material combination enables much higher forces than in [146]. As base garments we used stretchable GripGrab UV sleeves and Nike Dri-Fit Pro Compression shirts. The different parts of the connecting structure were cut with a Trotec 300 laser cutter and glued onto base garments following marks taken on an experimenter wearing the garment. Next, pressure buttons are riveted at locations where ES clutches connect. Finally, ES clutches are fixed onto the garments using pressure buttons and connected with thin wires to a custom voltage power supply powered by a USB power bank and controlled by Bluetooth (see Fig. 6.1b). The overall modular system can accommodate different sizes of clutches and slight variations in body sizes.

6.4 EVALUATION

We conduct a multi-faceted evaluation of our method showing results for different types of motions and garments in simulation, a mechanical force study, and a VR pointing task.

6.4.1 *Automatic Designs*

We show a range of designs produced by our method for a variety of motions and garment types. For all experiments, we use a common rest pose with the body in an A-pose, and sample from a set of motions that include *Arms Forward, Arms Raise, Arm Flexion, Arm Extension, Bend Forwards* (see Fig.6.2). Three garments are designed using our tool to cover a variety of body sites: a short-sleeve shirt, an arm-sleeve, and a long-sleeve shirt. All clutches are placed manually on the garments, typically over high strain energy areas of the garments (see Fig. 6.1a). We set the following standard BESO parameters for all experiments: evolutionary rate $ER = 1.5\%$, maximum material added per iteration $AR = 1.5\%$, material interpolation $p = 1.6$. Similarly, we set the material budget to $A^* = 15\%$ for all examples except for arm flexion and extension, where we use $A^* = 20\%$. As our primary metric, we use the *relative energy density*

$$\rho(\gamma, \mathbf{d}) = \frac{E_{\text{garment}}(x^*(\gamma, \mathbf{d}), \mathbf{d}) \cdot A_{\text{dense}}}{E_{\text{garment}}(x^*(\gamma, \mathbf{1}), \mathbf{1}) \cdot A_{\text{opt}}} \,, \tag{6.5}$$

i.e., the ratio between energy density for the optimized and fully dense designs. As the optimization progresses, we expect to see a widening gap in this metric between active and inactive clutch states (see Fig. 6.5 for a visualization).

SINGLE-MOTION DESIGNS We begin by showing results for the single motion cases of our method. We target Arm Flexion with a single (8cm) clutch on the elbow, and Arm Extension also with a single (8cm) clutch on the inside of the forearm. We show two separate results in Fig. 6.2 a, and b. Relative to the fully dense design, we see that energy density *increases* to 1.14 for flexion, and 1.74 for extension when clutches are active. When clutches are inactive, relative energy density *decreases* to 0.34 and 0.54, respectively.

Next, we target single motions on the upper body using three clutches of 15cm length. Fig. 6.2 shows results for Arms Forward, Arms Raise, and Bend Forwards, with increases in relative energy density of 2.13, 1.51, and 2.47 respectively. For deactivated clutches, we observe that relative energy density decreases to 0.48, 0.66, and 0.73 for each design.

MULTI-MOTION DESIGNS The ability to resist multiple motions with a single design is an important step towards programmable active kinesthetic garments. We used our method to create three such designs, starting with an arm sleeve design (Fig. 6.3) that combines Flexion and Extension. It uses the same 20% material budget as in the single motion designs, but now this material must be distributed to balance performance for two distinct motions. The optimized design achieves relative energy densities of 0.88 and 1.27 for Flexion and Extension, respectively, which is 77% and 73% of the corresponding single-motion designs. For perspective, when evaluating the single-motion designs for Flexion/Extension on the Extension/Flexion motion, the relative efficiency is only 2%/5%. These results are not unexpected as Flexion and Extension are orthogonal motions such that designs optimized for only one of them are ineffective for the other one.

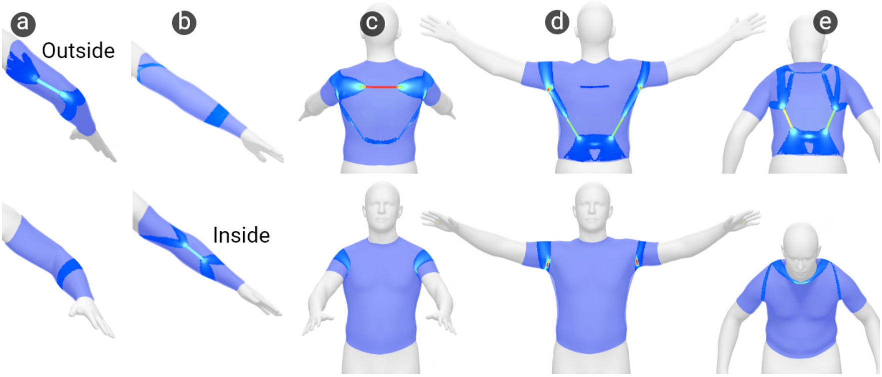

FIGURE 6.2: Single-Motion designs for (a) Arm Flexion, (b) Arm Extension, (c) Arms Forwards, (d) Arms Raise, and (e) Bend Forwards. Color coding indicates energy density.

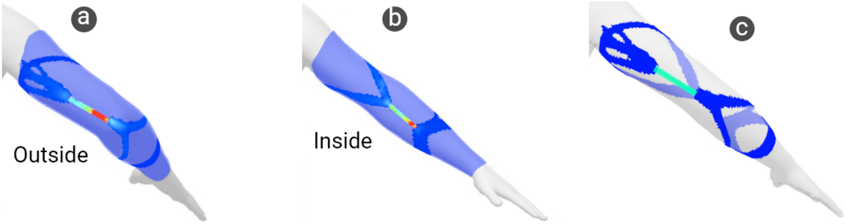

FIGURE 6.3: Multi-motion design for simultaneously optimized for (a) Arm Flexion and (b) Arm Extension. This design effectively integrates the single motion designs of Fig. 6.2 into an intertwined structure (c).

Our second design is a shirt that combines three upper body motions as shown in Fig. 6.4. Many of the features observed in the single-motion versions can be seen here, with clutches linking disconnected reinforcements. It is worth noting that each of these motions leads to a distinct load path (light green/yellow) running through at least one of the clutches. We also compare the performance of the multi-motion design to the single-motion versions. As can be seen in Table. 6.1, the multi-motion design is within 83%, 72%, and 65% as efficient as the single-motion designs, and yet using the same material budget. The performance of the single-motion designs on motions for which they were not optimized is, again, significantly lower.

Additionally, for each motion we show the progress plots of the evolutionary optimization in Fig. 6.5. As our automatic design method removes material, we see a clear separation in relative energy density for active and inactive states for all three motions. In the *inactive* mode, the relative energy densities of the garment for each motion are *decreased* by 0.62, 0.72, and 0.5, showing that our method is able to consistently achieve its minimization objective.

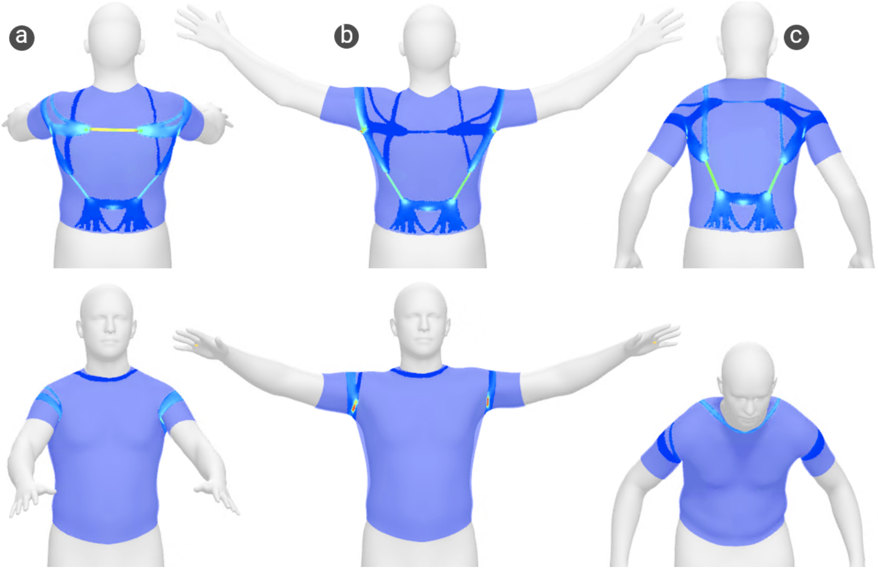

FIGURE 6.4: Active kinesthetic shirt designed for the three motions: (a) Arms Forwards, (b) Arms Raise, and (c) Bend Forwards. Strain energy density is shown in color-coding with increasing intensity from *dark blue* to *red*.

Our final example investigates the scalability of our method to more complex scenarios involving five clutches and five motions. The performance of this design exhibits relative energy density increases of 1.41, 0.85, 1.62, 0.64, and 1.16 for the motions Arms Forwards, Arms Raise, Bend Forwards, Arm Flexion, and Arm Extension, respectively. These numbers are comparable to the results obtained for our other multi-motion garments, especially as the allotted per-motion coverage has decreased overall. In general, the more motions a given design supports with the *same* material coverage target (i.e. 15%), the material available per motion will decrease and thus be less energy-dense in the ON state. In this case, the material coverage target can be increased, or the designer can sample from an earlier progression step with higher coverage.

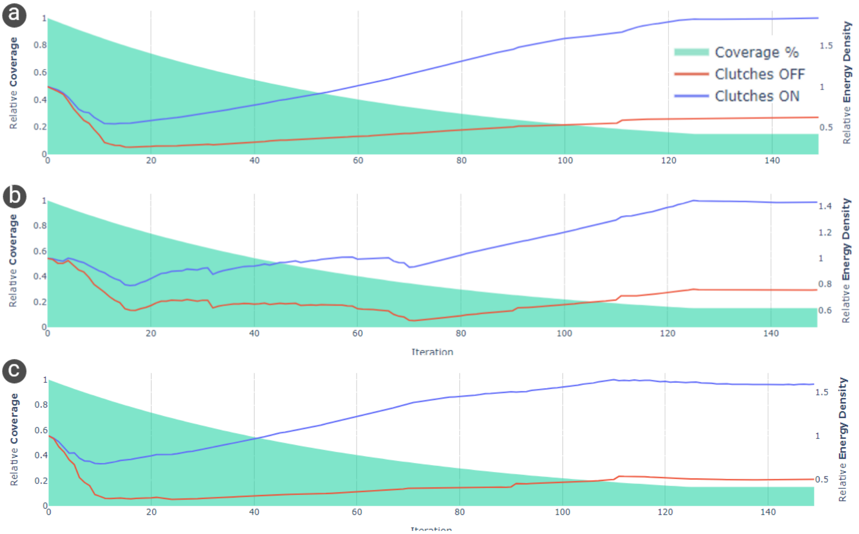

FIGURE 6.5: Progress of the evolutionary optimization algorithm for the Shirt Design for (a) Arms Forwards, (b) Arms Raise, and (c) Bend Forwards.

6.4.2 *Comparison to Manual Designs*

We conducted a pilot study to provide a manual baseline for our automatically generated designs. A central question in this context is whether users converge towards particular designs and if those designs exhibit features found in automatically generated ones. We recruited six participants (5M, 1F), two of whom were experts in structural optimization techniques (P2, P3). Using our interactive tool, we asked users to 'draw' stiff material on garment meshes, connecting a set of already placed clutches. Participants were asked to distribute material in such a way as to maximally resist the set of specified motions when clutches are activated. Each participant created two designs, a 2-clutch, 2-motion arm sleeve, covering no more than 20% of the available area, and a 3-clutch 3-motion shirt with a coverage budget of 15%. Each of these designs corresponds to an automatically generated designs shown in the previous section. The secondary goal of minimizing energy when the clutch is inactive was not assigned.

The manually-created results shown in Fig. 6.7 exhibit large variety in their designs. While most examples can be expected to perform reasonably, none of them resembled the automatically generated designs. Compared

Evaluated On	Optimized For			
	Forwards	Raise	Bend	All
Arms Forwards	2.13	0.46	0.34	1.77
Arms Raise	0.62	1.51	0.14	1.08
Bend Forwards	0.35	0.81	2.47	1.60

TABLE 6.1: Comparison of garments optimized for a single motion against a garment optimized for all three motions. A higher number corresponds to an increase in relative energy density when clutches are active.

to the fully dense version, manually-created designs were only 0.48x and 0.27x as energy-dense for the arm sleeve and shirt, respectively. Automatic designs, on the other hand, showed a 1.1x and 1.48x higher energy density. We can see that in the case of designing for a larger number of motions, the effectiveness of user designs drops drastically, while automatically generated designs can maintain a relatively high energy density. Looking at only designs from expert users, we see relative average energy densities of 0.54 for the sleeve, and 0.34 for the shirt, still much lower than our automatic designs. Non-expert designs on the other hand had average relative energy densities of 0.44 for the sleeve and 0.23 for the shirt, showing a much larger drop in performance for the more complex shirt design. Thus, automatic design methods can be especially useful for such users. Table 6.2 summarizes these findings.

6.4.3 *Physical Validation*

We seek to quantify the resistive force of our automatically designed garments under the motions for which they were optimized, and compare them against Manual-Design counterparts. We selected two designs for fabrication - the multi-motion arm sleeve and multi-motion short-sleeve shirt. We fabricated both, the designs produced by our automatic method and the corresponding manually-designed garments. For the shirt, we selected the clearly highest performing garment, which was from P2, while for the sleeve, we selected the design from P1. This sleeve design represents a common (line) design seen in literature [106, 157, 158], while having similar performance as other designs.

In order to best isolate the impact of the connecting structure, we replace clutches with flexible plastic strips that connect to a force sensor as shown

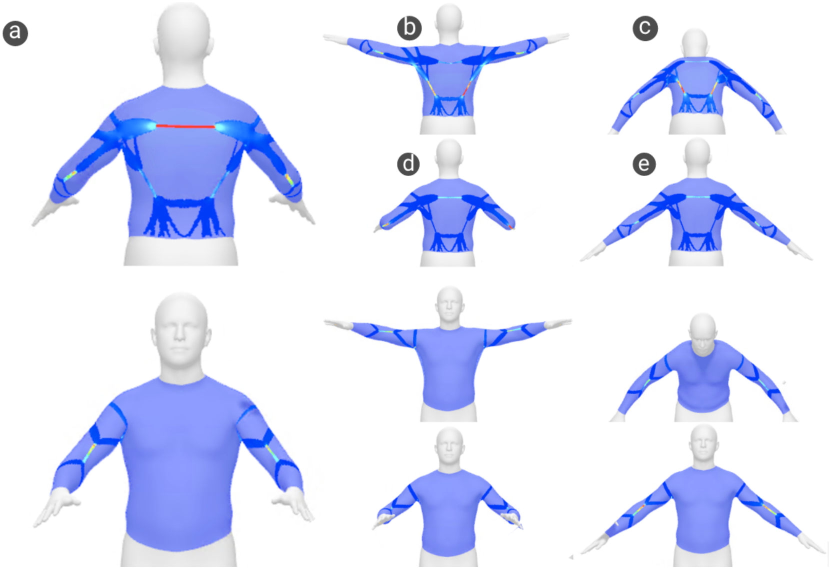

FIGURE 6.6: Active kinesthetic long-sleeve shirt with five clutches designed for five motions as indicated.

in Fig. 6.8. For the Arms Forward motion, we mount the force sensor in the upper back, while for the Arms Raise and Bend Forward motions, we mount it on the bottom left clutch location. The target motion is then slowly performed by the experimenter wearing the garment (three trials per motion), while the force is measured using a 10kg DYLY-108 force sensor with an HX711 load cell amplifier (see the Video Figure for visual demonstration). The results shown in Fig. 6.8 indicate that, relative to the manual design, the designs generated by our method were on average two times and up to four times more efficient in terms of force output. These measurements confirm our observations made on simulation results in which, as for the experimental case, the largest difference in relative energy density was observed for the Arms Raise motion.

6.4.4 *User Evaluation*

To quantitatively evaluate the ability of our active kinesthetic garments to efficiently block motion, we conduct a user study based on a VR pointing task in which participants were asked to reach targets from a predefined set of locations within their reach. The hypothesis that we seek to test is that,

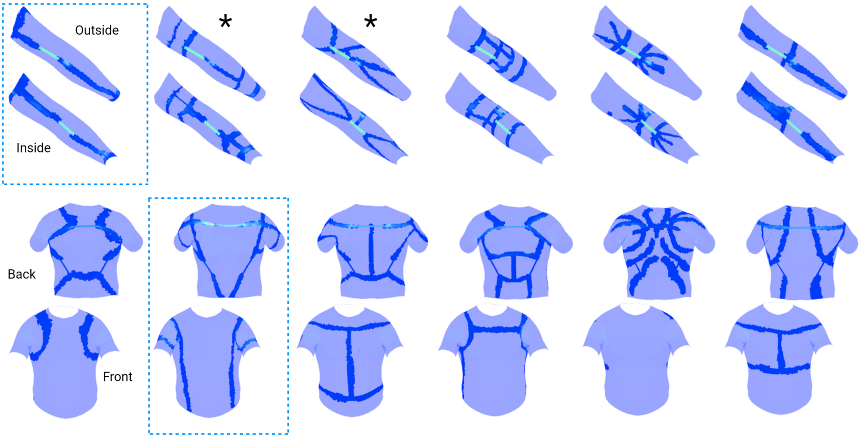

FIGURE 6.7: Manually-designed garments for 2-motion arm sleeve (top), and 3-motion shirt (bottom) for participants P1 (left) to P6 (right) with clutches shown in green. Note the large variance among the designs, particularly in the shirt case. The sleeve from P1 and the shirt from P2 were chosen for fabrication. * Denotes expert designed garments.

when wearing our optimized designs, users generally need more time to reach targets when clutches are active compared to when they are inactive. A secondary hypothesis is that our automatically generated designs lead to higher blocking efficiency than a user-generated baseline.

PROCEDURE AND SETUP Six healthy adult subjects (*M*=28.1; *SD*=4.14;) were recruited. Since we only fabricated one size of our designs, participants were all male and similar in size to the template STAR mesh. All participants wore noise cancelling headphones. The procedure and tasks were described and an introduction to the garments and the active components was given. After donning the garments (shirt and sleeve), clutches were attached and adjusted according to participant size to achieve sufficient pre-tension. The left hand of participants was rested on a tripod such for stability. Participants were then introduced to the VR setting and asked to practice touching the spherical targets with and without clutch activation until they felt comfortable proceeding. The study was implemented in Unity 2021 using a Meta Quest 2 relying on the built-in hand-tracking functionality.

STUDY DESIGN We use a within-subject design with two independent variables: *Feedback Type {via Auto-Designed, via Manual-Designed, Visual Only}*

Garment / Motion	Auto	P1	P2*	P3*	P4	P5	P6
Sleeve / Flex	**0.88**	0.50	0.53	0.54	0.67	0.26	0.58
Sleeve / Ext	**1.28**	0.37	0.66	0.44	.55	0.21	0.45
Shirt / Forward	**1.77**	0.11	0.45	0.14	.19	0.08	0.14
Shirt / Raise	**1.08**	0.10	0.37	0.07	.14	0.02	0.19
Shirt / Bend	**1.60**	0.38	0.75	0.27	.64	0.32	0.47

TABLE 6.2: Performance summary of manually-created designs. We report the energy density of the garment relative to the fully dense design. Note that the automatic design has the highest energy densities across all motions. * Denotes designs by expert users.

and *Target Placement: {Forward, Raise, Bend, Flex, Ext}*. Each target is placed to elicit a specific motion from the user, and is color-coded to 4 to participants which target they should touch (see Fig.6.9). As a dependent variable, we measure *Time*, which starts automatically when the participant's hand leaves the starting position (white sphere), and ends as soon as they touch it again. The main task was to touch a given target in one continuous ballistic back-and-forth motion at a natural speed. For each target placement, three trials were collected for a total of 30 trials, one half with clutches active, the other half with clutches deactivated (Visual). The order of clutch activation was randomized and participants were not told if the clutch was on or off. The order of the feedback type was also randomized. At the end of the study, participants were free to comment on their experience using each garment design.

RESULTS The mean time to reach targets were 1.68s ($\sigma = 0.77$) for the Auto-Designed condition, 1.32s ($\sigma = 0.43$) for Manual-Designed, and 1.33s ($\sigma = 0.43$) for Visual. A longer reach time indicates more impact on the participant's ability to reach the target. The full results are visualized in Figure 6.10. A two-way repeated-measures ANOVA resulted in a significant effect on feedback type ($F(2,5) = 29.82$, $p < .001$), target placement ($F(4,5) = 34.45$, $p < .001$) and interaction ($F(8,5) = 5.72$, $p = .004$). We conducted a Holm-corrected post-hoc test and found significant differences for feedback type. Our Automatic Design method significantly

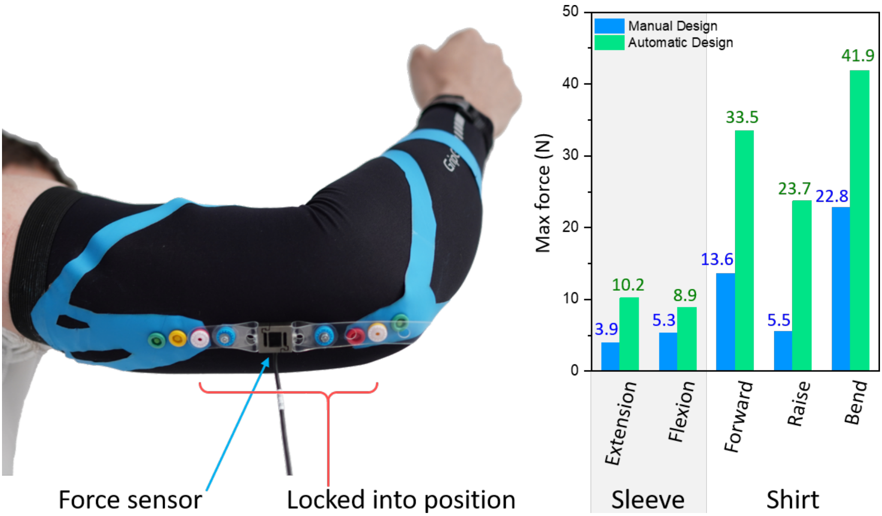

FIGURE 6.8: Physical force measurements. Left: experimental setup with force with clutch replaced by a stand-in equipped with a force sensor. Right: maximum force (N) readings when blocking different movements as labeled for the manually-designed (blue) and automatically generated (green) shirt and sleeve garments.

impacted participant movement time compared to both Manual Design feedback ($p < .001$) and Visual feedback ($p < .001$). We found no significant difference between Manual Design feedback and Visual feedback. When looking at times across target placements, our Automatic Design method significantly impacted participant movement time for the Bend and Raise motions when comparing to both Visual and Manual Design baselines (both $p < .001$).

From these results, we see a trend that the automatically designed garments performed better in terms of limiting user motion, particularly when the motions involved larger movements in the upper body.

We observe that our Auto-Designed garments performed substantially better in larger motions than the Manual-Design counterparts, results which are in-line with both simulated and force-characterization data.

The exception is the Forward motion, where we observed a less substantial impact, possibly due to the fact that participants could twist their body to reach that target. The low performance of the Flex and Ext methods could be due to the fact that we use smaller ES clutches for these motions, and the force was too small compared to the force produced by larger motions (see

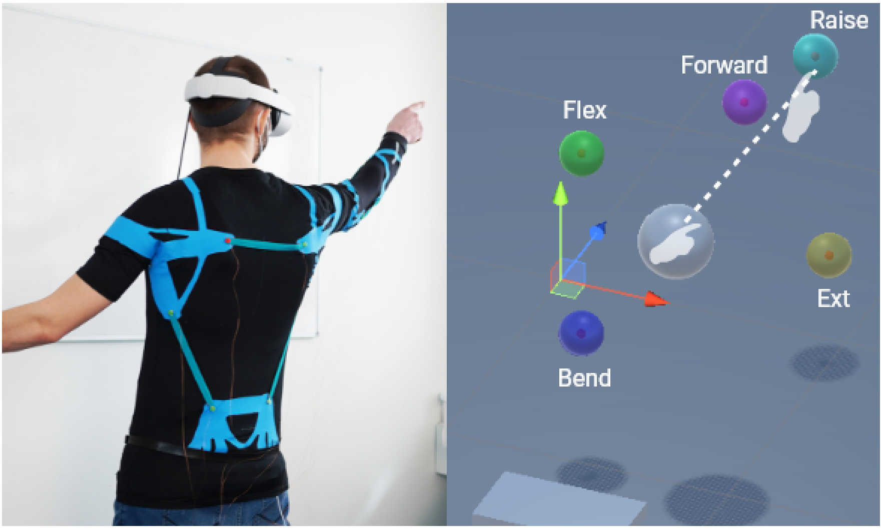

FIGURE 6.9: User study setup. Participant wearing Auto-Designed garment and reaching for target (left), and their corresponding motion in the virtual environment (right).

Fig. 6.8), and thus, below a critical threshold that would have an impact on user motion. Thus, our first hypothesis was confirmed for two of the three larger upper body motions.

What is surprising is that the performance of the Manual-Design baseline was nearly indistinguishable from the Visual baseline, even for larger motions. In relation to this, two participants commented that they had trouble perceiving any resisting effects of the Manual-Design.This shows that, even with the same active components, our optimization-based approach for designing connecting structures can indeed make the difference between a system having clear or negligible impact on user motion.

6.5 EXAMPLE APPLICATIONS

We show four applications enabled by the ability of active kinesthetic garments to selectively and dynamically engage clutches with a single design.

WORKPLACE TRAINING When picking up a virtual box, we activate all ES clutches on the shirt to provide stability to the upper-body, preventing

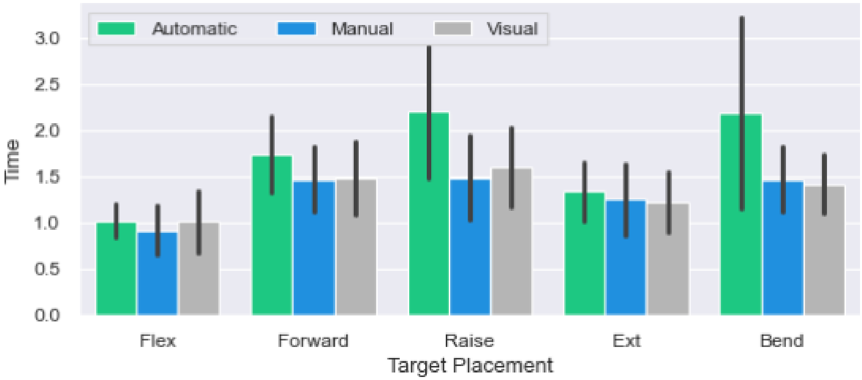

FIGURE 6.10: User study results showing the average trial time for each feedback type and target location.

the arms from going through the box, and preventing the user from overly bending their back. More complex motion control could also provide further training and guidance in combination with a complex control loop (i.e. using body pose as an input).

POSTURE CORRECTION Bad posture is a very common problem when sitting at a desk, and many posture correcting shirts already exist to help this issue. However, only active kinesthetic garments can periodically allow the user to go into a slouching posture on demand, in addition to keeping other limbs completely unrestricted (i.e. elbow).

RESISTANCE TRAINING Multiple clutches can be selectively activated to resist a target motion, the upper clutch in the case of arm extension, and the lower clutch in the case of arm flexion. The opposing clutch is meanwhile disabled, to prevent full arm-locking. This shows how a single garment can be re-configured at run-time for resisting multiple motions, potentially encompassing a user's entire workout.

VR GAMING VR immersion can be increased significantly by providing physical forces when users make contact with the world. In this game, a user practices hitting tennis balls out of the air, and only the upper back clutch is activated on contact, noting that the elbow clutch remains off and does not prevent natural elbow bending during such sports movements.

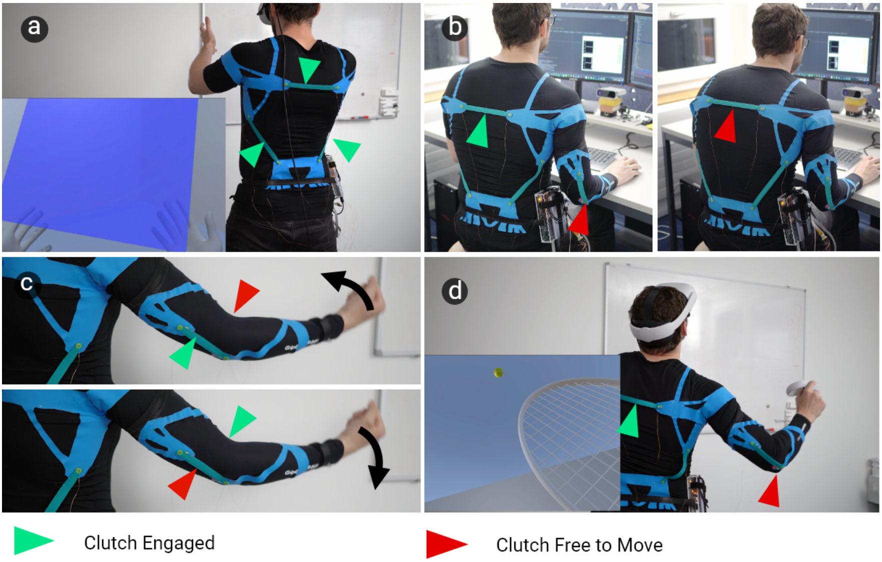

FIGURE 6.11: Applications in (a) Workplace Training, (b) Posture Correction, (c) Resistance Training, and (d) VR Gaming.

6.6 DISCUSSION AND FUTURE WORK

Our user study results indicate that automatically-designed active kinesthetic garments were able to have a significant impact on user motion, whereas the manually designed counterparts could not meet this threshold, indicating the need for automated methods to assist designers in such tasks.

EMERGENT STRUCTURAL PROPERTIES We found in our evaluation three emergent structural properties: 1) no connecting material is isolated from the main structure (no disparate island) 2) all active components are at junctures of connecting material, and 3) overlapping, yet distinct load paths are created for each specific motion. When comparing designs using our dual-objective directly to the single compliance minimization objective in [146] (b), we find that these same properties do not emerge (See Fig. 6.12). Each property plays an important role — for example, if clutches are not at junctures, then their activation will have no effect on the user. Similarly, unbalanced load paths and islands of disconnected material may degrade performance for particular motions and comfort respectively. Users on the other hand performed well in terms of connecting clutches, but struggled to balance load paths, leading to very poor performance in particular motions.

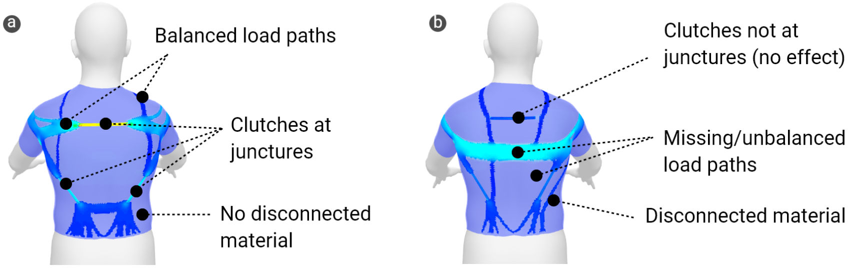

FIGURE 6.12: Emergent structural properties for the multi-motion shirt designed with dual-objective minimization (a) vs single compliance minimization from [146] (b).

LIMITATIONS The main work of the designer in our tool is in the manual placement of ES clutches. As the number of active components and motions grows, the requirement for manual clutch placement may become more and more challenging. Our method can be extended to optimize for ES clutch placement, thereby freeing designers from this task, potentially increasing the relative efficiency of the design. Our method is also limited to a simple mode of activation, where clutches are either all active or inactive. However, clutch activations can be controlled individually and continuously through voltage input that affect the degree of resistance. Accounting for these degrees of freedom during design could further improve efficiency and allow for more targeted resistance to selected motions.

While our active kinesthetic garments are fully wearable and mobile, they do not have any sensing capabilities. Integrating sensing could be done via capacitive sensors, which could be optimized based on the same strain-maximization principle used for connecting clutches.

Our study was limited to only six male participants and the type of feedback collected was mainly quantitative. Our method could be used to create garments for female users and even *personalized* garments by simply changing the β and gender parameters in the SMPL/STAR models. Richer VR interaction opportunities can be explored in the future by moving beyond simple button presses and object intersections, for example, by integrating body-pose sensing into the control loop.

6.7 SUMMARY AND OUTLOOK

In this chapter, we presented a computational approach for automatic design of active kinesthetic garments that resist user-specified body motion on-demand. As our core technical contribution, we cast the design of reinforcing structures that connect and anchor individual clutches as an on-body topology optimization problem and introduced a novel objective term that encourages maximum resistance of the garment when clutches are active while minimizing interference with body motion when they are inactive. Our experiments indicate that our designs are highly effective and consistently and significantly outperform user-created designs.

The structure optimization techniques developed here have the potential to be useful in the routing and placement of other types of active components such as actuators and sensors. By laying out a theoretical and algorithmic basis for this central problem, we hope that our work will serve as a step toward computational design of highly integrated multi-modal wearable interfaces in the future.

7

CONCLUSION

7.1 SUMMARY

The focus of this thesis is to advance the state-of-the art in wearable haptics. Wearable technologies in general are advancing at a rapid pace, and they pose a unique challenge - the balancing of *performance* and *wearability* factors, which are often not in alignment. In this thesis, we took a two-fold approach to this problem - first by moving towards a *compliant first* hardware design, and second by leveraging computational methods to enable a more foundational approach to balancing these factors through high-level objectives. The two approaches are complementary and necessary: infinitely compliant and performant haptic actuators (both kinesthetic and tactile) are upper-bounded by their grounding and force exertion onto the body, while optimizing the layout and connections of existing rigid actuators with stiff materials will likely lead to discomfort or impairment (limitation of movement) regardless of placement. We therefore make contributions to both the *fabrication* and *design* challenges posed by wearable haptics.

In the first two chapters (3,4), we developed two hand-worn systems for Tactile and Kinesthetic feedback in VR, respectively. Both systems take a compliant-first approach, using a textile glove as a soft-substrate. The main contributions of the tactile glove in chapter 3 are the miniaturized bi-stable pin type actuators leveraging electromagnetic latching. These actuators further enabled 1) an increased density of actuators to saturate the Palmar (inner) side of the hand, and 2) two modes of operation - via low (1Hz) and high (up to 200Hz) frequencies. The ability to render a wide range of frequencies while sufficiently stimulating many distinct areas of the hand opened up new abilities - users were able to distinguish textures, localize objects in the virtual space using only their tactile sense.

A much more difficult problem in hand-worn VR haptics has been how to physically impede finger flexion when making contact with virtual objects. This physical constraint is at the heart of many crucial daily manipulations, from grasping large objects, to fine grained manipulation of smaller objects, and ultra-precise manipulation of tools in critical tasks such as surgery. Chapter 4 contributes the first-known *on-hand* compliant braking mechanism based on electrostatic clutching. Each clutch (or brake) is only

8g while blocking up to 20N, and is crucially made of a flexible steel substrate. Because of these properties, the glove is able to provide resistance (in VR) in particular grasps while allowing the hand to assume a wide range of poses. This work presented us to first challenge of integrating such high flexibility while balancing for performance, that was addressed through a re-configurable system using 3D fabricated parts and simple Velcro attachments. This work allowed us to gain insights and understand the limitations of a manual approach of integrating flexible haptics onto a flexible and highly dexterous areas of the body and was the seed that inspired our follow-up work.

In the last two chapters (5,6), we investigated an often overlooked aspect of designing wearable haptics - the design effort itself. In contrast to manual design approaches, we take a physical modeling perspective in combination with FEM simulation, allowing more rapid iteration and opening up inverse design opportunities. We can split these two works into their respective targets - chapter 5 aims to design *Passive* wearable kinesthetic haptics, while chapter 6 extends this to *Active* kinesthetic haptics. In chapter 5, we first develop the notion of *Kinesthetic Garments* - these are garments that are specifically designed to resist a given motion. While such garments already exist (e.g. back-support shirts that force users in an upright pose), we generalize this to any body site and any type of motion. The main contribution in this chapter are the framing of computationally designing the most performant garments by using topology optimization (TPO). This idea, similar to standard TPO, is to remove material from low-energy areas of the garment, thus increasing the energy-density of the garments for particular motions, which in-turn increases the effort that the user must expend to perform it. On-body TPO thus allows designers to specify only their high-level objectives (motion, body-site), and get back a fully designed garment *automatically*. We further leveraged the properties of our evolutionary TPO algorithm to enable designers to pick from designs with different amounts of material. Kinesthetic garments were shown to perform better (preferred by users) across a range of garments, when compared to baseline designs.

Chapter 6 brings together the contributions of compliant electrostatic (ES) clutches developed in chapter (4) with the physical modeling design approach in chapter 5. The main contribution of this chapter are the embedding of such active components into the overall physical model of kinesthetic garments and leveraging topology optimization in order to design efficient *connecting and grounding* structures for these components.

We develop a modified dual-objective function that takes the state of ES clutches into account - when they are off and when they are on. This provides a timely opportunity - the ability of our method to design for both performance and wearability. In their off-state, we *minimize* the energy of the clutches and connecting structures, while in the on-state, we *maximize* it. This results in a garment that is comfortable to wear and does not otherwise resist motion, until it is activated on-demand, where it is optimized to resist motion as much as possible. Notably, the active braking ability of the garments also enabled the optimization with respect to different motions, resulting in a connecting structure that contains multiple load paths that are shared between motions. When compared to user-designed variants of the same garments, our automatically designed garments could be up to 10x more efficient in simulation, and were shown to resist up to 4 times more force. Most surprisingly, in our final user VR study, we found no significant differences between the feedback provided by user-designed garments and no feedback at all, while our automatically designed garments were found to be significantly better than no-feedback and feedback via manually-designed garments for 2 of the 4 motions tested. This indicates that not only are automatically designed kinesthetic garments more performant than manual ones, but that using modeling and principled automatic design approaches may actually be required to reach performance levels that users can feel. In other words, the ES clutches, despite having very high performance themselves, were not the limiting factor in this scenario. Rather, the design of the connecting and grounding structure played a more significant role.

7.2 OUTLOOK AND FUTURE WORK

The work presented here may serve as inspiration for a number of directions. Here we will explore several of these directions, from more direct follow-up work, to gaining insight into applying our approaches beyond haptics.

7.2.1 *Computational Haptics*

The computational frameworks developed in this thesis are largely oriented towards the optimization of *kinesthetic* criteria that are closely aligned with strain energy maximization. One aspect that can be developed further is the optimization of *cutaneous* or tactile haptics that would require developing new types of body models. Examples of such models would include soft-tissue outer layers [51, 52] that could deform with applied forces. A model that takes into account the rendering aspect of cutaneous actuators [161] can benefit the design of the tactile actuator itself, for example, by optimizing the angles and form-factor required to render particular stimuli. Furthermore, cutaneous and kinesthetic actuators can be modeled in the same system, to maximize not only strain energy, but the intended perception of users. Such a model would require linking two disparate areas - of perception and physical modeling, a difficult but perhaps highly rewarding direction. For example, when resisting arm flexion, rather than maximizing for strain energy density in a kinesthetic garment arm-sleeve, we can maximize the *torque* induced on the muscle fibers in the biceps. Other optimization criteria can also include pressure and pressure distribution [14], which are linked to particular perceptions and bring cutaneous and kinesthetic haptic models together. Computational haptics as a generalized approach can also benefit from modeling *control* and understanding what a user might do in the near time horizon, for example through delivering feedback that guides users rather than forcing them along paths [162, 163]. Applying this approach to our automatic design algorithms where we currently assume that all clutches are on or off at the same time, we can remove this limitation and instead attempt to move the user towards a particular motion or pose by turning on only specific clutches at specific times.

7.2.2 *Combining with Sensors and Control*

We presented our work with an eye towards wearability in general, as the methods and approaches used to create our haptic devices and methods are readily transferable to other wearability spaces. In particular, body-mounted sensing systems can be found everywhere, and run into largely the same wearability and performance requirements as haptics. For example, maximizing strain-energy in body-mounted clutches could benefit capacitive strain sensors to induce a more clear signal [164]. The vision for robotic garments [13] based on compliant textiles [12] outlines the need for sensors and actuators working in tandem to assist and enhance users in a variety of tasks. Remarkably, ES clutches have recently been demonstrated to have the ability to sense their configuration [149], thus allowing their use as *both* brakes and sensors. A computational approach that combines both sensing and actuation and manages to align their respective performance and wearability constraints may be highly beneficial across a number of applications from [31], from enhancing our sports performance to assisting elderly mobility.

7.2.3 *Wearability-for-All*

As our last point of future work, we move away from academia and explore how industry might approach wearability from a product perspective. While personalization is an often discussed future path for many wearable technologies, it requires a large effort on the part of the user, and is not really feasible to actually deliver in large quantities and at a reasonable cost. A rather different approach may be undertaken - where the effects of wearable topologies and designs on performance and comfort may be investigated across large populations. A simple example would be the Apple Watch strap, which is made of particular materials and has a range of designs that must be made to *fit*, be comfortable, and have enough contact with the skin to measure the user's heart rate and other data (performance consideration). How many strap variations are needed, how many types of materials, and how many sizes should Apple produce in order to minimize the number variations, while covering as much of the population as possible? Similar questions can also be asked about in-ear audio devices regarding their size and form-factor and physical effects on the ear canal, VR and AR headsets regarding their strap designs, and almost any wearable technologies produced for the mass-market. A principled computational approach to modeling such wearables on the body allows the rapid investigation of design variables across such populations, and as shown in this thesis, potentially also opens inverse design opportunities. An example of such an approach could, instead of producing simple sizing options for wearables, produce form-factors and shapes that best capture the variance in the population itself, an optimization in the PCA space of body shapes and physical properties. Such an approach would also necessitate the gathering of large datasets of users and the topology and soft-tissue behaviors of the respective body sites [52]. There is still a long way to go in the creation of accurate simulations of human bodies in action in a way that truly captures the variances across populations, and the corresponding design space of most wearable for sensing, actuation and other modalities remains fairly open for exploration. We leave this exciting work for future explorers.

A

APPENDIX

A.1 PROCEDURAL CLUTCH GENERATION

ES clutches are defined by the following variables: a *starting point*, a *surface direction*, and a *length and width*. The *starting point* is defined using barycentric coordinates (u_e, v_e) on a particular element e of the garment mesh. The *surface direction* is a vector in barycentric space $(\overrightarrow{u_e}, \overrightarrow{v_e})$ from the starting point to another barycentric coordinate on the same element e.

We start by tracing out a piece-wise linear path of the desired length in the direction of $(\overrightarrow{u_e}, \overrightarrow{v_e})$ until an edge is encountered, whereby the vector is converted to Euclidean space $\in \mathbb{R}^{3n}$ and rotated to lie on the surface of the next triangle e_i. This is repeated until the length of the vector is exhausted. Two endpoints are produced, one at the starting point, and one at the last barycentric coordinate of where the path finishes.

From this path, a mesh is triangulated by creating center vertices at edge intersections and projecting side vertices to the left and right of the path based on $\overrightarrow{e_n} \times \overrightarrow{p_{xyz}}$, the cross product of the element normal and the path direction in world space respectively. This is scaled by the *width* parameter. The resulting mesh has rest vertices $\mathbf{q} = (\mathbf{q}_1, \ldots, \mathbf{q}_m) \in \mathbb{R}^{3n}$ and $\bar{\mathbf{q}} = (\bar{\mathbf{q}}_1, \ldots, \bar{\mathbf{q}}_m) \in \mathbb{R}^{3n}$ when deformed. We give special treatment to the side vertices of the two endpoints by walking them in an orthogonal direction to the main path using the same walking algorithm outlined above.

The endpoints q^c of the clutch mesh (3 at each end) are connected to the garment mesh using simple quadratic penalty functions, which allows for firm attachment. The full path walking and meshing algorithm is fast enough to work in real-time, allowing for rapid user placement and repositioning of ES clutches.

A.2 GARMENT-ON-BODY MODEL

As our garment model, we use a compressible neo-Hookean material model [32] adapted with a relaxed energy under wrinkling as in [146]. This allows the garment to wrinkle under compression without producing geometric artifacts. This results in the garment energy $E_{\text{garment}}(\mathbf{x}, d)$, which is a function of the garment design d, and the deformed nodal positions $\mathbf{x}$. We similarly convert the discrete body mesh to a continuous implicit signed distance field [137], resulting in the energy $E_{\text{body}}(v, \mathbf{x})$, which pushes back on the garment vertices $\mathbf{x}$ away from the body. This allows the garment to smoothly slide on top of the body and to lift-off from its surface. To attach the garment to the body in specific areas, we introduce a simple coupling potential, $E_{\text{attach}} = \frac{1}{2}k(x^c - x_c^v)^T(x^c - x_c^v)$, attracting elements of the garment mesh x^c to corresponding elements v_c on the body mesh. As the garment mesh is initialized from the SMPL mesh, for more accurate simulation, we subdivide the garment mesh until it has 16x the resolution of the base SMPL template mesh.

A.3 ON-BODY TOPOLOGY OPTIMIZATION

To design passive reinforcement structures, Vechev et al. use a bi-directional evolutionary structural optimization (BESO) algorithm [39, 43] to solve the constrained optimization problem with a single objective,

$$\begin{aligned} \mathbf{d}^* = \arg\max_{\mathbf{d}} \quad & E_{\text{garment}}(x^*, \mathbf{d}) \\ \text{s.t.} \quad & \sum_e A^e d^e = A^* \ , \quad \mathbf{f}(x^*) = \mathbf{0}. \end{aligned} \tag{A.1}$$

The goal of this formulation is to find an optimal per-element material assignment $\mathbf{d}^*$ that maximizes the energy of the garment in its equilibrium state $\mathbf{x}^*$ while satisfying constraints on force equilibrium, $\mathbf{f}(x^*) = \mathbf{0}$, and material budget, $\sum_e A_e d_e = A^*$. The strain energy of the garment is defined per element as

$$E_{\text{garment}} = \sum_e t^e A^e W^e_{\text{garment}}(x^*, d^e) \ , \tag{A.2}$$

where W^e_{garment} is the elemental strain energy density, and t^e, A^e are the thickness and area of the element respectively.

BIBLIOGRAPHY

1. Huang, Y., Kaufmann, M., Aksan, E., Black, M. J., Hilliges, O. & Pons-Moll, G. Deep inertial poser: Learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Transactions on Graphics (TOG)* **37**, 1 (2018).
2. Glauser, O., Wu, S., Panozzo, D., Hilliges, O. & Sorkine-Hornung, O. Interactive hand pose estimation using a stretch-sensing soft glove. *ACM Transactions on Graphics (TOG)* **38**, 41 (2019).
3. Zarate, J. J., Gudozhnik, O., Ruch, A. S. & Shea, H. *Keep in Touch: Portable Haptic Display with 192 High Speed Taxels* in *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (ACM, Denver, Colorado, USA, 2017), 349.
4. White, O., McIntyre, J., Augurelle, A.-S. & Thonnard, J.-L. Do novel gravitational environments alter the grip-force/load-force coupling at the fingertips? *Experimental brain research* **163**, 324 (2005).
5. Gibson, J. J. Observations on active touch. *Psychological review* **69**, 477 (1962).
6. Günther, S., Makhija, M., Müller, F., Schön, D., Mühlhäuser, M. & Funk, M. *PneumAct: Pneumatic Kinesthetic Actuation of Body Joints in Virtual Reality Environments* in *Proceedings of the 2019 on Designing Interactive Systems Conference* (2019), 227.
7. Rognon, C., Mintchev, S., Dell'Agnola, F., Cherpillod, A., Atienza, D. & Floreano, D. Flyjacket: An upper body soft exoskeleton for immersive drone control. *IEEE Robotics and Automation Letters* **3**, 2362 (2018).
8. Benko, H., Holz, C., Sinclair, M. & Ofek, E. *NormalTouch and TextureTouch: High-fidelity 3D Haptic Shape Rendering on Handheld Virtual Reality Controllers* in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (ACM, Tokyo, Japan, 2016), 717.
9. Kim, J., Lee, G., Heimgartner, R., Revi, D. A., Karavas, N., Nathanson, D., Galiana, I., Eckert-Erdheim, A., Murphy, P., Perry, D., *et al.* Reducing the metabolic rate of walking and running with a versatile, portable exosuit. *Science* **365**, 668 (2019).

10. Lee, S., Karavas, N., Quinlivan, B. T., LouiseRyan, D., Perry, D., Eckert-Erdheim, A., Murphy, P., Goldy, T. G., Menard, N., Athanassiu, M., *et al. Autonomous multi-joint soft exosuit for assistance with walking overground* in *2018 IEEE International Conference on Robotics and Automation (ICRA)* (2018), 2812.
11. Pacchierotti, C., Sinclair, S., Solazzi, M., Frisoli, A., Hayward, V. & Prattichizzo, D. Wearable haptic systems for the fingertip and the hand: Taxonomy, review, and perspectives. *IEEE transactions on haptics* **10**, 580 (2017).
12. Yin, J., Hinchet, R., Shea, H. & Majidi, C. Wearable soft technologies for haptic sensing and feedback. *Advanced Functional Materials* **31**, 2007428 (2021).
13. Sanchez, V., Walsh, C. J. & Wood, R. J. Textile technology for soft robotic and autonomous garments. *Advanced Functional Materials* **31**, 2008278 (2021).
14. Montes, J., Thomaszewski, B., Mudur, S. & Popa, T. Computational design of skintight clothing. *ACM Transactions on Graphics (TOG)* **39**, 105 (2020).
15. Liu, Z., Han, X., Zhang, Y., Chen, X., Lai, Y.-K., Doubrovski, E. L., Whiting, E. & Wang, C. C. L. Knitting 4D Garments with Elasticity Controlled for Body Motion. *ACM Trans. Graph.* **40** (2021).
16. Vechev, V., Zarate, J., Lindlbauer, D., Hinchet, R., Shea, H. & Hilliges, O. *Tactiles: Dual-mode low-power electromagnetic actuators for rendering continuous contact and spatial haptic patterns in vr* in *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2019), 312.
17. Hinchet, R., Vechev, V., Shea, H. & Hilliges, O. *Dextres: Wearable haptic feedback for grasping in vr via a thin form-factor electrostatic brake* in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (2018), 901.
18. Vechev, V., Zarate, J., Thomaszewski, B. & Hilliges, O. *Computational Design of Kinesthetic Garments* in *Computer Graphics Forum* **41** (2022), 535.
19. Vechev, V., Hinchet, R., Coros, S., Thomaszewski, B. & Hilliges, O. *Computational Design of Active Kinesthetic Garments* in *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (2022), 1.

20. Culbertson, H., Schorr, S. B. & Okamura, A. M. Haptics: The Present and Future of Artificial Touch Sensation. *Annual Review of Control, Robotics, and Autonomous Systems* **1**, 385 (2018).
21. Hayward, V., Astley, O. R., Cruz-Hernandez, M., Grant, D. & Robles-De-La-Torre, G. Haptic interfaces and devices. *Sensor Review* **24**, 16 (2004).
22. Johansson, R. S. & Flanagan, J. R. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience* **10**, 345 (2009).
23. Scilingo, E. P., Bianchi, M., Grioli, G. & Bicchi, A. Rendering softness: Integration of kinesthetic and cutaneous information in a haptic device. *IEEE Transactions on Haptics*, 109 (2010).
24. Lederman, S. J. & Klatzky, R. L. Hand movements: A window into haptic object recognition. *Cognitive psychology* **19**, 342 (1987).
25. Heller, M. A. Active and passive touch: The influence of exploration time on form recognition. *The Journal of general psychology* **110**, 243 (1984).
26. Klatzky, R. L., Loomis, J. M., Lederman, S. J., Wake, H. & Fujita, N. Haptic identification of objects and their depictions. *Perception & psychophysics* **54**, 170 (1993).
27. Pont, S. C., Kappers, A. M. & Koenderink, J. J. Haptic curvature discrimination at several regions of the hand. *Perception & Psychophysics* **59**, 1225 (1997).
28. Follmer, S., Leithinger, D., Olwal, A., Hogge, A. & Ishii, H. *inFORM: Dynamic Physical Affordances and Constraints Through Shape and Object Actuation* in *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (ACM, St. Andrews, Scotland, United Kingdom, 2013), 417.
29. Choi, I., Hawkes, E. W., Christensen, D. L., Ploch, C. J. & Follmer, S. *Wolverine: A wearable haptic interface for grasping in virtual reality* in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on* (2016), 986.
30. Quinlivan, B. T., Lee, S., Malcolm, P., Rossi, D. M., Grimmer, M., Siviy, C., Karavas, N., Wagner, D., Asbeck, A., Galiana, I., *et al.* Assistance magnitude versus metabolic cost reductions for a tethered multiarticular soft exosuit. *Science robotics* **2**, eaah4416 (2017).

31. Shi, Y., Dong, W., Lin, W. & Gao, Y. Soft Wearable Robots: Development Status and Technical Challenges. *Sensors* **22**, 7584 (2022).
32. Bonet, J. & Wood, R. D. *Nonlinear Continuum Mechanics for Finite Element Analysis* 2nd ed. (Cambridge University Press, 2008).
33. Nocedal, J. & Wright, S. J. *Numerical optimization* (Springer, 1999).
34. Skouras, M., Thomaszewski, B., Coros, S., Bickel, B. & Gross, M. Computational design of actuated deformable characters. *ACM Transactions on Graphics (TOG)* **32**, 1 (2013).
35. Umetani, N., Kaufman, D. M., Igarashi, T. & Grinspun, E. Sensitive Couture for Interactive Garment Modeling and Editing. *ACM Trans. Graph.* **30** (2011).
36. Bendsoe, M. P. & Sigmund, O. *Topology optimization: theory, methods, and applications* (Springer Science & Business Media, 2013).
37. Duysinx, P. & Bendsøe, M. P. Topology optimization of continuum structures with local stress constraints. *International journal for numerical methods in engineering* **43**, 1453 (1998).
38. Deng, H., Vulimiri, P. S. & To, A. C. An efficient 146-line 3D sensitivity analysis code of stress-based topology optimization written in MATLAB. *Optimization and Engineering*, 1 (2021).
39. Huang, X. & Xie, Y. Convergent and mesh-independent solutions for the bi-directional evolutionary structural optimization method. *Finite elements in analysis and design* **43**, 1039 (2007).
40. Aremu, A., Ashcroft, I., Hague, R., Wildman, R. & Tuck, C. *Suitability of SIMP and BESO topology optimization algorithms for additive manufacture* in *2010 International Solid Freeform Fabrication Symposium* (2010).
41. Huang, X. & Xie, Y.-M. A further review of ESO type methods for topology optimization. *Structural and Multidisciplinary Optimization* **41**, 671 (2010).
42. Edwards, C., Kim, H. & Budd, C. An evaluative study on ESO and SIMP for optimising a cantilever tie—beam. *Structural and Multidisciplinary Optimization* **34**, 403 (2007).
43. Huang, X. & Xie, Y. M. Bi-directional evolutionary topology optimization of continuum structures with one or multiple materials. *Computational Mechanics* **43**, 393 (2009).

44. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G. & Black, M. J. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* **34**, 248 (2015).

45. Romero, J., Tzionas, D. & Black, M. J. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (TOG)* **36**, 245 (2017).

46. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A. A., Tzionas, D. & Black, M. J. *Expressive body capture: 3d hands, face, and body from a single image* in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2019), 10975.

47. Kingma, D. P. & Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).

48. Osman, A. A., Bolkart, T. & Black, M. J. *Star: Sparse trained articulated human body regressor* in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16* (2020), 598.

49. Pons-Moll, G., Romero, J., Mahmood, N. & Black, M. J. Dyna: A model of dynamic human shape in motion. *ACM Transactions on Graphics (TOG)* **34**, 120 (2015).

50. Casas, D. & Otaduy, M. A. Learning nonlinear soft-tissue dynamics for interactive avatars. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* **1**, 10 (2018).

51. Kim, M., Pons-Moll, G., Pujades, S., Bang, S., Kim, J., Black, M. J. & Lee, S.-H. Data-driven physics for human soft tissue animation. *ACM Transactions on Graphics (TOG)* **36**, 54 (2017).

52. Pai, D. K., Rothwell, A., Wyder-Hodge, P., Wick, A., Fan, Y., Larionov, E., Harrison, D., Neog, D. R. & Shing, C. The human touch: measuring contact with real human soft tissues. *ACM Transactions on Graphics (TOG)* **37**, 58 (2018).

53. Hoffman, H. G. *Physically Touching Virtual Objects Using Tactile Augmentation Enhances the Realism of Virtual Environments* in *Proceedings of the Virtual Reality Annual International Symposium* (IEEE Computer Society, Washington, DC, USA, 1998), 59.

54. Hinchet, R., Vechev, V., Shea, H. & Hilliges, O. *DextrES: Wearable Haptic Feedback for Grasping in VR via a Thin Form-Factor Electrostatic Brake* in *Proceedings of the ACM Symposium on User Interface Software and Technologies (UIST)* (ACM, Berlin, Germany, 2018).

55. Hummel, J., Dodiya, J., Center, G. A., Eckardt, L., Wolff, R., Gerndt, A. & Kuhlen, T. W. *A lightweight electrotactile feedback device for grasp improvement in immersive virtual environments* in *2016 IEEE Virtual Reality (VR)* (2016), 39.
56. CyberGlove Systems Inc. *CyberTouch Glove* `http://www.cyberglovesystems.com/cybertouch`. Last accessed: 18.03.2017.
57. Schorr, S. B. & Okamura, A. M. *Fingertip tactile devices for virtual object manipulation and exploration* in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), 3115.
58. Son, B. & Park, J. *Tactile Sensitivity to Distributed Patterns in a Palm* in *Proceedings of the 20th ACM International Conference on Multimodal Interaction* (ACM, Boulder, CO, USA, 2018), 486.
59. HaptX. *HaptX Glove* `https://www.haptx.com`. Last accessed: 03.05.2018.
60. Pece, F., Zarate, J. J., Vechev, V., Besse, N., Gudozhnik, O., Shea, H. & Hilliges, O. *MagTics: Flexible and Thin Form Factor Magnetic Actuators for Dynamic and Wearable Haptic Feedback* in *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (2017), 143.
61. Withana, A., Groeger, D. & Steimle, J. *Tacttoo: A Thin and Feel-Through Tattoo for On-Skin Tactile Output* in *Proceedings of the ACM Symposium on User Interface Software and Technology* (2018).
62. Bermejo, C. & Hui, P. A survey on haptic technologies for mobile augmented reality. *arXiv preprint arXiv:1709.00698* (2017).
63. Moehring, M. & Froehlich, B. *Effective manipulation of virtual objects within arm's reach* in *Virtual Reality Conference (VR), 2011 IEEE* (2011), 131.
64. NeuroDigital Technologies. *Gloveone Glove* `https://www.neurodigital.es/gloveone/`. Last accessed: 29.03.2017.
65. Plexus. *Plexus Gloves* `http://plexus.im/`. Last accessed: 29.03.2018.
66. Rekimoto, J. *Traxion: a tactile interaction device with virtual force sensation* in *Proceedings of the 26th annual ACM symposium on User interface software and technology* (2013), 427.
67. Kim, H., Yi, H., Lee, H. & Lee, W. *HapCube: A Wearable Tactile Device to Provide Tangential and Normal Pseudo-Force Feedback on a Fingertip* in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018), 501.

68. Culbertson, H., Walker, J. M., Raitor, M. & Okamura, A. M. *WAVES: a wearable asymmetric vibration excitation system for presenting three-dimensional translation and rotation cues* in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), 4972.
69. Johansson, R. S. & Vallbo, A. B. Tactile sensory coding in the glabrous skin of the human hand. *Trends in Neurosciences* **6**, 27 (1983).
70. Whitmire, E., Benko, H., Holz, C., Ofek, E. & Sinclair, M. *Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller* in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018), 86.
71. Yem, V. & Kajimoto, H. *Wearable tactile device using mechanical and electrical stimulation for fingertip interaction with virtual world* in *Virtual Reality (VR), 2017 IEEE* (2017), 99.
72. Kim, H., Kim, M. & Lee, W. *HapThimble: A Wearable Haptic Device Towards Usable Virtual Touch Screen* in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (ACM, Santa Clara, California, USA, 2016), 3694.
73. Yang, T.-H., Kim, S.-Y., Kim, C. H., Kwon, D.-S. & Book, W. J. Development of a miniature pin-array tactile module using elastic and electromagnetic force for mobile devices (2009).
74. Duvernoy, B., Farkhatdinov, I., Topp, S. & Hayward, V. *Electromagnetic Actuator for Tactile Communication* in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications* (2018), 14.
75. Dagnall, H. *Exploring Surface Texture: A Fundamental Guide to the Measurement of Surface Finish* (Taylor Hobson, 2003).
76. Fu, W., van Paassen, M. M. & Mulder, M. *The Influence of Discrimination Strategy on the JND in Human Haptic Perception of Manipulator Stiffness* in *AIAA Modeling and Simulation Technologies Conference* (2017), 3668.
77. Cornsweet, T. N. The staircase-method in psychophysics. *The American journal of psychology* **75**, 485 (1962).
78. Vallbo, A. B. & Johansson, R. S. Properties of cutaneous mechanoreceptors in the human hand-related to touch sensation. *Human Neurobiology* **6**, 3 (1984).
79. Jung, J., Youn, E. & Lee, G. *PinPad: touchpad interaction with fast and high-resolution tactile output* in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), 2416.

80. Klatzky, R. L., Lederman, S. J. & Metzger, V. A. Identifying objects by touch: An "expert system". *Perception & psychophysics* **37**, 299 (1985).
81. Feix, T., Romero, J., Schmiedmayer, H.-B., Dollar, A. M. & Kragic, D. The grasp taxonomy of human grasp types. *IEEE Transactions on Human-Machine Systems* **46**, 66 (2016).
82. Pacchierotti, C., Chinello, F., Malvezzi, M., Meli, L. & Prattichizzo, D. *Two finger grasping simulation with cutaneous and kinesthetic force feedback* in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications* (2012), 373.
83. CyberGlove Systems Inc. *CyberGrasp Glove* `http://www.cyberglovesystems.com/cybergrasp`. Last accessed: 12.03.2017.
84. Gu, X., Zhang, Y., Sun, W., Bian, Y., Zhou, D. & Kristensson, P. O. *Dexmo: An inexpensive and lightweight mechanical exoskeleton for motion capture and force feedback in VR* in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), 1991.
85. ContactCI. *Maestro Glove* `http://www.contactci.com`. Last accessed: 03.05.2018.
86. Choi, I., Culbertson, H., Miller, M. R., Olwal, A. & Follmer, S. *Grabity: A Wearable Haptic Interface for Simulating Weight and Grasping in Virtual Reality* in *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (2017), 119.
87. Minamizawa, K., Fukamachi, S., Kajimoto, H., Kawakami, N. & Tachi, S. *Gravity grabber: wearable haptic display to present virtual mass sensation* in *ACM SIGGRAPH 2007 emerging technologies* (2007), 8.
88. Massie, T. H., Salisbury, J. K., *et al. 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* **55** (1994), 295.
89. Araujo, B., Jota, R., Perumal, V., Yao, J. X., Singh, K. & Wigdor, D. *Snake Charmer: Physically Enabling Virtual Objects* in *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction* (ACM, Eindhoven, Netherlands, 2016), 218.
90. Festo. *HexoHand Glove* `https://www.festo.com/group/en/cms/10233.htm`. Last accessed: 04.04.2018.
91. Bouzit, M., Burdea, G., Popescu, G. & Boian, R. The Rutgers Master II-new design force-feedback glove. *IEEE/ASME Transactions on Mechatronics* **7**, 256 (2002).

92. Zubrycki, I. & Granosik, G. Novel haptic device using jamming principle for providing kinaesthetic feedback in glove-based control interface. *Journal of Intelligent & Robotic Systems* **85**, 413 (2017).
93. Choi, I., Corson, N., Peiros, L., Hawkes, E. W., Keller, S. & Follmer, S. A Soft, Controllable, High Force Density Linear Brake Utilizing Layer Jamming. *IEEE Robotics and Automation Letters* **3**, 450 (2018).
94. Blake, J. & Gurocak, H. B. Haptic glove with MR brakes for virtual reality. *IEEE/ASME Transactions On Mechatronics* **14**, 606 (2009).
95. Winter, S. H. & Bouzit, M. Use of Magnetorheological Fluid in a Force Feedback Glove. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **15**, 2 (2007).
96. Lelieveld, M. J., Maeno, T. & Tomiyama, T. *Design and development of two concepts for a 4 DOF portable haptic interface with active and passive multi-point force feedback for the index finger* in *ASME 2006 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference* (2006), 547.
97. Sense Glove BV. *Sense Glove* `https://www.senseglove.com`. Last accessed: 03.05.2018.
98. Plooij, M., Mathijssen, G., Cherelle, P., Lefeber, D. & Vanderborght, B. Lock Your Robot: A Review of Locking Devices in Robotics. *IEEE Robotics & Automation Magazine* **22**, 106 (2015).
99. Wang, D. M., Hou, Y. F. & Tian, Z. Z. A novel high-torque magnetorheological brake with a water cooling method for heat dissipation. *Smart Materials and Structures* **22**, 025019 (2013).
100. Saito, T. & Ikeda, H. Development of Normally Closed Type of Magnetorheological Clutch and its Application to Safe Torque Control System of Human-Collaborative Robot. *Journal of Intelligent Material Systems and Structures* **18**, 1181 (2007).
101. Shafer, A. S. & Kermani, M. R. *Design and validation of a Magneto-Rheological clutch for practical control applications in human-friendly manipulation* in *2011 IEEE International Conference on Robotics and Automation* (IEEE, 2011), 4266.
102. Johnsen, A. & Rahbek, K. A physical phenomenon and its applications to telegraphy, telephony, etc. *Journal of the Institution of Electrical Engineers* **61**, 713 (1923).
103. Fitch, C. J. Development of the Electrostatic Clutch. *IBM Journal of Research and Development* **1**, 49 (1957).

104. Niino, T., Egawa, S. & Higuchi, T. *High-power and high-efficiency electrostatic actuator* in *[1993] Proceedings IEEE Micro Electro Mechanical Systems* (IEEE, 1993), 236.

105. Niino, T., Egawa, S., Nishiguchi, N. & Higuchi, T. *Development of an electrostatic actuator exceeding 10 N propulsive force* in *[1992] Proceedings IEEE Micro Electro Mechanical Systems* (IEEE, 1992), 122.

106. Diller, S., Majidi, C. & Collins, S. H. *A lightweight, low-power electroadhesive clutch and spring for exoskeleton actuation* in *2016 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2016), 682.

107. Aukes, D. M., Heyneman, B., Ulmen, J., Stuart, H., Cutkosky, M. R., Kim, S., Garcia, P. & Edsinger, A. Design and testing of a selectively compliant underactuated hand. *The International Journal of Robotics Research* **33**, 721 (2014).

108. Karagozler, M. E., Campbell, J. D., Fedder, G. K., Goldstein, S. C., Weller, M. P. & Yoon, B. W. *Electrostatic latching for inter-module adhesion, power transfer, and communication in modular robots* in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2007), 2779.

109. Prahlad, H., Pelrine, R., Stanford, S., Marlow, J. & Kornbluh, R. *Electroadhesive robots—wall climbing robots enabled by a novel, robust, and electrically controllable adhesion technology* in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (2008), 3028.

110. Choi, I., Ofek, E., Benko, H., Sinclair, M. & Holz, C. *CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality* in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018), 654.

111. ElKoura, G. & Singh, K. *Handrix: animating the human hand* in *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2003), 110.

112. PetapicoVoltron. *Peta-pico-Voltron Project* `http://petapicovoltron.com/`. Last accessed: 12.07.2018.

113. Samyn, P., Quintelier, J., Schoukens, G. & De Baets, P. *Sliding properties of polyimide against various steel and DLC-coated counterfaces* in *Proceedings of of BALKANTRIB 05* (ed Ivkovic, B.) (Kragujevac, Serbia and Montenegro, 2005), 506.

114. Liu, Z., Han, X., Zhang, Y., Chen, X., Lai, Y.-K., Doubrovski, E. L., Whiting, E. & Wang, C. C. Knitting 4D garments with elasticity controlled for body motion. *ACM Transactions on Graphics (TOG)* **40**, 1 (2021).
115. Ortiz, J., Poliero, T., Cairoli, G., Graf, E. & Caldwell, D. G. Energy efficiency analysis and design optimization of an actuation system in a soft modular lower limb exoskeleton. *IEEE Robotics and Automation Letters* **3**, 484 (2017).
116. Moore, E., Porter, M., Karamouzas, I. & Zordan, V. *Precision control of tensile properties in fabric for computational fabrication* in *Proceedings of the 2nd ACM Symposium on Computational Fabrication* (2018), 1.
117. Wang, C. C. & Tang, K. Pattern computation for compression garment by a physical/geometric approach. *Computer-Aided Design* **42**, 78 (2010).
118. Bartle, A., Sheffer, A., Kim, V. G., Kaufman, D. M., Vining, N. & Berthouzoz, F. Physics-driven pattern adjustment for direct 3D garment editing. *ACM Trans. Graph.* **35**, 50:1 (2016).
119. Wang, C. C. & Tang, K. Woven model based geometric design of elastic medical braces. *Computer-Aided Design* **39**, 69 (2007).
120. Kwok, T.-H., Zhang, Y.-Q., Wang, C. C., Liu, Y.-J. & Tang, K. Styling evolution for tight-fitting garments. *IEEE transactions on visualization and computer graphics* **22**, 1580 (2015).
121. Martínez, J., Dumas, J., Lefebvre, S. & Wei, L.-Y. Structure and appearance optimization for controllable shape design. *ACM Transactions on Graphics (TOG)* **34**, 1 (2015).
122. Schumacher, C., Thomaszewski, B. & Gross, M. *Stenciling: Designing Structurally-Sound Surfaces with Decorative Patterns* in *Computer Graphics Forum* **35** (2016), 101.
123. Liu, H., Hu, Y., Zhu, B., Matusik, W. & Sifakis, E. Narrow-band topology optimization on a sparsely populated grid. *ACM Transactions on Graphics (TOG)* **37**, 1 (2018).
124. Skouras, M., Thomaszewski, B., Coros, S., Bickel, B. & Gross, M. Computational Design of Actuated Deformable Characters. *ACM Trans. Graph.* **32** (2013).
125. Bruns, T. E. & Tortorelli, D. A. Topology optimization of non-linear elastic structures and compliant mechanisms. *Computer methods in applied mechanics and engineering* **190**, 3443 (2001).

126. Huang, X. & Xie, Y. Topology optimization of nonlinear structures under displacement loading. *Engineering structures* **30**, 2057 (2008).
127. Zhang, Y. & Kwok, T.-H. Customization and topology optimization of compression casts/braces on two-manifold surfaces. *Computer-Aided Design* **111**, 113 (2019).
128. Zhang, X., Fang, G., Dai, C., Verlinden, J., Wu, J., Whiting, E. & Wang, C. C. *Thermal-comfort design of personalized casts* in *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (2017), 243.
129. Ye, Q., Guo, Y., Chen, S., Lei, N. & Gu, X. D. Topology optimization of conformal structures on manifolds using extended level set methods (X-LSM) and conformal geometry theory. *Computer Methods in Applied Mechanics and Engineering* **344**, 164 (2019).
130. Di Natali, C., Sadeghi, A., Mondini, A., Bottenberg, E., Hartigan, B., De Eyto, A., O'Sullivan, L., Rocon, E., Stadler, K., Mazzolai, B., *et al.* Pneumatic quasi-passive actuation for soft assistive lower limbs exoskeleton. *Frontiers in Neurorobotics* **14**, 31 (2020).
131. Park, S. J. & Park, C. H. Suit-type wearable robot powered by shape-memory-alloy-based fabric muscle. *Scientific reports* **9**, 1 (2019).
132. Masia, L., Hussain, I., Xiloyannis, M., Pacchierotti, C., Cappello, L., Malvezzi, M., Spagnoletti, G., Antuvan, C., Khanh, D., Pozzi, M., *et al. Soft wearable assistive robotics: exosuits and supernumerary limbs* 2018.
133. Gholami, M., Rezaei, A., Cuthbert, T. J., Napier, C. & Menon, C. Lower body kinematics monitoring in running using fabric-based wearable sensors and deep convolutional neural networks. *Sensors* **19**, 5325 (2019).
134. Barone, V. J., Yuen, M. C., Kramer-Boniglio, R. & Sienko, K. H. *Sensory garments with vibrotactile feedback for monitoring and informing seated posture* in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)* (2019), 391.
135. Mahmood, N., Ghorbani, N., Troje, N. F., Pons-Moll, G. & Black, M. J. *AMASS: Archive of motion capture as surface shapes* in *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), 5442.
136. Yu, T., Zheng, Z., Guo, K., Zhao, J., Dai, Q., Li, H., Pons-Moll, G. & Liu, Y. *Doublefusion: Real-time capture of human performances with inner body shapes from a single depth sensor* in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2018), 7287.

137. Öztireli, A. C., Guennebaud, G. & Gross, M. *Feature preserving point set surfaces based on non-linear kernel regression* in *Computer graphics forum* **28** (2009), 493.
138. Skouras, M., Thomaszewski, B., Kaufmann, P., Garg, A., Bickel, B., Grinspun, E. & Gross, M. Designing inflatable structures. *ACM Transactions on Graphics (TOG)* **33**, 1 (2014).
139. Pipkin, A. C. The relaxed energy density for isotropic elastic membranes. *IMA journal of applied mathematics* **36**, 85 (1986).
140. Liu, D. C. & Nocedal, J. On the limited memory BFGS method for large scale optimization. *Mathematical programming* **45**, 503 (1989).
141. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J. & Chintala, S. in *Advances in Neural Information Processing Systems 32* 8024 (Curran Associates, Inc., 2019).
142. Sigmund, O. A 99 line topology optimization code written in Matlab. *Structural and multidisciplinary optimization* **21**, 120 (2001).
143. Ghabraie, K. An improved soft-kill BESO algorithm for optimal distribution of single or multiple material phases. *Structural and multidisciplinary optimization* **52**, 773 (2015).
144. Jinyun, Z., Yi, L., Lam, J. & Xuyong, C. The Poisson ratio and modulus of elastic knitted fabrics. *Textile Research Journal* **80**, 1965 (2010).
145. Kwok, T.-H., Wan, W., Pan, J., Wang, C. C., Yuan, J., Harada, K. & Chen, Y. *Rope caging and grasping* in *2016 IEEE International Conference on Robotics and Automation (ICRA)* (2016), 1980.
146. Vechev, V., Zarate, J., Thomaszewski, B. & Hilliges, O. *Computational Design of Kinesthetic Garments* in *Computer Graphics Forum* **40** (2022), 101.
147. Al-Fahaam, H., Davis, S. & Nefti-Meziani, S. *Wrist rehabilitation exoskeleton robot based on pneumatic soft actuators* in *2016 International Conference for Students on Applied Engineering (ICSAE)* (2016), 491.
148. Al Maimani, A. & Roudaut, A. *Frozen suit: designing a changeable stiffness suit and its application to haptic games* in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017), 2440.
149. Hinchet, R. & Shea, H. High Force Density Textile Electrostatic Clutch. *Advanced Materials Technologies*, 1900895 (2019).

150. Chen, B., Ma, H., Qin, L.-Y., Gao, F., Chan, K.-M., Law, S.-W., Qin, L. & Liao, W.-H. Recent developments and challenges of lower extremity exoskeletons. *Journal of Orthopaedic Translation* **5**, 26 (2016).
151. Shen, Y., Ferguson, P. W. & Rosen, J. in *Wearable Robotics: Systems and Applications* 1 (Elsevier, 2019).
152. Schiele, A. & Hirzinger, G. *A new generation of ergonomic exoskeletons - the high-performance X-Arm-2 for Space Robotics Telepresence* in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2011), 2158.
153. Barnaby, G. & Roudaut, A. *Mantis: A Scalable, Lightweight and Accessible Architecture to Build Multiform Force Feedback Systems* in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (ACM, New York, NY, USA, 2019), 937.
154. Asbeck, A. T., De Rossi, S. M., Galiana, I., Ding, Y. & Walsh, C. J. Stronger, smarter, softer: Next-generation wearable robots. *IEEE Robotics and Automation Magazine* **21**, 22 (2014).
155. Fang, C., Zhang, Y., Dworman, M. & Harrison, C. *Wireality: Enabling Complex Tangible Geometries in Virtual Reality with Worn Multi-String Haptics* in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (ACM, New York, NY, USA, 2020), 1.
156. Delazio, A., Nakagaki, K., Lehman, J. F., Klatzky, R. L., Sample, A. P. & Hudson, S. E. *Force jacket: Pneumatically-actuated jacket for embodied haptic experiences* in *Conference on Human Factors in Computing Systems - Proceedings* **2018-April** (ACM, New York, NY, USA, 2018), 1.
157. Ramachandran, V., Schilling, F., Wu, A. R. & Floreano, D. Smart Textiles that Teach: Fabric-Based Haptic Device Improves the Rate of Motor Learning. *Advanced Intelligent Systems* **3**, 2100043 (2021).
158. Ramachandran, V., Macchini, M. & Floreano, D. Arm-wrist haptic sleeve for drone teleoperation. *IEEE Robotics and Automation Letters* (2021).
159. Zehnder, J., Li, Y., Coros, S. & Thomaszewski, B. NTopo: Mesh-free Topology Optimization using Implicit Neural Representations. *Advances in Neural Information Processing Systems* **34** (2021).

160. Muthukumarana, S., Messerschmidt, M. A., Matthies, D. J., Steimle, J., Scholl, P. M. & Nanayakkara, S. *Clothtiles: A prototyping platform to fabricate customized actuators on clothing using 3d printing and shape-memory alloys* in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (2021), 1.

161. Verschoor, M., Casas, D. & Otaduy, M. A. Tactile rendering based on skin stress optimization. *ACM Transactions on Graphics (TOG)* **39**, 90 (2020).

162. Langerak, T., Zárate, J. J., Vechev, V., Lindlbauer, D., Panozzo, D. & Hilliges, O. *Optimal control for electromagnetic haptic guidance systems* in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (2020), 951.

163. Langerak, T., Zárate, J. J., Lindlbauer, D., Holz, C. & Hilliges, O. *Omni: Volumetric sensing and actuation of passive magnetic tools for dynamic haptic feedback* in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (2020), 594.

164. Tavassolian, M., Cuthbert, T. J., Napier, C., Peng, J. & Menon, C. Textile-based inductive soft strain sensors for fast frequency movement and their application in wearable devices measuring multiaxial hip joint angles during running. *Advanced Intelligent Systems* **2**, 1900165 (2020).