TRITA TRITA-CSC-A-2016:06
ISSN 1653-5723
ISRN KTH/CSC/A--2016/06--SE
ISBN 978-91-7595-882-8

www.kth.se

JONAS FORSSLUND Preparing Spatial Haptics for Interaction Design

KTH 2016

DOCTORAL THESIS IN HUMAN-COMPUTER INTERACTION
STOCKHOLM, SWEDEN 2016

Preparing Spatial Haptics for Interaction Design

JONAS FORSSLUND

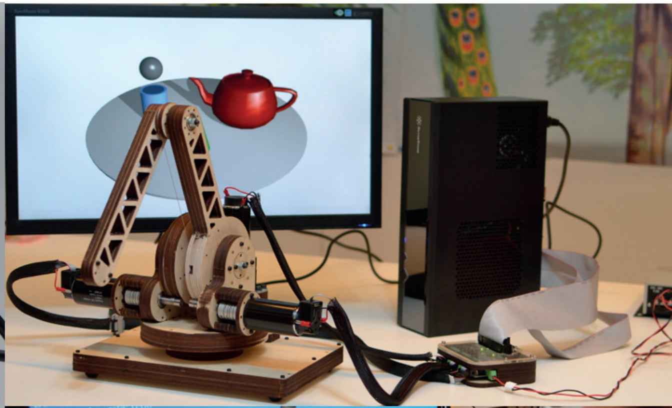

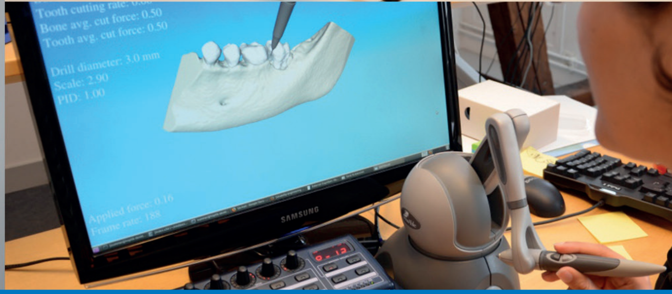

KTH ROYAL INSTITUTE OF TECHNOLOGY
SCHOOL OF COMPUTER SCIENCE AND COMMUNICATION

KTH Computer Science
and Communication

Preparing Spatial Haptics for Interaction Design

JONAS FORSSLUND

PhD Thesis
Stockholm, Sweden, 2016

TRITA TRITA-CSC-A-2016:06
ISSN 1653-5723
ISRN KTH/CSC/A--2016/06--SE
ISBN 978-91-7595-882-8

KTH School of Computer Science and Communication
SE-100 44 Stockholm
SWEDEN

Akademisk avhandling som med tillstånd av Kungliga Tekniska Högskolan framlägges till offentlig granskning för avläggande av doktorsexamen kl. 14 den 6 april 2016 i F3.

Tryck: Universitetsservice US AB

Abstract

Spatial haptics is a fascinating technology with which users can explore and modify 3D computer graphics objects with the sense of touch, but its application potential is often misunderstood. For a large group of application designers it is still unknown, and those who are aware of it often have either too high expectations of what is technically achievable or believe it is too complicated to consider at all. In addition, spatial haptics is in its current form ill-suited to interaction design. This is partly because the properties and use qualities cannot be experienced in an application prototype until a system is fully implemented, which takes too much effort to be practical in most design settings. In order to find a good match between a solution and a framing of a problem, the designer needs to be able to mould/shape/form the technology into a solution, but also to re-frame the problem and question initial conceptual designs as she learns more about what the technology affords. Both of these activities require a good understanding of the design opportunities of this technology.

In this thesis I present a new way of working with spatial haptic interaction design. Studying the serially linked mechanism from a well-known haptic device, and a force-reflecting carving algorithm in particular, I show how to turn these technologies from an esoteric engineering form into a form ready for interaction design. The work is grounded in a real application: an oral surgery simulator named Kobra that has been developed over the course of seven years within our research group. Its design has gone through an evolutionary process with iterative design and hundreds of encounters with the audience; surgeon-teachers as users and potential customers. Some ideas, e.g. gestalting authentic patient cases, have as a result received increased attention by the design team, while other ideas, e.g. automatic assessment, have faded away.

Simulation is an idea that leads to ideals of realism; that e.g. simulated instruments should behave as in reality, e.g. a simulated dental instrument for prying teeth is expected to behave according to the laws of physics and give force and torque feedback. If it does not, it is a bad simulation. In the present work it is shown how some of the realism ideal is unnecessary for creating meaningful learning applications and can actually even be counter-productive, since it may limit the exploration of creative design solutions. This result is a shift in perspective from working towards constantly improving technological components, to finding and making use of the qualities of modern, but not necessarily absolute cutting-edge, haptic technology.

To be able to work creatively with a haptic system as a design resource we need to learn its material qualities and how - through changing essential properties - meaningful experiential qualities can be modulated and tuned. This requires novel tools and workflows that enable designers to explore the creative design space, create interaction sketches and tune the design to cater for the user experience. In essence, this thesis shows how one instance of spatial haptics can be turned from an esoteric technology into a design material, and how that can be used, and formed, with novel tools through the interaction design of a purposeful product in the domain of dental education.

Sammanfattning

Att förbereda 3D-Haptik för interaktionsdesign

3D-haptik är en fascinerande teknologi med vilken användare kan utforska och modifiera tredimensionella datorgrafik-objekt med känseln, men dess användningspotential är ofta missförstådd. För flertalet applikationsutvecklare är tekniken fortfarande till stor del okänd, och de som känner till den har antingen alltför höga förväntingar av vad som är tekniskt möjligt, eller uppfattar 3D-haptik som alltför komplicerat för att vara ett gångbart alternativ. Dessutom är 3D-haptik i sin nuvarande form tämligen omoget för interaktionsdesign. Detta beror till stor del på att en applikationsprototyps egenskaper och användarkvaliteter inte kan upplevas innan ett system är implementerat i sin helhet, vilket kräver alltför stora utvecklingsresurser för att vara praktiskt försvarbart i de flesta designsituationer. För att uppnå en bra matchning mellan ett användarbehov i en viss situation och en potentiell lösning behöver en designer kunna å ena sidan formge och finjustera tekniken, och å andra sidan vara öppen för att ifrågasätta och ändra problemformulering och konceptdesign när hen lär sig mer om vilka möjligheter tekniken erbjuder. Båda dessa aktiviteter kräver en god förståelse för vilka designmöjligheter som en viss teknik, eller material, erbjuder.

I den här avhandlingen presenterar jag ett nytt sätt att arbeta med interaktionsdesign för 3D-haptik. Genom att studera i synnerhet den seriellt länkade mekanismen som återfinns i en vanligt förekommande typ av 3D-haptikenhet, och en kraftåterkopplande skärande/borrande algoritm visar jag hur man kan omvandla dessa teknologier från att vara en svårtillgänglig ingengörskonst till en form som är mer redo för interaktionsdesign. Denna förberedelse resulterar i ett slags designmaterial, samt de verktyg och processer som har visat sig nödvändiga för att effektivt kunna arbeta med materialet.

Forskningen är grundad i en verklig tillämpning: en simulator för käkkirurgi vid namn Kobra, som har utvecklas under sju år inom vår forskargrupp. Kobras utformning har genomgått en evolutionär utvecklingsprocess med iterativ design och hundratals möten med målgruppen; lärarpraktiserande käkkirurger och studenter som användare och potentiella kunder. Därvid har några designidéer, t.ex. gestaltning av patientfall, av designteamet fått utökad uppmärksamhet medan andra idéer, t.ex. automatisk gradering, har tonats ned.

Simulering är i sig självt en idé som ofta leder till ett ideal av realism; till exempel att simulerade instrument ska uppföra sig som i verkligheten, det vill säga ett simulerat tandläkarinstrument för att hävla (bända) tänder förväntas följa fysikens lagar och ge återkoppling i form av av både kraft och vridmoment. Om detta inte uppfylls betraktas simuleringen som undermålig. I det aktuella arbetet visas hur delar av realism-idealet inte är nödvändigt för att skapa meningsfulla lärandeapplikationer, och att det till och med kan vara kontraproduktivt eftersom det begränsar utforskande av kreativa designlösningar. Ifrågasättandet av realsimidealet resulterar i ett perspektivskifte vad gäller simulatorutveckling generellt, från att ensidigt fokusera på vidareutveckling av enskilda tekniska komponenter, till att identifiera och dra nytta av kvaliteterna som redan erbjuds i modern haptisk teknik.

För att kunna arbeta kreativt med ett haptiksystem som en designresurs behöver vi lära känna dess materialkvaliteter och hur, genom att ändra grundläggande parametrar, meningsfulla upplevelsekvaliteter kan moduleras och finjusteras. Detta kräver i sin tur

nyskapande av verktyg och arbetsflöden som möjliggör utforskande av det kreativa designrummet, skapande av interaktionssketcher och finjustering av gestaltningen för att tillgodose användarupplevelsen.

I grund och botten visar denna avhandling hur en specifik 3D-haptik-teknologi kan omvandlas från att vara en svårtillgänglig teknologi till att vara ett designmaterial, och hur det kan användas, och formas, med nyskapande verktyg genom interaktionsdesign av en nyttoprodukt inom tandläkarutbildning.

Acknowledgements

First and foremost I would like to thank my supervisor Eva-Lotta Sallnäs Pysander for her support, intellectual discussions, and encouragement to find my academic passion and go with it, even if it was uncharted territory. Next autumn it will be 10 years since I first approached her and asked if she would like to be my supervisor, at that time for my Master's thesis. I am indebted for all work she has done over the years, for the thesis but also for supporting and contributing to the Kobra project, for challenging me intellectually but never losing faith in me. Thank you and hope we can do interesting projects together also onwards!

A big thank you also goes to my co-supervisors Karl-Johan Lundin Palmerius, Ylva Fernaeus and Jan Gulliksen. KJ has been involved as long as Eva-Lotta, and has helped me to retain a solid technical ground in the work, even when my mind drift to concern more abstract design aspects. Ylva contributed by introducing me to a lot of the work on materiality and bringing attention to what the interesting findings in my work are regarding interaction design. Jan I thank for helping me focus and making sure that the thesis finally got settled.

I would also like to thank Petra Sundström for her excellent job as opponent of my licentiate thesis, and adding energy! She also contributed with a key idea; that the tools we create may be for our own use in the specialised design trade we chose to engage in, in my case haptic interaction design. I am honoured to be able to thank Karon MacLean, for agreeing to be my opponent and travel so far for my defence. The same goes for my committee; Sile O'Modhrain, Andreas Pommert and Charlotte Magnusson. An extra thanks to Charlotte, who gave invaluable feedback at my final seminar, and Cristian Bogdan, for reading my manuscript and asking good questions.

A large part of my doctoral studies was carried out as a visiting researcher at Stanford University. I am forever indebted to Kenneth Salisbury for letting me work in his lab during some of the best two years of my life. The fantastic environment was also enhanced by working with Mike Yip, with whom I made the first version of WoodenHaptics, and the rest of the lab; Reuben Brewer who taught me hands-on robotics design and at some point, when I doubted what do academically, said "if you want to make a haptic device, you should make a haptic device", Sonny Chan who became a dear friend and inspired me so much, and whom I enjoyed discussing everything with over a coffee in the lab or at excursions in sunny California, François Conti, Adam Leeper, Sarah Schvartzman, Billy Nassbaumer and Cédric Schwab, even undergrads programming robotic coffee runs contributed a lot to my understanding of haptics and what you can do if you are persisted and attentive. I should not forget to also thank the physicians and their exceptional engagement in our prototype development: doctors Nikolas Blevins, Rebeka Silva and Sabine Girod, thank you!

Back at KTH I was thrilled to find the working environment being transformed form a regular office space to a super-creative lab. I believe this change is much thanks to Ylva Fernaeus, and Kia Höök, who, among other things, gladly found the finance for "my' laser-cutter, and of course the merging of Mobile Lifers and other interaction designers into the environment. Without naming them all in fear of forgetting someone I wish to express my gratitude to them all, for letting me work with them in this fantastic research jungle. I have to especially thank Jordi Solsona, who not only happily joined in my stumbling steps in making electronics, but whose academic work, I think,

resonate very well with what is presented in this thesis. My many discussions with Anders Lundström has also been fruitful and always a pleasure. The same can be said for the many discussions over lunch in the "Blue Kitchen" with colleagues from all over the Media Technology and Interaction Design (MID) department.

The Kobra simulator would not have been what it is without the strenuous work by Martin Flodin, who contributed in all aspects of design, software development and not the least in joining me on road-trips to trade fairs with a simulator prototype in the trunk. Marcus Åvall did much of the professional design of the visuohaptic models, and my understanding of tools and workflow is much thanks to the privilege of working with him. Hans Forsslund, my dear father and oral surgeon who introduced me to the domain from the beginning, has contributed in many ways including interpreting patient cases, tweaking haptics and graphics, and hands-on woodworking! Many more deserve credit than I can find space for, but I have at least to mention the support and independent research on simulator usage by Bodil Lund and Annika Rosén at Karolinska Institutet. Ulrika Dreifaldt Gallagher, Helena Forsmark and colleagues at HiQ, Daniel Evestedt with colleagues at SenseGraphics, Anna Leckström, Ebba Kierkegaard, Johan Acevedo, Holger Ronquist and Martha Johansson, they have all contributed and have all been a pleasure working with. Ioanna Ioannou and Sudanthi Wijewickrema at University of Melbourne have been long-term contributors to the forssim software project, and we share many fun stories of the struggle with making surgical simulators, at both sides of the globe. The perspective of simulation case scenarios, the tuning tools and other ideas were conceived much thanks to our collaboration.

My greatest support however, whose company made me survive any periods of writers block and doubt, and with whom I have enjoyed far more periods of wonderful moments and adventures is my dear Anna Clara, and my family Titti, Hans, Ola and Annika.

Stockholm, March 2016
Jonas Forsslund

List of Publications

The thesis is composed of a summary and the following original publications, reproduced here with permission. Paper A and D are unpublished manuscripts.

Paper A

Forsslund, J., Sallnäs, E.-L. and Fernaeus, Y. Designing the Kobra Oral Surgery Simulator Using a Practice-Based Understanding of Educational Contexts. Manuscript submitted to European Journal of Dental Education.

Paper B

Forsslund, J., Yip, M., and Sallnäs, E.-L. (2015). Woodenhaptics: A starting kit for crafting force-reflecting spatial haptic devices. In Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction. Presented at TEI, Stanford, USA, 2015. Doi: 10.1145/2677199.2680595.

Paper C

Forsslund, J. and Ioannou, I. (2012). Tangible sketching of interactive haptic materials. In proceedings of Sixth International Conference on Tangible, Embedded and Embodied Interaction. Presented at TEI, Kingston, Canada 2012. 10.1145/2148131.2148156

Paper D

Forsslund, J., Sallnäs, E.-L. and Fernaeus, Y. Designing the Experience of Visuohaptic Carving. Manuscript submitted to Designing Interactive Systems 2016.

Paper E

Forsslund, J., Chan, S., Selesnick, J., Salisbury, K., Silva, R. G., and Blevins, N. H. (2013). The effect of haptic degrees of freedom on task performance in virtual surgical environments, Studies in Health Technology and Informatics, Volume 184: Medicine Meets Virtual Reality 20, pages 129 - 135. Presented at MMVR, Los Angeles, USA 2013.

The Author's Contribution to the Publications

This work was done as part of several research projects, at both KTH Royal Institute of Technology and at Stanford University where the author spent two years of his five years of PhD studies. The following summarise the contributions I have made to each attached paper and the underlying work.

Designing the Kobra Oral Surgery Simulator Using a Practice-Based Understanding of Educational Contexts

This research-through-design paper traces the seven years of design and development of an oral surgery simulator named Kobra. The results show how creative interaction design can be used to gestalt authentic surgical scenarios and discuss how the simulator design supports teacher-student collaboration and teaching. I have been the lead designer and developer of the simulator but with the support of a team and external consultants. The most recent patient cases, i.e. interactive exercises, were given form by a professional 3D artist. The analysis has been done together with the co-authors, while the text has been mostly written by myself with extensive feedback and support from the co-authors.

WoodenHaptics: A Starting Kit for Crafting Force-Reflecting Spatial Haptic Devices

This paper covers the design, discussion and evaluation of a novel haptic device named WoodenHaptics that is packaged as a *starting kit* where designers can quickly assemble a fully functional spatial haptic device and explore the design space of variations. The results show that non-specialist designers can assemble the device under supervision, that its performance is on par with high-quality commercial devices and what some variants of the device look like. The device was developed by the second co-author and myself during my two-year research visitor position at Stanford University, with support from the robotics lab we were in. The device kit was subsequently refined and rebuilt at KTH in Stockholm by myself. The electronics were improved with the assistance of Jordi Solsona. The user study on perceived performance was designed and largely performed by the third author. The technical performance study was performed by myself.

Tangible sketching of interactive haptic materials

This paper was a result of a joint project by myself and the co-author concerning how to explore and tune the haptic properties of digital objects for use in surgery simulation and similar applications. The result shows how a tangible music controller was re-purposed for real-time tuning of the properties and thereby to enable quick creation of interactive sketches that can be used to understand the "material", or be used to get feedback from stakeholders. The application stems from a need that we both had, in our two different universities, for developing a dental simulator and a temporal bone simulator respectively. The development and paper writing was conducted by both authors equally.

Designing the Experience of Visuohaptic Carving

This paper introduces the notion of *visuohaptic carving* as a useful design resource in various applications including, but not limited to, surgery simulation. To be a design resource, it is argued, there needs to be a reusable component, i.e. a software library, tools for forming the user experience and an efficient workflow that supports the creation of different interactive scenes that use the resource in question. A library with the necessary haptic algorithms has been implemented along with prototype tools and associated workflow. The application of these to the Kobra simulator project and the analysis constitute the results showing its usefulness. The library was developed by myself with external collaborators. The prototype tools and workflow were developed by myself with feedback from the collaborating 3D artist. The analysis was done in collaboration with the co-authors, while most of the text was written by myself with significant contributions of the co-authors.

The Effect of Haptic Degrees of Freedom on Task Performance in Virtual Surgical Environments

Haptic devices that can provide both directional and rotational force feedback are rare and expensive, which has motivated investigation of how much effect the rotational torque feedback gives compared to cheaper alternatives. Furthermore, there has been a misconception that multi-degree haptic-rendering algorithms are useful only if torques can be displayed by the haptic device. An experiment was therefore set up to test three different conditions with twelve human subjects performing tasks in two different virtual environment scenes. The study was conducted by me at Stanford University, with the support of the co-authors. The study was designed primarily by myself, while the test application was primarily developed by the second author. The analysis was carried out by me, while the text was written collaboratively by all co-authors.

Contents

Chapter 1

Introduction

While we are used to interacting with computers using vision, and to some degree audition, technical advancement has enabled the addition of *haptic* interaction, or the interaction with the sense of touch. Most people are familiar with the vibrations synthesised by mobile phones and other electronic devices designed to, e.g., alert their users of incoming messages without interfering with other sensory channels. This thesis is concerned with the bi-directional counterpart where users can explore and modify virtual shapes in three-dimensional space through the sense of touch, i.e. through *spatial haptic interaction*. Despite being around for almost 20 years, computerised spatial haptics has not yet met its full potential for improving interaction in real world applications [Wright, 2011]. Spatial haptics has been quite inaccessible for interaction design practitioners. This thesis will explore this topic and show how spatial haptics can be prepared for interaction design, in particular a kind that is applied in simulations for teaching surgical procedures.

Learning surgery is traditionally a kind of situation that heavily relies on hands-on practice under supervision. The mantra "see one, do one, teach one" is often used to describe the general educational approach. The advent of computer-based simulation technologies and spatial haptic technologies has opened up opportunities for developing products that can be used for improving the learning situation, not least by eliminating the patient risks involved in novices operating on live humans. We call these products surgery simulators.

An integral activity in developing any product is deciding which technologies it should use, in what way, how it should look and behave, what form it should take, how the users should interact with it and so on. A traditional approach in engineering is to do requirements engineering [Sommerville, 2004] through, e.g., field studies, interviews, observations etc. with the goal of forming system requirements. The requirements should be well defined and not ambiguous. The development project then shifts into a technical design phase where a prototype is defined and implemented to meet these requirements. It is important to not change the requirements in this phase; the developers should only try to meet or exceed them. If the requirements are not met, the whole process should be iterated until the requirements finally are met.

Design practice inspired by other design fields has recently gained increasing interest in the larger field of human-computer interaction (HCI), and applying a design approach to a simulator development seems to have many benefits. However, to work design-wise with the components of the simulator, in particular with the haptic interface, these technologies need to be what I call *prepared for design*. The current knowledge about developing haptic interfaces for synthetic touching and carving poorly support a design approach because:

1. There are no articulations of what the key qualities and affordances of this technology give in concrete, real applications, and there is little knowledge relevant for design, i.e. that clearly explains what use experiences we can expect to get and how these can be achieved and modulated (altered, tuned) with reasonable development effort, and what the trade-offs are.
2. Developers have to fully implement a system in order to experience what is possible and feasible. In contrast with many screen-based interaction systems, there are no good representational prototype methods that work sufficiently like paper prototyping does for some conventional user interfaces.
3. The range of devices is limited and those that exist provide very different levels of quality, e.g. stiffness, but there is no possibility of changing the qualities of these devices to find a good match between device and use situation.

1.1 Objective

The purpose of this thesis is to investigate what preparations are needed to effectively work with the interaction design of the haptic modality of advanced interactive products. The idea that technology needs to be prepared for interaction design has previously not been widely explored, although research on kits, tools and materialities in HCI arguably points in that direction. Therefore part of the thesis will be dedicated to arguing why it indeed is important and grounded in design experiences from the development of a surgery simulator. This will culminate in the development of a set of design resources, tools and associated practices based on proven technologies, i.e. known haptic-rendering methods and hardware principles, but catering for the needs of interaction design. Their usefulness is then investigated by applying them to the design of the haptic modality of a real-world surgery simulator.

1. Why is it important to prepare haptic technology for interaction design?
2. How can spatial haptic technologies be prepared for interaction design?
3. How can novel design resources, tools and associated practices for spatial haptic interaction design be leveraged for surgery simulation design?

1.2 Context of Research

This thesis is about supporting interaction design activities. Therefore it is important to clarify what is actually meant by design in this context.

The word *design* can have different meanings in different contexts and to different people. Although they sometimes overlap, I have come across three major different meanings: engineering design, integral design and styling design. These three categories should not be taken as defining all kinds of design, nor what the essence of design is. That is beyond the scope of this thesis, but the interested reader is advised to start exploring the philosophy of design in, e.g., [Lawson, 2005], [Brown et al., 2008] and [Nelson and Stolterman, 2012], and of practical knowledge in general in e.g [Molander, 1993]. Design practice has been subject to study as well, perhaps most well-known is Donald Schön's observations of student design work in architect education leading to the famous notion of *design as a reflective conversation with the situation* [Schön, 1984, Chap. 3]. In his chief example the situation in questions was an architectural challenge of designing a school building on a particular piece of land, that featured a particular slope. The student draw and tested various ways of layouts of the building, while continuously judging and evaluating the work, directly or with the help of her teacher. The *conversation* she was said to have was thereby with the situation of the sloping land, or with the tangible sketch she was making, in other words the *material* she was directly manipulating. This idea has been applied to software, for example in Terry Winograd's compilation "Bringing Design to Software" [Winograd et al., 1996], that also features an interview with Schön [Chap. 9][Winograd et al., 1996]. The material in question can be digital [Dearden, 2006], and even haptic sketches [Moussette, 2012], as will be discussed further in this thesis.

In traditional engineering terms, a development process starts with gathering and forming system requirements, a process called requirements engineering [Sommerville, 2004]. These requirements specify what the system should do, and what constraints are put on the solution. One can easily imagine the requirements for a bridge, with requirements for spanning a particular river, where the constraints are that it should hold one hundred cars with an average weight of two tonnes. In software engineering, it may be a search engine that should handle millions of multiple users and conducting for each of them a database lookup within 200 milliseconds. These requirements and constraints are used in the next phase of the development process, called the design phase, where a (usually only one) solution is formed that meets those requirements and complies with the constraints. The solution is then implemented[1] and tested in order to verify the solution against the initial requirements. The whole process can then be iterated, which is the basis for the original user-centred design process[2]. In reality, the phases of development are more integrated, and the specifications can be more or less rigid depending on the application. In some situations, such as an airplane control system or in healthcare, formal methods and strict requirements formulations are critical and motivated by the large costs and efforts that are involved. For other systems, the requirements definition and solution formation are more integrated. The point is that, in engineering lingo, there is still an important distinction between activities that belong to defining what the system should do, and what the solution should be like. The design, i.e. the technical solution, should never breach the requirements.

---

[1] the design and implementation is usually mixed too

[2] defined by ISO 13407

The architect Bryan Lawson [Lawson, 2005] paints another view of design in *How Designers Think - The design process demystified.* Here, design is a radically integrated process that goes back and forth between sketching potential solutions and need-finding combined with identifying formal constraints (in architecture there are many regulatory constraints), adding the designers own personal touch and more. It is inherently creative and allows for influences and inspiration from any source. The process is as much about problem-solving as about problem-setting, questioning the original task set by the client. This approach can seem very messy but it is exactly this messiness that in practice has resulted in innovative and good design. This view of design is inclusive and covers professionals such as architects, fashion designers and engineers, as well as amateurs decorate their living rooms. This multi-faceted view of design is also found in Winograd's early exploration of what design applied to software constitutes [Winograd et al., 1996].

Design is also used for form-giving and the styling of products. The foundation for styling is aesthetic sensitivity, and a professional designer is usually expected to have a degree in fine arts, e.g. an MFA (Master of Fine Arts), or some other artistic training. When an object in popular culture is referred as "designed" or as a "designer-product", what is meant is that particular attention has been paid to its form and style, which have sometimes been prioritised over more technical aspects such as power, efficiency etc. Form and style should not be seen as merely decoration; a good form is essential for ergonomics, and a good style clearly communicates the function of the product and how it can be used. In addition, form and style can signal qualities of the product, its producer (branding) and project qualities to the owner (you are what you wear). This is referred to as product semantics. Anna Ståhl [Ståhl, 2014] shows the power of this kind of design with the example of a research product called Affective Diary. This product consists of two parts: a body-worn device that logs heartbeats throughout the day, and a desktop application that visualises the sensor readings in a style that evokes reflection in an open-ended way using hand-drawn figures that represent different values. The discussion central to her work is the styling, not the holistic design of the product, which would include discussing the mapping of sensor values to figures among other technical aspects. Another example of discussions where the term "design" mainly refers to form and style over product design is a passage in Brunnström's (ed.) book on 20th century Swedish Industrial Design history [Brunnström, 1997], where particular designs of radios for domestic use are discussed. When a designer is named and the design is discussed, it is mainly about the shape and material of the enclosure and less about the design of the audio qualities[3].

In many research disciplines it is common to talk about study design, where, e.g., a questionnaire and procedures are designed to study some phenomenon. In traditional human-computer interaction, some apparatus, sometimes called a prototype, is often designed as a vehicle for experimental study of an isolated phenomenon, e.g. how quickly and accurately a user can move a mouse cursor from point a to point b, dependent on the size of the target [MacKenzie et al., 1991]. Design is also used as a research approach to exploring what something novel could be like. The question is then centred around how

---

[3] There are many other examples in that book where they do discuss design beyond form and style; for example, the design of fridges.

to design *for* x, where x is some aspect of particular interest to the researcher. Examples include *Designing for the Pleasure of Motion* [Moen, 2006], *Designing for Interaction Empowerment* [Ståhl, 2014], *Designing for Well-Being* [Ilstedt Hjelm, 2004] and *Designing for Children's Creative Play with Programming Materials* [Fernaeus, 2007]. Design can also be used to support enquiry into larger contexts, creating knowledge that is intended to reach far beyond how to design utility products. The designed artefacts may then spur discussion on, e.g., environmental concerns [Broms, 2014]. Design has even been used to create artefacts explicitly without any predefined purpose, just to see people's reaction, from which conclusions are drawn [Gaver et al., 2009].

In contrast to these works, the present thesis is not primarily concerned with designing *for* a particular domain or end, but takes its basis in a particular technology. At the same time, it is not the concern of the thesis to advance the technical state of the art either. The focus is to prepare advanced haptic technology for integrative design as discussed above. The aim is thus that interaction designers can investigate the design space and reformulate requirements in a much more direct fashion than if they were forced to engage in advanced technical problem-solving or rely on specialised engineers for realisations of prototypes.

1.3 Main Results

Haptic interaction design has been shown to greatly benefit from the possibilities of working directly with the material, without relying on artificial representations as is common in, e.g., low-fi prototyping [Moussette, 2012]. To prepare for design explorations in non-trivial target mediums, two general requirements need to be fulfilled. First, the technology needs to be prepared as a design resource (or "material"), which essentially implies encapsulating complex nuances and exposing design-relevant properties. Second, tools with which the design resource can be formed need to be created or re-purposed.

The main contributions of this thesis are two-sided. On one side, a particular subset of spatial haptic technology is transformed from an esoteric technology into a resource suitable for design explorations. This is done through the construction of a modular and modifiable physical haptic device whose performance is on par with commercial devices but which is still open for design variations. The workbench where the device is located becomes a tool for hardware design. A software library enables the creation of three-dimensional carving experiences, and a tool for tuning the experience of carving is proposed. The software tool is integrated into a workflow that leverages the skills and tools of professional 3D artists in the design of interactive environments. The parameters that can be tuned are directly derived from the internal workings of the rendering algorithms and mechanical reality, e.g. stiffness, carving rate and scale.

On the flip side, a fully functional haptic-enabled surgery simulator has been designed and developed. In effect, this simulator development has acted as a *principal driving problem*[4] motivating and generating requirements for the material and tool development. The research-through-design work of the simulator development has itself yielded design

---

[4] Frederick Brooks of UNC Chapel Hill famously used a long-term driving problem of molecular docking for his group's work on virtual reality and haptics; see, e.g., [Brooks Jr, 1996].

knowledge, in particular in terms of which role creative haptic interaction design can serve for the teaching of surgery. It was observed that surgical scenarios could be gestalted in the simulator and made relevant for teaching, not because they were super-realistic, but because they were linked to real practice and supported real-life tutoring between surgeon-teacher and learner.

Another contribution is the result of a controlled experiment with human participants that shows that employing a more advanced (6-DoF) haptic rendering algorithm improves task performance in some virtual environments. The most interesting result was that the performance increase remained even if a device without torque feedback was employed. It has previously been a common misconception that to benefit from a 6-DoF algorithm one has to use a torque-feedback capable haptic device. The study results shows that 6-DoF algorithms actually can be used with benefit together with under-actuated devices, i.e. cheaper devices that reads position and orientation but only exert directional forces.

1.4 Structure of the Thesis

The intent of this introductory chapter has been to define what kind of design work the thesis work is intended to support (Holistic, integrative design: 1.2 Context of Research). The methods used to approach the objective of finding what is required to support this design practice when using spatial haptics technologies have been discussed (1.1 Objective, and 3 Research Process), as well as high-level description of the results (1.3 Main Results). This is followed by a short summary of the attached papers (1.5 Short Summary of Papers). The papers themselves, found as appendices to the thesis body text, are recommended reading material and contain additional images and information that may complement the body text.

Chapter two introduces the background of this research and related work. A short introduction is given to the human sense of touch, particularly the kind of active touch that spatial haptic technologies cater for (2.1 Haptic Perception). These technologies, presented in a historical context, are introduced in 2.2 Core Technologies, which covers both hardware and software aspects. As the contributions of my work are related to creating tools for haptic interaction design, will a full section be dedicated to previous work in this domain (2.3 Tools for Haptic Interaction Design). This section will also cover ways in which haptic technologies have been packaged for designers or in other ways been made more accessible. Finally, as related work, will the application domain of surgery simulation be presented, with particular focus on design and use of surgical simulators in dental education (2.4 Surgery Simulation). This section will also present the Kobra simulator, including previous published results from studies of its various prototypes. This is because the simulator itself and its effect on dental education are not primary to the aim of this thesis but rather is the simulator used to motivate and drive the research.

Chapter three covers the research projects that have been undertaken in order to investigate the research questions. These projects are the Kobra simulator, Tuning of Visuohaptic Carving properties, WoodenHaptics and a study on haptic rendering degree of freedom effect on user performance.

Chapter four presents the research contributions of the thesis. This will not cover all results and contributions made during the thesis work, but a selected focus on what was found most interesting: the transformation of the technologies into tools and resources for interaction design (4.1 Tools and Resources for Spatial Haptics Interaction Design) and how these can been applied, with benefit, in a real-world surgery simulator design project (4.2 Interaction Design for Surgery Simulators). The latter section is also used do describe what role interaction design may play in advancing surgery simulation state of art.

The body text of the thesis will end in a discussion (Chapter 5, Discussion), that will discuss the work on a higher level and reflect on the research questions introduced in the introduction (1.1 Objective). In particular it will be discussed why *preparing* technology is a interesting perspective that motivates further attention in the field of Human-Computer Interaction. The chapter will also include limitations of the present work and conclusions.

1.5 Short Summary of Papers

The following summarize each paper that together with the body text make up the thesis. The papers are re-printed in full as appendices A-E.

Designing the Kobra Oral Surgery Simulator Using a Practice-Based Understanding of Educational Contexts

This research-through-design paper traces the seven years of design and development of an oral surgery simulator named Kobra. The results show how creative interaction design can be used to gestalt authentic surgical scenarios and discusses how the simulator design supports teacher-student collaboration and teaching.

WoodenHaptics: A Starting Kit for Crafting Force-Reflecting Spatial Haptic Devices

This paper covers the design, discussion and evaluation of a novel haptic device named WoodenHaptics that is packaged as a *starting kit* where designers can quickly assemble a fully functional spatial haptic device and explore the design space of variations. The results show that non-specialist designers can assemble the device under supervision, that its performance is on par with high-quality commercial devices and what some variants of the device look like.

Tangible sketching of interactive haptic materials

This paper presents a novel tool for sketching and tuning haptic properties of digital objects for use in surgery simulation and similar applications. The result shows how a tangible music controller was re-purposed for real-time tuning of the properties and thereby enables quick creation of interactive sketches that can be used to understand the “material” or present to stake-holders.

Designing the Experience of Visuohaptic Carving

This paper introduces the notion of *visuohaptic carving* as a useful design resource in various applications including, but not limited to, surgery simulation. To be a design resource, it is argued, there needs to be a reusable component i.e. a software library, tools for forming the user experience and an efficient work-flow that support creation of different interactive scenes that use the resource in question. A library with necessary haptic algorithms has been implemented along with prototype tools and associated work-flow. The application of these to the Kobra simulator project and two other applications, together with the analysis constitutes the results showing its usefulness.

The Effect of Haptic Degrees of Freedom on Task Performance in Virtual Surgical Environments

Haptic devices that can provide both directional and rotational force feedback are rare and expensive why it is motivated to investigate how much effect the rotational torque feedback gives compared to cheaper alternatives. Furthermore, there have been a misconception that multi-degree haptic rendering algorithms only are useful if torques can be displayed by the haptic device. An experiment was therefore set up to test three different conditions with twelve human subjects performing tasks in two different virtual environment scenes.

Chapter 2

Background and Related Work

2.1 Haptic Perception

In general, engineering and designing haptic interaction with computers is a large endeavour and requires special purpose robotics hardware. Why then go through so much trouble to support this sense when much of the everyday computing tasks can be accomplished with visual feedback alone? There are several answers to this question. One is that the application designers simply put may have a deep desire, a *desiderata*, to provide their users with a rich visceral interaction [Moussette, 2012]. Another answer is that the haptic sense, as will be discussed shortly, actually has a set of unique properties that can be leveraged for practical reasons in the interaction with a computer. Last but definitely not least, might not the haptic sense actually be of much more importance to humans, in comparison with the other senses, than what is commonly thought? Gabriel Robles-De-La-Torre [Robles-De-La-Torre, 2006] has rhetorically asked, “What would be worse? Losing your sight or your sense of touch?” and referred to two actual cases where patients had indeed lost large parts of their haptic sense due to nerve damage. One of them, Mr Waterman, who also featured in the BBC documentary “The Man Who Lost His Body”, had completely lost his proprioception from the neck downwards as a result of an autoimmune response to a virus infection attacking exactly those nerves that carry the information of limb position and touch sensation to the brain. In fact, Mr Waterman could still sense pain and temperature, and he could command his muscles to move. The problem was that without feedback the limbs would just drift away as he started moving them. Over the years he learned to move and even walk, but only by planning and executing each motion actively and under direct view. Any activity that required both cognitive load and fine-motor control, such as taking the minutes at a meeting, required constant switching between listening and cautiously controlling his handwriting [Robles-De-La-Torre, 2006]. The haptic sense is clearly something to take seriously and well worth the attention of interaction designers.

The haptic sense, or more precisely, the human haptic system, involves both sensory receptors and higher level cognition [Lederman and Klatzky, 2009]. When we explore the objects of the world through the sense of touch, sensory information is derived from both *cutaneous* receptors in the skin and *kinaesthetic* receptors in the muscles, tendons and joints. Sometimes haptic technology refers to the provision of one-directional stimuli, e.g. applying vibrations to the skin. This is useful for getting our attention without disturbing us or when other senses are occupied [MacLean, 2000]. This kind of haptics is, from the human perspective, *passive*, in that the stimulus is invariant to our motion. When humans explore everyday objects with the haptic system to form a mental representation of their properties such as shape, size, weight, surface texture and compliance, they do so through *active* touch. In fact, humans have developed several *explorative procedures* that are commonly used depending on what property is being examined. Weight is, for example, estimated best by lifting and wielding the object rather than holding it still. The exact shape of an object is best determined by following its contours with one or several fingers. Even when our interaction with the world is tool mediated, i.e. when holding onto a probe or a pencil and touching objects with that, the *contour-following* explorative procedure is effective. Most of the information is then received from the kinaesthetic receptors, but vibrations from the tool interaction and the skin shear it may cause, is registered by cutaneous receptors in the skin that also contribute to the perception. This human ability enables the construction of haptic *interfaces* where the user holds onto a tool but, instead of exploring everyday objects with it, can explore computer-generated ones. This is achieved through mechanically coupling the tool, which hereafter will be referred to as the *manipulandum*, to a robotic arm that will exert the forces that correspond to the forces reflected when the tool is pushed against real objects.

2.2 Core Technologies

The interaction of concern in this thesis is, at its most fundamental level, between a human-operated tool and one or several three-dimensional virtual objects residing in the memory of a computer. A precise definition can be challenging since the objects in question can either be virtual representations of real objects, or totally imaginative, and yet we will throughout the thesis use language such as “touching”, “seeing” and “carving”. As in the famous painting by René Margritte depicting a pipe subtitled *Ceci n’est pas une pipe*, “this is not a pipe”, these objects are only residing in our mind. This fact, however, does not disqualify a desire to give them form and use technology through which they can be perceived by our senses. It can therefore be meaningful to refer to them as objects, keeping in mind that their existence and material properties are at the same time immaterial and, through transducers, physical.

Practically, it may be more fruitful to use the term *computer graphics (CG) objects*, because of its familiarity and the fact that the study of computer haptics in computer science, as noted by Chan [Chan, 2014], shares several similarities with the study of computer graphics. It is only the rendering methods that are different. Geometric modelling, i.e. the way objects are represented mathematically, is fundamental both for visual and haptic displays. The creation of three-dimensional CG objects has a long tradition in the movie and computer game industry as well as in medical visualisation and many other fields.

The rest of this chapter will present the core technologies needed to touch and carve

CG objects. First, a short introduction to object representations in the field of computer graphics will be given. It serves two purposes: to define exactly what representations are suitable for carving and haptic rendering, and to give an account of how these are created in a professional way. These are the objects that will be interacted with through the mediation of a rigid tool, and the subsequent sections will describe how the interaction is materialised.

In order to create the sensation of touching the objects with a rigid tool, a physical link to the human is needed. This can be achieved with a spatial haptic device that has a *manipulandum* that the user holds on to and that can resist motion when a representation of the manipulandum - its avatar - comes into contact with the virtual objects. The ability to resist motion comes from the ability of these devices to exert computer-controlled forces onto the manipulandum. Thereby they become transducers of computational information; in other words a force display, in an analogy with visual displays [Salisbury et al., 2004]. These devices can be of different size and have different motion capabilities (e.g. whether they support rotations or not), ergonomics and force-producing capabilities. The devices commonly available today have a historical background to their looks and capabilities, which is important to the discourse. Contrary to what may first be thought, the oldest devices were more advanced than the newer, but that made them also very complex and expensive. This historical background supports the forthcoming discussion on complexity and sufficiency of realism.

The general process of computing the forces for display to the haptic device is subject to the field of *computer haptics*, which includes computing forces for conveying information, e.g. for visualisation [Palmerius et al., 2008]. The particular task of rendering contact with CG objects sorts under the subfield of *haptic rendering*. Computing the resulting forces of interaction between the user-controlled avatar and CG objects is not a trivial task, and needs to be completed in a short time, usually within one millisecond to guarantee stability of the haptic device. Different algorithms have been proposed of varied complexity and sophistication. The purpose is to give an overview of the problems involved and why some methods can be considered feasible to implement by a software engineering generalist, while others require highly specialist competence and effort. In addition, they will introduce the concept of *stiffness*, which is shared by practically all rendering algorithms, and which, together with haptic hardware, gives the relative hardness feeling peculiar to present-day spatial haptic interaction.

Finally, in order to carry out tasks like carving, the notion of *interaction techniques* is introduced, and how the carving has been used in the fields of computer graphics and haptics. The purpose is to show that carving, although under various labels, has been proposed both for visualisation and sculpting with imaginative tools, and for realism-aspiring simulation for surgical training in particular. Various algorithms of different levels of sophistication have also been proposed for this task. One important aspect that will be introduced is that different regions of a CG object can be designed to have different perceived carving hardnesses.

Representation and Creation of Solid CG Objects

An *object* can, as in everyday language, refer to a lump of physical matter such as a rock, a house or a ball. It can also refer to Margritte's pipe. In computer graphics, *geometric modelling* is the process of creating *representations* of object shapes in a format suitable for a computer [Foley et al., 1994]. The objects of concern in this thesis are solid, and thus pertain to the area of *solid modelling*, i.e. the representation of volumes completely surrounded by surfaces. These can be represented in different ways, e.g. a ball can be represented analytically with the mathematical definition of a sphere with a certain radius, or approximated with a collection of polygons (small flat surfaces) that bounds the volume, called a *polyhedron*, also referred to as a watertight polygon mesh. A polyhedron then in turn relies on mathematical descriptions of the small surfaces, the polygons, consisting of vertices (points) and edges (lines), which are referred to as geometric *primitives*. A *CG object* is then defined as a collection of geometric primitives organised in a hierarchy, and is stored together with all its numerical data, e.g. co-ordinates of its vertices [Foley et al., 1994]. It is worth highlighting, as Foley et al. do, that "when there is no preexisting object to model, the user creates the object in the modeling process; hence, the object matches its representation exactly, because its only embodiment is the representation" [Foley et al., 1994, p. 322]. In other cases there is always an approximation.

Most common are polygonal CG objects that only model the surface of an object. Interesting carving experiences also require the modelling of the non-homogeneous inside of the object. This implies that a way to represent solid objects is needed. Furthermore, a representation needs to be compatible with visual and haptic-rendering algorithms and suitable for carving. For these reasons it is usually more appropriate with a representation based on *spatial partitioning*, in particular a regular 3D grid of volume elements, *voxels*. In a spatial-occupancy representation each voxel contains only a Boolean value, i.e. the voxel either belongs to the object or is treated as free space. It enables very efficient look-up for, e.g., collision detection. The downside is that resolution is limited by the voxel size, and if not enough voxels are used it may look pixelated like a zoomed-in bitmap image.

Alternatively, a voxel may contain a value, which mathematically may represent a point sample value of a "smoother" object encoded by some band-limited signal [Engel et al., 2004, p. 3]. In practice, this means storing the equivalence of a grey-scale colour value in each voxel, e.g. from full black outside to full white inside, and allows for reconstructing a surface of the same grey values "between" the sample points, i.e. an *iso-surface*. This surface can be visually rendered either through direct volume rendering methods based on tracing races of virtual photons, or by constructing an intermediate polygon mesh through, e.g., Marching Cubes [Lorensen and Cline, 1987] and then rendering that.

The sources of CG objects can roughly be divided into human-made models and real-world acquisitions through imaging techniques [Riener and Harders, 2012]. The latter objects are acquired by scanning real objects, e.g. through computed tomography, where x-ray attenuation is recorded in a 3D grid. The former are usually created by a *3D artist* using interactive modelling programs, which fundamentally place primitives such as points and lines in spaces and arrange them in a hierarchy. The last two decades or so have seen a tremendous improvement not only in rendering techniques but also in the sophistication of interactive modelling programs and the professionalisation of the users, as is evident in job descriptions and emerging specialised education programmes for 3D artists[1] [Vaughan, 2011].

It is possible to translate from one representation to another. A polyhedron may be sampled or *voxelised* into a voxel volume. A computed tomography 3D image may be decomposed into structures through *segmentation*, a process where each voxel belonging to a structure of interest is assigned a label stored in an adjacent label volume [Preim and Bartz, 2007, pp. 95-96]. This can be achieved by manually "painting" areas of interest slice by slice, or through various automatic or semi-automatic methods, all with their respective benefits and costs in terms of accuracy, manual labour time etc. Specialised software for the work has been developed [Schiemann et al., 1992, Yushkevich et al., 2006], but can safely be said to be far from mature, for generating CG objects, as professional polygon modelling programs used by 3D artists. The clear benefit of segmenting is that the resulting label volume can easily be used to represent an object with several layers or tissues, e.g. a tooth can be modelled with a solid layer of dentin, covered by enamel and with pulp and nerves, and the shapes can be derived from a CT image as a template.

Spatial Haptic Devices

Haptic devices can in general be classified according to which part of the sense of touch they primarily support; vibrotactile devices stimulate cutaneous receptors in the skin, while kinaesthetic devices stimulate the kinaesthetic receptors in the muscles, tendons and joints. Vibrotactile devices, today ubiquitous in mobile phones and elsewhere, are generally one-directional in that they normally only act as an output channel without direct user input. Kinaesthetic haptic devices are, however, bi-directional, and it is through active human input *and* output that they can support haptic explorations. A spatial haptic device, then, is a kinaesthetic device that tracks a *manipulandum* (handle) in space, and has the means to restrict its motion or exert directional forces on the same.

Haptics as a human-machine interface has a long history, if we look outside the field of human-computer interaction. Force-reflecting remote-controlled manipulators were constructed as early as 1945 in the field of teleoperation, in particular for handling hazardous materials in the nuclear industry. These so called master-slave manipulators consist of two mechanically and electrically coupled arms, separated by a thick wall with a window through which the operator can see the manipulator in action. Being bilateral, or *force-reflecting*, any motion or force applied to the master is reflected on the slave and vice versa [Bejczy, 1980].

These early non-computerised tools relied on kinematically identical manipulators that allowed for a direct mapping between joints of the respective manipulators. To avoid this dependency, the kinematics had to be computationally converted from one manipulator to the other. This was the focus in one of the projects at the NASA Jet Propulsion Laboratory around 1980. Through attaching force and torque sensors to the end-effector of the remote

---

[1] E.g. University of Skövde three-year programme in Computer Game Development - Graphics, and two-year higher vocational education programme in 3D Graphics at FutureGames, Stockholm, Sweden

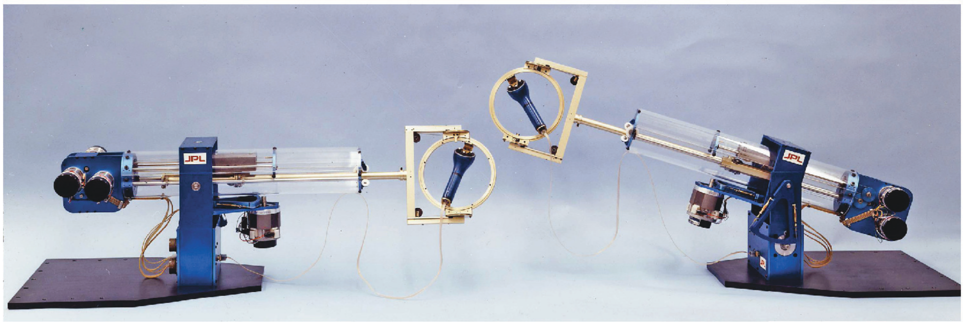

Figure 2.1: Two NASA JPL/Stanford Force-Reflecting Hand Controllers, circa 1989. Courtesy NASA/JPL-Caltech, www-robotics.jpl.nasa.gov (accessed 2016-02-11).

manipulator, and constructing a novel *general-purpose force-reflecting hand controller* capable of sensing position and orientation and applying forces and torques fed from the remote manipulator, the interaction became computer-relayed instead of directly coupled [Bejczy, 1980]. In this respect the user's manipulator, named the Stanford- (or Salisbury-) JPL Force Reflecting Hand Controller (figure 2.1), designed by Kenneth Salisbury and John Hill in the mid 1970s at Stanford Research Institute on contract from NASA JPL, became one of the first computer-controlled spatial haptic devices [Sherman and Craig, 2002].

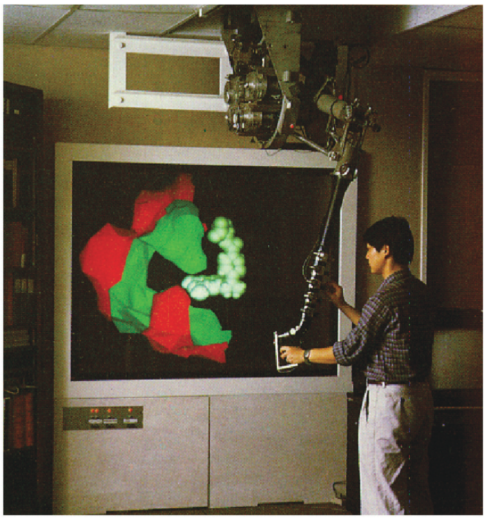

Figure 2.2: GROPE III, an Argonne ARM for haptic display used with a molecular docking application [Brooks Jr et al., 1990]. Image used with permission from Association for Computing Machinery, Inc.

An early account of using haptic devices originally developed for teleoperation, for interacting with computational models, was the long-term GROPEHaptic project at the University of North Carolina [Brooks Jr et al., 1990]. The nuclear remote manipulator they used was ceiling-mounted and used together with a large display where a standing user could explore a model of molecular docking complete with forces, albeit at low haptic update rates (figure 2.2). While users of GROPE could adapt to moving a manipulator in a workspace on the metre-scale (arm motion), they noted it would be simpler and more economical with a smaller device that provided a centimetre-decimetre scale (wrist and finger motion) and which also would be less tiring to use [Brooks Jr et al., 1990].

It was with this background that Thomas Massie under Kenneth Salisbury's supervision designed the *three degree of freedom force-reflecting haptic interface* that became the commercially successful and widely distributed Personal Haptic Interface Mechanism (Phantom) [Massie and Salisbury, 1994]. Compared to the earlier remote control masters it had a hand-scale workspace and no torque feedback, and in this respect a much simplified, cleaner design. It did have sensing of position and orientation, but providing only force feedback, and no torque feedback, making it an *asymmetric* haptic device [Barbagli and Salisbury, 2003]. The Phantom was far from the only haptic device at the time; indeed, 40 pre-dated devices were identified by Margaret Minsky [Minsky, 1995], who herself did pioneering work on haptic texture-rendering on a novel joystick-like haptic device. The fact that the Phantom was mass produced and contemporary to a boom in computational capabilities, as well as growing multi-disciplinary interest in haptics, contributed to its status as being close to an archetype of a spatial haptic device. Today a small range of commercial haptic devices is available on the market, some of which are depicted in figure 2.3. They span a cost range between a few hundred euros and several tens of thousands of euros, and a more or less corresponding range in fidelity and capabilities in terms of sensed and actuated *degrees of freedom*, or DoF, referring to the number of dimensions the manipulandum can be moved/rotated and pushed/twisted respectively.

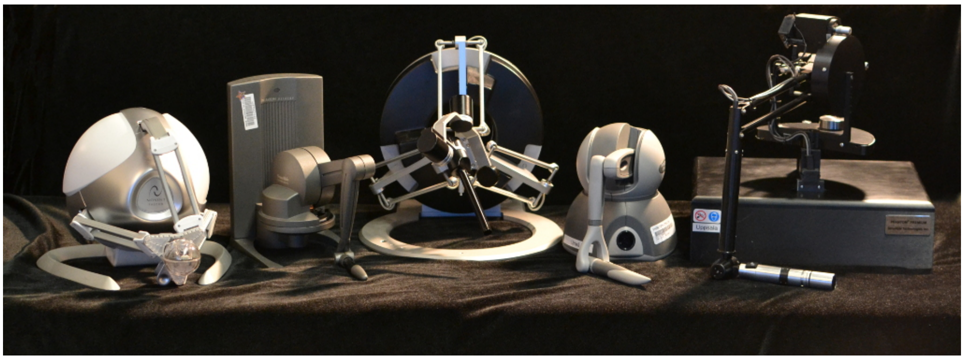

Figure 2.3: Commonly available haptic interfaces hardware. From left to right: Novint Falcon (3/3-DoF), Geomagic Phantom Desktop (3/6-DoF), Force Dimension Omega (3/6-DoF), Geomagic Phantom Omni (3/6-DoF) and Geomagic Phantom Premium (6/6-DoF)

For most application designers, the haptic device is treated as a black box. The designer is restricted to using one of the available pre-made devices. The choice of haptic device for a particular application has quite a high impact on the application's user experience. In certain circumstances it would therefore be meaningful to design and produce a custom device in order to get a certain resolution; e.g. to meet specifications derived from the nature of microsurgery [Salisbury et al., 2008]. However, engineering a high-quality haptic device is a large endeavour, requiring mechanical, electrical and computational know-how as well as tacit construction knowledge found only in few specialised robotics labs.

Quality Criteria for Haptic Devices

Massie and Salisbury have listed three main criteria applicable to haptic devices for use with virtual objects [Massie and Salisbury, 1994].

First, *free space must feel free*, meaning that ideally the user should not notice that the manipulandum is attached to anything restricting its motion in space. In reality, all mechanisms have some internal friction. In addition, the user may experience exaggerated weight and inertia (i.e. the motion-direction-resisting feeling of a heavy but weight-supported object), or backlash (i.e. the feeling of a gear transmission alternating between free motion and gear teeth resistance).

Second, *solid virtual objects must feel stiff.* Real-world *stiffness*, defined as the ratio of force over displacement, is very high for non-elastic solid objects. A wood plank deflecting 1 mm with a 10 kg weight on top corresponds to a stiffness constant, or *k-value*, of 100,000 N/m. Fortunately, much lower stiffness rendered by a haptic device is acceptable to perceive an object as relatively stiff. The original Phantom could render a stiffness of 3500 N/m, and users reported that stiffness of 2000 N/m could represent a solid wall [Massie and Salisbury, 1994]. Several devices cannot render such stiffness without causing stability issues, e.g. the well-used Phantom Omni can only render 800 N/m, making "hard" objects feel "mushy hard" as noted by Moussette [Moussette, 2012]. The ability to render high stiffness is an effect of the structure and material the device itself is made of, the quality of actuators and sensors and the control loop.

Third, *virtual constraints must not be easily saturated*, meaning that if pushing on a virtual wall or solid object with an increasing force, it should not suddenly let go, causing the manipulandum to "fall through". Even if constraint-based haptic rendering can avoid fall-through, it does so from a computational perspective. If a wall is rendered with some stiffness k, it will reflect a force that increases linearly with penetration depth x, but only up to the maximum force that the device can generate. What limits the maximum force in general is the motors; they become *saturated* at some limit torque, and if it is prolonged they can get overheated. A sufficiently powerful motor is therefore required for solid constraints.

Srinivasan and Basdogan have added two important criteria to the list, namely a) that there *should be no unintended vibrations*, bringing attention to the fact that unwanted vibrations are a common issue or trade-off in the design of haptic hardware, and b) that the interface *should be ergonomic and comfortable to use* since discomfort and pain supersede all other sensations [Srinivasan and Basdogan, 1997].

There are multiple ways to achieve high-quality haptic feedback according to the criteria above, although there are always trade-offs among them. Excluding non-contact-based spatial haptics based on, e.g., magnetic levitation or ultrasound, devices can broadly be classified according to mechanical structure and control paradigm. Serially linked manipulators, like the Phantom, consist of links and joints mounted, as the word suggests, serially, compared to parallel manipulators like the Falcon and Omega (figure 2.3). The control of the manipulator can be based either on *admittance* control or *impedance* control. Admittance control systems have a force sensor measuring the force the user is applying to the manipulandum and move the manipulator accordingly, while impedance control systems read the position and output a force when in contact with objects. All devices in figure 2.3 are impedance controlled. The HapticMaster is an example of an admittance-controlled device [Van der Linde et al., 2002]. The rest of this thesis will focus on serially linked impedance-controlled devices, since these are the most common. The relative structural simplicity, wide use and the fact that the same device exists in several similar but experientially different variants make the Phantom a particularly suitable object of analysis for the study of spatial haptic hardware for interaction design.

Principles Behind Phantom

Fundamentally, the Phantom, or any Phantom-like haptic device in general, can be described as a *mechanical manipulator* that consists of three actuated rigid links plus a base, connected by three revolute joints in a chain [Craig, 2005]. The Phantom also possesses a set of three additional passive links and revolute joints that form the *gimbal*, where the manipulandum is attached [Massie, 1993]. The manipulandum can be a thimble or stylus. Through sensing the angle of each joint and knowing the length of each link, the position and orientation of the manipulandum in Cartesian space $(x, y, z, \alpha, \beta, \gamma)$ can be determined through the mathematical construct *kinematics*. The actuated links are driven by computer-controlled actuators (motors) through mechanical *power transmission*. The torque to apply to the respective motor to exert a force vector to the manipulandum (F_x, F_y, F_z) can be determined mathematically via the principle of virtual work [Craig, 2005, p. 164].

Each of the three active joints is actuated by a direct current motor and uses wire rope for mechanical power transmission and gear reduction. This design has a number of benefits, including zero backlash due to avoidance of gears, low friction and good backdrivability, i.e. it is easy to move the manipulandum about even when the power is off.

Haptic Rendering of Solid CG Objects

In this thesis we are concerned with a haptic interaction system that enables the user to explore the shape of a CG object by moving a virtual sphere “attached” to the centre of rotation of the manipulandum and feel the repelling forces from contacts with the object. Ideally this force increases the more the user pushes, keeping the object impenetrable. As the user moves the manipulandum, they form a mental image of the shape of the object. This strategy is one of several *explorative procedures* humans use to understand the properties of a physical object using touch [Lederman and Klatzky, 2009].

The computational method that makes this sensation possible is called *haptic rendering* [Salisbury et al., 2004]. Haptic rendering can be drawn from control theory, collision detection and handling, interaction techniques and computer animation. The word *rendering* leads the reader to think of its analogue in computer graphics, where rendering is the process of representing graphical objects on a visual display. Correspondingly, haptic rendering represents objects on a *haptic display*, a synonym for a haptic device. However, since the haptic sense is inherently bi-directional (both sensing and actuating), so is the haptic device, and thus it is also the haptic-rendering algorithms' responsibility to move the avatar that corresponds to the manipulandum's position and orientation. In addition, a haptic-rendering algorithm has to act as a *controller* in its control theory sense, keeping the physical manipulandum stable and safe.

6-DoF and 3-DoF

The original Phantom had a thimble where the user put their fingers and whose rotation point effectively was co-located with the user's fingertips. This way the device afforded direct interaction - with one point - with the virtual environment. Present-day Phantoms, and all devices in figure 2.3, possess a manipulandum that the user holds with the hand. Consequently, the reasonable representation is that of a tool interacting with the environment, it is *tool-mediated*, like touching objects with a screwdriver or a wrench. Furthermore, this constitutes a full 6-DOF rigid body interaction between tool and environment, and the problem of computing the correct configuration and forces as results of 6-DoF contact (and detecting those contacts) poses a significant challenge. Approaching and solving this problem is referred to as *6-DoF haptic rendering* [Otaduy et al., 2013]. It has only recently been solved and requires knowing and handling sophisticated mathematical constructs which lead to a large implementation effort. The research community therefore naturally began with the more approachable problem of supporting haptic interaction with just the tip of a virtual tool, such as a sharp pencil, that effectively translates into interacting with the environment through a single contact point or sphere. These methods did not have to consider rotations or torque, and were therefore labelled as *3-DoF haptic rendering* [Otaduy et al., 2013, Forsslund et al., 2013].

Direct Rendering and Virtual Coupling

The bi-directional nature of a haptic system implies that a haptic-rendering algorithm is responsible for two tasks [Lin and Otaduy, 2008]:

1. Compute the configuration (i.e. position and orientation) of the on-screen avatar, given the configuration of the physical manipulandum and constrained by the virtual environment.
2. Compute and display (i.e. communicate to the haptic device) the forces resulting from contact between the avatar and the virtual environment (CG objects).

The most straightforward method to determine the avatar position and orientation is to assign it directly to the position and orientation of the manipulandum. This is referred to as *direct rendering* and reduces the haptic-rendering problem to only having to solve task 2 [Lin and Otaduy, 2008]. The prototypical example is that of rendering a virtual wall with a penalty force proportional to the penetration depth, e.g. $F = kx$, where x is the displacement and k is a *stiffness constant* [Ruspini et al., 1997]. The stiffness constant can be chosen arbitrarily within the stable limits of the haptic device and the update rate of the haptic algorithm. If k is set too high or the update rate too low, the manipulandum will either vibrate or be "kicked out" of the surface with an exaggerated force. A too low value, however, will feel rather mushy or spongy.

Direct rendering can be used to relatively easily implement haptic interaction between a spherical avatar, such as the tip of a dental drill, and a CG object with a cubic voxel volume representation. The method proposed by Agus et al. involves computing the penetration depth of the avatar and the surface area for determining friction, which both can be derived from computing the intersecting volume between avatar and object [Agus et al., 2003]. They have observed that if voxels are treated as cubes it would be computationally expensive to calculate the intersection. If instead they were interpreted as small spheres of the same volume as the cubic voxels, then the volume can be computed as a direct function of the distance, given that the radii of voxel-sphere and avatar-sphere are known. The sum of all fully or partly intersecting voxel-spheres gives the total intersecting volume. The corresponding surface normal can be estimated by a normalised weighted sum of the vectors from each sphere to the centre of the avatar sphere. Agus et al. have used these values as a step in calculating the forces using a physically motivated model; however they can also be used directly by letting the force magnitude be equal to a constant times the intersecting volume, and the direction be the normal. This has been done in the implementation by Forsslund et al. and has proved to be a sufficient method for application in a surgery simulation prototype [Forsslund et al., 2009]. This simplified method has, however, a number of drawbacks of which the premier one is that the avatar is not guaranteed to stay on the surface. In fact, if the user pushes the whole avatar sphere into the object, it may pop out anywhere. It therefore becomes extra important that a large enough force and stiffness can be displayed by hardware, since if the user pushes harder it will pop through. The direct coupling between manipulandum and avatar means that it will always look like the avatar is partly sunk into the object when the intent is only to explore its surface. This can be partly compensated for by visually displaying a slightly smaller sphere than that which is used for haptics.

Virtual Coupling

The principle of virtual coupling separates the motion of the manipulandum and the on-screen avatar, but connects them with a virtual spring and, optionally, a viscous damper. Through collision detection, constraints can be formulated that guarantee that the avatar always stays on the surface of objects. As the user pushes the manipulandum through the surface, the spring stretches and the corresponding force is reflected to the user. The same principles as those of the virtual wall apply, meaning that stiffness can be specified but is limited by the stability of the device. Early work on 3-DoF constraint-based rendering of polygonal surfaces through virtual coupling includes [Zilles and Salisbury, 1995] (under the name “god-object”) and [Ruspini et al., 1997] (“proxy”). The principle has also been extended to 6-DoF rendering of both polygonal [Ortega et al., 2007] and volume-embedded [Chan, 2011] surfaces.

Since the visual sense is more dominant than the haptic sense, the separation of manipulandum and avatar has the additional benefit of making objects be perceived as more stiff and impenetrable than what they mechanically are from the haptic device perspective. In other words, while the user may push the manipulandum slightly into the surface, the visual avatar will show a full stop.

Stable and stiff 6-DoF haptic rendering is an ongoing research problem, although several algorithms have been proposed. Several aspects make it challenging, including the need for high update rates for stability reasons (at least a 1000Hz for the central control). Furthermore, most 6-DoF and constraint-based algorithms rely on sophisticated mathematics and methods rarely encountered by the average software developer.

Carving

What has been discussed so far in this chapter are technologies that enable digital representation of solid objects and tool-mediated haptic interaction with their surfaces. This is rarely sufficient for a purposeful and useful application. To enable the user to perform tasks with the system, certain methods need to be implemented that interpret the user’s action and perform the task. These methods are called *interaction techniques* [Bowman et al., 2004]. Most people are familiar with interaction techniques in the 2D desktop metaphor, such as moving an arrow-shaped cursor with a 2D input device called the mouse, hovering over the icon of a folder and double-clicking the mouse button to open up a window showing its contents. The field of 3D user interface design correspondingly deals with 3D interaction techniques, such as selection and manipulation of objects whether such tasks are completed with 2D input devices via widgets or with direct manipulation using 3D input devices [Bowman et al., 2004]. It may be already worth mentioning here that the notion of an *interface* can be problematic if too much attention is paid to the interaction that is happening through this interface at the expense of off-line interaction that, as will be described in subsequent chapters, can be at least as important to design well.

Carving is an interaction technique that has been used for virtual sculpting [Galyean and Hughes, 1991, Wang and Kaufman, 1995]. Avila & Sobierajski have presented a method for carving (they use the word “melting”) with haptic feedback, as part of a volume visualisation system [Avila and Sobierajski, 1996]. The force feedback is posed as particularly useful to understanding spatial structures and to using the device for input in modifying the visibility of different structures. Example images from their system include drawing and cutting into a human skull.

Carving has subsequently been used for simulating surgery [Pflesser et al., 2002], and the field is starting to focus on developing physically correct models [Petersik et al., 2003, Agus et al., 2003, Chan, 2014]. An important aspect of these works is that different structures in the CG object (i.e. skull) can be assigned different densities, meaning that some structures take a longer time to carve; they are perceived to be of a harder material.

As with the section on haptic rendering, there also exist simplistic methods for performing the carving deformation. The implementation by Forsslund et al. is inspired by Agus et al., but uses just an 8-bit counter per spherical voxel, which is decreased with an amount equal the hardness or *cut ratio* factor defined for the segment this voxel belongs to [Forsslund et al., 2009].

2.3 Tools for Haptic Interaction Design

Haptic interaction design is still a young field, but it has been established that designing for the haptic sense requires unique considerations and sensitivity for the modality [MacLean and Hayward, 2008, Moussette, 2012]. Sketching and prototyping are central to any design discipline [Lawson, 2005, Buxton, 2007, Nelson and Stolterman, 2012]. It is through creating representations and the materialisation of ideas that the designer can have a *conversation with the design situation* [Schön, 1984]. Digital materials are also used to this end [Dearden, 2006, Löwgren, 2007, Fernaeus and Sundström, 2012]. Much interaction design deals with the conceptual, i.e. what should happen when, as an effect of what etc. Conceptual design can often be sketched out with pen and paper in the form of storyboards, scenarios or mock-up interfaces. One can easily imagine drawing a scenario where a character, sitting in a meeting, feels her phone vibrating three times to indicate a message from her mother. Many examples exist where haptics is used to communicate symbolic meaning, e.g. requesting attention in critical settings where other senses are occupied [MacLean, 2008]. All haptic interaction, symbolic or not, also has a qualitative dimension, a rich subtleness of nuances that is difficult to capture in words. Consider the following quote from Donald Norman's *Emotional Design* [Norman, 2005, p. 79]:

> "Just turn the knob," I'm told, as something is thrust into my hands. I find the knob and rotate it. It feels good: smooth, silky. I try a different knob: it doesn't feel as precise. There are dead regions where I turn and nothing seems to happen. Why the difference? Same mechanism, I am told: the difference is the addition of a special, very viscous oil. "Feel matters," a designer explains, and from the "Tech Box" appear yet more examples: silky cloth, microfiber textiles, sticky rubber, squeezable balls - more than I can assimilate at one experience.

What Norman illustrates is the power of *visceral design*. The two knobs may be used to do the same thing, but the feeling of one is much more satisfying. Imagine the impact it can have on the volume control of a high-fidelity stereo. He also illustrates the practice in the industrial design agency he visits; they carry a "Tech Box" filled with exemplars for use in their designs, a *tool* for inspiration and sense-based exploration.

In contrast to the knob in the quote above are the haptic experiences that are subject to design, in this thesis primarily *synthetic* and computer controlled. This implies that they are only materialised and experienceable when fully implemented. How, then, have designers approached designing for synthetic haptic experiences? The following sections will first cover the use of representations, where substitute materials are used to craft prototypes of various fidelity that can give an idea of the final experience. Some accounts will be given as to why the use of representations is inherently limited for haptic interaction design. Second, some tools that have been proposed in the literature will be described that let the designer work closely with the realisable synthetic experiences. Finally, toolkits of components that have been proposed as useful for designers to create haptic experiences of will be discussed.

Representing Haptic Experiences

Prototypes can be classified according to their *fidelity*, i.e. how well the representation reflects the finished product. They can also be classified according to scope, i.e. whether all features are covered (horizontal prototype) or whether just a particular aspect is implemented (vertical prototype). A high-fidelity prototype is usually closer to the finished product, and user evaluation of such a prototype is expected to better predict the end result. A low-fidelity prototype, on the other hand, is expected to require less time and cost to produce, which is why they are often favoured in early development phases and for quick generation of multiple alternatives [Bjelland and Tangeland, 2007].

Miao et al. are working on designing tactile 2D “graphical” user interfaces for use by the blind. They are designing for a particular 120x60 pin display, and propose a method to create paper prototypes using embossing printers [Miao et al., 2009]. The users feel a representation of the application prototype in another material (embossed paper) rather than the target material (metallic pins). This is motivated by the assumption that it is quicker to change the paper mock-ups according to users’ comments [Miao et al., 2009]. I argue that, while it has its benefits, including the storage of mock-ups in binders and easy multiplication, there are two major drawbacks to this approach. First, the mock-ups have to be fabricated with a special machine, and therefore lack the directness of pen-and-paper prototyping. Second, the similarity between embossed paper and metallic pin array may not be as close as anticipated. Apart from the obvious, in that paper feels different to metal, it is also the case that the paper is static and uni-directional (“output” only), while the particular pin array display they design for is refreshable and bi-directional in that it also has touch sensors. It thereby supports input gestures, although the user has to press a peripheral button for it to distinguish between user reading and user actions [Prescher et al., 2010]. O’Modhrain et al., who themselves are visually impaired, argue that it is critical that *transcribers*, i.e. designers working with converting conventional media to tangible media for the visually impaired, accurately understand and utilise the rendering capabilities of the device to be used [O’Modhrain et al., 2015]. This is a precise art that requires selecting and matching the most important information to the perceptual channels available to the blind user, and doing this within the constraints of the materials, digital or not, available for rendering [O’Modhrain et al., 2015]. In other words, what is a good match for embossed paper output may not be good for a bi-directional pin array display and vice versa.

Kern [Kern, 2009] suggests using everyday objects, such as fruits of varying stiffness, as representations for haptics requirements gathering, especially in dialogue with domain stakeholders. The assumption again is that once the requirements are known, they can be engineered for. To some extent it can be useful to utilise props to learn the desired range of motions and for establishing a common ground in how objects ideally should feel. But I would argue that there is a large risk that the engineered solution deviates from the feeling of a fruit, is too costly or in other ways substantially differs from the designer's plausible naive intention. It is also unclear what quality of the fruit should be engineered for: the surface deformation, the overall stiffness or the smell. In any case, if it so turns out that the synthetic feeling does not match the feeling of the fruit, there is no clear prescription of what to do next.

Noting the great dependency of the nuances of haptic properties in the user experience of haptic interfaces, Bjelland and Tangeland discourage using low-fidelity prototypes [Bjelland and Tangeland, 2007]. They recommend prototyping through technology substitution, i.e. using analogue mechanical devices or modifications of already existing electronic products. As a case study they prototyped a ship's throttle controller with force and vibrotactile feedback. The prototype was created through physically modifying a commercial low-cost force-feedback steering wheel and adding transducers from a force-feedback mouse. Haptic feedback effects were designed using graphical representations of the vibrotactile waveforms. The prototype helped the development team by providing immediate feedback on changes and formed a common vocabulary of the haptic effects. However, they also note that it was difficult to predict the user experience of the final system. While it is not explicitly stated, it can be assumed that the final product was expected to be developed with other, plausibly more robust and high-quality, components. The value of fine-tuning the feeling for low-quality transducers should therefore be questioned. The design software could probably be re-used for the high-quality device though.

I argue that the largest problem with the approach of using representations is that it can give false impressions of what is feasible or even possible to implement. It also shifts the responsibility of a good user experience from the designer to the engineer in charge of the implementation. As noted by Fernaeus and Sundström [Fernaeus and Sundström, 2012], there is a risk of trivialising the technological choices required for a good design. If the result is bad, it is sometimes implicitly understood that it is just a technical bug or mishap that could have been resolved if only the engineer had been better at *their* practice, and that this practice should not be of concern to the designer [Fernaeus and Sundström, 2012]. Critique of this view has motivated a *material move* in HCI, where more emphasis is placed on the importance of digital *materials knowledge* for the designer. Digital materials are here understood as both software and hardware; in other words the parts that a product is made of. This insight can be used for synthetic haptic design in two different ways, as will be discussed in the following two sections. The designer can get direct feedback from the *target material* during a design exploration through the use of particular design tools, created for the purpose, or the designer can adhere to *design through making*, in particular through the use of ready-made components and kits that help in getting quick feedback on user experience, also in the target material.

Haptic Design Tools

Computer-controlled haptic systems have a digital or signal part that itself is subject to design. This is well understood for vibrotactile devices; not only is the binary action “it’s vibrating!” interesting, but also the frequency, rhythm and pattern [MacLean, 2008].

Swindells et al. [Swindells et al., 2006] describe a software tool for designing the force feedback of a 1-DoF haptic knob. It includes a waveform editor, a palette of short effects (haptic icons) and the possibility of combining them into patterns in a similar fashion to a music sequencer. A particular example application is given: a fan control knob with four “clicks” felt at different angles, with a particular resistance between them. This is a good example of designing the fan knob feeling and getting direct feedback in the process, given that the final fan knob is using the same components.

Ledo et al. [Ledo et al., 2012] have developed a tool for interactively designing the vibrotactile haptic feedback of a tabletop hand controller. The hand controller, named Haptic Tabletop Puck, is used on a large tabletop touch screen and can, in contrast to the screen itself, provide haptic feedback. The optically tracked puck has a pressure sensor on top of a servo motor-controlled vertical rod placed on top of the puck. The rod’s height and compliance can be controlled in the software. The puck also has a servo motor pushing a rubber plate against the table that can be used to control mechanical friction. These features can be programmed directly, or can be tuned in a “Behaviour Lab” application. Different areas of the tabletop can then be assigned different effects or rod compliance. The Behaviour Lab allowed the developers to explore and feel available forms of haptics feedback before developing a complete application [Ledo et al., 2012]. Schneider and MacLean developed a device similar in style to a musical instrument that simultaneously controls two separate vibrotactile bracelets [Schneider and MacLean, 2014]. A user and a friend could then wear one bracelet each and both users could feel the vibrotactile feedback “played” by the user holding the instrument. This allows for improvisation and sharing of experiences without having to rely on other forms of representations.

De Felice et al. [De Felice et al., 2009] have created an authoring tool for assigning audio and haptic effects and properties to pre-existing virtual 3D environments for use by blind subjects. The environment resembles an office floor plan with rooms and doors, where the geometry has been defined using conventional 3D authoring tools. A scenario designer using De Felice’s tool can select, e.g., a door and assign it a relative stiffness and friction, and also add the effect of vibration when it is touched as a means for the blind user to distinguish the door from the walls. The resulting scene description is stored as an XML file. The stated purpose of the tool is that a domain expert without technical virtual reality skills can design the haptic and audio parts of the virtual environment. The device they use is a Phantom Omni, although others are supported. The design tool’s user interface is two dimensional and uses context menus. They report on a workflow where the scene was first defined using the authoring tool, then evaluated together with a user and then edited again. I argue that since the execution was thereby decoupled from the design, it added a time delay to each design iteration that may have limited the ability to get a feel for the material through direct experimentation. Haptics was also viewed as an add-on “special effect” to the pre-made geometric model.

Bone cutting rate: 0.55
Tooth cutting rate: 0.06
Bone avg. cut force: 0.50
Tooth avg. cut force: 0.50

Drill diameter: 3.0 mm
Scale: 2.90
PID: 1.00

Applied force: 0.16
Frame rate: 188

Figure 2.4: Designer tuning haptic-rendering parameters using a physical controller

Forsslund and Ioannou [Paper C] have presented a tool (figure 2.4) for sketching haptic carving applications, catering in particular for the exploratory design of haptic-rendering properties that affect the feeling of different layers of CG objects, i.e. the haptic “materials” they are made of. The central haptic properties are object scale, stiffness and carving rate, along with visual properties of colour and transparency. The design tool enables designers to tune the properties and feel the result in real time. This is accomplished with a tangible mixer-board-like input device where each property is associated with one slider or knob. These can be used with one hand while the other hand is holding on to the haptic device’s manipulandum, directly feeling the result of changes in the CG objects’ material properties. Different variants can be stored and recalled directly from the controller, which can be used in dialogue with domain experts for quick exploration of alternatives.

Panëels et al. have developed a visual programming tool for prototyping haptic visualisation applications. It is targeted at non-programmers and generates Python code to be interpreted by H3D API, a toolkit discussed later in this chapter [Panëels et al., 2010]. The focus of the tool is not on designing the actual haptic feeling, but logical behaviour such as adding and removing magnetic lines as a result of keystrokes. The designer does this through dragging and dropping boxes in a 2D GUI to form logical paths. Other advocates of visual programming are Rossi et al. [Rossi et al., 2005], who propose avoiding C++ programming through the use of a rather sophisticated multi-application, multi-computer set-up using Simulink, a professional mechanical/electrical engineering tool. The example application is a sphere stretched to an oval. These tools may be useful to some, but I argue that they offer little value to a modern interaction designer for several reasons. First, a modern interaction designer can fairly be expected to know how to implement event handling using common high-level programming languages, and learning a new one, even if it is visual, risks costing more time than it frees. Second, the focus seems to be on events and actions rather than the feeling, helping little in this regard.

Toolkits for Crafting Haptic Applications

Toolkits have been highlighted as important instruments for interaction design in several areas of human-computer interaction. Phidgets [Greenberg and Fitchett, 2001] are a popular example of a toolkit for physical interaction design. As the name hints, they are an extension of the concept of the graphical user interface (GUI) widgets to a physical interaction design: physical widgets. A widget is a ready-made component, e.g., a slider or a button that can be easily integrated into a GUI. Phidgets relieve the designer of solving many technical nuances of how to light up a diode, control a step motor or receive input from switches, light sensors and the like. Compared to, e.g., Arduino, a stand-alone platform typically used in physical computing, Phidgets are connected directly to a desktop computer and as such extend the interaction beyond desktop computing. Greenberg et al. show how Phidgets have been used by students in creative design explorations through various examples.

Interfaces such as Phidgets support designers beyond facilitating the realisation of their design ideas. They can facilitate *sketching in hardware*, where non-committal design explorations can be created hands-on [Moussette and Dore, 2010]. One of the benefits is that designers thereby get a direct feel for the qualities that the components afford. This can be used to gain a *heightened sensitivity* to a design material such as haptics, but only if the designer is working actively with and through the material [Moussette and Banks, 2011]. Moussette shows commitment to this conviction with his *Simple Haptics* concept [Moussette, 2012]. This approach builds upon creating and exploring real haptic experiences, not as representations of some future product, but as an end in itself. Examples include making hand-held boxes made of laser-cut plywood with a motor spinning an unbalanced wheel of various weights and shapes. Moussette's *simple haptics proposition* is succinctly described by himself in the following quote [Moussette, 2012, p. 215]:

> Simple haptics consists in a simplistic, rustic approach to the design of haptic interactions. It advocates an effervescence of direct perceptual experiences in lieu of technical reverence and dutiful attention to empirical user studies. Simple haptics boils down to three main traits: 1) a reliance on sketching in hardware to engage with haptics; 2) a fondness for basic, uncomplicated, and accessible tools and materials for the design of haptic interactions; and 3) a strong focus on experiential and directly experienceable perceptual qualities of haptics.

I argue that while tools and material are described by Moussette as uncomplicated and accessible, they are so from the perspective of a rather modern, up-to-date, physical interaction designer. This includes light Java programming skills and fluency in using platforms such as Arduino, electrical components such as step-motors and manufacturing tools such as laser-cutters. The simple haptics proposition does not require any particular tool custom-made for haptics. Theoretically any tools and materials could be used. The strength of simple haptics lies in the simplicity in which the medium is approached rather than the simplicity of the actual components used. The WoodenHaptics [Paper B] project described below aims to bring some of the simple haptics philosophy to advanced haptics.

Knörig et al. [Knörig et al., 2009] have developed a software tool for electronics prototyping, making it less dependent on electronics engineering proficiency, improving robustness and facilitating collaboration. Mellis and Buechley [Mellis et al., 2011, Mellis and Buechley, 2012] show how a modular kitchen radio presented as *open source hardware* can assist designers in crafting their own radio, changing only those parts pertinent to the designers' particular interest. In another project they show how conductive ink can be used to create interactive compositions of micro-controllers and paper-compatible materials that build upon established papercrafting practices [Mellis et al., 2013]. Hartmann et al. [Hartmann et al., 2006] have developed a tool for the design of electronic products that integrate design, user testing and analysis in one system.

Work to support *personal fabrication* [Gershenfeld, 2008] can also be of the utmost usefulness for the designer who elects to work closely with the physical materials of the final product rather than with representations. Personal fabrication with digital tools allows for shorter iterations than with mass-production methods. Tools like 3D-printers and laser-cutters have often been labelled "rapid prototyping tools" for their use in corporate product development, but they can also be used when the "prototype" is the end result. The digital nature of these tools, as compared with handcraft tools, enables quick alternations since the drawings can be changed and the fabrication is relatively cheap and quick. This also facilitates dissipation and community building, since digital drawings can be shared easily. Recent work in the field includes bringing the designer into close contact with the fabrication tool, rather than first working with a Computer Assisted Design program on a computer [Mueller et al., 2012]. Others have made special purpose software for enabling the design and fabrication of, e.g., a piece of furniture without extensive carpentry skills [Saul et al., 2011].

Some kits have been developed for teaching and learning haptic engineering. Shaver and MacLean [Shaver and Maclean, 2005] have developed the Twiddler, an affordable 1-DoF rotational device consisting of an electronics box and a motor that can be attached to a conventional computer and programmed by students. Its design is fully documented for reproduction and teaching of its inner workings. Slightly more advanced but with the same teaching goal is the 1-DoF Haptic Paddle, which extends the motor with a link through a cable transmission similar to the Phantom [Okamura et al., 2002]. Since its conception around 1996, it has evolved and is now used in several schools, including for on-line teaching, where kits are sent to students for their web-guided assembly [Richard et al., 1997, Morimoto et al., 2014]. The intended use of the Haptic Paddle is as a laboratory tool in engineering dynamic systems, i.e. formulating equations of motion and applying control theory. It was reportedly very welcomed by students in that they could get a direct feel for the otherwise abstract concepts. In contrast to kits like Phidgets, it is not intended to encapsulate technical details; rather the opposite, making them all transparent for learning, experimentation and modification. As a learning platform it is not intended for direct use in the design of complete applications.

WoodenHaptics is a 3-DoF haptic device similar to the Phantom, packaged as a starting kit for design explorations [Paper B]. Instead of teaching technical details, it is designed to encapsulate these in modules; e.g. the electronics box that is connected between the computer and the mechanism. The equations of motion etc. have already been solved for and the device is ready for use in software applications through a common application programming interface in the same manner as commercial devices. The designer may switch components or change link lengths, and only has to modify corresponding values in a text file. The device is designed with modifiability in mind; e.g., it is easy to swap motors since they are attached through flexible couplings that are clearly accessible from the outside rather than being embedded inside the device like most commercial devices. Since the kit itself is open source, it opens up for deeper modification, including its electronics box, for those designers who are so interested, but it is not necessary for most applications. Being open source, it supports designers in the same way as Mellis and Buechley's kitchen radio, encouraging modification of the parts pertinent to the designers' interest [Mellis and Buechley, 2012]. The key point of WoodenHaptics, however, is that different designs can be crafted and tried out easily. The kit will be discussed in greater depth in chapter 4.

Software Toolkits

The predominant way of encapsulating computer software technology for use by application developers is through software libraries, software development kits (SDKs) and application programming interfaces (APIs). They generally facilitate software construction through providing *abstraction layers* where the designer does not have to handle the details of lower levels. The flip side is loss of control, and sometimes understanding, of what is going on. The designer is also restricted to the programming language or interface that the toolkit developers have decided on. The benefits of software toolkits usually outweigh the alternatives for a number of reasons, including:

1. They implement and encapsulate complex algorithms so the designer does not need to implement them from scratch.
2. They usually provide example programs that can be used as inspirational bits [Sundström et al., 2011] and, if well-commented code for the examples is provided, be useful in showing how a technique can be implemented with reasonable effort.
3. They can provide functionality for getting started, e.g. setting up a window, drawing simple graphic elements, keyboard handling and so on, thereby saving time and preparing the designer for experimentation instead of having to engage in marginally relevant problem-solving.
4. Depending on architectural structure and licensing terms, they can be used as the basis for extensions, e.g. implementation of new haptic-rendering algorithms.

5. The API itself can, if it is open source, be used for studies in software architecture.

An early software toolkit for spatial haptics was the C++-based General Haptics Open Software Toolkit (GHOST) developed by the manufacturer of Phantom [Burdea and Coiffet, 2003]. GHOST has no own visual-rendering capabilities. It provides a synchronised haptic-rendering loop and maintains a *scene-graph*[2] that the user can populate with rigid objects to be rendered. A scene-graph is a popular and powerful way of structuring scene elements in a hierarchy; e.g. a teapot object can be placed on a table, and if the table is moved the teapot follows. The successor to GHOST is the OpenHaptics Toolkit, which is split into a device-level library and a library that integrates with the low-level graphics API OpenGL [Itkowitz et al., 2005]. It is still rather low-level and so requires substantial C++ programming to create a useful application. In addition, despite its "open" name, it is neither open source nor does it support other haptic devices. The need to support different devices in a uniform way and for use in a haptics course motivated the development of the open source CHAI3D toolkit in 2002 [Conti et al., 2007]. While this toolkit requires the use of C++, it abstracts OpenGL calls and is designed to be easy to get started with through numerous well-documented examples and by intentionally being a small library. The provision of boilerplate code (template code with comments of where to change things) for custom haptic devices facilitates implementing support for the WoodenHaptics device in CHAI3D. While CG objects can be dynamically loaded in CHAI3D, the scene and its interaction have to be hard-coded or provided by other means.

H3D API is an open source haptics and graphics toolkit that can be used on multiple abstraction levels. It relies on the extendibility of X3D, a web standard for describing 3D scenes in a human-readable, hierarchical graph in a text file, usually XML [Daly and Brutzman, 2007]. H3D API is distributed with an executable application that loads a user-defined XML file. This shifts development from imperative programming in, e.g., C++ to declarative programming of the scene, its objects and relations, similar in style to editing HTML files for web browsers. Most of the features of the X3D standard ISO-19775 are supported. This allows for compatibility with textbooks, e.g. [Brutzman and Daly, 2007], and X3D editors. X3D can expect an increase in popularity through being compatible with modern web browsers without plug-ins [Behr et al., 2010]. Behaviour can be programmed using the X3D "route" feature or, in H3D, through Python bindings. The X3D standard only covers visual aspects, but allows for extensions, which H3D uses to describe haptic-rendering properties of objects. Stiffness is one such property, and it is per default set at a relative value between 0 and 1, where 1 is the maximum stiffness the haptic device can handle[3]. H3D API is a large toolkit building on top of several smaller libraries and a sophisticated architecture. This can, I argue, make extensions such as implementing new algorithms less straightforward than in CHAI3D. The declarative nature also requires working in the corresponding paradigm for handling dynamic behaviour, something that can be new to developers more experienced in imperative programming (e.g. Java).

Forsslund et al. have developed an extension to H3DAPI called *forssim* [Forsslund et al., 2009]. This API implements a volume haptics algorithm inspired by [Agus et al.,

---

[2] in the case of GHOST it was restricted to being a tree, i.e. no node could have multiple parents

[3] H3DAPI 2.3 Doxygen Documentation, class Smooth Surface, www.h3dapi.org accessed 2015-09-24

2003], eventually also a constraint-based algorithm [Chan, 2011], a graphical rendering method [Wijewickrema et al., 2013] and some other techniques such as support of surgical procedure "game logic" through a basic state machine.

CHAI3D and H3DAPI both support several different devices in an application agnostic way, meaning that an application developer can write his/her application without explicitly specifying which haptic device it should be used with. If the developer so wishes, the virtual workspace can even be scaled to account for differences in hardware workspace, meaning that a device with a small workspace, such as the Novint Falcon, can be used to touch the same CG object (e.g. a teapot) scaled down, as a Phantom Omni can, occupying its whole workspace. The use of normalised workspaces and relative stiffness properties can be misleading to designers. Camille Moussette notes in a design exploration with H3DAPI and the Phantom Omni that while he had defined a virtual surface to be of maximum stiffness, it still felt mushy [Moussette, 2012]. As mentioned earlier, the stiffness is by default defined in relation to the maximum stiffness supported by the particular haptic device employed. I therefore argue that it is imperative for designers to consider both software and hardware aspects simultaneously when designing haptic applications.

Based on several design cases, in particular a dental simulator, Forsslund et al. [Paper A] argue that having access to ready-made technical components is required, but not sufficient, for effective hands-on design exploration. They present *visuohaptic carving* as an interaction technique and design resource for use in multiple applications [Paper D]. Designing the synthetic feeling of carving requires both a force-reflecting haptic device, a collection of algorithms for haptic and visual rendering of CG objects and their deformation, and a visual display. The algorithmic need can be fulfilled with H3D and Forssim libraries, and the hardware bought off the shelf or built using, e.g., WoodenHaptics [Paper B]. Creating a particular application where the user can interact with something more interesting than, e.g., boxes and spheres requires authoring CG objects, converting them to appropriate voxel format for rendering, and setting the right parameters in the text file defining the scene. Doing this manually, i.e. through command-line calls and a process that relies on editing text-files, starting the application, testing, closing, and re-editing the files, is error prone and impedes creative exploration. Therefore a combination of ready-made components (forssim) and specific design tools (figure 2.4 and scripts) have been developed. The tools can then be used to tune parameters, and scripts can be used for file conversions that let 3D artists work with tools they are familiar with. The design of the tools and workflow is not speculative in terms of what can be created with them, but directly supports the real-world surgery simulator project. The tools, and corresponding strategy as well as the design of the Kobra oral surgery simulator [Paper A] will be discussed further in chapter 4.

2.4 Surgery Simulation

The first section in this chapter discussed haptic technology that enable users to touch, see and carve CG Objects. This was followed with a discussion on how design tools and the collection of technological components have been proposed and used for turning the aforementioned technologies into prototypes, systems and applications by designers. This section will explore a popular application of spatial haptics: surgery simulation. The design and development of a fully functional oral surgery simulator [Paper A] has served as a vehicle motivating and informing the other parts of the thesis why a discussion of this domain is highly relevant.

Surgery simulators and their uses have been the subject of different discourses. The development of novel simulators has predominantly been technology-driven, with advances in computer science and engineering [Chan, 2014, Agus et al., 2003, Cormier et al., 2011, Riener and Harders, 2012]. Indeed, much of the technological development discussed in chapter 2 is directly motivated by its application to surgery simulation. Simulators have also been subject of studies on learning efficacy, i.e. whether training with a simulator has improved the skills or knowledge of the learner [Joseph et al., 2014], and of efficiency, i.e. whether teaching with simulators requires less resource in terms of human instructors or time than alternative methods [Bakker et al., 2010]. The usage of simulators in teaching contexts has been the subject of science and technology studies [Prentice, 2005, Johnson, 2004, Hindmarsh et al., 2014, Johnson, 2007]. Simulator design has rarely been discussed as a subject in its own right, but can be found as a part of most technological papers.

The purpose of this chapter is to introduce work related to surgery simulator design and some results from the aforementioned fields of study that are relevant to the design of simulators. First, a brief survey of studies and scholarly thought on the actual usage of simulation in healthcare will be given. Then a number of simulators developed for the dental education domain will be examined more closely. These simulators build upon much of the technology discussed in chapter 2, and the technology contribution of each will therefore not be repeated. The focus will instead be on their design. The survey of simulators constitutes the related work of the simulator named Kobra (figure 2.5), which I myself have been the lead designer and developer of [Paper A].

Simulation in healthcare

The use of simulation in healthcare spans a large spectrum of practices and medical disciplines, including role-playing based simulations with passive mannequins for team training, part-task trainers with physical props for, e.g., practising suturing, interactive web-based scenarios, and professional actors imitating patients and relatives for both diagnostics and communication training [Levine et al., 2013]. In virtually all these settings the simulator equipment is only one aspect among many that influences the quality of the educational experience. I argue that understanding the simulator's use in context can support designers in two key ways: for need-finding and for leveraging the context and human instructors in such a way that the costs or complexity of the equipment, i.e. the simulator, can be reduced. Qualitative studies in social sciences can be instrumental to this end.

Rystedt and Sjöblom [Rystedt and Sjöblom, 2012] analyse the simulation practice of two different simulators: one desktop GUI-based anaesthesia simulator and one full-body mannequin-based trauma team training simulator. The first one resembles a typical point-and-click software application, with icons, symbols and numerical displays showing the state and condition of the "patient". The user can administer a specific amount of drug

Figure 2.5: The Kobra Oral Surgery Simulator. Illustration of instructor and student solving a patient case.

injection through dialogue boxes, and observe values for blood pressure, heart rate and the like. These resemble patient-monitoring equipment found in an operating room. The participants (two students and a teacher) who were observed using these abstract controls discussed the actions in medically relevant ways not necessarily found on the screen, e.g. the dilation of blood vessels as a result of increased provision of anaesthetic gas. The second simulator, a full-body mannequin, was used as part of team training in response to the urgent care of a hospital-admitted car accident patient. A large part of the simulation required active participation and role-playing, where some participants were new to the scenario and others were assisting in playing it out. One scenario involved insufficient information. The doctor who had been called to the site would in routine practice have been informed of the patient's state by first visiting the radiological department, but in this scenario he had no such information. The participants instantly made up a story where this could have been likely, and continued the scenario. Rystedt and Sjöblom show through these examples, based on observations and dialogue analysis, how all participants in both these simulation practices take active part in creating the learning experience, focusing on relevant similarities with clinical practice and ignoring what is irrelevant:

> "In order to go on with the simulation as a simulation of a clinical practice presupposes, though, that the participants continuously constitute and reconstitute *what* the simulation is a simulation of, what the relevant *similarities* are for furthering the activity, and, finally, what *differences* between the simulation and perceived referent are actually irrelevant for the situation at hand." [Rystedt and Sjöblom, 2012].

It is thus evident that some differences between the simulator practice and clinical practice are perfectly accepted and sometimes even welcomed, to, e.g., not distract participants with everything in an operating room but only a subset relevant for the learning activity at hand.

The instructors have a particular role in creating clinically relevant practice out of the simulator practice. One way instructors do this, as observed by Johnson et al. [Johnson et al., 2004, Johnson, 2009], is through *reconstitution* of the absent patient's body. Johnson observed how an instructor using a minimally invasive surgery[4] simulator for knee surgery explains how the corresponding patient's body is oriented in relation to the simulated view on the monitor. The instructor was observed using his own knee, positioning it and pointing in which directions the camera was oriented. This way a more complete body was reconstituted than what was presented by the simulator alone. Johnson points out that even if one could imagine equipping the simulator with a mannequin leg, that would not necessarily be required since the instructors could leverage reconstitution [Johnson, 2004].

Furthermore, simulator practice is highly *situated* in the context of a teaching hospital [Johnson, 2004]. This entails that instructors, who are themselves usually practicing surgeons, relate the simulated activity to clinical practice in a broad sense. This can be telling anecdotes and reassuring struggling students that something that is difficult in a simulator can be difficult in the clinic too. Part of the experience is to initiate the student into a professional community of practice, i.e. turning the apprentice into a professional surgeon. In this regard it is not a problem, Johnson argues, that, e.g., an instructing surgeon gets called away in the middle of a session; it can rather be seen as a feature illustrating the professional reality of working in a hospital. Johnson also suggests that using props such as wearing scrubs can be used to strengthen the simulated activity. There are, in other words, many things that the users can do to ensure that the otherwise technical practice in the simulators is transformed into medically relevant practice [Johnson, 2004].

---

[4] "key-hole" surgery where the surgeon operates with long instruments and a camera inserted in body through smaller incisions than in open surgery

Dentistry educators have been early adopters of hands-on simulators, beginning with mechanical ones where students can drill with real equipment on synthetic teeth, to computer haptics-based virtual reality simulators [Gottlieb et al., 2013]. Hindmarsh et al. [Hindmarsh et al., 2014] have observed the use of two different dental simulators, one mechanical and one prototype visuohaptic simulator (hapTEL, examined below), with a focus on how participants discuss and explore the differences between simulator practice and clinical practice. The dialogue between teacher and student has been recorded and analysed. Hindmarsh et al. show how the teacher explains the benefit of resting a finger in such a way (e.g. on teeth adjacent to the operation) so that if the patient suddenly moves the student still has control. This is helping the student to prepare for real clinical practice and not just the simulator practice where there will be no sudden movements of the fixed mannequin. Another example is how the teacher corrects a student who is not using the water spray, by questioning and explaining how a corresponding real tooth's nerve could be hurt by the otherwise high temperature generated by the drill. The "absent body of the patient is invoked explicitly in assessing the student's practice" [Hindmarsh et al., 2014], whereby students are assessed and taught in relation to the clinical socio-material environment rather than the simulator practice alone. Similar observations were made with the visuohaptic simulator; rather than teaching how to remove brown pixels on a screen, the dialogue concerned removal of caries in a particular way that minimises the risk of bacterial infection.

Simulator Design for Use in Dental Education

Most dental schools today have laboratories where students can practice teeth preparation using mechanical simulators. These simulators consist of a workbench with real dental drills and other instruments and a mannequin head with disposable synthetic teeth. The students are instructed in how to prepare the teeth, i.e. using the dental drill to carve a shape for applying a filling. They can later show the result by handing the carved plastic tooth to an instructor for grading [Buchanan, 2001, Gottlieb et al., 2013].

The DentSim simulator (Image Navigation, New York, USA) (figure 2.6 left) is based on such a mechanical dental simulator with real drills and plastic teeth. The system optically tracks the position and orientation of dental drills in relation to the synthetic teeth and can thereby reconstruct corresponding virtual motion and carving for display on a screen [Gottlieb et al., 2013]. This enables real-time quantitative measures of student performance and instructional feedback during the tooth preparation. A study comparing use of the computer-assisted system with one with only the mechanical part shows that instructor time can be reduced by a factor of five thanks to the computer-based instructions and feedback [Jasinevicius et al., 2004]. This simulator benefits from providing haptic feedback "for free" through the analogue contact between drill and plastic teeth. The downside is that it requires disposables, and that exercises are limited to the range of teeth available. Usually the head lacks bone, which would be required for surgery.

Simodont Dental Trainer (Moog Industrial Group, Amsterdam, Netherlands, figure 2.6 right) is a haptic-enabled virtual reality-based dental simulator developed during the same time period as the Kobra. It uses an authentic-looking dental handpiece (drill) attached to

Figure 2.6: Left: DentSim simulator uses a passive mechanical simulator base coupled with optical tracking and on-screen visual feedback. Right: Simodont Dental Trainer. A wide-spread commercial dental simulator. DentSim video frame from image-navigation.com/mental-image accessed 2016-02-12, used with permission. Images of Simodont courtesy of Moog Industrial Group, used with permission.

a custom-designed admittance-controlled [Van der Linde et al., 2002] haptic device and a projector-based 3D display system with projected image co-located with the handpiece giving the user wearing polarising glasses the sensation of touching the virtual teeth where they are seen [Bakker et al., 2010]. The system also offers a touchscreen to the side for control, a Spacemouse (3Dconnexion, Munich, Germany), foot pedals, hand rest and an adjustable stand. It uses a 3-DoF haptic-rendering algorithm. Software with education content, or so-called Courseware, is developed by the Academic Centre for Dentistry in Amsterdam (ACTA) [Bakr et al., 2013]. A study of a faculty's impression of Simodont reports that it is not realistic but has educational potential. They also report on some technical limitations regarding, e.g., the hand rest position, and that they do not think the simulator can fully replace teachers [Bakr et al., 2013]. A study of the ability to transfer skills gained with the Simodont to a standardised drilling task has been performed by Bakker et al. [Bakker et al., 2010], showing that it works as well as training with mechanical simulators but requires less direct supervision time.

In 2007, the Voxel-Man simulator originally developed for temporal bone surgery was adapted for apicectomy, an oral surgery procedure [Von Sternberg et al., 2007]. This sim-

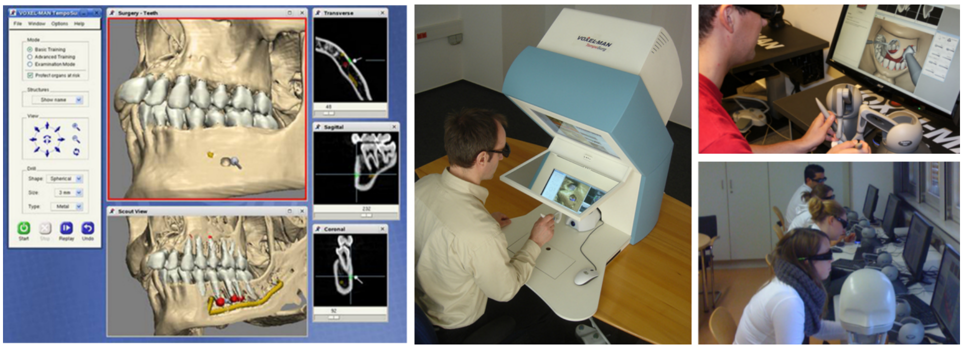

Figure 2.7: Voxel-Man simulators. Left: visualising the nerves through transparent bone [Pohlenz et al., 2010]. Center: Styled enclosure. Right: Desktop PC set-up. Images courtesy of Voxel-Man Group, University Medical Center Hamburg-Eppendor, and Elsevier, Oxford, UK, used with permission.

ulator had a co-located stereographics set-up, an Omni haptic device and a GUI combined with the actual 3D display. Three different training levels were provided, where the basic level showed transparent bone and highlighted artificial pathologies and nerves, as well as windows showing CT image planes (figure 2.7). The highlights and transparent bone are disabled in the second training level, and the CT images are disabled in the final examination level. In 2010, the simulator had matured to also include a styled enclosure (figure 2.7 center) [Pohlenz et al., 2010]. They decided to simulate an apicectomy (a tooth root treatment procedure) because it is a common procedure and was determined to be suitable for simulation. Although surgical extraction of teeth was identified as the most commonly performed surgery procedure in dentistry, it was deemed not suitable for simulation because of complex movements. This simulator also has a capability of magnifying the operating field up to 20 times [Von Sternberg et al., 2007].

As of writing, in 2015, Voxel-Man mainly promote[5] a version of the dental simulator that uses a non-colocated 3D display in a more traditional PC set-up (figure 2.7 right). I argue that this can be understandable in the light of our own work in that making a custom enclosure is a large endeavour itself, especially since it absolutely requires making it aesthetically pleasing and it is heavy to transport.

Another aspect of the Voxel-Man simulator worth mentioning is how it extends the research group's pioneering work in interactive medical visualisation [Höhne et al., 1995]. In fact, "VOXEL-MAN" is also the name used for their interactive 3D anatomy atlas (figure 2.8). This application's main purpose is not to photo-realistically represent the human body, but to link spatial regions of it with a knowledge base, e.g. the name and function of organs. A book-based anatomical atlas is static by nature and therefore only minimal annotation and view perspectives can be produced on a page. Comparatively, a computerised knowledge database can store multiple attributes on a per-voxel basis that can be recalled

---

[5] http://www.voxel-man.com/simulator/dental/ accessed 2015-10-05

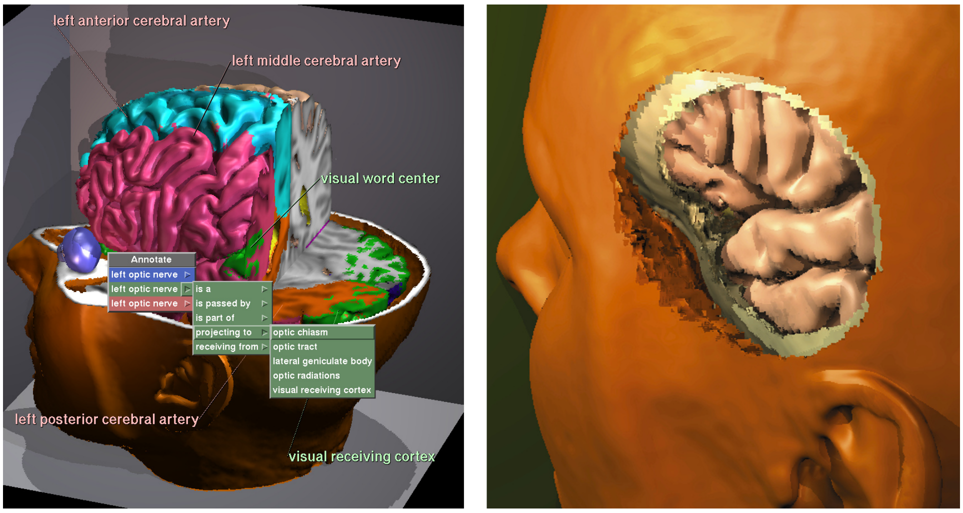

Figure 2.8: VOXEL-MAN interactive medical atlas from 1995. Interactive tools enable arbitrary cutting and selection of informative views [Höhne et al., 1995]. Images courtesy of Voxel-Man Group, University Medical Center Hamburg-Eppendor, used with permission.

interactively. The full dataset can therefore, if corresponding tools are implemented, be explored interactively and for various purposes. Höhne et al. [Höhne et al., 1995] show how cut-planes, enabling and disabling of layers, and also free-form cutting can be used to explore and, to a certain extent, simulate some surgical interventions [Pflesser et al., 2002]. In this respect it can be noted that unlike simulators designed to accurately mimic the actions of physical tools such as dental drills, these anatomical atlases provide several tools, e.g. cut-planes that have no physical correspondence. The apicectomy simulator mentioned above (figure 2.7) illustrates, I argue, in a powerful way the combination of mimicking bone drilling through the use of a haptic device on the one hand and creative visualisation techniques of, e.g., transparent bone on the other.

VirTeaSy is a simulator for dental implants [Cormier et al., 2011]. It has a planning phase, when cases are presented with x-ray slices, and a surgical phase where the surgery is performed. The latter consist of a set-up with a tracked, head-mounted display and a Virtuoso 6D haptic device. In addition there is a clinical case database [Cormier et al., 2011]. VirTeaSy optionally shows a cross where students should drill, a recommended angle, depth and warning colours for overheated drilling. The students' drilling is recorded and the result can be viewed in the scan mode. The teacher has an interface where he or she can view and interact (zoom etc.) on a separate screen. Students can switch between the two phases, and view the surgical outcome in the planning phase retrospectively. Cormier et. all stress the importance of multidisciplinary competencies in the design of the simulator. The design is grounded in video recordings of surgery and self-confrontation interviews. The implementation has been selective, based on pedagogical value: "Only the important information (needed for learning) has been implemented in the simulator, and all elements not specific to implantology or not essential have been set aside" [Cormier et al., 2011]. Joseph et al. [Joseph et al., 2014] have studied the learning contribution of the VirTeaSy simulator for third-year dental students. Three groups were compared: 20 students without simulator training, 20 students who received simulator training and 20 faculty dentists. All were given the task of drilling in a physical model that was judged in terms of position and angle deviation from a perfect drilling. The results show that the students became better over time with increasing the number of training sessions with the simulator. Most notable was the difference in variance in position and angle deviation between the simulator group and the non-simulator group [Joseph et al., 2014].

Wang et al. [Wang et al., 2012] have designed and developed a simulator for three non-drilling dental procedures: pocket probing, calculus detection and calculus removal. The first procedure is performed with a thin, cylindrical instrument with millimetre marks that is inserted into the pocket between the tooth and the gingiva (gums). A Phantom Desktop, a standard monitor (later versions use a co-located display), a 3-DoF haptic-rendering algorithm and a visually rendered full head with an open mouth where the mandible teeth are clearly shown from above. One of the reasons for developing the simulator was to "understand the necessary requirements". They adopted the concept of *construct validity*, in which the simulator is evaluated in terms of its ability to reflect actual skill levels, i.e. if the simulator can distinguish between novices and experts, then it is accurately simulating something clinically relevant. In addition to this evaluation, they used a questionnaire for feedback on the level of realism of the simulator. Based on these studies they conclude that the simulator has two design limitations. First, the cheek occludes correct reading of the probing tool, which is met by suggesting that "tongue and cheek should be deformable bodies and a mirror should be used to deform" and that a 6-DoF haptic algorithm and haptic device with torque feedback should be used to avoid penetration of the probing instrument. Since these technical advances are very complex, I argue that there may be other design solutions that would solve the occlusion issue, e.g. making the cheek translucent or deforming the cheek without haptic feedback of that action.

Tse et al. [Tse et al., 2010] have developed hapTEL (figure 2.9), a simulator for pre-clinical training, based on user requirements found with an earlier prototype [San Diego et al., 2008]. This simulator had to be cost-effective in order to allow for large-scale educational evaluation, which required a limited series of units to be produced. Therefore compromises between quality and the number of systems had to be made from the start. An evaluation of several devices informed the decision to produce 12 simulators based on a modified Falcon (total cost: GBP 4,000) and two simulators based on Omega. The low-cost Falcon, lacking orientation sensing, was modified to hold a real dental handpiece through magnetic coupling and was tracked by a linked arm at the rear of the handpiece, where the cord of a real handpiece goes. Evaluations of real use by 48 students of the twelve units of the second prototype show, among other things, that the internal friction and mass of the haptic device is considered too high - some students even used two hands to operate it. Conclusions for future work include more sophisticated haptic-rendering algorithms and rubber cheeks to limit the range of motions to be closer to that of reality. They also discuss the importance of different colours and the learning potential of, e.g.,

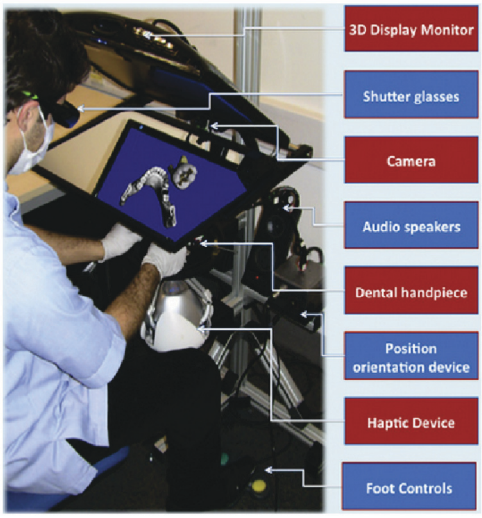

Figure 2.9: HapTEL simulator. Used with permission of the author.

transparency for showing structures in a way unavailable to traditional settings.

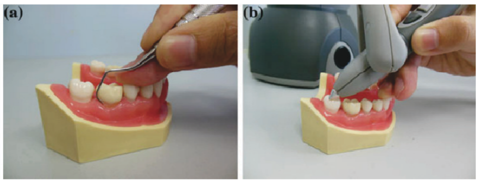

Figure 2.10: Manipulandum shape affects usage in relation to other objects. From [Wang et al., 2014]. Used with permission of the author.

The selection of which haptic device to use in a surgery simulator (and other applications) is an important design decision. A first look at a set of commonly available devices (figure 2.3) shows different form factors that constrain how they can be integrated into a simulator set-up. The shape of the manipulandum affects how it can be used in relation to physical objects (figure 2.10).

The Kobra simulator (figure 2.5) developed by Forsslund et al. [Paper A] over several years and iterations has been designed to support teaching of oral surgery procedures. The simulator features a co-located stereoscopic display large enough for two simultaneous users (e.g. the learner and the teacher), a silicon mannequin head, a Phantom Desktop haptic device and two foot pedals. The complete simulator is housed in a professionally designed enclosure. The simulation computer is started by a single button at the front of the enclosure and boots directly into the application, which contributes to the intended feeling of a coherent "product" rather than a desktop computer with add-ons. The user can select among various patient cases, i.e. particular procedure exercises modelled from particular anonymised patients who have undergone this procedure in real life, using the graphical user interface on a touch screen pad computer placed at the side on the simulator.

Each patient case consists of three key components: the computer graphics objects that make up the patient's anatomy, the dental instruments' form and function, and a state machine through which the procedure is progressed. The CG objects are divided into interactive and non-interactive objects. The non-interactive objects of, e.g., the face, soft tissue and protective cloth provide a context for the surgical scene and procedure. These objects, represented by static textured polygonal meshes, are non-interactive in the sense that they are not used for collision detection and thus no haptic feedback is provided. The interactive objects are the jaw and teeth models, which originate from the original model patient and can be carved and rendered both visually and haptically. The dental instruments may be a dental drill, an elevator (a screwdriver-like instrument) or an excavator (similar to an elevator but used for removal of infected tissue). The simulator currently only supports 3-DoF sphere-based haptic rendering, so all haptic interaction is confined to the tip of the respective instrument. Nevertheless, the procedures can be carried out with variable realism, through *gestalting* rather than *simulating* the procedures, something that will be discussed in later chapters. The final key component of the patient case, the state machine, is used to this end. Through keeping track of the amount of material removed in certain regions of interest decided by the case designer, it is possible to trigger propagation of a basic state machine and masking out of pre-defined regions of the object. This way procedures such as dividing a tooth and taking it out in parts can be simulated; when the tooth has been cut through deeply enough and the elevator is applied with a certain force, then the tooth part is masked out and removed.

The first prototype of the Kobra simulator was developed following a user-centred design process, with initial technical investigations in parallel with field studies, interviews and observations in operating rooms [Forsslund, 2008]. This led to design decisions such as the use of context mesh and interactive objects, and the focus of the simulation on one step in a particular procedure: surgical extraction of wisdom teeth. The prototype was evaluated using the co-operative evaluation method [Monk et al., 1993] with four experienced surgeons, which led to improvements in hand support, orientation of the patient with respect to the operator, and more. Development and co-operative evaluation were thereafter intermixed in iterations, focusing on different aspects; e.g. the visual feedback [Flodin, 2009]. In total, five co-operative evaluation sessions with senior dentists and one with dental students have been held [Paper A]. In addition, one particular full-scale study has been conducted where two copies of the simulator were used in two teaching sessions totalling 2x30 students and three teachers [Lund et al., 2011]. The simulators were used by the students, often in pairs, but under the guidance of the teachers. Observations have indicated that when used in this way it seems more useful than what had been observed in single use. The teachers had a richer dialogue with the students than simply how to carry out the technical task in the computer. This will be discussed further in later chapters as well. Besides the observations, there was a quantitative questionnaire study recording the students' impressions of the learning session. The results show that most of the students rated the simulator's realism as a 4 on a 6-grade scale, that the majority wanted more training with it, that they gained new knowledge from practicing in the simulator and that 73% of the students strongly agreed that practice in the simulator should be included in the course [Lund et al., 2011]. Another study investigated the role of the instructor by way of a between-group experiment with groups of eight students who were subjected to practice in the simulator with either no feedback, feedback from another student, feedback from a technician or feedback from a surgeon. The result shows that the best performance was given with a surgeon as instructor [Rosen et al., 2014]. In addition to these studies, there has been much informal feedback from dental school faculty at private demonstrations and public exhibitions at four trade fairs [Paper A]. Finally, four copies of the simulator have been installed since 2013 and are in use at the Riga Stradiņš University.

Chapter 3

Research Process

The prevailing research approach in this work can be characterised as *research through design* [Zimmerman et al., 2007]. The findings presented in the paper *Designing the Kobra Oral Surgery Simulator Using a Practice-Based Understanding of Educational Contexts* follow the central idea of research through design, in that results in *design knowledge* relate to what a designer of such a simulator would benefit from considering. Apart from this paper, however, the focus of the thesis is on what needs to be done in a more technical sense to *prepare* haptic technology for interaction design. This includes implementing well-known algorithms and encapsulating them in such a way that designers can explore the user experiences they afford. While not necessarily novel in terms of technical sophistication, this work is grounded in the needs made apparent by the Kobra design work. In addition, to test a hypothesis regarding the effect of using fully actuated haptic devices and six degree of freedom haptic rendering algorithms, a controlled experiment has been performed.

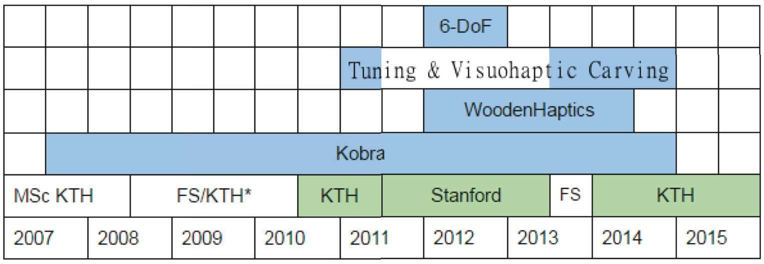

Figure 3.1: Temporal distribution of projects. Green periods are part of the PhD. FS/KTH* Part-time as research engineer at KTH and self-employed at Forsslund Systems. Projects not included in thesis are excluded.

The research presented in this thesis can be sectioned into four partly overlapping research projects; the Kobra simulator (paper A), WoodenHaptics (paper B), Tuning of visuohaptic carving properties (paper C and D), and Haptic rendering fidelity for surgery simulation (6-DoF, Paper E). The temporal distribution of these projects is shown in figure

3.1. Several iterations of the initial Kobra prototype development cycle with field studies, design, implementation and user testing have been conducted as part of a sharp product development project. The development work of the Kobra simulator was used as basis for the analyses and the requirements formulation of tools and resources presented in this thesis. In the following each project will be described and how these make up the parts of the overall research through design process.

3.1 Developing the Kobra Simulator

For this project the on-going commercial and open-source driven product development project Kobra has been used as a vehicle for the research. There have been many user-centered activities within the product development project, and many of them involve empirical data gathering with prototypes and users, with the explicit goal of advancing the design. In addition there have been research activities where the Kobra simulator has been used in context to serve other researcher's empirical needs, such as inquiry into the efficacy of simulator training and how the training can be integrated in dental education curriculum. Furthermore, the research presented in this thesis adds a retrospective high-level analysis to these which includes the tools and technologies implemented as objects of study as well. In order to have the right expectations it is important to classify the activities according to the purpose they originally served. These are:

1. Activities pertinent to a user-centred design process, e.g. field studies, design and realization of prototypes and cooperative evaluation with domain experts. Materials gathered here are analysed to the extent deemed necessary to inform the next design iteration.

2. Studies of the use of the simulator in order to answer axillary research questions such as the efficacy of simulator training, the impact of different instructors or students acceptance rate. These studies are only indirectly motivated by improving the current design, and primarily view the simulator prototype as a research object. Observations by the developers during these studies have however also contributed to the iterative design.

3. The retrospective analysis of the data gathered above plus annotations, photos, videos, anecdotal feedback from exhibitions, public documents created by the developers and more have been used for the holistic reflection of the design process, with the purpose of creating design research knowledge regarding important aspects of surgery simulator design.

4. The retrospective analysis of the design resources like for example the software library *forssim* and tools created during the product development such as a tool for tuning the haptic rendering properties. These design resources have also been applied in other projects such as interactive art.

The product development process of the Kobra simulator is covered in paper A. Following a pre-study with both technical investigations and field studies including interviews and surgery observations, an early prototype was developed and it was subsequently subject to a cooperative evaluation [Monk et al., 1993]. The results of this evaluation informed the development process of the Kobra simulator that followed. A total of six iterations and five cooperative evaluations were performed spanning the years 2007 through 2014. The development process has covered all aspects of the Kobra simulator, including the styling of the casing, graphical user interface and virtual patient cases. The development work was further informed by observations during studies of category two. For example, in one study 60 students were using the simulator under guidance of a senior surgeon [Lund et al., 2011]. It was observed how the dialogue between student and teacher included aspects of the procedure and patient treatment not directly present in the simulation. The simulator prototypes was also exhibited at dental education conferences, shown to faculty at various dental schools and experimentally included in a continuous education course. This resulted in significant amount of informal feedback from the target user group. In addition to these, the product development went through a particular phase, when four simulators were produced and distributed to a dental school in Latvia. This phase included creation of new patient cases (exercises) with a professional 3D artist. All this was reflected upon, and presented in papers A and D.

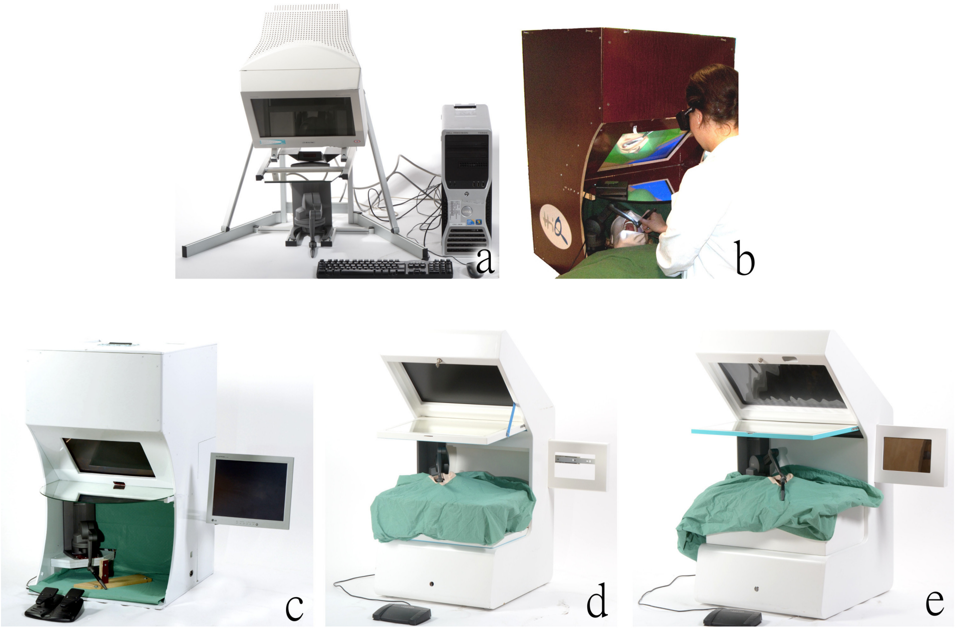

Figure 3.2: The Kobra simulator in different development stages. a) using available co-location setup (2007), b) first enclosure (2008), c) painted version (2009), d-e) present design (2011-2014).

The process that was used for analysing the design was as follows. First a diary, or history log book, was created where development milestones and related studies, events or focused design work were ordered chronologically. This document included excerpts from video recordings, product descriptions, screen-shots and references to academic and other publications made during the period. This material helped recalling why certain decisions were made, and when. For example an event was recalled, when pictures of an early prototype (figure 3.2b) had been presented to faculty of an dental school that was interested in the product. Even though the primary focus at that point in time was on getting feedback for further developing the software, the home-made looking case was too distracting to direct the discussion on the software requirements. Therefore it was decided to already at this stage initiate styling of the case (figure 3.2d-e). These reflections was then written down and clustered into themes, together with relevant notes from literature, e.g. [Johnson, 2004]. The resulting themes, or *design aspects* are presented in paper A. Their purpose is to inform simulator designers of important aspects that the authors think should be considered, but not necessarily be solved in the same way as the Kobra.

During the product development, the first priority was to implement the fundamental algorithms and to achive an acceptable level of simulation fidelity i.e. the haptic rendering was stable and the visual rendering capabilities included shading and colours [Flodin, 2009, Wijewickrema et al., 2013]. In the next stage, it became evident that either a lot of time should be invested in advancing the technology i.e. 6-DoF haptic rendering, or time should be spent on forming what was developed so far into a more useful and meaningful learning experiences. The observations and questionnaire results from studies with surgeon teachers and students had shown that even a rudimentary prototype was appreciated, when it was used in a meaningful context of teacher-driven learning [Lund et al., 2011]. This lead to questioning the often not articulated assumption that improving *realism* as the primary goal of the development and instead opening up for alternatives using the same technological base [Forsslund, 2011]. In order to explore and understand what those alternatives could be, how they look and feel, and what opportunities they afford, two new research projects were initiated; one that investigates the hardware properties, i.e. WoodenHaptics, and one that focuses on the design resource (or digital material) named visuohaptic carving, discussed below. The haptic device used in Kobra has some drawbacks that motivated this investigation. First, it lacked screw holes or similar for fixating the device to the simulator, which, in addition to its “stand-alone” looks, makes it suboptimal as a component of an integrated system, which it is not designed for. Second, there is no obvious way of modifying the hardware, for example replacing the manipulandum with an authentic dental drill. Third, it is relatively expensive, costing more than the other components of the system combined. What more, it is not obvious what actually causes the quality difference between this device (Phantom Desktop) and its low-cost alternative (Phantom Omni) of same manufacturer with similar function. These were the reasons for investigating the hardware in depth from a design perspective, but also, as will be argued in the sections on visuohaptic carving, from a software perspective, since when it comes to spatial haptics are hardware and software tightly coupled, especially in terms of stiffness and virtual object size.

3.2 Developing Spatial Haptic Hardware: WoodenHaptics

The WoodenHaptics project served two major purposes. Initially it was a learning process in which I, who had no prior electromechanical design experience, set out to build my own spatial haptic device in order to understand all aspects of the technology. For instance, commercial devices can vary significantly in fidelity; the Phantom Omni and Phantom Desktop are two similar devices that have different stiffness capabilities, but as a software designer it is difficult to understand why. Building a device helped me understand how the devices are constructed and made it possible to experiment with alternative designs. Engineering a spatial haptics device is however a large endeavour only feasible in highly specialised robotics labs. Fortunately I had the opportunity to conduct this research in such a lab; the Salisbury BioRobotics lab at Stanford University. This leads to the second major purpose of the thesis; to investigate if and how this kind of tacit knowledge can be packaged so that designers can use it without access to the competencies of a sophisticated lab. One way to do this is through the concept of toolkits, something that has been theorized by von Hippel in *User toolkits for innovation* [Von Hippel, 2001] and *Democratization of innovation* [Von Hippel, 2003]. He argues that needs of a domain is sticky, i.e. difficult to transfer to a manufacturer of a particular kind of solutions. Applied to haptic devices, it can be the case that manufacturers could make a custom made device for a particular domain, but do not know enough about it to motivate costs and risks. The users, or in this case, designers, would know about the domain but not enough about trade-offs of different solutions. That knowledge is internal to the manufacturer and equally sticky. Toolkits made by a manufacturer have been shown to be able to bridge this gap through putting easy to use design kits in the hands of domain experts, and, in von Hippel's cases, still make use of the manufacturers machines for final production, with economic benefit to both [Von Hippel, 2001].

The Salisbury lab has an extensive tradition of building high-quality functioning robot prototypes as part of their research. Much of this knowledge was shared in the course *CS235 Applied Robot Design for Non-Robot Designers* introduced for the first time in 2012. The motivation for providing the class is perhaps best captured by the professor in the following quote:

> What motivated this class? For a long time I have been a mechanical engineer and worked with computer scientists. One thing that always concerns me, is that sometimes, not always, CS folks view a robot as a black box; "We can wrap some code around it and it will work". That is actually not true at all. If that robot got back-drive in it, or friction, or natural frequencies that are wrong there are many mechanical (issues) that makes it very difficult to make the robot do what you want to do. (...) This class is beginning from ground zero, how to build machinery, robots in this case, that works. But more importantly, lets say you are in a position to supply (i.e. acquire) a robot, and you look at the spec sheet and it says: inertia x, back-lash y, and acceleration z. What does that mean? Is that good or bad? The next level is what mechanical

> properties should they have? Should they be compliant? To the lowest level; how do you build such a robot?[1]

The course consisted of lectures and weekly, directed, individual hands-on construction projects, culminating in one larger project of choice. The content covered all aspects of actual making a functioning device, including sourcing parts, Computer Assisted Design, laser-cutting plywood, 3D-printing, cutting and grinding steel shafts and assembly in a professional way without use of glue or other adhesives. The necessary tools and production equipment was available under guidance. This direct access to fabrication resources has been labelled *personal fabrication* [Gershenfeld, 2008] and is now an active research area in itself. Personal fabrication was important since it allowed for short iterations. WoodenHaptics was born out of the final project, where I teamed up with Michael Yip, and used additional resources of the lab; motors that was detached from an abandoned prototype, along with lab-bench motor control equipment. The development of the device was bottom-up, i.e. step by step I learned to read output from the encoders and set voltage signals on the computer to communicate with the power amplifiers. The structure was designed around the specific motors at hand, and elements of the design were borrowed from the course assignments, for instance the cable tensioning mechanism and the first link design. The control of the device was implemented in C++ using only the Chai3D and the Digital Acquisition Card's libraries. The kinematics was derived with assistance of Adam Leeper, a course assistant of professor Mitiguy's ME331a Advanced Dynamics course, and the software Motion Genesis. The device worked well according to our subjective experience. The question remained if this device, built using scrap components in the Salibury lab, could be replicated elsewhere? Therefore the project shifted from one-off prototype building to kit construction. Replacement components were purchased which resulted in an up-to-date Bill of Materials (BoM). Theoretically, it should be sufficient to provide the BoM and blueprints of the wooden parts. To verify this, the device was later replicated at the Media Technology and Interaction Design department at KTH, which at the time lacked the fabrication facilities of the Salisbury lab. The reconstruction of the device required substituting several components that was difficult to acquire from the original sources in USA or proved prohibitively expensive. It was also necessary to equip the lab with a laser-cutter and necessary hand-tools. The positive side-effect of this was that a complete list of vital items for a personal fabrication lab capable of producing the haptic device could be formulated. Additional practicalities proved necessary to overcome, such as the selection of wood species; pine was too hard to cut, so the material used in the end was birch. Step by step the usability of the assembly was improved, through embedding the electrical couplings in a custom-made electronics box, for instance. The result was a set of assembly-ready parts, and easy-to-connect electronics.

To verify that WoodenHaptics was not only easy to assemble at a new site, but possible to assemble for a robotics construction novice, two interaction design researcher with no prior experience of this kind of construction were assigned with assembling the device,

---

[1] Video capture of first lecture, available on http://www.youtube.com/watch?v=Pk1ou6C4jWg accessed 2015-12-20.

under supervision. The process was recorded on video, lasted in total 11 hours, and proved successful. The device was also subject to an experiment where 10 participants compared the user experience of the device with that of three commercial devices (Phantom Omni, Phantom Desktop and Novint Falcon). To show the ability to make variations of the device design one smaller version and one with a lathe-crafted link was made. Furthermore, measurements of the device in terms of friction, workspace dimensions, stiffness and theoretical maximum force was gathered, both for the original and the smaller variant. The details of these studies and measurements are presented in Paper B. As is further elaborated on in the paper and in chapter 4, WoodenHaptics is designed for user customization and exploration. Being able to switch motors with relative ease is one such feature that support designers' own exploration of what impact the motor qualities have on the final haptic user experience. This was something that I, as a designer, saw a need for in my own work designing applications such as the Kobra simulator.

3.3 Tuning of Visuohaptic Carving Properties

The starting point for this project was the need and desire to explore the design space of the haptic technology that had been implemented in the *forssim* software library (see chapter 4.1 Visuohaptic Carving). The project arose from collaboration between KTH and University of Melbourne. Both had been independently developing simulators, for oral surgery and temporal bone (ear) surgery, respectively. The two groups used and contributed to forssim, and had noticed that changing haptic rendering parameters affected the user experience and that there is a non-trivial relationship between the haptic device employed, the size and stiffness of the CG objects touched and the size of the carving sphere. It was observed that experimenting with these properties was possible, but it required that the user changed values in a document and re-starting the simulation between iterations. This took too long time for the user who could not remember the touch sensations and this in turn made tuning of the haptic feedback very hard. This motivated the development of a tool for more direct adjustments of values, for both exploration of opportunities and for fine-tuning. With inspiration from audio tuning applications and earlier unrelated explorations, the tangible midi-controller BCF2000 was selected as a potentially suitable interface for directly tuning the parameters. The midi-controller BCF2000 has several channels that independently can be controlled by a slider or a knob, and transmits their values to the computer over USB. Through a Python script the values are routed to the visual and haptic rendering parameters. This way a working haptic "studio" could be set up without much effort, and no excessive software development. Additional benefits were that the BCF2000 could store and recall a number of pre-sets, i.e. the user could easily create different combinations and save them in the device. This opened up for new forms of experimentation, or sketching the user experience of a particular visuohaptic carving application. It became for example possible to scale the CG object up, in this case a jaw model, and to tune the transparency of the bone slightly so that the roots of the teeth could be shown, and to adjust the carving rate of the root and bone so that one could feel the harder root when carefully removing bone. To get this right (i.e. feel "nice") with the low-cost Omni haptic device the enlargement of the jaw model was necessary, which of course broke the realism in terms of 1-to-1 mapping between real teeth and the virtual in terms of size. However, this could then be saved as one of several sketches and later shown to surgeons who themselves could tune it as well. In that way, these sketches could influence the decision on what extent of realism that is necessary and desired, based on what qualities the technology could offer. The reflection on the utility of this tool is presented as one of the contributions in Paper C.

The tool was later used to tune the parameters of the haptic and visual rendering in both simulators, as well as in an art project called *immaterial materials*. Since this art project didn't have any constraints on realism it was more open to maximizing the user experience, and that enabled us to be more creative with the settings. The art project also made use of the forssim library, although other solutions had been possible, and it was shaped around using exactly the technology that had been developed for the Kobra surgery simulator. This planted the seed for reflecting on how this piece of technology, the particular haptic rendering method implemented in forssim, acted as a malleable design resource. Several things contributed to this concept. First, as the artist approached us, who were software engineers, about doing a haptics-based project, we suggested doing something using the code-base from the Kobra simulator. What computer graphics objects and how they were presented was up to the artists judgement but we showed how to do it, often working together side-by-side on the same computer, or the two developers together with the artist on the side. The objects that the artist created or acquired were represented by meshes, and command-line tools were used for voxelising these.

Two years later we got a request to develop new patient cases for the Kobra simulator. This actualized the need for a systematic process for creating them. The previous case had been created through a tedious segmentation process where a cropped computed tomography volume hade been "painted" slice by slice, and the model was fitted with a mesh face that lacked several anatomical features. This was good enough for a prototype but not for production. The project team was expanded with a professional 3D artist. This time the 3D artist was fully involved in the patient case-making process instead of only making and delivering some elements, e.g. face meshes. The intent for using this way of working was to test to what degree the 3D artist could work independently, and how to utilise his professional skills and tools in the process of creating patient cases. Furthermore, the intent was to investigate what more might be required, in terms of novel tools, processes or engineering support to enable a 3D artist to work independently with haptic enabled patient case creation. It became evident that the slice painting segmentation process and the editing of inadequately commented text-files were frustrating and insufficient. We therefore had to work out a more suitable process together. The resulting workflow model is presented as the third piece of the contributions from this thesis, and it is presented in more detail in chapter 4.

3.4 Evaluating 6-DoF versus 3-DoF Haptic Rendering

The larger project that this study belongs to is that of surgery rehearsal and planning, at the Salisbury BioRobotics lab at Stanford University. Much of the work in the lab was centred around developing novel haptic rendering methods suitable for patient-specific surgical rehearsal practises [Chan, 2011]. The domain was primarily microsurgery of the middle ear canal, where navigating narrow spaces with a surgical drill or probe is part of the task. Not only was a novel simulation application proposed, but also the aim of the greater project was to advance the state of art in computer science [Chan, 2014]. Since fully actuated 6-DoF haptic interface and corresponding rendering carry significant financial and computational costs it was motivated to investigate which level of fidelity or realism actually would be required to achieve surgical simulation objectives. Besides comparing task performance with two different haptic rendering algorithms, this project aimed to investigate the effect of using the advanced 6-DoF algorithm together with an under actuated device, which would significantly reduce hardware costs of a system.

For this study a within-group experiment with twelve subjects was designed. The materials used was a software application developed by the lab, where two different haptic rendering algorithms could be run, either one that only considers collisions with a rotational invariant sphere placed at the tip of the avatar (3-DoF), or one that considers the full rigid body of the avatar (6-DoF). The physical setup consisted of a Phantom Premium 6-DoF haptic device, a 3D TV, and a 3D rendering capable computer. The torque feedback of the haptic device could be enabled or disabled as well. Two different virtual environment scenes were developed, where the user should navigate with a probe and touch small spheres without excessive contact with surrounding material. The system kept a score of the time and number of errors per scene and what stimuli was currently tested, i.e. one of the three independent variables; sphere rendering, full rigid body rendering, and rigid body rendering with disabled torque output. The complete details of the experiment are covered in paper E. The results showed that there were no significant differences between displaying torque feedback or not, but that 6-DoF haptic rendering significantly improved the users performance. This can be of use in a future version of the Kobra simulator, in that it strongly suggests that the user experience can be improved without a hardware cost premium, if a 6-DoF haptic rendering algorithm is implemented in the software.

Chapter 4

Research Contributions

This chapter will present the research results which, taken together, explain how spatial haptic technologies can be prepared for interaction design work and applied in a real-world design case. The way the technologies are prepared is through encapsulation of technical nuances into formable design resources, tools for forming the resources into applications and a suggested way of working with these tools and resources. The design resources are *WoodenHaptics* [Paper B] and *visuohaptic carving* [Paper D]. WoodenHaptics is a novel haptic device that can be used as the basis for user-side design explorations. It can be used as-is or adapted for a particular use case by an application designer, and allows for exploring the user experience of using different motors, materials and dimensions without engaging in extensive engineering problem-solving. Visuohaptic carving is an intangible design resource that at its core consists of a software library called *forssim* [Forsslund et al., 2009] that implements algorithms for haptic rendering and carving of solid, multi-layer computer graphics objects. Providing only the library, even when exposed for use in a high-level declarative programming environment, has proved to be insufficient for practical design work, i.e. efficient prototyping and exploration of user experience. Therefore an interactive design tool has been developed (in two variants) where interaction designers can explore the haptic-rendering properties that affect the user experience in real time, and tune them for the particular object and haptic device employed [Paper C]. To put the tool into the context of practical use, a specific workflow - a kind of design process - was developed. This process involved prototype file conversion tools in the form of semi-automatic scripts that allowed a professional 3D artist to leverage his skills in using the modelling tools he was accustomed to throughout the process. In addition to forming the objects and tuning their visual and haptic properties, an interactive scene can be defined in declarative language, complete with some rudimentary event handling that enables the design of interactive scenarios.

The second part of the chapter will present how visuohaptic carving has been applied in the design of the oral surgery simulator Kobra [Paper A]. This research-through-design study investigated what constitutes a useful surgery simulator beyond mimicking the interaction of surgical instruments and human tissue as realistically as possible. The project has resulted in the articulation of several important aspects and the argumentation of why these are worth the closer attention of simulator designers. Most central to the discussion of the thesis is how the relatively modest simulation technologies implemented in *forssim*, i.e. visuohaptic carving, were used to *gestalt* specific surgical procedures that had been performed on real patients by surgeons in a teaching hospital. Several patient case scenarios have been implemented and delivered as per the request of the same teaching hospital and remain in use as of this writing. Among the conclusions it has been noted, based on this inductive research, that as a grounded theory, it *may* be more fruitful to focus on supporting the cognitive aspects of surgical proficiency, i.e. reasoning on how to surgically treat a particular patient, rather than fine motor skills training. In summary, the contributions put forward in this thesis are:

1. Design resources for interaction designers who would like to engage in purposeful haptic interaction design. These are WoodenHaptics and Visuohaptic Carving.
2. WoodenHaptics, a starting kit for crafting haptic devices, that can be used as a design resource in developing haptic enabled interactive systems. It is shown how this device can be produced with personal fabrication methods and how it can be altered to explore the design space of various work-spaces, motors or industrial design.
3. Visuohaptic Carving, a conceptual design resource that is complete with a ready-to-use software library, tools for exploring and tuning the user experience, and a work-flow that leverages professional tools and skills of 3D artists.
4. The appropriation of Visuohaptic Carving along with interaction design to *gestalt* rather than true-to-life simulate authentic patient cases, and how this can support teaching of surgical procedures.
5. Evidence for that under-actuated haptic devices paired with full 6-DoF haptic rendering can be a viable alternative to premium priced 6-DoF devices, in some situations.

4.1 Tools and Resources for Spatial Haptic Interaction Design

The design resources presented below are WoodenHaptics and visuohaptic carving respectively.

WoodenHaptics

The WoodenHaptics starting kit (figure 4.1) has been developed as a bridge between highly specialised engineering and the need for hands-on interaction design. It does so by providing a *starting kit* consisting of all the software and hardware components needed to assemble a high-quality 3-DoF spatial haptic device and use it like other devices through a common high-level application programming interface and example software applications. Once the device is assembled, which has been timed to take 11 hours for a novice robotics designer under guidance, the designer can begin modifying the device to suit a particular application and to learn its material qualities [Paper B].

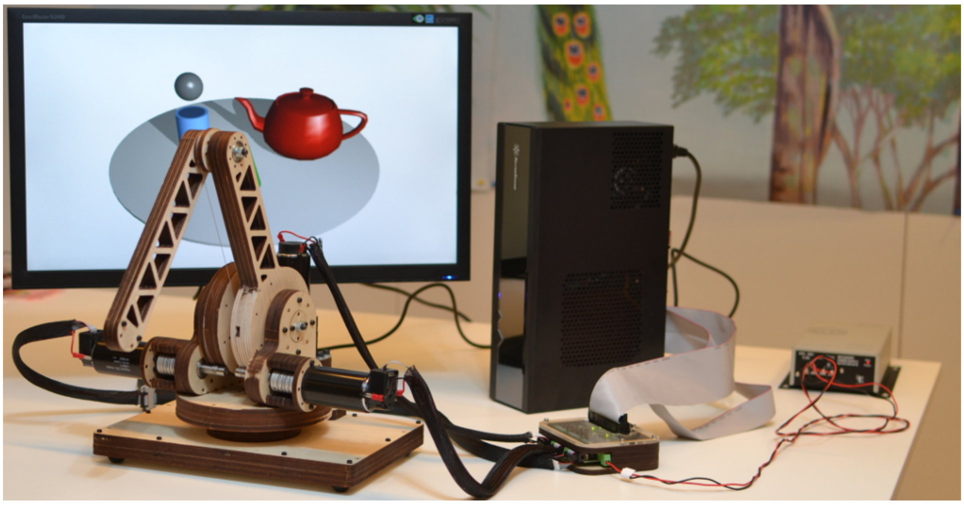

Figure 4.1: Workbench where WoodenHaptics starting kit is being used.

Engineering a custom-made spatial haptics device has been a large endeavour which is only feasible in highly specialised robotics labs that have the electromechanical and computational know-how as well as the fabrication resources. Interaction designers have therefore previously been restricted to writing software for use with a few pre-made devices available on the market.

There are several design decisions behind the WoodenHaptics kit that make it suitable for design explorations. The primary ones are: 1) the encapsulation of electronics, 2) a configurable software module and 3) the default structural design (the device itself) that embeds best practice and tacit knowledge regarding material and component selection and construction. Furthermore, it is designed to be functionally transparent in that it is easy to see how it works mechanically, e.g. by following the wire-rope that transmits the mechanical power from the motors to the respective link motion.

The encapsulation of several technical nuances into easy-to-use components reduces the designer's problem-solving activities. The externally powered electronics box (figure 4.1) interfaces between the computer on one side and the motors and encoders (the motor shaft angle sensor) on the other side. Minor details, such as the use of standard connectors that cannot be plugged in the wrong direction and a physical, labelled case, yield quick and fail-safe connection and disconnection. This fact should not be underestimated, since alternative lab-bench style connectors result in an immobile set-up and individual connections broken by accidents may require hours of troubleshooting.

A significant part of any robotics project is to formulate equations of motion and implement its control in software. The WoodenHaptics software module, technically delivered as a software patch to the Chai3D API, provides a solution that works with the default design out of the box, yet is modifiable on two levels. On a lower level is the documented source code available for those interested in control theory. For most designers, it is, however, unnecessary to work on that level, but it remains important to be able to change control-dependent variable values, such as link lengths. For instance, if longer links are used, a larger workspace can be achieved. This change needs to be reflected in the control software. WoodenHaptics therefore provides a simple configuration file where the designer can modify variables such as link lengths and motor characteristics. In addition, the file specifies the maximum force and stiffness supported, which can be found experimentally by the designer by, e.g., increasing stiffness until vibrations or other issues become apparent. These changes can all be done without editing code.

The use of laser-cut plywood for structural elements in WoodenHaptics makes it fast and easy to modify on different levels: in the computer-assisted design (CAD) model, on the flat sheet drawings or directly with handcraft tools. The choice of wood may seem less rigid than other materials, but since several stacks of plywood sheets are used it is actually rather stiff and robust, yet lightweight and suitable for self-threading screws and press-fitting of ball bearings. The rigidity of the construction is also reflected by the results of a user study where the users have rated the feeling of WoodenHaptics as being close to the more expensive and higher quality Phantom Desktop than the more common Phantom Omni haptic device [Paper B]. Furthermore, the designer is encouraged to experiment with other materials to learn their advantages and disadvantages.

The modules and parts can all be manufactured using personal fabrication methods. This entails that the kit itself can be distributed digitally in the form of schematics, blueprints and parts lists for reproduction by a third party. The use of a permissive open-source license also allows and encourages the modification and improvement of the components themselves. This can be useful for designers who want to push the boundaries of what the current kit affords, when so motivated.

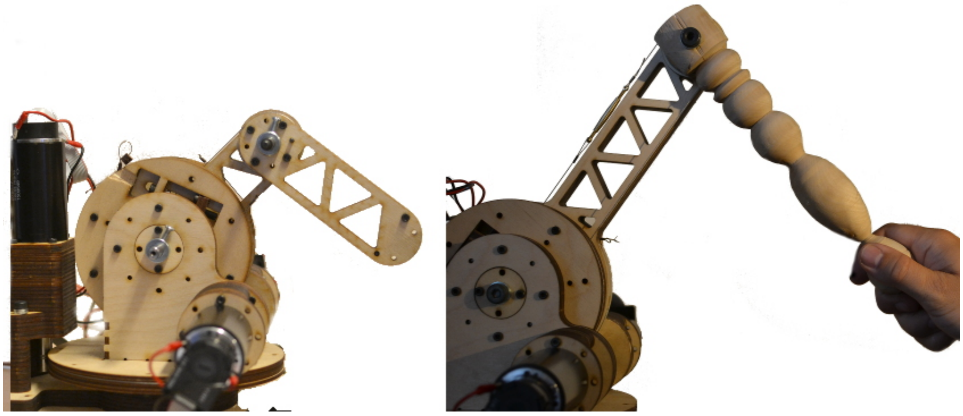

Figure 4.2: Variations of the WoodenHaptics device: shorter links which yield smaller workspace but higher maximum force, and a handcrafted link.

Taken as a whole, WoodenHaptics can be used as a workbench for hands-on exploration and designing through making [Moussette and Banks, 2011]. Figure 4.2 shows two variants: one where the workspace has been reduced, which results in larger maximum forces or reduced motor size demands, and one variant where the last link has been handcrafted using a lathe. The use of flexible couplings and external access makes it easy to replace motors, as compared with most mass-produced commercial devices where the motors are embedded deep inside the device.

Visuohaptic Carving

It is important that visuohaptic carving, as a design resource that consists of a computational module (i.e. the software library), is matched with design tools and a practice (i.e. a workflow). These elements will be described further below.

Forssim Software Library

As mentioned in the background, H3D API is one of the major open-source haptic application programming interfaces available as of the time of writing. It provides its users with a high-level declarative programming interface, i.e. a designer can declare a scene using a text editor, in a similar fashion to editing HTML code (figure 4.3). Its standard distribution allows for visualisation of polygonal and volumetric data, and primarily haptic interaction with the former. A designer who would like to provide other interaction modes, such as carving of objects, needs to implement low-level algorithms in C++ before these can be accessed on the declarative level. The *forssim* software library is a relatively small extension to H3D API which implements such key algorithms for visuohaptic carving [Forsslund et al., 2009, Chan, 2011, Wijewickrema et al., 2013].

The inclusion of the library means that a designer gets access to new tags; e.g. <VolumeModel>, which acts as a container for the voxel-based CG object the user can carve into, and <ADrillForce>, which implements a haptic-rendering algorithm. A few haptic properties can be defined for a multi-layered (segmented) solid CG object: its scale, stiffness and carving rate. The carving rate can be defined to be different for different layers [Paper D]. In addition, one can specify the location and orientation of the object in space through providing the coordinates and rotation matrices. All these properties are specified by numerical values entered directly into the text file (figure 4.3). For setting up example scenes this can be feasible, but it quickly becomes inconvenient to iterate. To see and feel the effects of changes, the designer has to start the application, try out, e.g., carving the object, then close the application, edit the text file and re-start again. What's more, the CG object files need to be prepared in a particular way to be compatible with the carving and rendering algorithms used. It is therefore not surprising that the Kobra project has seen a need for making tools that support the design of scenes and haptic material properties, which will be discussed in the following section.

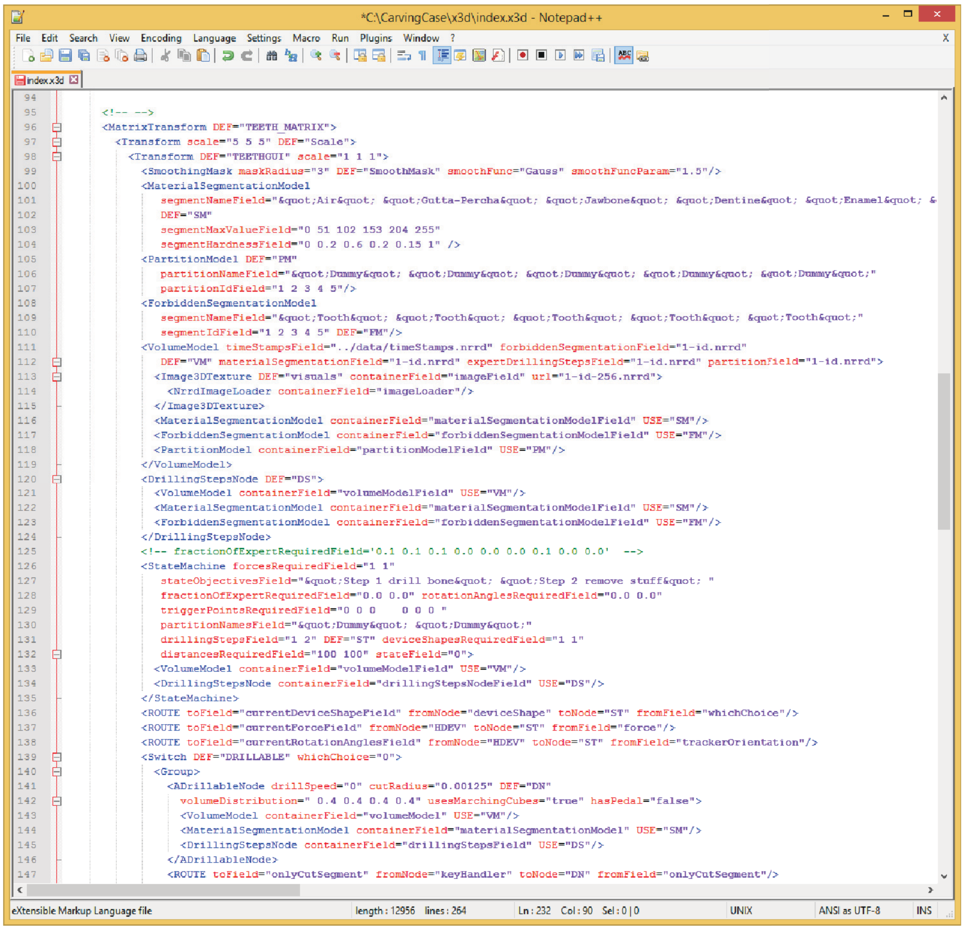

```

<!-- -->
<MatrixTransform DEF="TEETH_MATRIX">
  <Transform scale="5 5 5" DEF="Scale">
    <Transform DEF="TEETHGUI" scale="1 1 1">
      <SmoothingMask maskRadius="3" DEF="SmoothMask" smoothFunc="Gauss" smoothFuncParam="1.5"/>
      <MaterialSegmentationModel
         segmentNameField="&quot;Air&quot; &quot;Gutta-Percha&quot; &quot;Jawbone&quot; &quot;Dentine&quot; &quot;Enamel&quot; &
         DEF="SM"
         segmentMaxValueField="0 51 102 153 204 255"
         segmentHardnessField="0 0.2 0.6 0.2 0.15 1" />
      <PartitionModel DEF="PM"
         partitionNameField="&quot;Dummy&quot; &quot;Dummy&quot; &quot;Dummy&quot; &quot;Dummy&quot; &quot;Dummy&quot;"
         partitionIdField="1 2 3 4 5"/>
      <ForbiddenSegmentationModel
         segmentNameField="&quot;Tooth&quot; &quot;Tooth&quot; &quot;Tooth&quot; &quot;Tooth&quot; &quot;Tooth&quot;"
         segmentIdField="1 2 3 4 5" DEF="FM"/>
      <VolumeModel timeStampsField="../data/timeStamps.nrrd" forbiddenSegmentationField="1-id.nrrd"
         DEF="VM" materialSegmentationField="1-id.nrrd" expertDrillingStepsField="1-id.nrrd" partitionField="1-id.nrrd">
        <Image3DTexture DEF="visuals" containerField="imageField" url="1-id-256.nrrd">
          <NrrdImageLoader containerField="imageLoader"/>
        </Image3DTexture>
        <MaterialSegmentationModel containerField="materialSegmentationModelField" USE="SM"/>
        <ForbiddenSegmentationModel containerField="forbiddenSegmentationModelField" USE="FM"/>
        <PartitionModel containerField="partitionModelField" USE="PM"/>
      </VolumeModel>
      <DrillingStepsNode DEF="DS">
        <VolumeModel containerField="volumeModelField" USE="VM"/>
        <MaterialSegmentationModel containerField="materialSegmentationModelField" USE="SM"/>
        <ForbiddenSegmentationModel containerField="forbiddenSegmentationModelField" USE="FM"/>
      </DrillingStepsNode>
      <!-- fractionOfExpertRequiredField='0.1 0.1 0.1 0.0 0.0 0.0 0.1 0.0 0.0'  -->
      <StateMachine forcesRequiredField="1 1"
         stateObjectivesField="&quot;Step 1 drill bone&quot; &quot;Step 2 remove stuff&quot; "
         fractionOfExpertRequiredField="0.0 0.0" rotationAnglesRequiredField="0.0 0.0"
         triggerPointsRequiredField="0 0 0    0 0 0 "
         partitionNamesField="&quot;Dummy&quot; &quot;Dummy&quot;"
         drillingStepsField="1 2" DEF="ST" deviceShapesRequiredField="1 1"
         distancesRequiredField="100 100" stateField="0">
        <VolumeModel containerField="volumeModelField" USE="VM"/>
        <DrillingStepsNode containerField="drillingStepsNodeField" USE="DS"/>
      </StateMachine>
      <ROUTE toField="currentDeviceShapeField" fromNode="deviceShape" toNode="ST" fromField="whichChoice"/>
      <ROUTE toField="currentForceField" fromNode="HDEV" toNode="ST" fromField="force"/>
      <ROUTE toField="currentRotationAnglesField" fromNode="HDEV" toNode="ST" fromField="trackerOrientation"/>
      <Switch DEF="DRILLABLE" whichChoice="0">
        <Group>
          <ADrillableNode drillSpeed="0" cutRadius="0.00125" DEF="DN"
             volumeDistribution=" 0.4 0.4 0.4 0.4" usesMarchingCubes="true" hasPedal="false">
             <VolumeModel containerField="volumeModel" USE="VM"/>
             <MaterialSegmentationModel containerField="materialSegmentationModel" USE="SM"/>
             <DrillingStepsNode containerField="drillingStepsField" USE="DS"/>
          </ADrillableNode>
          <ROUTE toField="onlyCutSegment" fromNode="keyHandler" toNode="DN" fromField="onlyCutSegment"/>
```

Figure 4.3: Illustration of the structure of an X3D-scene as viewed in a text editor. Each tag represents an object node in the scene-graph. <VolumeModel> and <ADrillForce> are two of the extensions provided by forssim. The file is both human and machine readable but editing in a text editor alone is insufficient for practical design work.

Workflow for designing visuohaptic carving scenes

So far the thesis has presented haptic technologies and how they can be packaged in a modularised hardware platform and a software library extension that exposes some functionality (visuohaptic carving) for inclusion in scene designs. The layout of these scenes, which gives their digital objects shape and surface properties, and defines the events of actions performed in the scenes, is an activity pertinent to the realm of interaction design rather than engineering. Making this distinction is rarely articulated in surgery simulator design research, or at least in the discourses presented in the background. This does not mean that dedicated tools and particular workflows have not been used when making "content" for the simulators; on the contrary, for commercial systems it would not be a stretch to assume that they have. The distinction between "engine" and "content" is perhaps most evident in modern 3D game development. In the history of computer game production there is a clear trend towards more and more design effort going into "level design", i.e. creating the worlds where the players interact. Specialised level editors have been developed where designers can create maps and assets, and these editors have become more and more sophisticated, and more emphasis is being placed on the creative role of the level designers to, e.g., not only make interactive worlds but to tell the *story* that the player participates in [Shahrani, 2006, Labschütz et al., 2011]. For the sake of argument, the analogy is made that the underlying H3D API and the forssim extension correspond to the "game engine", while the design of a scene corresponds to the "level design". The creation of levels requires authoring CG objects, placement of these in a world, experimentation, scripting of behaviour and so on, which together form a workflow or pipeline where several programs are used for different parts, and results from one program (e.g. a CG object) are used as input in another (level editor). The scene, or patient case, in the Kobra simulator follows a similar pattern. In the following, the workflow used for the production of the most recent surgical scenes used in the Kobra simulator will be described.

1. A client dental school sends a CD-ROM disk with a pre-operative computed tomography scan of a recent, anonymised patient. Attached to the disk is a note of the medical problem and the intervention that the surgeon has performed and which they want gestalted in the simulator.

2. The CT image is loaded in *MeVisLab*[1], an image-processing application, where coarse filtering and surface extraction are performed. One or several iso-surfaces, denoting the transition from lower-density to higher-density sample values, are then exported as polygonal meshes. The bone surface is stored in one file, and the enamel of the teeth in another. These meshes are rather coarse, have holes and unwanted artefacts such as "spikes" caused by metal-containing dental fillings. They still provide a good image of the particular patient's unique anatomy.

3. The polygonal meshes are imported into professional interactive modelling programs: *3D Studio Max*[2] and *ZBrush*[3]. These are typical software applications that are the tools of the trade for professional 3D artists, and which they have built up a mastery and skill in using over the years. The 3D artist simultaneously designs a face, based on stock photo images, in a pose and facial expression that follow the position and tools of the surgical procedure, and the interactive anatomy (teeth and bone). For example, consider the pulled-up lip and the exposure of the bone in figure 4.4. The bone and teeth are well integrated in the artistically sculpted face, and a non-interactive wound hook is shown to pull up the lip and gingiva, which, as a

---

[1] http://www.mevislab.de/

[2] http://www.autodesk.com/products/3ds-max/overview

[3] http://pixologic.com/

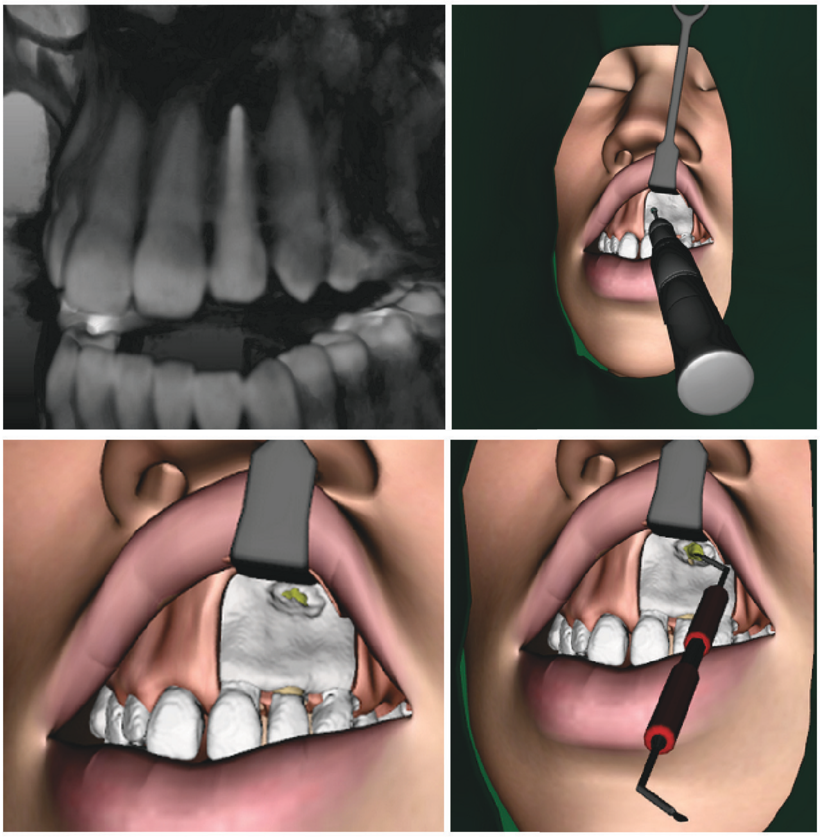

Figure 4.4: Gestalting apicectomy in the Kobra simulator. Upper left: original CT scan. Upper right: first step in the procedure. Note how the non-interactive face mesh is designed based on how the wound hook is deforming the lip and, as a side effect, the nose. Bottom: removal of infected tissue with an excavator.

side-effect, partly pulls up and deforms the nose. The non-interactive face model is exported with textures for direct inclusion in the simulator scene. The bone, teeth (enamel, dentin, pulp) and, in figure 4.4, infected root tissue are, to the extent they were possible to create at this stage, exported as surfaces to a voxelisation program: *3D-coat*[4].

4. The voxelisation program takes as input the tissue meshes created in the previous step, and samples them in order to export them as voxel volumes[5].

---

[4] http://3d-coat.com/

[5] Representing the object in the band-limited signal sense, i.e. as greyscale values; see background for the difference between these and the original binary definition.

5. The tissue layers are re-joined into a segmentation map volume, where each voxel is identified as belonging to either "air" or any of the pre-defined tissues: "bone", "enamel", "dentin", "pulp/nerve" and "infected tissue". This is currently achieved through the use of scripts calling the command-line tool *teem unu*[6].

6. The segmentation map is loaded in another image-processing software program that is capable of direct voxel editing: *itk-snap*. Here any fine tweaking of the shapes can be performed. The program is also used to mark regions where the end-user should carve, regions to avoid and mask regions that will be cleared as a result of the procedure, e.g the crown part of a tooth that will be removed after sectioning and prying with the elevator tool.

7. The scene objects are included via their file names in the X3D scene using a text editor (figure 4.3). Some part of this process can be assisted with *Blender*[7], an open-source 3D modelling tool with direct X3D support. The state transitions of the procedure simulation are coded in this file as well, e.g. that the end-user has to remove sufficient amount of bone in a particular region in order to progress. The resulting collection of files can then be loaded by the simulator executable.

8. Finally, the "material properties" of the various tissue segments are fine tuned with a novel interactive tool developed for the purpose. This tool, which will be discussed further in the next section, allows the designer to instantiate the scene as if run by the simulator, and tune visual, haptic and carving parameters while simultaneously seeing and feeling the result. When satisfied, the designer can save the scene, which then is ready for deployment and user testing.

The present process constitutes a prototype workflow in that it still requires manual actions to maintain it. For example, the voxelisation process causes the inter-object position information to be lost, so bone and teeth may not be aligned correctly. This is currently resolved by placing all objects in a reference box that is later removed, both actions being performed through the execution of scripts. Other conversion tricks involve changing bit depth after saving files in itk-snap. These conversions currently require keeping track of volume sizes, resolution and so on. The resolution loss that is caused by converting the "grey-scale" volume output from the voxelisation program into a binary label volume (material/no material) may have unwanted effects such as the disappearance of thin tissues, and that what is seen in the voxelisation program is not what one finally gets. Fine-tuning afterwards on a per-voxel basis and filtering scripts can be used to compensate for this to some extent. A serial production-ready pipeline should involve custom-made GUI tools that act as glue between the aforementioned professional tools and simplify the process. It needs to be easy to step back and forth in the workflow, and edit a file in one step without having to edit and execute scripts. The need for setting up a workflow and creating custom tools that streamline content creation and interaction design has been somewhat overlooked

---

[6] http://teem.sourceforge.net/unrrdu/

[7] http://www.blender.org

by the haptics and surgery simulation literature, but has been proved essential in our work with the Kobra simulator. I therefore argue that articulating this piece of the puzzle is one of the contributions of this thesis. Further work should consider how the workflow can be improved further.

Tools for tuning the user experience of visuohaptic carving

The last step in the workflow is dependent on the possibility of getting immediate sensory feedback on the changes made to the visuohaptic digital material. Even though no low-level source code needs to be changed, thanks to the X3D abstraction, it is not, as has become evident in the practical design work of making scenes for the Kobra simulator, sufficient to rely on performing the tuning by changing numerical values in a text editor. The editing-launching-testing-closing cycle is too inefficient to get the “just right” colour and carving resistance. Therefore a custom GUI-based editing tool has been developed.

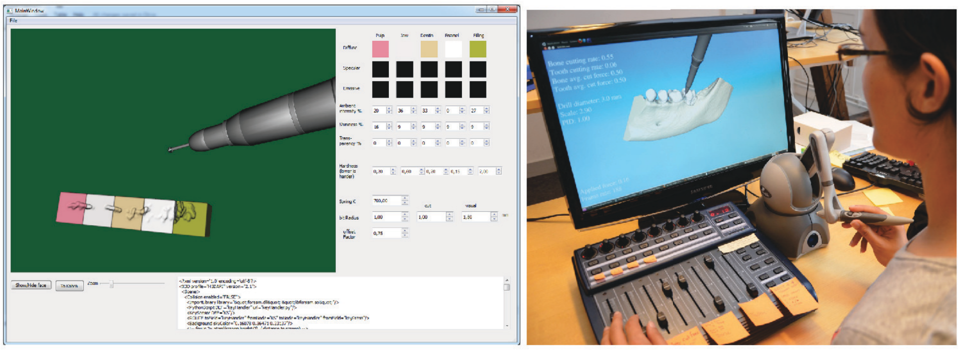

Figure 4.5: Two different tools made by the author for tuning the visual and haptic properties of the digital material influencing the carving experience. Left: GUI of the visuohaptic carving scene-tuning tool. Right: sketching tool with tangible controller

With this tool a designer can load X3D scenes with voxel-based CG objects featuring up to five segments (layers), and then define and tune the visual and haptic properties of each segment. In figure 4.5 (left) a set of five different boxes has been loaded. Each box can have different values. The visual properties follow the X3D standard and are Diffuse, Specular and Emissive colour values. These can be set with a colour palette familiar from the operating system used, and ambient intensity, shininess and transparency can be set with spinners or by entering a value in the range of 0 to 100. In addition, a hardness (carving rate) can be specified for each object. Applied to scenes created for the Kobra simulator, each segment represents human tissue: enamel, dentin, pulp, bone and filling. The drill’s properties can be set too, which includes spring constant (stiffness), the radius of the avatar used in the haptic algorithm, the radius of the avatar used in the carving algorithm and the visual avatar radius. The purpose of having different dimensions of the carving and haptic radii is that if a constraint-based algorithm is used, then a same-size carving avatar will actually never carve into the object since the haptic avatar is restricted from penetrating the surface. If the penalty-based algorithm is used, it will on the other hand always penetrate slightly, which is why a visual and cutting avatar that is smaller than the haptic avatar can be used to reduce the experience of having the drill penetrating the surface when the user is just touching it slightly. The designer can tweak and tune these parameters and instantly explore how they affect the user experience.

While this tuning tool was designed to be used in the last fine-tuning stage of a scene design, it can be used as its prototypical ancestor, which will be discussed next, in the very early stages of an application's design. This other prototyping tool (figure 4.5 right) was designed to support *sketching* with the digital material of visuohaptic carving, as a way to explore the design space early on, before settling on which hardware device to use and which level of realism and fidelity is required or desired [Paper C].

In order to avoid having to use the mouse and keyboard to change values, a tangible controller with motorised sliders and knobs was used (figure 2.4). The tangible controls enabled the designer to use one hand for tuning a value, without looking, while feeling the effect with the other. This resulted in a closer connection with what the property actually implied in terms of experience. In addition to stiffness and carving rate, *size* was seen as a highly significant contribution to the haptic experience. The subjective impression was that when the object was scaled up it could be much better perceived and, e.g., carving away bone without hurting underlying teeth felt “good” even with the lower-cost Omni haptic device. The research contribution lies not in determining whether this indeed is true (although that is a worthy subject on its own), but in the empowerment that the tool gives the designer to subjectively explore the effect of changing size and other parameters.

4.2 Interaction Design for Surgery Simulators

The previous section proposed that visuohaptic carving and WoodenHaptics can be seen as design resources, as well as proposing what is required to prepare the underlying technologies and how an associated workflow practice may look. In this section it will be shown how visuohaptic carving, its tools and workflow have been applied in the design of patient case scenes in the Kobra simulator. In particular, it will be discussed why designing with a starting point in the availability of these design resources is different and novel.

Important aspects for surgery simulator design

The *research through design* work with the Kobra oral surgery simulator has resulted in the articulation of a number of aspects that have been deemed to be of particular importance to consider in the design of a successful simulation-based teaching tool in the context of dental education [Paper A]. These aspects are: realism and surgical relevance, the social setting of surgery teaching, visual and haptic aesthetics, and the qualities of the physical design. These will not all be elaborated on here. What is important for the present discussion is that the final design of Kobra supported the teaching of oral surgery procedures through interactive gestalts (creative representations) of authentic surgical patient cases. Rather than striving for the most realistic representation possible, the design focus was on presenting the unique patient and the steps of the appropriate procedure that the student should take, under the supervision of an experienced surgeon-teacher.

A role for creative haptic interaction design in surgery simulation

One of the results of the Kobra project is that the same design resources can be used to gestalt different procedures which, at first sight, seem to require more advanced technology. What has overcome the perceived limitations of the technology is the creative interaction design. Some examples will be given that show how this is done, and which work to prove that interaction design can have an important role in advancing the state of the art of surgery simulation.

Figure 4.4 shows two different steps of an apicectomy, a surgical procedure that is performed when a previously root-filled tooth has been infected at the root apex. It begins with opening up the soft tissue gingiva covering the bone, carving the bone with the dental drill to access the root apex, removing the infected tissue that surrounds the root apex and replacing it with another material. The simulation begins with the correct gingiva tissue flap prepared (figure 4.4 top right) and the student may start removing bone in the correct location. This requires a good understanding of the anatomy and translating between the 2D x-ray image (figure 4.4 top left) and the 3D representation. When the infected tissue is exposed (figure 4.4 bottom left), the student can switch to an excavator tool and clean the cavity. Technically this action is implemented exactly the same as carving with the drill, but only affects voxels belonging to the infected tissue segment. In reality, the infected tissue is connected and carefully separating it from the bone cavity is a rather delicate process, an interaction for which generating the "correct" forces would be very difficult. Nevertheless this simplified representation was accepted by the client dental school. It was still possible to detect whether all of the tissue was removed and haptically inspect the cavity afterwards.

A second example is the surgical extraction of impacted teeth (figure 4.6). This case is challenging in that it requires removal of two teeth at once and the operator needs to be extra careful in navigating the patient-specific anatomy. The procedure involves removing surrounding bone for exposure and sectioning the teeth prior to extraction. Then the teeth are extracted with the elevator tool if they are loose and divided into small enough pieces. In particular, this use of the elevator with a loose tooth is complex to simulate true to nature. This fact has led other simulator designers to avoid the task completely. For example, Pohlenz et al. [Pohlenz et al., 2010] write that "tooth extraction or surgical removal, although the most commonly performed surgical procedure in dentistry, could not be reproduced with this model because the complex movements and the resulting forces cannot currently be adequately simulated". In the Kobra case the elevator is only visually different from the inactive dental drill. The same 3-DoF haptic algorithm governs its motion and feedback. This limited realism does not hinder students from operating the tool differently. It has been observed, for example, that students change their grip in order to use it as they have been taught [Paper A]. In a real procedure it is sometimes required to drill for a while, then probe with the elevator if the tooth is loose enough, then go back to drilling and so on. The same is possible in the Kobra scene, where the state machine,

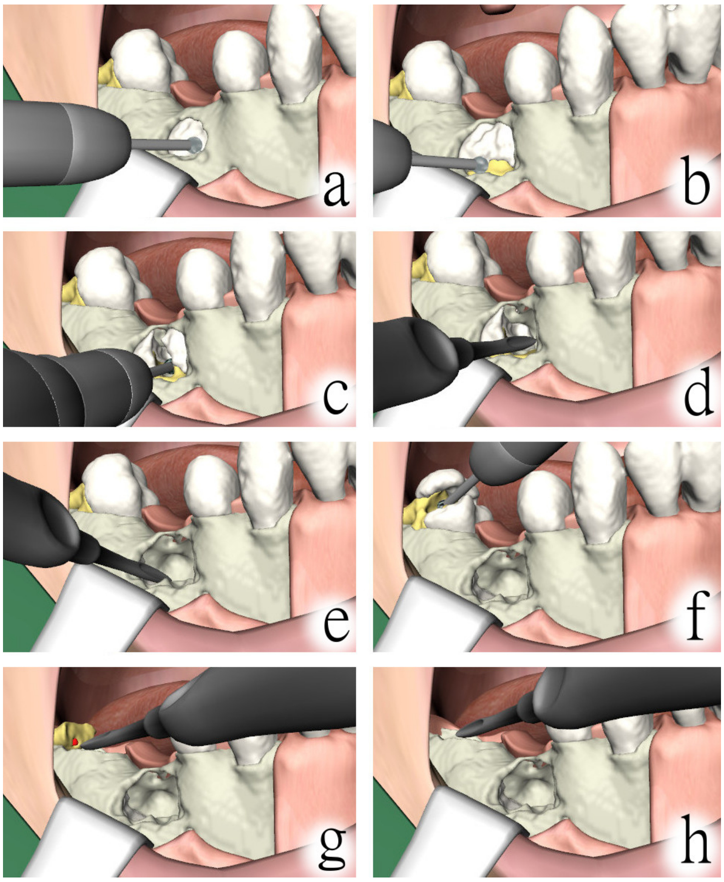

Figure 4.6: Cropped screen-shot from the simulation. This scenario gestalts surgical extraction of one impacted (a-e) and one third molar tooth (f-h), and is based on a real patient case. The proximity to another impacted tooth and the mandibular nerve adds challenge to the case. The user begins by dissecting bone without carving into the impacted teeth to get sufficient visibility and access (a). Then, the user alternates between sectioning the teeth with a surgical drill (b,c,f) and cracking/extracting them with the elevator (d,e,g,h).

triggered by the elevator, governs if enough bone has been removed to extract the tooth in order to progress. Taken together, this design allows for the performance of many of the steps in the procedure in figure 4.6: carving in the correct regions, sectioning teeth and prying with the elevator. Technically the simulation is modest, and especially the haptic feedback of the elevator, but it has been possible to form an educational experience anyway. Herein lies the role of the interaction design for surgery simulators; through creative use of the design resources, in this case visuohaptic carving, a surgical procedure can be gestalted even if all forces cannot be adequately simulated.

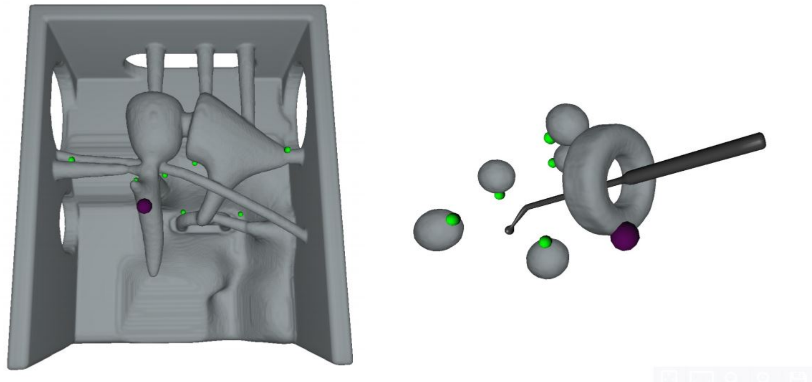

Figure 4.7: The task of the study was to touch all the green spheres with the instrument shown to the right without pushing it excessively into the surroundings. Left: “Ear” scene. Model of middle ear anatomy. Right: “Port” scene with a narrow corridor.

The value of asymmetric haptic feedback

As discussed more thoroughly in the background, any real tool-mediated interaction is naturally 6-DoF. It is an obvious limitation that the present haptic interaction model in the Kobra simulator is restricted to a single contact sphere situated at the tip of the instrument. With the elevator in particular this can be disturbing when the prolonged blade and shaft are allowed to penetrate any surface. To overcome this issue the whole rigid body of the instrument needs to be considered for collision detection and resolution. In effect, a 6-DoF haptic algorithm is required. It has been a common misconception, however, that a 6-DoF algorithm necessarily requires a fully actuated 6-DoF haptic device, i.e. a device that can generate torque feedback. These devices are much more complex and costly than the under-actuated ones that have 6-DoF sensing but only 3-DoF actuation. For this reason there was motivation to measure the effect of 6-DoF rendering and torque display, and a within-group experimental study with twelve subjects was performed [Paper E]. The task was to navigate a probe in two static virtual environments and touch a number of points without excessive contact with the surrounding environment. All subjects performed the task in two different scenes and with three haptic-rendering modes: 3-DoF sphere-based, 6-DoF rigid body without torque feedback, and 6-DoF rigid body with torque feedback. Completion time and number of errors were recorded, and subjective perceived performance was gathered by a questioner.

The results show a significant difference between 3-DoF and 6-DoF rendering, but no significant difference between displaying torque and no torque. The implication for design is that it can be worthwhile to invest in using a 6-DoF algorithm even when budget or other constraints make using a fully actuated device impractical.

Chapter 5

Discussion

Returning to the discussion initiated in the background chapter on surgery simulation usage in practice, it is now possible to discuss a more multi-faceted view of haptic technology development for surgery simulation. I postulate that a *naive technical view* of simulator development holds that the only, or at least primary, task at hand when designing a simulator should be to study and measure elements of nature, i.e. the interaction forces of surgical instruments, and then replicate these as faithfully as possible in a machine. Advancing the state of the art in simulator development is then restricted to advancing the level of realism with which tissues and instruments are represented in the simulator, and the designer is restricted to formulating requirements. This view can sometimes be read between the lines in the literature. Interestingly enough, there are several simulators in the related work, including the Kobra simulator, which deviate from the naive technical view, e.g. the work by the VOXEL-MAN group [Pohlenz et al., 2010]. It is especially evident with the inclusion of features that are non-naturalistic, e.g. the ability to render tissues transparently. The fact that simulators have been used as teaching equipment, building on resemblances of situations in the operating room partly enacted by the elements of the equipment and partly enacted by the participants [Johnson, 2004], also supports a more multi-faceted view. Freeing simulator design from technically mimicking reality opens the way for a multitude of opportunities, which obviously also puts new demands on the designer to come up with novel solutions.

In this thesis I have shown how established spatial haptic technology can be used to *gestalt* surgical procedures through interaction design. The procedures gestalted in the Kobra simulator, however, have not been realised in just any medium or material. They are gestalted using what in this thesis is referred to as *visuohaptic carving*. Alternatively, one can imagine designing a simulator for oral surgery using only a screen-based point-and-click interface, which is familiar both to interaction designers and programmers. As noted by Rystedt and Sjöblom [Rystedt and Sjöblom, 2012], the goal of the design is to be *relevant* rather than necessarily *realistic* in the strict sense of the word. The haptic medium, however, offers unique opportunities compared to the point-and-click alternatives. Designing with visuohaptic carving is therefore something different and constitutes a different practice with its own set of tools, resources and workflow. This argumentation can now be rephrased in the light of the research questions introduced in chapter 1, starting with the most particular.

How can novel design resources, tools and associated practices for spatial haptic interaction design be leveraged for surgery simulation design?

By using the design resource *visuohaptic carving* in an application, an interaction designer can gestalt surgical interventions in a highly interactive manner in the form of interactive patient cases. This may help teachers explain important spatial relationships to students, and students can themselves practice hands-on the critical steps of a procedure. The patient cases presented in this thesis represent surgical interactions that at first sight would have required more advanced technologies. Some of these technical limitations have been overcome by creative interaction design, using the presented tools and work-flow. Examples include the use of carving for simulating excavation of infected tissues, and haptic rendering confined to a tip-located sphere for prying out impacted teeth. It has been argued why creative interaction design requires access to suitable tools and materials i.e. design resources, with which designers can sketch and form prototypes even in early phases of development. Therefore is the provision of design resources, tools and associated practises proposed to benefit surgery simulator development.

How can spatial haptic technologies be prepared for interaction design?

One way to prepare spatial haptic technologies is to turn them into design resources. What constitutes a design resource in this context is encapsulation of technical nuances while exposing important properties for forming and tuning the interaction experiences. This is done for haptic hardware with the WoodenHaptics starting kit, and for software with the implementation of visuohaptic carving in forssim, the software library. The library was however not sufficient to be an effective design resource on its own, but requires custom tools and well planned work-flow, which also were created and discussed in this thesis.

Why is it important to prepare haptic technology for interaction design?

That haptic interaction design can be useful is supported by the designs brought forward in this thesis. But why is it important to single out *preparation* as something essential in this design work? It has been mentioned that implementing advanced haptic rendering algorithms and engineering haptic devices is challenging and time consuming. Nevertheless is a fully implemented system required to feel the actual result of a design. This makes it difficult to predict what can be created and how it eventually will feel. Well-known design methods such as paper prototyping works well for systems where there is a strong link between the anticipated result and the paper sketch. In other words, when it is well understood how the system can be implemented based on the prototype alone. Even when paper prototyping is reserved for conceptual design may lack of access to functional haptic design resources limit the designers ability to draft useful proposals, in particular since there are no clear resistance against sketching naive solutions that cannot be feasibly implemented, something that has in general been referred to as cargo cult design [Holmquist, 2005].

Preparing for interaction design is in this thesis postulated as something slightly different than some previous work that focuses on introducing the underlying engineering concepts to newcomers [Hayward and MacLean, 2007]. While some high-level understanding of haptics is essential for creating good design, can the access to design resources, such as WoodenHaptics, relieve designers from handling every single detail that makes up the technical solution. The downside from not learning the details is obviously the limitation in solution space, why the approaches naturally complement each other. With this said, it is worth noting that WoodenHaptics and visuohaptic carving is not created only to bring advanced haptic technologies to a wider design community of non-engineers. In fact, as an engineer, I built these tools and resources primarily to use them myself in my design work. The need for them in the design work of the Kobra simulator is also suggesting that their value persists even when the one who create the tools and use them are the same.

Materiality

The present discussion resonates with the contemporary discourse in HCI regarding *materiality* [Fernaeus and Sundström, 2012]. A materiality perspective acknowledges the unique properties of each "digital material" which the interaction designer turns into a product. The word material can be confusing, since in everyday language it can be first thought to be restricted to passive lump of matter, but should actually be seen as a cultural entity, as eloquently put by Solsona [Solsona Belenguer, 2015]:

> A material perspective is not a property of things-in-themselves, but manifests itself when combined with knowledge of how to shape a material and the skills required to do so. Without knowledge and skill, there is matter rather than material. For example, wood becomes a material when you know how to work with it; otherwise it is just wood and difficult to use for anything. Hence, material is not a physical manifestation, but instead, and this is how we engineers would benefit from this approach, a material manifests itself when combined with accumulated knowledge of the material and contexts in which it is used.

A material is then, by this definition, matter plus knowledge. The design resources, WoodenHaptics and implementation of visuohaptic carving in the forssim library, that have been described in this thesis can then be considered digital materials, insofar they are coupled with a meaningful creative practice. The WoodenHaptics starting kit and the tools for tuning the haptic rendering properties are instrumental to supporting this creative practice. It should be no logical barriers for including software libraries in what can constitute a material, even though its matter part is restricted to a string of ones and zeros. The benefit of using Solsona's definition is that we can begin talking about how well a particular digital technology, e.g. a software library, acts as a material for interaction design and what particular knowledge and skills are required. There are many software libraries available that has a steep learning curve, even for experienced developers. The idea of *preparing* technology that I have used in this thesis, can be one approach that engineers and computer scientists can take if they wish to turn their components into a malleable interaction design material.

Limitations and future work

The present work does not investigate *how well* the Kobra simulator works or how well the interactive patient cases gestalt the surgical procedures they intend to support the learning of. In other words, there are no studies of the learning impact of this technology. The aim, however, has not been that, but to present plausible designs, i.e. designs that are grounded in empirical design work, and, especially, what has been required to prepare for such design work. Nevertheless, to what extent this technology and interaction design are useful and how realistic the representations need to be and so on will remain an open question.

Another limitation that warrants future work is how well received WoodenHaptics and visuohaptic carving and their related tools and workflow will be in the interaction design community. To what extent can interaction design practitioners pick up the design resources and make use of them? What competencies are required and how long will it take them to master them? How refined need the tools be?

The haptic rendering discussed in this thesis is primarily restricted to sphere-based carving. There are many more features described in computer science literature on how to simulate friction, textures, and complex avatar shapes, not to mention soft tissue deformation and cutting. Obvious future work include exploring how these too can be turned into design resources and what knowledge and tools are necessary to effectively work creatively with them, i.e. turning them into digital materials for interaction design.

Bibliography

[Agus et al., 2003] Agus, M., Giachetti, A., Gobbetti, E., Zanetti, G., and Zorcolo, A. (2003). Real-Time Haptic and Visual Simulation of Bone Dissection. *Presence: Teleoperators and Virtual Environments*, 12(1):110–122.

[Avila and Sobierajski, 1996] Avila, R. S. and Sobierajski, L. M. (1996). A haptic interaction method for volume visualization. In *Visualization'96. Proceedings.*, pages 197–204. IEEE.

[Bakker et al., 2010] Bakker, D., Lagerweij, M., Wesselink, P., and Vervoorn, M. (2010). Transfer of manual dexterity skills acquired in the simodont, a dental haptic trainer with a virtual environment, to reality: a pilot study. *Bio-Algorithms and Med-Systems*, 6(11):21–24.

[Bakr et al., 2013] Bakr, M. M., Massey, W., and Alexander, H. (2013). Evaluation of simodont® haptic 3d virtual reality dental training simulator. *International Journal of Dental Clinics*, 5(4).

[Barbagli and Salisbury, 2003] Barbagli, F. and Salisbury, K. (2003). The effect of sensor/actuator asymmetries in haptic interfaces. In *Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS'03)*, HAPTICS '03, pages 140–, Washington, DC, USA. IEEE Computer Society.

[Behr et al., 2010] Behr, J., Jung, Y., Keil, J., Drevensek, T., Zoellner, M., Eschler, P., and Fellner, D. (2010). A scalable architecture for the html5/x3d integration model x3dom. In *Proceedings of the 15th International Conference on Web 3D Technology*, pages 185–194. ACM.

[Bejczy, 1980] Bejczy, A. K. (1980). Sensors, controls, and man-machine interface for advanced teleoperation. *Science*, 208(4450):1327–1335.

[Bjelland and Tangeland, 2007] Bjelland, H. V. and Tangeland, K. (2007). User-centered design proposals for prototyping haptic user interfaces. In *Proceedings of HAID07*, pages 110–120, Berlin, Heidelberg. Springer.

[Bowman et al., 2004] Bowman, D. A., Kruijff, E., LaViola, J. J., and Poupyrev, I. (2004). *3D User Interfaces: Theory and Practice*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA.

[Broms, 2014] Broms, L. (2014). Storyforming: Experiments in creating discursive engagements between people, things and environments.

[Brooks Jr, 1996] Brooks Jr, F. P. (1996). The computer scientist as toolsmith ii. *Communications of the ACM*, 39(3):61–68.

[Brooks Jr et al., 1990] Brooks Jr, F. P., Ouh-Young, M., Batter, J. J., and Jerome Kilpatrick, P. (1990). Project gropehaptic displays for scientific visualization. In *ACM SIGGraph computer graphics*, volume 24, pages 177–185. ACM.

[Brown et al., 2008] Brown, T. et al. (2008). Design thinking. *Harvard business review*, 86(6):84.

[Brunnström, 1997] Brunnström, L. (1997). *Svensk industridesign: En 1900-talshistoria*. Norstedt.

[Brutzman and Daly, 2007] Brutzman, D. and Daly, L. (2007). *X3D: extensible 3D graphics for Web authors*. Morgan Kaufmann.

[Buchanan, 2001] Buchanan, J. A. (2001). Use of simulation technology in dental education. *Journal of Dental Education*, 65(11):1225–1231.

[Burdea and Coiffet, 2003] Burdea, G. and Coiffet, P. (2003). Virtual reality technology. *Presence: Teleoperators and virtual environments*, 12(6):663–664.

[Buxton, 2007] Buxton, B. (2007). *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan Kaufmann.

[Chan, 2011] Chan, S. (2011). Constraint-based six degree-of-freedom haptic rendering of volume-embedded isosurfaces. In *IEEE World Haptics*.

[Chan, 2014] Chan, S. (2014). *Haptic Rendering of Medical Image Data for Surgical Rehearsal*. PhD thesis, Stanford University.

[Conti et al., 2007] Conti, F., Morris, D., Barbagli, F., and Sewell, C. (2007). Chai 3d. *Invited talk at MMVR 2007, available at http://www.chai3d.org/papers/CHAI3D-MMVR-2007.pdf.*

[Cormier et al., 2011] Cormier, J., Pasco, D., Syllebranque, C., and Querrec, R. (2011). Virteasy a haptic simulator for dental education. In *The 6th International Conference on Virtual Learning*, volume 156, pages 61–68.

[Craig, 2005] Craig, J. J. (2005). *Introduction to robotics: mechanics and control*, volume 3. Pearson Prentice Hall Upper Saddle River.

[Daly and Brutzman, 2007] Daly, L. and Brutzman, D. (2007). X3d: extensible 3d graphics standard. *IEEE Signal Processing Magazine, Nov. 2007.*

[De Felice et al., 2009] De Felice, F., Attolico, G., and Distante, A. (2009). Configurable design of multimodal non visual interfaces for 3d ve's. In *Haptic and Audio Interaction Design*, pages 71–80. Springer.

[Dearden, 2006] Dearden, A. (2006). Designing as a conversation with digital materials. *Design Studies*, 27(3):399 – 421.

[Engel et al., 2004] Engel, K., Hadwiger, M., Kniss, J. M., Lefohn, A. E., Salama, C. R., and Weiskopf, D. (2004). Real-time volume graphics. In *ACM Siggraph 2004 Course Notes*, page 29. ACM.

[Fernaeus, 2007] Fernaeus, Y. (2007). *Let's Make a Digital Patchwork: Designing for Childrens Creative Play with Programming Materials*. PhD thesis, Department of Computer and Systems Sciences, Stockholm University.

[Fernaeus and Sundström, 2012] Fernaeus, Y. and Sundström, P. (2012). The material move how materials matter in interaction design research. In *Proceedings of the Designing Interactive Systems Conference*, pages 486–495. ACM.

[Flodin, 2009] Flodin, M. (2009). Betydelsen av skuggning vid volymrenderad visualisering i en multimodal simulator för operativ extraktion av visdomständer. Master's thesis, KTH Royal Institute of Technology.

[Foley et al., 1994] Foley, J. D., Van Dam, A., Feiner, S. K., Hughes, J. F., and Phillips, R. L. (1994). *Introduction to computer graphics*. Addison-Wesley Reading.

[Forsslund, 2008] Forsslund, J. (2008). Simulator för operativ extraktion av visdomständer. Master's thesis, KTH Royal Institute of Technology.

[Forsslund, 2011] Forsslund, J. (2011). Is realism the most important property of a visuohaptic surgery simulator? In *Society in Europe for Simulation Applied to Medicine annual congress*.

[Forsslund et al., 2013] Forsslund, J., Chan, S., Selesnick, J., Salisbury, K., Silva, R. G., and Blevins, N. H. (2013). The effect of haptic degrees of freedom on task performance in virtual surgical environments. In *Studies in Health Technology and Informatics, Volume 184: Medicine Meets Virtual Reality 20*, pages 129 – 135.

[Forsslund et al., 2009] Forsslund, J., Salinas, E.-L., and Palmerius, K.-J. (2009). A user-centered designed foss implementation of bone surgery simulations. In *EuroHaptics conference, 2009 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009. Third Joint*, pages 391–392.

[Galyean and Hughes, 1991] Galyean, T. A. and Hughes, J. F. (1991). Sculpting: An interactive volumetric modeling technique. In *ACM SIGGRAPH Computer Graphics*, volume 25, pages 267–274. ACM.

[Gaver et al., 2009] Gaver, W., Bowers, J., Kerridge, T., Boucher, A., and Jarvis, N. (2009). Anatomy of a failure: how we knew when our design went wrong, and what we learned from it. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2213–2222. ACM.

[Gershenfeld, 2008] Gershenfeld, N. (2008). *Fab: the coming revolution on your desktop–from personal computers to personal fabrication*. Basic Books.

[Gottlieb et al., 2013] Gottlieb, R., Vervoorn, J. M., and Buchanan, J. (2013). Simulation in dentistry and oral health. In *The Comprehensive Textbook of Healthcare Simulation*, pages 329–340. Springer.

[Greenberg and Fitchett, 2001] Greenberg, S. and Fitchett, C. (2001). Phidgets: easy development of physical interfaces through physical widgets. In *Proceedings of UIST*, pages 209–218. ACM.

[Hartmann et al., 2006] Hartmann, B., Klemmer, S. R., Bernstein, M., Abdulla, L., Burr, B., Robinson-Mosher, A., and Gee, J. (2006). Reflective physical prototyping through integrated design, test, and analysis. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, pages 299–308. ACM.

[Hayward and MacLean, 2007] Hayward, V. and MacLean, K. E. (2007). Do it yourself haptics: part i. *Robotics & Automation Magazine, IEEE*, 14(4):88–104.

[Hindmarsh et al., 2014] Hindmarsh, J., Hyland, L., and Banerjee, A. (2014). Work to make simulation work: Realism, instructional correction and the body in training. *Discourse Studies*, 16(2):247–269.

[Höhne et al., 1995] Höhne, K. H., Pflesser, B., Pommert, A., Riemer, M., Schiemann, T., Schubert, R., and Tiede, U. (1995). A new representation of knowledge concerning human anatomy and function. *Nature medicine*, 1(6):506–511.

[Holmquist, 2005] Holmquist, L. E. (2005). Prototyping: generating ideas or cargo cult designs? *interactions*, 12(2):48–54.

[Ilstedt Hjelm, 2004] Ilstedt Hjelm, S. (2004). *Making sense: design for well-being*. PhD thesis, KTH Royal Institute of Technology.

[Itkowitz et al., 2005] Itkowitz, B., Handley, J., and Zhu, W. (2005). The openhaptics toolkit: a library for adding 3d touch navigation and haptics to graphics applications. In *Eurohaptics Conference, 2005 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2005. World Haptics 2005. First Joint*, pages 590–591. IEEE.

[Jasinevicius et al., 2004] Jasinevicius, T. R., Landers, M., Nelson, S., and Urbankova, A. (2004). An evaluation of two dental simulation systems: virtual reality versus contemporary non-computer-assisted. *Journal of dental education*, 68(11):1151–1162.

[Johnson, 2004] Johnson, E. (2004). *Situating simulators: The integration of simulations in medical practice*. PhD thesis, Linköping University.

[Johnson, 2007] Johnson, E. (2007). Surgical simulators and simulated surgeons: Reconstituting medical practice and practitioners in simulations. *Social Studies of Science*, 37(4):585–608.

[Johnson, 2009] Johnson, E. (2009). Extending the simulator: Good practice for instructors using medical simulators. In Dieckmann, P., editor, *Using Simulators and Simulations for Education, Training and Research*. Pabst.

[Johnson et al., 2004] Johnson, E., Ström, P., Kjellin, A., Wredmark, T., and Felländer-Tsai, L. (2004). Evaluating instruction of medical students with a haptic surgical simulator: The importance of coordinating students' perspectives. *Journal on Information Technology in Healthcare*, 2(3):155–163.

[Joseph et al., 2014] Joseph, D., Jehl, J.-P., Maureira, P., Perrenot, C., Miller, N., Bravetti, P., Ambrosini, P., and Tran, N. (2014). Relative contribution of haptic technology to assessment and training in implantology. *BioMed research international*, 2014.

[Kern, 2009] Kern, T. A. (2009). *Engineering haptic devices*. Springer.

[Knörig et al., 2009] Knörig, A., Wettach, R., and Cohen, J. (2009). Fritzing: a tool for advancing electronic prototyping for designers. In *Proceedings of TEI*, pages 351–358. ACM.

[Labschütz et al., 2011] Labschütz, M., Krösl, K., Aquino, M., Grashäftl, F., and Kohl, S. (2011). Content creation for a 3d game with maya and unity 3d. *Institute of Computer Graphics and Algorithms, Vienna University of Technology*.

[Lawson, 2005] Lawson, B. (2005). *How designers think: the design process demystified (4th ed.)*. Routledge.

[Lederman and Klatzky, 2009] Lederman, S. and Klatzky, R. (2009). Haptic perception: A tutorial. *Attention, Perception, & Psychophysics*, 71(7):1439–1459.

[Ledo et al., 2012] Ledo, D., Nacenta, M. A., Marquardt, N., Boring, S., and Greenberg, S. (2012). The haptictouch toolkit: enabling exploration of haptic interactions. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*, pages 115–122. ACM.

[Levine et al., 2013] Levine, A. I., DeMaria Jr, S., Schwartz, A. D., and Sim, A. J. (2013). *The comprehensive textbook of healthcare simulation*. Springer Science & Business Media.

[Lin and Otaduy, 2008] Lin, M. C. and Otaduy, M. (2008). *Haptic rendering: foundations, algorithms, and applications*. CRC Press.

[Lorensen and Cline, 1987] Lorensen, W. E. and Cline, H. E. (1987). Marching cubes: A high resolution 3d surface construction algorithm. In *ACM siggraph computer graphics*, volume 21, pages 163–169. ACM.

[Löwgren, 2007] Löwgren, J. (2007). Interaction design, research practices, and design research on the digital materials. *Available at webzone.k3.mah.se/k3jolo.*

[Lund et al., 2011] Lund, B., Fors, U., Sejersen, R., Sallnäs, E.-L., and Rosén, A. (2011). Student perception of two different simulation techniques in oral and maxillofacial surgery undergraduate training. *BMC medical education*, 11(1):82.

[MacKenzie et al., 1991] MacKenzie, I. S., Sellen, A., and Buxton, W. A. (1991). A comparison of input devices in element pointing and dragging tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 161–166. ACM.

[MacLean, 2000] MacLean, K. E. (2000). Designing with haptic feedback. In *Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on*, volume 1, pages 783–788. IEEE.

[MacLean, 2008] MacLean, K. E. (2008). Haptic interaction design for everyday interfaces. *Reviews of Human Factors and Ergonomics*, 4(1):149–194.

[MacLean and Hayward, 2008] MacLean, K. E. and Hayward, V. (2008). Do it yourself haptics: Part ii [tutorial]. *Robotics & Automation Magazine, IEEE*, 15(1):104–119.

[Massie, 1993] Massie, T. (1993). Design of a three degree of freedom force-reflecting haptic interface. Master's thesis, MIT.

[Massie and Salisbury, 1994] Massie, T. H. and Salisbury, J. K. (1994). The phantom haptic interface: A device for probing virtual objects. In *ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems, vol. 55, no. 1, pp. 295-300.*

[Mellis and Buechley, 2012] Mellis, D. A. and Buechley, L. (2012). Case studies in the personal fabrication of electronic products. In *Proceedings of the Designing Interactive Systems Conference*, pages 268–277. ACM.

[Mellis et al., 2011] Mellis, D. A., Gordon, D., and Buechley, L. (2011). Fab fm: the design, making, and modification of an open-source electronic product. In *Proc. TEI*, pages 81–84. ACM.

[Mellis et al., 2013] Mellis, D. A., Jacoby, S., Buechley, L., Perner-Wilson, H., and Qi, J. (2013). Microcontrollers as material: crafting circuits with paper, conductive ink, electronic components, and an untoolkit. In *TEI*, pages 83–90.

[Miao et al., 2009] Miao, M., Köhlmann, W., Schiewe, M., and Weber, G. (2009). Tactile paper prototyping with blind subjects. In *Haptic and Audio Interaction Design*, volume 5763 of *LNCS*, pages 81–90. Springer.

[Minsky, 1995] Minsky, M. D. R. (1995). *Computational haptics: the sandpaper system for synthesizing texture for a force-feedback display*. PhD thesis, Queen's University.

[Moen, 2006] Moen, J. (2006). *KinAesthetic movement interaction: designing for the pleasure of motion*. PhD thesis, KTH Royal Institute of Technology.

[Molander, 1993] Molander, B. (1993). *Kunskap i handling*. Daidalos. A recent edition in English is available as Molander, B. (2015) The Practice of Knowing and Knowing in Practices.

[Monk et al., 1993] Monk, A., Wright, P., Haber, J., and Davenport, L. (1993). Apendix 1–cooperative evaluation: A run-time guide. *Improving your human-computer interface: a practical technique, Prentice-Hall*.

[Morimoto et al., 2014] Morimoto, T. K., Blikstein, P., and Okamura, A. M. (2014). [d81] hapkit: An open-hardware haptic device for online education. In *Haptics Symposium (HAPTICS), 2014 IEEE*, pages 1–1. IEEE.

[Moussette, 2012] Moussette, C. (2012). *Simple Haptics - Sketching Perspectives for the Design of Haptic Interactions*. PhD thesis, Umeå Institute of Desgin.

[Moussette and Banks, 2011] Moussette, C. and Banks, R. (2011). Designing through making: exploring the simple haptic design space. In *Proceedings of TEI11*, pages 279–282, New York, NY, USA. ACM.

[Moussette and Dore, 2010] Moussette, C. and Dore, F. (2010). Sketching in hardware and building interaction design: tools, toolkits and an attitude for interaction designers. *Proc. of Design Research Society*.

[Mueller et al., 2012] Mueller, S., Lopes, P., and Baudisch, P. (2012). Interactive construction: interactive fabrication of functional mechanical devices. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*, pages 599–606. ACM.

[Nelson and Stolterman, 2012] Nelson, H. G. and Stolterman, E. (2012). *The design way: Intentional change in an unpredictable world, second ed.* Educational Technology.

[Norman, 2005] Norman, D. A. (2005). *Emotional Design: Why We Love (or Hate) Everyday Things*. Basic Books.

[Okamura et al., 2002] Okamura, A. M., Richard, C., Cutkosky, M. R., et al. (2002). Feeling is believing: Using a force-feedback joystick to teach dynamic systems. *Journal of Engineering Education*, 91(3):345–350.

[O'Modhrain et al., 2015] O'Modhrain, S., Giudice, N., Gardner, J., and Legge, G. (2015). Designing media for visually-impaired users of refreshable touch displays: Possibilities and pitfalls.

[Ortega et al., 2007] Ortega, M., Redon, S., and Coquillart, S. (2007). A six degree-of-freedom god-object method for haptic display of rigid bodies with surface properties. *Visualization and Computer Graphics, IEEE Transactions on*, 13(3):458–469.

[Otaduy et al., 2013] Otaduy, M., Garre, C., Lin, M. C., et al. (2013). Representations and algorithms for force-feedback display. *Proceedings of the IEEE*, 101(9):2068–2080.

[Palmerius et al., 2008] Palmerius, K. L., Cooper, M., and Ynnerman, A. (2008). Haptic rendering of dynamic volumetric data. *IEEE Transactions on Visualization and Computer Graphics*, 14(2):263–276.

[Panëels et al., 2010] Panëels, S., Roberts, J. C., Rodgers, P. J., et al. (2010). Hitproto: a tool for the rapid prototyping of haptic interactions for haptic data visualization. In *Haptics Symposium, 2010 IEEE*, pages 261–268. IEEE.

[Petersik et al., 2003] Petersik, A., Pflesser, B., Tiede, U., Höhne, K.-H., and Leuwer, R. (2003). Realistic haptic interaction in volume sculpting for surgery simulation. In *Surgery Simulation and Soft Tissue Modeling*, pages 194–202. Springer.

[Pflesser et al., 2002] Pflesser, B., Petersik, A., Tiede, U., Höhne, K. H., and Leuwer, R. (2002). Volume cutting for virtual petrous bone surgery. *Computer Aided Surgery*, 7(2):74–83.

[Pohlenz et al., 2010] Pohlenz, P., Gröbe, A., Petersik, A., Von Sternberg, N., Pflesser, B., Pommert, A., Höhne, K.-H., Tiede, U., Springer, I., and Heiland, M. (2010). Virtual dental surgery as a new educational tool in dental school. *Journal of Cranio-Maxillofacial Surgery*, 38(8):560–564.

[Preim and Bartz, 2007] Preim, B. and Bartz, D. (2007). *Visualization in medicine: theory, algorithms, and applications*. Morgan Kaufmann.

[Prentice, 2005] Prentice, R. (2005). The anatomy of a surgical simulation: The mutual articulation of bodies in and through the machine. *Social Studies of Science*, pages 837–866.

[Prescher et al., 2010] Prescher, D., Weber, G., and Spindler, M. (2010). A tactile windowing system for blind users. In *Proceedings of the 12th international ACM SIGACCESS conference on Computers and accessibility*, pages 91–98. ACM.

[Richard et al., 1997] Richard, C., Okamura, A. M., and Cutkosky, M. R. (1997). Getting a feel for dynamics: Using haptic interface kits for teaching dynamics and controls. In *1997 ASME IMECE 6th Annual Symposium on Haptic Interfaces, Dallas, TX, Nov*, pages 15–21.

[Riener and Harders, 2012] Riener, R. and Harders, M. (2012). *Virtual Reality in Medicine*. Springer.

[Robles-De-La-Torre, 2006] Robles-De-La-Torre, G. (2006). The importance of the sense of touch in virtual and real environments. *Ieee Multimedia*, (3):24–30.

[Rosen et al., 2014] Rosen, A., Eliassi, S., Fors, U., Sallnäs, E.-L., Forsslund, J., Sejersen, R., and Lund, B. (2014). A computerised third molar surgery simulator–results of supervision by different professionals. *European Journal of Dental Education*, 18(2):86–90.

[Rossi et al., 2005] Rossi, M., Tuer, K., and Wang, D. (2005). A new design paradigm for the rapid development of haptic and telehaptic applications. In *Control Applications, 2005. CCA 2005. Proceedings of 2005 IEEE Conference on*, pages 1246–1250. IEEE.

[Ruspini et al., 1997] Ruspini, D. C., Kolarov, K., and Khatib, O. (1997). The haptic display of complex graphical environments. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 345–352. ACM Press/Addison-Wesley Publishing Co.

[Rystedt and Sjöblom, 2012] Rystedt, H. and Sjöblom, B. (2012). Realism, authenticity, and learning in healthcare simulations: rules of relevance and irrelevance as interactive achievements. *Instructional science*, 40(5):785–798.

[Salisbury et al., 2008] Salisbury, C., Salisbury, J., Gillespie, R., and Blevins, N. (2008). A microsurgery-specific haptic device for telerobotic medical treatment. In *ANS Joint Emergency Preparedness and Response & Robotic and Remote Systems, Topical Meeting*, volume 6.

[Salisbury et al., 2004] Salisbury, K., Conti, F., and Barbagli, F. (2004). Haptic rendering: introductory concepts. *IEEE computer graphics and applications*, 24(2):24–32.

[San Diego et al., 2008] San Diego, J. P., Barrow, A., Cox, M., Harwin, W., et al. (2008). Phantom prototype: exploring the potential for learning with multimodal features in dentistry. In *Proceedings of the 10th international conference on Multimodal interfaces*, pages 201–202. ACM.

[Saul et al., 2011] Saul, G., Lau, M., Mitani, J., and Igarashi, T. (2011). Sketchchair: an all-in-one chair design system for end users. In *Proc. TEI*, pages 73–80. ACM.

[Schiemann et al., 1992] Schiemann, T., Bomans, M., Tiede, U., and Hoehne, K. H. (1992). Interactive 3-d segmentation. In *Visualization in Biomedical computing*, pages 376–383. International Society for Optics and Photonics.

[Schneider and MacLean, 2014] Schneider, O. S. and MacLean, K. E. (2014). Improvising design with a haptic instrument. In *Haptics Symposium (HAPTICS), 2014 IEEE*, pages 327–332. IEEE.

[Schön, 1984] Schön, D. A. (1984). The reflective practitioner: How professionals think in action.

[Shahrani, 2006] Shahrani, S. (2006). Educational feature: A history and analysis of level design in 3d computer games pt. 1 and pt. 2. accessed 2016-03-02 from www.gamasutra.com/view/feature/131083/educational_feature_a_history_and_.phpw.

[Shaver and Maclean, 2005] Shaver, M. and Maclean, K. (2005). The twiddler: A haptic teaching tool for low-cost communication and mechanical design.

[Sherman and Craig, 2002] Sherman, W. R. and Craig, A. B. (2002). *Understanding virtual reality: Interface, application, and design*. Elsevier.

[Solsona Belenguer, 2015] Solsona Belenguer, J. (2015). *Engineering through Designerly Conversations with the Digital Material: The Approach, the Tools and the Design Space*. PhD thesis, KTH Royal Institute of Technology.

[Sommerville, 2004] Sommerville, I. (2004). *Software Engineering. International computer science series*. Addison Wesley.

[Srinivasan and Basdogan, 1997] Srinivasan, M. A. and Basdogan, C. (1997). Haptics in virtual environments: Taxonomy, research status, and challenges. *Computers & Graphics*, 21(4):393–404.

[Ståhl, 2014] Ståhl, A. (2014). *Designing for Interactional Empowerment*. PhD thesis, KTH Royal Institute of Technology, Media Technology and Interaction Design. TRITA-CSC-A, ISSN 1653-5723; 2014:20.

[Sundström et al., 2011] Sundström, P., Taylor, A., Grufberg, K., Wirström, N., Solsona Belenguer, J., and Lundén, M. (2011). Inspirational bits: towards a shared understanding of the digital material. In *Proc. CHI*, pages 1561–1570.

[Swindells et al., 2006] Swindells, C., Maksakov, E., MacLean, K. E., and Chung, V. (2006). The role of prototyping tools for haptic behavior design. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2006 14th Symposium on*, pages 161–168. IEEE.

[Tse et al., 2010] Tse, B., Harwin, W., Barrow, A., Quinn, B., San Diego, J., and Cox, M. (2010). Design and development of a haptic dental training system - haptel. In Kappers, editor, *Haptics: Generating and Perceiving Tangible Sensations*, volume 6192 of *LNCS*, pages 101–108. Springer.

[Van der Linde et al., 2002] Van der Linde, R. Q., Lammertse, P., Frederiksen, E., and Ruiter, B. (2002). The hapticmaster, a new high-performance haptic interface. In *Proc. Eurohaptics*, pages 1–5.

[Vaughan, 2011] Vaughan, W. (2011). *Digital Modeling*. New Riders.

[Von Hippel, 2001] Von Hippel, E. (2001). Perspective: User toolkits for innovation. *Journal of product innovation management*, 18(4):247–257.

[Von Hippel, 2003] Von Hippel, E. (2003). *Democratizing innovation*. MIT Press.

[Von Sternberg et al., 2007] Von Sternberg, N., Bartsch, M., Petersik, A., Wiltfang, J., Sibbersen, W., Grindel, T., Tiede, U., Warnke, P., Heiland, M., Russo, P., et al. (2007). Learning by doing virtually. *International journal of oral and maxillofacial surgery*, 36(5):386–390.

[Wang et al., 2014] Wang, D., Xiao, J., and Zhang, Y. (2014). Application: A dental simulator. In *Haptic Rendering for Simulation of Fine Manipulation*, pages 131–160. Springer.

[Wang et al., 2012] Wang, D., Zhang, Y., Hou, J., Wang, Y., Lv, P., Chen, Y., and Zhao, H. (2012). idental: a haptic-based dental simulator and its preliminary user evaluation. *Haptics, IEEE Transactions on*, 5(4):332–343.

[Wang and Kaufman, 1995] Wang, S. W. and Kaufman, A. E. (1995). Volume sculpting. In *Proceedings of the 1995 symposium on Interactive 3D graphics*, pages 151–ff. ACM.

[Wijewickrema et al., 2013] Wijewickrema, S., Ioannou, I., and Kennedy, G. (2013). Adaptation of marching cubes for the simulation of material removal from segmented volume data. In *Computer-Based Medical Systems (CBMS), 2013 IEEE 26th International Symposium on*, pages 29–34. IEEE.

[Winograd et al., 1996] Winograd, T., Bennett, J., De Young, L., and Hartfield, B. (1996). *Bringing design to software*. ACM Press New York.

[Wright, 2011] Wright, A. (2011). The touchy subject of haptics. *Communications of the ACM*, 54:20–22.

[Yushkevich et al., 2006] Yushkevich, P. A., Piven, J., Cody Hazlett, H., Gimpel Smith, R., Ho, S., Gee, J. C., and Gerig, G. (2006). User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128.

[Zilles and Salisbury, 1995] Zilles, C. B. and Salisbury, J. K. (1995). A constraint-based god-object method for haptic display. In *Intelligent Robots and Systems 95.'Human Robot Interaction and Cooperative Robots', Proceedings. 1995 IEEE/RSJ International Conference on*, volume 3, pages 146–151. IEEE.

[Zimmerman et al., 2007] Zimmerman, J., Forlizzi, J., and Evenson, S. (2007). Research through design as a method for interaction design research in hci. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 493–502. ACM.

Paper A

Designing the Kobra Oral Surgery Simulator Using a Practice-Based Understanding of Educational Contexts

Author 1, Author 2, and Author 3

Abstract—Surgery simulation is a core application area of computer haptics and simulation technologies, giving aspiring surgeons the opportunity to practice hands-on using complex manual actions before encountering real patients. The design of the haptic feedback is an important aspect of developing such tools, but the design of a surgery simulator involves also many other aspects. This paper presents a long-term case of designing and iteratively developing an oral surgery simulator named Kobra. Based on feedback from surgeons, students and curriculum developers, as well as through insights from actual design work, particular aspects of the design that support learning have been identified and articulated. Based on experience of designing simulator exercises originating from authentic patient cases it is shown how simulation techniques can be appropriated to support oral surgery teaching, through targeted interaction design. The study highlights important aspects to consider for further design work in this domain, i.e. the value of realism and surgical relevance, the social setting of teaching surgery, content authenticity, and the physical qualities of the simulator.

Index Terms—Surgery Simulation, Oral Surgery, Interaction Design, Haptics

✦

1 Introduction

COMPUTER-BASED DENTAL SIMULATION with haptic feedback is an active research area spanning multiple disciplines [1], [2], [3]. Wang et al recently published an extensive survey of twelve dental simulators concluding that there are still much room for improvements and open research questions, for instance which degree of realism is required to provide effective training [4]. This paper presents the design rationale and the findings of the *research-through-design* that have been conducted during the design and development of one of the simulators identified in Wang et al's survey; the Kobra Oral Surgery simulator (Fig. 1). Research-through-design is a research approach in contemporary Human-Computer Interaction that allows interaction designers to contribute knowledge gained during a project in a structured way [5]. This way, open-ended and underconstrained problems such as how to design a useful dental simulator, can be addressed, without being restricted to evaluating already developed systems, or being confined to incremental improvements to the systems technical components. Furthermore, this research stance usually emphasizes the importance of technology being grounded in real practice, which adds important complexities for the designers and engineers involved in developing these advanced technology-intense systems. Importantly, grounding designs in real practice may emphasize other qualities or design challenges than brought to the fore in the more technology-oriented research areas. The increased focus on studying and supporting real user practices has for instance emphasized aspects of shared and social activity around computer systems and the tendency that different people may interact with and make use of the same system in different ways [6], [7].

Fig. 1. Illustration of instructor and learner collaborating in solving a patient case in the Kobra Oral Surgery Simulator.

Dentistry and related fields have received particular attention in the technical domain of surgery simulation, partly because the professional practice is tool mediated and mainly deals with drilling or carving in hard tissue, which means it can use existing

---

- *Author 1, Author 2 and Author 3 is with the department of X at the University Y. City, Country. E-mail: xx; yy; zz @ university*

Manuscript submitted February 28, 2016;

haptic technology without requirement of advanced soft-tissue deformation simulations. Several algorithms for haptic and visual rendering of bone erosion have been developed [8], [9], [10], [11]. Together with the continuous commoditization of high-performance computing hardware, haptic devices, stereoscopic displays and software frameworks, these technological advancements have enabled design of simulators for training of surgical procedures in an immersive, hands-on manner. However, there is much more to simulator design than technological improvements of its components. In order to be a valuable resource for learning, surgery simulators need to be designed to support a particular educational context. While several promising simulators for dental or oral surgery procedure have been developed [12], [13], [14], [15], it remains open what aspects of them are the most important in such situations and how the haptic technology can best support the learning activities.

In this paper we contribute to this on-going discussion on what constitutes a useful simulator through an analysis of the design, development and user feedback of the oral surgery simulator named *Kobra*. The work has been involved an interplay between detail and wholeness, materials and texture [16], where trade-offs have had to be made between realism and implementation effort among many other aspects including aesthetics and creation of teaching content.

Through numerous encounters with users, potential customers, students and teachers a design has evolved over time to a state that it is now in permanent use at a dental school. Aspects of the simulator that seemed more important at each development stage have been prioritized and developed further while other aspects have been toned down. Based on these experiences we argue that when designing a surgery simulator it is important to consider: 1) the relevance of realism, 2) the social setting of surgery teaching, 3) the authenticity of simulation content, 4) aesthetics of graphics and haptics, and 5) the offline qualities of the physical design.

2 BACKGROUND

To arrive at a *practice-based* perspective of dental simulation it is useful to review both what simulators have been produced and how they have been actually *used* in dental education practice. This section will cover some of the simulators developed in academic and corporate settings, the technology common to several of them, and how they have been designed and used. For a more extensive survey of simulators and their technical features see Wang et al [4].

Most dental schools rely on training on passive mechanical simulators, where students can practice drilling with real tools on disposable synthetic teeth. The DentSim simulator (Image Navigation, New York, USA) extend such a mechanical simulator with optical tracking of instruments, whereby a virtual representation can be displayed on a monitor [3]. Drilling in the real synthetic teeth corresponds to virtual teeth being carved in synchronization. This enables real-time quantitative measures of student performance and instructional feedback during the tooth preparation, which have been shown to reduce the required human instruction time by a factor of five [17]. Haptic feedback is provided mechanically “for free”. The downside is that it requires disposables, the mannequin lacks bone for surgery, the material is far from the hardness of real teeth and patient variety is limited to available range of synthetic teeth. Furthermore, the physical drilling requires access to water and sewage and pose some safety concerns. Other hands-on training involve animal cadavers which have vastly different anatomy, or deceased humans which leads to ethical concerns. The prevailing form of hands-on training is thus on humans in an apprentice setting, which pose a patient risk and is limited to the range of patients entering the clinic where the apprenticeship takes place.

Besides hands-on training, screen-based clinical reasoning simulations that rely on symbolic actions and treatment planning of virtual patients, have been developed and combined in training sessions with haptics-based simulations [18], [19]. Symbolic simulations, compared with hands-on simulations, rely on resemblances of the clinical situations rather than realism [20] which can be very effective, but provide little or no hands-on training. Computer haptics-based surgery simulators can however be designed to support learning in both clinical reasoning and hands-on practice at the same time.

2.1 Enabling Technology

The technical advancement of computer haptics has made it possible to design computer-based simulators that eliminates the need for disposable teeth and replaces the surgical instrument with a force-reflecting haptic device and the plastic teeth with a virtual environment, usually rendered visually to a stereoscopic display. The central interaction in dental and oral surgery is probing and carving of hard tissue of varying density. Simulating this requires an anatomy model, a model of the drill, visual and haptic rendering algorithms and a method for material erosion. Agus et al [9] developed a penalty based haptic rendering method where the bone was represented with a volumetric rectilinear sample grid (voxels), treating each sample as a sphere and thereby computing the interaction volume between bone and a spherical drill as series of sphere-sphere intersection calculations. The reflecting force is then proportional to the distance to the penetration depth of the summarized intersecting volume. Carving was implemented through a method where work on each voxel was registered in a counter until it was fully removed. Petersik et al [8] developed a multi-point proxy-based haptic rendering method were a set of collision detection points are placed on the surface of the drill and used to compute interaction forces with an implicit surface, allowing for sub-voxel resolution. They also introduced improved methods for visualization and carving [21].

Other previous work on haptic rendering and bone dissection is Morris et al [22]. Related work on 6-dof haptic rendering and patient-specific surgical rehearsal [10], [23] has pushed the boundaries for surgery simulator design. All work mentioned so far have both technical (algorithmic) contributions and design contributions in terms of how the technology is adapted and applied towards a particular surgical area. We will now turn our attention to the design of simulators that utilizes these or similar technologies.

2.2 Simulator Design

The design of the simulator named VOXEL-MAN, has been adapted to several surgical areas including the oral surgery procedure of Apicectomy [24]. These simulators originates from the group’s pioneering work on interactive anatomical atlases [25], thereby extending the naturalistic surgeons-view with visualization techniques such as cut-planes, colouring of important tissues and label displays [21]. The design of the Apicectomy simulator features a graphical user interface with windows showing both

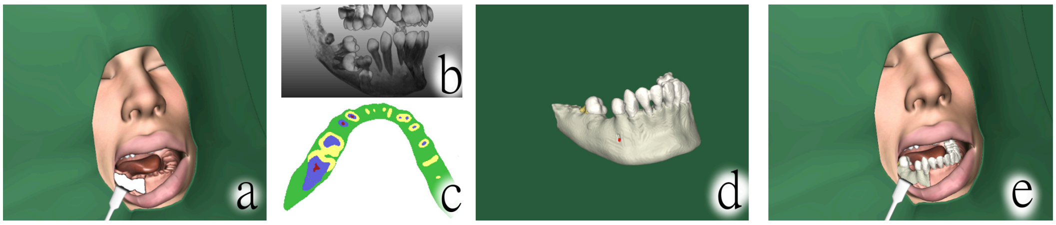

Fig. 2. Visual and haptic model of a impacted tooth extraction case. a) non-interactive polygon mesh, b) original CT-image, c) slice of the remodeled jaw fragment, showing color-coded tissue segments used for haptic rendering of different hardness and color, d) jaw fragment as viewed in simulator, e) standard view in simulator.

the surgical view and computed tomography (x-ray) image planes which helps the student map between the radiology data and the reconstructed surgical view. Three different training levels were provided, where the basic level showed transparent bone and highlighted artificial pathologies and nerves as a form of assistance that was enabled or disabled depending on training level. The VOXEL-MAN simulators have been presented with stereoscopic rendering both in a desktop setup and in a mirrored manner where haptic and visual displays are co-located and embedded in a stylish enclosure [26]. The Simodont Dental Trainer (Moog, Amsterdam, Netherlands) is a simulator explicitly designed to replace conventional mechanical dental simulators, and its focus is thus on fundamental tasks such as preparing teeth for fillings. The system features a 6-dof sensing authentic looking dental handpiece with admittance controlled 3-dof force feedback [27]. A projector-based 3D display gives a high resolution, co-located stereoscopic view to one user wearing polarizing glasses. The whole system is embedded in a styled red-and-white hight-adjustable enclosure that also features a touchscreen as a system interface, an additional 6-dof relative positioning input device, foot pedals and hand rest [28]. The haptic environment is complemented by an educational content system [14]. VirTeaSy is a simulator developed for dental implants. It has a windows-based planning phase where cases are presented with x-ray slices, and a haptic-enabled surgical phase where the procedure is carried out. The system includes a tracked head-mounted display, a Virtuoso 6D haptic device, and a clinical case database [29]. VirTeaSy optionally shows a cross where students should drill, recommended angle, depth and warning colors for overheated drilling. The students drilling is recorded and the result can be viewed in the scan mode. The teacher has an interface where she can view and interact, eg. zoom, on a separate screen. Students can go back and forth between the two phases, including viewing surgery results in the planning view retrospectively. The student is able to distinguish between bone of four different densities using the sense of touch. Wang et al [13], [30] have designed and developed a simulator for non-drilling dental procedures such as pocket probing which is performed with a thin, cylindrical instrument with millimeter marks, that is inserted in the pocket between a tooth and the gingiva. Tse et al [12] developed hapTEL, a simulator for pre-clinical training, based on user requirements found with an earlier prototype [31]. This simulator had to be cost-effective in order to allow for large-scale educational evaluation, which required a series of units to be produced. Therefore compromises between quality and number of systems had to be made from the start. An evaluation of several devices; Phantom Omni, Phantom Desktop, Phantom Premium 1.5, Falcon, Omega 3, led to the decision to produce twelve simulators based on a modified Falcon and two simulators based on Omega. Falcon, lacking orientation sensing, was modified to hold a real dental hand-piece through magnetic coupling, and tracked by a linked arm at the rear of the hand-piece where the cord of a real hand-piece goes.

Previous work also include explicit requirements gathering, for example Ioannou et al [32] have studied the haptic and visual cues used by experienced dentists and measured the forces they used.

2.3 Use of Simulators in Surgery Training

Dental and surgery simulators have been studied from the perspectives of training efficacy and how they are used in real life. The Simodont has been subject to studies that show that its efficacy of training is comparable to mechanical simulators while reducing the time needed for supervision [28], although faculty impressions are that it would not completely replace the teachers [14]. Experimental studies of the learning contribution of the VirTeaSy simulator in a standardized plaster drilling task showed that students became better over time with increasing number of training sessions with the simulator [15]. Evaluations of real use of the hapTEL Falcon-based simulator showed among other things that the internal friction and mass of the haptic device was considered too high - some students even used two hands to operate it. Conclusions for future work was more sophisticated haptic rendering algorithms and rubber cheeks to limit range of motions closer to that of reality [12]. Parallel to these studies there have been research in *social and technology studies*, that investigate the actual use of simulators, e.g. the dialogue between student and teacher [33], how the student is transformed into a professional surgeon [34] and how simulations are real - as in real, worthwhile training - even if the simulator equipment used bear only a few resemblances of the situation the student is training for [20].

Effective use of simulators for training requires not only a good simulator, but a well thought-through practice led by an experienced medical professional who can turn the simulator practice into medical practice. This implies on one hand that we should be aware of that simulators only cover a few elements of medical practice, but on the other hand that the teachers are capable and willing to complement and appropriate the simulators for a relevant learning experience. Medical professionals who plan simulator training sessions can direct the learning opportunities through integrating the simulator in a larger medical context and through reconstitution of patient bodies [35].

3 THE ORAL SURGERY SIMULATOR, KOBRA

The purpose of the Kobra simulator is to *support teaching* of oral surgery procedures, in first hand to students in general dentistry. A common procedure is surgical extraction of third molars where drilling into the jawbone is necessary to access impacted teeth. This kind of procedures are among the most advanced a general practicing dentist performs, yet not advanced enough for being remitted to specialists. Students therefore request to practice these procedures during their education, but are rarely allowed due to costs and patient risks.

Fig. 3. Passive silicon and plaster mannequin and hand tool supports the activity in the simulator.

The Kobra simulator has a passive silicone mannequin with an empty mouth that is slightly enlarged to accommodate the manipulandum joint of the haptic device (Fig. 3). The role of the mannequin is to familiarize the operator and to provide the important hand support. Through the mirrored display an image of teeth is ideally projected inside the mannequins mouth, but sometimes in the vicinity in which case the role of the mannequin is solely hand support. The visualized teeth can be felt using the manipulandum of the haptic device, that gives a directional force feedback while tracing the surface of teeth and bone (Fig. 2e). If the user presses a foot pedal the drill starts and varying resistance can be felt, as material is removed from anatomical structures such as bone, enamel and dentin. A side pad computer acts as the systems user interface for all textual information, keeping the 3D environment clutter free.

The simulator is designed to run in an embedded fashion which imply, from a user perspective, that the simulation computer behaves like a computer in an integrated product rather than as a host of a desktop software application. Practically, it automatically boots into a full-screen application after the user presses the single start button on the front. The user can then log in to his or her personal account as prompted on side screen (Fig. 4a). The user can choose between several *patient cases* each with its own x-ray image, amanesis (problem description), and treatment plan (Fig. 4b). Selecting one case gives the option of starting the simulation or play back of previously recorded sessions. The patient case is downloaded from the Internet and the patient is presented in stereoscopic 3D, draped in surgical cloth and with any required surgical flap already prepared for, so the user can begin at the bone or tooth dissection step of the surgery. The user is presented with a virtual drill (Fig. 5c) that can be switch to a secondary instrument (Fig. 5a or 5b) by pressing the right foot pedal. During simulation the amount of material removed is displayed on the side screen (Fig. 4c), grouped by different tissues and areas that should be avoided (neighboring teeth and nerves). The user has also the option of hiding the face model and rotating the jawbone fragment (Fig. 2d) to view the surgical area from different angles.

A key aspect of the simulator is the different exercises (patient cases) that have been designed based on computed tomography scans of live patients. A surgical procedure is encoded in each case with a state machine (Fig. 7) where the user has to remove a certain amount of material in a specific area in order to progress to the next state. The progression is triggered by probing or applying a small force with a secondary instrument within proximity of a point defined by the designer of respective case. The progression might involve automatic removal of a segment, defined by a separate image mask. This allows for e.g. removal of crown and root if the user has burred enough to crack and separate them with the elevator (Fig. 5 top). All instruments use the same 3-dof haptic rendering algorithm where the tip of the instrument is modeled as equidistant collision points [10] of a sphere with a diameter of roughly 2 mm. Together this allows for conducting a multi-step surgical procedure (Fig. 6).

3.1 System Architecture

Kobra consist of two computers. The main computer located at the base of the simulator is a desktop personal computer (Intel Core i5 2.4 Ghz) with a Quad Buffered Stereo enabled graphical processing unit (Nvidia Quadro 4000) running Ubuntu 12.04 32 bit[1] operating system. The side computer is a touch pad (ASUS Transformer Pad 10.1) running the default Android operating system and web browser, displaying the HTML/JavaScript based custom made graphical user interface. The side computer communicates with the main computer via a dedicated wireless router. The main computer runs a custom made launching software written in C++ using Qt libraries that responds to the pad computers HTTP requests and downloads, starts or stops the simulation application. The simulation, which is built using the H3D API extended with the custom made forssim library, loads an XML file that defines the scene objects and state machine. The simulation application also responds to HTTP requests with a data dictionary of amount of material removed etc for display on the pad. Finally a web application was developed to host the cases, the pad front-end interface and any recorded sessions.

3.2 Visual Rendering

As can be seen in figure 1, the simulator has a monitor oriented 45 degrees toward the user. The user looks at it through a horizontal mirror, causing a virtual image plane appear 45 degree down from the mirror i.e. in a diagonal plane from upper back to lower front of the space under the mirror. The mannequin's mouth intersects this plane. Together with stereoscoping rendering and shutter glasses this enable projection of the virtual image of head and dental drill co-located with the mannequin and manipulandum. In this way both static models (artist-made textured polygon meshes surrounding the operating area) and dynamic models (those that can be carved) are presented in the same scene. The virtual environment is defined in metric cartesian coordinates originating from the lower back left corner of the space under the mirror.

[1] due to compatibility with the calibration routines of the haptic device drivers

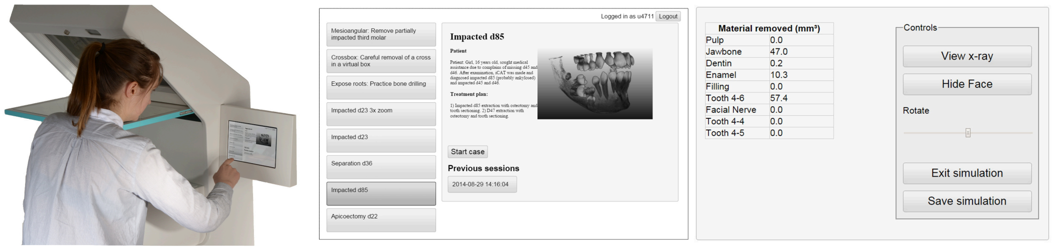

Impacted d85

Patient

Patient: Girl, 16 years old, sought medical assistance due to complains of missing d45 and d46. After examination, iCAT was made and diagnosed impacted d85 (probably ankylosed) and impacted d45 and d46.

Treatment plan:

1) Impacted d85 extraction with osteotomy and tooth sectioning. 2) D47 extraction with osteotomy and tooth sectioning.

Previous sessions

2014-08-29 14:16:04

Material removed (mm³)	
Pulp	0.0
Jawbone	47.0
Dentin	0.2
Enamel	10.3
Filling	0.0
Tooth 4-6	57.4
Facial Nerve	0.0
Tooth 4-4	0.0
Tooth 4-5	0.0

Fig. 4. The Graphical User Interface. a) interaction with the tablet computer b) case selection, c) display of how much volume is removed from each tissue layer and execution controls.

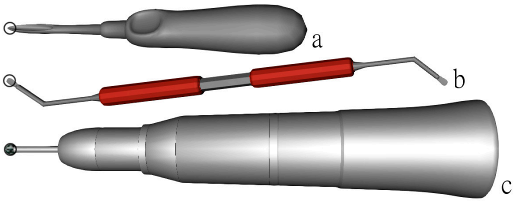

Fig. 5. Simulated dental instruments; a) elevator used for cracking teeth and prying roots, b) excavator used to remove infected tissue, c) surgical drill. For collision detection and haptic rendering, each instrument is modeled as a sphere located at the tip of respective instrument, illustrated with a circle. (a) and (b) are secondary instrument.

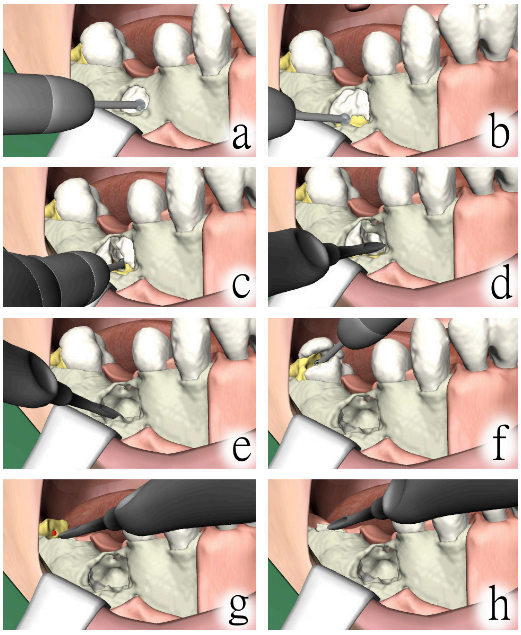

Fig. 6. Cropped screen-shot from the simulation. This scenario gestalts surgical extraction of one impacted (a-e) and one third molar tooth (f-h), and is based on a real patient case. The proximity to another impacted tooth and the mandibular nerve adds challenge to the case. The user begins with dissecting bone without carving into the impacted teeth to get sufficient visibility and access (a). Then, the user alter between sectioning the teeth with surgical drill (b,c,f) and cracking/extracting them with the elevator (d,e,g,h).

This means that a designer can physically measure the position to virtual objects placed in this space, and the objects will appear where expected. However, considerations need to be taken since the further an object is placed from the virtual diagonal image plane, the larger separation of left and right image which can be uncomfortable for the user to fuse [36].

3.3 Haptic Interaction

The simulator utilizes a Phantom Desktop haptic device, which gives input in position plus orientation and can render directional force feedback. As the user touches the segmented bone model with the drill avatar, a force is computed using Chan's algorithm [10]. In short, this algorithm models a number of contact points on the surface of the spherical drill tip, and upon collision between these points and the virtual tissue, it constrain the avatar (visual representation) to move on the surface, and calculates a repelling force proportional to the now penetrating manipulandum, towards the avatar. This force is rendered to the haptic device as long as contact remains, and increases in magnitude with penetration depth and a stiffness factor k.

While drilling, the amount of tissue material removed is proportional to time in contact, on a per voxel basis, and a factor set by the designer for each tissue, meaning that enamel takes longer time to drill than bone using the same force. Both the stiffness k and tissue "hardness" constants are tuned with the help of surgeons.

4 DESIGNING THE KOBRA SIMULATOR

In this section the activities that have influenced the design of the Kobra simulator are described. The simulator development has gone trough six iterations spanning the years 2007 through 2014. Each iteration has resulted in a prototype that has been subject to formal or informal evaluations, which in turn have impacted the work in the following iteration. The activities can be structured around 1) probing the target audience with prototypes and field studies, 2) implementing enabling technology motivated by design needs, 3) evolve a physical design with the help of professional designers and carpenters, 4) make tools for tuning and scaffolding workflow 5) design patient cases with help of professional 3D artists.

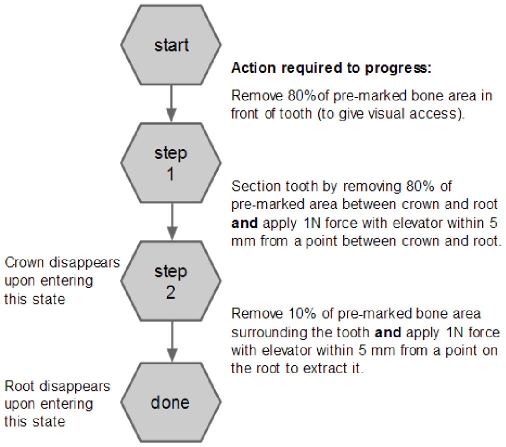

Fig. 7. Scenario propagation example for a certain tooth extraction.

4.1 Encounters With Target User Group

The starting point for the work was a pre-study with two parallel tracks; field-studies and technical investigations. The field-studies included observations of live surgery with follow-up interviews with the teaching surgeons, sitting in on tooth extraction lectures, studying domain literature [37] and experimenting by drilling pig jaw cadaver, which is an established teaching method. The pre-study resulted in an analysis and synthesis of design recommendations. A typical patient case was decided on, based on what was the least surgically advanced procedure among those cases that still required bone dissection. This is the mesioangular (forward-tilted) partly impacted third molar. Dental literature [37] and interviews show that educators prefer to discuss a surgical procedure in terms of discrete steps and it is important that the students learn which these are and how they should perform them. Out of these steps, the one prioritized for simulation was bone drilling to get sufficient access to the impacted tooth in order to be ready for sectioning and subsequent extraction in parts. It was noted that one thing that makes the operation difficult is that the operator naturally have limited view and movement freedom because the tooth is located in the very rear of the mouth, and that the cheek should not be touched at all by the drill in order to avoid patient harm. This was manifested in the design by modeling a non-interactive polygon mesh of the face, combined with a cropped voxel volume derived from a computed tomography scan of a patient provided by the dental school.

The first prototype was evaluated using the cooperative evaluation method [38] with four experienced surgeons. The task was to remove sufficient bone to be able to begin tooth sectioning. Among the outcomes was the need for hand support, correction of the view angle since the head was placed face forward but the operation is usually performed from the side, and the unwanted side effect of the noise from the haptic device was worrying the operator since low frequency noise sounds like the drill bit is about to break. Technical improvements and refined design was then intermixed with cooperative evaluations in the subsequent iterations.

Throughout the project five cooperative evaluation sessions with senior dentists and one with dental students have been conducted at different development stages. One version of the prototype has also been subject to a study on the student acceptance rate of simulator practice as part of the curricula with two sessions totaling 2x30 students. This study showed that student acceptance of the simulator was very high [19]. One experimental study was conducted comparing different instructors (no feedback, another student, a technician and a surgeon) with eight students per group. This showed that the best results were achieved with a surgeon as instructor [39]. In addition there have been at least three private demonstrations to faculty members of various dental schools, and demonstrations at public exhibitions at three occurrences of the annual meeting of Association of Dental Education in Europe and at one occurrence of the Annual Swedish Dental Fair.

4.2 Implementation of Enabling Technology

The aim of the initial technical investigation was to identify and become acquainted with resource readily available to the design team such as co-located stereoscopic display, haptic devices (Phantom Desktop, Phantom Omni), software libraries (H3D API) and estimation of the implementation effort of various algorithms for haptic rendering and material erosion. Candidate technologies have been tested and appropriated using a bricolage approach [40].

Since the premier focus was on interaction design it was not feasible to engage in extensive algorithmic development why a relatively straight-forward method inspired by Agus et al. [9] was selected and implemented in a dynamic library extension to H3D named *forssim* [41]. Forssim, as an open source project, has since evolved to support a proxy-based haptic algorithm based on Chan et al. [10], a multi-surface marching cubes-based visual rendering method [11], a state machine mechanism that let users complete a procedure step by step, and methods to record and play-back procedures.

4.3 Evolution of a Physical Design

The initial setup consisted completely of off-the-shelf components, including a co-located display (Fig. 8a). The aim was to make a coherent looking product, even though mostly off-the-shelf components were used. To achieve that goal, prototype enclosures were made in wood that embedded the computer, 3D monitor, haptic device, pedals and mannequin (Fig. 8b). Based on feedback from some dentistry faculty members that a home-made looking wooden box was not suitable for incorporation in a well respected university, the form of the simulator had to be improved. Professional furniture designers and carpenters were contracted to polish the prototype (Fig. 8c). As flat 3D monitors and pad computers became affordable, the design evolved (Fig. 8d), and updated to the final design (Fig. 8e) to encompass new monitor dimensions and minor tweaks.

4.4 Formulation of a Workflow and Tools

In order to work effectively with authoring content for the simulator, i.e. creating several patient cases, a workflow had to be formulated. The interactive solid model (Fig. 2d) was designed in several steps. First, a computed tomography (CT) scan of a patient (Fig. 2b) was imported in an image processing software for rough segmentation and conversion into polygon representations of jaw, enamel, dentin and pulp. These was subsequently decimated, and imported into a polygon sculpting software where the 3D

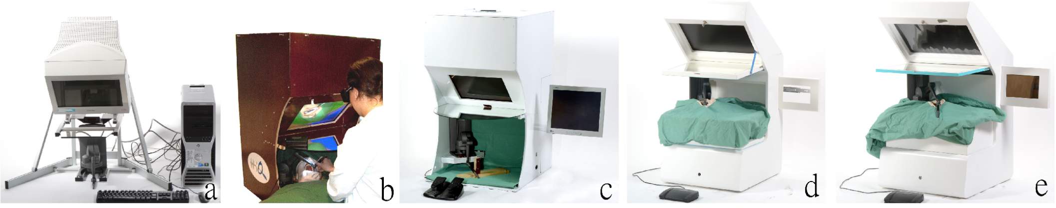

Fig. 8. Evolution of the physical design, from collection of stand-alone components (a), to home-made looking enclosure (b-c) to final design (d-e).

Bone cutting rate: 0.55
Tooth cutting rate: 0.06
Bone avg. cut force: 0.50
Tooth avg. cut force: 0.50

Drill diameter: 3.0 mm
Scale: 2.90
PID: 1.00

Applied force: 0.16
Frame rate: 188

Fig. 9. Tool for interactive tuning visual and haptic rendering properties

artist could address scan imperfections by removing flashes from fillings, filling hollow areas, and separating coalesced teeth. Models of the different tissue layers were then sampled (voxelized) and merged to one multi-layered model. This model was finally inspected and edited slice-by-slice (Fig. 2c). Several of these steps required creation of scripts for converting, merging and splitting files.

The non-interactive polygon face mesh (Fig. 2a) was designed in a professional polygon modeling software, using high-resolution orthographic stock photos as templates and for texturing. This model was then modified to fit together with the polygon representation of the jaw described above. These models were placed in a virtual scene defined in an XML document, together with visual and haptic properties, tool models (Fig. 5) and scenario logic. For this step, an interactive tool (Fig. 9) was created. The tool allowed for real-time tuning of carving rate, drill size and color attributes [42].

4.5 Design of Patient Cases

Having a workflow set up with corresponding tools enabled a practice of creating multiple interactive patient-cases for the simulator. A team consisting of a professional 3D artist, an interaction designer, and an oral surgeon gestalted four typical oral surgery procedures, based on interpretations of real clinical cases. In dialogue with a client dental school, patient CT-scans were acquired and interpreted, models created and areas to avoid and steps to take in the procedure was encoded following the workflow described above. A typical case is depicted in figure 6. Each case could also include models of its own instruments (Fig. 5). High-level programming of these tools was possible, such as removing infected tissues with the excavator (Fig. 5 middle), using the same carving algorithm as for the drill, but only allowing it to carve voxels belonging to the infected tissue segment. In summary, what a patient case constitutes is a set of instruments, an anatomical model and a procedure.

5 RESULTS: ANALYSIS OF DESIGN PROCESS

In this section we present an analysis of the cohesive design, development and evaluations. An extensive amount of data was gathered consisting of a chronological description of the development history, research papers, video tapes from user studies, marketing materials and e-mail discussions.

Particularly interesting themes were identified and related to work by others and theory. These cover realism and surgical relevance, the social setting of surgery teaching, content authenticity, visual and haptic appearance of on-screen representations, and the qualities of the physical design.

To guide our analysis and grounding it in a perspective of practice, we have made use of the four shifts in ideals that such as approach entails, according to Fernaeus *et al* [7]. The first theme concerns a shift from an "information-centric" to an "action-centric" perspective on interaction, and how that affected the understanding of our design in terms of use qualities striven for. The second shift concerns a broadened focus from studying properties of the system, to instead aim at supporting qualities of the activity of using a system, in our case teacher-led education at a dental school. The third theme concerns a shift towards supporting shareable use, rather than primarily individual use settings. The last theme concerns the shift towards multiple and subjective interpretation of how to use new technological artifacts. We discuss how these themes are reflected in our design, and its implications for the design of other applications specific to the domain of dental education.

5.1 Realism and Surgical Relevance

A simulator needs to accurately represent relevant aspects of a surgery situation, but not necessarily by photo-realistic graphics and physics. Already early versions of the Kobra simulator proved meaningful in the hands of a teacher instructing a student. The teacher covered up for faults in the simulator e.g. a cropped tooth was called a non-issue and was instead labeled as part of the patients history. Rystedt & Sjöblom's discussion [20] on rules of relevance and irrelevance is applicable to the Kobra simulator. They are studying a simulation of anestesiology in the form of a 2D desktop software symbolically showing the status of the patient and numbers from various measurements. Despite being purely symbolic it was relevant to the education of nurses and deemed realistic for that purpose.

In Kobra, one design decision was to visually display a face around the area subject to surgery. This indirectly constrains the manipulation of the virtual drill and illustrates that the student cannot drill from "any angle they so wish", but have to plan their surgery from the context of cheek and tongue that might restrict their view and be a physical obstacle. This face was however never simulated haptically or collisions detected in any other way (e.g. for assessment), but was still respected by students. Completely ignoring the effort of rendering this soft tissue - which is known to require significant implementation effort - saved significant development time without, in this context, necessarily be a problem.

One assumption was that simulated teeth should be rendered in natural size visually and haptically. The design explorations in the Kobra project show that scale of objects together with haptic hardware qualities affect the user experience. Alternative designs could have lead to benefits, such as use of cheaper haptic hardware [42]. However, changing scale of the haptic representation raise new questions of how it should be displayed visually as well as size of the mannequin if co-location is to be preserved, why the original conceptual design, i.e. co-location and high-quality haptic device, was left intact. The fact that dental instruments to their nature and use in dentistry are inherently six degree of freedom, was assumed to also require six degree of feedback simulation and haptic feedback which is challenging to implement. The idea was that if the system did not provide this, the tool should not be simulated at all. The use of the elevator to, by rotation in a tooth incision, cracking and separating the crown from the root, would then be expected to compute advanced physics collision response and torques. However a much simplified simulation solution proved to be well appreciated. What are discussed here are different types of realism. In addition to natural realism, we can talk about situation realism in how the simulator or some elements of the simulator can be designed to be relevant to the context of surgery teaching. Authentic patient cases that the teacher has actually had own experience of operating on were very well received. The simulation is then just one part of a larger story of how this case was treated. The case is then very real in terms of it being a real human being operated upon in contrast to a random artificial tooth without further context anchoring.

Some features of the Kobra simulator were included because they improved the instructor's ability to discuss something important with the student, while being radically unrealistic. One example is the ability to hide the face and only show the jaw (Fig. 2d) and rotate the jaw to see the result of the carved area from the tongue side. This feature was extensively used by the teachers to show the danger of drilling too deep with risk of popping through and hurting the tongue. There are still minimum levels of realism required. Fall-trough was in early version experienced by some surgeons that applied high forces to the device, where high is larger than half the spherical drill penetrating the bone surface. Fall-through breaks the illusion of touching a solid object [43]. This issue was eliminated when switching to a constraint-based algorithm [10]. However, other issues emerged as the avatar sometimes could get stuck in some parts of the model.

This observation shines new light on the question of what part of the surgery procedure is most suitable for simulation, both from a medical education perspective, what is new to the students and what do they already know, and technical perspective, what is technically feasible to simulate with sufficient fidelity? This requires investigations both into the procedure of the surgery task at hand and what the possibilities are of the candidate technology, which affects the whole design process.

Tse et al [12] discuss the importance of different colors and the learning potential of e.g. transparency to show structures in a way unavailable in traditional settings. These are "unrealistic" features that have been designed into the VOXELMAN simulators, and clearly support training of cognitive rather than motor skills. Case based learning, supported by creative visualizations and interaction design thereby seem to be a recurring theme of useful simulator design and worth of further investigation.

Wang et al adopt the concept of construct validity [44] to their simulator, in which the simulator is evaluated in terms of its ability to reflect actual skill levels, i.e. if the simulator can distinguish between novices and experts, it is simulating something clinically relevant accurately. In addition to this evaluation, they used a questionnaire for feedback on the perceived level of realism of the simulator. Based on these studies they conclude that the simulator has two design limitations. First, the cheek occludes correct reading of the probing tool which is met by suggesting that "Tongue and cheek should be deformable bodies and a mirror should be used to deform" and that a 6-DoF haptic algorithm and haptic device with torque feedback should be used to avoid penetration of the probing instrument. Our work suggests that there could be other fruitful solutions that do not depend on availability of such technologies.

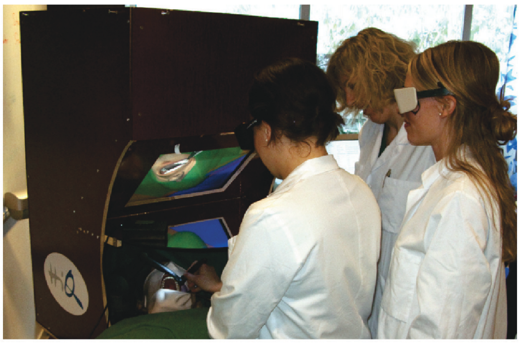

Fig. 10. Two dental students collaborate under supervision of a surgeon-teacher.

5.2 Social Setting of Surgery Teaching

Observations showed that teachers and students collaborated and discussed surgery topics beyond the specific simulation task (Fig. 10). They used the simulator and the simulated patient as a focal point for those discussions, for example how to work safely. In that way the simulation supported dialogue, it became an illustration that helps to form discussion of a larger topic. The form of the simulator supported a large enough display to enable good stereoscopic view for two simultaneous users, something that was often used by the student-teacher pair. The "cheap" elements incorporated in the simulator (cloth, wound hook) came from a suggestion by a surgeon, with the idea of familiarizing the surgeon.

An initial assumption was that the simulator should primarily be used for self-study by a single student at a time. Therefore it was assumed that it was important that the system gave the user feedback, recorded sessions for later teacher assessment and logged quantitative results giving a quantitative assessment of skill. However, generating quantitative reports on amount of bone and teeth removed and in what areas, rarely caught the interest of users. When asked if the simulator has a learning guide the response that it did not but was designed to be used under supervision was accepted. Over time it became clear that actually involving instead of replacing the teachers was much more fruitful at least in these still early generations of surgery simulators.

Studies by Johnson et al [45] of practice with a minimally invasive surgery simulator revealed that instructors were essential in order to *reconstitute* medical practice out of the simulation practice. Medical practice is what a surgeon does in an operating room with a unique human patient, including all factors that make the operator a professional surgeon working in a professional environment. Simulator practice without a professional surgeon as instructor will often be limited to technical exercises. Johnson observed that the surgeons related elements of the simulator to relevant and real medical practice, through telling anecdotes and more. This is important because surgery is not only about technical skill, it is as much about good surgical judgment and ability to make quick decisions in demanding situations.

Reconstitution is however even more powerful beyond bringing general professionalism to the simulation training session. Through their instructions, surgeon-teachers help reconstitute a complete patient out of the limited local part that is simulated in the machine. The minimally invasive surgery simulator Johnson observed did not have a physical representation of a patients knee, which was the subject of operation, and students had difficulties in orienting the camera and tools accordingly. Through pointing to the surgeons own knee in the position of the patient (temporarily acting, or reconstituting, the patient) in relation to the simulator, and talking in terms of both simulator use and surgery practice, the teacher could align students perspectives to match the reconstituted surgical practice [45]. The teaching surgeon has a great effect on making the simulation practice better. In the same sense it is important to design the simulator training session. This is not necessarily done by the same people that design the simulator, but nevertheless it will affect how meaningful the training activity will be. The simulator is intended to help learning within a specific cultural environment, it is highly situated [34]. When a student missed the briefing part by the senior surgeon prior to the simulation training that student was observed to be completely off compared to the other students, according to the teacher.

5.3 Content Authenticity

The first, and for a long time the only, patient case created for the simulator was relatively generic, a typical partly impacted third molar. While the case was based on a CT scan of a patient, it was designed without any particular patient in mind. The two most recent cases made were however directly designed from authentic patients who had undergone surgery at the dental school which uses the simulator. Rather than being a generic training task designed by the simulator developers, these cases were given by the simulator users (surgeon-teachers) as an implementation task, on their own initiative. They had treated these patients at the clinic, had taken CT scans and performed the procedure. The scans was provided to the developers on CD-ROM. The anonymized images were converted into a format suitable for 3D modeling, cleaned up and segmented into different tissue components, and surgical steps encoded. Since the client expected a solution within a fixed amount of time there was not an option to investigate implementation of new simulation techniques, why existing technology had to be used to come up with a solution, together with targeted interaction design.

One of these two cases are quite illustrative in how a real, particular surgical procedure became gestalted in the simulator. A 16 year old girl sought medical assistance due to complaints of two missing teeth. The CT-scan (Fig 2b) showed an impacted primary tooth that never had erupted, and the third molar was hindering the permanent teeth to erupt. The treatment plan was to extract these two teeth with initiatory bone carving. The design team had to make a decision on which tooth that should be extracted first, and the senior surgeon in the team recommended the impacted. The case (Fig. 6) was later tested at the dental school by two surgeon teachers; one junior and one senior professor. The junior surgeon began extracting the molar, and when pointed out that this might not work since the procedure was strictly encoded to begin with the impacted tooth, she sought confirmation from the professor if this was correct. The professor looked at the case and recalled that she had actually herself treated this patient, and remember beginning with the impacted, and explained to the junior surgeon the benefits of starting with the more difficult of the two teeth. In this way, the dialogue between the two surgeons shifted to focus on the patient the case was based upon and the correct treatment plan rather than to which degree of realism it was represented in the simulator.

The result was that using authentic patient cases, rather than completely invented ones, as basis for the simulator exercises brought additional benefits for teaching. The surgeon, especially if she had performed the particular procedure herself, could weave in the simulator exercise in a story about how the actual patient was treated and vice versa.

5.4 Visual and Haptic Appearances of Models

Improvements in the visual models made by the contracted professional 3D artist were well received by the target group. Best results were attained when the artist was given complete control and responsibility of the patient case, supported by a developer. Tuning of the user experience could be done in several steps - the colors and textures of the face in the 3D modeling software used, the colors of the interactive bone in a custom made software, and finally the whole experience tested and positioned in the simulator. The graphics was improved in several stages. Manual segmentation and implementation of coloring allowed for explicit visualizing of structures. Manual segmentation also allowed for the case designer to deviate from the underlying CT scan since the aim is a training case - not necessarily to exactly replicate the anatomy of the patient (contrary to patient-specific surgery rehearsal). Switching to a marching cubes-based renderer was fruitful since that simplified handling of lights and shading as the resulting mesh was rendered in the same manner as the surrounding non-interactive meshes, which also made it possible to define the visual appearance in a congruent manner. The real lift in visual quality was when a professional 3D artist could work with the material - both the surrounding meshes and the interactive jaw model, using tools he was used to.

Taken together the fine-tuning of graphics, haptics (and enclosure) was all received as great improvements, sometimes to a larger extent than any extra software features. Fine-tuning did not necessarily take much time from the project but yielded good results, if done by professionals. It was however important to empower the professionals with tools that enabled them to do the fine-tuning. It was not sufficient to provide terminal commands and text-files for editing parameter values. Not because they would not be able to learn to use them, but because even editing text-files and then launching a simulations that might take a minute to start, give too slow feedback on what a change imply and is way too error-prone for the vast number of changes that was made. The tools developed for tuning (Fig. 9) proved useful in learning what the parameter values meant in terms of user experience. Tuning is, after all, about making small changes and perceiving the result over and over again.

5.5 Off-line Qualities of Physical Design

The generative idea behind the setup and enclosure was twofold. First, as a long-term strategy it was decided to start with the best equipment we could get hold of, get something that the users like, and eventually start optimizing for cost. Second, it was an early desire to create a cohesive product that was easy to use and maintain, and an all-in-one product signal this. Effort was put towards making a product that did not feel computeresque [46] with cables and user interfaces that feels “stuck on” but instead feel integrated to the product.

The enclosure was improved by the furniture designers in three steps, first the brown box (Fig. 8b) was polished by filleting corners and painting it glossy white among some other small details that made it suitable for exhibition. Then a completely new design was made, with well chosen colors to make it fit in the healthcare domain, and finally this design was updated with some detailed improvements such as embedding the infrared 3D glass synchronization emitter. It was evident that industrial design style aesthetics had great impact on how the product was perceived even in prototype stages, in how it signal trustworthiness and identification. It was observed that good hand-support not only helps stable control of the mainipulandum (dental drill) but also guides the surgeon-in-training to an effective posture and body orientation. The physical design of the haptic device, especially the manipulandum, is important as shown by Wang et al [30], where the bulky manipulandum of the Phantom Omni restricted optimal hand positioning in a dental task. The less intrusive physical design of the Phantom Desktop’s manipulandum was another reason why it was selected over the Omni for use together with the mannequin (Fig. 3).

Making a physically large unit also had impact on the development process. It was heavy to move, and we had to arrange logistics to ship it around for demonstrations, user testing and exhibitions. Using a custom-made enclosure not only put demand on constructing it, but that the industrial design is on par with expectations of advanced and expensive equipment.

The mannequin in the simulator is “non-existent” from a computer-centered point of view, i.e. it is not part of any input or output. However, it is a vital part of the simulator both in terms of aesthetics and function (hand-support). The mannequin is usually draped in authentic green surgical protection cloth, as in a real operation. Mannequins are interactive, even without computers, and are more common than simulators for simulation in healthcare. This resonates with observations by Fernaeus [7] that an input/output model of tangible interaction is not sufficient to cover even the most important aspects of interaction.

6 DISCUSSION

Early feedback from students indicated that practicing anatomy was as important if not more important than fine motor skills, as one student mentioned “its not the very technical (movements) that is difficult, its where one should drill” and “anatomy is one of those things they expect us to know, but you feel so insecure anyway”. These statements indicate that there are aspects of surgical proficiency that requires training but not necessarily need to be simulated with the highest level of fidelity, but require well designed content.

Technical implementation and design tools are prerequisites to interaction design. It is a recurring strategic dilemma for simulator developers, especially when using an agile development process, as to when to further implement enabling technology and when to design and fine-tune with what have been implemented so far. In order to do interaction design you need to have something to make the design out of; a design material. Early prototypes in this project therefore was naturally focused on implementing enabling technology, like support for rendering of explicit segments, shading and colors. For haptics, a penalty-based method that only supported sphere-shaped interaction was selected for its straightforward implementation qualities. This algorithm was shown to be useful even for simulating other tools as discussed above. With the advent of more advanced algorithm that support 6-dof multipoint collision detection [10], [47], further design opportunities emerges, like covering the whole instrument with collision points to support haptic interaction with the whole instrument. This in turn requires new design tools and practice for tuning and catering for meaningful user experiences.

In this project it was not feasible to make custom haptic hardware, mainly because of limited competencies and time. Recent work on open platforms for haptic hardware could potentially overcome limitations of off-the-shelf device [48]. In this project, the qualities or user experiences afforded by the available off-the-shelf devices had to suffice, and make the best experience out of. One such quality that was possible to work with was the relative stiffness. The design experience of Kobra showed a trade-off between stability and stiffness, where too high stiffness caused vibrations and too low made the surface spongy. Tuning this parameter was one part of the design work. Another was tuning of cut-rate for different tissues which translated in feeling of hardness.

Tools were developed for fine-tuning the haptic experience including modulation of scale, stiffness and cut rate in a hands-on manner using sliders in a graphical and tangible interface [42]. The need for tuning confirms earlier observations by Morris et al who let an otologist within the group tune their drill parameters [22]. It can be anticipated that development of design and tuning tools for spatial force feedback used in surgery simulators will follow the interest, in research and practice, of design tools developed for its vibrotactile and 1-dof force feedback counterparts [49], [50].

Creation of patient cases required setting up a workflow from CT-scans, to segmentation, voxelization and composition with polygon artwork. The process needed to be iterable with as short cycles as possible, preferably in real-time. While it was manageable to use scripts, custom authoring tools would have been more useful for this task.

Oral surgery education and practice is diverse throughout Europe, and an area worth further attention for harmonisation [51]. A simulator with a growing library of patient cases have the potential of to help harmonisation through exposing the users in different countries to the same set of virtual patients. Working with cases in a simulator could potentially, as a form of practice-based learning, also be enhanced further with active student reflection as suggested by Woodman et al [52].

7 CONCLUSIONS

In general, our work seek to contribute to an ongoing discussion on what constitutes a useful spatial haptics based simulator for teaching and learning of relevant elements of oral surgery, and how such a simulator can be fully realized. This overarching problem spans so many disciplines, e.g. dentistry, product semantics, ergonomics, engineering, learning, cognition, computer science and philosophy of knowledge construction, that no single field could capture its richness in isolation. Through our research-through-design work we have opened up some of this richness to further disciplined investigation. We can now ask ourselves where it seems most fruitful to dig deeper. One finding is that while there is a great need for psycho-motor skill training it might be more fruitful, with the current state of accessible haptic technology, to support surgical tasks of cognitive nature such as case-based problem solving, where spatial haptics still seem to play an important role in assisting and mediating a learning-by-doing and master-apprentice teaching approach. Using authentic cases opens up for new multimodal teaching strategies. A teacher could e.g. present a case within a theory lecture, showing a video of the procedure being applied in a particular case and finally let the students perform the procedure in a simulator on the very same case. Exactly how this kind of learning can best be supported is however subject to future work. In summary four major conclusions can be drawn from the work presented in this paper.

1) Given the haptic technology employed it was more fruitful to support learning cognitive skills rather than motor skills regarding surgery proficiency. This is related to the role the simulator was shown to have as a mediating tool between student and teacher, and that authentic patient cases seem to play an essential role in this.
2) Quantitative assessments was not an absolute necessity for the simulator to have a role in surgery education, and its role can instead be fulfilled by an active teacher.
3) Creative interaction design of patient cases, where tools and actions were simulated with limited realism, was accepted by surgeons. This shows that not only novel realism-enhancing technology, but also interaction design, can make simulators more relevant for surgery teaching.
4) The haptic technology had to be prepared for interaction design beyond encapsulating it in an API, through establishing a workflow and creating design tools which makes it possible to tune the haptic experience. Better tools that streamline the process and offer more possibilities for interaction designers, including support for novel rendering algorithms, are suggested for future work.

CONFLICT OF INTEREST

Fist author is majority owner of the research spinn-off company commercializing the Kobra simulator. The research and analysis is carried out as part of his PhD studies. The other authors have no affiliation with the company.

ACKNOWLEDGMENTS

Furniture designers Anna Leckström and Kolbrún Leósdóttir designed the final enclosure, 3D-artist Marcus Åvall designed several patient cases mentioned in this paper with support by oral surgeon Hans Forsslund. Software engineer Martin Flodin with colleagues at HiQ Stockholm contributed with software development. Research colleagues Sonny Chan, Ioanna Ioannou and Sudanthi Wijewickrema contributed haptic and visual rendering algorithm implementations used in the simulator. Daniel Evestedt with collagues at SenseGraphics have been very supportive especially in the early phases of the project and for the development of underlying H3D API and so has Karl-Johan Lundin Palmerius of Linköping University. We also wish to extend our thanks to collaborators at Riga Stradins University and Karolinska Institutet for invaluable feedback and material for creating the patient cases.

REFERENCES

[1] T. R. Coles, D. Meglan, and N. John, “The role of haptics in medical training simulators: a survey of the state of the art,” *Haptics, IEEE Transactions on*, vol. 4, no. 1, pp. 51–66, 2011.
[2] R. Riener and M. Harders, *Virtual Reality in Medicine*. Springer, 2012.
[3] R. Gottlieb, J. M. Vervoorn, and J. Buchanan, “Simulation in dentistry and oral health,” in *The Comprehensive Textbook of Healthcare Simulation*. Springer, 2013, pp. 329–340.
[4] D. Wang, T. Li, Y. Zhang, and J. Hou, “Survey on multisensory feedback virtual reality dental training systems,” *European Journal of Dental Education*, 2015.
[5] J. Zimmerman, J. Forlizzi, and S. Evenson, “Research through design as a method for interaction design research in hci,” in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2007, pp. 493–502.
[6] K. Kuutti and L. J. Bannon, “The turn to practice in hci: Towards a research agenda,” in *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 2014, pp. 3543–3552.
[7] Y. Fernaeus, J. Tholander, and M. Jonsson, “Towards a new set of ideals: consequences of the practice turn in tangible interaction,” in *Proceedings of the 2nd international conference on Tangible and embedded interaction*. ACM, 2008, pp. 223–230.
[8] A. Petersik, B. Pflesser, U. Tiede, K. Höhne, and R. Leuwer, “Realistic haptic volume interaction for petrous bone surgery simulation,” in *CARS 2002 Computer Assisted Radiology and Surgery*. Springer, 2002, pp. 252–257.
[9] M. Agus, A. Giachetti, E. Gobbetti, G. Zanetti, and A. Zorcolo, “Real-Time Haptic and Visual Simulation of Bone Dissection,” *Presence: Teleoperators and Virtual Environments*, vol. 12, no. 1, pp. 110–122, Feb. 2003. [Online]. Available: http://dx.doi.org/10.1162/105474603763835378
[10] S. Chan, F. Conti, N. H. Blevins, and K. Salisbury, “Constraint-based six degree-of-freedom haptic rendering of volume-embedded isosurfaces,” in *World Haptics Conference (WHC), 2011 IEEE*. IEEE, 2011, pp. 89–94.
[11] S. Wijewickrema, I. Ioannou, and G. Kennedy, “Adaptation of marching cubes for the simulation of material removal from segmented volume data,” in *Computer-Based Medical Systems (CBMS), 2013 IEEE 26th International Symposium on*. IEEE, 2013, pp. 29–34.
[12] B. Tse, W. Harwin, A. Barrow, B. Quinn, M. Cox *et al.*, “Design and development of a haptic dental training system-haptel,” in *Haptics: Generating and Perceiving Tangible Sensations*. Springer, 2010, pp. 101–108.
[13] D. Wang, Y. Zhang, J. Hou, Y. Wang, P. Lv, Y. Chen, and H. Zhao, “idental: a haptic-based dental simulator and its preliminary user evaluation,” *Haptics, IEEE Transactions on*, vol. 5, no. 4, pp. 332–343, 2012.
[14] M. M. Bakr, W. Massey, and H. Alexander, “Evaluation of simodont® haptic 3d virtual reality dental training simulator,” *International Journal of Dental Clinics*, vol. 5, no. 4, 2013.

[15] D. Joseph, J.-P. Jehl, P. Maureira, C. Perrenot, N. Miller, P. Bravetti, P. Ambrosini, and N. Tran, "Relative contribution of haptic technology to assessment and training in implantology," *BioMed research international*, vol. 2014, 2014.

[16] M. Wiberg, "Methodology for materiality: interaction design research through a material lens," *Personal and Ubiquitous Computing*, vol. 18, no. 3, pp. 625–636, 2014.

[17] T. R. Jasinevicius, M. Landers, S. Nelson, and A. Urbankova, "An evaluation of two dental simulation systems: virtual reality versus contemporary non-computer-assisted," *Journal of dental education*, vol. 68, no. 11, pp. 1151–1162, 2004.

[18] N. Zary, G. Johnson, and U. Fors, "Web-based virtual patients in dentistry: factors influencing the use of cases in the web-sp system," *European Journal of Dental Education*, vol. 13, no. 1, pp. 2–9, 2009.

[19] B. Lund, U. Fors, R. Sejersen, E.-L. Sallnäs, and A. Rosén, "Student perception of two different simulation techniques in oral and maxillofacial surgery undergraduate training," *BMC medical education*, vol. 11, no. 1, p. 82, 2011.

[20] H. Rystedt and B. Sjöblom, "Realism, authenticity, and learning in healthcare simulations: rules of relevance and irrelevance as interactive achievements," *Instructional science*, vol. 40, no. 5, pp. 785–798, 2012.

[21] B. Pflesser, A. Petersik, U. Tiede, K. H. Höhne, and R. Leuwer, "Volume cutting for virtual petrous bone surgery," *Computer Aided Surgery*, vol. 7, no. 2, pp. 74–83, 2002.

[22] D. Morris, C. Sewell, N. Blevins, F. Barbagli, and K. Salisbury, "A collaborative virtual environment for the simulation of temporal bone surgery," in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2004*. Springer, 2004, pp. 319–327.

[23] S. Chan, F. Conti, K. Salisbury, and N. H. Blevins, "Virtual reality simulation in neurosurgery: technologies and evolution," *Neurosurgery*, vol. 72, pp. A154–A164, 2013.

[24] N. Von Sternberg, M. Bartsch, A. Petersik, J. Wiltfang, W. Sibbersen, T. Grindel, U. Tiede, P. Warnke, M. Heiland, P. Russo *et al.*, "Learning by doing virtually," *International journal of oral and maxillofacial surgery*, vol. 36, no. 5, pp. 386–390, 2007.

[25] K. H. Höhne, B. Pflesser, A. Pommert, M. Riemer, T. Schiemann, R. Schubert, and U. Tiede, "A new representation of knowledge concerning human anatomy and function," *Nature medicine*, vol. 1, no. 6, pp. 506–511, 1995.

[26] P. Pohlenz, A. Gröbe, A. Petersik, N. Von Sternberg, B. Pflesser, A. Pommert, K.-H. Höhne, U. Tiede, I. Springer, and M. Heiland, "Virtual dental surgery as a new educational tool in dental school," *Journal of Cranio-Maxillofacial Surgery*, vol. 38, no. 8, pp. 560–564, 2010.

[27] R. Q. Van der Linde, P. Lammertse, E. Frederiksen, and B. Ruiter, "The hapticmaster, a new high-performance haptic interface," in *Proc. Eurohaptics*, 2002, pp. 1–5.

[28] D. Bakker, M. Lagerweij, P. Wesselink, and M. Vervoorn, "Transfer of manual dexterity skills acquired in the simodont, a dental haptic trainer with a virtual environment, to reality: a pilot study," *Bio-Algorithms and Med-Systems*, vol. 6, no. 11, pp. 21–24, 2010.

[29] J. Cormier, D. Pasco, C. Syllebranque, and R. Querrec, "Virteasy a haptic simulator for dental education," in *The 6th International Conference on Virtual Learning*, vol. 156, 2011, pp. 61–68.

[30] D. Wang, J. Xiao, and Y. Zhang, "Application: A dental simulator," in *Haptic Rendering for Simulation of Fine Manipulation*. Springer, 2014, pp. 131–160.

[31] J. P. San Diego, A. Barrow, M. Cox, W. Harwin *et al.*, "Phantom prototype: exploring the potential for learning with multimodal features in dentistry," in *Proceedings of the 10th international conference on Multimodal interfaces*. ACM, 2008, pp. 201–202.

[32] I. Ioannou, L. Stern, E. Kazmierczak, A. C. Smith, and L. Z. Wise, "Towards defining dental drilling competence, part 2: a study of cues and factors in bone drilling," *Journal of dental education*, vol. 74, no. 9, pp. 941–950, 2010.

[33] J. Hindmarsh, L. Hyland, and A. Banerjee, "Work to make simulation work:realism, instructional correction and the body in training," *Discourse Studies*, vol. 16, no. 2, pp. 247–269, 2014.

[34] E. Johnson, "Situating simulators: The integration of simulations in medical practice," Ph.D. dissertation, Linköping University, 2004.

[35] ——, "Extending the simulator: Good practice for instructors using medical simulators," in *Using Simulators and Simulations for Education, Training and Research*, P. Dieckmann, Ed. Pabst, 2009.

[36] D. F. McAllister, *Stereo computer graphics and other true 3D technologies*. Princeton University Press, 1993.

[37] L. J. Peterson, E. Ellis, J. R. Hupp, and M. R. Tucker, *Contemporary oral and maxillofacial surgery*. Mosby St Louis, MO:, 1998.

[38] A. Monk, P. Wright, J. Haber, and L. Davenport, "Apendix 1–cooperative evaluation: A run-time guide," *Improving your human-computer interface: a practical technique, Prentice-Hall*, 1993.

[39] A. Rosen, S. Eliassi, U. Fors, E.-L. Sallnäs, J. Forsslund, R. Sejersen, and B. Lund, "A computerised third molar surgery simulator–results of supervision by different professionals," *European Journal of Dental Education*, vol. 18, no. 2, pp. 86–90, 2014.

[40] A. Vallgårda and Y. Fernaeus, "Interaction design as a bricolage practice," in *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 2015, pp. 173–180.

[41] J. Forsslund, E.-L. Sallnäs, and K.-J. Palmerius, "A user-centered designed foss implementation of bone surgery simulations." in *WHC*, 2009, pp. 391–392.

[42] J. Forsslund and I. Ioannou, "Tangible sketching of interactive haptic materials," in *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*. ACM, 2012, pp. 111–114.

[43] T. H. Massie and J. K. Salisbury, "The phantom haptic interface: A device for probing virtual objects," in *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems*, vol. 55, no. 1. Chicago, IL, 1994, pp. 295–300.

[44] V. Neufeld and G. R. Norman, *Assessing clinical competence*. Springer New York, 1985, vol. 7.

[45] E. Johnson, P. Ström, A. Kjellin, T. Wredmark, and L. Felländer-Tsai, "Evaluating instruction of medical students with a haptic surgical simulator: The importance of coordinating students' perspectives," *Journal on Information Technology in Healthcare*, vol. 2, no. 3, pp. 155–163, 2004.

[46] T. Djajadiningrat, S. Wensveen, J. Frens, and K. Overbeeke, "Tangible products: redressing the balance between appearance and action," *Personal and Ubiquitous Computing*, vol. 8, no. 5, pp. 294–309, 2004.

[47] D. Wang, X. Zhang, Y. Zhang, and J. Xiao, "Configuration-based optimization for six degree-of-freedom haptic rendering for fine manipulation," *Haptics, IEEE Transactions on*, vol. 6, no. 2, pp. 167–180, 2013.

[48] J. Forsslund, M. Yip, and E.-L. Sallnäs, "Woodenhaptics: A starting kit for crafting force-reflecting spatial haptic devices," in *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 2015, pp. 133–140.

[49] M. J. Enriquez and K. E. MacLean, "The hapticon editor: a tool in support of haptic communication research," in *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2003. HAPTICS 2003. Proceedings. 11th Symposium on*. IEEE, 2003, pp. 356–362.

[50] C. Swindells, E. Maksakov, K. E. MacLean, and V. Chung, "The role of prototyping tools for haptic behavior design," in *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2006 14th Symposium on*. IEEE, 2006, pp. 161–168.

[51] M. Nasseripour, C. Hervé, and J.-P. Meningaud, "Oral surgery in the european union: challenges of diversity in training and practice," *European Journal of Dental Education*, 2015.

[52] T. Woodman, B. Pee, H. Fry, and E. S. Davenport, "Practice-based learning: emerging professional characteristics, self-concepts and patterns of knowing in dental training," *European Journal of Dental Education*, vol. 6, no. 1, pp. 9–15, 2002.

Paper B

WoodenHaptics: A Starting Kit for Crafting Force-Reflecting Spatial Haptic Devices

Jonas Forsslund
Royal Institute of Technology
Stockholm, Sweden
jofo02@kth.se

Michael Yip
Stanford University
Stanford, CA, USA
mcyip@stanford.edu

Eva-Lotta Sallnäs
Royal Institute of Technology
Stockholm, Sweden
evalotta@kth.se

ABSTRACT

Spatial haptic interfaces have been around for 20 years. Yet, few affordable devices have been produced, and the design space in terms of physical workspace and haptic fidelity of devices that have been produced are limited and discrete. In this paper, an open-source, open-hardware module-based kit is presented that allows an interaction designer with little electro-mechanical experience to manufacture and assemble a fully working spatial haptic interface. It also allows for modification in shape and size as well as tuning of parameters to fit a particular task or application. Results from an evaluation showed that the haptic quality of the WoodenHaptics device was on par with a Phantom Desktop and that a novice could assemble it with guidance in a normal office space. This open source starting kit, uploaded free-to-download online, affords *sketching in hardware*; it "unsticks" the hardware from being a highly-specialized and esoteric craft to being an accessible and user-friendly technology, while maintaining the feel of high-fidelity haptics.

Author Keywords

Guides; do-it-yourself; open-source; open-hardware; spatial haptics; force-feedback; haptic device; hardware sketching; interaction design

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation (e.g. HCI): User Interfaces (D.2.2, H.1.2, I.3.6) Haptic I/O

PART I: INTRODUCTION

Spatial haptic interfaces are grounded human interface devices that track a physical manipulandum (handle) in space, and provides the means of reflecting a directional force on that manipulandum and consequently the user. With the device and appropriate haptic rendering algorithms, an end-user can explore a virtual environment through virtual coupling between the manipulandum's position and a representative avatar [24]. As more applications have 3D user interfaces [1], spatial haptics becomes increasingly useful for feeling the shapes of occluded objects, collisions, object stiffness and inertias, surface textures, and so on [24]. While vibrotactile haptics are ubiquitous in all cellular phones, the 3D force-displaying counterparts, spatial haptics, are not as widely present. Devices for spatial haptics have been used for surgical simulation [3], physical rehabilitation [2] and for hand tool guidance [9]. Applications in other domains could benefit from including spatial haptics, but this requires finding a good match between the qualities of the device employed and the application. The common devices commercially available (Figure 2) represent only a limited design space in terms of fidelity, price and capabilities (e.g. workspace dimensions and maximum force). Therefore, application specific devices have sometimes been developed, such as for simulation of micro-surgical bone drilling [22]. However, engineering a haptic device is still a large commitment and only feasible in highly specialized robotics labs that have the mathematical and mechanical know-how to realize and achieve high quality haptics in their design.

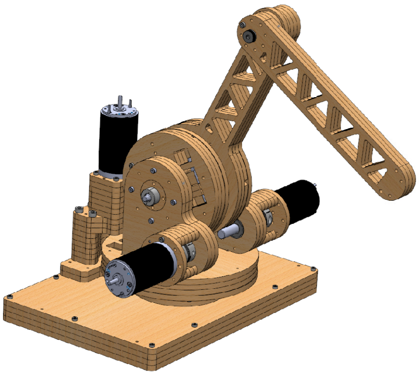

Figure 1. The completed wooden haptics device.

Research and advances in haptics tend to be focused on technological refinements and little attention is directed towards holistic design and aestetics of the devices. As a response to this, Moussette has proposed the intersection of Design and Haptics as a new field of study [17]. This approach is centered around hands-on workshops [16], exploration, making and sketching in hardware with simple haptics [18] in order to get a heightened sensitivity to haptics. The benefits that follow from this approach can be extended to spatial haptics.

TEI 2015, January 15–19, 2015, Stanford, California, USA.

ACM 978-1-4503-3305-4/15/01 ...$15.00.
http://dx.doi.org/10.1145/2677199.2680595

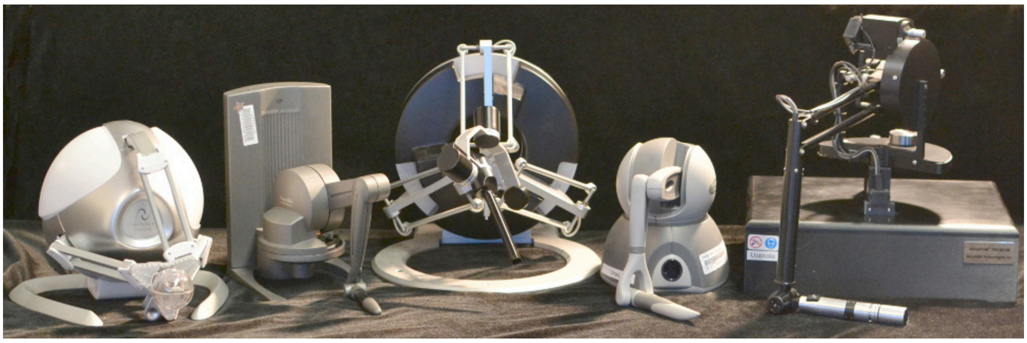

Figure 2. Common spatial haptic devices. From the left: Novint Falcon, Phantom Desktop (now 3D Systems Geomagic Touch X), Force Dimension Omega, Phantom Omni (now 3D Systems Geomagic Touch), and Phantom Premium 6-DOF (now 3D Systems Geomagic Phantom Premium).

In this paper we introduce a new spatial haptic device that is designed and packaged as a kit in a way that a designer can re-configure or re-design it, thus adapting it to different domains or requirements. This kit, called WoodenHaptics, is designed to provide comparable haptic fidelity to commercial devices. We carefully chose to encapsulate certain technical details (e.g. the electrical system), whereas others are very visible (the mechanical structure and wire rope power transmission); this is intended to help designers focus more on designing for their application rather than problem solving through mechanical and electrical nuances and details. Through these efforts we have reduced the "stickiness" [29] of constructing spatial haptic devices to a level where a designer can design, create, assemble and modify their own version of the device. We evaluated its ease of adoption on non-specialists and first-time builders (under guidance) to show that it provides a valuable medium for "hardware sketching" in spatial haptics.

Background

With the advent of Massie's [10] force-reflecting device in 1993, spatial haptics as a multi-purpose human-computer interaction interface became popular through the commercialized Phantom series [11] still available on the market today (figure 2). While all these devices can read a spatial position and render a directional force back to the user through the manipulandum, the experience and quality of the forces/movement is quite different, something that is also reflected in the price tag that ranges from $300 to over $20,000 USD. Other devices with equivalent functionality from the user's perspective have since appeared, but the market is far from being as diverse as that of computer mice and joysticks.

A great deal of fundamental theory for building a haptic device has been described in e.g. [7] and [6]. However, bridging the gap from reading the fundamentals to constructing a fully-functional 3D spatial haptic device of the prototypical Phantom [10, 11] is still a daunting task to a common interaction designer and is only feasible for an expert roboticist. Much practical and tacit knowledge is required to actually make a high-fidelity haptic device, since it relates to making a correct combination of design choices, ranging from which type of motors, what mechanical structure, which control paradigm and even which type of screws to choose. Then the parts need to be located and purchased, which can be very time-consuming and confusing. Furthermore, the robotics literature describing the mathematics required to operate the haptic device [7, 4] might be overwhelming in scope and content to the electro-mechanical novice.

Kits and Tools for Design

Kits and tools for design through making and crafting is an active research area in TEI/HCI [5, 8, 13, 14, 25, 26]. Phidgets [5] is used to simplify development of physical interfaces through providing "everyday programmers" with a kit of pre-made electronic physical widgets. Toolkits have been described as particularily instrumental to sketching in hardware [19]. Software tools have been developed explicitly targeting designers without production training in electronics [8] and furniture design [26]. Even the notion of an "untoolkit" has been proposed as a conceptual tool to leverage existing standard materials and components in new artifacts [14]. Open source hardware designed for personal fabrication has been described as an approach to support design of different aspects of electronic products, since the designer only has to modify those parts of the design pertinent to the designers interest and still get a working product [13]. Other kits such as the Hapkit [15] is constructed with goal of teaching engineering concepts per se through hands-on experience [21, 27].

In this paper, WoodenHaptics is presented as an open source "starting kit" for material exploration, design and realization of application specific force reflecting haptic devices. This distinguishes our kit from a toolkit where combinations of provided parts yield many designs (we only provide one reference design in the kit itself), and an untoolkit where none of the modules from the kit goes into the final design. The intended audiences for the kit are interaction design studios and HCI researchers, especially for cases where applications require different form factors (e.g. length of arms), and other properties (e.g. maximum force) that off-the-shelf devices won't meet, something we in our own practice have seen a need for. Our primary aim is to support professional design

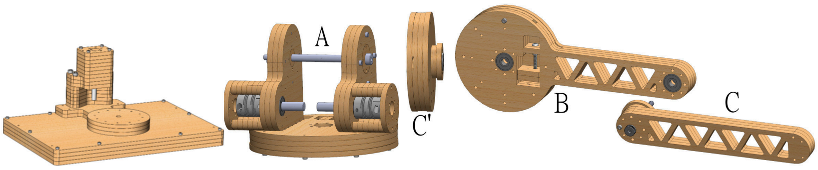

Figure 3. The parts needed to be assembled by user: the base driving the body A, that in turn drives body B and, indirectly through C', body C.

exploration of the interaction qualities that follows from modifications of the reference design. A high-quality spatial haptic interface requires careful attention to its computational and physical materials, which can *only* be experienced as a whole with a fully assembled device. Therefore, a designer who decides to deepen her engagement beyond software will have to set up a workbench for explorations. WoodenHaptics serves as a complete kit for starting such explorations.

PART II: KIT DESCRIPTION AND USAGE

A pair of interaction design researchers without training in mechanical engineering were provided with the kit consisting of a complete set of hardware components that make up a full spatial haptic device. This included all the pre-cut plywood parts, screws, bearings and all other mechanical components (Figure 4). The kit came with three motors (Maxon RE40) with pre-mounted encoders, and an electronics box (Figure 5) that connects to a 48V lab power supply and a standard PC equipped with a Data Acquisition interface (DAQ, Sensoray S826). The kit requires only a limited set of tools: hex keys, a steel wire crimping tool and snippers, a torch and an arbor press (Figure 8); a list of these tools and where to purchase them are available online. Software required to operate the device was included as well. Thus, the builder can immediately run available demo programs and proceed to application-specific development.

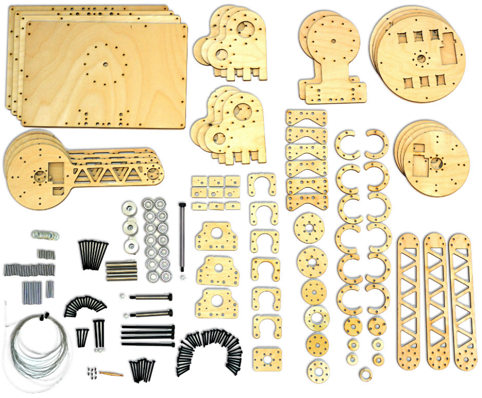

Figure 4. The parts included in the kit. Not shown here are the three motors, electrical cables and the electronics box (Figure 5), and configurable software that completes the kit.

Assembly

The entire structural pieces are manufactured from a laser cutter out of 6 mm (or approximately 1/4 inch) plywood. To form stiff three-dimensional parts from the flat sheets, several layers are stacked and held together with screws. All holes in the plywood parts are adjusted with sub-millimeter precision such that all screws can self-tap (self-thread) the holes, allowing for quick assembly and disassembly. Stacked parts are aligned by inserting dowel pins (precision cylindrical pins) with an arbor press before adding screws. Bearings are press-fit as well using the arbor press. In fact, there is no use of bondants or adhesives, resulting in a visually and mechanically clean, quickly disassemble-able and reconfigurable device. The kit comes with instructions on how to assemble the main bodies, as well as video documentation.

The bodies A, B and C (figure 3) form the three links or *degrees of freedom* (DOF) that together enable the tip of the device (P in figure 6) to be moved left/right, up/down and in/out. Each DOF is coupled independently to a motor through wire rope. The angle of each DOF is a fixed ratio to the rotation of the motor shaft, and therefore the angles are measured by the *encoders* mounted on each motor (figure 5).

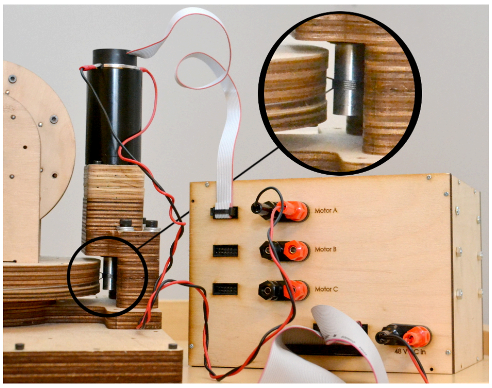

Figure 5. The first degree of freedom motor connected with power and encoder wires to the electrical interface. The close-up view shows the aluminum capstan and wire rope coupling.

Cabling

The kit utilizes cable drive for all its transmissions: a strong steel wire rope transmits the power from each motor to its own respective link. Figure 5 shows a standard cable drive transmission used in all degrees of freedom. The motor shaft is attached to the *capstan*, which is a shaft for a cable to wrap around and grip. The cable makes 5 wraps around the capstan and is terminated at both ends. The cable needs to be taut to grip the capstan, which is done at the termination by either tightening or loosening a screw. For the last link, a turnbuckle is used to maintain a taut cable. Now, for each body, when the capstan is rotated with the cable gripped firmly to it, the body is then rotated; alternately, when the body is rotated, the capstan is subsequently rotated. This completes the transmission assembly, allowing for the motors and the driven axis to not require collocation. This allows for gearing up of the motor torques for achieving larger forces without using gearboxes, as well for easy replacement of motors. The reasons for these design choices are further discussed in Part III.

Electrical system

The kit comes with three high-quality motors, each driving a respective degree of freedom. The designer only has to connect the encoder to the electronics box (that routes them to the computer), and each motor power cable to respective output of the electronics box (figure 5). Two ribbon cables connects the electronics box with the Sensoray S826 board on the PC.

The motors chosen are more powerful than is common in the devices pictured in figure 2. They are specified for allowing a max continuous current of 3.16A safely, and we have limited the maximum current to 3 ampere. This means that the user will not have to worry about electrical heat, burning, etc, which is the case when the motors are overdriven in short periods of time, which is common practice otherwise.

Software Configuration

The kit is complete with a working open-source software module for the mechanical design that comes with the kit. If a dimension have been changed by the user or tuning of the experience is desired, the user can easily modify a variable in a text file to represent this change. The variables of interest to change are: the diameter of each capstan and body, the length of each link and the mass and mass center of each body. This effectively is equivalent to changing the gearing of the motor, and changing the size of the workspace, respectively. The design also affords the easy replacement of motors with different motors, but the user will then need to adjust the torque/current ratio as defined by their motor datasheet. The maximum stiffness and damping of the complete device can be found retroactively by experimenting and adjusting the values accordingly.

PART III: FUNDAMENTALS AND THEORY

This section describes how the kit was developed and the design principles/considerations involved. We cover here briefly, the mechanical structure, the kinematics and control theory applied and why certain design decisions were made to support easy user fabrication and modification in particular. We are knowingly only addressing one kind of mechanical structure (the serially linked) and one control type (impedance control). Alternatives are parallel mechanical structure like e.g. Novint Falcon, and admittance control [28] that requires force sensing.

Haptic fidelity and transparency

The power transmission of a haptic device is a critical component that needs to be designed carefully as it transmits the forces and velocities from software to the hand of the user. Haptic fidelity is achieved by having *transparency* in the system – that is, desired forces and velocities defined in software accurately match forces and velocities delivered to the user. Three major contributions that reduce the transparency of a system is friction (resulting in diminished haptic perception), backlash (resulting in chatter in the motors and the device), and physical compliance (resulting in a loss of ability to percieve stiff environments).

Although motor and gearbox combinations are commercially much more common, cable drive transmission is the standard for haptic devices because it provides a near frictionless transmission and has no backlash, which no gearbox can achieve. The choice of cable is also an important factor: a cable with high flexibility will provide greater transparency as the users will not percieve the forces required to “bend” and “unbend” the cable as the capstan rolls. Therefore, uncoated stainless steel cables with high count of individual steel fibers (we use a cable of 0.54 mm diameter, with fibers in a 7 × 7 configuration, more is recommended if available) present a viable option. The grip of the cable on the capstan increases exponentially as the cable wraps around, and therefore even a few turns will immediately prevent the cable from slipping. In practice, 5 turns is more than enough to prevent any slipping between the capstan and the cable. This principle is also how the final link’s cable transmission (using the cable loop and turnbuckle) works without slipping.

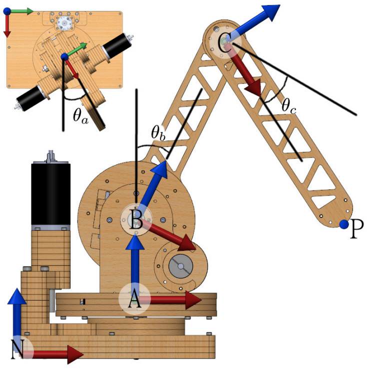

Figure 6. The device is a serially linked mechanism, where the angle of A, B and C uniquely defines point P in the base reference frame N.

Another major design choice for achieving high haptic fidelity was to mount the second and third axis motors on body A (figure 6). This choice is highlighted by three intentional benefits: the simplified and shorter cable routing affords better transparency in the system, placement of the motors allow for easy access, removal, and installation for other motors of different sizes, and shorter cable routing reduces the chance of the transmission de-cabling. Failure in the cable transmission (e.g. cable snap or comes loose) is thereby localized to its own small section of the device, while allowing for the rest of the device to remain in working order.

The third major contribution to achieving transparency in the system is having a physically stiff device. Increasing stiffness (ie. reducing compliance) in the device's structure is done by increasing the second moment of inertia of each link (e.g. making link wider so they do not twist), improving the joint stiffnesses (e.g. by increasing shaft diameters, increasing distance between shaft bearings that hold the shafts straight), and using a stiff material. Because plywood is a layered composite, it is in fact quite stiff and yet still reasonably light; it is also soft enough for self-tapping holes and very minor misalignments that all contribute to making the device more accessible and forgiving to build, without sacrificing substantial haptic fidelity.

Finally, each motor to capstan combination is connected through a flexible shaft coupler, which acts to not only reduce friction caused by misalignments in the axes of the motor and the capstan, but also serves as an easy way to swap out different motors and find the best motor for an application without performing any disassembly of the cable transmission. This serves to promote hardware sketching on the actuator side.

Mathematical description and analysis

In order for the spatial haptic interface to be useful, the position of the user's hand and thus the end-effector of the device must be known in space. This is achieved by measuring the angle of the motor shafts using encoders, and doing the forward kinematics to map motor angles to cartesian-space position. Below we provide only a very brief overview of the mathematics involved in achieving haptic feedback; all the details regarding the position and force mappings discussed in this section are derived from [4] using standard techniques for robot manipulators.We derive the manipulandum position through forward kinematics, which in this case is a classic 'RRR' configuration manipulator; that is, it has three moving links which are serially-linked through revolute(R) joints (figure 6), and can be calculated as follows:

$$\begin{aligned} P_x &= \cos\theta_a(L_b\sin\theta_b + L_c\cos(\theta_b+\theta_c)) \\ P_y &= \sin\theta_a(L_b\sin\theta_b + L_c\cos(\theta_b+\theta_c)) \\ P_z &= \qquad L_b\cos\theta_b - L_c\sin(\theta_b+\theta_c) \end{aligned} \tag{1}$$

where L is the length of each body to the next and $\theta_{a,b,c}$ are the angles of respective body. To give a force $\mathbf{F}$ at the manipulandum, the body torque τ is computed as:

$$\tau = \mathbf{J}^\top \mathbf{F}, \tag{2}$$

where $\mathbf{J}$ is called the Jacobian matrix, and is the first partial derivative of the forward kinematics (1) with respect to the body angles θ [4]. A final necessity to account for is the weight of the manipulandum: without compensating for the manipulandum weight, the user will have to hold up the device's weight in their hands. To compensate for gravity, the weights of the three links as well as their centers of gravity are estimated, and motor torques to counter gravity forces are applied. For ease of use, the kit's software module allow for tuning of the link parameters while masking the mathematics involved to solve for force, position, and gravity.

Electrical system

The electrical system has two purposes: to drive the motors and to measure their angular position. The torque of the motor used is proportional to the current that is driven through it, not the voltage it is supplied. Therefore a current or torque controller (in our case Maxon ESCON 50/5) is connected between a generic power supply and the motor.

It is worth mentioning that the components used (motors, amplifiers, encoders and acquisition card) are of professional lab quality and should not be confused with hobbyist counterparts. While efforts to replace them with lower cost alternatives are very welcome, one has to be careful in preserving the precision needed. For example, the delay has to be less than 1 ms and the resolution and quality of D/A converter sufficient [23]. However, this also brings to the surface the potentials of this starting kit, as it allows users to explore what their haptic tolerance is for lower-cost alternatives.

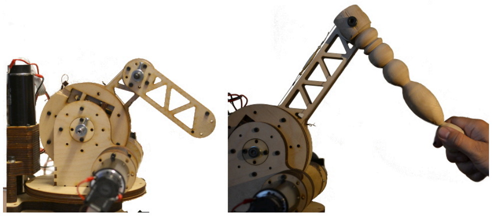

Figure 7. Exploring designs: Mini-woody with smaller workspace and larger forces, and a different handle arm crafted using a lathe.

Variations

While the starting kit provides everything needed to complete a functioning device, the intention is to invite the designer to modification and variants of the design. Below we highlight a few interesting areas worthy of exploration:

Workspace

The user can very easily try different sizes (lengths) of the body, and experience the difference in scaling up or scaling down their reachable workspace and the haptic perception. Figure 7 depicts a smaller version, that also, as a direct consequence, can render larger forces (table 1).

Motors and Encoders

The user can switch between using high-cost, high-quality motors and encoders, or a low-cost alternatives. This allows the designer to identify the specific factors and limits of haptic fidelity (e.g. the backlash from a geared motor versus ungeared motor, the cogging or friction from a $20 hobby shop motor versus a $300 motor). Effects of motor size can also be investigated.

Material

Plastics (such as acrylic) are as easy to cut as plywood, and comes in different colors for the designer to experiment with, but can be brittle. They also tend to be heavier, which have to be supported with more motor torque for gravity compensation. Aluminum is lightweight and stiff, but needs to be cut using special equipment (warter-jet cutter) and requires threading holes separately. Solid and composite wood choice can provide different stiffness and weight tradeoffs. Physical stiffness, the inertia of the device, and even Visual appeal can be explored by using different materials. Figure 7 shows a variant where one part is hand-fabricated from solid wood using a lathe.

Add-ons

A user may add buttons, sensors or even vibrotactile actuators on the manipulandum, which can further improve perception of textures [12]. Different grips or end-attachments that interface with the user can be explored.

PART IV: EVALUATION

Three aspects of the reference design and starting kit were evaluated. First, to what extent could someone without robotics training or access to a sophisticated lab use the kit, assemble the device and make it work? Second, how does the device compare to commercially-produced haptic devices? And finally, what are the technical properties?

Assembly Workshop

In order to investigate the feasibility of the kit, the extent it can reduce stickiness, and the level of instructions required, a workshop was held with two researchers previously inexperienced in robotics construction. While the first version of the device was built in a robotics lab (in the US), this workshop was held in a interaction design lab on another continent (in Europe). All the parts in the kit seen in figure 4 and motors, electronics box and a tool set of hex keys, cable crimper, cutter and arbor press were presented to the users. A computer with the software installed to run the completed device with a virtual environment was provided.

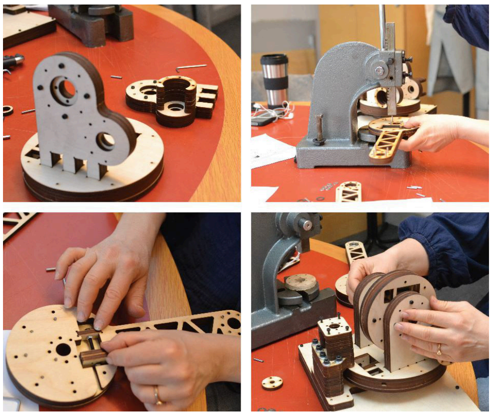

Figure 8. User assembly of the wooden haptics device as part of the evaluation of the feasability of the starting kit.

The instruction material was in the form of print-outs of the bodies with notations of screw location and so on. The users took turns building while being supervised by one of the authors, and guidance was provided when deemed necessary. The assembly sessions were videotaped and notes and build times recorded.

It took the users approximately 11 hours distributed over 4 sessions to successfully assemble the complete WoodenHaptics device. Most of the guidance required was initially regarding the proper use of the tools, order of insertion of pins and screws and so on, and during the process of cabling (the process of mounting the wire rope). These instructions could be readily provided in a detailed construction document as is provided in Lego® sets. Since only one copy of the kit was present and thus some errors would be irreversible, care was taken to make sure stacking of similar parts and assembly order was performed correctly. Despite that, one part got placed in a flipped direction, but could still work in a satisfactory manner. Some bearings were not inserted until parts added later made it cumbersome, and had to be mounted with mallet (with risk of damaging them). Two outer diameters were wrongly adjusted by the authors and had to be reduced with sand paper, since cutting new parts was not an option (no laser cutter in the lab, and long turn-around for ordering). Finally, some extra tools were used: a pair of tweezers during the cabling, and the mallet when the arbor press was too small. With this said, the final device worked as well as the one built previously.

Perceived quality of finished assembly

An evaluation was conducted regarding the quality of the haptic feedback, comparing the Phantom Desktop, Phantom Omni, Novint Falcon and our assembled version of the starting kit, hereafter named “Woody” for short. For this evaluation, 10 participants (5 women and 5 men, mean age 32) were recruited among students and faculty. The same virtual environment[1] was used for each device that were placed side-by-side, which enabled the user to go back and forth between the devices (Figure 9).

Figure 9. The setup of the evaluation showing the four different devices and the demo application used in the test.

---

[1] Chai3D version 3.0 example 13 - Primitives, www.chai3d.org

The evaluation consisted of two parts. The participants first task was to rate the experienced quality of the haptic feedback using the devices one at a time. Experienced quality was evaluated in a questionnaire using a seven point Likert-type scale with four items asking about to what degree the haptic feedback felt being of *high quality*, *precise*, *smooth* and *distinct*. In the analysis, the four items were summed into one dimension measuring subjective experience of haptic feedback quality. The questionnaire was filled in after experiencing each haptic device respectively. In the second part, participants compared all devices again and reported which one of the commercial devices they felt most closely matched that of Woody in haptic quality on their own terms.

The results, as the average summed ratings (std. dev.) of each device was as follows. Desktop 25.8 (2.10), Omni 19.5 (5.7), Falcon 11.2 (6.3) and Woody 24.4 (3.37). The participants thus rated Woody's experience between Phantom and Omni. It resonates with the second part of the evaluation when the participants were explicitly asked to draw a line between Woody and the device they thought was most similar to Woody. The results show that 7 of the participants thought that Woody was most similar to the Desktop, 3 reported that Woody felt the most similar to the Omni.

Changing haptic properties when varying design

In order to demonstrate the ability to modify haptic performance through design, different versions of the device were constructed with the starting kit as a base. Woody (B=20 cm, C=22cm) and a "mini" variant (B=10 cm, C=12 cm) were created to explore the trade-off betweeen workspace and force-feedback properties (Figure 7). Stiffness was empirically tuned at a level where no vibrations was felt while interaction were at a maximal stiffness (5 N/mm). The results are presented in Table 1. For comparison, maximum force values for Omni and Falcon were added from [22] and friction and workspace measurements were taken for these devices as well. A lower bound for the workspace was found by placing a virtual sphere in the virtual origin and increasing its diameter while maintaining that the whole sphere can easily be touched from the outside. Peak force was found for Woody and Mini-woody by calculating, but not physically outputting, the largest force that could be applied in x, y, and z direction until at least one motor saturates. The workspace was swept and the lowest value was recorded. Since the motors were specified to handle more than 3 ampere nominally the continuous force and peak force is the same. Back-drive friction was measured using a hand-held digital force gauge (FG-5000A-232, Lutron Electronics, Taiwan) slowly moved from one side of the workspace (as defined above) to the other passing the origin, sideways, inwards and upwards. Ten measurements were taken and averaged for each direction and device. Gravity compensation was enabled if available. Omni, which lacked active gravity compensation was supported by a 2 m thin string from the ceiling in sideways and inwards measurements.

This study demonstrated the exploration of trade-offs in changing workspace dimensions with forces and friction. The commercial devices provided the static performance mark for which the modified designs were compared against. It can be seen that there is no one device that provides the largest workspace, forces, or minimum friction altogether, showing the strengths and limitations of each design and the benefit for hardware sketching.

	Woody	Mini-w.	Omni	Falcon
workspace	200+	80+	100+	60+
peak force	9.9+	19.0+	3.3	8.9+
cont. force	9.9	19.0	0.88	8.9
friction	0.6/0.7/0.9	0.6/1.0/0.9	0.2/0.4/1.1	1.2/3.6/1.3

Table 1. Varying haptic properties through different design changes, and comparison to commercial devices.

PART VI: DISCUSSION

We have shown how the WoodenHaptics starting kit can be an engaging spatial haptics device testbed without many of the sticky issues usually involved in the craft. We have furthermore demonstrated that high haptic fidelity was achievable using WoodenHaptics, on par with commercial devices. For an interaction designer, the WoodenHaptics toolkit serves to:

- help the designer understand the fundamentals of the mechanism (e.g. it shows clearly how three motors work together to generate one force vector at the end point of the manipulandum).
- enable the designer to incorporate the device into their own projects quickly and easily without being a electromechanical expert.
- enable the designer to explore the user experience (by objectively changing/tuning certain parameters or replace components such as motors).
- establish a common language between designers and experienced hardware engineers. The designer can now say "can you make a device like this, but smaller?" or "what could or needs to change for us to manufacture something like this for our application?"

With WoodenHaptics, a designer can create variations of a serially-linked 3-DOF grounded spatial haptic device. The constraints imposed by the kit frees the designer from solving many electrical, computational and mechanical problems since these have already been solved; it instead allows the user to innovate in terms of motor choices, workspace dimensions, physical material, aesthetics and extended functions like buttons. As personal fabrication of parts becomes easier, e.g. through direct interaction with a laser cutter [20] or software tools [26], designers can quickly explore different variations that can optimize their haptic experience for a particular application.

Common haptic devices and application programming interfaces sometimes give wrong expectations of what experiences they actually support. For example, Mousette [17] noted that "hardware hard is relative" from his experiments with a commercial haptic device where a virtual object specified to be of maximum stiffness still yielded a sensation he refer to as "mushy hard". It is likely that he would have had a different experience with a device equipped with more powerful motors, or if the developers had used another terminology than "hardness" to describe the feature. By crafting with WoodenHaptics one can learn, experience, quantitatively define, *and alter* mushiness or other unarticulated haptic experiences.

WoodenHaptics is not intended to replace off-the-shelf devices and is not necessarily cheaper. Instead it offers unique opportunities for dedicated designers as a workbench for exploring the experiential qualities of new designs. The components mentioned here carry a cost of about 3000 USD. Future work, especially on the electronics side, will lead to significantly lower costs (e.g. use of custom circuit board).

Finally, designers are encouraged to share their experiences and designs with the community, and improve upon the kit itself. This allows them to go beyond the original constraints set by the kit and its modules when they are ready. We also expect simpler versions for e.g. 2-DOF planar device to benefit from our modules. Bringing spatial haptic device design to a larger audience allows them to share more perspectives on both what a designs should look like and how they should be evaluated.

ACKNOWLEDGMENTS

WoodenHaptics was initially developed at Stanford University. We are thankful to Reuben Brewer and Ken Salisbury for introducing many of the concepts discussed in this paper in CS235: Applied Robot Design for Non-Robot Designers. We would also like to thank Adam Leeper and Ylva Fernaeus.

REFERENCES

1. Bowman, D. A., Kruijff, E., LaViola, J. J., and Poupyrev, I. *3D User Interfaces: Theory and Practice*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 2004.
2. Broeren, J., Rydmark, M., and Sunnerhagen, K. S. Virtual reality and haptics as a training device for movement rehabilitation after stroke: a single-case study. *Archives of physical medicine and rehabilitation 85*, 8 (2004), 1247–1250.
3. Chan, S., Conti, F., Salisbury, K., and Blevins, N. H. Virtual reality simulation in neurosurgery: technologies and evolution. *Neurosurgery 72* (2013), A154–A164.
4. Craig, J. J. *Introduction to robotics*, vol. 7. Addison-Wesley Reading, MA, 1989.
5. Greenberg, S., and Fitchett, C. Phidgets: easy development of physical interfaces through physical widgets. In *Proceedings of UIST*, ACM (2001), 209–218.
6. Hayward, V., and MacLean, K. E. Do it yourself haptics: part i. *Robotics & Automation Magazine, IEEE 14*, 4 (2007), 88–104.
7. Kern, T. A. *Engineering haptic devices*. Springer, 2009.
8. Knörig, A., Wettach, R., and Cohen, J. Fritzing: a tool for advancing electronic prototyping for designers. In *Proceedings of TEI*, ACM (2009), 351–358.
9. Lee, J. Haptic intelligentsia, http://studio- homunculus.com/portfolio/ haptic-intelligentsia-human-prototyping-machine, apr 2014.
10. Massie, T. Design of a three degree of freedom force-reflecting haptic interface. Bsc thesis, MIT, 1993.
11. Massie, T. H., and Salisbury, J. K. The phantom haptic interface: A device for probing virtual objects. In *Proc. Haptic Symposium 1994*, vol. 55, Chicago, IL (1994), 295–300.
12. McMahan, W., and Kuchenbecker, K. J. Displaying realistic contact accelerations via a dedicated vibration actuator. In *WorldHaptics*, IEEE (2009), 613–614.
13. Mellis, D. A., and Buechley, L. Case studies in the personal fabrication of electronic products. In *Proceedings of the Designing Interactive Systems Conference*, ACM (2012), 268–277.
14. Mellis, D. A., Jacoby, S., Buechley, L., Perner-Wilson, H., and Qi, J. Microcontrollers as material: crafting circuits with paper, conductive ink, electronic components, and an untoolkit. In *TEI* (2013), 83–90.
15. Morimoto, T. K., Blikstein, P., and Okamura, A. M. Hapkit: An open-hardware haptic device for online education. In *Haptics Symposium, 2014 IEEE* (Feb 2014), 1–1.
16. Moussette, C. Learn to make, make to learn: Reflections from sketching haptics workshops. *Design and semantics of form and movement* (2012), 180.
17. Moussette, C. Simple haptics: Sketching perspectives for the design of haptic interactions.
18. Moussette, C., and Banks, R. Designing through making: exploring the simple haptic design space. In *Proc. TEI*, ACM (2011), 279–282.
19. Moussette, C., and Dore, F. Sketching in hardware and building interaction design: tools, toolkits and an attitude for interaction designers. *Proc. of Design Research Society* (2010).
20. Mueller, S., Lopes, P., and Baudisch, P. Interactive construction: Interactive fabrication of functional mechanical devices. In *Proceedings of UIST*, ACM (New York, NY, USA, 2012), 599–606.
21. Okamura, A. M., Richard, C., and Cutkosky, M. R. Feeling is believing: Using a force-feedback joystick to teach dynamic systems. *Journal of Engineering Education 91*, 3 (2002), 345–349.
22. Salisbury, C., Salisbury, J., Gillespie, R., and Blevins, N. A microsurgery-specific haptic device for telerobotic medical treatment. In *ANS Joint Emergency Preparedness and Response & Robotic and Remote Systems, Topical Meeting*, vol. 6 (2008).
23. Salisbury, C. M. *Haptic hardware: evaluation, design and placement*. PhD thesis, Stanford, 2009.
24. Salisbury, K., Conti, F., and Barbagli, F. Haptic rendering: introductory concepts. *IEEE computer graphics and applications 24*, 2 (2004), 24–32.
25. Sankaran, R., Ullmer, B., Ramanujam, J., Kallakuri, K., Jandhyala, S., Toole, C., and Laan, C. Decoupling interaction hardware design using libraries of reusable electronics. In *Proc. TEI*, ACM (2009), 331–337.
26. Saul, G., Lau, M., Mitani, J., and Igarashi, T. Sketchchair: an all-in-one chair design system for end users. In *Proc. TEI*, ACM (2011), 73–80.
27. Shaver, M., and Maclean, K. The twiddler: A haptic teaching tool for low-cost communication and mechanical design.
28. Van der Linde, R. Q., Lammertse, P., Frederiksen, E., and Ruiter, B. The hapticmaster, a new high-performance haptic interface. In *Proc. Eurohaptics* (2002), 1–5.
29. Von Hippel, E. Perspective: User toolkits for innovation. *Journal of product innovation management 18*, 4 (2001), 247–257.

Paper C

Tangible Sketching of Interactive Haptic Materials

Jonas Forsslund
School of Computer
Science and Communication
KTH Royal Institute of Technology
jofo02@kth.se

Ioanna Ioannou
Health and Biomedical
Informatics Research Unit
University of Melbourne
ioannoui@unimelb.edu.au

ABSTRACT

The activity of sketching can be highly beneficial when applied to the design of haptic material interaction. To illustrate this approach we created a design tool with a tangible hardware interface to facilitate the act of haptic material sketching and used this tool to design an anatomy exploration application. We found this approach particularly efficient in designing non-visual properties of haptic materials. The design tool enabled instant tactile perception of changes in material properties combined with the ability to make on-the-fly adjustments, thus creating a sense of pliability.

Author Keywords

Sketching, Design, Haptics, Interaction Design, Tangible Interfaces

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*Haptic I/O*; D.2.2 Software Engineering: Design Tools and Techniques—*User interface*

INTRODUCTION

Haptic applications enable users to feel virtual objects using a haptic device such as the Sensable Phantom Omni (Figure 1). Shape, position and hardness of virtual objects is conveyed through a combination of haptic and visual rendering. Haptic rendering algorithms such as that presented by Agus et al. [1] enable virtual materials like bone and teeth to be felt and manipulated with a virtual drill. The virtual drill itself has material properties which affect the tactile feedback experienced by users as they drill. The non-visible nature of haptic material properties creates a challenge in designing haptic interactions, as designers cannot see the effects of these properties, they can only feel them.

The design of haptic materials is commonly an iterative process, where designers make adjustments and proceed to interact with materials to feel the effects of those adjustments.

Figure 1. Designer using a PHANToM Omni device and MIDI controller interface to sketch the material properties of a dental anatomy exploration application

To assist in this process, some applications feature a graphical user interface with sliders and other components which allow users to adjust haptic rendering parameters (Figure 2). The utility of these interfaces is limited by the time and effort required to implement them and the fact that in order to operate them, users must interrupt their manipulation of haptic materials and shift their focus to GUI components, thus interrupting their workflow.

There exists a need for flexible and efficient design platforms that facilitate the exploration of solutions with satisfying haptic qualities [2]. The activity of sketching has been used in some haptic applications as a way to explore a design space. Miao et al. [7] used paper prototyping to evaluate tactile interfaces for the visually impaired. De Felice et al. [3] created an authoring tool where a virtual world could be designed interactively. While these studies propose creation and usage of sketches as a way to explore the design space and test ideas, they do not explore the design of the haptic interaction per se. We believe that sketching can be an effective and efficient approach to interactive haptic material design.

Sketching as a design activity plays two valuable roles. The more obvious role is in generating quick, cheap, throw-away prototypes suitable for exploring ideas in early stages of de-

TEI 2012, February 19 - 22, 2012, Kingston, Ontario, Canada.

velopment. The second important role of sketching is the act of sketching itself. When sketching the (graphical) designer does not merely draw a mental picture on paper (or other material), they receive immediate visual feedback and the sketch itself evolves. The sketch "talks back" to the designer [4]. We intend to bring the same level of interactivity to haptic material sketching. Schön's concept of the designer having a conversation with his/her sketch has been applied to digital materials [4] and has transcended into the design of tailored haptic feedback devices, an approach known as "sketching in hardware" [8].

Our approach brings the concept of sketching to the design of haptic applications by allowing the designer to create and explore various sketches of haptic material properties, based on run-time parameter adjustments using a tangible interaction device. We introduce a suitable sketching tool and illustrate its usage with a design study of a dental anatomy exploration application. The primary focus of this paper is not limited to sketching as a means of generating early prototypes or even sketching with physical materials, but the possibility of sketching the haptic material itself.

DENTAL ANATOMY EXPLORATION APPLICATION

Surgical simulation is a common application area for haptics research. Surgical simulators aim to teach theoretical aspects of a surgical procedure and improve trainees' technical skills through practice. Designing and implementing a stable, satisfactory realistic simulation remains a challenge both in terms of hardware, as well as algorithms and application design. Challenges include the elimination of device vibrations and calculation of haptic feedback forces with less than 1 ms delay. Different anatomical materials need to be designed with distinct physical properties which affect the tactile feedback experienced by users. These challenges make surgical simulation a suitably challenging design case for our study of haptic material sketching.

To demonstrate the use of sketching in designing haptic material interaction, we present our experience in developing a tool to facilitate the act of sketching, and using this tool to design an interactive dental anatomy exploration application with haptic feedback. The aim of this application is to help students make the transition from theoretical learning to hands-on practice by enabling them to interact with an anatomically accurate jaw model using their tactile sense.

The objective of the sketching approach in this case is twofold. First, to generate interaction ideas and identify sets of parameters which yield an interesting and educational user experience. Second, and at least equally important, to enable the designer to sketch the haptic material itself, i.e. the material properties of the jaw bone, teeth and drill. Materials do not necessarily have to be realistic for this application; in fact it could be useful to present exaggerated material differences similar to the way medical illustrations exaggerate anatomical colors.

Process overview

Our approach was comprised of the following steps:

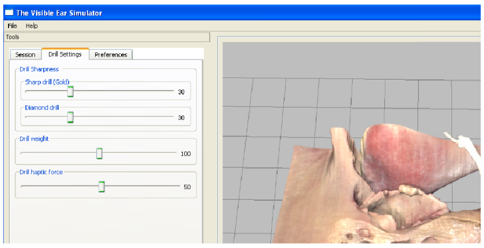

Drill Sharpness
Sharp drill (Gold) 30
Diamond drill 30
Drill weight 100
Drill haptic force 50

Figure 2. Part of screenshot showing a conventional user interface used to tune the look and feel of a temporal bone surgery simulator [9]

1. *Parameter choice:* We used domain knowledge gained from the literature and from interactions with domain experts to choose a set of parameters that manipulate haptic and visual aspects of the application.
2. *System implementation:* We designed and implemented a sketching tool such that the parameters selected in step 1 were easily adjustable. This was done by attaching the parameters to an interactive interface, making them modifiable at run-time.
3. *Sketch generation:* We used the sketching tool to interactively adjust parameter values and produce a set of sketches with desirable properties, such as high haptic stability or highly accurate representation of haptic material boundaries. In addition, we used this process to evaluate the parameter space and make adjustments such as adding / removing parameters and modifying parameter ranges. The set of sketches produced by this step was the haptic application equivalent of the set of design mockups that graphic designers commonly present to their clients.

Choosing design parameters

Choosing an appropriate set of parameters to vary during sketching requires insight in the form of technical knowledge of application algorithms as well as knowledge of the application domain. In our case, both authors have prior experience in the design and implementation of haptic dental training simulators and have worked extensively with domain experts. This experience was of great assistance in constructing an appropriate set of parameters.

Given the exploratory nature of the task, it is unreasonable to expect that the best set of parameters will be chosen from the very beginning, therefore we used an iterative approach. Parameter sets were chosen, implemented in our sketching tool, and evaluated during a sketching session. Any parameters found to be ineffective or with an inappropriate value range were modified in the next iteration until we reached an ideal set. To facilitate this iterative process, we designed our sketching tool such that parameters could be modified with relative ease.

The final set of sketch parameters is shown in Table 1. The chosen parameters alter both the look and feel of the simulation. Some parameters were chosen from the beginning,

Parameter	Description
Burr size	The diameter of the cutting burr. Range: 0.2 mm - 20.0 mm
Jaw size	The size of the jaw model as a ratio of the original size. Range: 0.5 - 6.8
Haptic stiffness	The overall stiffness presented by the haptic device when touching an object. Range: 0.5 - 3.7
Bone and tooth cutting rate	A number representing the rate at which each material is removed during every time cycle. The smaller the rate, the less material is removed each cycle and vice versa. Range: 0 - 1
Bone and tooth average cutting force	A value (in newtons) representing the force that the user is expected to apply when cutting each material. The higher the value, the more force the user will have to apply to remove material. Range: 0 N - 3 N
Bone and tooth transparency	The transparency percentage of each material. Range: 0% - 100%
Bone and tooth color	The HSV color value of each material. The Hue, Saturation and Value of a material color are each represented by one MIDI control, thus requiring three controls per material color. Range: 0 - 1

Table 1. Description of chosen haptic application parameters

while others were added as the need arose. The range for each parameter was adjusted iteratively during the sketching process.

Parameters such as burr size and jaw size were chosen based on knowledge of the chosen haptic device's position resolution and stability characteristics. Past experience has shown that there is a trade-off between realistic jaw size and haptic rendering fidelity. The size of a tooth is relatively small compared to the workspace of the most common haptic feedback device, the PHANToM Omni. The smaller the tooth size, the smaller the motions that will be made during drilling. As motion size approaches the device's position resolution, we begin to lose haptic rendering accuracy. With a larger jaw model, we use more of the device's workspace, which results in increased haptic fidelity.

Other haptic parameters were chosen to enable detailed fine-tuning of the way in which haptic feedback is rendered by the device. Haptic stiffness was chosen because we know that there is a limit to the stiffness our haptic device can render while maintaining overall stability. Therefore it is necessary to experimentally find the ideal stiffness, both during touching and drilling.

The bone and tooth cutting rate parameters were chosen to enable fine-tuning of the difference in how fast each material can be drilled. We found that a simple time-based cutting rate was not sufficient to fully represent the differences in material hardness, so we added the parameters representing the bone and tooth average cutting force. These parameters represent the force that is expected when cutting each material and can be calculated based on real life force measurements if realism is desired. The amount of material removed during drilling is increased or decreased based on how much force the user applies compared to the value of the parameter. The combination of the material cutting rate and average force parameters enabled detailed fine tuning of material hardness rendering during drilling. One of the differences between novice dentists and expert oral surgeons is their ability to differentiate material boundaries, particularly with differences in material hardness [5]. Thus it was important for our application to facilitate the exploration of material boundaries both visually and haptically. However, in contrast to surgical simulations, in our anatomical exploration application we do not seek to ground the hardness properties in physical attributes but in perception, just as medical illustrations do not necessarily use colors derived from nature.

Finally, we introduced parameters to vary the transparency and color of each material. We chose to vary these parameters in order to explore less realistic colors that help highlight material differences. Varying the transparency can also help understand the anatomical relationship of teeth and bone, such as how deep the teeth reach inside the bone.

Hardware and software

To allow for rapid creation of sketches we created the setup shown in Figure 1. Each rendering parameter is linked to a slider or knob of a USB-connected Behringer BCF2000 MIDI-controller such that the user of the system can interactively modify all parameters at run-time, while interacting with the Sensable PHANToM Omni haptic device. The open source project Forssim (http://dev.forsslundsystems.se) was used as a basis for the software. This system was chosen because it had most of the required functionality already implemented and it is built on the H3D API (http://www.h3d.org) which provides access to a wide range of haptic and visual parameters. The haptic rendering algorithm is a modified version of the Agus volume-sampling algorithm [1], which enables a direct rendering of the interaction of a spherical drill and a segmented volume model of bone and teeth, derived from Computed Tomography images. The system runs on a Linux-based PC (Intel Xeon 3.2 GHz CPU, 4GB RAM, nVidia Quadro 4000 graphics).

Creating design sketches

To develop a set of sketches for our anatomical exploration application, we ran the sketching tool and varied the parameters using the MIDI controller, until a pleasing result was achieved. Once a good parameter combination was found, it was saved as one of many pre-sets on the MIDI controller. Saved pre-sets could easily be recalled using the MIDI controller, which has motorized sliders that are automatically set to correct positions based on the pre-set being loaded.

DISCUSSION

The sketching tool and workflow presented above facilitated iterative, interactive sketching of haptic materials for a specific application. This experience with haptic sketching has been very positive. Although both authors have several years of experience in haptic simulation programming, we had not previously had the opportunity to manipulate the parameters of haptic rendering algorithms in such direct manner. The sketching process provided an intuitive understanding of the effects of parameter variations on haptic rendering which we had not experienced with previous approaches. We were able to quickly and efficiently sweep the application parameter space to identify the parameter values that produce a good educational experience.

The MIDI controller proved to be a highly effective user interface for our sketching tool. The use of a tangible interface enabled us to manipulate haptic materials without having to look away and stop drilling. Multiple parameters could be changed at once, using two or more fingers, and feedback was instantly received both visually and proprioceptively. These features of the design tool created a strong sense of pliability [6].

Furthermore, the use of a tangible interface eased the process of creating the sketching tool itself by eliminating the cumbersome effort of having to link simulation parameters to a graphical interface. Additionally, the MIDI controller's ability to save pre-sets of parameters provided an easy way to store several different sketches of the application without additional programming effort.

That said, it should be noted that creating the sketching tool was non-trivial. The choice of parameters, their ordering, as well as their mapping to sliders and knobs required prior knowledge of rendering algorithms and an understanding of the constraints they impose on the application. For this reason, the development of the design tool is best left to experienced simulation programmers. However, once the tool has been developed, it is quite possible to involve non-technical people in the sketching process, provided they are presented with a simplified and clearly labelled set of parameters that intuitively relate to what they are seeing on the screen and feeling through the haptic device. In particular, direct expert involvement could enable designers to capture tacit knowledge such as how a bone should feel when drilled. The sketching tool could allow both experts and end-users to evaluate a range of sketches to identify the most desirable ones and further fine-tune sketch parameters interactively. This is a novel concept in the design of haptic applications.

Generalizing the approach

Given the success of applying the sketching approach to the example presented here, there is merit in suggesting that such an approach could be helpful in other haptic application domains.

The steps of choosing a set of parameters, iteratively creating a sketching tool, and generating a set of design sketches can easily be applied to other domains. For example, following the success of the design study presented in this paper, one of the authors proceeded to apply the same approach to the design of a temporal bone surgical simulator. As in our example, a set of suitable application-specific parameters was chosen and linked to the MIDI controller, which enabled the creation of a set of sketches. Preliminary evaluation of these sketches by otolaryngology experts has yielded a highly positive response. Another possible application is haptic sculpting, where the sketching parameters could control the effects of various sculpting gestures.

CONCLUSION

We have introduced the idea of sketching as a means for designing haptic material interaction. The case study presented above illustrates how this approach can be applied successfully to provide unique insight into the design space of a haptic application and empower designers in creating desirable haptic interactions. Our experience in developing the sketching tool and using it to generate sketches was so positive that we have began applying this approach to other haptic projects, such as temporal bone surgical simulation. We see potential in applying the steps followed in this design study to sketch haptic material interaction in a wide range of applications.

REFERENCES

1. M. Agus, A. Giachetti, E. Gobbetti, G. Zanetti, and A. Zorcolo. Real-Time Haptic and Visual Simulation of Bone Dissection. *Presence: Teleoperators and Virtual Environments*, 12(1):110–122, Feb. 2003.
2. H. V. Bjelland and K. Tangeland. User-centered design proposals for prototyping haptic user interfaces. In *Proceedings of HAID07*, pages 110–120, Berlin, Heidelberg, 2007. Springer.
3. F. De Felice, G. Attolico, and A. Distante. Configurable design of multimodal non visual interfaces for 3d ve's. In *Haptic and Audio Interaction Design*, volume 5763 of *LNCS*, pages 71–80. Springer, 2009.
4. A. Dearden. Designing as a conversation with digital materials. *Design Studies*, 27(3):399 – 421, 2006.
5. I. Ioannou, L. Stern, E. Kazmierczak, A. Smith, and L. Wise. Towards defining dental drilling competence, part 2: a study of cues and factors in bone drilling. *Journal of Dental Education*, 74(9):941, 2010.
6. J. Löwgren. Toward an articulation of interaction esthetics. *New Review of Hypermedia and Multimedia*, 15(2):129–146, 2009.
7. M. Miao, W. Khlmann, M. Schiewe, and G. Weber. Tactile paper prototyping with blind subjects. In *Haptic and Audio Interaction Design*, volume 5763 of *LNCS*, pages 81–90. Springer, 2009.
8. C. Moussette and R. Banks. Designing through making: exploring the simple haptic design space. In *Proceedings of TEI11*, pages 279–282, New York, NY, USA, 2011. ACM.
9. M. S. S. Sorensen, J. Mosegaard, and P. Trier. The visible ear simulator: a public PC application for GPU-accelerated haptic 3D simulation of ear surgery based on the visible ear data. *Otology & Neurotology*, 30(4):484–487, June 2009.

Paper D

Designing the experience of Visuohaptic Carving

1st Author Name
Affiliation
City, Country
e-mail address

2nd Author Name
Affiliation
City, Country
e-mail address

3rd Author Name
Affiliation
City, Country
e-mail address

ABSTRACT

This paper introduces an interaction technique called visuohaptic carving, and a strategy for how it can be utilized concretely by designers. Visuohaptic carving is well suited for visualization of multi-layered computer graphics objects where the aim is to illustrate and explore spatial relationships between segments and structures within three-dimensional objects. Possible applications include anatomy exploration, cut-away illustrations and interactive art. Through our work with designing applications that uses visuohaptic carving, we have identified three key requirements as a strategy for making visuohaptic carving an effective design resource: 1) a ready-made but mutable implementation, 2) tools to directly form and tune the implementation in terms of scale, stiffness and carving rate, and 3) formulation of a work-flow practice.

Author Keywords

Haptics; interaction design; interaction technique; tools.

ACM Classification Keywords

H.5.2. User Interfaces Haptic I/O; H.5.2 User Interfaces prototyping; D.2.2 Design Tools and Techniques.

INTRODUCTION

Interactive 3D applications are undergoing a transformation from being highly specialized professional tools to becoming more of an everyday application for home use e.g. for creating content for 3D printing. One major challenge when it comes to interacting with on-screen 3D objects is that standard interaction devices, i.e. touch screen, mouse, keyboard and trackpad, are all essentially designed for interaction with the flat surfaces of standard 2D user interfaces. Further, these techniques do not generate much of haptic feedback, i.e. they do not provide a tactile sense of touch back in response to interaction with on-screen objects. In the research area of haptic interaction design and 3D user interfaces [22, 3, 29], a range of devices has been developed over the years, enabling tactile sensations, and thus more purposeful manipulation of materials and objects displayed on screen. The primary technologies that enable such visuohaptic manipulations are the force-reflecting haptic feedback device [23], coupled with haptic rendering algorithms [29] specialized for touching, manipulating and modifying complex computer graphics shapes [1, 24, 6, 27]. Like with 3D printing hardware, these devices have recently started to become more accessible due to e.g. open source starting kit versions that can be assembled and customized in DIY fashion [12]. This, together with the promises of the rich interaction qualities of interacting with the sense of touch, has led to a growing interest in studying haptic interaction design, and how to make this technology more accessible as a resource for the design of interactive systems [22, 25].

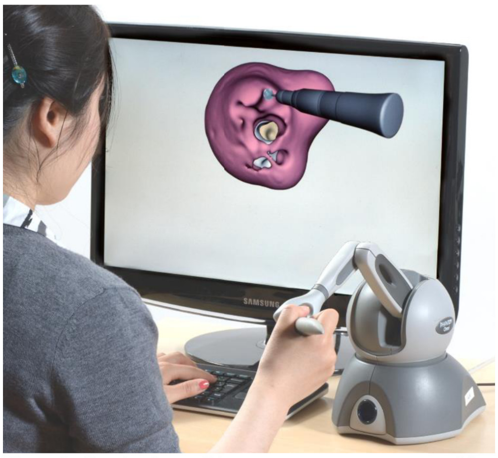

Figure 1. Carving in a solid, multi-layered computer graphics object with a virtual drill using force-reflecting haptic feedback in the process

In this paper we introduce an interaction technique that we have chosen to call *visuohaptic carving*, and a strategy for creating applications that lets designers work with this as a

Submitted for review to DIS 2016

Figure 2. Overview image and stills from the interactive art installation Immaterial Materiality and Virtual Structures by Martha Johansson.

design material. The visuohaptic carving technique lets users carve in computer graphics objects using a force-reflecting haptic device (figure 1), revealing inner structures and feeling resistance in the process. This carving technique is most commonly used in surgery simulation [1, 6], but as will be shown, visuohaptic carving has useful applications beyond that domain, including interactive art and different forms of interactive visualization settings.

How to actually go about developing interactive 3D applications that provide visuohaptic carving remains a challenge, especially from a perspective of design. While tactile experience is a natural part of designing with physical materials, the sensations that result from interacting with computerized haptics are less well understood. To address these challenges, we propose in this paper a strategy for designing the feeling of visuohaptic carving in a way that is open for different contexts and constraints. Instead of relying on representations such as low-fi prototypes for subsequent implementation by a highly specialized engineer the approach is intended to give the designer the freedom to explore the design space hands-on, or as put by Moussette and Banks [26], work *with and through the material*.

The contribution of this paper is a description of a set of design tools that integrates these in a professional 3D authoring ecosystem. Together with a reusable implementation in the form of a software library, they form a strategy towards designing visuohaptic carving applications, which we have found useful in our own practice. We have identified three key requirements as a strategy for making visuohaptic an effective design resource: 1) a ready-made but mutable implementation, 2) tools to directly form and tune the implementation in terms of scale, stiffness and carving rate, and 3) formulation of a work-flow practice.

BACKGROUND

From a computer graphics perspective, the possibility to deform 3D objects by removing material, similar to carving, is in itself an interesting research topic, which has been explored extensively through a range of studies and explorations since the late 1990's, with or without haptic feedback [14, 38, 2]. One of the now most commonly explored interaction devices designed specifically for three-dimensional manipulations is the force-reflecting haptic feedback device [23], which is also the focus in the present work. Force-reflecting haptic feedback means that as the user pushes against a surface of a 3D Computer Graphics (CG) object, the device pushes back with an equal opposite force, which enables a convincing experience of touching and tracing the surface of virtual objects. The experience of engaging with digital objects with such a device is therefore very much a physical experience, which is also a central aspect of form giving through carving or molding in ordinary physical space. The device has therefore been suggested as potentially useful in e.g. industrial design, although early studies (e.g. 32) indicated challenges both in terms of maturity of the technology and in its relationship to existing design practice. Instead the device is now more commonly used for visualisation and simulation, e.g. for dental and surgery education [28, 1]. One specific challenge in the design of haptic experiences for force-reflecting haptic feedback devices is the exact tuning of the sensed material properties of virtual objects. Apart from the visual appearance of on-screen objects and materials, computer-controlled haptic systems thereby have a second signal part that is subject to design. To address this challenge, designers have explored methods for gathering meaningful input from clients, for instance in initial requirement analysis in user-centered design processes. Kern [18] suggests that specialized haptics engineers ask clients to describe the intended tactile experience with reference to common items such as fruits, springs and various materials. Another approach is to measure the forces and motion of a real action and use those numerical values as guides for the design [16]. While the requirements specification approach is usable in some situations, it may also hinder the designer from exploring the range of possible synthetic haptic experiences, or as what is sometimes referred to as a "conversation with the material" [31]. To approach haptics as a specific design material in need of exploration, Moussette [26, 25], showed how using accessible components such as motors, wood and straight-forward programming platforms opens up a design space where designer can directly experience what kind of haptic interactions that are achievable, and to build up a certain sensibility for the material [25].

	Pulp	Jaw	Dentin	Enamel	Filling
Ambient intensity %	20	36	33	0	27
Shininess %	16	9	9	9	9
Trans-parency %	0	0	0	0	0
Hardness (lower is harder)	0,20	0,60	0,20	0,15	2,00

		cut	visual	
Spring K	700,00			
bit Radius	1,00	1,00	1,50	mm
offset Factor	0,75			

```
<?xml version='1.0' encoding='utf-8'?>
<X3D profile="H3DAPI" version="2.1">
 <Scene>
  <Collision enabled="FALSE">
   <ImportLibrary library="&quot;forssim.dll&quot; &quot;libforssim.so&quot;"/>
   <PythonScript DEF="keyHandler" url="keyHandler.py"/>
   <KeySensor DEF="KS"/>
   <ROUTE toField="keyHandler" fromNode="KS" toNode="keyHandler" fromField="keyPress"/>
   <Background skyColor="0.16078 0.36471 0.23137"/>
```

Figure 3. **A) A surgeon teacher illustrating a surgical technique to medical students in an anatomy lecture. B) Mannequin used for hand support and for reference and context of dentistry. C) Screenshot from tool for setting and tuning the parameters modulating the feeling of carving, for a particular CG object. D) Midi-controller interface for tuning haptic sensation.**

These and similar explorations around tools and kits [34, 19, 10, 30, 12], suggest an approach based on design-through making [26], which let designers sketch, tune and feel different expressions using haptic technology before committing to formulate strict requirements for a particular system. Yet, how to concretely approach visual and haptic experiences together in a design process is still a major challenge. For vibrotactile feedback it is as least, if not more, interesting to design how something vibrates as when it should vibrate. Different signals with various patterns, rhythms and frequencies have been used to convey semiotic meanings [22]. Software tools have been developed where signals can be improvised, designed and experienced interactively [34, 19, 30]. One tool was developed specifically for sketching haptic material properties for use in surgery simulation [10]. The present strategy will build upon this work, in that a tuning tool play an essential, but is not the only, role in an overall strategy for designing the user experience of visuohaptic carving.

THE INTERACTION TECHNIQUE

In general, carving can be defined as the act of using tools to shape something from a material by scraping away portions of that material. The technique can be applied to any material that is solid enough to hold a form even when pieces have been removed from it, and yet soft enough for portions to be scraped away with available tools [40]. In our work the aim is not to simulate any particular carving tools, but enabling users to scraping away portion of the synthetic haptic material to reveal inner structures and gain a better understanding of shape and relative location of a computer graphics objects' internal features. This can be useful in several scenarios, and three will be presented in the following. Together the examples serve to show the versatility of visuohaptic carving.

Example one: surgery simulation

Our first example concerns the use of visuohaptic carving as an educational resource in a specific case of dentist education. The use value of students in this case being able to practice interactively using haptic feedback means that they can develop a sense of the phenomenological differences between drilling and in other ways manipulating different materials and tissues. When trilling in a tooth, it is essential for students to learn that there will be a felt difference in sensation when drilling through the enamel, dentin or bone, and approximately what that will be. Being able to see, along with the haptic feedback, is then an

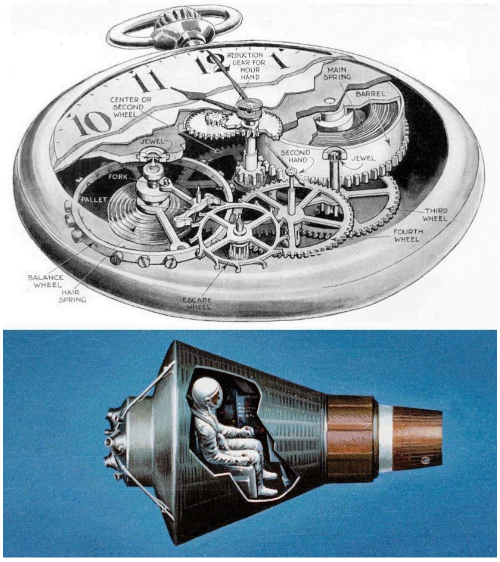

Figure 4. Examples of cut-away illustrations.Top: Pocketwatch by B.G. Seilstad. Bottom: Mecury spacecraft by D. Metzer. Images in public domain.

important added educational value, although in real surgery you may often be left to the tactile experience alone.

Another important insight from this particular use case, was that feedback from numerous user studies and review of other simulators [37, 28, 8, 35, 17] indicated that one really important aspect of the simulator design is the gestalt of different patient cases that represent anatomical and pathological differences among individuals and how they should be treated. This is potentially more important to get right than striving for uttermost realistic graphics, haptics and physics. This can be achieved if attention is paid to the crafting and tuning of the virtual scene, its scenario, CG Objects and the haptic interaction, regarding stiffness, scale and carving rate. When working with this case, we, as interaction designers, realized that the professional surgeons and dentists needed to be actively involved in shaping the haptic experience of such cases. Tuning and carving are therefore two aspects of the interaction that needs testing together, in close loops.

Example two: expressive carving

Our second example case for visuohaptic carving concerns more creative or expressive uses, e.g. for sculpting of objects for 3D printing, or for artistic uses of various kinds.

We have provided visuohaptic carving as a design resource to an external artist for use in an interactive art installation (figure 2). The purpose of the art piece was to have fascinating experiences related to what is real and virtual and what is hidden within abstract objects. For example one could cut into the model of a brain and find a small baby model. The artistic value of this is out of scope for the paper, but in the development work itself we had to prepare the material and *tune it to fit the device* (Phantom Omni) which meant some limitations to resolution and stiffness. It is worth noting that while only carving algorithms where used, different tools where visually represented giving different subjective experiences; for instance compare carving a yellow block with a clay tool to cutting a heart with a scalpel.

Example three: Cut-away visualisation

Cut-away drawings (figure 4) are a well-known illustration technique for presenting inner structures of an object while providing a surrounding context, and has been adopted for computer graphics objects [20]. Visuohaptic carving could be appropriated for both creation of cut-away illustrations and applications where the end-user can freely explore an object through cutaway operations. Example applications of such visuohaptic carving could be illustration of anatomy for surgery theory.

An example of this use of visuohaptic carving occurred spontaneously in one of our studies when we observed a surgeon using a simulator originally designed for pre-operative planning, but this time used for teaching anatomy. In this setting, students were sitting in a lecture room, wearing 3D glasses and following the teacher`s display projected on a large screen (figure 3A). The teacher then used the haptic device to control a virtual drill as in a surgical rehearsal scenario. However, instead of using it as a single-user simulator for patient specific surgery rehearsal, which it was designed for, he used the tool in a creative way, by selecting a very large drill and scooping away large parts of the scull to provide a good view for his students. In this, he was not simulating something that he would ever do in reality, but rather used visuohaptic carving for the interaction and selection of presentation material. The visuohaptic experience was an essential part of being able to do this effectively, although the haptic part of the experience was never shared or presented to the students.

STRATEGY FOR CREATING APPLICATIONS

The strategy we propose for creating applications that features visuohaptic carving presupposes access to one or more spatial haptic devices (e.g. the Phantom Omni pictured in figure 1), a standard or stereoscopic monitor, a specific software platform, and specific design tools. In this section the strategy will be illustrated with the platform we used and extended, and the prototype tools we developed together with a specific workflow that integrates with a professional 3D object authoring ecosystem.

The software platform used is H3D API (http://www.h3dapi.org) used with our extension library named *forssim* (http://dev.forsslundsystems.se), are both open source. Forssim implements a set of necessary algorithms for achieving visuohaptic carving, namely 1) a haptic algorithm that reads the position of the device, computes contacts between a computer graphics object and a virtual tool, and returns a force to the device [11, 7], 2) a carving algorithm that deforms the object when an activated tool is in contact, and 3) a visual rendering algorithm that continuously renders an image of the object as it is being carved [39]. H3D API comes with a light-weight executable called H3DLoad that can load an xml-based document that describe the layout and contents of virtual scene in a similar fashion as a html document describes a web page [9]. The scene document describes the initial spatial layout of the object subject to carving, the tool the user carves with, and the visual and haptic parameters modulating the sensation of carving. The document also host high-level scripting with which events can be programmed, for example masking out certain predefined segments when sufficient amount of material have been removed in another specified area. This is for instance used in an oral surgery simulator to gestalt extraction of a tooth when sufficient bone has been removed surrounding the tooth.

The computer graphics object needs to be in a format supported by the algorithms. The most common representation of 3D graphics objects is with a polygon mesh and thus only the surface is modeled. In comparison, the objects for carving should support non-homogeneous solid representations. This means that inner structures can be modeled, e.g. a tooth could be modeled with enamel, dentin and pulp, or the inner layers as seen in figure 1. Technically the models are usually represented with rectilinear 3D grid of samples values also known as a voxel volume. As the object is going to be explored with a physical device which has absolute measurements, not relative such as the mouse, it needs to be defined in physical units (e.g. centimeters).

The tool consists of a hardware part referred to as a manipulandum, and a screen-based representation for visual and haptic rendering respectively. The physical form of the manipulandum, the handle the user holds on to, varies between haptic devices. Different forms and materials can be used, especially if a custom-made device is constructed. The digital representation for visual rendering can be an arbitrary CG object, but for haptic rendering a simpler representation is commonly required. In our work the tool is represented by a rotational invariant sphere which means that the orientation and shape of the shaft is ignored for haptic feedback and carving. With more complex algorithms the tool could be modeled as a point cloud covering the surface of the whole tool [24, 6]. As with the object subject to carving, the virtual carving tool needs to have a physical measure.

The parameters modulating the feeling of carving are parameters of the haptic and carving algorithms, the CG object and the virtual drill, that can be altered to get different behavior and user experience. The premier parameters used to modulate the haptic experience in this project are scale, stiffness and carving rate.

When actively carving, the experience of carving is not a crisp direct contact between tool and object but a virtual spring contact. The stiffness of this contact is definable by the designer but is dependent on which device used. Set it too high and it will cause vibrations. Set it too low and the experience will be of touching something very mushy. The spring is stretched until maximum force can be given from the device i.e. the motors are saturated. Friction in the device will also mask out some forces which are then not perceived fully. As a consequence, small details can be difficult to feel using some devices, especially cheaper ones. The designer is therefore encouraged to scale object to get a high quality feeling with good perception of details and no vibrations. Second, the carving rate can be adjusted as a good way to experience different layers of the CG object as being of different material hardness. Again, it is useful to be able to modify these properties in real-time since it is only when you experience the feeling of touching the results, that you can understand the effects of the changes made. Therefore a tool for tuning these properties (figure 3C and 3D) was developed.

Exploring the larger design space

When a fully functional visuohaptic carving application has been created, including relevant CG objects, it is possible to start exploring the larger design space of how the carving feels as a result of switching or altering components. For example, using different haptic devices or working hands on with the components internal to a haptic device such as with the WoodenHaptics device [12], will change the feeling especially how stiff the CG object can be rendered. It is also worth exploring hardware in combination with different sizes and resolutions of CG objects, as voxels of the magnitude of 1 mm can be felt with a decent device. Furthermore, different haptic rendering algorithms can be explored, given that they are implemented as easily replaceable components. The FORSSIM library implements two haptic algorithms. The first is a penalty-based volume-sampling algorithm inspired by Agus et al [1]. It gives a rather smooth feeling even over rugged surfaces, but has the downside that the user risks falling through the solid object if more than half of the tool sphere is penetrating the surface. The second algorithm, which is based on Chan et al [6], is constraints-based and thus always keep the tool on the top of the surface. However, in our implementation and with rugged binary sampled models some "nervousness" can be experienced. One is not clearly better than the other, but depends on use context, and hardware, why experimentation is required to arrive at a good combination. Application designers also need to take into consideration the environment where the visuohaptic carving takes place. Hand support is often important from an ergonomic perspective, and surrounding props can also be used as constraints or give context to the procedure. In figure 3B a natural sized mannequin is placed in the workspace of the haptic device. This is good for hand-support and to provide the user with a frame of reference, but it also limits the scale for which teeth can take if rendered inside the mouth of the mannequin. This is one example of the kind of trade-offs the application designer has to consider.

DISCUSSION

Early in the project, before the development of tools for tuning and testing the haptic experience, the content authoring process was tedious and unreliable. CG objects for a dental simulator were modeled from computed tomography images previously segmented manually, slice by slice. Then the color and carving rate for these structures where entered as numbers in a text file. The only way to design was to guess appropriate numbers, start the software, try it, close the software, adjust and restart. The painting was also time consuming and the 3D result was unsmooth, which also meant we had to do filtering which in turn made e.g. small structure disappear. The strategy we used to address this situation involves incorporation of a range of professional tools, and 3D artists, in combination with custom made editing and tuning.

Interactive applications that involve what we here refer to as visuohaptic carving, have previously been developed predominantly in surgery simulation [1, 24, 6, 27] or virtual sculpting [21]. The main contributions of those works have lied in the advancement of the state of the art in the algorithms that make such applications possible, as well as in particular useful design cases. In this paper, our aim was to extend this work first by articulating visuohaptic carving as a concept that enables discussion, and reflection, among designers who may engage with this and similar technologies [15]. Secondly, we have articulated what we see as core findings from our explorations, which we hope will enable further developments and also facilitate designers to bring this technology into their practices.

We have shown, that the main concept of visuohaptic carving is feasible with existing technologies, and that tools that enable designers to engage experientially with these technologies is a critical element for allowing effective form giving within this domain.

Design practice and judgment is both about getting the design right and getting the right design [5]. Applied to visuohaptic carving, this means getting the touch experience just right with careful adjustment of nuances analogous to how a graphical designer would tweak colors, and to find out where, when and how this material is a suitable part of a solution or other options are better.

Vallgårda and Redström [36] discuss Computational Composites, which we find suits this haptic technology very well: both hardware and software qualities are essential parts of the experience. For stiffness for instance, there is a clear interplay between on one hand the qualities of the motors and structure and the other hand the stiffness set in software. If a penalty-based algorithm is used there is risk for falling through the object when more than half the spherical avatar is penetrating the surface of the object. If both the manipulated graphic object and the tool are small, very high stiffness need to be set so that the force become large enough "within time" before this happens (the force is proportional to penetration depth). However, the more affordable devices with weak motors cannot provide the stiffness or force required resulting in vibrations or unstable behaviour. There is thus a clear interplay between the computational and physical elements of the "material".

To understand early on what possibilities the medium offers to a designer, creation of so called "inspirational bits" can be a useful approach [33]. Furthermore, as the designer become more acquainted with the design material, in this case the haptic technology, certain sensitivity to its specific properties is expected to be developed. The medium specificity hypothesis [13] suggests that each medium has its own optimal use, or rather that an artist/author should choose the medium to work with/in that most suits the intention. It is expected that visuohaptic carving has unique possibilities that are worth exploring further.

CONCLUSION

In this paper the concept of visuohaptic carving has been articulated, grounded in our work on a working dental simulator and an art installation, as well as elaborating on the potential application of cut-away illustrations. Realizing appropriate haptic sensations in a system that implements visuohaptic carving depends on several factors including the structural qualities of the hardware, rendering algorithms and geometric properties of the graphical object subject to carving. The fact that the felt sensation of carving is dependent on these factors does not mean that there is one single optimal feeling that can be engineered for. Instead, visuohaptic carving can be realized in multiple ways, with different technologies depending on design. Since the sensations can only be experienced in a completed system, it is difficult to predict how it will feel. External constraints such as hardware budget, the surrounding physical environment, and transportability of the system, need to be considered as well. Active hands-on experimentation is therefore needed in order to work with visuohaptic carving from a design perspective, in how to arrive at appropriate somaesthetic qualities.

REFERENCES

1. M. Agus, A. Giachetti, E. Gobbetti, G. Zanetti, and A. Zorcolo. Real-Time Haptic and Visual Simulation of Bone Dissection. Presence: Teleoperators and Virtual Environments, 12(1):110–122, Feb. 2003.

2. R. S. Avila and L. M. Sobierajski. A haptic interaction method for volume visualization. In Visualization'96. Proceedings., pages 197–204. IEEE, 1996.

3. D. A. Bowman, E. Kruijff, J. J. LaViola Jr, and I. Poupyrev. 3D user interfaces: theory and practice. Addison-Wesley, 2004.

4. D. Brutzman and L. Daly. X3D: extensible 3D graphics for Web authors. Morgan Kaufmann, 2007.

5. B. Buxton. Sketching User Experiences: Getting the Design Right and the Right Design: Getting the Design Right and the Right Design. Morgan Kaufmann, 2010.

6. S. Chan, F. Conti, N. H. Blevins, and K. Salisbury. Constraint-based six degree-of-freedom haptic rendering of volume-embedded isosurfaces. In World Haptics Conference (WHC), 2011 IEEE, pages 89–94. IEEE, 2011.

7. S. Chan, F. Conti, K. Salisbury, and N. H. Blevins. Virtual reality simulation in neurosurgery: technologies and evolution. Neurosurgery, 72:A154–A164, 2013.

8. J. Cormier, D. Pasco, C. Syllebranque, and R. Querrec. Virteasy a haptic simulator for dental education. In The 6th International Conference on Virtual Learning, volume 156, pages 61–68, 2011.

9. L. Daly and D. Brutzman. X3d: extensible 3d graphics standard. IEEE Signal Processing Magazine, 24(6):130, 2007.

10. J. Forsslund and I. Ioannou. Tangible sketching of interactive haptic materials. In Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction, pages 111–114. ACM, 2012.

11. J. Forsslund, E.-L. Salinas, and K.-J. Palmerius. A user-centered designed foss implementation of bone surgery simulations. In EuroHaptics conference, 2009 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009. Third Joint, pages 391–392, 2009.

12. J. Forsslund, M. Yip, and E.-L. Sallnäs. Woodenhaptics: A starting kit for crafting force-reflecting spatial haptic devices. In Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction, pages 133–140. ACM, 2015.

13. V. Fuchsberger, M. Murer, and M. Tscheligi. Materials, materiality, and media. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pages 2853–2862. ACM, 2013.

14. T. A. Galyean and J. F. Hughes. Sculpting: An interactive volumetric modeling technique. In ACM SIGGRAPH Computer Graphics, volume 25, pages 267–274. ACM, 1991.

15. K. Höök and J. Löwgren. Strong concepts: Intermediate-level knowledge in interaction design research. ACM Transactions on Computer-Human Interaction (TOCHI), 19(3):23, 2012.

16. I. Ioannou, E. Kazmierczak, L. Stern, A. C. Smith, L. Z. Wise, and B. Field. Towards defining dental drilling competence, part 1: a study of bone drilling technique. Journal of dental education, 74(9):931–940, 2010.

17. D. Joseph, J.-P. Jehl, P. Maureira, C. Perrenot,

N. Miller, P. Bravetti, P. Ambrosini, and N. Tran. Relative contribution of haptic technology to assessment and training in implantology. BioMed research international, 2014, 2014.

18. T. A. Kern. Engineering haptic devices. Springer, 2009.

19. D. Ledo, M. A. Nacenta, N. Marquardt, S. Boring, and S. Greenberg. The haptictouch toolkit: enabling exploration of haptic interactions. In Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction, pages 115–122. ACM, 2012.

20. W. Li, L. Ritter, M. Agrawala, B. Curless, and D. Salesin. Interactive cutaway illustrations of complex 3d models. In ACM Transactions on Graphics (TOG), volume 26, page 31. ACM, 2007.

21. Y.-H. Lu, W.-T. Wang, R.-H. Liang, and M. Ouhyoung. Virtual sculptor: A feature preserving haptic modeling system. In Proc. of ACM International Workshop on Immersive Telepresence (ITP2002), 2002.

22. K. E. MacLean. Designing with haptic feedback. In Robotics and Automation, 2000. Proceedings. ICRA'00. IEEE International Conference on, volume 1, pages 783–788. IEEE, 2000.

23. T. H. Massie and J. K. Salisbury. The phantom haptic interface: A device for probing virtual objects. In Proc. Haptic Symposium 1994, volume 55, pages 295–300. Chicago, IL, 1994.

24. W. A. McNeely, K. D. Puterbaugh, and J. J. Troy. Six degree-of-freedom haptic rendering using voxel sampling. In ACM SIGGRAPH 2005 Courses, page 42. ACM, 2005.

25. C. Moussette. Simple haptics: Sketching perspectives for the design of haptic interactions. 2012.

26. C. Moussette and R. Banks. Designing through making: exploring the simple haptic design space. In Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction, pages 279–282. ACM, 2011.

27. G. T. O'Neill, J. William, and W.-S. Lee. Haptic-Based 3D Carving Simulator. INTECH Open Access Publisher, 2010.

28. P. Pohlenz, A. Gröbe, A. Petersik, N. Von Sternberg, B. Pflesser, A. Pommert, K.-H. Höhne, U. Tiede, I. Springer, and M. Heiland. Virtual dental surgery as a new educational tool in dental school. Journal of Cranio-Maxillofacial Surgery, 38(8):560–564, 2010.

29. K. Salisbury, F. Conti, and F. Barbagli. Haptic rendering: introductory concepts. Computer Graphics and Applications, IEEE, 24(2):24–32, 2004.

30. O. S. Schneider and K. E. MacLean. Improvising design with a haptic instrument. In Haptics Symposium (HAPTICS), 2014 IEEE, pages 327–332. IEEE, 2014.

31. D. Schön and J. Bennett. Reflective conversation with materials. In Bringing design to software, pages171–189. ACM, 1996.

32. B. Sener, P. Wormald, and I. Campbell. Evaluating a haptic modelling system with industrial designers. In proceedings of the EuroHaptics International Conference, pages 8–10, 2002.

33. P. Sundström, A. Taylor, K. Grufberg, N. Wirström, J. Solsona Belenguer, and M. Lunde´n. Inspirational bits: towards a shared understanding of the digital material. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pages 1561–1570. ACM, 2011.

34. C. Swindells, E. Maksakov, K. E. MacLean, and V. Chung. The role of prototyping tools for haptic behavior design. In Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2006 14th Symposium on, pages 161–168. IEEE, 2006.

35. B. Tse, W. Harwin, A. Barrow, B. Quinn, M. Cox, et al. Design and development of a haptic dental training system-haptel. In Haptics: Generating and Perceiving Tangible Sensations, pages 101–108. Springer, 2010.

36. A. Vallga˚rda and J. Redström. Computational composites. In Proceedings of the SIGCHI conference on Human factors in computing systems, pages 513–522. ACM, 2007.

37. N. Von Sternberg, M. Bartsch, A. Petersik, J. Wiltfang, W. Sibbersen, T. Grindel, U. Tiede, P. Warnke, M. Heiland, P. Russo, et al. Learning by doing virtually. International journal of oral and maxillofacial surgery, 36(5):386–390, 2007.

38. S. W. Wang and A. E. Kaufman. Volume sculpting. In Proceedings of the 1995 symposium on Interactive 3D graphics, pages 151–ff. ACM, 1995.

39. S. Wijewickrema, I. Ioannou, and G. Kennedy. Adaptation of marching cubes for the simulation of material removal from segmented volume data. In Computer-Based Medical Systems (CBMS), 2013 IEEE 26th International Symposium on, pages 29–34. IEEE, 2013.

40. Wikipedia. Carving — Wikipedia, the free encyclopedia, 2016. [Online; retrieved 17-January-2016].

Paper E

The Effect of Haptic Degrees of Freedom on Task Performance in Virtual Surgical Environments

Jonas FORSSLUND [a,1], Sonny CHAN [a,1], Joshua SELESNICK [b],
Kenneth SALISBURY [a,c], Rebeka G. SILVA [d], and Nikolas H. BLEVINS [b]
[a] *Computer Science Department, Stanford University*
[b] *Department of Otolaryngology, Stanford University*
[c] *Department of Surgery, Stanford University*
[d] *Oral & Maxillofacial Surgery Section, VA Medical Center, San Francisco*

Abstract. Force and touch feedback, or haptics, can play a significant role in the realism of virtual reality surgical simulation. While it is accepted that simulators providing haptic feedback often outperform those that do not, little is known about the degree of haptic fidelity required to achieve simulation objectives. This article evaluates the effect that employing haptic rendering with different degrees of freedom (DOF) has on task performance in a virtual environment. Results show that 6-DOF haptic rendering significantly improves task performance over 3-DOF haptic rendering, even if computed torques are not displayed to the user. No significant difference could be observed between under-actuated (force only) and fully-actuated 6-DOF feedback in two surgically-motivated tasks.

Keywords. surgical simulation, haptics, haptic rendering, task performance

1. Introduction

What degree of haptic fidelity must a surgical simulator have in order to optimally achieve its objective? The inclusion of force and touch feedback, or haptics, plays a significant role in the realism of many virtual reality surgical simulations. Research in novel haptic interfaces and force rendering algorithms has continued to enhance the fidelity of instrument control and manipulation in surgical simulators. While it is clear that sophisticated devices and rendering techniques can deliver a more realistic experience, they may do so at prohibitive financial or computational expense. Additional effort is still required to improve our understanding of the impact of haptic fidelity on the efficacy of virtual reality simulators [1,2]. In the present work, we specifically examine consequences for task performance of using different numbers of degrees of freedom (DOF) of force feedback.

Laparoscopic surgery simulators are currently the most mature application of virtual reality surgical simulation, and this specialty appears to be the one for which the role

---

[1] These authors contributed equally to this work. Correspondence to: Stanford BioRobotics Laboratory, 318 Campus Drive, Stanford, CA 94305; E-mail: sonny@cs.stanford.edu

of haptic feedback has been rigorously evaluated. Studies have shown that haptic feedback improves performance on laparoscopic tasks in a virtual environment [3] and has a positive effect on skills training [1], especially for surgical tasks in which forces play an important role (eg. stretching, grasping, cutting) [4,5].

Little is known about the effect of the quality of haptic feedback on performance. In fact, it is possible that haptic feedback can even be detrimental. For example, learning surgical practices with an unrealistic model can allow the surgeon-in-training to use techniques that would be impossible or even dangerous in real surgery [6], which may even lead to negative training transfer [4,5]. Intuitively, the overall success of a simulator is dependent on how well the haptic feedback reflects relevant real forces experienced by the surgeon while performing the surgical task [2,7].

2. Background

Several authors have begun to investigate the effect of haptic feedback fidelity in various applications. Kim et al. [5] tested two force response models of different accuracy for material elasticity in a laparoscopic surgery simulation. Although they observed that results with an approximate model were similar to the high-fidelity model, laparoscopic surgery involves manipulating constrained instruments that limit the surgeon's haptic sense [2], which makes it more difficult to perceive small differences in force. Wagner et al. performed an experiment to demonstrate the effect of varying degrees of force feedback for a blunt dissection task using a tele-operated robot [8].

Exploring a complex environment using a rigid instrument is a 6-DOF interaction involving both forces and torques. Wang and Srinivasan first attempted to characterize the role of torque feedback on a subject's ability to determine a virtual object's distance through making contact with a long, thin rod [9]. Verner and Okamura designed a simple tracing and drawing task where the subject used a virtual pencil with varying combinations of force and torque feedback [10]. They found that, for such a task, force feedback significantly improved user performance, but the addition of torque feedback did not yield significant improvement over forces alone. Weller and Zachmann showed that 6-DOF haptic devices outperformed their 3-DOF counterparts in terms of intuitiveness of control and quality of force feedback in a competitive object collection game [11].

Well-known algorithms for haptic rendering are primarily 3-DOF in that they compute output forces based on device position only, and have no concept of orientation or torque. They permit haptic interaction only through a point or a rotationally-invariant sphere. A number of surgical simulators in which the surgeon manipulates a rigid virtual instrument, such as a scalpel or surgical drill, have been developed based on these 3-DOF haptic rendering principles [12,13,14]. In contrast, a 6-DOF haptic rendering algorithm computes both forces and torques from the position and orientation of the device. These methods use the entire virtual instrument's geometry for collision and contact handling.

A common misconception is that use of a 6-DOF haptic rendering algorithm requires a fully-actuated 6-DOF haptic device. Such devices carry a significant cost premium due to mechanical design challenges that need to be overcome and the high cost of parts. Today, many commercially-available haptic devices are asymmetric in that they have a different number of sensors than actuators (motors) [15]. A common kind (e.g. SensAble's Phantom Omni) senses 3D position and orientation (6-DOF), but provides only directional force feedback (3-DOF).

3. Research Questions

As a step toward informing the level of haptic realism and fidelity required to achieve surgical simulation objectives, we study the effect of haptic feedback degrees of freedom on task performance. In minimally invasive surgery or microsurgery, the surgeon must often work through narrow corridors while avoiding excessive force or accidental incursions that can cause trauma to surrounding tissue or sensitive structures [8]. We designed a surgically-motivated interaction task that involves similar precise positioning of a virtual instrument in kinematically constrained environments to reflect this condition.

We aim to compare the effect of 3-DOF haptic rendering to that of 6-DOF haptic rendering, and within the latter we also compare its effect when rendered on a fully-actuated (force and torque output) versus an under-actuated (force only) haptic interface. We use *sphere* rendering to refer to a 3-DOF method that computes haptic feedback through a sphere centered at the tip of the instrument (e.g. [12,13,14]). Our 6-DOF rendering algorithm [16] treats the virtual instrument as a full rigid body for collisions and contact, and we henceforth refer to it as *r-body* rendering. Our hypotheses are then:

H1 *R-body* haptic rendering improves task performance over *sphere* haptic rendering.
H2 *R-body* haptic rendering on a force and torque display improves task performance over rendering on a force-only display.

4. Methods & Materials

An experimental study with a within-group design was conducted to measure the effect of three variants of haptic feedback (sphere rendering, under-actuated r-body, and fully-actuated r-body) on task performance in two surgically-relevant virtual scenes. The presentation order of the two scenes, and then of the three haptic rendering variants within a scene, was randomized. The experiment, including a pre-study questionnaire and a semistructured debriefing interview, lasted 90 minutes. A five-minute break was mandated midway through the experiment.

One virtual scene used for the study was a model of middle ear anatomy (Figure 1a), inspired by our ongoing work in otologic surgery simulation. The other was a synthetic scene modeled to emulate similar constraints that may be encountered in other surgical procedures (Figure 1b), where the instrument must be passed through a small, round port. A number of small targets were placed at various locations within the scenes, and the task was to touch all of the targets (in any order) using the tip of a virtual probe while avoiding excessive contact with obstacles in the environment.

The three variants of haptic feedback were compared as the independent variable in this study. With sphere rendering, only contact with the tip of the probe results in force feedback. The subject would experience no additional haptic feedback if the shaft of the instrument were to collide with obstacles in the environment. With r-body rendering, under-actuated display is emulated on the same 6-DOF haptic device by simply discarding the computed torques, thus controlling for differences of other device characteristics.

Task performance was measured in terms task completion time and the number of errors made. An error was defined as exceeding 5 mm of incursion of the instrument into another structure. In all variants, we provided a form of sensory substitution (or "visual haptics" [2]) by coloring the probe yellow for small penetrations (>2 mm), then

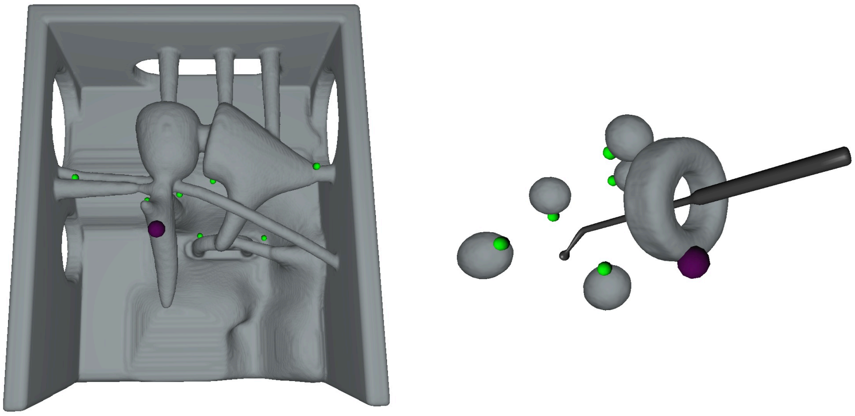

(a) “Ear” scene. Model of middle ear anatomy.
(b) “Port” scene with a narrow corridor.

Figure 1. Virtual environment scenes used in the study. The objective of the task is to touch all the small spherical targets using the virtual probe instrument shown in (b).

orange (>3.5 mm), and finally red when the error threshold is exceeded. To complement measured performance, perceived performance was captured by a questionaire and a semistructured interview.

4.1. Apparatus

A stereoscopic 3D virtual environment was created to conduct the experiment. Within the environment, the subject controls and manipulates the virtual probe instrument using a Phantom Premium 1.5/6-DOF haptic device (Figure 2). Virtual scenes were scaled (including the middle ear) to a size of roughly 20 cm to fit the workspace and spatial resolution capabilities of the device. Stiffness of haptic rendering was set to 500 N/m of displacement. Torsional stiffness for 6-DOF interaction depends on the inertia of the virtual instrument, and amounted to approximately 1.8 Nm/rad for the probe. Visual feedback was provided in stereoscopic 3D through an LG 32” television with passive circular polarizing glasses. Task completion time and number of errors made during each trial were automatically recorded by the software application.

4.2. Procedure

Twelve subjects (8 males, 4 females aged 19-41, mean 25) participated in the experiment. One subject was left-handed, for which the virtual scenes were mirrored. Five subjects were medical students and four have had clinical experience. Subjects were compensated with two movie tickets for their participation.

Subjects were instructed on the use of the haptic interface, first within the manufacturer’s test application, then in an unscored pre-study scene in our application, until they were familiar with haptic exploration procedures and force feedback. The subjects were given written instructions of the task including instructions to complete the task as quickly and with as few errors as possible. In addition, the instructions reassured that no error would be recorded for contact resulting in a warning (yellow or orange) level, and that the participant should explore error boundaries during the practice sessions.

(a) Experimental setup showing the Phantom Premium 1.5/6-DOF haptic device and 3D television.

(b) A subject's grasp of the device handle and the corresponding virtual probe instrument.

Figure 2.

Each subject completed multiple measured trials of the task in both scenes and with all three variants of haptic feedback. The subject was asked to practice under each condition for about four minutes or until s/he felt ready. Then the subject repeated the task for five minutes while measurements of time and errors were taken. An average of 8 trials per condition were recorded for every subject.

A written questionnaire was administered after each condition session. Perceived difficulty was measured as the sum of the answers to two 7-degree Likert scale questions, one regarding the difficulty of hitting the targets and the other of avoiding collision with the surrounding environment. Perceived benefit of haptic feedback was also measured on a 7-degree Likert scale as the answer to the question, "Did you perceive that the haptic feedback was assistive in helping you to complete the task?" The experiment ended with an interview regarding the participant's experience of the haptic feedback variants.

5. Results

Analysis of the data with paired t-tests (all having $df = 11$) showed significant differences between sphere and r-body rendering. Apart from perceived performance, no significant differences between fully-actuated and under-actuated display were observed.

5.1. Task Performance

The analysis was based on comparing the average of each subject's result for one condition and scene with the same subject's average result for each of the other two conditions (within-subject, paired t-test). The average for all measurements is reported in table 1.

Task completion was significantly faster with r-body rendering compared to sphere rendering using both the fully-actuated ($t{=}7.0, p{<}0.001$) and under-actuated ($t{=}7.8, p{<}0.001$) display in the port scene. No significant time differences were found in the ear scene. Significantly fewer errors were made with r-body rendering compared to sphere rendering using both the fully-actuated ($t{=}6.5, p{<}0.001$) and under-actuated ($t{=}6.6, p{<}0.001$) display in the port scene as well as the ear scene ($t{=}3.8, p{=}0.002$; $t{=}3.3, p{=}0.003$ respectively). No significant differences were found between fully-actuated and under-actuated display in terms of completion time or errors for either scene within a 95% confidence interval.

Table 1. Mean values and standard deviation of task completion time (sec.), errors, and questionnaire results concerning perceived difficulty (range 2-14) and perceived benefit (1-7). *U* indicates under-actuated display.

	port scene			ear scene		
	sphere	r-body^U	r-body	sphere	r-body^U	r-body
Task completion	39.5 (12.2)	27.0 (7.8)	23.9 (9.1)	43.3 (16.3)	40.8 (13.8)	38.7 (8.9)
Task errors	5.3 (3.9)	0.8 (1.3)	0.6 (1.2)	2.8 (2.5)	1.3 (2.0)	0.9 (1.2)
Total measurements	88	118	136	81	84	88
Perceived difficulty	11.3 (2.2)	7.0 (2.7)	5.6 (2.3)	9.1 (2.5)	7.5 (3.6)	6.5 (2.0)
Perceived benefit	1.9 (1.5)	4.4 (1.2)	5.1 (1.2)	2.5 (1.0)	4.1 (1.3)	4.8 (0.8)

5.2. Perceived Performance

The task was perceived to be significantly more difficult with sphere rendering than with r-body rendering using the fully-actuated display in the ear scene (t=3.9, p=0.001). In the port scene, the sphere rendering was perceived to be more difficult than the r-body rendering using both fully-actuated (t=8.0, p<0.001) and under-actuated display (t=7.7, p<0.001). Whatmore, using the fully-actuated display was perceived as less difficult than the under-actuated display (t=2.3, p=0.022). Percevied benefit was signficantly higher for r-body rendering compared to sphere rendering using both the fully-actuated (t=4.4, p<0.001) and under-actuated (t=4.4, p<0.001) displays in the port scene, as well as in the ear scene (t=6.7, p<0.001; t=3.2, p=0.004 respectively).

The interviews revealed that all subjects recognized the difference between sphere and r-body rendering, but only some could tell or articulate any difference between the fully-actuated and under-actuated display. One participant described that the fully-actuated variant *"felt more smooth when I brushed the tool over the surface, but I could not really tell (...) it felt like it gave more graded feedback."* Others perceived the fully-actuated rendering to give harder or earlier feedback. One subject particularly liked the fully-actuated rendering: *"There was something about [it] that made it make a little bit more sense, a little bit more intuitive. Maybe it was the resistance on it, maybe something else."* However, another subject felt that there was something *"weird"* with the fully-actuated variant and thus preferred the under-actuated r-body rendering which *"felt to me like I had lot of control. Maybe it had a slower response in how I rotated it. This one [fully-actuated] I felt was too fast and too hard to control."*

6. Discussion

The results of our study show that 6-DOF haptic rendering, where the full geometry of the virtual instrument is used for collision detection and contact handling, allows subjects to complete an instrument positioning task in a constrained virtual environment with fewer errors and sometimes faster. In addition, the task was clearly perceived as easier and the user experience superior to that provided by 3-DOF haptic feedback.

Subjects performed the poorest with the sphere-based haptic rendering even though visual warnings were provided before an error was made. This may indicate that sensory substitution of this form is inferior to real, high-fidelity haptic feedback. Apart from user experience, no significant difference in performance was observed between under-actuated and fully-actuated display. The use of fully-actuated devices may still have a greater effect when applied to other tasks, or with more complex geometry. A greater contribution may also become apparent when using higher-fidelity devices, such as those with improved inertia, friction, or stiffness.

A surgical simulator that provides the user with a realistic visuohaptic experience is postulated to be of utility as a rehearsal or teaching environment for rare or technically difficult surgical procedures. The results of our study motivate 6-DOF haptic rendering as a valid approach for simulation of dexterous manipulation tasks, such as those encountered in many types of surgery, regardless of whether or not torque can be displayed.

7. Acknowledgments

This research was supported in part by NIH Grant R01-LM010673 and the Department of Veterans Affairs, Veterans Health Administration, Office of Research and Development, Rehabilitation Research and Development Service. The authors thank Eva-Lotta Sallnäs Pysander of the Royal Institute of Technology for her valuable feedback on study design.

References

[1] P. Ström et al., Early exposure to haptic feedback enhances performance in surgical simulator training, *Surgical Endoscopy* **20** (2006), 1383–1388.

[2] O. A. J. van der Meijden, and M. P. Schijven, The value of haptic feedback in conventional and robot-assisted minimal invasive surgery and virtual reality training: A current review, *Surgical Endoscopy* **23** (2009), 1180–1190.

[3] L. Panait et al., The role of haptic feedback in laparoscopic simulation training, *The Journal of Surgical Research* **156** (2009), 312–316.

[4] M. K. Chmarra et al., Force feedback and basic laparoscopic skills, *Surgical Endoscopy* **22** (2008), 2140–2148.

[5] H. K. Kim, D. W. Rattner, and M. A. Srinivasan, The Role of Simulation Fidelity in Laparoscopic Surgical Training, In Proc. *MICCAI* (2003), 1–8.

[6] K. S. Lehmann et al., A prospective randomized study to test the transfer of basic psychomotor skills from virtual reality to physical reality in a comparable training setting, *Annals of Surgery* **241** (2005), 442–449.

[7] P. H. Cosman et al., Virtual reality simulators: Current status in acquisition and assessment of surgical skills, *ANZ Journal of Surgery* **72** (2002), 30–34.

[8] C. R. Wagner, N. Stylopoulos, and R. D. Howe, The role of force feedback in surgery: analysis of blunt dissection, In Proc. *IEEE Haptics Symposium* (2002), 68–74.

[9] S. Wang, and M. A. Srinivasan, The role of torque in haptic perception of object location in virtual environments, In Proc. *IEEE Haptics Symposium* (2003), 302–309.

[10] L. N. Verner, and A. M. Okamura, Force & torque vs force only feedback, In Proc. *IEEE World Haptics* (2009), 406–410.

[11] R. Weller, and G. Zachmann, User Performance in complex bi-manual haptic manipulation with 3 DOFs vs. 6 DOFs, In Proc. *IEEE Haptics Symposium* (2012), 315–322.

[12] M. Agus et al., Real-time haptic and visual simulation of bone dissection, *Presence* **12** (2003), 110–122.

[13] D. Morris et al., Visuohaptic simulation of bone surgery for training and evaluation, *IEEE Computer* (2006), 48–57.

[14] P. Trier et al., The visible ear surgery simulator, *Studies in Health Technology and Informatics* **132** (2008), 523–525.

[15] F. Barbagli, and K. Salisbury, The effect of sensor/actuator asymmetries in haptic interfaces, In Proc. *IEEE Haptics Symposium* (2003), 140–147.

[16] S. Chan et al., Constraint-based six degree-of-freedom haptic rendering of volume-embedded isosurfaces, In Proc. *IEEE World Haptics* (2011), 89–94.