Improving Haptic Sensations

Final Report

22nd of October 2019

Innovate UK Project No: 7238
Competition: Emerging and Enabling Round 3
Internal project codename: Elephant

Lead author: Dr Orestis Georgiou (OG)
Second authors: Salvador Catsis (SC), Dr Jerzy Dziewierz (JD) and Dr Robert Malkin (RM)

Content

1 Executive Summary

This final report summarises the work done over the past 11 months on the InnovateUK supported project codenamed Elephant. The project has successfully developed novel algorithms for the generation of ultrasonic mid-air haptic surfaces, rather than haptic focal points. These new algorithms were optimised and tested; however current hardware is not sufficiently powerful to fully exploit their advantages. Going forwards, IP protection will be sought, and the algorithms will be tested on next generation hardware due to be sufficiently powerful to achieve haptic surfaces. Further, this research project has opened up the doors to a potentially new field of acoustic imaging using ultrasound.

2 Introduction

2.1 Project Scope

Ultraleap (previously Ultrahaptics) has been developing mid-air, non-contact haptic feedback systems for human-machine interfaces by using focused high-intensity ultrasound. The state-of-the-art Ultraleap technology creates a variety of haptic sensations using a collection of focus points that are rapidly animated to form a haptic image similar to how an old CRT monitor works.

This project has developed fundamentally new underlying haptic technology to allow a much greater range of sensations to be generated. Namely, the project has developed algorithms that create complex ultrasonic fields that represent haptic surfaces in a single instance, analogous to how LCD monitors work, rather than in a recursive and iterative manner. Thus, the codename *Elephant* was chosen for this project since elephants have a large surface area.

In more detail, the proposed project has investigated algorithms that solve for a complex acoustic field and can efficiently transform this field in space thus saving in computational effort. For example, instead of a single haptic point output, the desired acoustic field now consists of a line or a curve haptic output, i.e., an increase by one in the spatial dimension of the haptic output.

The new results have been simulated, implemented in hardware prototypes, and disseminated across the *Research, Engineering* and *Product* teams of the company.

Going beyond the project's 11-month timeline, we will investigate patent opportunities and the implementation of the new algorithms in next generation hardware products.

2.2 Work Activities

In order to deliver the project's main result, namely the development of new and improved haptic algorithms, we have undertaken a large variety of high-quality technical work.

Specifically, we have:

1) performed a literature review to understand the background and state-of-the-art;
2) mathematically investigated different approaches inspired by Fourier optics and medical imaging;
3) defined and implemented tests in acoustics simulation software;
4) designed and implemented new solver algorithms in simulation software;
5) boosted the performance of said algorithms using OpenCL parallel processing architectures;
6) tested the algorithms in simulation;
7) transferred the algorithms onto hardware;
8) performed acoustic measurement and thermal camera studies; and
9) frequently presented our findings to internal Engineering teams, Ultraleap management, and to project review meetings with InnovateUK.

2.3 People, Training and Resources

In order to deliver the project's work activities listed above, SC, a recent and exceptional graduate from the University of Bristol was brought on board to deliver the majority of the technical work, under the close supervision and guidance of the more experienced scientists JD, OG and RM.

In total, the project actively involved 8 technical and 2 non-technical staff. SC was 100% dedicated to this project while everyone else had a supporting and supervisory role. Namely, management, reporting, supervision and adoption activities were led by JD and OG. Resource tracking and financial reporting activities were led by the *Finance Team* and *Project Management Team* using internal project management and time tracking software. Algorithm development and testing activities were led by SC. Acoustic measurements were supported by the *Acoustics Team*. Software development activities were supported by the *Software Team*. Perceptual user testing activities were supported by the *User Experience Team*. Patent advice was obtained from the Ultraleap *Intellectual Property Committee*.

SC had weekly mentoring sessions with JD who is a Senior Scientist at Ultraleap with plenty of supervisory experience. SC also attended an acoustics training event in London organised by the UK Acoustics Network (UKAN), an acoustics summer-school and workshop in Powys, Wales, and the International Congress on Industrial and Applied Mathematics (ICIAM) that took place in Valencia, Spain. All these training and mentoring activities resulted in the development of technical and soft skills for SC as he further advances his career as an independent researcher.

SC was given full access to a company laptop for day-to-day research and development coding activities (MacBook Pro 2018) and to a high-performance laptop (Alienware 2018) equipped with the GTX1080 graphics card for parallel processing. Almost all code was written using the Python programming language. Further, SC utilised a thermal imaging camera for flare monitoring, and a 1/8-inch Brüel & Kjær (B&K) ultrasonic microphone mounted on a repurposed 3D printer to produce acoustic field scans.

2.4 Implementation

In order to properly plan and manage the progress of the project and to delineate tasks into smaller engineering efforts, we have broken down the workplan into seven work packages, each with a work package leader as follows:

WP1: Management and Reporting (JD and OG)
Months 01 to 11
Coordinate, lead decision-making and reporting, data management, training activities, ensure delivery of project goals, and facilitate good working relationships. This work package spans the whole project.

WP2: Planning & Research (SC)
Months 01 to 03
Ensure project novelty and use case through an exhaustive literature review of the state-of-the-art technology both within Ultraleap database and knowledgebase, but also externally through am interdisciplinary search of the Web of Science (formerly ISI Web of Knowledge). Understanding of the current solver and research into existing techniques. Define a high-level solver architecture that will be compatible with fundamental principles and limitations of existing hardware technology platform.

WP3: Solver implementation (SC)
Months 04 to 11
Develop new or improved algorithms for the efficient creation of haptic surfaces. Implement and validate these algorithms in a software simulation package (Alpha release). Optimise these algorithms in a software simulation package (Beta release). Implement and validate these algorithms in hardware (Final release).

WP4: Experimental Validation (SC)
Months 06 to 11
Use acoustic measurement equipment and any other apparatus necessary for the experimental analysis and validation of the performance of the new Solver.

WP5: Perceptual testing (SC)
Months 09 to 11

Enlist a group of people to try out and evaluate the new haptic solver. The conclusions will then feed back to the implementation team to improve the development of the algorithm.

WP6: Exploitation plan (OG and JD)

Months 06 and 11

Continuously evaluate the innovation and intellectual property opportunities arising from WP3 activities. Lead the planning of dissemination and exploitation activities such as paper presentations to the academic community, patent applications, and algorithm implementation in existing and future products.

WP7: Adoption (JD)

Month 11

Document and knowledge transfer of results and new findings across the company.

Furthermore, we have identified a key set of milestones that would help us track progress and minimise risks of completing the project.

Key Milestones

M1.1-4: Review meetings with project officer
M2.1: Literature review and Solver architecture
M3.1: Prototype in Python and basic tests operational
M3.2: Detailed module tests and visualization demonstrated
M4.1: Experimental validation complete
M5.1: Perceptual testing complete
M6.1: Exploitation plan update
M7.1: Adoption plan activities

Management Approach

The project used the Agile project approach in meeting the above milestones. The project team held frequent stand-ups to highlight work complete, upcoming work and any issues or perceived risks. OG and JD attended technical team lead (attended by all team leads across the company) meetings with the rest of the company where significant updates were reported and discussed.

Risk Management

A risk register was maintained and updated by JD. This was reviewed by the project officer during review meetings. Only one risk has materialized:
R4: New algorithm does not provide a better haptic sensation compared to single point.

The risk was identified as early as the project kick-off meeting, and a mitigation strategy was devised. The reason for this eventuality was that the current generation of the ultrasound hardware was not powerful enough (powerful being used as a proxy for generated sound pressure) to produce the required haptic output. A single focal point is adequately generated by the 256-transducer array used by Ultraleap. In order to generate curves and surfaces, many more transducers are required. Such large hardware is still in development by the Engineering team and is expected to be delivered by end Q4-2019. To that end, the new haptic algorithms have been tested and verified in simulation and have also been experimentally tested using existing hardware prototypes using microphones and thermal cameras since the output was not sufficiently strong to generate a tactile sensation perceivable by a human hand. These initial measurement and testing results are encouraging. We will therefore perform a full evaluation of the new algorithms on new larger hardware platforms when they are available in 2020.

3 Technical Summary

3.1 Acoustic Surface Field Reproduction

This work is concerned with developing a commercially suitable algorithm for creating a desirable distribution of acoustic energy across space.

The output of this algorithm is a set of complex numbers referred to as *activation coefficients.* These are amplitudes and phases that a set of ultrasonic transducers should operate at in order to produce the desired spatial distribution of acoustic field intensity.
The operation of solving for the activation coefficients is also named “inverse propagation problem”, or “backwards (in time) propagation problem”

During this project, it has been observed that the problem of creating a desired acoustic intensity distribution across space can be represented as a classic mathematical problem of matrix inversion operation, where components of the matrix are complex numbers, and the matrix is not necessarily square nor fully or uniquely invertible.

Since the solution to the backward propagation problem is fundamentally not unique nor exact, one also needs to forward propagate (in time) the acoustic waves emanating from the collection of ultrasound transducers in order to see what happens. This can be done using our pre-existing simulation software, with example result shown in Figure 1 below.

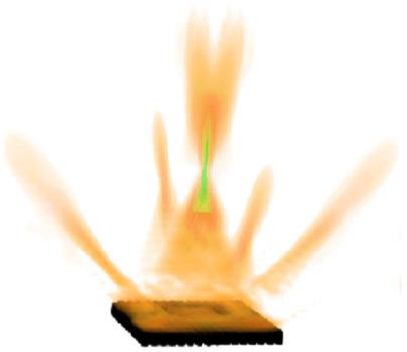

Figure 1: Acoustic field with a single-point focus at 20 cm above an array of 256 ultrasonic transducers. Black – transducers; green - focal point location within main lobe, orange-side lobes and grating lobes.

As can be seen, apart from the desired focal point(green), a series of side and grating foci are also generated in the acoustic field (orange). These are secondary and unwanted effects that can be treated to some extent by employing more advanced excitation techniques.

Having the problem formalised as matrix inversion, numerous strategies for finding a solution exist and there is a large body of existing literature on optimisation theory.

However, the classic, text-book solutions for matrix inversion only deal with most typical, abstract problems, and are not oriented towards the specific engineering problem of dealing with ultrasound beam forming. For example, a typical text-book solution assumes that it is enough to minimize the linear difference between desired and actual solution, while ignoring any other effects.

In our case, our choice of strategy is constrained by the additional physical and engineering requirements that:

- We maximize the overall intensity in the high-intensity areas, but keep the transducer output power within its physical limits
- We maximize the contrast between high intensity areas and low-intensity areas,
- We minimize the effect of grating foci or move them to locations where they become non-consequential
- We find a solution in deterministic time and with the least computational resources
- We deal with complex numbers, rather than real-only numbers.

We thus proposed and tested the following algorithms:

A1. Iterative projection algorithm based on Fourier optics
A2. Pseudo-exact propagation algorithm
A3. Constrained optimisation
A4. Gradient descent
A5. Stochastic gradient descent with momentum
A6. Time reversal
A7. Least-squares by singular value decomposition with Tikhonov regularisation
A8. Pseudoinverse with power iteration

Additionally, as a computation cost reduction step, a “Solve, Store, Recall, Transform” (SSRT) method has been developed. The first step in the S.S.R.T approach is to compute the activation coefficients required to produce a given acoustic field using one of the methods above. The next step is to store these activation coefficients in a library. The final step involves loading these coefficients into memory and then using approximate translation and scaling algorithms to translate them around three-dimensional space rather than having to re-solve every time.

Translation, and rotation algorithms were developed to facilitate the computational savings associated with algorithms A1-A8. These would appropriately apply a slight phase shift to the activation coefficients of a particular solution causing the haptic curve to be displaced as desired as seen in Figure 2. The trade-off between accuracy and compute performance was investigated in detail. One example of the study of z-axis translation algorithm is presented in Figure 2.

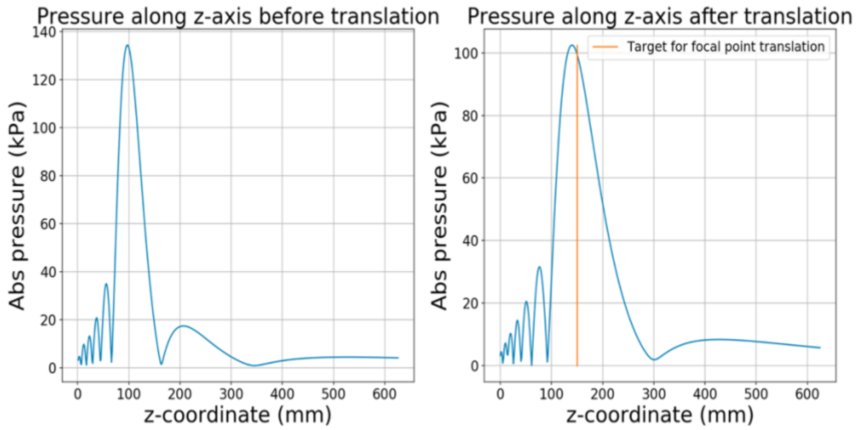

Figure 2: Results of approximate translation in the z plane using the algorithm

All these algorithms have been detailed in a separate report. They have been implemented in Python as an extension module for HandyBeam; a fast acoustic simulation tool developed by Ultraleap. The code implementation is available on our internal code repository, while the acoustic simulator (but not the haptic surface algorithms A1-A8) can be downloaded from: https://github.com/ultraleap/HandyBeam

3.2 Acoustic Simulations and Testing

In 1678 Huygens proposed a model where each point on a wave front may be regarded as a source of waves expanding from that point. HandyBeam uses Huygens principle along with OpenCL/GPU acceleration to visualise the acoustic field. For the purpose of this project, various new virtual measurement tools (e.g., sampling grids, focal spot size measures, side grating foci measures, etc.) were developed and included in the most recent build of HandyBeam.

More specifically, the following test packages were developed:

- A rectilinear and a hexagonal sampling grid with an arbitrary origin and an arbitrary orientation. A Lambert azimuthal equal-area projection sampling grid. These sampling grids can be used to specify the desired surface acoustic field for the new solvers.

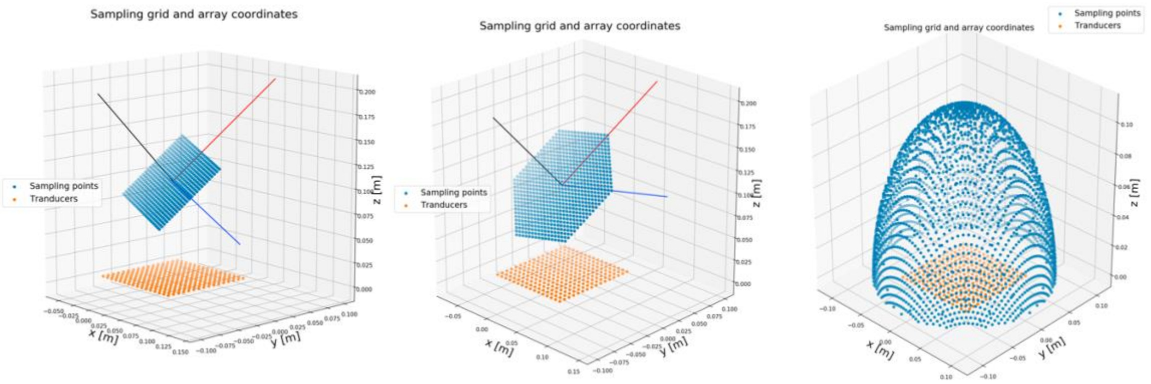

Figure 3: rectilinear, hexagonal and Lambert sampling grids above a 256-transducer array.

- Acoustic field **metrics** such as the focal spot area (FSA), the focal spot contrast (FSC), the focal to grating-foci peak pressure ratio (FGFPR), the focal to total intensity ratio (FTIR), the array focus indicator (AFI) in 2D and 3D, the array microphone interference indicator (AMII), the array haptic indicator (AHI) in 2D and 3D, the haptic surface contrast (HSC), the haptic surface efficiency (HSE), and the haptic surface variance (HSV).
- 3D visualisation package for visualising the whole acoustic field.

These were later used to assess the quality and performance of the acoustic surface field algorithms described in the previous subsection. Some of the results are shown in Figure 4 for the generation of an acoustic surface that looks like a spiral and an equal sign.

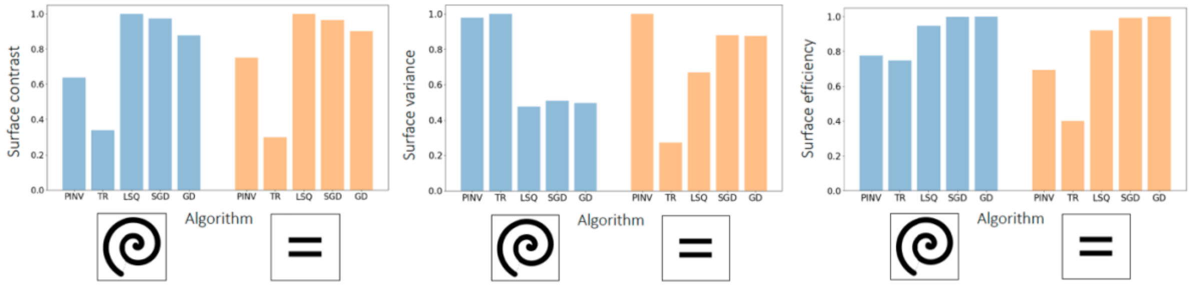

Figure 4: Comparison of different metrics for a spiral and equal sign acoustics field

3.3 High performance implementation

It is possible that future mid-air haptic products use a dedicated digital signal processor (DSP) or graphics processing unit (GPU) style processor, instead of a field-programmable gate arrays (FPGA). To prepare for this possibility, we developed an OpenCL implementation of selected algorithms A4-A8. OpenCL (Open Computing Language) is a framework for writing programs that execute across heterogeneous platforms consisting of central processing units (CPUs), GPUs, DSPs, FPGAs and other processors or hardware accelerators. Importantly, OpenCL facilitates for parallel processing, a method of simultaneously breaking up and running program tasks on multiple microprocessors, thereby reducing processing time. Here, we have allocated each sampling point and each source point to separate processor threads thus achieving a linear speed-up in computing time.

3.4 Experimental validation

Current mid-air haptic technology uses 256 ultrasonic transducers to focus acoustic energy onto a single point situated on the user palm or fingers (see Figure 5).

Figure 5: Ultrasonic phased array generating a focus point

The maximum sound pressure level that can be concentrated to a single point by current hardware is approximately 162 dB sound pressure level ((SPL) – re: 20µPa), or equivalently a sound intensity of 15.8 kW/m2. The perceptible threshold of the human skin is estimated to be close to 150 dB SPL, or equivalently a sound intensity of 1.0 kW/m2. Note that the area occupied by a single focal point is approximately the size of the operating wavelength, which at 40 kHz corresponds to ≈8mm, thus giving an area of about 50 mm2. Each line comprising the equal sign (=) seen in Figure 4 has width 8 mm and length 50 mm, thus giving an area of about 800 mm2, which is essentially 16 times larger than a focus. A back of the envelope calculation of the sound pressure level loss when spreading the same acoustic energy over an area that is 16 times larger than that of a focus point gives a -12 dB loss. In other words, the maximal achievable sound pressure level at any part of the “equal sign” haptic surface would be similar to the minimum perceivable tactile threshold by humans.

More powerful hardware (either through greater transducer number of single transducer output) is therefore needed in order to adequately output the haptic surfaces of algorithms A1-A8 such that they can be perceptually felt by humans. Alternatively, more energy efficient algorithms could also be evoked. The former is already part of the internal engineering roadmap and is due to be delivered for internal testing by the end of Q4-2019. The latter involves the intelligent use of the power output of the device as to achieve better haptics using less power. One way to do this in the case of the equal sign is to “blink” each of the two parallel out of phase to each other. Thus, at any given instance of time, the hardware is outputting a single line of area 400 mm2. Touch being a time-averaged sensory input cannot resolve this blinking and thus continues to “see” the full equal sign. Through this technique we have therefore achieved a 200% efficiency gain in power use. This however is still not sufficient for a meaningful haptic user study of the performance of the new haptic algorithms as was advised by our user experience team.

Milestone M5.1: Perceptual testing complete was therefore not completed within the short time frame of the project but remains as an internal task for 2020 when new, larger and more powerful hardware becomes available. We have identified this eventuality as a risk, fairly early in the project. Unfortunately, little control over the possible mitigation strategy was afforded by the Elephant team, since the fabrication and testing of larger and more powerful hardware is a separate and much more complex internal project that we had little influence over.

We therefore found alternative ways of experimentally validating the algorithms we have developed. Namely, we have used thermal imaging (see Figure 6) and acoustic measurements (see Figure 7).

It is well known that an acoustically absorbent material will increase in temperature when it is insonified with an acoustic wave. We utilise this effect as an indirect method of measuring the acoustic pressure field generated by the state-of-the-art hardware running haptic surface algorithms.

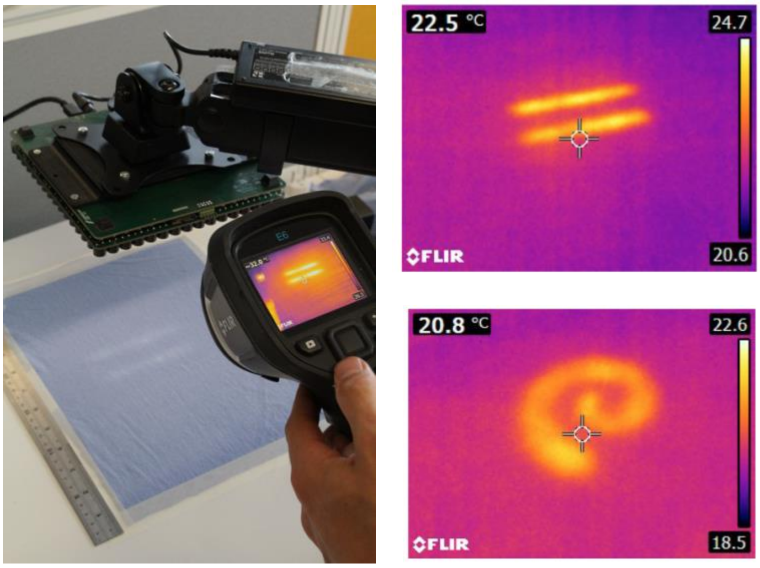

Figure 6: Experimental set up of the thermal imaging stage. Left: An equal symbol projected onto temperature sensitive cotton. Right: An equal and spiral surface temperature profile on thermal cotton.

Another way to measure the acoustic pressure field generated is to directly place a microphone in some volume above the array and record the measured pressure. We therefore repurposed a 3D printer into an acoustic measurement stage by replacing the printer head with a 1/8 B&K microphone and programming the microphone to move over a discretised volume with a given velocity. Notice that the microphone field measurement is much more granular due to the time required to acquire the experimental data.

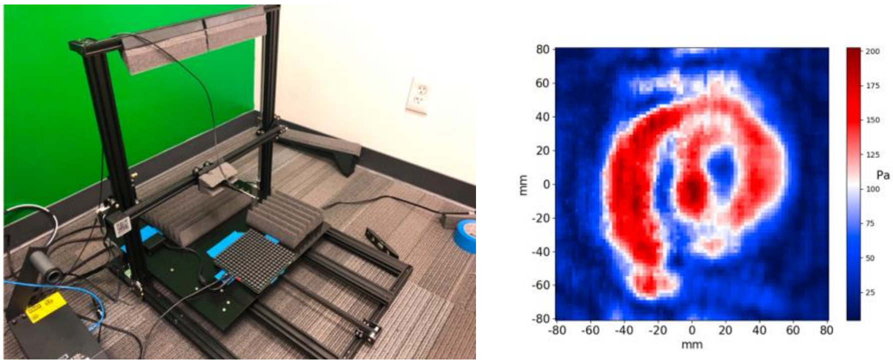

Figure 7: The acoustic pressure field generated by the LSQ algorithms as measured by a microphone scan.

4 Exploitation and Adoption

The key opportunity identified through this project is that of fundamentally changing the underlying algorithms that drive Ultraleap non-contact haptic feedback systems such that larger surfaces and volumes could be rendered directly. It is the analogous step-up in technology as that of going from a CRT to and LCD screen, but for mid-air haptics. This haptic innovation allows for more complex and realistic objects to be generated and hence is directly applicable to high growth markets such as virtual reality systems in location base entertainment venues as well as for simulation and training. Both these use cases do not impose small form factor and low power restrictions, and are therefore good candidates for the deployment of large ultrasound arrays, capable of utilising the haptic surface algorithms developed herein.

As discussed in the previous sections, we have developed the necessary algorithms, simulated them, tested them, enhanced them using high-performance computing, and experimentally validated them on real hardware. . Further advancements are needed in the effective power output of ultrasonic hardware for our algorithms to be fully usable in commercial context. Therefore, direct exploitation of the algorithms currently remains on the horizon.

In the short term however, the company aims to write up the novel algorithms and supporting evidence into a patent filing, thereby securing the innovative intellectual property we have developed. After the patent application is made, we will also seek to write up our findings in the form of an academic publication (the majority of which has been completed already by SC) such as to share our newly gained knowledge and to invite peer review from the scientific community.

Internal dissemination of the project's results has been achieved through a companywide presentation by SC. The presentation summarised the content of this report, and also included some technical details regarding the implementation of the algorithms and also where the relevant code could be accessed on our local repository. The presentation was attended by both Engineering and Product team members therefore achieving knowledge and technology transfer within different departments and at different levels. It was noted that while large haptic surfaces could not be adequately outputted by current hardware, small haptic surfaces (non-point like) were possible and should be explored by the Capabilities team and tested by the User Experience team.

5 What comes next?

There are currently three potential avenues for further research and development that have sprung from the foundational and innovative work done in the current project. Each of these options has different priority rankings and resource requirements:

1. Elephant 2.0 - further optimize and test the algorithms developed but with an additional energy constraint
2. Elephant 2.1 - optimize and test the algorithms developed on newly available hardware platforms
3. Explore non-haptic applications of the newly developed algorithms, e.g., medical/industry ultrasound imaging.

While projects 1 and 2 form incremental progress and therefore can be pursued internally, project 3 is of higher risk and requires expertise that we currently do not have in the company. Therefore, project 3 may be a good candidate project to seek external funding from national or international funding bodies. At this time, one such funding source has been identified, and the project is internally known as "ultrasonic camera for dusty, optically opaque environments".