HAPTIC ENHANCEMENT OF SENSORIMOTOR LEARNING FOR CLINICAL TRAINING APPLICATIONS

Earle Spencer Jamieson

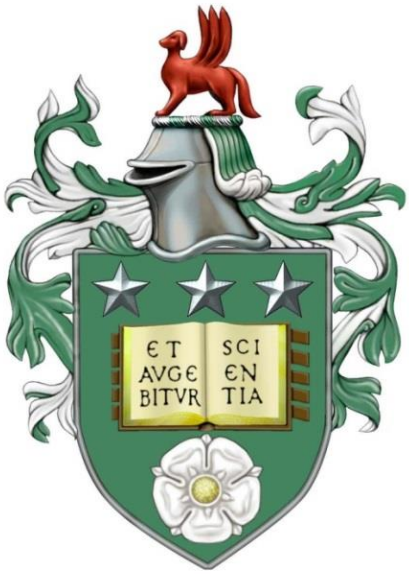

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

The University of Leeds, UK

School of Mechanical Engineering & School of Psychology

Leeds, UK

December 2015

INTELLECTUAL PROPERTY AND PUBLICATIONS

The candidate confirms that the work submitted is his own, except where work which has formed part of jointly-authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

The following publications were obtained during the course of the PhD project.

Jamieson, E. S., Chandler, J. H., Culmer, P., Manogue, M., Mon-Williams, M., & Wilkie, R. M. (2015). Can virtual reality trainers improve the compliance discrimination abilities of trainee surgeons? In 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC) (pp. 466–469).

- This paper is based on data presented in Chapter 3.

Jamieson, E. S., Wilkie, R.M., Culmer, P., Bingham, G., Mon-Williams, M (2014). The effect of guidance forces in the learning of a sensorimotor task. In 2014 24th Annual Meeting of the Society for the Neural Control of Movement (NCM).

- Poster abstract describing preliminary findings of a subset of data presented in Chapter 5, Experiment 1.

ACKNOWLEDGEMENTS

This PhD has been one of my greatest challenges. There are a number of people who I would like to thank for making the experience possible and enjoyable.

First and foremost, I would like to thank my supervisors, Mark Mon-Williams, Peter Culmer and Richard Wilkie, for their guidance and support throughout this journey. Looking back to when I started, I was quite naive about what I had ahead of me. Mark has taught me that there is (nearly) always a story to be told, as long as you're thinking straight (using a sensible framework). Pete's ability to engineer the perfect solution will never cease to impress me. Richard was always friendly and patient, even when my 'stories' didn't make much sense…

Faisal Mushtaq deserves a special mention for going beyond the call of duty (on several occasions) to read and criticise my work when I needed it the most, and to provide moral support during those difficult times.

Megan Wood and Imogen Crook were organised and efficient at helping me with data collection.

Thanks to James, Zahra, William, Evan, Oscar and Callum for their help and generally for being around when it was time to unwind, although I sometimes do wonder whether their sole motive was to get a mention in my thesis acknowledgments…

Thanks to my parents, Yvonne and Steve. As the first person in my family to go to university, I felt that there were times when they wondered why I would want to put myself through such an 'ordeal'. Sometimes, even I found it difficult to explain the reasons… But they supported me nevertheless.

ABSTRACT

Modern surgical training requires radical change with the advent of increasingly complex procedures, restricted working hours, and reduced ‘hands-on’ training in the operating theatre. Moreover, an increased focus on patient safety means there is a greater need to objectively measure proficiency in trainee surgeons. Indeed, the existing evidence suggests that surgical sensorimotor skill training is not adequate for modern surgery. This calls for new training methodologies which can increase the acquisition rate of sensorimotor skill. Haptic interventions offer one exciting possible avenue for enhancing surgical skills in a safe environment. Nevertheless, the best approach for implementing novel training methodologies involving haptic intervention within existing clinical training curricula has yet to be determined. This thesis set out to address this issue. In Chapter 2, the development of two novel tools which enable the implementation of bespoke visuohaptic environments within robust experimental protocols is described. Chapters 3 and 4 report the effects of intensive, long-term training on the acquisition of a compliance discrimination skill. The results indicate that active behaviour is intrinsically linked to compliance perception, and that long-term training can help to improve the ability of detecting compliance differences. Chapter 5 explores the effects of error augmentation and parameter space exploration on the learning of a complex novel task. The results indicate that error augmentation can help improve learning rate, and that physical workspace exploration may be a driver for motor learning. This research is a first step towards the design of objective haptic intervention strategies to help support the rapid acquisition of sensorimotor skill. The work has applications in clinical settings such as surgical training, dentistry and physical rehabilitation, as well as other areas such as sport.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

There are a number of abbreviations used throughout the thesis. These are defined within the text but are also defined here for ease of reference.

VR	Virtual reality
CNS	Central nervous system
HAT	Haptic assessment toolkit
CSI	Compliance simulation interface
MIS	Minimally invasive surgery
LS	Laparoscopic surgery
MBL	Model-based learning
MFL	Model-free learning
DLL	Dynamic link library
API	Application program interface
JND	Just noticeable difference
KR	Knowledge of results
M	Mean
SD	Standard deviation
SEM	Standard error of the mean

CHAPTER 1

INTRODUCTION

1.1 Overview

Human sensorimotor learning is a continuous process that allows us to adapt to new or changing environments, acquire new movement skills and perceptive abilities, and recover from debilitating conditions (Tresilian, 2012). Motor control (involuntary and volitional) affects every moment of an individual's existence, but the skills many of us take for granted (picking up a cup, walking through a doorway) have taken years for each individual to develop fully. Whilst the need for motor control to interact with the environment might seem obvious, there are some situations in the modern world that require an exceptionally high degree of skilled motor action that go beyond the 'generic' skills learnt during childhood. As a consequence these skills require a substantial amount of additional time to learn. Medical practice (and surgery in particular) is one such discipline. The acquisition of sensorimotor skills is important across various clinical settings (Drucker, Prieto, & Kao, 2012; Hamdorf & Hall, 2000), including, but not limited to, minimally invasive surgery and dentistry. In these contexts, the sensorimotor system of the practitioner must be trained to generate appropriate motor commands in response to a perceptual input: effective movements need to be accurate, precise and time-efficient.

Surgical training was traditionally in the form of the 'see one, do one, teach one' apprenticeship-style program (Barnes, 1987, p. 19). Trainees would observe an experienced surgeon performing a specific procedure on a patient before attempting it themselves, and their skills were then gradually developed over time through practice on live patients. In the interest of patient safety, regulations were gradually put in place to decrease the amount of training performed on live patients, especially during the early stages of training (Parsons, Blencowe, Hollowood, & Grant, 2011). In addition, the European working time directive (EWTD) has recently stipulated the need to regulate working hours, to ensure that no doctor can work more than 48 hours per week in Europe. The consequences of this legislation have been profound, most notably, decreasing the time available to train in the operating theatre (Ahmed et al., 2014). Furthermore, the 1970's heralded a revolution in surgery with the introduction of minimally invasive surgery (MIS), which has since become the preferred method for a number of surgical procedures. Its popularity has been driven by the fact that this form of surgery leads to reduced access trauma, less pain, better cosmesis and faster patient recovery while maintaining equivalent clinical outcomes to open surgery (Larsen et al., 2009). However, the advent of MIS brought with it further challenges for surgical training due to the increased complexity of the operating environment, such as viewing the operation site via a remote 2D screen (limiting field of view and depth information) and using long levers that reverse the direction of motion of the tool tip, and limit haptic feedback (Derossis, Bothwell, Sigman, & Fried, 1998). All combined, these changes have resulted in trainees effectively receiving less training to perform increasingly complex procedures. Unfortunately, there is evidence suggesting that this has led to insufficient training, resulting in a negative impact on patient care (Grantcharov & Reznick, 2008).

The issues outlined above have resulted in a critical need for surgical training analogues that allow trainees to learn the necessary sensorimotor skills for performing operative techniques. Sensorimotor learning inherently requires 'hands-on' practice (as discussed in Section 1.4), which has led to the use of simulations of the surgical environment. The ambition of programmes to address this shortage in trainee practice hours has been to maximise the efficiency of available training time to result in the greatest educational (and consequentially, performance) benefit. Simulation has been the primary focus of these types of interventions as it allows prolonged practice, thus encouraging the automation of motor skill acquisition as a result of extensive practice in a controlled

environment that, crucially, has no direct consequence on patient safety. Virtual reality (VR) simulation systems provide a virtual model of the surgical environment, allowing the trainee to interact with three dimensional computerised images to facilitate sensorimotor interaction and learning through the provision of visual and haptic (force) feedback. Indeed, VR training systems are becoming increasingly popular in surgical training (and in other clinical areas such as dentistry), and there is evidence indicating that motor skills acquired on VR systems can transfer to real surgery (Sturm et al., 2008). However, it is important to investigate new ways in which surgical training can be further optimised.

One intriguing possibility of increasing the rate of surgical skill acquisition that leverages the inherent properties of modern VR training systems is 'haptic intervention' (Sigrist, Rauter, Riener, & Wolf, 2012). This involves applying forces to the limb during the learning process with the aim of increasing the rate of sensorimotor learning (Reinkensmeyer & Patton, 2009). However, whether or not such methods are effective has been a matter of debate in recent years. Not least, a consolidated theoretical account for such interventions is yet to be established. This highlights the need for new, systematic and empirical investigations into the effects of haptic intervention strategies on the acquisition rate of sensorimotor skill.

1.2 Modern surgery

The sensorimotor skills needed in laparoscopic surgery (LS) are complex in comparison to traditional open surgery (see Figure 1.1). The surgeon sees a two-dimensional representation (on a visual display) of the three-dimensional abdominal cavity and has to manipulate tissue skilfully without many of the visual depth cues that are present under direct observation (binocular vision is lost), whilst using instruments that significantly impair dexterity and tactile sensation (Culmer et al., 2012). Indicating the potential impact of this increased level of difficulty, there was an increase in the rate of complications following the initial introduction of LS (Kirk, 2002). The UK's National Patient Safety Agency (Catchpole et al., 2009) highlights technical problems (i.e. surgeon factors) as the most important element for patient harm in LS. Research indicates an average of four potentially consequential errors per operation (Tang, Hanna, Joice, & Cuschieri, 2004).

In addition to this, there is a relatively long learning curve associated with LS (Wattiez et al., 2002).

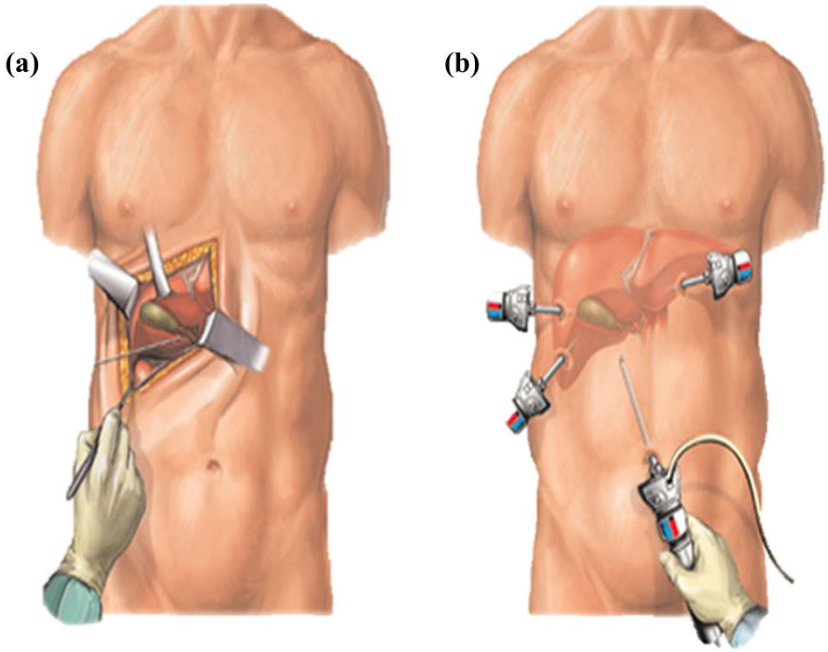

Figure 1.1. (a) Traditional open surgery whereby a large incision is made, retracted and the site is viewed directly. (b) Laparoscopic surgery (LS) uses multiple small incisions with inserted ports whereby various long-handled tools can be inserted. The lower port site shows the laparoscopic camera that provides an image of the site that is displayed on a remote screen.

Within a LS procedure, common tasks will include tissue manipulation, assessing tissue health, cutting and suturing. As outlined earlier, successful execution of these tasks relies on highly developed sensorimotor skills which can take years to develop (Larsen et al., 2009). One way in which the challenges associated with LS have been addressed is through robot-assisted MIS, whereby the surgeon doesn't directly control the movement of the tool, rather the movements are sensed and then a robotic arm performs the required movement. Such a system allows the relationship between the surgeon's movement and the movement of the tool to be altered, for example to make a large hand movement translate to a small movement of the tool if very fine adjustments were needed. The da Vinci® System (Intuitive Surgical, Inc.), for example, supports precise control and highly dexterous movements of surgical tools via an intuitive hand-controlled interface, and it also provides additional depth information via stereo vision (rather than a standard flat 2D screen; (Gomes, 2011). Unfortunately, high costs, long setup times and absence of haptic feedback are all factors associated with the use of robot-assisted MIS systems

relative to traditional MIS, which has meant that they are generally reserved for procedures requiring a high level of precision and dexterity, such as brain and prostate surgery (Turchetti, Palla, Pierotti, & Cuschieri, 2011). These factors have meant that traditional MIS is still the default approach to certain classes of surgical procedures.

The perceptual and motoric difficulties intrinsic to MIS are likely to require longer periods of training, however there is in parallel a pressure to reduce working hours and surgical costs as outlined in the next section.

1.3 The current state of surgical training

Surgery has traditionally been a craft in which skills and techniques were handed down from the 'master' to their apprentice in an informal apprenticeship-style program, often referred to as the 'see one, do one, teach one' style of teaching. Since then, training has changed by becoming more formalised and with a greater emphasis on patient safety (Aggarwal, Hance, & Darzi, 2004). In 1997, a structured training curriculum, named the intercollegiate surgical curriculum (ICSP) was introduced which laid out the skills that surgeons were required to learn. This was effective to a certain degree as it monitored progress and reformed the structure of training (JCST, 2012). However, modern surgical trainees have less time to learn about their craft than their peers did. In fact the Royal College of Surgeons calculates the time spent training has reduced from 30,000 hours to just 8,000 in recent years (Philips, 2003).

Several factors have contributed to these issues, including restrictions on working hours, the growing number of techniques and specialties that surgeons are required to learn (e.g. a move from open surgery to MIS techniques), moral and ethical restrictions resulting in a reduced hands-on experience with real patients (especially during initial years of training), as well as a reduced availability of suitable training analogues (e.g. cadavers). Indeed, one study found that only 34% of surgical trainees believed they were given adequate training (Bann, Datta, Khan, Ridgway, & Darzi, 2005). This is unsurprising given the fact that the number of operations available for training is 38% less than the recommended amount (Crofts, Griffiths, Sharma, Wygrala, & Aitken, 1997). This has highlighted the need to transform the way surgery is taught, with focus shifting from quantity to the quality of training that trainees receive.

In the UK, the training syllabus is divided into four core areas: specialty-based knowledge, clinical judgement, technical and operative skills and professional skills and behaviour (ISCP, 2013). Of these, technical and operative skills are considered by many the most difficult to master (Grantcharov, Bardram, Funch-Jensen, & Rosenberg, 2002). Indeed, it can take years for a surgeon to progress from the early (explicit 'cognitive') stages of learning to executing precise and efficient movements with minimal cognitive effort ('automation'). This highlights the need to investigate how such skills can be learned more efficiently. First, however, it is important to clearly define what is meant by 'skill', and to outline the current methods used in sensorimotor skill assessment.

1.3.1 Assessment of surgical skill

Skill assessment is vital in surgical training. From the early foundation years through to consultant level, skills in each area are assessed at different stages of the syllabus. Traditionally, assessing performance of surgeons involved assessing their theoretical knowledge, through structured questions whilst neglecting to measure technical proficiency (Moorthy, Munz, Sarker, & Darzi, 2003). Additionally, direct outcome measures were often employed: for example, a surgeon's abilities would be correlated with operative mortality, and for those using slightly more sophisticated metrics - length of stay in hospital. Although such measures have a degree of face validity, they are confounded by patient variables (e.g. the patient may be obese and therefore be a more difficult case). More experienced surgeons tend to take on more difficult cases which naturally have more chance of complications. In more recent times, video assessments of technical ability have been used. Here, an operation is videotaped and later rated by expert reviewers. Whilst this approach has shown significant benefits over previous methods, one issue with this technique is that it is largely subjective from the assessors' viewpoints. Furthermore, although video assessments are able to distinguish large differences in performance, evidence suggests that subtle differences may not be detected and so may not accurately assess surgical proficiency (van Hove, Tuijthof, Verdaasdonk, Stassen, & Dankelman, 2010).

Today, the most widely used measure of surgical proficiency is the 'objective structured assessment of technical skill' (OSATS; (Martin et al., 1997)). The OSATS involves trainees performing several different techniques, such as tissue handling, on various bench models. Although this provides a general measure of skill, the assessment of

surgical proficiency is still based on the subjective opinion of the observer. Therefore, the assessments lack an objective appraisal of ability (Paisley, Baldwin, & Paterson-Brown, 2001) and are costly in terms of manpower (Martin et al., 1997). Motion tracking systems have also been developed to try to assess surgical skill: the 'advanced Dundee endoscopic psychomotor trainer' (ADEPT), a motion tracking system which records duration and accuracy of movement (Kitagawa, Dokko, Okamura, & Yuh, 2005), and the Imperial College Surgical Assessment Device (ICSAD), a measure of hand motion efficiency are two examples. Although it uses validated measures, ADEPT has been criticised for having limited difficulty so it does not effectively capture performance of more experienced trainees (Grantcharov et al., 2002) and ICSAD for requiring a standardised technique which means it is not robust to stylistic differences in technique (Darzi & Mackay, 2002). Nevertheless, objective assessment of skill has become an important aspect of surgical training, since it is an important component of producing useful feedback for correcting errors (Kopta, 1971). Objective measurement of motor control skills are starting to appear within the screening tests used as entry requirements for surgical training. These tests assess the potential candidates' visuospatial capabilities to ensure they are capable of learning the complex skills involved in LS. This is typically done by assessing the ability to manipulate objects in abstract tasks such as stacking sugar cubes using laparoscopic probes (Seymour et al., 2002). The current trend in LS skill assessment techniques is the use of measurement systems designed to obtain objective metrics of performance as skill progresses. Indeed, it is common for surgical trainees to develop surgical skills of 'baseline surgical dexterity' on models (simulators designed to replicate the surgical environment or tasks as well as measure performance) rather than on living patients (McCaskie, Kenny, & Deshmukh, 2011).

1.3.2 Surgical simulation

The drive to reduce patient risk has resulted in the use of simulators whereby trainees can safely practise their skills without endangering the health and wellbeing of patients. For context, it is useful to review the general process used during a surgical procedure and the types of competencies that are required. Consider a laparoscopic cholecystectomy (gall bladder removal). After preparing the operating theatre and surgical tools and positioning the anaesthetised patient, four incision locations are marked on the outside of the abdomen. Incisions are made using a scalpel and the laparoscope (camera) and laparoscopic instruments are inserted (this involves breaking through several layers of

tissue). Upon internally identifying the region of interest (ROI) visually, tissue structures (fat, organs and/or muscles) are grasped and manipulated using laparoscopic graspers whilst simultaneously using laparoscopic scissors to gradually expose, isolate and excise the ROI from surrounding tissues. The remaining open internal structures are then sealed by cauterisation (if there is bleeding) and/or suturing (this involves manipulating a curved needle using a laparoscopic grasper), the tools are removed and the external incisions are sutured. Within these steps, there are a number of elements which can be classified into procedural (i.e. knowledge-based and/or decision making) skills, and motor skills. For instance, knowledge of the human anatomy is critical for identifying the ROI and knowing what structures can be manipulated and/or cut (these tasks are predominantly procedural), whilst navigating the environment, manipulating tissue structures, cutting and suturing primarily require complex technical (motor) skills.

Simulators are designed to train some or all of the major skills that are needed to perform a procedure, such as the ones outlined above (Tsuda, Scott, Doyle, & Jones, 2009). There have been a number of different simulators in use, each of which presents different benefits and limitations. Table 1.1 provides an overview of the types of simulations that are currently available, outlining their advantages, disadvantages and the situations in which they have been most effective.

Table 1.1. Overview of simulation methods used in surgical training (adapted from (Reznick & MacRae, 2006).

Simulation	Advantages	Disadvantages	Best use
Cadavers	High fidelity, true anatomy, can practice entire operations	Cost, availability, single use, tissue compliance, infection risk	Advanced procedural knowledge, dissection, continuing medical education
Bench models	Low cost, portable, reusable, minimal risks	Low fidelity, basic tasks, down time	Basic skills for novice learners, discrete skills
Virtual reality simulators	Reusable, data capture, minimal setup time	Cost, maintenance, downtime	Basic laparoscopic skills, endoscopic and transcutaneous procedural skills

Cadavers present a close alternative to the anatomy of live human patients (i.e. they are 'high fidelity'), whilst bench models (e.g. an orange) are low fidelity because they do not accurately represent the visual and haptic feedback associated with the surgical environment. This makes cadavers a preferable choice for training. However, the cost and ethical issues involved generally means that cadavers are only used in later years of training to teach advanced skills. Simulations that are more commonly involved in the development of motor skills are bench models and VR simulators. Whilst bench models (e.g. a box with sugar cubes inside – a common setup used for practicing LS manipulation tasks) offer a low cost, portable and reusable solution, again they don't match the fidelity that is now possible with VR systems. This is the main factor that has made VR systems so popular in modern clinical training (Yiannakopoulou, Nikiteas, Perrea, & Tsigris, 2015).

1.3.3 Virtual reality surgical training

The 1980s saw VR researchers highlighting the prospect of the user being able to experience a desired environment, such as the use of flight simulators (Otaduy & Lin, 2006). Later, in the 1990s, researchers developed the idea of using virtual reality headsets in order to rehearse surgery (McCloy & Stone, 2001). Initially, as the simulators developed, VR technology was only used for the simulation of basic LS skills which used

abstract graphics, i.e. low-fidelity simulations (Grantcharov, 2008). More recent developments have enabled trainee surgeons to practise whole procedures, including the simulation of a range of pathological states and anatomical variations. This was the start of 'high-fidelity' VR systems.

Modern VR surgical training systems are designed to simulate the 'look and feel' of surgery as accurately as possible, with the aim to provide trainees with the procedural and sensorimotor skills needed to operate safely on real patients. These are often equipped with a 2D or 3D display, mechanical interfaces (e.g. laparoscopic or dental probes) which are used to navigate the virtual environment, and provide haptic feedback to the user to simulate interactions with tissue or other objects. The benefit of VR trainers in contrast to performing actual surgical procedures is the significantly reduced patient risk and the ability to give trainees the opportunity to train on the same procedure repeatedly at a convenient time and location away from the constraints of attending a particular procedure in an operating room (Yiannakopoulou et al., 2015).

Some examples of commercially available VR training systems are shown in Figure 1.2. The LapMentor (Simbionix) is perhaps the most commonly used LS training system. It gives surgical trainees the ability to undergo Basic Skills Training and Procedural Training (von Websky et al., 2012). Basic Skills Training includes abstractions of general tasks that are required in most full LS procedures. These skills focus on manipulating the camera, practicing hand-eye coordination, clip application, bimanual manoeuvring and displacement of objects. Procedural Training simulates an entire specific LS procedure (e.g. a cholecystectomy). Trainees have the opportunity to choose from a range of patient anatomies to vary the task environment. Immediate feedback from a virtual teacher is available and the trainee's progress can be monitored over time due to the LapMentor's ability to plot learning curves. VR training systems generally aim to provide a platform for prescribing basic LS skills training and procedural context, and to objectively assess performance. Other examples of VR training systems are the LapVR (CAE Healthcare), ProMIS (CAE Healthcare), LapSIM (Surgical Science), MIST-VR (Mentice), and SEP (SimSurgery) – see Figure 1.2 for details. Due to their reproducibility, ease of set up, scalability, and relatively high simulation fidelity compared to other modelling techniques such as box trainers (visual and haptic feedback are reported to be more representative of reality), these methods are considered a valuable tool for surgical training (R. Hart & Karthigasu, 2007).

Figure 1.2. Examples of four commercially available VR LS training systems: a) the LapMentor (Simbionix), b) LapVR (CAE Healthcare), c) LapSim Haptic and d) LapSim Non-Haptic Systems (Surgical Science). VR training systems often come with a display monitor, laparoscopic tool interface and auxiliary inputs such as a touchscreen, computer mouse and/or keyboard for configuring the training environment and procedures. Higher quality VR training systems come with built-in haptic feedback to allow the trainee to experience the force feedback associated with interactions with biological tissue and other objects.

In parallel to developments in LS training, there have been other clinical areas that have similarly benefited from the use of VR systems. Two such examples are dentistry and physical rehabilitation (see Figure 1.3). Whilst there has been some degree of success in physical rehabilitation systems for the training of basic sensorimotor skills (such is the case for neurologically injured patients; (Selzer, Clarke, Cohen, Duncan, & Gage, 2006)), this has not been the case for areas requiring more complex skills like dentistry and LS.

(a)

(b)

Figure 1.3. (a) The Simodont Dental Trainer (MOOG) provides dental simulation and procedural training with built-in courseware and assessment tools. (b) The InMotion ARM (Interactive Motion Technologies) is designed to deliver sensorimotor therapy to neurologic patients via forces which adapt and challenge the patient's ability.

The focus so far for LS VR training systems has been on maximising their fidelity to replicate the look and feel of surgery as closely as possible. Whilst this has been shown to be beneficial for obtaining task-relevant contextual and procedural information (Otaduy & Lin, 2006), one question that remains is whether VR systems could be used more effectively for training sensorimotor skills. A system which provides this capability could have a significant impact in any area requiring an individual to learn complex motor patterns - not just LS, but dental training, physical rehabilitation, as well as sports such as tennis or golf or even learning musical instruments (McCaskie et al., 2011). To investigate potential techniques that could be employed to improve motor learning, it is useful to review the theories that underpin sensorimotor learning and control.

1.4 Principles of sensorimotor control and learning

According to the classic Fitts and Posner theory of sensorimotor skill acquisition (Fitts & Posner, 1979), learning can be divided into three main stages: the *cognitive* stage, the *integrative* stage and the *autonomous* stage. Learning to drive a car is a skilled motor task that most adults learn in their late teenage years and is a useful example of sensorimotor skill acquisition. The initial cognitive stage is characterised by erratic performance as the trainee is required to learn the mechanics of the task (e.g. learning the position of the gears, which pedal is the clutch and which is the brake, and how far to push the clutch when changing gears). At this stage, performing the task requires significant cognitive effort. With prolonged practice, the trainee progresses into the integrative stage, at which point performance becomes more refined as the learner is able to apply their knowledge (e.g. the gears are selected smoothly, and the driver concentrates on the road ahead, checking mirrors and indicating before braking). When performance reaches a plateau, the autonomous stage has been reached; the task is no longer overly demanding and it can be carried out with virtually no cognitive effort. This qualitative description of the learning structure can be considered with respect to devising optimal surgical training regimes: trainees would ideally (from a patient's safety perspective) learn outside the operating room until the autonomous stage is reached. Nevertheless, such qualitative descriptions provide little useful information about the frameworks that underlie sensorimotor learning.

A more formal approach is needed to explore novel methodologies of optimising practice conditions. Fortunately, there have been significant advances in our understanding of the mechanisms involved in sensorimotor control and learning. Before reviewing these, it is worth considering the complexity involved with effectively interacting with the environment. The sensorimotor system is able to overcome many issues to allow skilful interactions with the environment (David W. Franklin & Wolpert, 2011). To perform an action, the system is required to quickly make a selection out of hundreds of muscles and joints and an infinite number of possible trajectories and speed profiles that can be used to perform an action. Having many ways to achieve the same goal is the issue of *redundancy*. The system must also be able to deal with *noise* which adds variability in estimating both the internal states of the body (e.g. hand position) and external states of the environment (e.g. the location of an object in space). Noise is present at all stages of the sensorimotor process, including sensory processing, planning, and in efferent motor commands. *Delays* must also be dealt with to compensate for the time taken in receiving afferent sensory information and the delay in muscles responding to afferent motor commands. *Uncertainty* is also present in the system due to incomplete knowledge about the environment or task (not knowing the weight of an object, for example). The *nonstationarity* of the system is another challenge which arises due to constant internal changes (e.g. muscle fatigue). Finally, the whole sensorimotor system is highly *nonlinear*, meaning that there are complex mappings between task goals and the afferent motor commands that are needed to achieve a required action. Rather remarkably, the system is able to skilfully overcome these issues, allowing humans to adapt to new environments and learn to carry out highly sophisticated actions to manipulate the world around them.

1.4.1 Sensory processing, planning and decision making

Humans have a remarkable ability to obtain relevant information from the world and implement actions to efficiently meet an objective. To do this, the CNS (central nervous system) is able to efficiently sample information from the environment using various sensory streams and filter out task-irrelevant information. To reduce the effects of inherent sensory noise and thus increase certainty, multiple sources of sensory information are often combined in a near-optimal statistical (Bayesian) fashion (Wolpert, Diedrichsen, & Flanagan, 2011). The human system is then able to define an objective to achieve a high-level goal based on the incoming sensory information, and choose an

effective action for achieving that objective from an infinite number of possibilities (Andersen & Cui, 2009).

1.4.2 Implementing action

Virtually any movement can be defined as a number of sub-actions which can then be divided further. At the lowest level, the human system can be thought of as having at its disposal a repertoire of elementary action units, or 'motor primitives' (Mussa-Ivaldi, Giszter, & Bizzi, 1994). Conceptually, motor primitives can be combined to create 'control modules' which can be recruited to define more complex movements (Tresilian, 2012). This hierarchical structure can be used to generate patterns that range from very basic movements (e.g. actuating one muscle in one degree of freedom), to highly complex behaviours such as controlling a number of muscles and joints with specific spatiotemporal characteristics (throwing a ball, for example).

There is evidence suggesting that control modules can be recruited during both voluntary (planned) and involuntary (reflexive) actions (C. B. Hart & Giszter, 2010). By virtue of the fact that reflexive actions are processed as an automatic response to a stimulus (i.e. without conscious processing, or any processing at all, via the brain; Courtine et al., 2009), it is likely that these control modules are, at least at some level, stored in the spinal cord and peripheral nervous system. During voluntary action, high-level commands are generated centrally (in the brain) and further processed at different stages of the nervous system to generate the commanded action (Tresilian, 2012). To identify a command that suitably meets a goal however, the actor must first possess a framework that defines the input/output relationship of the motoric architecture, with context of the outside world in which the action will take place. Such a framework has been referred to as an internal model or models which can be used to predict the motor commands needed to achieve a desired state. This topic is discussed later. Integral to this process, there are distinct control strategies that the human system is thought to use for implementing movement under a number of different internal (the body) and external (the environment) conditions and situations.

1.4.2.1 Control strategies

Three control mechanisms (impedance, feedback and predictive control) are thought to be interrelated and recruited in parallel to balance reward (e.g. accuracy) and effort (e.g.

energy expenditure) – this is commonly referred to as the 'accuracy/effort trade-off' (Emken, Benitez, Sideris, Bobrow, & Reinkensmeyer, 2007).

Impedance control

This control strategy describes the modulation of limb joint impedance to increase accuracy by reducing the effects of external disturbances to the human system (Selen, Franklin, & Wolpert, 2009). This is primarily achieved through muscular co-contraction around a joint, resulting in a 'stiffer' joint. Another strategy which can be categorised under impedance control is the tonic stretch reflex, whereby a muscle length threshold can be set centrally and monitored via low-level (peripheral) feedback to control muscle actuation (Tresilian, 2012). At a cost of increased energy expenditure, impedance control is effective when the limb is exposed to novel environments to help maintain stability, a situation that is often encountered during the early stages of skill acquisition. Indeed, as motor skill level improves, limb impedance has been shown to decrease, at which point other control mechanisms may become more dominant (Milner & Franklin, 2005).

Feedback control

Feedback control uses an error signal (the difference between the desired and actual outcome) to make corrections during the execution of a movement (Wolpert, Miall, & Kawato, 1998). Because feedback control needs to determine the actual outcome using a combination of sensory signals such as vision and kinaesthesia, there are inevitable delays in the transmission of signals through the neurophysiological system. Use of just feedback control tends to result in either slow movements, or fast and jerky movements. Some responses to stimuli can occur relatively quickly (some reflex mechanisms occur at a timescale of 10-40 ms) whilst centrally processed responses accrue significantly longer time delays (i.e. in the order of hundreds of milliseconds, or more; (Wolpert & Flanagan, 2010). To carry out fast, smooth movements that are not subject to such time delays requires the use of predictive models of control.

Predictive control

Predictive or feedforward control makes use of an inverse model, an internal representation of the motor system and the environment, to estimate the motor commands required to produce an action (Shadmehr, Smith, & Krakauer, 2010; Wolpert, 1997). The critical advantage of a feedforward mechanism is that actions can be generated rapidly without the constraint of feedback error.

Movement typically involves a composite model involving feedforward and feedback mechanisms: any errors that arise due to incorrect predictions from the inverse model can be compensated with feedback control, and also to inform future predictive actions. This process is reviewed in more detail next in context of internal mechanisms that help drive learning and adaptation.

1.4.2.2 Internal models of sensorimotor control

Motor commands are generally described as being generated via the use of two distinct and interacting conceptual mechanisms: inverse and forward models (Wolpert et al., 1998). As discussed above, inverse models (see Figure 1.4) define the motor commands that are needed to achieve a state (e.g. a position and velocity of the arm), whilst forward models predict the consequences (feedback sensory information) of such motor commands using the current state (e.g. of the arm) and a copy of the efferent signal. Such sensory predictions can be used to make fast corrections using the pre-existing controller (inverse model) before sensory feedback is available (this has been modelled using the 'Smith predictor' control scheme; (Desmurget & Grafton, 2000), and integrated with actual sensory feedback to optimise state estimates and enhance perception (Miall & Wolpert, 1996; Shadmehr et al., 2010).

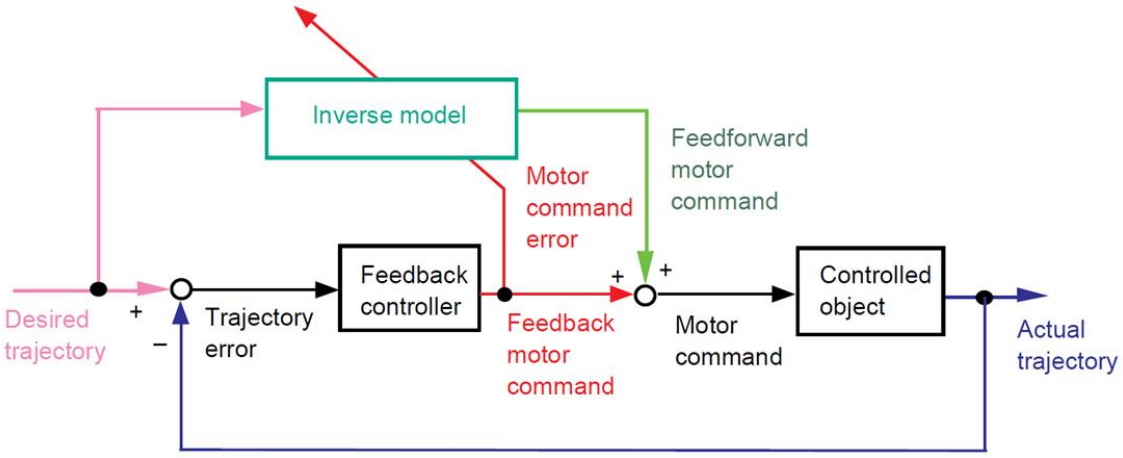

Figure 1.4. The feedback-error-learning model (reproduced from Wolpert et al., 1998). Feedforward and feedback motor commands are used to generate an action. Errors are compensated using a feedback controller. This error-corrective process acts to gradually tune the inverse model, leading to smoother, more accurate and faster movements. This is because movements become controlled by the predictive component of the model and require less corrective action which is slow in comparison.

The process of tuning internal models during interactions with the environment forms the basis of sensorimotor learning. To better understand how to optimise the training of skilled actions (such as MIS) it is important to consider further how this tuning takes place - this is considered further in the next section.

1.4.3 Learning from action

In error-corrective learning, discrepancies between the desired and actual sensory information drive changes of the inverse model, and thus, future motor commands. In this way, perception is intrinsically linked to action. The intimate relationship between feedback and feedforward control mechanisms can be observed during the acquisition of a novel task, such as when learning to ride a bicycle: movements are initially jerky and inaccurate (due to the influence of feedback mechanisms making large, slow corrections). As the inverse model becomes tuned to perform the action, smaller errors occur during predictive control and the need for large feedback corrections decreases. As a result actions become faster, smoother, and more accurate as skill level increases. Linking back to the theme of surgical training, such progressions are often exhibited during the learning of laparoscopic surgery, with experienced practitioners exhibiting faster, smoother and more accurate (i.e. small amplitude deviations from a desired trajectory) movements than novices (Oropesa et al., 2011).

Learning a new skill (such as efficiently interacting with the surgical environment using laparoscopic instruments) is thought to involve two fundamentally distinct and interacting conceptual processes, model-based (MBL) and model-free (MFL) learning (Haith & Krakauer, 2013). MBL is a relatively fast adaptation process involving sampling the dynamics of the environment and adapting internal models to solve the new task (Huang, Haith, Mazzoni, & Krakauer, 2011). MFL is a slower process involving trial and error to sample the environment and directly identifying successful policies to perform an action effectively (i.e. tuning of an inverse model).

MBL and MFL are thought to occur simultaneously during the learning process (Dayan & Daw, 2008). MFL processes are thought to be predominant at the later stages of learning, once MBL has configured the human system in such a way that it performs in the ‘ball-park’ of what is needed. However, to achieve MBL, it is not sufficient for the human system to possess previous ‘models’ that are in some way related to the novel task. Rather, a mapping of the *relationship* between the behaviours needed for the two tasks is

needed. This mapping can be related to what has been referred to in the literature as the 'meta-parameter' in structural learning (Braun, Mehring, & Wolpert, 2010).

1.4.4 Structural learning and the parameter space

Structural learning has been accredited with the remarkable ability of humans to quickly adapt to new environments, for example, learning to ride a motorcycle with minor difficulty having previously learned to ride a bicycle. Performing any given task requires a certain combination of parameters that define the actions needed to perform the task. We will refer to this as the 'parameter space'. Relative to riding a skateboard, for example, it would seem that riding a bicycle would require a similar parameter combination to riding a motorcycle. This is because the physics involved with performing one task are highly relevant to the other, and vice-versa. In support of this, there is evidence showing that motor task variation (i.e. moving from one task to another similar one or, in other words, deliberately changing the parameter space) induces structural learning, which facilitates generalisation to new, yet similar, tasks (Daniel A. Braun, Aertsen, Wolpert, & Mehring, 2009).

To further analyse the concept of parameter space, consider a simplified analogy using two parameters (Figure 1.5): an archer firing an arrow at a target (ignoring horizontal angle and the effect of wind). To hit the target, the archer must find a combination of force and tilt angle which generates the required trajectory. A novice archer will begin by trying different combinations of force and tilt angle and iteratively adjusting the two parameters based on the trajectory of the arrow. Gradually they will approach a suitable combination which will result in the arrow hitting the target. For any given target, a range of solutions may exist for a particular target position. For a different target position, a different combination of parameters will be needed. Again, the archer must explore the parameter space until a suitable combination of parameters is found.

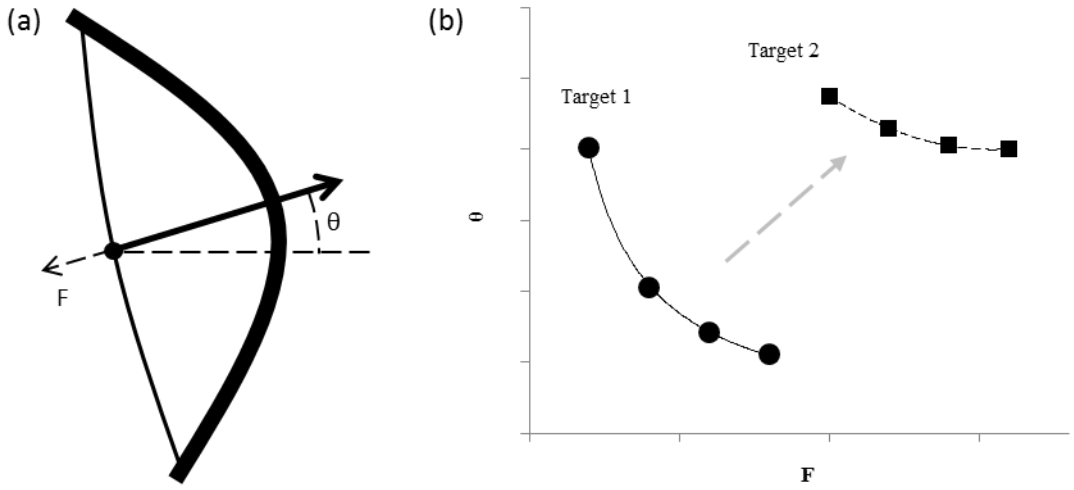

Figure 1.5. Illustration of the parameter space for an analogous and simplified two-dimensional task, showing: a) an illustration of the tilt angle (θ) and applied force (F) parameters when firing an arrow using a bow, excluding horizontal angle; b) a range of solutions made up of different combinations of force and tilt angle are possible for any target position. If the target position is changed, the parameter relationship shifts. To hit a new target (target 2), the learner must find a suitable new combination of θ and F.

Once the archer has experienced the effects of varying the two parameters involved in the task, they are able to build a structure that represents the relationship between the two parameters in terms of the goal (hitting a target at any given distance). By simply modifying a meta-parameter, the learner is able to quickly adapt to a new task (see Figure 1.6).

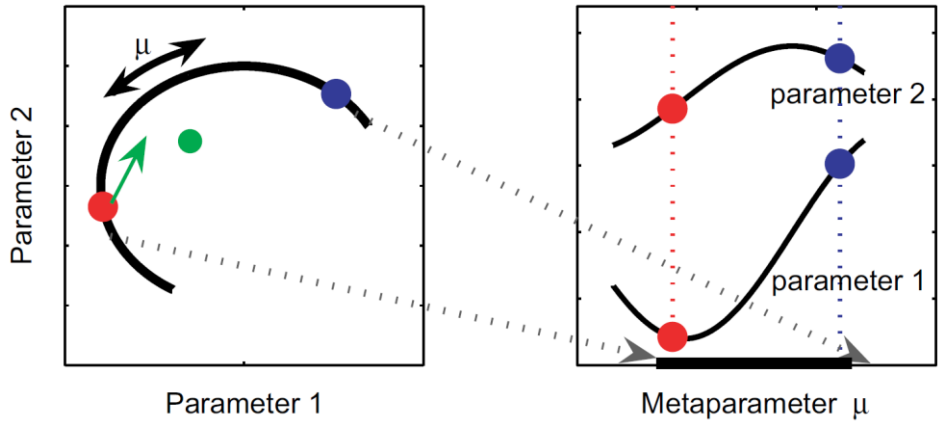

Figure 1.6. Parametric learning requires exploration across the full (multi-dimensional) parameter space to reach a suitable solution, whilst structural learning uses past experiences to infer the relationship between the variables (the meta-parameter), essentially reducing the problem to one dimension, thus greatly simplifying the learning process (reproduced from Braun et al., 2009).

It is important to note that, for simplicity, this is an abstract example. Complex human action, which often requires the synchronised control of tens or hundreds of muscles, conceivably involves operating in a parameter space comprised of hundreds or thousands of dimensions (Braun et al., 2010). This highlights the complexity that the sensorimotor system must contend with in order to quickly learn what actions are best suited to new tasks.

1.4.5 Learning optimal behaviour

To maximise performance efficiency, the motor control system must minimise task complexity and cope with uncertainty by allowing variability in redundant (task-irrelevant) dimensions while minimising variability in the dimensions that have a significant impact on performing the task (Todorov & Jordan, 2002). Thus, a key component to motor learning is finding which errors detract from the goal and what strategy can be used to minimise the effects of such errors. Whilst error-corrective learning is relatively fast and capable of producing a near-zero average error (i.e. it is accurate), the variability that results from this strategy alone can be large (i.e. it is not precise). This is because, whilst the human system has found a solution to the problem, it is not necessarily the optimum solution (i.e. the solution with the lowest combined variable and constant error - to be both precise and accurate).

With reference to the parameter space example (Figure 1.5), an error-corrective learning process provides an efficient way to derive a combination of parameters that suitably meet the task demands. Let's assume that the solution found through structural learning was the one requiring the largest force and lowest vertical tilt solutions shown in Figure 1.5. Whilst the solution appears to achieve the goal (hitting the target) there may be better solutions. A larger tilt angle and a lower power output could also result in hitting the target but with lower overall energy expenditure. Thus, it is not the *optimal* solution (in terms of power output). But once the human system has arrived at a local minima (i.e. the best solution in that region of the parameter space) it may be difficult for the motor system to move away to find the optimal solution since learning a more effective strategy will require further exploration of the parameter space that could result initially in increased errors. One solution would be to provide additional guidance to impose external constraints on the human system to move it around the parameter space towards the most effective strategy.

‘Reinforcement learning’ has been described as a learning process which, in the context of our parameter space example, involves exploring alternative parameter combinations to find alternative (and potentially more effective) solutions to a problem (Dayan & Daw, 2008). However, external stimuli are often required to induce such changes: one way in which reinforcement learning has been achieved is through the provision of knowledge of results (KR) to inform the learner of the effectiveness of a strategy on achieving the overall goal. Such feedback can therefore be used as guidance for future attempts, thus improving response to a stimulus in a relatively long-lasting way (Salmoni, Schmidt, & Walter, 1984). Another process which may possibly occur in parallel with reinforcement and error-corrective learning processes is ‘use-dependent learning’, a mechanism that is driven through repetition of an action without the need for any clearly defined goal or objective. The presence of this process has been shown empirically during reaching tasks (Diedrichsen, White, Newman, & Lally, 2010). With practice, consequent movements become biased towards previous movement patterns and less variable (Verstynen & Sabes, 2011).

It seems likely that teaching an optimal strategy to achieve a goal will involve a combination of error-corrective, reinforcement and use-dependent learning. In clinical training, the growing emphasis on VR surgical training systems should be seen as an opportunity to investigate new methods of doing this which may have not been previously available. One method that has been a matter of discussion over the past decade is the notion of applying forces to a learner’s limb with the aim of accelerating learning. This topic is the focus of this research.

1.5 Haptic enhancement of sensorimotor learning

The advent of robotic technology has stimulated a range of research investigations to determine the best way to support the human sensorimotor learning processes through the provision of haptic feedback. Haptic feedback technology encompasses systems that deliver forces to the body to simulate or augment the forces associated with interaction with objects (Salisbury, Conti, & Barbagli, 2004). Early applications of haptic technology were seen in virtual environments starting with flight simulators and master-slave teleoperated robotic devices (Salisbury et al., 2004). The early 1990’s saw a new wave of haptic technology, that exploited the propensity of the human CNS to integrate visual

and motor information, using a combination of haptics synchronised with graphical interfaces (Brooks,Jr., Ouh-Young, Batter, & Jerome Kilpatrick, 1990). In these environments virtual objects dynamically changed their geometry as a result of the applied forces. Haptic feedback is often complemented by visual feedback in this way due to the way in which vision complements touch and is generally more effective for extracting spatial and contextual information (Ernst & Banks, 2002).

The control architecture of haptic devices is generally split into two types: impedance and admittance controlled. Impedance haptic devices simulate mechanical impedance by measuring linear or angular position and generating a force or torque; admittance haptic devices simulate mechanical admittance by measuring force or torque and generating a linear or angular position (Otaduy & Lin, 2006; Salisbury et al., 2004). Impedance haptic devices such as the PHANTOM range (SensAble) are most common because they are generally simpler and cheaper to produce than admittance devices, such as the HapticMASTER (MOOG FCS), which are better suited to applications requiring high forces and a large workspace (van der Linde, Lammertse, Frederiksen, & Ruiter, 2002).

1.5.1 Implementation of haptic feedback

Modern haptic systems have been used to simulate complex environments such as the forces associated with performing a surgical procedure (Botden & Jakimowicz, 2009). Haptic feedback can be implemented in many ways. Often, a mechanical analogue like a mass-spring-damper system is used (Culmer, 2007; Hogan, 1984), as described by the mass-spring-damper equation (1.1), where F is the output force, k is a spring stiffness, c is a damping coefficient, m is a simulated mass, P is the current position and t is time. An approach like this one allows for the implementation of virtual objects with position and time dependent properties (see Figure 1.7).

$$\boldsymbol{F}(t) = k(P_2(t) - P_1(t)) + c\left(\dot{P}_2(t)\right) + m\left(\ddot{P}_2(t)\right) \tag{1.1}$$

The simplicity of this method allows computationally efficient rendering of haptic environments.

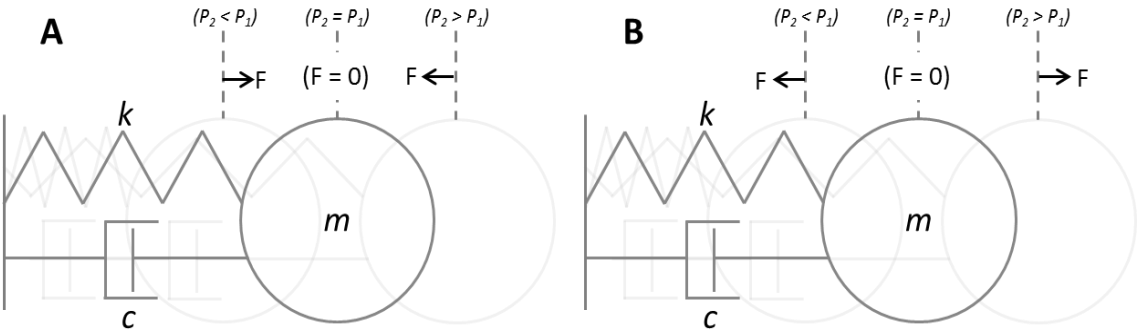

Figure 1.7: Illustrations of a one-dimensional haptic model using the mass-spring-damper method to create a) position-stabilising (error reduction) and b) destabilising (error augmentation) forces.

Simple damping and inertial properties have been used in this way to implement force field effects, such as the 'curl' force field which has previously been used to study the sensorimotor learning process (F. Huang, Patton, & Mussa-Ivaldi, 2007). For more complex implementations of haptic environments, other methods are available. Within the field of surgical simulation, more complex viscoelastic models are needed to mimic the mechanical properties of biological tissue with relatively high fidelity (Brouwer, Mora, & Laroche, 2007). Alternatively, real-time finite element computations have been effective for implementing virtual surfaces of irregular geometry and complex spatiotemporal characteristics though this is at the expense of complex modelling and significantly higher computational power (Cotin & Delingette, 1998; Sedef, Samur, & Basdogan, 2006).

1.5.2 Current use of haptic feedback technology in sensorimotor learning

High fidelity simulation plays a critical role in providing trainees with representative feedback related to the simulated environment. The underlying assumption is that they allow the learner to form internal models which can be easily translated to the real world with minimal adaptation and/or further learning. However, whilst simulation has almost certainly played an important role in developing the sensorimotor skills of trainees (van der Meijden & Schijven, 2009), there has been little progress in leveraging the full potential of haptic feedback technology to increase the rate of skill acquisition.

Theoretically, there are a number of approaches that could be taken with the aim of optimally exploiting the sensorimotor system and increasing learning rates. Options include guiding the arm towards optimal movement strategies (i.e. reinforcement learning), or allowing the user to repetitively sample information from a novel environment to build robust internal models for navigating said environment (use-dependent learning). One example of a training procedure which may induce reinforcement and use-dependent learning is one reported in a study where trainees were required to continuously compare tissue samples of similar compliances (stiffness; Teodorescu, Bouchigny, & Korman, 2013). This task is a critical skill in a variety of clinical settings such as surgery and dentistry. Over several sessions (with knowledge of results) feedback was provided informing participants whether their selection at each comparison was correct or incorrect (Teodorescu, Bouchigny, & Korman, 2013). Improvements over time (as quantified by the perceptibility of smaller compliance differences) were attributed by the authors to 'haptic perceptual learning'. However, they fail to differentiate whether these changes were due to systematic adjustments to the probing actions (kinematic behaviour has previously been identified as an important factor which informs the perception of compliance; (Karadogan, Williams, Howell, & Conatser, 2010), to an increased sensory ability (e.g. sensitisation of the haptic senses; (Tresilian, 2012), or a combination of both of these factors. Further research is needed to investigate the role of reinforcement and use-dependent learning for informing actions over time in clinical settings such as this one. This issue is addresses in Chapters 3 and 4 of this thesis.

Most research in the area of haptic feedback technology for sensorimotor learning has so far focussed on error-corrective learning strategies ((Cesqui, Aliboni, et al., 2008; Chen

& Agrawal, 2012, 2013; Conditt, Gandolfo, & Mussa-Ivaldi, 1997; F. Huang et al., 2007; Melendez-Calderon, Masia, Gassert, Sandini, & Burdet, 2011; Patton & Mussa-Ivaldi, 2004; Patton, Mussa-Ivaldi, & Rymer, 2001; Patton, Stoykov, Kovic, & Mussa-Ivaldi, 2006; Reinkensmeyer & Patton, 2009; Takahashi, Scheidt, & Reinkensmeyer, 2001). This is perhaps due to haptic devices lending themselves to well-defined strategies of adjusting execution error (augmenting or reducing execution error of a task using a haptic feedback device). These interventions can be broadly classified into two categories (see Figure 1.7): error reduction and error augmentation. Examples of the most pertinent work in this area are reviewed next.

1.5.2.1 Error reduction

Error reduction strategies act to provide assistance, supporting the learner to perform a task. A haptic device attached to the learner's limb provides forces which can assist the movement, with the level of possible assistance varying from very small forces providing subtle 'nudges' to behaviour through to full guidance whereby the participant's limb (typically the arm) is essentially passive.

Lüttgen and Heuer carried out an experiment where subjects practised drawing circles with the velocity profile of ellipses (fast on sections of low curvature, slow on sections of high curvature) whilst receiving assistance from a robotic device (Lüttgen & Heuer, 2012b). Whilst the robotic device provided some assistive forces in the direction of the moving target, some level of active control was needed from the subject. The assisted group performed better than the control group during practice on three spatiotemporal metrics. After practice, all of the improvements disappeared with the exception of timing modulation. These findings were confirmed in a similar study by the same authors (Lüttgen & Heuer, 2012a) and also in another study which showed that assistance helped to improve temporal aspects, yet not directional errors of cursive handwriting and putting tasks (Basteris & Sanguineti, 2011). Despite these findings that support haptic guidance for sensorimotor learning, the mechanisms and effects of haptic assistance are not fully understood. It has been proposed, however, that assistance is only useful during the initial stages of learning (i.e. to learn basic aspects of the tasks such as a temporal control strategy), but ineffective once these basic characteristics have been learned (Sigrist et al., 2012).

1.5.2.2 Error augmentation

In contrast to error reduction strategies, error augmentation has been suggested to be most effective at more advanced stages of skill acquisition (Cesqui, Aliboni, et al., 2008). The function behind this technique is to artificially increase errors that arise during the execution of a movement, and thereby intensify the learning process. There have been a number of studies investigating error augmentation strategies. Some of these are outlined below.

In one study (Patton et al., 2001), subjects were required to reach a target whilst moving their arm through a novel viscous force field applied via a robotic device (see Figure 1.8). The initial errors were large but were followed by rapid adaptation, as evidenced by a sharp decrease in execution error (the mean distance of the device from the target path). Upon removing the force field, there was a pronounced after-effect whereby movement errors were opposite in sign to those exhibited before initial adaptation, indicating the presence of error-corrective learning. Similar results involving implicit learning (i.e. the subjects learned movements with minimal instructions) are reported in (Patton & Mussa-Ivaldi, 2004) and (Patton et al., 2006), the latter of which showed that there were no differences in after-effects between healthy and brain injured (hemiparesis due to stroke) subjects. A major disadvantage of implicitly training movements using this technique, however, is that after-effects are not long-lived because de-adaptation occurs once the force field is removed, thus rendering this technique ineffective for the development of long-term skill.

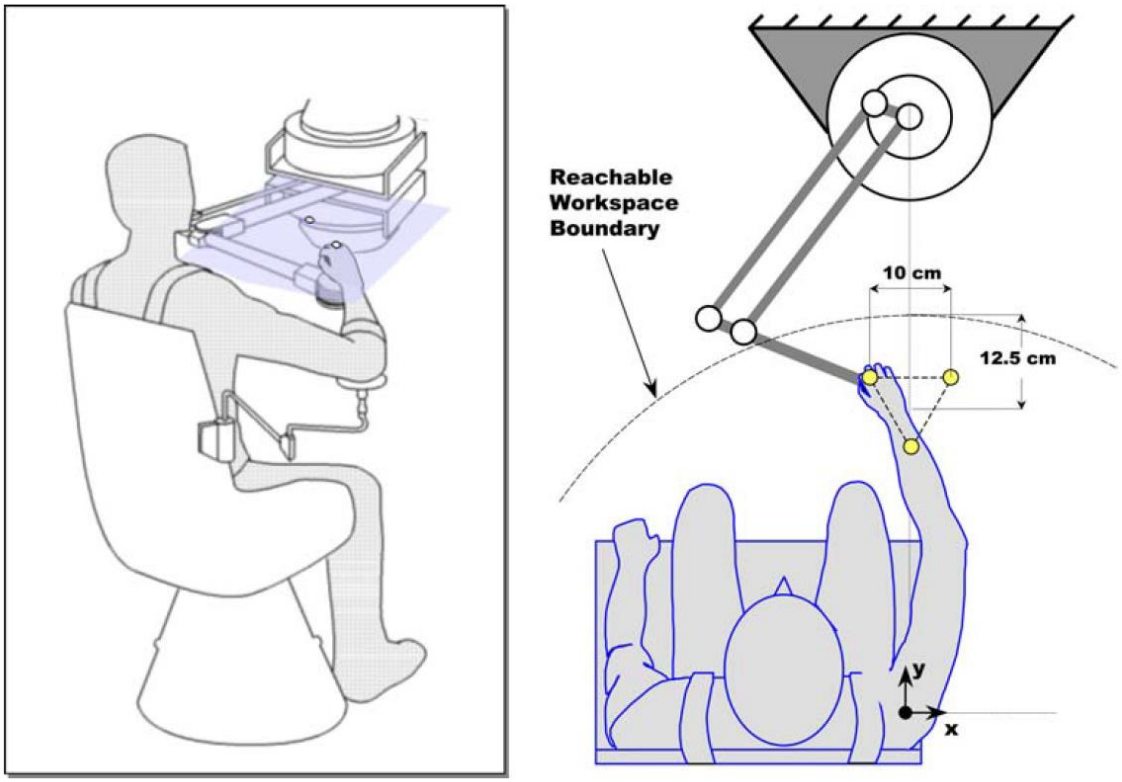

Figure 1.8. Experimental setup used in error augmentation studies testing the effects of novel force fields on learning a sensorimotor task, showing the robotic device and workspace (Patton et al., 2006).

Huang et al. (2007) compared the performance (deviations from a visual path) during a circle tracing task under stable and unstable task dynamics (unstable dynamics were induced using negative damping). The unstable group were able to improve more (they exhibited less trajectory error after training) than the stable group. The authors concluded that the increased performance improvements could be due to increased parameter space exploration induced by unstable dynamics: the unstable damping caused larger velocities and accelerations, thus potentially forcing greater interaction with the inertial characteristics of the task. Complementing these findings, (Patton, Wei, Bajaj, & Scheidt, 2013) tested the effects of different levels of augmented error feedback on adaptation to a 30° visuomotor rotation with healthy adults. Augmentation performed better than a control, but findings suggest that too much augmentation can be detrimental.

1.5.2.3 The roles of assistance and disturbance forces

It seems that there are some differences in the efficacy of assistance (error reduction) and disturbance (error augmentation) for promoting the learning of skilled movements. In an attempt to investigate these differences, several studies have served to contrast the effects of the two techniques. Milot, Marchal-Crespo, Green, Cramer, & Reinkensmeyer (2010) compared haptic disturbance and assistance in a timing task, a computerised pinball-like game, and increased or reduced errors by applying forces to the wrist to retard or accelerate movement. Both the error reduction and error augmentation groups showed improvements, but more skilled subjects showed greater benefits from error augmentation. In a similar study, (Heuer & Rapp, 2011) compared target assistance (an attractive force that pulled the subject towards the target), path assistance (they were pulled towards the path that the target moved along), no assistance and path repellence (forces pushed the hand away from the path) during a visuo-motor rotation adaptation task. After training, direction errors were greater for the assistance conditions than no assistance. The lowest error was shown for the path repellence (error augmentation) group. Similar conclusions were reached in another study (Chen & Agrawal, 2012) which compared 'assist-as-needed' trajectory to repelling (disturbance) forces during a joystick-controlled wheelchair driving task. The assist-as-needed condition provided more assistance when errors were high, and less assistance when they were low.

Despite some evidence supporting the effectiveness of error adjustment forces in accelerating motor learning, there is still debate over whether this is the optimal method and what is the mechanism of action (Reinkensmeyer & Patton, 2009; Sigrist et al., 2012). It has previously been suggested that error augmentation leads to corrective actions which are generally larger and are required more often (Patton & Mussa-Ivaldi, 2004; Reinkensmeyer & Patton, 2009). However, it is unclear why larger and more frequent error corrections could aid learning. One possibility is that larger corrections induce faster learning through greater exploration of the parameter space (F. Huang et al., 2007). Conversely, faster learning effects have previously been attributed to increased 'attention' during these tasks (Sigrist et al., 2012). Further work will be needed to investigate and dissociate these factors. This is considered in the final experiment presented in Chapter 5 of this thesis.

One major criticism of error adjustment strategies is that they effectively alter task dynamics, resulting in the learning of a different task (Winstein, Pohl, & Lewthwaite,

1994). However, some of the findings discussed above suggest otherwise. Whether or not the nervous system is able to filter the superimposed error adjustment forces to obtain information about the underlying task, and to what extent, is a question that merits further investigation. This topic is also addressed in Chapter 5.

In summary, current evidence indicates that haptic guidance supporting error reduction may help during initial stages of learning, and in particular for improving timing aspects of tasks, whilst error augmentation may intensify learning via enhanced error-correction mechanisms and/or induce greater exploration of the parameter space. These findings suggest that, whilst error adjustment may be effective under some circumstances, the optimal method of haptic guidance is likely to depend on the difficulty of the task as well as the proficiency of the participant. However, there is currently no consensus on the best way of determining the most effective method. Thus, further investigations are needed to define the optimal way of using error adjustment techniques (such as the ones described above) to accelerate motor learning.

1.6 Discussion

Interacting with the environment is a two-way process involving obtaining information and implementing suitable behaviours to manipulate said environment. In laparoscopic surgery, information about the environment may be obtained via a combination of visual, haptic and auditory signals. For instance, the force response of a tissue (as measured through a handheld laparoscopic probe along with visual cues) can give an indication of changes in the tissue's structures, and/or indicate the presence of anomalies such as tumours. It seems likely then, that the quality and quantity of information available will depend on the suitability of the action(s) generated for improving the quality and quantity of information available to the CNS. This topic is addressed in Chapters 3 and 4. Furthermore, Accurately and efficiently controlling a laparoscopic probe is critical to safely interact with tissues: the surgeon is required to operate in a complex environment in which movements are mirrored, attenuated and/or amplified (due to the trocar effect). In addition, the operating environment is constantly changing, i.e. due to differences in patient physiology. The need for trainee surgeons to learn increasingly complex procedures, coupled with constraints on available training time, highlights the need to investigate how the learning of skilled behaviour can be accelerated. In moving from

traditional open surgery to laparoscopic surgery there has been a significant increase in the complexity of the sensorimotor skillset needed to perform successful operations. This fact, coupled with restrictions on working hours, has meant that trainees now effectively need to learn more within a shorter time period. This topic is addressed in Chapter 5.

The growing trend of using VR systems in clinical training is an opportunity to explore novel ways of increasing the rate of sensorimotor skill acquisition of trainees. The current focus of these systems so far has been to replicate the surgical environment as closely as possible. However, there is no system in place that intervenes with the trainee's actions or training procedure with the aim of increasing the learning rate. Applying precisely controlled forces to a limb during task execution may provide the best opportunity for rapidly enhancing the sensorimotor learning process. A useful framework for conceptualising this process may be to consider the learner's exposure to the parameter space with the aim to optimally challenge them during practice. Achieving a framework which outlines the implementation of such intervention strategies could play a critical role in future clinical training systems.

The advent of haptic feedback devices has greatly facilitated research investigating the role of haptic feedback technology in motor learning, as they provide an ideal platform to integrate computing with the application of forces to a subject's limb during the execution of movements (Klein, Spencer, & Reinkensmeyer, 2012; Patton et al., 2001). So far, there has been some evidence supporting error adjustment forces as a functional means of accelerating motor learning via enhanced error-corrective processes (Milot et al., 2010; Patton & Mussa-Ivaldi, 2004; Reinkensmeyer & Patton, 2009; Squeri, Basteris, & Sanguineti, 2011). However, this matter is still open to debate. There is a need for more new approaches that investigate the underlying learning mechanisms involved to fully exploit the potential of such interventions.

To this end, there is a clear research need to develop novel, flexible control systems that enable the development of tailored visuohaptic environments, allowing researchers to test the effects of different control paradigms under varying task conditions. To achieve this requires a multidisciplinary approach bridging medicine, neuroscience and engineering. This research has the potential to better establish the fundamental principles of sensorimotor control and learning, and thus obtain better understanding of the approach required in developing accelerated motor learning environments.

This thesis addresses the above points across two themes. The first theme (covered in chapters 3 and 4) investigates novel ways of inducing reinforcement and use-dependent learning via two studies that test compliance discrimination skills within an active-perceptive framework. The second theme (chapter 5) examines the effects of inducing error-corrective processes in haptic interventions on the development of novel sensorimotor skills. Chapter 2 addresses the need for novel tools to investigate the role of haptic intervention techniques on motor learning. The general tools and methodologies used throughout the experiments in the thesis are described. Also detailed is the development of two novel tools for implementing visuohaptic tasks and objectively capturing performance within robust experimental frameworks: the compliance simulation interface (CSI, used in Chapters 3 and 4) and the haptic assessment toolkit (HAT, used in Chapter 5). Chapter 3 examines the effects of long-term, repetitive training with knowledge of results (performance feedback) on a compliance discrimination task using a handheld tool. Chapter 4 investigates the effects of training of this critical skill on a simulated real-world task. Across three experiments, Chapter 5 investigates the role of error augmentation strategies in sensorimotor learning within a framework of parameter space exploration.

CHAPTER 2

GENERAL EXPERIMENTAL METHODS

ABSTRACT This chapter describes the general methodologies used throughout the studies reported in this thesis. Chapter 1 outlined the urgent need for experimental tools to investigate novel clinical training methodologies for sensorimotor learning. The aim was to develop novel tools to enable the implementation of visuohaptic environments within a robust experimental framework. To meet the objectives of the research, two novel software toolkits were developed. First was the compliance simulation interface (CSI), which enables the implementation of virtual surfaces with a specified compliance distribution. Second was the haptic assessment toolkit (HAT), a configurable development platform that allows non-programmers to implement full experimental procedures involving complex visuohaptic environments. An overview of the apparatus, general operating principles and standard experimental setup are given for both systems.

2.1 Introduction

Virtual reality (VR) training systems are fast becoming the option of choice for the development of critical sensorimotor skills in areas such as dental and laparoscopic surgery, and in physical rehabilitation (Bakr, Massey, & Alexander, 2013; K, S, S, Je, & M, 2012; Seymour, 2008). The use of VR systems reduces the need for practicing on cadavers and live subjects during the early stages of skill acquisition, thus placing an intermediary safeguard between the patient and inexperienced trainee surgeons. Recent technological and commercial advances in this area have meant that VR systems are quickly becoming superior to mechanical systems (such as trainer boxes, which include a set of accessories to simulate surgical environments) due to their superior robustness, repeatability, accuracy and precision.

VR training systems generally include visual and haptic feedback that is congruent with the environment that they simulate. Indeed, the focus of these has so far been to replicate real-world conditions. In the case of dental and laparoscopic surgery, trainees complete simulated tasks and procedures that are relevant to their discipline. In physical rehabilitation, robotic systems are used in place of the actions of the physical therapist, namely, manipulating the learner's limb to achieve an action by providing guiding or resistive forces (Jackson et al., 2007).

It is clear that the role of visuohaptic feedback is an essential part of the experience provided to the learner, allowing them to obtain rich information about their environment as well as learning to navigate through it effectively. The relatively slow rate of acquisition of these skills has been of recent concern in clinical areas. Fortunately however, there is some evidence suggesting that visuohaptic systems could play an important role for increasing the acquisition rate of sensorimotor skills (Reinkensmeyer & Patton, 2009). However, the effectiveness of such methodologies is still open to debate, and our understanding of the mechanisms that underpin the learning processes involved is limited (Sigrist et al., 2012). This highlights the need for further investigations in this area. To achieve this, there is a need for robust tools that are capable of generating congruent visual and haptic feedback as well as capturing objective performance data within well-defined experimental protocols.

The experiments discussed in this thesis can be divided into two distinct themes: the first theme (covered in chapters 3 and 4) studies the human ability of compliance assessment using a handheld tool, a critical skill in a multitude of clinical settings. This required a software-controlled interface for the simulation of object surfaces of varying compliance through a haptic device with a built-in handheld tool, as well as capturing of kinematic (movement) data. The focus of this work was on fine (finger and wrist) movements and so relatively low forces and a small workspace were suitable. The second theme (discussed in chapters 5 and 6) is on the use of haptic forces that either enhance or reduce error during gross arm movements. Such movements are analogous to those needed in laparoscopic surgery and physical rehabilitation. In contrast to the previous theme, this work required a relatively large force and workspace, capable of manipulating the mass and movement range of the arm. A flexible software interface was needed to quickly configure and implement complex visuohaptic environments with software-defined force functions, without the need of low-level programming.

To address the requirements of the two themes discussed above, two novel tools were developed for the research in this PhD and are detailed in this chapter. First is the compliance simulation interface (CSI), an interface that allows the simulation of one- and two-dimensional compliant surfaces for tool-based interactions, as well as automatically capturing essential kinematic information about the active behaviours employed to interact with the virtual surfaces. Second is the haptic assessment toolkit (HAT), designed for the study of the effects of force intervention on motor learning. It is a high-level development platform that enables an experimenter to configure, run and analyse experiments with visuohaptic feedback.

2.2 The compliance simulation interface (CSI)

The work in chapters 3 and 4 required a way of generating haptic feedback to simulate the force response of soft, human-like tissues when using a handheld probe. More specifically, the experiment described in Chapter 3 required a method of successively comparing two virtual objects of different compliances. The compliance difference needed to be controlled programmatically within a defined experimental protocol. Logging of position data in the vertical axis was required to carry out kinematic analyses of the probing movements. The experiment described in Chapter 4 required the simulation

of three-dimensional tissue volumes with or without embedded tumours. In contrast to the requirements of Chapter 3, the stiffness distribution along the surface of the simulated tissue was not homogeneous (i.e. the tumour was represented by a region of significantly greater stiffness than the surrounding tissue). This section describes the general work carried out to implement these requirements. First, a breakdown of the aims and objectives is given, followed by detailed technical specification for the CSI, and finally a description of the system.

Most of the basic technical developments for the CSI were carried out during a team Master's project investigating the effect of haptic augmentation on the perception of simulated tumours (see Appendix 1). Within this, the role of the author was modelling and rendering of the haptic environment. The main changes made during this PhD were scaling the functionality of the dynamic link library (dll) from containing a selection of pre-defined surface/tumour combinations to directly specifying Gaussian parameters (see Section 2.2.4). Hence, the flexibility of the system was significantly increased, allowing for the implementation of any tissue/tumour combination for which modelling data are available.

2.2.1 Aim and objectives

The overall aim was to develop a framework to enable the implementation of virtual compliant volumes of either homogeneous or, for the case of surfaces with embedded tumours, variable compliance across the surface plane. A handheld tool was to be used to interact with the virtual environment whilst movement kinematic data were captured for further analyses. To achieve this aim, a number of key objectives were outlined:

- Identify and procure a haptic device with an attached tool, allowing free rotation of the wrist and fingers, and capable of generating an upward force in the vertical (Normal) axis. The force output range should enable simulation of the force responses of soft human tissue. The size of the workspace should allow hand and finger movements.

- Achieve two-way communication between the haptic device and a computer and bespoke control program (i.e. to set force parameters of the device, and read position), allowing integration into a full experimental environment. The system bandwidth should allow the stable generation of forces to simulate the compliant tissue-like surfaces.
- Programmatically define functions to generate stiffness in the vertical (Normal to the surface) axis. This should allow the simulation of spherical inclusions of differing stiffness and size to be embedded within soft tissue.
- Develop software routines to allow programmatic control over the behaviour of the haptic device (i.e. to initialise, operate and shut down the device).

2.2.2 Technical specifications

The aims and objectives were used to create the main technical specifications for the CSI, as detailed below.

2.2.2.1 Haptic device

Range of movement

To achieve the required range of movement (rotation of the wrist and finger movements), the device should enable movement in 5 degrees of freedom (up/down, left/right, in/out, side to side rotation, and forwards/backwards rotation).

Force output

The Young's modulus of human liver has been reported to be approximately 1 kPa, whilst that of tumours is approximately 75 kPa (Mueller & Sandrin, 2010). For a simulated probe size of 10 mm (this is typical of laparoscopic tools) and an indentation of 15 mm directly above a simulated tumour, a maximum force output of 1.8 N in the Normal axis was required (see Appendix 1 for the calculations).

Workspace size

The experiments required hand and finger movements of 20 mm in the Y axis, and a square area of 100x100 mm in the X and Z axes. Thus, this is the minimum required workspace size for the device.

2.2.2.2 Control software

The control software should allow communication with the haptic device and all other peripheral components of the system (monitor, mouse and keyboard).

2.2.2.3 System bandwidth

Haptic feedback

In order to generate smooth and realistic forces, a minimum update rate of 500 Hz is required (Salisbury et al., 2004).

Visual feedback

The visual scene should be updated at a minimum rate of 30 Hz to achieve smooth and realistic visualisations.

2.2.2.4 Physical setup

The system should allow for participants to sit at a desk with the monitor in front of them, and the haptic device where it could be reached comfortably by the right hand.

2.2.3 The haptic device

The PHANTOM OMNI® (SensAble Technologies, Geomagic), shown in Figure 2.1, is a low cost device that is part of the PHANTOM range of haptic devices. Designed for interactions with soft virtual objects via a handheld tool, it is a 6 degrees-of-freedom (DoF, 3 active and 3 passive) device with a built-in handheld tool and a workspace of 160 x 70 x 120 mm in the frontal (X, left-right), longitudinal (Y, up-down) and sagittal (Z, in-out) axes, respectively. It is capable of a maximum output force of 3.3 N in any axis, and a maximum stiffness of 1.26, 2.31 and 1.02 N/mm in the X, Y and Z axes, respectively ("Dental Lab Home," n.d.). From these specifications, it is clear that this device meets all of the minimum technical specifications for the CSI system's haptic device (refer to Section 2.2.2).

Figure 2.1: The PHANTOM OMNI® (SensAble) is a commercially available haptic device with a built-in handheld gimbal (tool). Three internal motors are used to adjust the force at the tooltip to simulate the feel of soft virtual objects. Position encoders are used to sense the position of the tooltip as the user navigates through a virtual 3D environment using wrist and finger movements.

The PHANTOM OMNI uses an impedance-control paradigm (see Figure 2.2). Impedance-controlled devices are highly stable and suitable for the implementation of very small forces. However, the haptic fidelity of impedance-controlled devices is somewhat limited due to their mechanical characteristics. The mass and friction of the device can be felt by the user, which means that they are generally lightly built and are highly backdriveable to maintain a realistic haptic environment (van der Linde et al., 2002). Thus, their mechanical design tends to limit impedance-controlled devices to applications requiring relatively low forces and small workspaces.

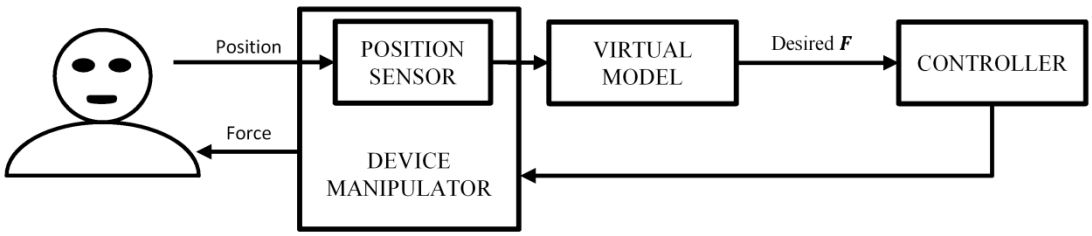

Figure 2.2. The impedance control paradigm is to output force (F) as a function of the device position. Position sensors are used to measure the position of the device's manipulator as the subject applies forces to move it. An internal model is used to calculate the force vector that would result from interactions with the virtual environment. A controller then adjusts the force vector applied by the device.

The force and stiffness range and workspace size of the PHANTOM OMNI makes it ideal for use in this application, as outlined previously in the technical specifications.

Communication with the device is achieved using the OpenHaptics Toolkit (SensAble Technologies), an application program interface (API) designed to work with all PHANTOM devices. It is open source and available for download from http://www.dentsable.com. The toolkit includes the PHANTOM device drivers and Haptic Device API (HDAPI). The OpenHaptics Toolkit enables the implementation of haptic effects (such as forces, springs and dampers) using C++ commands. With reference to our description of the impedance-control paradigm, the HDAPI toolkit acts as the Controller, taking a desired force vector and sending electronic signals to the device actuators to implement the haptic environment.

2.2.4 Generation of haptic stimuli

A schematic description of the CSI is shown in Figure 2.3. A custom dynamic link library (DLL) was developed to enable two way communications with the OpenHaptics HDAPI library to read position from, and write a stiffness value to, the haptic device. With reference to our diagram of the impedance control paradigm (Figure 2.2), hapsurf.dll acts as the 'virtual model' which calculates a force vector as a function of position based on pre-defined functions to simulate a compliant object's surface.

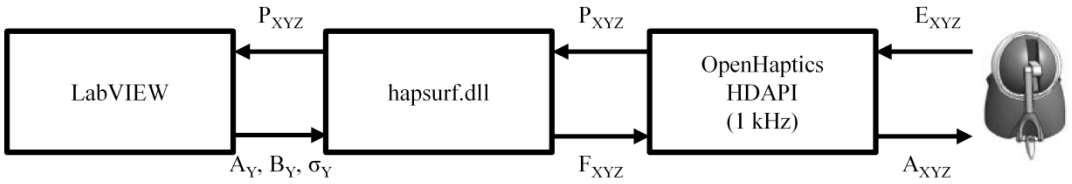

Figure 2.3. The CSI is comprised of a series of LabVIEW functions that communicate with a bespoke dynamic link library (hapsurf.dll) to access the OpenHaptics HDAPI functions. This achieves programmatic behaviour of the device. LabVIEW communicates with hapsurf.dll to write the Normal (i.e. acting in the Y axis) force parameters (A_Y, B_Y and σ_Y) of a surface and read the 3D position vector (P_{XYZ}) of the device. hapsurf.dll uses the surface model parameters to calculate the force vector at the device tool tip as a function of the device's position. The OpenHaptics HDAPI generates the required force vector (F_{XYZ}) specified in hapsurf.dll by sending electronic control signals (I_{XYZ}) to the device actuators and receiving position feedback (O_{XYZ}) at a rate of 1 kHz.

In line with Objective 4, the requirements for the virtual model were to allow the implementation of either a homogenous or variable stiffness along the plane of a virtual object's surface. The latter is required to represent the force response of an object with an embedded inclusion (specifically, representing a section of tissue with an embedded tumour). Previous work carried out during a fourth year engineering undergraduate Master's team project investigated the force response of an elastic object's surface with embedded inclusions. Computational finite element analysis (FEA) was used to parametrically calculate the force response of the surface to indentations with a spherical probe at various locations (details on the modelling process are available in Appendix 1). Surface stiffness, tumour stiffness, tumour size, and tumour depth were parametrically varied and a response force surface was produced for each. A Gaussian function, described in equation (2.1), was then fitted to each surface to calculate force, F, as a function of peak force, A, relative to an offset (baseline) force, B, the distance between the end-effector and inclusion centre, x_r, and the function's width, σ. This was deemed an effective method for modelling volumes of tissue with embedded tumours during the preliminary work described in Appendix 1 (see Figure 2.4).

$$F = Ae^{-\frac{x_r^2}{2\sigma^2}} + B \tag{2.1}$$

The baseline term (B) represents the force of a volume of tissue (i.e. a flat force response), whilst the peak force and exponential terms act to superimpose the response of an embedded inclusion.

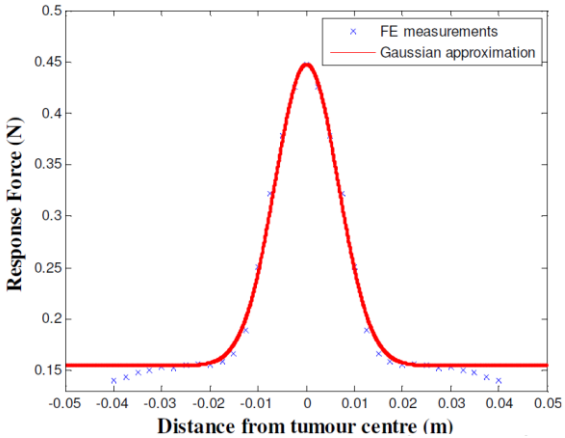

Figure 2.4. Gaussian approximation to FEA force data for a 12 mm diameter embedded tumour with 10 mm indentations. RMSE = 0.0036 N (Reproduced from (Chandler, J., Dickson, M., Jamieson, E., Mueller, T., Reid, T., Unpublished).

The response force profiles of various surface/inclusion combinations are shown in Figure 2.5. A linear spring model (f = kx, where f is the output force, k is the stiffness coefficient and x is indentation depth) is then used to scale the force as a function of indentation depth. This method provides a simple means of specifying a surface with homogeneous stiffness (i.e. by using an A value of 0) and a positive value for B, whilst a positive value of A will result in the simulation of an inclusion. Variable x_r is calculated by accessing the OpenHaptics Toolkit, allowing the smooth update of forces and realistic haptic rendering. Implementation of different surfaces is achieved by sending parameters A, B and σ (defined above) to hapsurf.dll (see Figure 2.3).

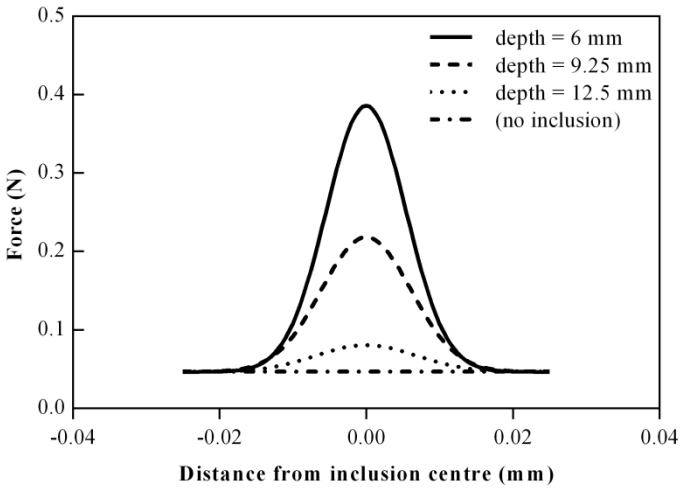

Figure 2.5. Response force profile of an object with a Young's Modulus of 1 kPa (similar to that of human liver), indented by 3 mm with a 10 mm diameter spherical probe of infinite Young's Modulus. Where included, a spherical inclusion of 12 mm diameter with a Young's Modulus of 75 kPa (similar to tumours typically found in human liver) is embedded at various depths beneath the surface. Depths quoted are from the object's surface to the centre of the inclusion, i.e. at a depth of 6mm the inclusion is flush with the surface.

LabVIEW (National Instruments, USA) is a graphical development platform with integrated functionality to interface with DLLs, generate visual and other (e.g. auditory) feedback, and process data. This made LabVIEW a suitable candidate for implementing the experimental protocols described in Chapters 3 and 4. To enable control of the haptic device from LabVIEW and thus complete the architecture needed to implement surfaces of defined compliances, a series of custom discrete LabVIEW functions were developed to read position, write (Gaussian) forcing parameters, and safely shut down the device.

2.2.5 Experimental setup

To promote consistency across the experiments presented in Chapters 3 and 4, a general experimental setup (shown in Figure 2.6) was defined. Participants sat in a chair in front of a 740 mm tall table. Visual stimuli (generated using LabVIEW software) were presented on a computer monitor located on the table, directly in front of the participant and approximately 600 mm from the table edge. The haptic device (PHANTOM OMNI®) was located to the right of the monitor, approximately 500 mm from the table edge (this distance was varied slightly depending on the participant's arm length and was moved for their comfort). Participants were instructed to hold the device tool in their right hand using a standard pencil grip. An arm support was placed under the right arm so that it rested 65 mm above the table. A wrist support was located on top of the arm support which heightened the wrist by an additional 20 mm to limit movements to wrist and finger actions.

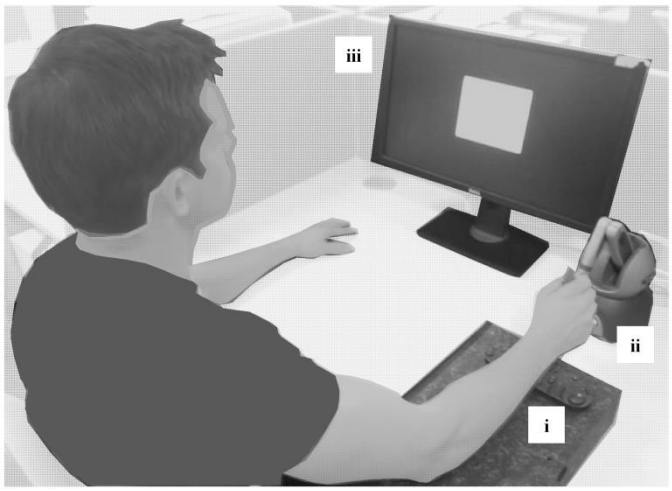

Figure 2.6. Diagram of the experimental setup used for all compliance discrimination tasks (Chapters 3 and 4), showing i) the arm and wrist supports, ii) the haptic feedback device, and iii) the position of the computer monitor and graphical display.

Validation of the haptic fidelity of this method was carried out in previous work (see Appendix 1). In accordance with the objectives, this architecture makes it possible to reliably simulate the compliance of soft surfaces with or without embedded inclusions. The force update rate (1 kHz) allowed for the implementation of smooth and realistic forces, whilst the maximum output stiffness (2.31 $N.m^{-1}$) met and exceeded the specifications required to simulate the force response of a liver.

2.3 The haptic assessment toolkit (HAT)

Previous research has suggested that haptic forces that act to manipulate the error signal could play an important role in increasing the rate of motor learning (Reinkensmeyer & Patton, 2009). However, the effectiveness and underlying mechanisms involved in this process are widely debated topics in the literature (Sigrist et al., 2012).

Chapter 5 investigated the role of haptic intervention on the rate of motor learning. The requirements of this work were to establish a reliable method of controlling different force-based intervention strategies that act to apply a force to the subject's arm during the execution of a task, whilst providing congruent visual feedback, within a robust and configurable experimental protocol.

D-Flow (Motek Medical) is one successful example of a high-level software interface that allows researchers to integrate different input and output hardware components to build an interactive, virtual environment where a human is part of a real-time feedback loop (Geijtenbeek, Steenbrink, Otten, & Even-Zohar, 2011). The system supports the integration of haptic devices and allows an operator to configure a visuohaptic environment, as well as to store kinematic data. However, D-Flow is limited to virtual models of objects (using planes, spheres, etc.) and effects (such as springs and dampers) to build an interactive haptic environment (van der Linde et al., 2002). These environments are dynamically fixed which significantly limits the form of haptic feedback algorithms that can be implemented. Another potential limitation of D-Flow is that it does not automatically generate objective measures of motor performance.

One novel solution which generates objective measures of motor performance is the Kinematic Assessment Tool (KAT) (Culmer, Levesley, Mon-Williams, & Williams, 2009). KAT is capable of measuring human movement in configurable visual-spatial tasks and automatically outputting performance measures. The system was designed specifically for the assessment of motor skills and has been used extensively within this domain, e.g. to objectively quantify handwriting performance (Flatters et al., 2014). Unfortunately however, KAT does not support the integration of haptic feedback. A solution is required, therefore, that combines D-Flow and KAT, in order to enable the construction of bespoke experimental procedures with haptic force fields, as well as

obtaining objective measures of human performance. As outlined above, the requirements for this system vary significantly from that of the CSI, and so a separate system is needed.

This section describes the development of the HAT. The toolkit integrates haptic and visual feedback, data acquisition, real-time processing and data management within a flexible software platform. It is designed to allow the experimenter to quickly and easily create novel visuohaptic environments for use within experimental procedures.

2.3.1 Aim and objectives

The overall aim was to develop a system that can be used by researchers to configure, run and analyse experiments using configurable haptic force fields and congruent visualisations. This gave rise to a number of key objectives, as outlined below.

- Identify and procure a haptic feedback system with a suitably-sized workspace to allow gross arm movements and a force output that is capable of accurately moving the mass of a resting (passive) human arm.
- Develop a software interface to control the behaviour of the haptic device.
- Simulate the behaviour of a mechanical dynamic (mass-spring-damper) system.
- Provide smooth visual feedback that is congruent with the task environment.
- Enable the integration of bespoke haptic feedback algorithms.
- Allow non-specialists (non-programmers and non-engineers) to configure novel visuohaptic environments via a robust user interface.
- Automatically store kinematic data, and generate objective performance measures.
- Define a method of implementing force fields which act to distort the force distribution along the device's physical workspace (this would serve to induce task novelty).
- Define a method of generating error adjustment forces, i.e. forces that push or pull the subject away from (error augmentation) or towards to (error reduction) a target position where the force changes as a function of the distance between the device and target positions.

- Develop software routines to allow programmatic control over the behaviour of the haptic device (i.e. to initialise, operate and shut down the device).

2.3.2 Technical specifications

A set of specifications were outlined for the HAT system based on the objectives outlined above.

2.3.2.1 Haptic device

Range of movement – the range of movement required for this system is relatively simple: it should allow for gross arm movements in three axes. Thus, the device should have three degrees-of-freedom to account for this.

Workspace size – The system is required to allow for gross arm movements. For this, an estimated minimum workspace size of 250 x 250 x 250 mm was required.

Force output – the force output of the device should enable moving of a passive arm around the workspace. Accounting for the instantaneous forces required to accelerate a passive arm (estimated to have a mass of approximately 50 N) at a maximum rate of 2 m/s^2, an estimated minimum output force of 100 N is needed (F = ma), excluding the force required to move the device end-effector.

2.3.2.2 Control software

The control software should allow communication with the haptic device and all other peripheral components of the system (monitor, mouse and keyboard).

2.3.2.3 System bandwidth

Haptic feedback – In order to generate smooth and realistic forces, a minimum update rate of 500 Hz is required (Salisbury et al., 2004).

Visual feedback - The visual scene should be updated at a minimum rate of 30 Hz to achieve smooth and realistic visualisations.

2.3.2.4 Development platform

The system should ideally provide the experimenter with a simple method of developing visuohaptic environments with bespoke haptic interventions to 1) systematically adjust execution error, and 2) implement force fields.

2.3.2.5 Physical setup

The physical apparatus should be setup so that the haptic device is located directly in front of a standing or seated participant, with visual stimuli shown clearly in front of them, preferably above the haptic device.

2.3.3 The haptic device

Previous researchers in this area have often developed bespoke haptic devices (Patton et al., 2006; Reinkensmeyer & Patton, 2009). However, this was outside the scope of this project and so a commercial device was needed. Based on the technical specification for the HAT, the HapticMASTER was selected (see Figure 2.7). It is a commercially available haptic feedback device specifically designed for haptic interactions with virtual environments involving gross arm movements, with a workspace of 0.36 m x 1 rad x 0.4 m in the X, Y and Z axes, respectively. It has three active degrees-of-freedom (DoF) and a position resolution of < 4 μm, force sensitivity of < 0.01 N and maximum output force of 250 N (van der Linde et al., 2002). This device meets and exceeds the minimum technical specifications outlined in Section 2.3.2.

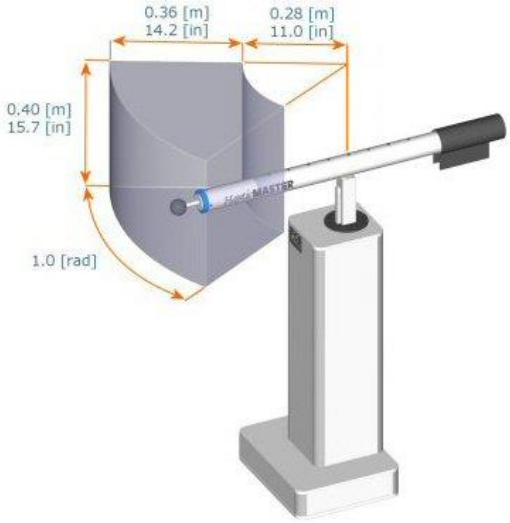

Figure 2.7. The HapticMASTER robotic arm and workspace (reproduced from "VR Laboratory - University of Twente," n.d.)

The HapticMASTER is an admittance-controlled haptic device, which works to measure force and output position (see Figure 2.8). In contrast to impedance control, this makes the HapticMASTER suitable for the implementation of large forces, as well as for compensation of friction and gravity. However, finite force sensing accuracy and resolution limits the implementation of small forces (van der Linde et al., 2002).

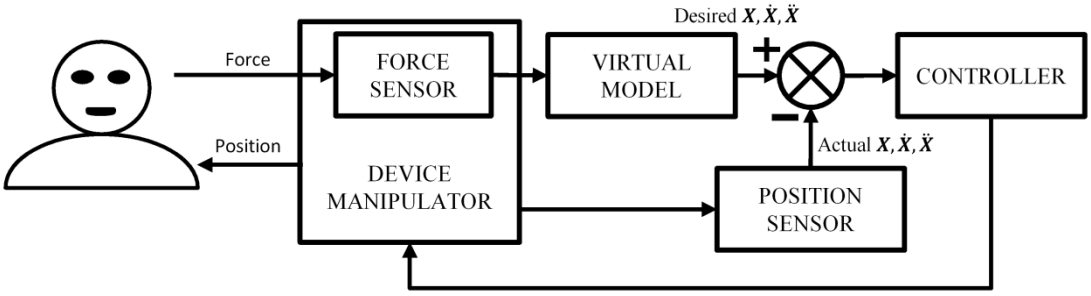

Figure 2.8. The admittance control paradigm is to measure force and output position. A force sensor measures force applied to the manipulator of the device. An internal model is used to calculate the force, velocity and acceleration (kinematics) that would result from the applied force to the virtual environment. A controller adjusts the device's position as a function of error (the difference between the device's actual and desired kinematics).

The HapticMASTER's force range and workspace makes it suitable for moving a fully passive (relaxed) arm, or counteracting forces applied by an active arm. This makes it suitable for this application, as outlined in the aim and objectives.

The device includes an API (the HapticAPI), and a DLL (hapticAPI2.dll) which enables access to the HapticAPI functions. This allows the creation of haptic effects (bias forces, springs and dampers) and objects (spheres, cubes, cylinders and toruses), as well as setting the global damping, inertial and frictional properties of the device. More information on the HapticAPI can be found in (MOOG, 2011).

2.3.4 The haptic device interface

A software interface was developed in LabVIEW to achieve two-way communication with the HapticMASTER device, allowing the creation of effects such as forces, springs and dampers, as well as the measurement of three-dimensional position in real-time. HapticAPI2.dll was used to access the built-in functions available with HapticAPI from LabVIEW. A schematic overview of the interface is shown in Figure 2.9.

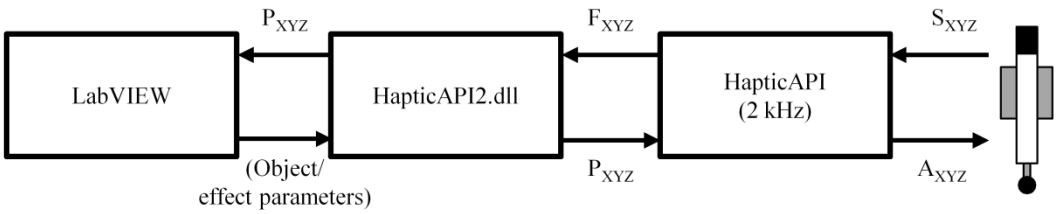

Figure 2.9. The LabVIEW-HapticMASTER interface is a series of LabVIEW functions which communicates with hapticAPI2.dll to write object or effect parameters to HapticAPI, and to read a 3D position vector (P_{XYZ}). Actuation signals (A_{XYZ}) are sent to the device motors, and sensory signals (S_{XYZ}) are received back to compute a force vector (F_{XYZ}). The HapticMASTER's integrated real-time controller runs at a rate of 2 kHz to achieve high-fidelity haptic feedback for any objects or effects that are specified from LabVIEW.

Bespoke functions were created to give access to the high-level functions of the device. This allows the generation of haptic effects and objects, as well as initialising and safely shutting down the device. The LabVIEW-HapticMASTER interface was used as the basis for creating haptic feedback in HAT.

2.3.5 Description of the HAT system

The HAT system is designed around the delivery of interactive visual-spatial experimental trials in which visual and/or haptic stimuli are coordinated with the movements of the haptic device. An illustrative example of the functionality of HAT within an experimental procedure is given in Figure 2.10.

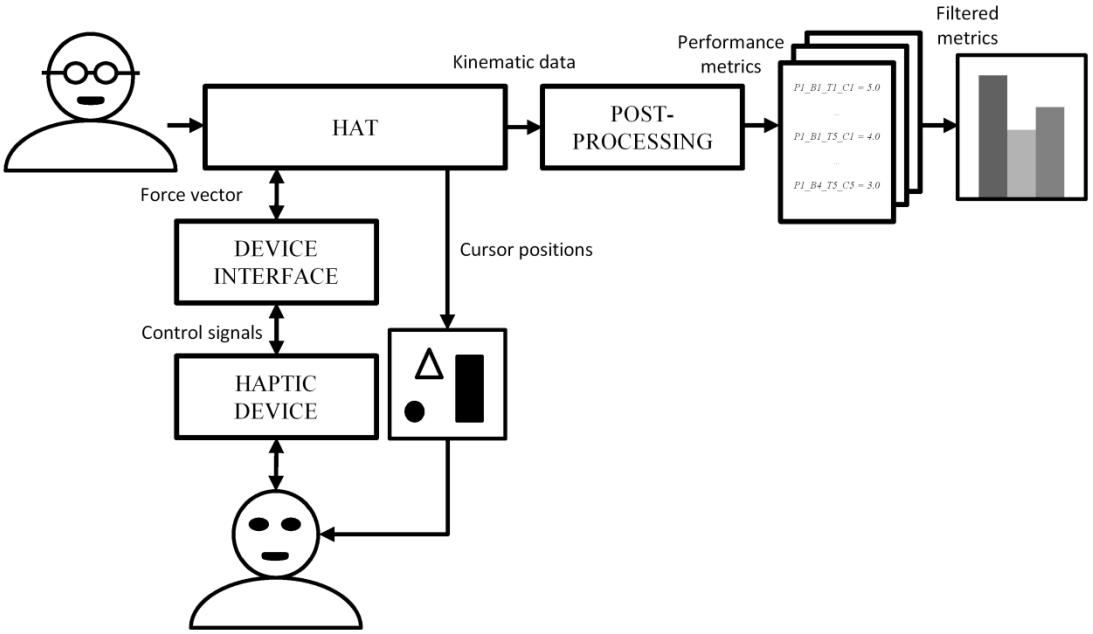

Figure 2.10. High-level illustration of the role of HAT within an experimental process. The experimenter (top left) inputs information into HAT to configure an experiment. HAT communicates with a haptic device via a software interface to send and receive haptic information (position/force) to provide haptic feedback to the subject. Visual information is also generated within HAT and presented to the subject. Raw kinematic data are recorded during the experiment and a post-processing utility is used to output data in a pre-specified format for further analyses.

HAT is designed to operate on any Windows operating system (Microsoft, v7 or later). It was developed and tested on a desktop computer (Intel quad-core i5-2400, 3.1 GHz, 4 GB RAM, Intel HD Graphics Family). Refer to Section 2.3.6 for a validation of the system's performance under this configuration.

2.3.5.1 Trial structure and design

The configuration of an experiment is built up as a hierarchical structure comprised of any number of movements along pre-specified spatial trajectories (defined as a set of Cartesian points within the workspace) with bespoke haptic force functions. Experiments are defined in five tiers, as illustrated in Figure 2.10: at the lowest level, a 'node' defines a point within the workspace, identified by Cartesian coordinates. Next, a 'component' defines the nodes of an individual movement. An example of a component and its constituent nodes can be seen in Figure 2.11 . It is defined as a single trajectory with target start and end locations, and any number of positions (nodes) in between. Table 2.1

describes the configurable parameters of each component relating to the haptic environment.

Table 2.1. List and description of haptic component parameters

Parameter	Description
Node coordinates	Array of Cartesian coordinates for nodes
Movement speed	Movement speed of the target cursor
Mass/spring/damper settings	Coefficients of mass, stiffness and damping
Force field algorithm	Reference to a force field algorithm sub-routine
Maximum force	Absolute operating force envelope of the force field
Workspace distortion	Type (sine/square/triangle/sawtooth wave), period (mm), phase offset (rad), amplitude (N), offset (N) and angle (rad).

Once the end node of the component is reached, the next component is loaded automatically. A 'trial' is a full cycle of a trajectory, as specified by individual components. A 'section' is made up of any number of trials. An 'experiment' is comprised of any number of sections (see Figure 2.11).

This hierarchical structure makes it easy to re-use whole or parts of other experimental sections to construct a complete experimental procedure using the built-in configuration utility (see Section 2.3.5.3). Nevertheless, the operator is able to configure parameters down to the level of individual components. Message prompts and images can also be displayed at any stage of the experiment to, for example, provide instructions to the subject.

2.3.5.2 Software architecture

The toolkit was developed using LabVIEW (version 2013, National Instruments). All computations relating to the visuohaptic environment are carried out within the software. Figure 2.11 shows a high-level illustration of the structure of a full experiment, as well as the processes involved in running an experiment.

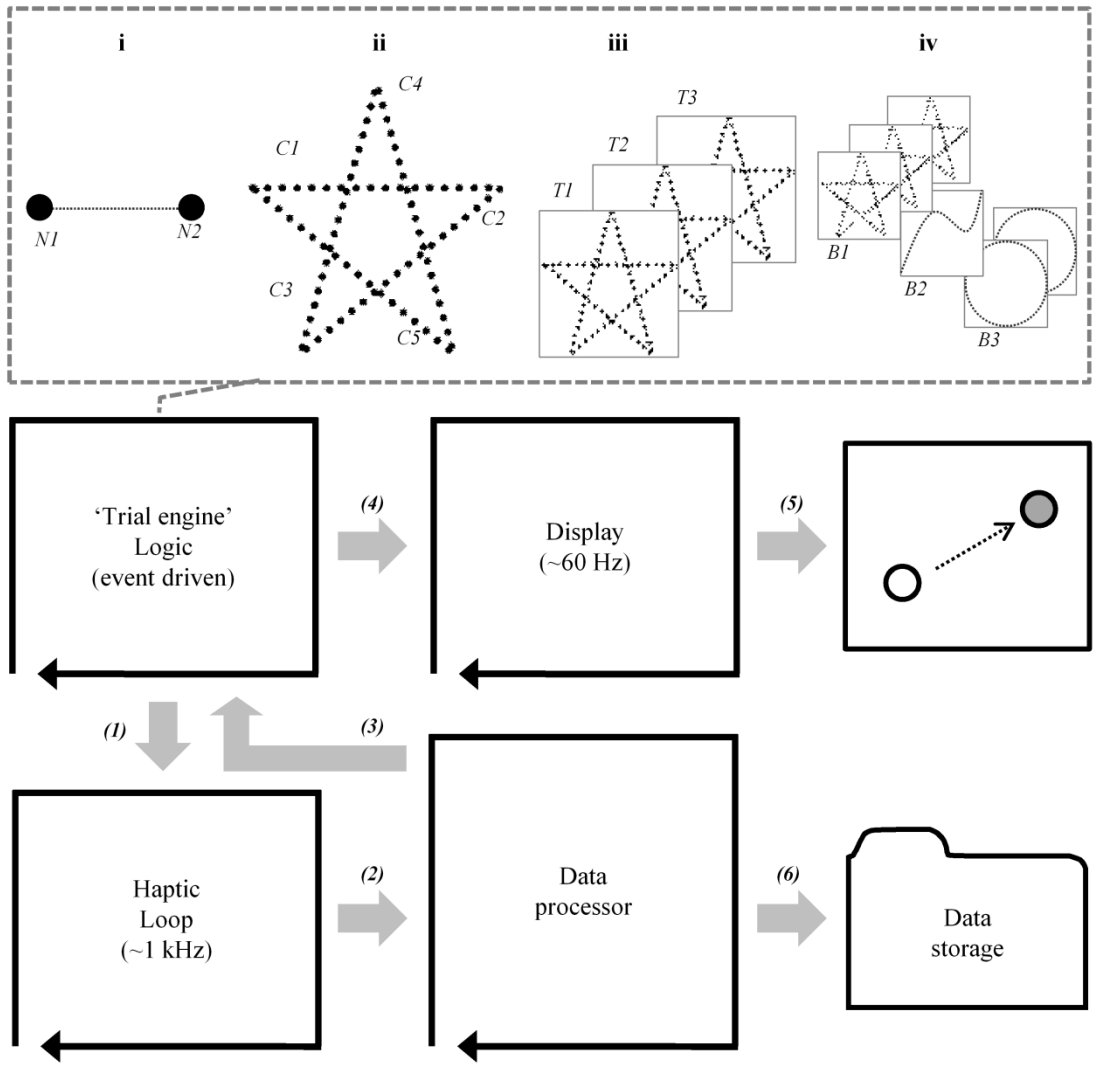

Figure 2.11. Illustration of the main processes carried out to implement a component movement. An experiment is configured as a five tier structure comprised of i) nodes (denoted 'Nx'), ii) components ('Cx'), iii) trials (Tx), and iv) sections (Sx). The trial engine uses the hierarchical experiment configuration to generate and send component parameters to the haptic loop (1). The haptic loop interfaces with the haptic device, generates forces and acquires position data at 1 kHz. Raw kinematic data are then sent to the data processor to transfer computational load away from the haptic loop (2). Task data (vectors of current position, velocity and applied force) are sent to the trial engine (3). Visual data are sent to the display module (4), which performs OpenGL rendering (5) and displays visualisations related to task information. Raw and processed data for each component are sent to the data storage module (6), where they are compiled and saved to file.

Upon initialisation, the trial engine loads stored configuration data containing the parameters for each component of the experiment. This defines the behaviour of the system throughout the experiment. After carrying out initialisation procedures (e.g. device initialisation, display instructions to subject, etc.), the parameters of the first

component movement are sent to the haptic loop to initiate the task. Upon reaching the end node of each component (to within a pre-specified tolerance), the trial engine loads and sends the parameters for the next component.

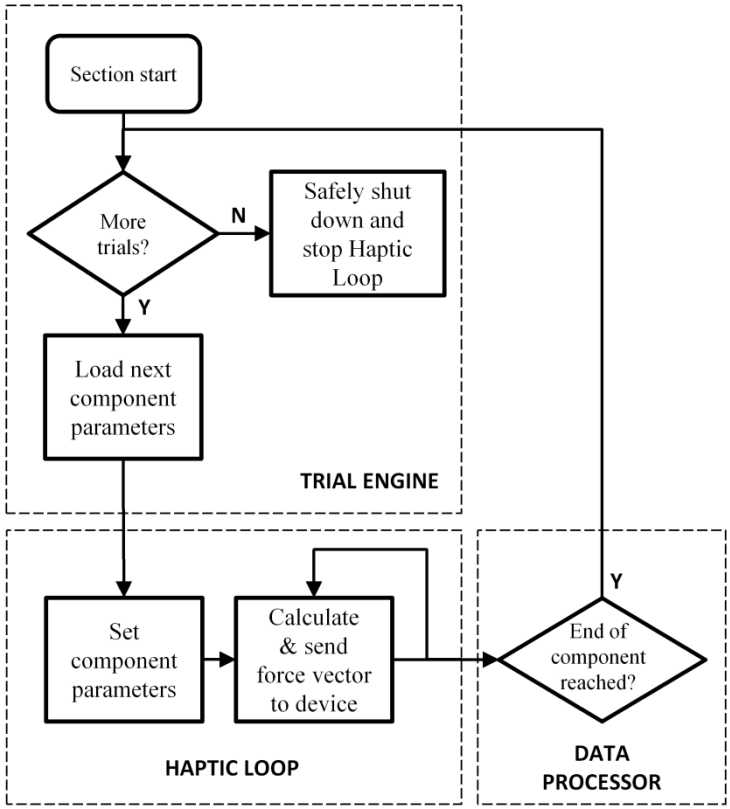

Figure 2.12. Simplified flowchart of the process used to implement the haptic feedback of a Section of an experiment, showing interactions between the trial engine, haptic loop and data processor. The Trial Engine loads new component parameters and sends them to the Haptic Loop. The Haptic Loop constantly calculates the force vector for the current component and sends it to the device until new component parameters are received, or a shutdown command is received. Current position data is sent from the Haptic Loop to the Data Processor, which monitors the current position against the end location of the current component. Upon the device reaching the end location of the component, an event is registered and sent to the Trial Engine. If more trials are present, the next component is loaded. Otherwise, the Section is ended by stopping the Haptic Loop.

There is an option to use an adaptive controller to dynamically adjust component parameters based on measured performance data (see the 'Data processing' section for more information). This process is repeated until all components have been completed. Any shut down procedures (e.g. subject debrief, device shutdown, etc.) are then performed and the experiment ends.

2.3.5.3 Experiment configuration utility

A configuration utility was implemented to allow the configuration of the haptic environment quickly and easily and without the need to access low-level code. The Test Configuration Utility allows the experimenter to create and save experiment configuration files, which can then be loaded by the system to set all parameters for an experiment. Figure 2.13 shows the user interface that is used to configure an experiment. The main functionality and user interface of the configuration utility was developed by Jack Brookes, an engineering undergraduate student, under the author's supervision.

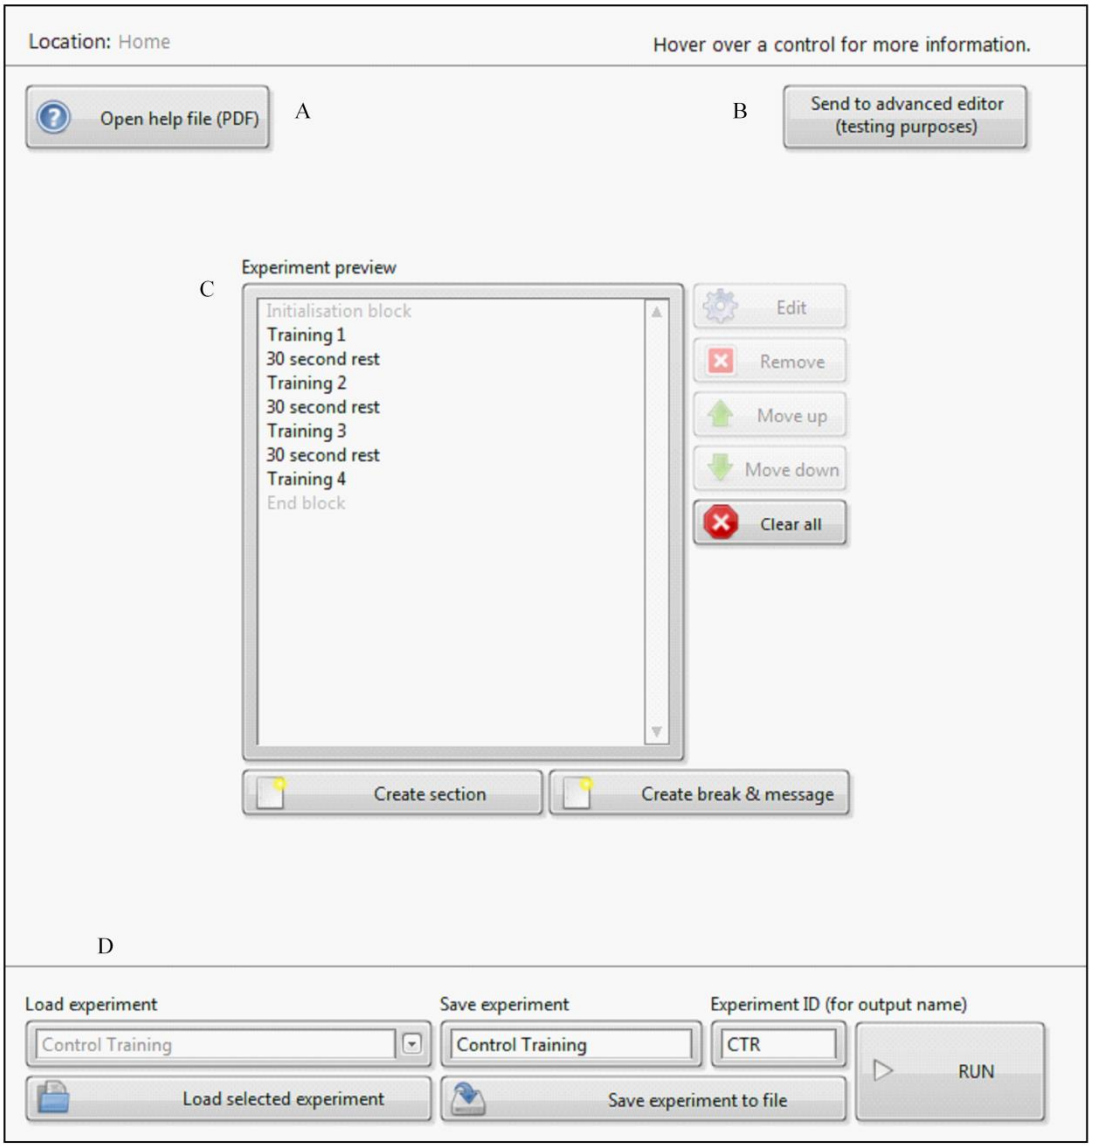

Figure 2.13. User interface of the configuration utility front panel, showing A) an option to load a help file explaining how to use the program, B) an option to enter the advanced editor (allows editing of the commands line by line), C) Preview of all experiment sessions and controls used to create, edit, remove, move and clear experimental sections, D) options for loading, saving and naming experiment configuration files.

Figure 2.14 shows a screenshot of the user interface that is used to create a section of an experiment (this is loaded upon pressing the 'Create section' button in the main front panel).

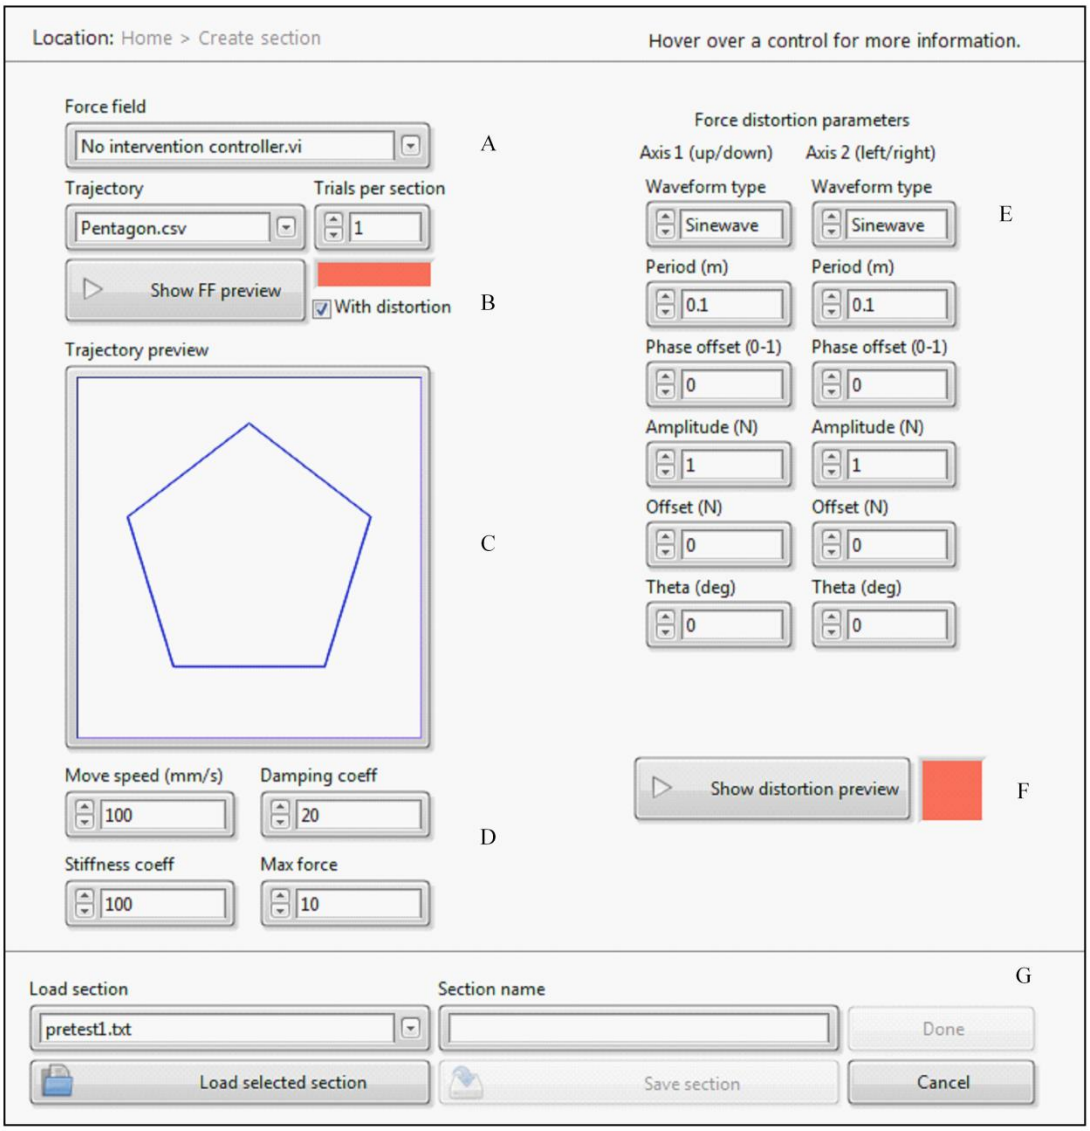

Figure 2.14. User interface of the section creator. A) options to select pre-defined error adjustment force algorithm sub-routine, and trajectory (path), B) option to show a preview of the noise signal as a function of the workspace (as shown in Figure 2.15) with an option of including or excluding the workspace distortion force field, C) a preview of the trajectory in the visual display, D) options to set the target movement speed as well as the stiffness and damping coefficients and maximum force applied by the haptic device (a safety feature), E) options for setting workspace distortion force field parameters (refer to Table 2.1), F) an option to show a preview of the workspace distortion force field with a selectable colour map, G) Option to load previously saved Sections for editing, as well as saving any new Section.

The Configuration Utility contains an option to display a 2D visual representation of the forces present in any trial (see Figure 2.15). This allows the operator to visually inspect the haptic environment at the experiment configuration stage, before implementing them on the device.

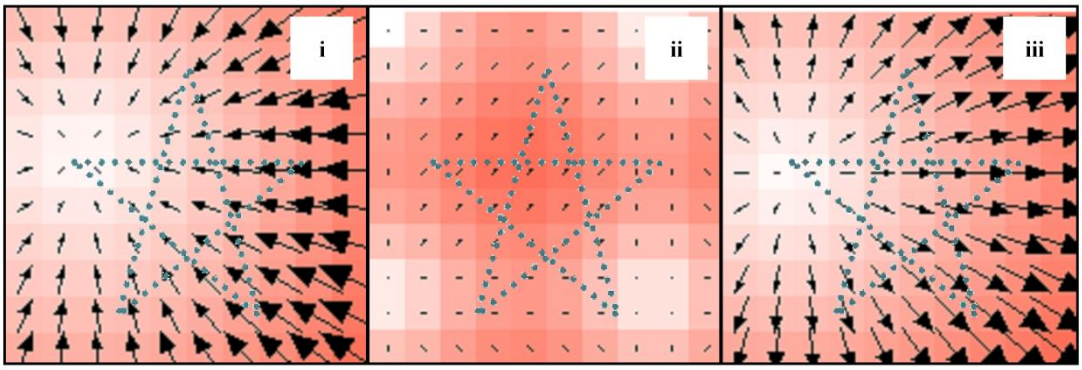

Figure 2.15. Example of a 2D visualisation of the haptic environment relating to a task with 'assistive' (i), 'workspace distortion' (ii), and 'disruptive' (iii) forces. Arrows indicate the direction and proportional magnitude of the force vector at discrete locations within the workspace. Relative magnitude is also represented using a colour map, where white = no force and dark red = high force. In this instance, forces are displayed for the first time point (the start location) of the first component of a pentagram path.

2.3.5.4 Haptic rendering

Achieving the minimum required force update rate of 500 Hz (refer to Specification in Section 2.3.2) with a minimum amount of jitter typically involves the use of a real-time controller which acts to implement pre-defined functions (i.e. the virtual model) that define the haptic environment. However, an important design objective was to develop a flexible method of controlling the behaviour of the device by enabling the implementation of bespoke force functions. Unfortunately, the HapticAPI functions are limited in terms of the flexibility of the haptic environments that can be achieved. Thus, there was a need for the 'virtual model' of the haptic environment to be specified outside of the device's built-in real-time computer.

In terms of haptic feedback, the experimental requirements of HAT were to generate forces as a function of workspace and device position (objectives 5 and 8-10). To achieve this functionality, it was necessary to update the force specified in hapticAPI2.dll from an external application using bespoke mathematical functions, whilst a minimum update rate of 500 Hz and minimum jitter (low enough to avoid any perceived 'jerky' behaviour)

were needed for high haptic fidelity. One way of reliably achieving this would be to use a real-time operating system to overcome the disadvantages of non-deterministic operating systems. However, this would require additional equipment and the use of new programming environments.

For the reasons outlined above, haptic feedback is generated directly from the HAT system. This uses the LabVIEW-HapticMASTER interface to read position and velocity data and write a single force vector per iteration of the haptic loop. A virtual model which defines the force commands sent to the haptic device is generated within the haptic loop (see Figure 2.11). The parameters of each component movement are loaded in parallel to the main haptic loop. Within the haptic loop, the virtual model defines 1) the dynamic properties of the device (such as inertial and global damping settings), 2) force fields which act to generate distortive forces as a function of position within the workspace (e.g. a sinewave which produces force in the Y axis as a function of position in the X axis), and 3) 'haptic noise' functions, forces which act to augment or reduce execution error (for instance, assisting or repelling forces can be generated as a function of positional deviation relative to the current component trajectory). Implementing the virtual model within the HAT enables operators to implement haptic environments of practically any level of complexity. However, due to the non-determinism of the Windows operating system, there is a potential risk that the required update rate is not achieved. The toolkit architecture has been specifically designed to address this issue by appropriately managing resources within the LabVIEW software and allocating sufficient computational power to the haptic rendering module. The performance of the haptic loop was tested and is reported in Section 2.3.6.

2.3.5.5 Visual rendering

Visual information was defined based on the requirements of the experiments discussed in Chapter 5. This included options to display a) instructional text and images, b) a target cursor and/or the current position of the device, c) the current component path, and d) a dotted line joining the two cursors to represent the current error vector.

A screenshot of the visual display is shown in Figure 2.16. It is rendered in OpenGL, which uses hardware acceleration to make use of the graphical processing unit (GPU) thus freeing up the PC central processing unit (CPU) for task logic, haptic feedback computations and device interaction in the haptic loop.

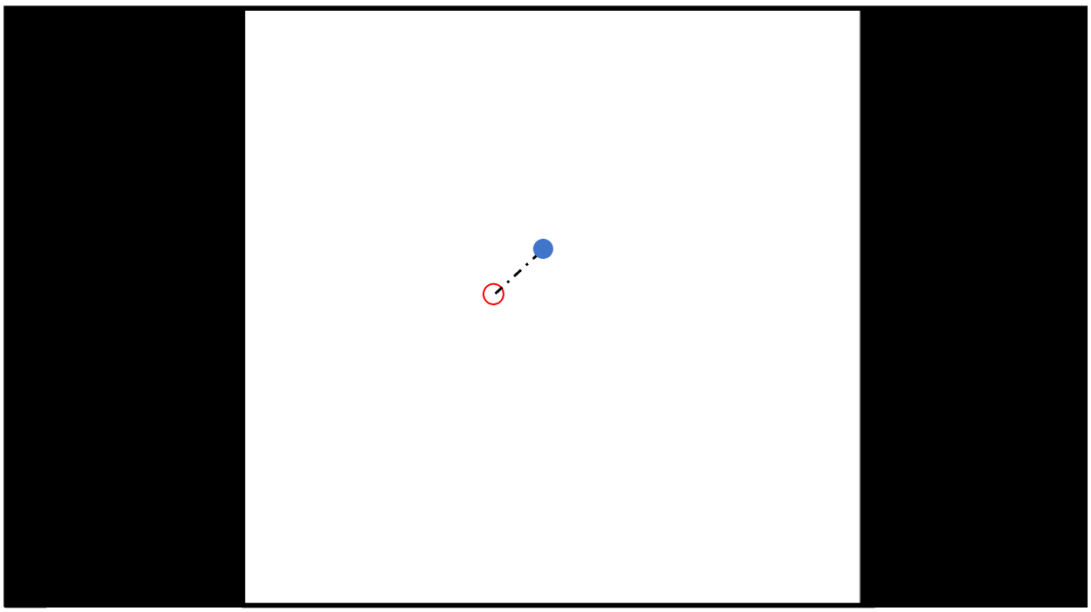

Figure 2.16. Screenshot of the visual display and example stimuli. The filled blue circle represents the position of the device end-effector within the workspace whilst the hollow red circle represents the target position. A configurable option allows the display of the current component. An optional dotted line is used in this case to visually show the magnitude of the error (a dashed line between the actual and target positions).

2.3.5.6 Data processing

To produce objective measures of performance within HAT, five kinematic metrics are calculated within the data processing module upon completion of each component. Four metrics are automatically generated to measure spatial and frequency indices of motor behaviour. These metrics have previously been used and validated successfully within the KAT system (Culmer et al., 2009) as objective measures of motor performance, making them suitable for this application.

Movement time

Skilled movements are commonly associated with an ability to move quickly and accurately (Schaverien, 2010). Movement time (T_M) is the total time taken to complete a movement. It is computed as the difference between the start and end times of a component.

Path length

Total distance travelled, or path length (L_P) is a measure of spatial accuracy, a simple measure of how well a demand trajectory is replicated. This can be defined using the pseudocode:

Input Data: Movement(Time, X, Y)
Output Data: Path_Length
For each sample in Movement - 1
Length(sample) = Abs_Error(Movement(Sample,X,Y), Movement(Sample+1,X,Y)
Path_Length = sum(Length)

Tracking error and standard deviation of the tracking error

Another measure of accuracy is tracking error (E_T). This is defined as the mean distance between the device and target cursors and provides information about how well a moving target is tracked along a spatiotemporal domain. The standard deviation of tracking error (SDE_T) provides information about how E_T varies over time (i.e. of how constant the error was):

Input Data: Reference(Time, X, Y), Movement(Time, X, Y)
Output Data: RMS Tracking Error, STD Tracking Error
For each sample in Movement
Tracking Error(sample) = Abs_Error(Reference(Sample, X, Y), Movement(Sample, X, Y))
RMS Tracking Error = RMS(Tracking Error(sample))
STD Tracking Error = STD(Tracking Error(sample))

Path error

Path error (E_P) is the mean nearest distance between the device cursor and the trajectory path:

Input data = Reference(Time,Y,Z), Movement(Time,Y,Z)
Output data = Path_Error
For each sample in Movement
Movement Error(sample) = Search for minimum distance from Movement(sample) to Reference
Cumulative_Movement_Error =Sum of Movement Error
Path_Error= Cumulative Movement Error/Number of Movement Samples

Normalised jerk

Motor learning generally leads to the generation of increasingly smoother movements. A measurement of 'smoothness' which has previously been used in the literature is Normalised jerk (J_N). This is the derivative of acceleration in time, and is defined by

equation (2.2), where T is movement time (T_M), L is path length (L_P), and j(t) is the triple derivative of position with respect to time.

$$J_N = \sqrt{\frac{T^5}{2L^2}\int_0^T j\,(t)^2 dt} \tag{2.2}$$

This metric is normalised with respect to time and distance which allows the trajectories of different durations and lengths to be compared. The measure is consequently unitless.

Hybrid measures of performance can be obtained by combining two or more of the output metrics. For instance, it is often more useful to consider a combination of speed and accuracy for the measurement of performance, as skilled movements tend to be faster and more accurate. This is often termed the speed/accuracy trade-off and can be obtained here by combining the outputs T_M and E_P and/or E_T. Nevertheless, it is important to note that the assessment of human performance is non-trivial and can be highly specific to task and/or environmental conditions (Fitts & Posner, 1979). HAT is an open and scalable architecture which allows for future integration of additional performance measures to meet the specific needs of other tasks.

2.3.5.7 Post-processing utility

Kinematic data files are processed using custom software to produce data in a usable format (see Figure 2.10). The post-processing utility allows the selection of what output metrics are required, and at what level or levels within the hierarchical structure of each session these should be saved (i.e. at a Trial, Section or Session level). Processed data are then saved in an open delimited text format which can be opened in a spreadsheet for further analysis.

2.3.6 System validation

The performance of the HAT system for generating haptic feedback was validated. The aim was to objectively assess the fidelity of the haptic feedback rendered by HAT through the HapticMASTER. This is particularly important considering the non-deterministic nature of the system, as discussed in Section 2.3.5.4. An experiment was carried out to 1) assess whether the servo loop updated at the minimum required rate of 500 Hz (see Specification in Section 2.3.2) to generate smooth and realistic forces under representative experimental conditions, and 2) to assess the fidelity of the system's

response to a step position input over time. This was achieved by comparing the desired and actual step response of a mass-spring-damper simulation.

2.3.6.1 Methods

A mass (m = 3 kg) - spring (k = 100 N/m) - damper (c = 2.5, 5, 10 and 20 Ns/m) algorithm, as described in equation (2.3) where X is the 2D position vector, was implemented using the HAT. These values were used as they are representative of those used in the experiments of Chapter 5.

$$F = m\ddot{X} + c\dot{X} + kX \tag{2.3}$$

The response of the full system was then tested using a step input of 0.1 m in the vertical axis. Additionally, to investigate temporal performance, the iteration period of the haptic loop was recorded over ten thousand iterations during the implementation of the condition c = 10 Ns/m.

2.3.6.2 Results

A histogram illustrating the temporal performance characteristics of the haptic loop is shown in Figure 2.17.

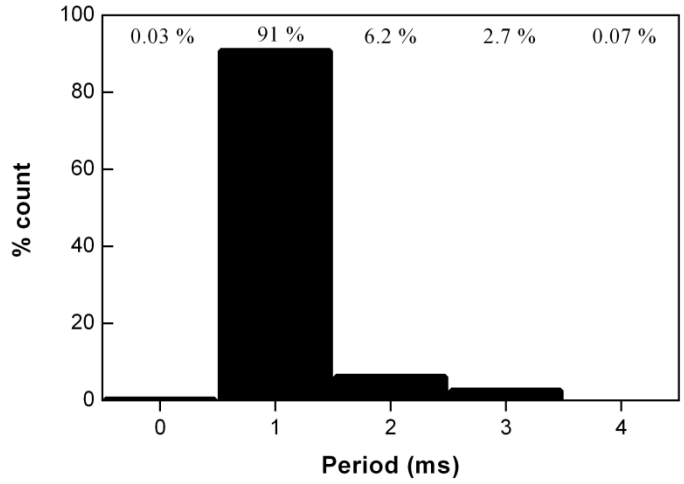

Figure 2.17. Haptic loop period (*M* = 1.11 ms, *SD* = 0.4 ms) and frequency of occurrence measured whilst rendering a mass-spring-damper system with settings m = 3 kg, k = 100 N/m, c = 10 Ns/m, over ten thousand iterations of the HAT servo loop. The largest recorded value was 4.5 ms.

Whilst the servo rate was not constantly below the recommended 500 Hz (2 ms period, as outlined in the Specification), the majority of iterations were. The servo rate was consistently above 200 Hz for all 10000 iterations. Considering the small percentage of deviations from the recommended loop rate (jitter; defined as the standard deviation of the mean loop rate), it is unlikely that the user would notice the effects of such small discrepancies in the haptic update rate. This suggests that the non-determinism of the operating system does not have a significant detrimental effect on the software's ability to communicate with the haptic device at an appropriate speed.

The performance of the mass-spring-damper simulation was next analysed. Figure 2.18 shows the response of the simulation to a step input of 0.1 m. The system was able to produce an output that was characteristic of the desired response (i.e. with a relatively small mean error of 2.1 %), indicating that the HAT is capable of haptically rendering a mass-spring-damper system with relatively high fidelity.

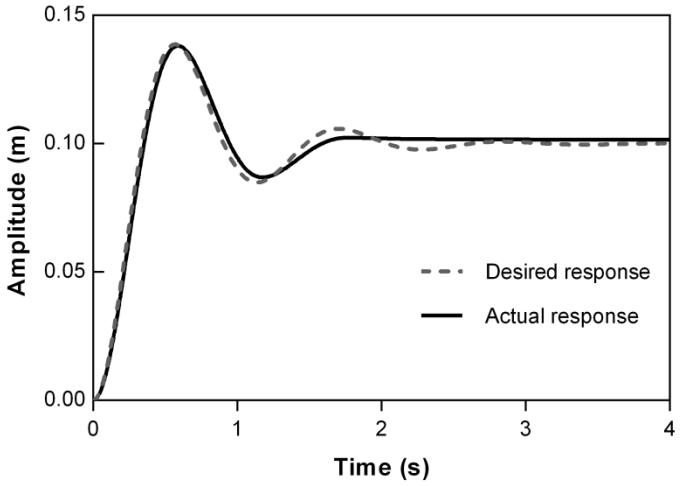

Figure 2.18. Desired versus actual responses of a mass-spring-damper simulation with settings m = 3 kg, k = 100 N/m, c = 10 Ns/m. Mean RMS error = 2.1 mm.

The observed errors in the simulation were due to two effects: 1) a phase lag and 2) greater attenuation of the actual response relative to the desired response (this is most pronounced from the third oscillation onwards). One contributing factor to these effects may be the interaction between the limitations of the virtual model (implemented from HAT) and servo rate (refer to Figure 2.17): consider the force response on the first iteration of the servo loop. This is the largest force present in the simulation, and it is next updated on the second iteration of the servo loop. However, the time between the start of the first

iteration and the start of the second will have errors due to the discretised response of the virtual model. The same effect will be in place between the second and third iterations, and all subsequent iterations. Thus, relative to an analogue (continuous) mass-spring-damper response, the discretised force response of the virtual model will result in an additive displacement error effect across a simulation. One way of addressing this issue could be to incorporate Proportional, Integral, Derivative (PID) control to account for the time response of the system. However, this is likely to require additional computational power, thus potentially decreasing the 'smoothness' of the response. This was not investigated further as it was deemed to fall outside of the scope of this work.

Another likely contributing factor to the observed errors is the inherent signal noise present at the force sensor of the HapticMASTER (a common disadvantage of admittance control as discussed in Section 2.3.3), whereby small errors in the measured force input to the system would interact with the system and therefore affect the response fidelity. Consider the system's natural frequency, ω_0, the square root of ratio of the stiffness, k, and mass, m:

$$\omega_0 = \sqrt{\frac{k}{m}} \tag{2.4}$$

The phase lag may be due to an effectively smaller ω_0, i.e. due to an error in the implementation of k and m. Next, consider the damping ratio of the system, ζ, the ratio of the damping factor, c, and twice the square root of the product of k and m:

$$\zeta = \frac{c}{2\sqrt{km}} \tag{2.5}$$

This relationship defines the oscillatory response of the system, which may contribute to the attenuation effects observed in Figure 2.18. Thus, it is likely that both observed error effects could be due to an error in one or more of the variables k, c and m. This raises the question what is the effect of changing system parameters on response fidelity? Figure 2.19 shows the error of the mass-spring-damper simulation for various damping values, showing that RMS error is inversely proportional to the damping value.

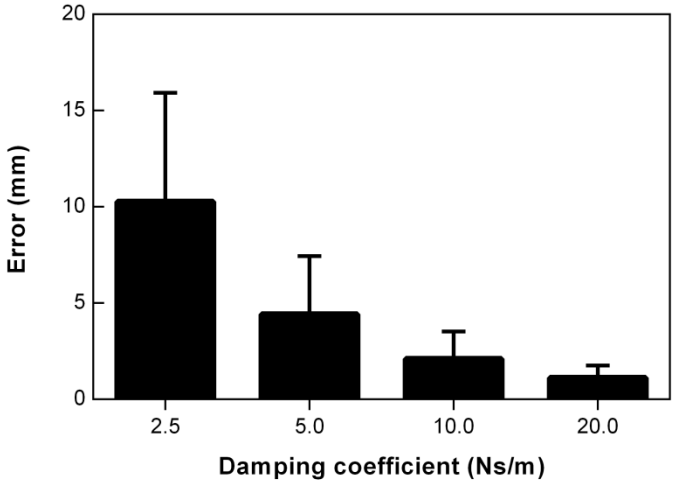

Figure 2.19. Mean RMS error between the desired and actual responses obtained for different damping (c) settings (m = 3 kg and k = 100 N/m for all cases). Error bars represent one standard deviation of the mean.

The overall performance of the system is a function of both the HAT software and the haptic device. Whilst the simulation of a dynamic system does not perfectly match the ideal response, it is capable of providing smooth and realistic forces. For the purposes of this PhD (as outlined in the objectives), the haptic environment will be required to intervene with the subject's movements by pushing or pulling their hand in a specific direction. Thus, high dynamic fidelity is not necessary, indicating that the performance of the HAT system is adequate for this application. However, these limitations should be considered in relation to the specific task requirements for which HAT may be used in the future.

2.3.7 Experimental setup

With the aim of promoting consistency across experiments, a standard experimental setup was specified for HAT (see Figure 2.20). Subjects were stood in front of the HapticMASTER within a marked safe zone located 5 cm outside of the operating workspace of the device. Participants held the end-effector (i.e. manipulator) with their right hand. They were not given any explicit instructions on how to hold the device. Visual stimuli were presented on a computer monitor located approximately 1.4 m above the floor (this was in line with the maximum displacement of the device in the vertical axis to avoid obstruction of the display by the device), and directly behind the HapticMASTER, approximately 2 m from the participant.

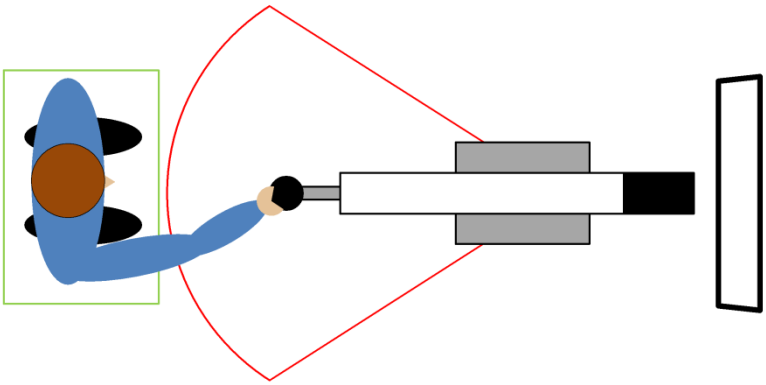

Figure 2.20. Plan view of the standard HAT experimental setup, showing the relative positions of the participant, HapticMASTER and monitor (see Figure 2.16 for a screenshot of representative visual stimuli). The marked green and red areas represent the assigned standing area of the subject and the operating region of the HapticMASTER, respectively. For safety, the red region should not be entered during operation of the HapticMASTER other than by the arm of the subject. This is to avoid any potentially dangerous collisions with the device.

2.4 Discussion

This chapter describes two novel tools that can be used to create bespoke visuohaptic environments, CSI and HAT. This enables the acquisition of experimental data via robust and semi-automated methodologies. These tools are the benchmark for all experimental protocols used throughout this thesis.

CSI enables the simulation of virtual compliant objects and the acquisition of kinematic data. CSI is designed for the implementation of simple virtual elastic models to represent the surface response force of compliant objects when they are interacted with using a handheld tool. One limitation of this is approach is that it does not consider time-dependent characteristics of objects which are common in human tissue (i.e. viscoelasticity). However, the focus of the experiments in this PhD was on human action and perception, and thus (considering the higher modelling complexity associated with alternative options) an approximate simulation of characteristic tissue properties has been deemed sufficient. CSI is compatible with any device in the PHANTOM range of haptic devices, which makes it possible to use it for a range of hand-device interfaces (manipulators), workspace sizes and output forces. However, the impedance control paradigm which these devices (and therefore CSI) use limits them to small forces and relatively small workspaces. CSI requires the development of LabVIEW functions to

control the behaviour of the device within an experimental protocol (the ones used for this PhD are described in Chapters 3 and 4).

The HAT system is a configurable, high-level platform that allows inexperienced users to configure visuohaptic environments to implement force-based intervention strategies during the execution of motor behaviour. Force feedback functions of virtually any degree of complexity are implemented through a high-level user interface. The HapticMASTER is best suited to arm movements and large force output (i.e. gross movements), allowing the device to manipulate the position of a resting arm, or counteract forces applied to the manipulator. HAT automatically outputs objective performance metrics that can be used to assess human performance, and a post-processing utility allows for the tabulation of experimental data into a useful format which can be used for later analyses. The HAT platform is open access and scalable, allowing for future adjustments of the software's functionality.

CHAPTER 3

COMPLIANCE DISCRIMINATION WITH A HANDHELD TOOL BEFORE AND AFTER TRAINING

ABSTRACT The effect of training on tool-based compliance discrimination ability and the kinematic variables associated with discrimination performance are investigated in this chapter. Within the context of an action-perception framework described in Chapter 1 it is well established that performance feedback is needed to make strategic adjustments to improve human perceptual abilities. However, to the author's knowledge there has been no work that specifically examines what movement strategies are most effective for assessing compliance with a handheld tool (a critical skill in areas such as laparoscopic surgery and dentistry), and indeed the effects of training on such strategies. The abilities of naïve participants to detect compliance differences with and without knowledge of results (KR), as well as the abilities of participants who had undergone repetitive training over several days, were investigated. Kinematic analyses were carried out to objectively measure the probing action. Untrained participants had poor detection thresholds, and no short-term effects of KR on performance were found. Intensive training substantially improved group performance. Probing action (in particular, slower movement execution) was associated with better detection thresholds, but training did not lead to systematic changes in probing behaviour.

3.1 Introduction

Humans perform skilful interactions with objects using a combination of visual and haptic information (Mon-Williams, Wann, Jenkinson, & Rushton, 1997). There are few domains where successful and skilful interactions of this type are as crucial as surgery and dentistry. During open surgery, vision and haptics can be used together to perform palpation (examination of tissue) and retrieve information that can confirm the location and extent of physiological anomalies such as lumps or tumours. This type of manual probing is highly sophisticated since a number of finger movements can be performed in sequence to provide multiple estimates of tissue compliance across a wide area. Unfortunately, there are many surgical situations where such manual probing is not possible. For example, laparoscopic (keyhole) surgery is increasingly being used because of the significant patient benefits (e.g. reduced recovery time and trauma; (Cuschieri, 1995) but one of the main difficulties with such techniques is the loss of high quality visual and haptic information (Culmer et al., 2012). Thus, palpation during laparoscopic surgery is fairly limited due to the weak haptic signals available from laparoscopic graspers and the necessity to carry out sequential probing. In contrast, dentists routinely use a handheld probe to explore the properties of the periodontium (the supporting tissues of the tooth) and tooth structure. Amongst other methods, it is common for dentists to use instruments such as a blunt dental probe to confirm the health of a tooth, as extensive tooth decay (dental caries) alters the compliance of the tooth's structure (Selwitz, Ismail, & Pitts, 2007). Likewise, the periodontal probe allows a dentist to determine the health of the periodontium, in part through detection of changes in the structure's compliance. Whilst the use of probing techniques to obtain information about dental health is common, it is also reported as being difficult to teach students and can take a long time to master (Drucker et al., 2012). To the author's knowledge, only one study (Teodorescu et al., 2013) has shown that repetitive training can lead to better compliance discrimination performance. This effect was attributed to haptic perceptual learning through short- and long-term gains. However, there was no investigation into whether these improvements were due to systematic changes in probing strategy, to an increased perception of compliance cues, or to a combination of the two. The present study was motivated by the fact that there is little information on the effects of training on perceptual thresholds for compliance discrimination with a handheld probe, nor the actions underpinning this process.

To determine how humans perceive properties of the world (such as compliance) it is useful to consider the role of active behaviour in perception. From an evolutionary perspective, the primary purpose of perception is to support the performance of skilled actions (such as avoiding predators or picking up food) rather than perception for its own sake (J. J. Gibson, 1986). A consideration of the nature of compliance perception using a handheld probe appears to indicate that this task fits perfectly within this 'active perception' framework. Tissue compliance can only be determined through interactions between the probe and the tissue, and the natural way to elicit this interaction is by performing a probing action that will itself shape the quality and quantity of information supplied to the central nervous system (Kaim & Drewing, 2011; Lederman & Klatzky, 1987). It seems that compliance perception requires the monitoring of force and displacement information during interactions with an object (Choi, Walker, Tan, Crittenden, & Reifenberger, 2005; Kaim & Drewing, 2011; Tan, Durlach, Beauregard, & Srinivasan, 1995; Tan, Durlach, Shao, & Wei, 1993). This is achieved by using a combination of cues from haptic receptors, as well as vision (Kuschel, Di Luca, Buss, & Klatzky, 2010; Sigrist et al., 2012; Srinivasan, Beauregard, & Brock, 1996; Tiest & Kappers, 2009; W. C. Wu, Basdogan, & Srinivasan, 1999). Thus, it is important to consider various visual-motor factors in the overall perception of compliance. To this end, empirical investigations have studied human compliance discrimination abilities under various conditions, including: active and passive interactions, with rigid and non-rigid surface objects, during direct (e.g. finger) and indirect (tool) interactions, and with constrained and unconstrained movements. Some of these are summarised next.

(Tiest & Kappers, 2008) showed that during finger interactions with deformable surface objects, active discrimination trials resulted in a just noticeable difference (JND) of 12%. Under passive conditions (when the compliant sample was brought into contact with the stationary finger) the compliance JND increased to 14%. This effect has been attributed to the contribution of kinaesthetic information towards the perception of compliance (thought to be small compared to that of cutaneous information), which agrees with findings from other studies (Friedman, Hester, Green, & LaMotte, 2008; Kuschel et al., 2010; Lederman & Klatzky, 2004). Further, the inclusion of congruent visual (Kuschel et al., 2010; W. C. Wu et al., 1999), and even auditory (LaMotte, 2000) information seems to improve compliance discrimination ability, indicating that the CNS integrates information from multiple sources to generate the final perception of compliance (Ernst

& Banks, 2002; Sigrist et al., 2012). (Tiest & Kappers, 2008) also showed that cutaneous information about geometrical changes of an object's surface obtained at the finger pad is an important cue for perceiving compliance: worse performance (the mean JND nearly doubled to 23%) was observed during active finger indentations of compliant objects with rigid surfaces (i.e. when no surface deformation occurred). This is in line with findings by (Srinivasan & LaMotte, 1995), who showed that tactile information alone is insufficient to encode the compliance of rigid surface objects. This was attributed to the fact that for a given net force, skin deformation is dependent on the compliance of deformable surface objects, but not of rigid surface objects. Thus, for the latter case, the CNS appears to rely mostly on kinaesthetic information to encode compliance. This might also be the case for tool interactions, in which there is no direct contact with the object and so cutaneous information is only available about the object's response in the form of Normal (for one finger interactions) or shear (for two fingers in a precision grip) forces (Friedman et al., 2008). One difference between tool and direct finger rigid surface interactions might be cues about the rate of change of force produced upon impact with the sample at higher movement speeds, i.e. tapping versus pressing (LaMotte, 2000). This raises two important questions: (i) what is the relationship between probing strategy and performance; and (ii) are there optimum movement patterns which can be used to maximise the information available to the CNS?

Previous work has found that haptic performance in compliance perception depends on the executed exploratory movements. (Kaim & Drewing, 2009) showed that maximum finger force and velocity are strategically adjusted to the expected compliance. In a later study (Kaim & Drewing, 2011), the same authors found that the application of higher forces resulted in lower JNDs (better performance). These results seem to indicate that optimal probing strategies might exist, dependant on the compliance of objects. Unfortunately, most of this work has so far focused on direct finger interactions, which cannot be generalised to tool-based interactions due to the active and perceptual differences between these tasks, as highlighted earlier. Whether this effect is also true for tool-based interactions, and indeed whether 'optimal' exploration settings are reached over short or long term periods is also yet to be established.

Finally, another factor in terms of exploration strategy is that of consistency across the probing actions used to assess the compliance of specimens. (Tan, Pang, & Durlach, 1992) measured JNDs for fixed and variable displacements during active pinching of a

rigid surface using an electromechanical device. For each sample, when the required displacement was reached, the force was switched off. Variable displacements between trials resulted in significantly reduced performances (JND = 5-15% with fixed displacements versus 22% for variable displacement). This suggests that constant pairwise indentation depth may lead to better discrimination performance, but further investigation is required to establish whether this is the case during 'normal' operating conditions (i.e. when displacement is controlled by the subject instead of by the experimental apparatus). Further, Srinivasan & LaMotte (1995) found that performance deteriorated for passive interactions with rigid surface objects when indentation speed and terminal force (the response force of an object at maximum indentation depth) were inconsistent between the two probing actions. This was not the case during active interactions. However, performance was measured using samples of discrete compliances, and thus small effects of this phenomenon may not have been captured in the active case. It was hypothesised that the task would be greatly simplified by generating identical movements between the compliant samples being compared, because perceptual information would then be directly comparable and differences could be detected without the need to have stored knowledge of absolute compliance values (e.g. a 'lookup table').

CSI (described in Chapter 2) was used to facilitate the measurement of the JND of compliance differences in adults across three groups, along with the kinematics of their movements. Group 1 consisted of untrained participants with no explicit feedback (knowledge of results, KR) provided after each trial (this mimics normal conditions where trainees must become attuned to intrinsic visual-motor feedback in the absence of performance feedback). Group 2 were given KR after each trial (a feature possible within VR training systems). Group 3 received intensive training over a week (whilst providing KR) to determine whether compliance discrimination can be improved through training with performance feedback. The relationship between probing strategy and performance was also examined, and whether there were systematic changes in probing behaviour during training.

3.2 Methods

3.2.1 Materials

The compliance simulation interface (CSI, described in Chapter 2) was used to develop custom software to enable the generation of forces through the haptic device. Visual stimuli was generated in LabVIEW and displayed on a computer monitor. The standard CSI experimental setup (detailed in Chapter 2) was used. Participants were positioned directly in front of the monitor and the haptic device was placed to the right of the monitor (as illustrated in Figure 3.1) where it could be reached comfortably with their right hand. All participants were right handed. Participants were instructed to rest their right arm and wrist on supports during the whole of the experiment. This was done to limit the probing strategy to only wrist and finger movements, thus promoting consistent behaviour across participants.

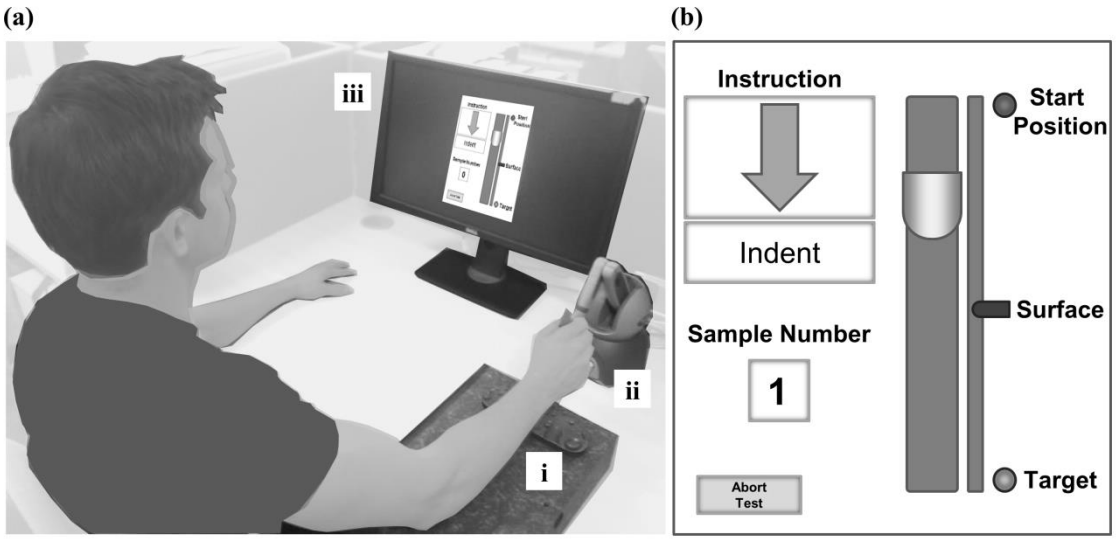

Figure 3.1. (a) Diagram of the experimental setup, showing the physical arrangement of i) the arm and wrist supports, ii) the haptic feedback device, and iii) the computer monitor showing a graphical display. (b) A close-up of the graphical display, where the instruction panel informs the participant to 'Indent' the sample, 'Await Instruction', or to 'Move to [the] start position'; 'Sample number' denotes which sample (1 or 2) is being displayed. The position of the stylus tip (the shaded square with a rounded bottom edge) is updated at a rate of 30 Hz and is shown relative to the virtual object's surface, the start position and the target.

3.2.2 Experimental design

Three groups were tested: novice – no knowledge of results (KR) (N), novice – with KR (NKR), and trained - with KR (TKR). First, between-subjects tests were carried out to assess the effects of KR on the performance of novices. A repeated measures design was then used to test the learning effect of the NKR group as they progressed through the twelve training sessions.

3.2.3 Task configuration

Participants actively probed two virtual samples, one after the other, and then judged which was the least compliant (stiffest). An adaptive staircase algorithm, the PEST (Parameter Estimation by Sequential Testing) was employed to generate the stimuli on a trial by trial basis. A correct or incorrect judgement in the previous trial was used to adjust the stimulus properties in the next, so that the staircase size and direction adapted on the basis of past trial performance (Lieberman & Pentland, 1982). This procedure converges on an individual's threshold more quickly and with greater precision than standard staircase methods (Leek, 2001). The threshold obtained at the end of each set of 100 trials was identified as the JND value. This number of trials was selected based on pilot trials to allow successful convergence on a JND value, whilst providing the same number of training trials to all participants. A constant baseline stiffness of 0.075 N/mm was used for one sample in each trial (stiffness is the inverse of compliance). The other sample had a stiffness value ranging from 0 to 0.1 N/mm above the baseline. This stiffness baseline and range were selected to fall within the force output capabilities of the haptic device, and are similar to the conditions encountered during palpation of soft human tissue such as liver (Mueller & Sandrin, 2010). The order of appearance (first or second) of the 'baseline' sample was randomised. The first trial of each session contained a 'baseline + offset' sample located in the middle of the offset range at 0.125 N/mm.

3.2.4 Participants

Nineteen participants (10 male, 9 female, aged: 21 to 28 years, $M = 23.4$, $SD = 2.56$) were recruited and randomly allocated to one of the three groups: Group N performed one compliance discrimination session. Group NKR also performed one session but with knowledge of results (KR): a green 'tick' was displayed for a correct response and a red 'cross' for an incorrect response. The procedure for Group TKR was identical to that of NKR (KR was provided), except that each participant completed twelve sessions over

four days (three sessions per day - morning, noon and late afternoon). All participants were Psychology or Engineering undergraduate or PhD students. The research was approved and conducted under the guidelines established by the School of Psychology, The University of Leeds Research Ethics Committee.

3.2.5 Procedure

On each trial, participants used the haptic device to indent each sample of simulated tissue before identifying which of the two was perceived as the stiffest. For each sample, participants were required to move the stylus downwards from a "start" position past the "surface" until the "target" indentation was reached, as indicated in Figure 3.2 (b). An auditory tone indicated when participants could start their movements. A higher pitched tone then indicated when they had reached the target. Forces were generated as a function of indentation depth to simulate a stiffness value once the surface had been passed, and a cursor on the monitor indicated the position of the stylus relative to the three reference positions (*start, surface* and *target*). Participants were required to vocally identify the stiffest sample ('one' for the first sample or 'two' for the second sample). The experimenter electronically recorded their response and then the next trial was presented. No explicit instructions were given to participants regarding what movement characteristics were expected (e.g. speed or acceleration were unconstrained). However all participants were required to hold the stylus of the robotic device in a precision grip (i.e. held between the tips of the thumb, index and middle finger) and they were informed of the essential vertical probing movements that would be required for the task. Movement was unconstrained in all dimensions. For each sample, the stimulus and response given for each pair of samples were recorded. All participants received a practice run of 25 trials prior to the start of the experiment to ensure that they understood, were familiar and were comfortable with the task. Each session consisted of 100 trials and lasted approximately 15 minutes.

3.2.6 Kinematic analysis

The movement kinematics of the probing actions over time were explored in order to determine whether there was a relationship between probing strategy and performance. Figure 3.2 illustrates the virtual probe and its relative position to the surface during the probing action.

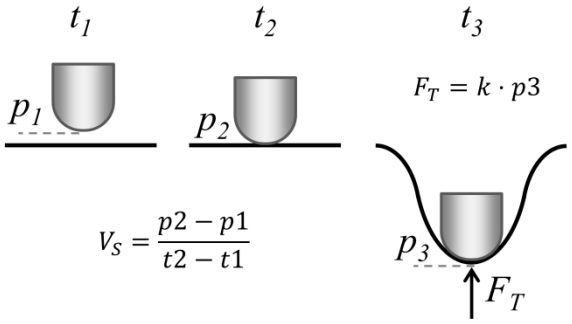

Figure 3.2. Illustration of the virtual probe at different stages of indentation, showing i) the probe at approximately 2 mm prior to contacting the virtual surface, ($p_1 \approx -2$ mm), ii) the instant at which the probe comes into contact with the surface ($p_2 \approx 0$ mm), and iii) the deepest indentation (p_3) during the probing action, when F_T is calculated. The variables t_X and p_X denote the time and relative position of the probe to the surface at each probing stage, respectively.

For each sample, time and position in the vertical axis were recorded at a frequency of 100 Hz. To objectively assess probing strategy, two kinematic metrics were determined for each sample: strike velocity (V_S) and terminal force (F_T). These variables are described in equations (3.1) and (3.2).

$$V_S = \frac{p_2 - p_1}{t_2 - t_1} \tag{3.1}$$

$$F_T = k.p_3 \tag{3.2}$$

V_S is calculated as the average velocity of the probe as it travels from position p_1 to position p_2. It is not influenced by interactions with the sample, making it useful for assessing movement behaviour independent of sample compliance (which varied from sample to sample). F_T is calculated as the product of the sample's stiffness coefficient, k, and the position at maximum indentation, p_3. Due to the finite resolution of the kinematic data (samples were acquired at approximately 10 millisecond intervals), the exact values of p_2 and p_3 were taken as those points that were closest to and above their respective specified locations of 0 mm and the overall maximum indentation depth, respectively. The variable p_1 was then selected as the closest position to p_2 that met the condition $p_1 - p_2 \geq 2$ mm. The time points (t_X) were specified as the time elapsed from the start of the trial until each corresponding p_X location was reached. Time measurements were accurate to within 0.5 milliseconds. With the aim to assess the effects of consistency of probing kinematics on performance, the absolute difference in V_S (ΔV_S) and F_T (ΔF_T) between

the first and second indentations were calculated. Note that due to the relationship between k and p_3, described in equation (3.2), the magnitude of JND (effectively a measure of the compliance difference between the two samples) would affect ΔF_T if p_3 was, to any extent, a controlled variable by the subject (i.e. ΔF_T would be proportional to JND). Thus, ΔF_T should only be assessed for trials where identical pairwise sample compliances are presented.

3.3 Results

The compliance JND values and mean kinematic metrics for all trials obtained for groups novice – no KR (N), novice - KR (NKR) and Session 1 of trained - KR (TKR) are shown in Table 3.1. The JND obtained for one participant in group N did not fall below 100% and so they were excluded due to failure to perform the task.

Table 3.1: Compliance JND and mean V_S, ΔV_S, and F_T across all trials for groups N, NKR and session 1 of TKR.

Participant	Group	JND (%)	V_S (mm/s)	ΔV_S (mm/s)	F_T (N)
1	N	18.9	58.2	21.0	1.03
2	N	45.7	152.5	51.4	1.23
3	N	25.0	30.1	13.2	0.92
4	N	33.0	84.1	26.6	1.26
5	N	27.4	45.3	13.3	1.11
6	N	43.7	90.5	33.5	1.54
7	NKR	39.8	52.0	18.6	1.05
8	NKR	21.7	36.7	12.0	1.00
9	NKR	39.8	132.7	34.8	1.45
10	NKR	63.3	33.8	14.1	1.22
11	NKR	15.7	47.8	22.9	0.83
12	NKR	28.7	60.4	20.9	0.96
13	TKR (1)	23.8	46.5	17.8	0.91
14	TKR (1)	23.8	45.1	11.8	0.91
15	TKR (1)	36.2	36.4	20.3	0.95
16	TKR (1)	28.7	62.2	24.6	1.05
17	TKR (1)	22.8	175.9	48.0	1.76
18	TKR (1)	23.8	46.8	13.3	1.07

An independent-samples t-test revealed no significant difference between the JNDs of group N and NKR ($t(10) = 1.145$, $p = .279$, $r = .34$), suggesting that KR had no significant effect on performance. There were also no reliable differences between NKR and Session 1 of TKR ($t(10) = -.31$, $p = .76$, $r = .098$).

To assess longer term training on JND performance, changes across each session for TKR were examined. The group's mean JND values and mean kinematic metrics of all trials obtained for each training session are shown in Table 3.2. Figure 3.3 shows the mean JND obtained at each training session for TKR. A repeated-measures ANOVA showed a significant main effect of Session on JND ($F(11,55) = 3.15$, $p = .002$, $\eta_P^2 = .39$), with

participants gradually improving compliance sensitivity over time from 26.6% to a best value of 12.1% in session 7 (though values then drifted to 16%, possibly reflecting exploration of different probing strategies).

Table 3.2. Kinematic metrics obtained for all participants in the TKR group across the four days of training.

Session no.	Day	Mean JND (%)	SEM JND (%)	V_S (mm/s)	ΔV_S (mm/s)	F_T (N)
1		26.5	2.1	68.8	22.6	1.10
2	1	25.7	3.2	76.6	25.7	1.16
3		20.7	4.1	82.0	26.7	1.09
4		21.7	3.6	63.1	22.0	1.05
5	2	17.3	1.8	48.9	16.5	0.98
6		15.3	1.8	42.3	14.6	0.94
7		12.1	2.3	41.2	15.7	0.92
8	3	15.6	2.4	43.5	17.8	0.93
9		16.9	2.3	50.1	19.6	0.99
10		14.9	3.8	51.4	20.3	0.99
11	4	17.1	4.3	51.4	20.0	0.96
12		15.9	2.3	51.4	19.6	0.96

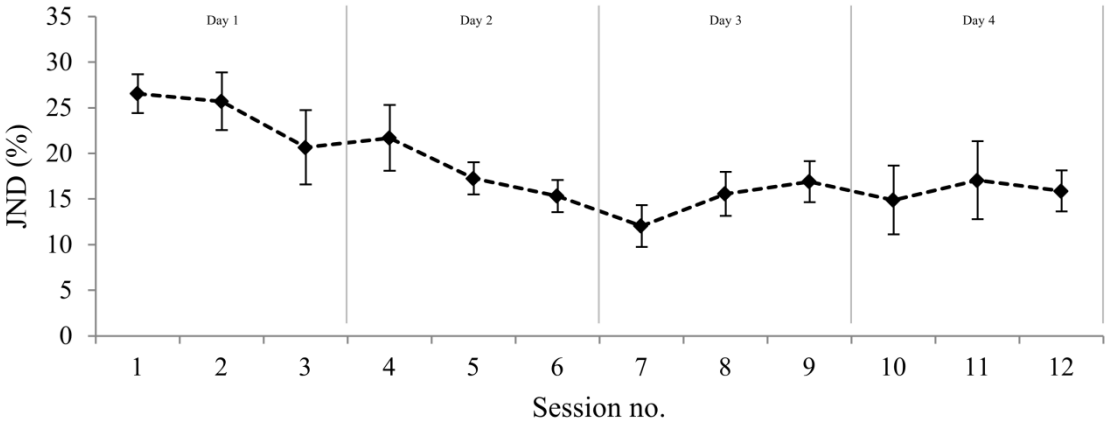

Figure 3.3. Mean JND obtained at each training session for TKR, showing a gradual improvement in performance over the four days of training. Error bars represent ± one standard error of the mean.

To determine whether changes in probing strategy could explain the improvements in compliance sensitivity, the kinematic metrics of TKR were examined. Plots of V_S, ΔV_S and F_T versus JND for each participant are shown in Figure 3.4, respectively. There was a significant positive correlation between V_S and JND whereby improved sensitivity to compliance differences was associated with slower probing velocity (r = .476, 95% BCa CI [.254, .674], p<.001). A similar relationship was also found between ΔV_S and JND (r = .526, 95% BCa CI [.309, .711], p<.001), and between F_T and JND, (r = .484, 95% BCa CI [.256, .695], p<.001). Two participants (P17 and P18) employed particularly fast probing actions, but if they were excluded from the correlations the pattern of results did not change. Also observed was a strong relationship between V_S and F_T, (r = .909, 95% BCa CI [.807, .953]), and V_S and ΔV_S, (r = .865, 95% BCa CI [.744, .931]) (ps < .001), consistent with slower movements leading to lower terminal forces and greater consistency of probing. To determine whether probing behaviour altered across sessions, a repeated-measures ANOVA was performed on each kinematic measure. These analyses revealed no significant main effect of Session on V_S (F(11,55) = 0.97, p = .49, η_P^2 = .16), ΔV_S (F(11,55) = 1.38, p = .21, η_P^2 = .22) or F_T (F(11,55) = 1.59, p = .13, η_P^2 = .24). Repeated-measures ANOVAs were carried out on the consistency measures, which failed to demonstrate a significant main effect of Session on ΔV_S (F(11,55) = 0.76, p = .68, η_P^2 = .13), or ΔF_T (F(11,55) = .78, p = .66, η_P^2 = .13). These results suggest that the group's JND improvements over the training period were not due to changes in probing strategy captured by V_S and F_T, nor by increased consistency in probing actions (as measured by ΔV_S and ΔF_T). It seems therefore that JND improvements may have instead been achieved by tuning into the appropriate perceptual information (Tresilian, 2012).

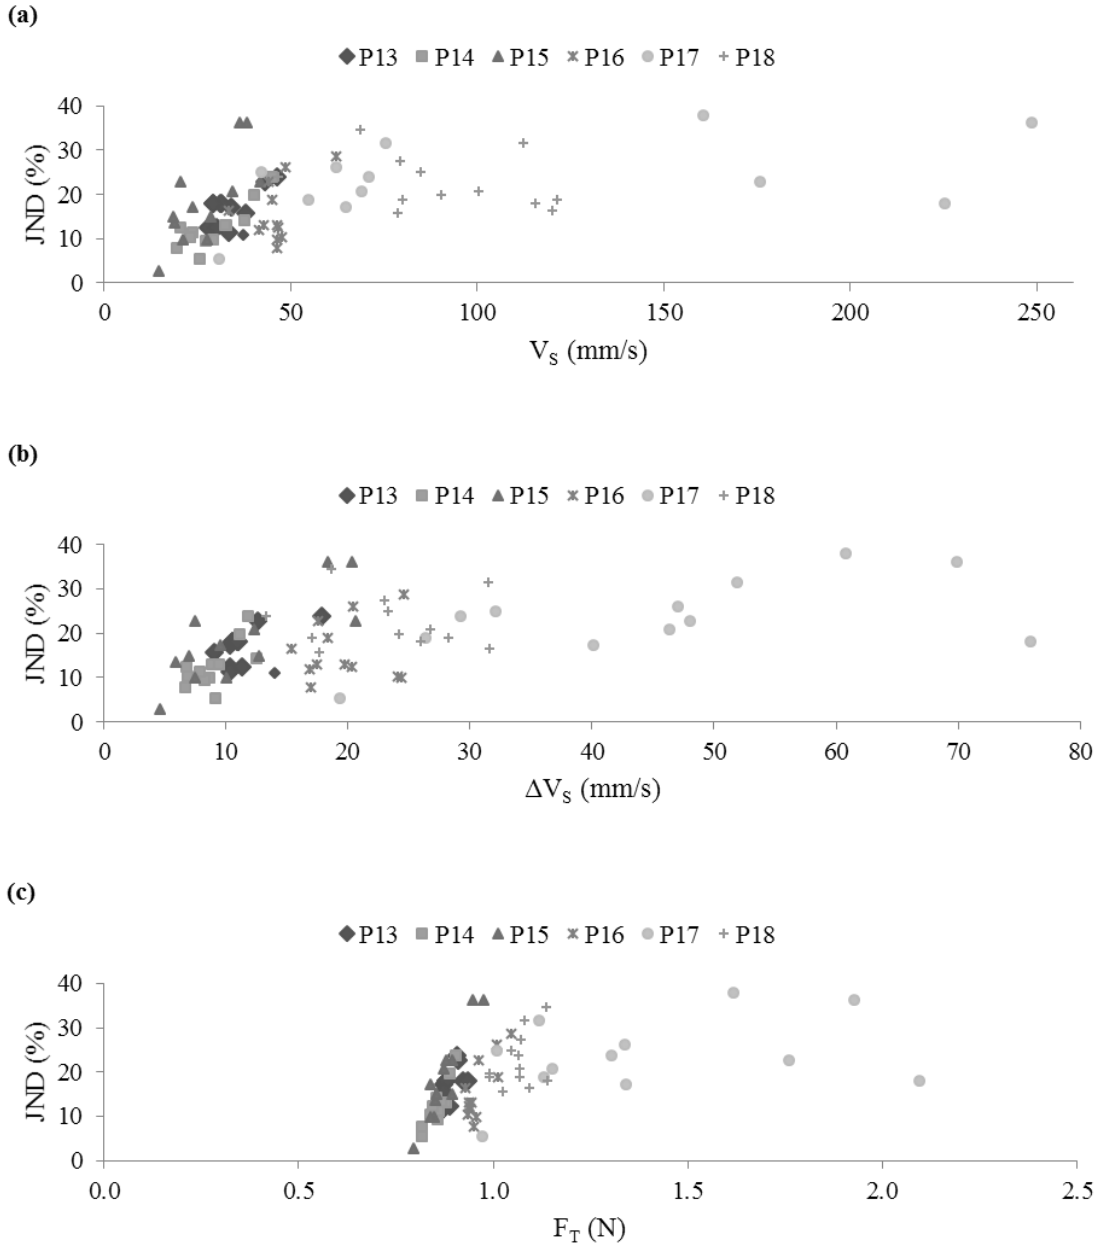

Figure 3.4. Plots of the mean (a) V_S, (b) ΔV_S and (c) F_T against JND, for all participants and all training sessions of TKR.

Finally, the extent to which individual differences explained the relationship between probing behaviours and JND performance was assessed. Figure 3.5 shows the correlation coefficient between V_S, ΔV_S and F_T and JND. Three individuals match the group level analysis – a clear relationship between all three kinematic metrics and JND performance. However there were also three individuals who displayed a weaker relationship between probing kinematics and JND and these individuals were those that exhibited the highest mean V_S, ΔV_S and F_T across all sessions.

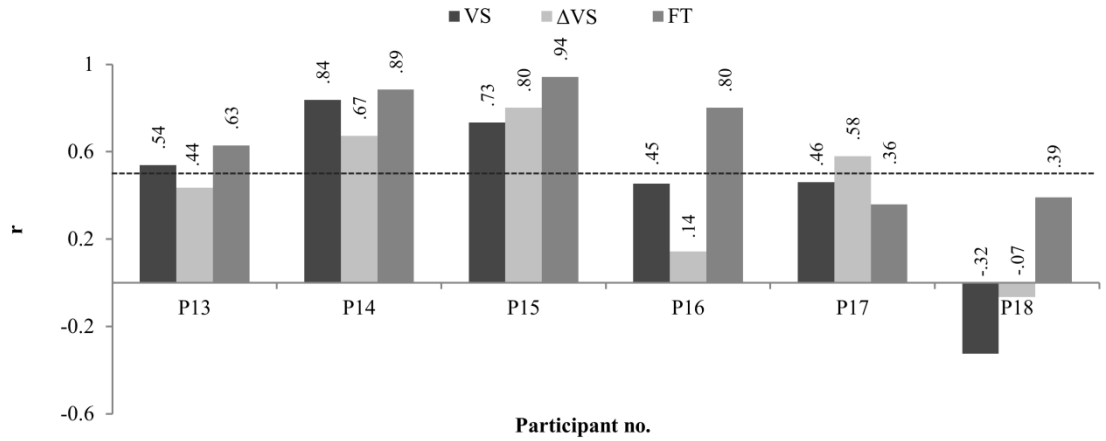

Figure 3.5. Correlation coefficients for each participant obtained for the mean V_S, ΔV_S and F_T against JND obtained during all training sessions. The horizontal lines indicate the significance threshold (r = .497), above which the correlation becomes significant at the p = .05 level.

With a group size of six it is difficult to draw firm conclusions about these differences, but it seems that there is a non-linear relationship between probing kinematics and performance, and that different probing strategies may have been adopted across individuals. For instance, the magnitude of V_S had a strong positive relationship with JND for P14, whereas for P18, there was a negative relationship suggesting that they may have tuned into information related to F_T (or another unmeasured variable).

3.4 Discussion

This chapter considered compliance detection within the theoretical framework of 'active-perception' where perception's primary goal is to support action and where humans obtain perceptual information through active interactions with the world. This framework would suggest that perceptual thresholds are a function of movement strategy so that a complete understanding of compliance detection requires an investigation into probing kinematics as well as perceptual sensitivities. To this end, the effects of long-term training on discrimination performance were investigated. Probing strategy was objectively measured to assess the relationship between exploratory strategy and discrimination abilities. The results suggest that novice participants are generally quite poor at discriminating compliance differences using a handheld tool. High values of JND can be generally associated with the execution of fast movement speeds, large interaction forces, and poor consistency in pairwise trial kinematics. The variables V_S, ΔV_S and F_T were all strongly correlated, which might indicate that fast movement speeds lead to

larger indentations (and therefore higher interaction forces), and a lesser ability to implement consistent movements. These findings would seem to disagree with those by (Koçak, Palmerius, Forsell, Ynnerman, & Cooper, 2011): active tapping (assumingly, tapping is a significantly faster movement than those employed in this study) of two virtual boxes by untrained participants resulted in a mean JND of 12.9%. However, in contrast to this study, participants were able to indent each sample as many times as they desired. Arguably this could enable them to capture more information about the objects under scrutiny, resulting in better performances than the ones observed here.

A comparison of the results from groups N and NKR suggests that KR does not have an immediate effect on compliance sensitivity, but the results from group TKR suggest that intensive training over several sessions with KR does. These results agree with a similar study by (Teodorescu et al., 2013). Both sets of findings imply that with training, humans can improve their ability to detect compliance differences using a handheld tool. The critical question is how do such improvements in perceptual abilities take place? One possibility is that when humans learn to perform a skilled action they become attuned to the available information (Wilson, Collins, & Bingham, 2005). Another is that they alter the behaviour that supplies such information (Wilkie et al., 2008), or some combination of the two. Findings from this study indicate that behaviour is indeed an important component in compliance detection (in this case via adopting an effective probing strategy). However, training did not lead to changes in probing strategy. This suggests that participants were able to improve their perceptual abilities (perhaps through sensitisation of the haptic senses; (Tresilian, 2012). Further, through visual inspection of Figure 3.5 it seems that weaker relationships between probing kinematics and JND are associated with individuals who employed a higher V_S, ΔV_S and F_T (with reference to Figure 3.4). This effect hints towards the presence of a non-linear relationship between probing strategy and JND: past a certain point in the magnitude of these variables, their influence on performance becomes weaker. It also seems that some of these variables may have different effects across individuals. This effect could explain some of the inconsistencies between individuals of groups N and NKR. However, further work is required to fully investigate the effects of individual differences in probing behaviours.

Investigating the relationship between action and perception in this way seems a logical way of deciphering the role of different sensory cues and their importance in generating the human perception of compliance, and indeed, for perceiving other phenomena. The

relationship between movement characteristics and compliance discrimination performance under natural operating conditions, i.e. manipulating a handheld probe within an unconstrained 3D space was assessed. Training over an extended period could be an effective method of improving performance in tasks requiring the discrimination of small compliance differences using a handheld tool. Nonetheless, beyond achieving increased perceptive abilities to changes in compliance, further research is needed to identify how to ensure that probing is performed within a temporal 'sweet spot' that avoids the disadvantages of acting either too fast or too slow. It must be the case that very slow movements impair sensitivity because slow movements are generally more difficult to execute smoothly, which may have implications for sensory performance (Nagasaki, 1989). In contrast, very fast movements must limit the perceptual information available from the probing action. Also understanding what training programmes could be effectively implemented to guide the subject towards these movement strategies will be critical for future applications. Finally, it is important to validate these training programs by assessing the transfer of augmented skills to more realistic and meaningful tasks, such as dental and laparoscopic procedures.

CHAPTER 4

THE ROLE OF ACTION-PERCEPTION IN A VIRTUAL TUMOUR DETECTION TASK

ABSTRACT The previous chapter investigated the role of active and perceptive mechanisms in the detection of compliance whilst using a handheld tool. One issue that remains is whether there are particular methods for training compliance perception that are more effective than others. Continuing on the theme of compliance discrimination abilities in surgery and dentistry discussed in Chapter 1, this chapter investigates the role of haptic versus kinematic training (Training) on a virtual palpation task (Test). There were three independent groups: haptic training (HT), kinematic training (KT), and a control (CT). Training was in the form of six sessions (of 15 minutes duration) over three days (two sessions per day). Test was a virtual palpation task that was representative of some components of a real-world medical procedure. During Training, HT and KT used the same haptic device as that used at Test, whereas CT completed a number of tracing tasks using a tablet PC and a stylus. Results showed a significant improvement from the first Test (pre-Training) to the second Test (post-Training) in palpation performance. However, the types of training employed did not influence the rate of improvement. It seems, therefore that all groups improved their palpation performance irrespective of training, and there was no clear advantage of a particular training regime. There were, however, initial differences between groups during Test 1 (pre-training, with Group CT being worse performers) that may have interacted with the training regimes.

4.1 Introduction

Compliance discrimination with a tool is an important skill in a multitude of surgical and other settings. However, this skill can take years to refine and master (Drucker et al., 2012), which raises the critical question: how can the rate of acquisition of this skill be increased?

The individual roles of action and perception in our detection of compliance are not fully understood. The previous chapter showed that the strategy employed to interact with an object affects our overall perception of the object's compliance, and that humans are able to improve their detection of small compliance differences over time. There are also individual differences in probing strategy, and slower speeds generally result in the best performance. However, training did not seem to lead to systematic changes in probing strategy. These findings indicate that humans are able to improve their ability to discriminate compliance differences by increasing sensitivity to compliance cues, but it is unlikely that individuals will spontaneously modify their active probing behaviours, even though such changes may have been highly beneficial in terms of increasing the available perceptual information. This chapter examines whether novel training regimes can be used to guide participants to use more effective particular probing patterns with the aim of improving the quality and quantity of information available to the CNS.

From a theoretical standpoint, a number of factors should be considered in the design of an environment that acts to guide optimal probing movements. As discussed in Chapter 1, motor learning is often described as an iterative process that leads to the construction and refinement of internal models. Within this framework, prior knowledge of the state of the environment and motor architecture (e.g. the inertia of a tool or end-effector) is used to objectively plan an action to achieve a desired goal (e.g. to indent an object at a specific speed; Wolpert, 1997). Thus, optimal motor learning conditions are highly task-dependent and so it is important that the training environment accurately represents the components of the task to be learned (e.g. the inertial and other physical properties of the tool used for training should be the same as that used in the real environment).

Two components that could be involved in our ability to assess compliance are haptic sensitivity and kinematic performance. It seems that increased haptic sensitivity can be achieved through repetitive training whereby objects with similar compliances are

repetitively assessed (see Chapter 3). The second component, kinematic performance, requires guidance towards a probing strategy or strategies that may result in improving the quality and quantity of available information relating to an object's compliance. It would seem, therefore, that increased compliance discrimination performance could potentially be achieved through training regimes that act to 1) increase perceptual sensitivity and 2) improve kinematic performance.

This study investigates the effects of two training interventions on the ability to discriminate compliance differences. The haptic training (HT) group were prescribed the just noticeable difference (JND) task described in the previous chapter. The kinematic training (KT) group were required to follow a target moving within a speed range that was designed to optimise compliance information, and so would be expected to promote good compliance discrimination performance. To independently assess the effects of kinematic training (i.e. in isolation to haptic sensitisation), force feedback was not provided to the KT group. To objectively assess the effects of these two training conditions on a meaningful, real-world procedure, a virtual palpation task was used to measure performance before (Test 1) and after (Test 2) Training. This task was designed to simulate the palpation of tissue with a handheld probe, a procedure that is used in a number of medical procedures. A control training (CT) group was also tested to provide a comparison when examining the relative performances of the HT and KT groups from Test 1 to Test 2.

4.2 Methods

This section describes the general methodology employed for the experiment, as well as a detailed description of each of the experimental conditions (Test, haptic training, kinematic training and control training).

4.2.1 Materials

A Phantom Omni haptic device was positioned in the same way as that discussed in General Methods (Chapter 2), but this time placed inside a 43x31x52 cm box, with an opening (16x31cm) for the participant to place their hand inside to hold the device gimbal. There was no physical contact between the box and the device or participants during the task. This was done to mimic more closely the environment in which a surgeon operates, whereby the ends of the laparoscopic probes are only visible on a 2D screen. This

eliminated the possibility that perceptual cues were obtained from the movement of the device during interactions with virtual objects. For the control task, the Clinical Kinematic Assessment Tool (CKAT) was used (Culmer et al., 2009). It consists of a tablet PC and a hand-held stylus on which 2D visuomotor tasks can be programmed. In this case, four tracing tasks which have been previously used to assess motoric abilities were employed (Flatters et al., 2014). An HP EliteBook 2760p tablet PC with an 11.42x8.35 inch screen, 1280x800 resolution, 125 dpi, 64-bit colour, 60 HZ refresh rate was used.

4.2.2 Experimental design

Three groups were tested: Haptic training (HT), kinematic training (KT) and control training (CT), all of which received separate training interventions. All groups completed the same virtual palpation task before (Test 1) and after (Test 2) Training. A within-subjects design was used to assess the effects of Training over time for the HT and KT conditions, and a mixed design was used to measure any relative changes between groups from Test 1 to Test 2.

4.2.3 Participants

Thirty unpaid participants aged 20 – 29 years (M = 23, SD = 2.26) were recruited for the study and randomly allocated to one of the three groups. All participants reported that they had normal or corrected vision, and that they were right-handed. No participants had received any surgical training and had never used the haptic device. The study was approved by The University of Leeds ethics committee and was performed in accordance with British Psychological Society (BPS) ethical guidelines. All participants provided their informed consent prior to the commencement of the study.

4.2.4 General procedure

Participants first read an instruction sheet detailing the requirements of the study and signed a consent form. They were asked to complete eight sessions over five consecutive days (see Figure 4.1). The first session (Test 1) was a virtual tumour palpation task, which required the exploration of a region of simulated tissue to detect differences in compliance. Test lasted approximately twenty minutes. Over the next three days, participants were required to complete six Training sessions: two per day with a minimum of 2.5 hours between any two consecutive sessions. The HT group completed a JND compliance discrimination task which was identical to that described in Chapter 3 for the

JND task with knowledge of results. The KT group completed a pursuit task which required actively aligning a cursor representing the position of the gimbal end (the 'device cursor') with a target cursor. HT and KT used the haptic device during Training by holding the device gimbal in a standard pencil grip. The CT group required participants to complete various tracing tasks using a tablet and stylus. The duration of each Training task was approximately 15 minutes. The second and final Test session (Test 2, which was completed by all groups) was identical to Test 1, excluding practice trials, and was carried out on day 5. Each task is described in more detail below.

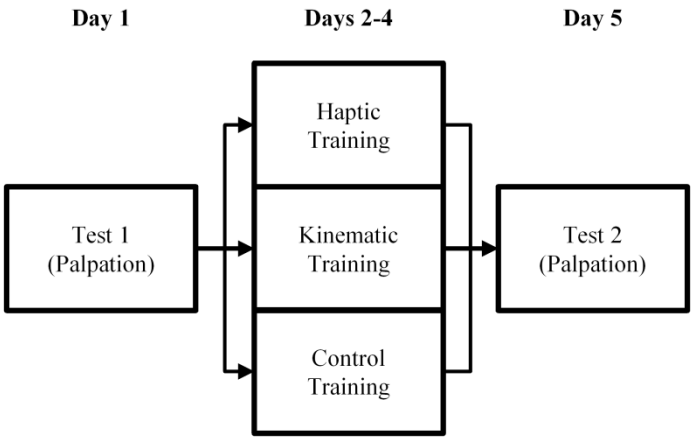

Figure 4.1. The study was conducted over five days. On day 1, a virtual palpation task was completed by all groups. This was Test 1. On days 2-4, participants were given Training (two sessions per day) as per their allocated group (haptic training - HT, kinematic training – KT, or control training - CT). Test 2 (which was identical to Test 1) was completed on day 5 by all groups.

4.2.5 Virtual palpation task (Test)

The virtual palpation task was completed by all participants before (Test 1) and after (Test 2) Training.

4.2.5.1 Task configuration

A visuohaptic environment was generated using the compliance simulation interface (CSI, described in Chapter 2). Within this system, a sample with embedded tumours was modelled using a Gaussian approximation, as described in Equation (4.1), where F_G is the output force in Newtons, A is the peak force value relative to the baseline, B, σ is the function width variable and x_r is the radial distance from the end-effector to the inclusion centre. A description of the modelling process can be found in Appendix 1.

$$F_G = Ae^{-\frac{x_r^2}{2\sigma^2}} + B \quad (4.1)$$

Tissue deformation was designed to approximate the physical visual response of human tissue. For this, a simple elastic model was used. Previously acquired experimental data were used to determine the modelling coefficients of different inclusions (see Appendix 1). Three different samples were implemented to simulate tumours with a 12 mm diameter at depths of 6 mm ('S1'), 9.25 mm ('S2') and 12.5 mm ('S3'). These combinations of parameters were selected after a preliminary study during a Master's project which the author was involved with, which tested the ability of untrained participants to find tumours of different sizes and at different depths (see Appendix 1). Most participants were able to detect the inclusion at a depth of 6 mm, whilst most were unable to detect it at 12.5 mm. These were selected with the aim of avoiding floor and ceiling effects.

4.2.5.2 Procedure

Participants were seated and asked to rest their right arm and wrist on supports (as per the general experimental setup described in Chapter 2). They were asked to hold the haptic device stylus as a pen and use it to probe the virtual sample displayed on the monitor using a vertical motion to find a hidden tumour. A screenshot of the virtual environment is shown in Figure 4.2.

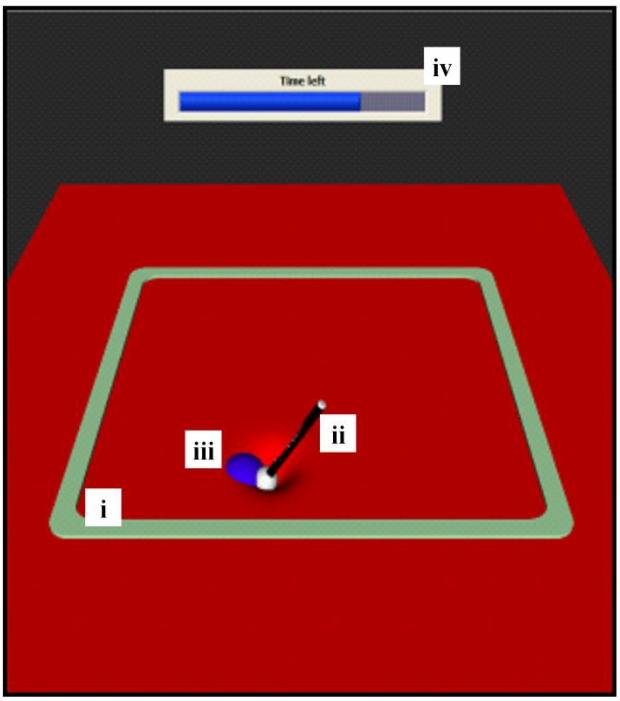

Figure 4.2. Screenshot of the visual stimuli given in the virtual environment, showing i) the probing area (the inside of the green square which measured 100 x 100 mm; ii) the virtual probe, consisting of a rod and spherical end; iii) the blue embedded tumour (shown visually only during the first practice trial); and iv) an indicator of the time remaining for the current trial. The deformation of the tissue to an indentation is also shown.

Participants were required to navigate the 3D space by controlling the stylus and vertically indenting the virtual tissue. Once indented, a maximum horizontal movement of 10 mm was allowed to more closely replicate the probing strategy adopted during internal tissue palpation (i.e. during an operation), whereby horizontal movements are small in order to avoid damage to tissue (Culmer et al., 2012). Any horizontal movements which exceeded the permitted threshold resulted in a loud auditory signal (a 'beep'), and a freezing of the visual scene until the probe had moved above the surface of the sample. The basic probing technique was demonstrated by the experimenter before practice trials began.

Upon detecting a tumour, participants were asked to place the probe as closely as possible to the centre of the tumour and press the dark grey button on the stylus of the device. Once a selection was made, a confirmation message was given and the next trial was presented. With the aim of minimising outliers due to random positive selections when a tumour had not successfully been detected, participants were given the option of selecting that no tumour was present. To make a 'no tumour present' selection, participants were

required to press the light grey button. After making a selection, participants were asked to vocally confirm their selection with the experimenter, who then confirmed or cancelled the selection using a keyboard input. This allowed participants to return to the trial if an accidental selection was made (e.g. if a button had been mistakenly pressed). The selection was recorded and the next sample was automatically loaded. Each trial had a maximum duration of 2.5 minutes, after which the trial would automatically end and a 'timeout' would be logged. A time bar at the top of the screen indicated the time remaining during each trial.

Participants were required to complete four practice trials (two sets of two trials with a 5 second rest in between), followed by nine experimental trials (5 trials followed by 4 trials with a 30 second rest in between). The force response of the samples in the practice trials was augmented to make them easier to detect. Additionally, in the first practice trial, the tumour was visually displayed as a blue sphere to allow participants to familiarise themselves with the haptic feedback associated with the inclusion.

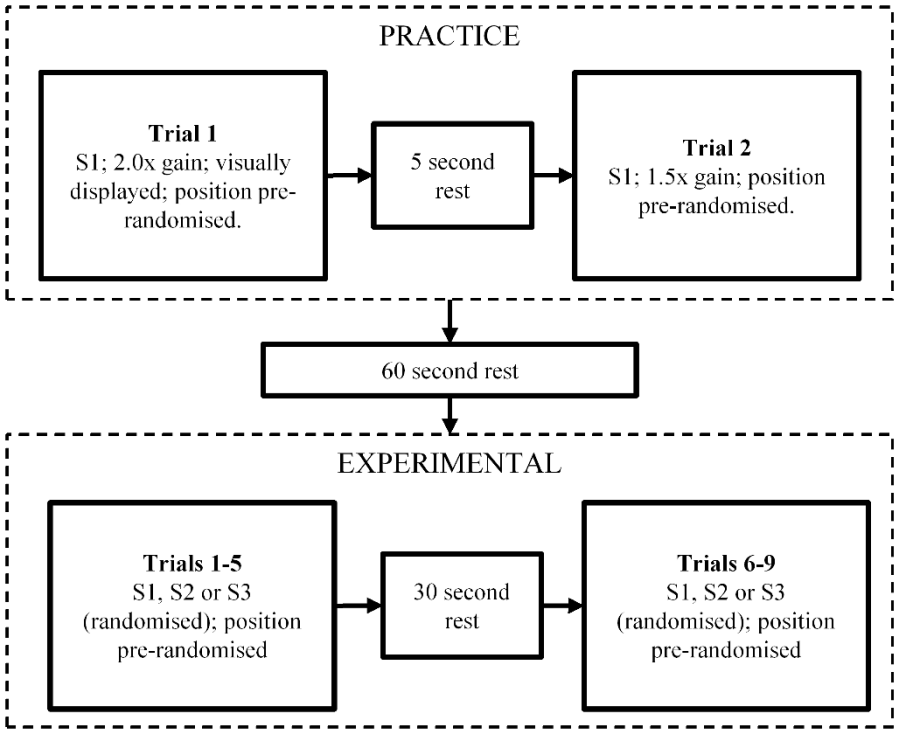

Figure 4.3. Diagrammatic overview of the Test sessions, which consisted of two practice trials (only for Test 1) and nine experimental ones. S1, S2 and S3 refer to the 12 mm tumour at 6, 9.25 and 12.5 mm from the sample's surface, respectively. The first practice trial was with S1, but with a 2.0x haptic gain to augment the difference between the baseline and peak forces. The tumour was visually displayed as a blue sphere. Followed by a 5 second rest, the second practice trial was again with S1 but with a 1.5x gain, and it was not visually displayed. After approximately 1 minute, the experimental trials were loaded. The order of tumours (S1, S2 or S3) was randomised across all trials so that each tumour would appear three times. The Cartesian positions of all tumours were pre-randomised (i.e. tumours were in the same location for each trial).

Practice was only given in Test 1. To avoid order effects, the order of appearance of each sample in the experimental trials was randomised. For both the practice and experimental trials, the position of the inclusions for the experimental trials was pre-randomised, whereby their locations within the 2D Cartesian space were pre-specified using a random number generator. The order was the same for both Test sessions.

4.2.5.3 Data capture

For each sample, the program recorded kinematic data in the form of x, y and z position with a time stamp (in ms, accurate to +/- 0.5 ms). These data were used to calculate the mean probing velocity (V_P) for each trial. The standard deviation of the probing velocity (SDV_P) at each trial was also calculated. Actual and selected positions of each tumour

were recorded when a positive selection was made. Also recorded was whether there was a timeout or 'no tumour present' (negative) selection.

With the aim of objectively assessing performance, outcome metrics were derived from the raw kinematic data. These were radial error (the distance from the location of the probe to the centre of the tumour), correct or incorrect selection (a correct selection was defined as one where the probe overlapped with the tumour i.e. when the centre of the end-effector was within 11 mm of the tumour centre, whilst an error above this value was defined as an incorrect selection), and movement time (the time taken to make a decision from the start of the trial when a positive selection was made). For each participant, the mean of each outcome metric at each sample (S1, S2 and S3) was calculated and used for the main analysis reported here. The metrics were: selection error (E_S, the mean radial error for all positive selections made) number of correct selections (S_C), selection time (T_S, the mean amount of time taken to make a positive selection), and a hybrid measure, selection error x selection time (ET_S). A speed-accuracy composite metric is commonly used to assess surgical motoric skill because it captures the two main task demands: namely performing quickly and accurately (Judkins, Oleynikov, & Stergiou, 2008).

Outliers were identified as those participants who were much more inaccurate than the rest of their group when a positive selection was made. First, the Z-score was calculated for each sample across all groups for the metric E_S at Test 1. The same was done for Test 2, but this time individually for each group. This was because any systematic differences between groups (after Training) could have resulted in the incorrect definition of an outlier in one group who may have performed better or worse than the other groups. Outliers were identified as those with a Z-score above 2.58 (i.e. a E_S value that was larger than 99% of the sample's distribution) and were classified as individuals who failed to perform the task at that specific stage in the experiment (possibly due to a momentary lapse in concentration, for instance). Using this method, there were three, two and six participants who were identified as outliers for S1, S2 and S3, respectively. Their data were consequently excluded from the analysis for that particular sample and Test session, but were included in the rest of the analyses where they did not exceed the critical Z-score value.

4.2.6 Haptic training (HT) condition

The compliance JND task was identical to the JND task described in Chapter 3. Each Training session consisted of 50 trials of the JND task with knowledge of results (an indication of whether a correct or incorrect selection was made). Outliers were again defined as those with a Z-score of 2.58 or above at each of the Training sessions. One participant was removed in this way.

4.2.7 Kinematic training (KT) condition

Haptic feedback was not provided in this task. Participants were required to follow a target cursor as accurately as possible for a number of trials at variable speeds by controlling the haptic device.

4.2.7.1 Task configuration

CSI was used to read the position of the device, and the visual stimuli were generated using LabVIEW. A screenshot of the visual display is shown in Figure 4.4.

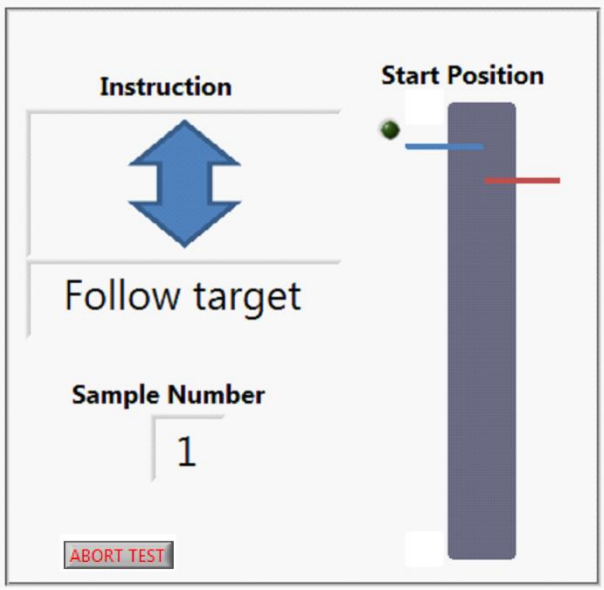

Figure 4.4. Screenshot of the visual stimuli window for the KT task, showing an instruction ('follow target' or 'move to start'), sample number ('1' or '2'), the target (red) and device (blue, representing the position of the tip of the device gimbal in the vertical axis) cursors, and an LED which served to highlight the start position upon completion of each sample. The length of the vertical slider was 20 mm.

The target cursor moved at various speeds between 10 and 50 mm/s along the vertical slider. This range was chosen based on findings from the previous chapter, which showed

that these speeds were representative of those that resulted in good performance. Target position was computed using a sinusoidal wave, as described by Equation (4.2), Where *P* is the position of the cursor along the vertical plane in mm and *f* is the frequency in Hz. To generate the required movement speed of 10-50 mm/s along the 20 mm slider, movement frequency, *f*, varied between 0.25 and 1.25 Hz. The length of the slider was the same as that used in the HT task.

$$P = 20\cos(f) \tag{4.2}$$

4.2.8 Procedure

To start the session, participants were required to move to the 'Start Position'. The target cursor then began to move and they were required to control the gimbal to align the device cursor with the target cursor as accurately as possible. For each sample there was one cycle (up, down, up movement). To maintain consistency with the HT task, two samples were presented per trial. The speed of each pair of samples was constant. For the first Training session there were 10 practice trials followed by 50 experimental trials. In all subsequent Training sessions, only the experimental trials were given.

4.2.9 Data capture

Radial error, the distance between the target and device cursors, was calculated at each time step and for each sample. The mean Tracking Error (E_T, the mean distance between the target and device cursors at each session) was used as a measure of performance. There were no outliers in the data (outliers were again defined as those individuals with a Z-score of 2.58 or above at each of the Training sessions for the metric E_T).

4.3 Results

The results can be split into the Test component and the Training component. First, performance at Tests 1 and 2 were examined, followed by an analysis of probing kinematics. Training performance over the six sessions was then examined to see if there were improvements. This was done to demonstrate whether participants were engaged with the tasks and actually learning during training.

4.3.1 Virtual palpation Test

To determine whether there was an effect of Training on Test performance, an analysis of performance at Tests 1 (pre-Training) and 2 (post-training) was carried out. Each sample (S1 = 12 mm tumour flush with the surface, S2 = 12 mm tumour 3 mm below the surface, S3 = 12 mm tumour 6 mm below the surface) was analysed independently to control for difficulty. The four performance metrics (S_C, E_S, T_S and ET_S) were used to objectively assess performance.

4.3.1.1 Palpation performance

Figure 4.5 shows each outcome metric and Surface, showing group performances at each Test session.

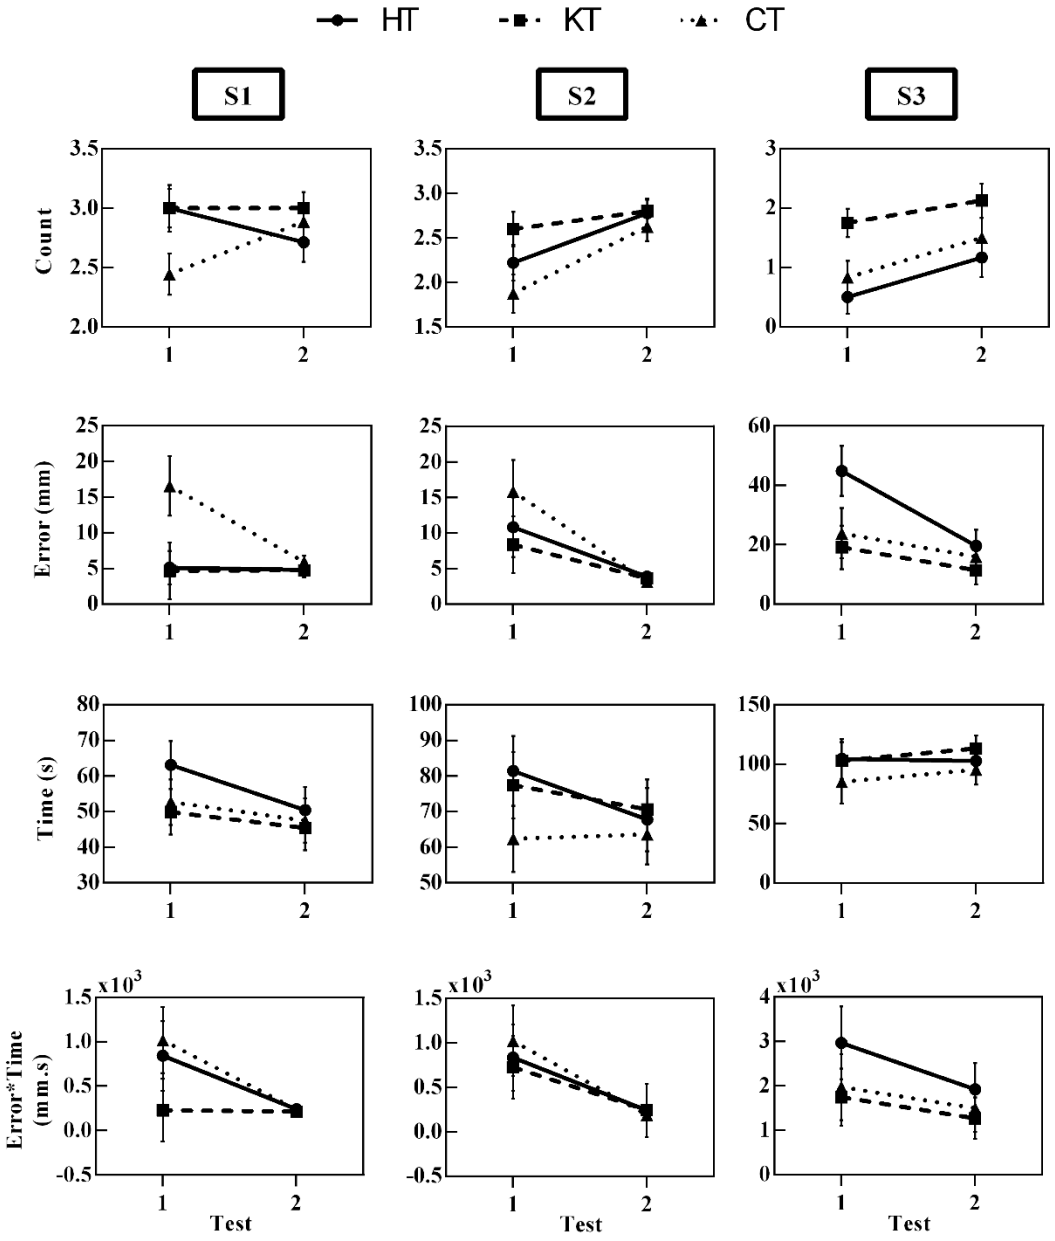

Figure 4.5. Plots of performance metrics for Tests 1 and 2 (virtual palpation task). From top to bottom: Number of Correct Selections (C_S), Mean Selection Error (E_S), Mean Selection Time (T_S) and Mean Selection Error × Mean Selection Time (ET_S), for samples S1 (12 mm tumour, flush with the surface) S2 (12 mm tumour, 3 mm below surface) and S3 (12 mm tumour, 6 mm below surface). Error bars indicate ± one standard error of the mean.

To investigate differences in performance between Test 1 and Test 2 (i.e. before and after Training), repeated-measures ANOVAs were conducted on each of the kinematic metrics (C_S, E_S, T_S and ET_S). There was a significant effect of Test for ET_S at S1 (S1_ ET_S), S2_E_S, S2_C_S, S2_ ET_S, S3_E_S and S3_C_S (all *p*s <.05). All other effects of Test were non-significant. There were no Test×Group interactions. These results suggest that, whilst

there may have been some improvements in performance from Test 1 to Test 2, there were no differences in the rate of this improvement between groups.

One-way ANOVAs for each metric revealed that there were several unexpected near-significant differences between groups at the first palpation session, in particular for C_S and E_S at S1. These are visualised in Figure 4.5. At S3, the difference was significant for C_S ($F(2, 19) = 6.474$, $p = .008$, $\eta_P^2 = .432$). Further analysis revealed that there were significant between-group differences in the mean number of 'no tumour present' selections made at session 1 ($F(2,86) = 11.139$, $p < .001$, $\eta_P^2 = .21$). Pairwise Bonferroni comparisons revealed a significant difference between the haptic training (HT) group ($M = 0.074$, $SD = 0.267$) and the control (CT) group ($M = 0.633$, $SD = 0.928$) (mean difference = -0.599, SE = 0.15, $p = .001$), and between the kinematic training (KT) group ($M = 0.000$, $SD = 0.000$) and CT (mean difference = -.0.633, SE = 0.146, $p < .001$). Another ANOVA revealed a significant between-group difference in number of timeouts ($F(2,84) = 3.781$, $p = .027$, $\eta_P^2 = .083$). Pairwise comparisons revealed a significant difference between HT ($M = 0.481$, $SD = 0.849$) and CT ($M = 0.033$, $SD = 0.183$) (Mean difference = 0.448, SE = 0.176, $p = .038$). These findings indicate the likely presence of systematic differences between groups during the first palpation session, whereby participants in groups HT and KT were less likely to make a 'no tumour present' selection than those in group CT. This bias was mirrored in the number of timeouts, i.e. groups HT and KT were more likely to timeout than group CT. Effectively, this was due to groups HT and KT more often searching the environment during the whole of the available time for each trial, whilst group CT were more likely to end trials early by making a 'no tumour present' selection.

4.3.1.2 Palpation probing strategy

Plots of palpation velocity (V_P) and within-participant variability in probing velocity (SDV_P) are shown in Figure 4.6.

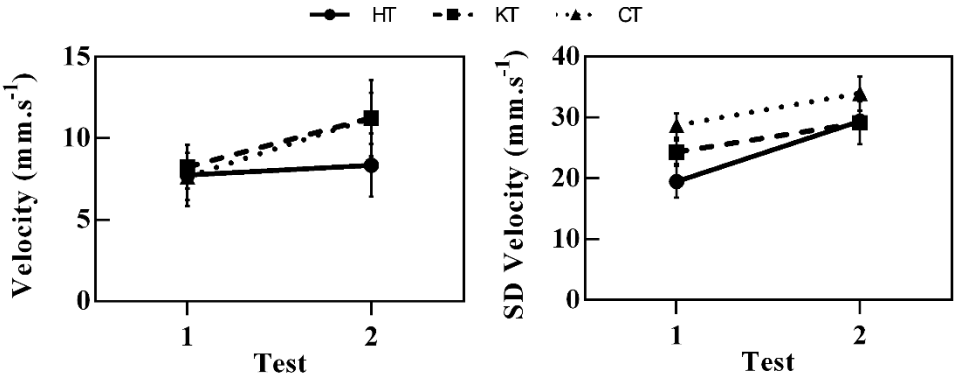

Figure 4.6. Plots of mean Velocity (V_P, left panel) and the Standard Deviation of the mean Velocity (SDV_P, right panel) for the haptic training (HT), kinematic training (KT) and control training (CT) groups, at Test 1 and Test 2. Error bars indicate the standard error of the mean.

There was a significant effect of Test on V_P ($F(1,26) = 6.789$, $p = .015$, $\eta_P^2 = .207$), with V_P increasing from 7.9 mm/s in Test 1 to 10.3 mm/s in Test 2. There was no Time×Group interaction ($F(2,26) = 1.106$, $p = .346$, $\eta_P^2 = .078$). The same pattern was observed for SDV_P with increased variability from Time 1 to Time 2 ($F(1,26) = 17.413$, $p < .001$, $\eta_P^2 = .401$). There was no significant Time×Group interaction ($F(2,26) = 1.38$, $p = .269$, $\eta_P^2 = .096$). These results suggest that probing strategy changed from Test 1 to Test 2 (V_P increased and it became more variable).

4.3.2 Haptic Training (HT) group

A plot of JND at each Training session for the HT group is shown in Figure 4.7.

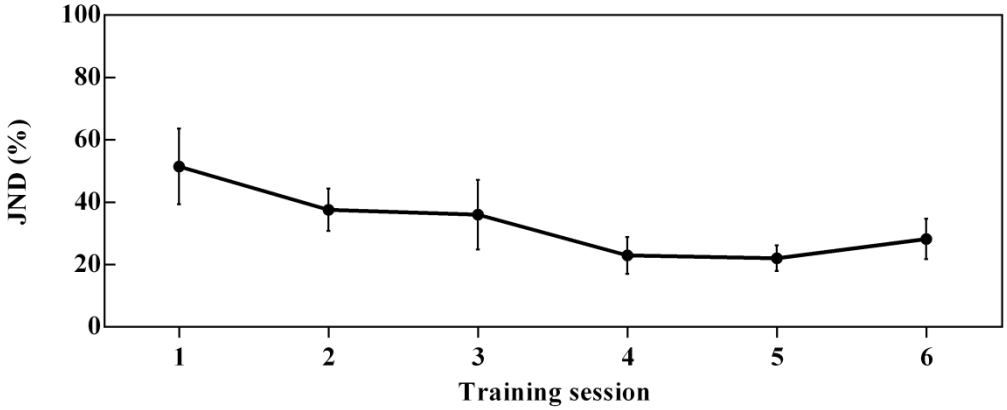

Figure 4.7. Plot of JND as a percentage of the difference between the two compared samples for the Haptic training (HT) group at each Training session. Error bars indicate the standard error of the mean.

Repeated-measures analysis of variance (ANOVA) revealed that there was a significant effect of Session on JND ($F(5,40) = 2.957$, $p = .023$, $\eta_P^2 = .27$). Further investigation of within-subject contrasts revealed a significant improvement from session 1 ($M = 51.52$ %, $SD = 36.42$ %) to session 4 ($M = 22.96$ %, $SD = 18.67$ %), ($F(1,8) = 8.154$, $p = .021$, $\eta_P^2 = .505$) and from session 1 to session 5 ($M = 22.08$, $SD = 13.01$), ($F(1,8) = 7.635$, $p = .025$, $\eta_P^2 = .488$). This suggests that subjects had significantly improved from session 4. The unexpected increase in JND at session 6 could be due to effects of fatigue or other unknown phenomena.

4.3.3 Kinematic Training (KT)

A plot of the mean tracking error (E_T) for each Training session is shown in Figure 4.8.

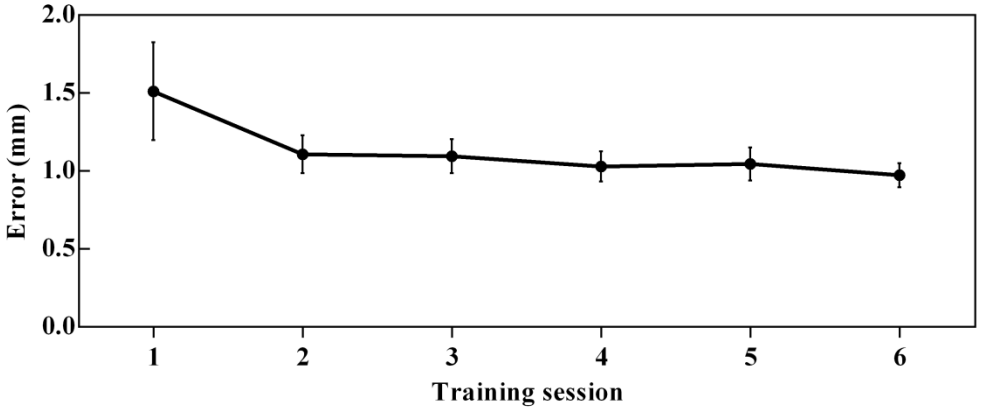

Figure 4.8. Plot of mean tracking error (Error in mm, E_T) at each Training session for the kinematic training (KT) group. Error bars indicate the standard error of the mean.

A repeated-measures ANOVA was carried out on all Sessions for the Kinematic Training group. Mauchly's Test of Sphericity indicated that the assumption of sphericity had been violated ($\chi 2$ (14) = 77.178, p = <.001), so Huynh-Feldt corrected values are reported. There was a near-significant effect of Session on performance (F(5,45) = 3.173, p = .096, η_P^2 = .261), with a mean error of 1.51 mm in session 1 and 0.97 mm in session 6. Pairwise comparisons revealed no significant differences between sessions. These results suggest that subjects were not able improve their performance over the training period.

4.4 Discussion

The aim of this study was to investigate how the CNS acquires essential compliance information about an object during tool-based interactions, and what intervention strategies may lead to improvements in our ability to detect small compliance differences. A three-dimensional virtual task that is analogous to a real world medical procedure was used to objectively test the effectiveness of two different training interventions. The Training sessions were designed with the aim of exploiting the potential effect of becoming attuned to compliance cues (for the HT task) and the effect of using appropriate probing strategies (for the KT task), thereby also indicating the way in which humans acquire and use information during execution of this task.

The HT condition required participants to repeatedly assess compliance differences that were close to their JND of this variable. In the previous chapter, this method was shown to successfully improve compliance discrimination over time. This effect was attributed

to an increased perceptual sensitivity, and not to systematic changes in probing strategy. However, findings in Chapter 3 also suggest that probing strategy plays an important role in our perception of compliance. With the aim of providing a practice environment that guides the subject towards movement patterns that maximise their perception of compliance, the KT condition required participants to repeatedly follow a moving cursor at a speed range that was representative of those that previously resulted in maximal compliance discrimination performance. The CT condition was not expected to promote learning and so should have allowed for a between-group comparison of the relative contributions of HT and KT towards improving compliance perception.

Results indicate that the HT group was able to improve during Training, and that all three groups were able to improve from Test 1 (pre-Training) to Test 2 (post-Training). Some improvement due to practice during Test 1 was expected across all groups. However, there was no difference in the rate of this improvement between groups. This nullifies our a-priori hypothesis that the HT and KT conditions would result in greater improvements than the CT group due to a) haptic sensitisation for HT and b) the use of effective probing strategies for KT. An analysis of the probing kinematics for the palpation task revealed that probing velocity (V_P) increased from Test 1 to Test 2 for all participants in all three groups. The within-participant variability also increased, which could simply be attributed to faster movement speeds resulting in more variable movements, a well-documented effect in the motor control literature (Latash, Scholz, & Schöner, 2002). An alternative explanation for this effect could be that more experienced subjects used more exploratory probing strategies, i.e. whereby more information about the environment could be gained through the use of a variable V_P. Thus, it seems that subjects were able to improve at Test 2 when a larger V_P was used. Further, findings from Chapter 3 (that probing strategy remains constant with time) are contradicted here, as evidenced by the increase in V_P for the HT and KT conditions. To examine what factors could give rise to this effect, it is useful to consider the differences from the HT and KT tasks to the virtual palpation task, which can be considered to be more complex in a number of ways.

First, the virtual palpation task is three-dimensional rather than one-dimensional, thus requiring additional movements in the lateral and longitudinal directions to enable a full exploration of the virtual sample. It seems possible that there was an influence of movement speed on the vertical probing action from movements in the other axes, i.e. that there was a correlation between the speed adopted in the lateral and longitudinal axes

to that in the probing direction. Thus, an increase in V_P in the vertical axis could be attributed to an increase in speed in the other axes. Second, it is possible to experience progressive changes in the compliance of the sample when moving from the outside edges of an inclusion towards its centre. Once the edge is detected, an effective strategy might be to adopt faster movement speeds to more quickly identify the centre of the tumour by confirming the direction in which the compliance increases. Such a strategy may become refined with practice. Third, the deformation of the sample's surface in response to indentation force could provide rich visual information relating to compliance. Hypothetically, the usefulness of such visual cues may vary as a function of indentation speed, giving rise to systematic adjustments in probing strategy to optimise the information that is available to the CNS. Finally, there is a finite time constraint to make a selection. This could result in faster movements due to a trade-off between exploring as much of the sample's surface as possible (to increase the possibility of a positive detection) and obtaining reliable information via an appropriate probing strategy. Indeed, if our ability to perceive compliance improves with time, an increase in probing speed by skilled subjects would seem a useful strategy for obtaining a similar level of information at each indentation as unskilled ones. The overall amount of information obtained at each trial is then increased through achieving more interactions within the available time. These factors lay outside the focus of this study but they should be considered in any future experiments that make use of these tasks.

Similarly to the probing strategy findings, performance increased from Test 1 to Test 2 for all groups, and there were no between-group differences in improvement rates. This initially suggests that none of the training interventions were more effective at improving this skill. However, there are issues with the experimental data, including between-group differences at Test 1, an effect that seems due to an increased number of 'no tumour present' selections for the CT group compared to the number of trial timeouts for HT and KT.

Upon reflection, it is possible that the KT task was not adequate for teaching an effective movement strategy for this particular task: the absence of haptic feedback did not allow for the generation of an internal representation of what movement strategy results in good discrimination performance. If this is true, it would be difficult for participants in the KT group to relate the probing strategies used during training to the virtual palpation task. Future investigations should include a Training condition that combines the HT and KT

tasks by providing congruent haptic feedback during training with a range of movement speeds. However, another aspect to consider is that of increasing the learning rate of appropriate probing strategy, which has obvious advantages in clinical settings (e.g. reduced training times). Furthermore, the virtual palpation task is an extremely simplified version of a real medical procedure, and whilst it may serve for initial validation of training interventions such as this one, care should be taken to systematically verify any findings during a real medical procedure before any conclusions can be made on the effectiveness of such interventions

Active behaviour involves sensory, cognitive and motoric mechanisms, all of which are coupled within the internal models that allow humans to fine tune their actions (i.e. to extract as much information as possible from the environment, for example(Wolpert et al., 2011). This learning process is thought to be driven by execution error, whereby error (e.g. the difference between a probe's desired and actual positions) is used to correct future movements. Applying forces to the hand to guide probing (or any other skilled action) may seem a sensible method of illustrating an effective strategy to the subject. However, according to the theory, these 'error reduction' forces would act to limit the amount of information available to the CNS to make future corrections. This gives rise to the following question: can execution error be adjusted to aid motor learning? This topic will be investigated in the following chapter.

CHAPTER 5

INFLUENCE OF SUPERIMPOSED ERROR ADJUSTMENT FORCES ON SENSORIMOTOR LEARNING

ABSTRACT This chapter addresses the notion of increasing the rate at which complex motor skills are learned. As outlined in Chapter 1, this work is critical for addressing the increasing complexity of surgical skills, and the limited training time that is available for learning said skills. Motor performance might be enhanced through haptic guidance but recent learning theories have led to the counter-intuitive hypothesis that disruption benefits motor learning. This chapter presents three experiments that investigate motor learning in workspaces with complex novel force fields in the presence of haptic assistance and disturbance. Experiment 1 showed that haptic guidance hindered learning. Experiment 2 explored generalisation across the workspace with three groups who experienced: (1) haptic assistance (error reduction); (2) no guidance; or (3) a constant disruptive force (error augmentation). Haptic assistance showed the worst learning whilst those exposed to disruptive forces evidenced the most training errors, but steepest training curves and best generalised learning. Experiment 3 revealed that a random combination of assistive and disruptive forces enhanced learning, but learning was impaired when workspace exploration was reduced. Taken together, these results demonstrate that humans can: (i) detect and rapidly adapt to a simple externally imposed force; and (ii) benefit from the presence of task-irrelevant disturbance due to increased workspace exploration.

5.1 Introduction

Neonates must determine the complex relationship between perceptual input and motor output in order to learn how to move their arms effectively. This process is repeated throughout life as humans acquire new skills or recover from injury. Technological advances have enabled robotic systems designed to accelerate such motor learning, for example in laparoscopic surgical training and stroke rehabilitation. Nevertheless, it is unclear how these devices might be optimised for enhanced learning. Engineering effective solutions requires an understanding of human motor learning – a process that can be conceptualised as involving two broadly interacting but distinct mechanisms; model free (MFL) and model-based (MBL) learning. These mechanisms represent qualitatively different computational approaches to learning and refining a skill (Sutton & Barto, 1998). MFL involves the refinement of movements based on the success or failure of prior interactions with the environment; a reinforcement driven trial-and-error process (Bush & Mosteller, 1953; Rescorla & Wagner, 1972; Thorndike, 1901). In contrast, MBL describes the creation of neural model(s) of task dynamics capable of computing optimal strategies in novel environments. These 'forward models' predict the consequences of actions given the state of the system (Miall & Wolpert, 1996). Whilst motor learning studies and theories have largely focused on MBL (Shadmehr & Krakauer, 2008), these processes appear to work co-operatively in the acquisition and refinement of skilled behavior, with MFL being a necessary precursor to MBL (Daw, Niv, & Dayan, 2005; Dayan, 2009; Fermin, Yoshida, Ito, Yoshimoto, & Doya, 2010; Gläscher, Daw, Dayan, & O'Doherty, 2010; V. S. Huang, Haith, Mazzoni, & Krakauer, 2011).

MFL and MBL both rest upon the ability of the system to identify task relevant error. This ability is compromised by the existence of task-irrelevant noise. A simple model of MFL shows that motor output on trial *n* is determined by the difference between the internal (*Z*) and external state (the environment, *U*). The error on this trial (Y_n) is the result of this difference plus the inherent noise (*E*):

$$Y_n = (U_n - Z_n) + E_n \tag{5.1}$$

Motor learning requires the system to change the internal state of subsequent movements (Z_{n+1}) in response to the discrepancy between the internal state and the desired movement (U_n - Z_n), where L is the learning rate:

$$Z_{n+1} = Z_n + L(U_n - Z_n) \tag{5.2}$$

Combining equations (5.1) and (5.2) gives:

$$Z_{n+1} = Z_n + L(Y_n - E_n) \tag{5.3}$$

Equation (5.3) shows that MFL requires task relevant error to be distinguished from noise. In this context, noise is an unavoidable and undesirable factor within motor learning. On this basis, haptic training devices have often applied assistive forces (Hesse, Schmidt, Werner, & Bardeleben, 2003; Krebs et al., 2008; Prasad et al., 2003). Nevertheless, there is evidence to suggest that learning can be accelerated through the application of *disruptive* forces (Cesqui, Aliboni, et al., 2008; Emken & Reinkensmeyer, 2005; Huang & Shadmehr, 2007; Lee & Choi, 2010; Reinkensmeyer & Patton, 2009; Schmidt & Bjork, 1992). In the account of MFL described above, it appears counterintuitive to apply a disruptive force - as this constitutes additional noise with which the system must contend. It is possible that this paradox can be reconciled if disruption benefits motor learning via its effect on MBL as predicted by recent theories of motor learning.

MBL requires the system to extract the invariant rules that govern a range of input–output mappings. The difficulty faced by the system relates to the large number of internal parameters that connect the sensory input to the motor output (the larger the number of parameters, the greater the 'dimensionality' of the parameter space). Structural learning theory (Braun, Mehring, & Wolpert, 2010) suggests that the motor system reduces the dimensionality of the parameter space by predicting the topology of the input–output mappings of tasks sharing a similar structure. This allows the system to restrict and control a subsection of the whole parameter space through adjustment of 'meta-parameters'. Braun et al. (2009) formalised such structural learning within a Bayesian framework. In this conceptualisation, a hidden variable (μS) can decrease the dimensionality of the parameter space associated with a novel environment. The task facing the nervous system is to extract μS so that the joint probability distribution can be computed across the workspace. This process requires the system to infer the structure between the hidden variable and the directly measurable processes (the observables) via

two steps: (i) computing the posterior probability P(S|X) of the structure (S) given the data (X), and (ii) computing the posterior probability P(μS|S, X) of the parameter μS (given S and X). It can be seen that providing more data (i.e. increasing X samples) will improve the posterior probability estimates and thereby potentially accelerate the structural learning process. It is proposed that low dimensional force disruption (i.e. where a meta-parameter can capture the change in an externally applied force field through a relatively simple adjustment) increases data sampling. Notably, structural learning predicts rapid compensation to a disturbance of the input–output topology if the disruption is at a low dimension of the parameter space. It follows that a low dimensional disruption has the potential to improve motor learning.

In line with this notion, Wu, Miyamoto, Gonzalez Castro, Ölveczky, & Smith (2014) demonstrated that the intrinsic movement variability associated with motor commands (from Z_n to Z_{n+1} to Z_{n+2} …) predicts individual rates of motor learning. Relatedly, (van Beers, 2009) has shown that the random effects of planning noise accumulate; in contrast to task-relevant errors which show close to zero accumulation (explained by effective trial-by-trial corrections). On these grounds, it has been argued that intrinsic movement variability leads to motor exploration, which sub-serves motor learning and performance optimisation. Indeed, the idea that action exploration can drive learning has long been mooted in theories of operant behaviour (Sutton & Barto, 1998) and human development (Bruner, 1973; E. Gibson, 1988; Thelen, 1989). These findings suggest that providing guidance may impair learning through error reduction but raise the intriguing possibility that haptic devices could help learning by adding disruption to the training process. These ideas were tested in three experiments.

5.2 Experiment 1: Active versus passive learning

The aim of experiment 1 was to test the hypothesis that error is necessary for learning. Subjects were asked to move their arms to follow a target in a force field inherent within a haptic robotic device (i.e. the inherent environment that the operator interacts with when controlling the device). One group was provided with no haptic assistance (the Active-Control group) whereas another group was guided through the requisite movements, with little need to deviate from the desired path (Guidance group). A third group, a control,

observed the device as it moved autonomously but did not physically interact with it (Vision group).

5.2.1 Methods

5.2.1.1 Materials

The haptic assessment toolkit (HAT, described in Chapter 2) was used to build the virtual visuohaptic environment and experimental framework used for the task. The HapticMASTER was used to deliver forces and record the position and velocity of the device's end-effector at a rate of approximately 1 kHz. The standard experimental setup described in Chapter 2 was used.

5.2.1.2 Participants

Twenty four right handed participants (6 female, aged 20 – 28, M = 24.9, SD = 3.9) were recruited and randomly allocated to one of three training groups. Each participant was required to attend one session of approximately 45 minutes. All participants were Engineering or Psychology postgraduate students at The University of Leeds and were not compensated for participation. The research was approved by, and conducted under the guidelines established by the School of Psychology Research Ethics Committee at The University of Leeds.

5.2.1.3 Experimental design

The experiment consisted of one session in which Pre-test, Training and Post-test blocks were completed. There were three groups: the Active-Control group were required to actively move the device end-effector along pre-specified trajectories by following a guide circle as accurately as possible. The Guidance group rested their hand on the end-effector whilst it guided them around the same paths. They were instructed to completely relax their arm. The Vision group visually observed the end-effector as it moved around the same paths, without touching the end-effector and with both arms resting by their sides. A within-groups design was used to analyse the performance of the Active-Control group as they progressed through training. A mixed design was used to assess performance at baseline (Pre-test), whilst a between-groups design was used to calculate improvements from Pre-test to Post-test for all groups.

5.2.1.4 Stimuli

Error adjustment forces were implemented using the mass-spring-damper model described in equation (5.4). A constant simulated end-effector mass of 3 kg was used for all conditions. For the Active-Control condition, stiffness and damping coefficients of 0 N/m and 1 Ns/m were used to simulate the feeling of moving a free object through a viscous fluid (e.g. air) in a null-gravity environment. The damping value was used to maintaining system stability. The Guidance and Vision conditions were created using stiffness and damping coefficients of 25 kN/m and 110 Ns/m (again, to maintain system stability) so that the end-effector would follow the target cursor with negligible deviations to the desired trajectory, regardless of any forces applied to the end-effector. Visual stimuli consisted of solid black paths on a white background and cursors (circles) to indicate the device and target positions of the end-effector. The current position of the end-effector was represented by a filled red circle (the 'device cursor'), while the target position was represented by a filled green circle (the 'target cursor') which moved along the path at a pre-defined speed. During trials when the device moved autonomously i.e. those requiring no input from the participant, only the target cursor was displayed because the actual and target positions were the same.

$$F = m\ddot{X} + c\dot{X} + kX \tag{5.4}$$

The paths used at each stage of the experiment are shown in Figure 5.1. The practice stage consisted of two trials (cycles) of simple paths. The paths were generated by a PhD student and collaborator on the project, Aaron Fath. The first path, P1, was an equilateral triangle. P2 was a square. The length of all sides of P1 and P2 was 0.3 m. Both started in the lower left corner. The triangle proceeded in a clockwise direction and the square proceeded counter-clockwise, both at a constant speed of 0.3 m/s. The purpose of the practice trials was to familiarise participants with the novel interface, but precautions were taken to prevent learning or priming from occurring during the practice trials. The starting location was chosen to be in the lower left because it was the farthest corner from the starting location of the experimental trials. The practice trials were run in opposite directions so that participants were not primed to expect a specific direction of motion during the experiment. The Pre-test, Training and Post-test paths consisted of sinusoids that were transformed such that they oscillated around an invisible guide circle. The start location for all experimental trials was at angle θ = 0; they proceeded in a counter-clockwise direction until arriving back at the start location. Each path changed every $\frac{\pi}{3}$

radians. This resulted in each path consisting of six regions, specified by the equation $\rho = r + A \cdot \sin(v\theta)$, where r was the radius of the guide circle, A was the amplitude of the sinusoid, and v was the sinusoid's spatial frequency. P3 varied amplitude at each landmark, P4 varied spatial frequency, and P5 varied both amplitude and spatial frequency.

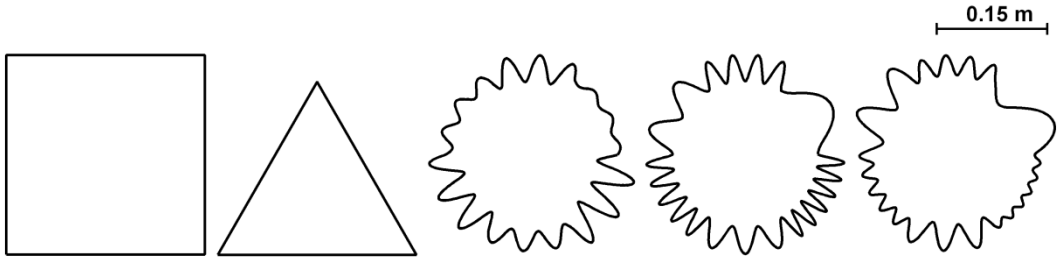

Figure 5.1. Illustration of the paths used in the experiment from left to right: square (P1), triangle (P2), amplitude-variable (P3), frequency-variable (P4), and amplitude- and frequency-variable (P5) paths. The 0.15 m scale is with reference to the HapticMASTER workspace, not the visual scene. P1 and P2 were used for practice paths. P5 was used in Pre-test and Post-test. The experimental trials sequentially alternated between paths P3 and P4.

The circle that the sinusoids oscillated around had a radius of 0.130 m for all paths. P3 held the spatial frequency constant at 18 cycles per polar cycle, but varied amplitude. In order, the six regions had amplitudes of 0.005, 0.010, 0.015, 0.020, 0.025, and 0.030 m. P4 held the amplitude constant at 0.020 m, but varied spatial frequency. In order, the six regions had spatial frequencies of 6, 12, 18, 24, 30, and 36 cycles per polar cycle. P5 varied amplitude and spatial frequency, using the same ranges of values as the other two paths. The frequency/amplitude pairs were matched such that the highest spatial frequency was paired with the smallest amplitude, the second highest spatial frequency with the second smallest amplitude, and so on. The six frequency/amplitude pairs were, in order, 6/0.030, 24/0.015, 12/0.025, 30/0.010, 18/0.020, and 36/0.005.

During active control, the guide speed was chosen to challenge participants to keep pace with it whilst staying on the path. There was considerable inertia (3 kg) to be accounted for during changes in direction, thus the speed of the guide was varied as a function of curvature in order to emulate realistic movement strategies. Angular speed varied between the six sections of each path, but each section had its own constant angular speed. The constancy of angular speed over the course of a cycle is what resulted in the guide slowing down for the turns. At a crest or trough, the path is perpendicular to the radius

and as such, a unit of angle corresponds to less path distance than it does elsewhere on the path. The angular speeds for each section were chosen such that all sections of all paths had a common mean path speed of 0.145 m/s. The sections of P3 had the following constant angular speeds, in order, 0.32π, 0.17π, 0.26π, 0.14π, 0.21π, and 0.12π rad/s. P4's angular speeds were 0.30π, 0.13π, 0.22π, 0.11π, 0.17π, and 0.09π rad/s. P5's angular speeds were 0.26π, 0.17π, 0.19π, 0.19π, 0.17π, and 0.26π rad/s. The differences in the angular speeds between paths were the result of differences in path lengths. Paths could not be given the same lengths while varying amplitude and frequency unless the radii of the paths' guide circles were changed appropriately. However, maintaining a constant radius across paths was deemed a more important constraint than maintaining a constant path length, as this may affect the dynamics of the task (i.e. a smaller radius would decrease the size of the operating workspace, and have a greater contribution towards the spatial curvature differential along the path). The path lengths for Paths 3-5 were 1.59, 1.96, and 1.45 m, respectively. Whilst movement was evaluated in the up-down and side to side planes, it was unconstrained in the forwards-backwards plane.

5.2.1.5 Task and procedure

Participants read an information sheet describing the experiment and were then given specific written instructions about how to complete the task, with an opportunity to ask the experimenter questions related to task requirements. All trials were performed with the right (dominant) hand. Participants were free to take a break and sit down after the completion of any trial, as needed. The Training trials took approximately 30 minutes to complete and the entire experiment took approximately 45 minutes. The Pre-test and Post-test blocks were identical for all three groups, and consisted of 3 trials of path 3 (P3) under active conditions. For all of the groups, the Training trials consisted of two trajectories, one for each of the other two paths. The order of presentation of the paths was counterbalanced, alternating between P4 and P5 at each trial pair. This allowed the examination of participant's ability to generalise their learning by requiring them to tackle simultaneous changes to frequency and amplitude (this was only required with P5), a feature that they did not encounter during Training. There were 30 Training trials in total (15 of each path). Once the experiment started, the device was initialised and the visual environment (see Figure 5.2) was displayed on the monitor.

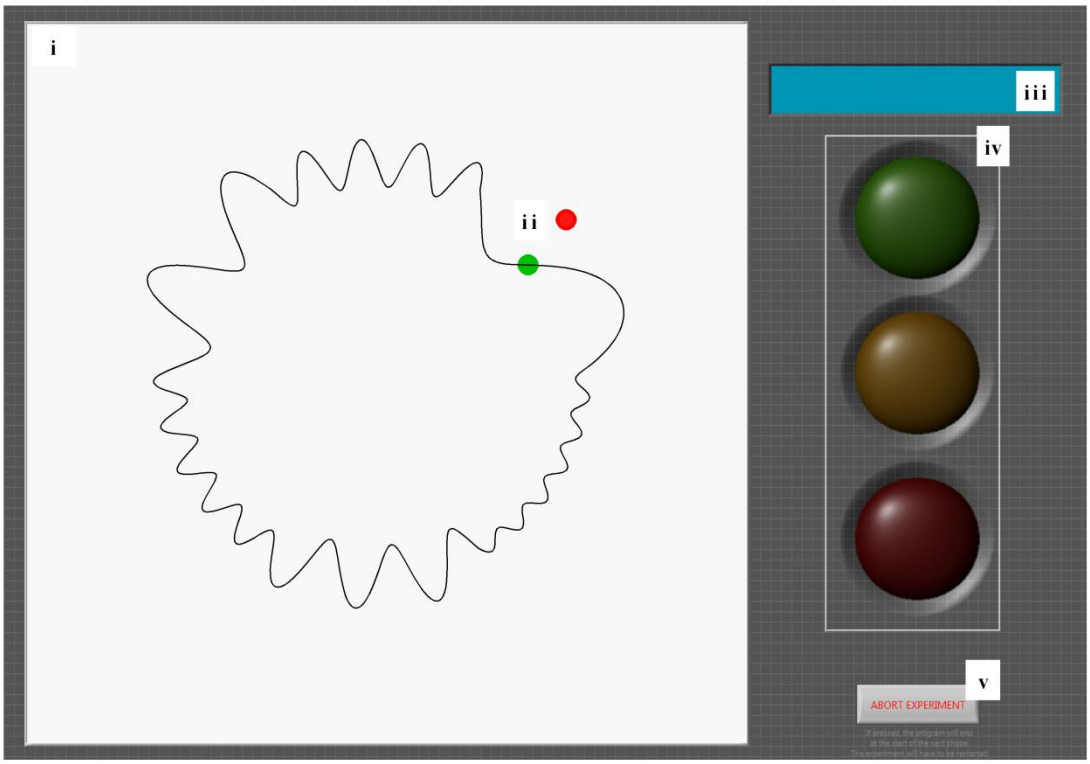

Figure 5.2. Screenshot of the visual environment, showing: i) the visual scene; ii) path P5 and target (green) and device (red) cursors; iii) the instructions panel; iv) a semaphore display, indicating whether participants should release the device (red), be ready to start the next trial by (with exception of the Vision condition) holding the device (yellow), or that the trial had begun (green).

During all trials that required holding the device (that is, all trials except those in the Vision Training condition), participants were instructed to 'hold the device when ready' in preparation for the next trial. Once the device was held, the path was displayed on the visual scene. After 1 second, the device moved in position control (along with the participant's hand) from the origin to the start ('home') location of the path for that trial. The home location was a hollow black circle that fitted around the device cursor. Then, a 3 second countdown was displayed in the instructions panel. When the countdown ended, the green light lit up and the device either became compliant (for all Pre-test and Post-test blocks and the Active-Control Training condition), or moved autonomously along the path (for the Guidance and Vision Training conditions). Participants in the Active-Control group were required to align the device cursor with the target cursor as accurately as possible. Participants in the Guidance group were instructed to leave their arm 'slack' while it guided them around the path. Participants in the Vision group were required to rest both arms by their sides whilst watching the visual scene to observe the relative

movements of the device and target cursors. To minimise fatigue, participants were given the option to rest and sit for an unlimited time between trials. The end of the trial was registered when both the target and device cursors returned to the 'home' location. The red light then lit up and participants were instructed to 'release the device'. Upon releasing the device, the yellow light lit up and participants were again instructed to 'hold the device when ready'. This process was repeated until all trials were complete. When all trials were complete, an 'experiment complete' message was displayed.

5.2.1.6 Data analyses

The HAT post-processing utility was used to generate the standard kinematic output metrics discussed in Chapter 2 from the raw experimental data. For each trial, extreme values were categorised as ones that were above 99% of the distribution (i.e. those with an absolute Z-score greater than 2.58) were removed for every metric (less than 2% of data). A mean average of Pre-test trials was subtracted from the mean average of Post-test trials to provide a measure of the learning rate (where lower scores indicate greater learning). Two participants failed to adequately follow task instructions and data from these subjects were also statistical outliers (z-scores > 2.58) and therefore removed before inferential statistics were conducted.

5.2.2 Results

To select a kinematic metric or metrics that objectively captured performance data that were specific to this task, it was important to consider the level of congruency between the objectives of the task and the available kinematic metrics. The instruction given to participants was to 'align the device cursor with the target cursor as accurately as possible'. Out of the available metrics (movement time, T_M, trajectory error, E_T, standard deviation of the trajectory error, SDE_T, path error, E_P, and normalised jerk, J_N), a measure of alignment accuracy between the two cursors is directly given by E_T, i.e. the mean radial distance between the device and target cursors at each trial (refer to Chapter 2). Thus, E_T was initially chosen as a suitable metric for objectively assessing performance on this task.

First it was important to determine whether common baseline performance was exhibited between groups at Pre-test. A mixed ANOVA was used. Mauchly's test indicated that the assumption of sphericity had been violated, $\chi^2(2) = 23.317$, $p < .001$, and therefore Greenhouse-Geisser corrected tests are reported ($\varepsilon = .658$). The results showed a main

effect of Time on E_T ($F(1.414, 28.287) = 5.673$, $p = .015$, $\eta_p^2 = .221$), no Time*Group interaction ($F(2.829, 28.287) = .422$, $p = .696$, $\eta_p^2 = .045$), and no between-subject differences ($F(1, 19) = .724$, $p = .498$, $\eta_p^2 = .071$). These results indicate that all groups performed at the same level and that they were able to improve during the Pre-test trials. One subject in the Active-Control group obtained a Z-score greater than 2.58 and thus they were consequently removed from the analysis due to a failure to perform the task.

A plot of each group's error reduction in E_T from Pre-test to Post-test (the 'normalised E_T') is shown in Figure 5.3. A One-Way ANOVA revealed that there were no significant differences between groups ($F(2,20) = .565$, $p = .577$, $\eta_p^2 = .053$).

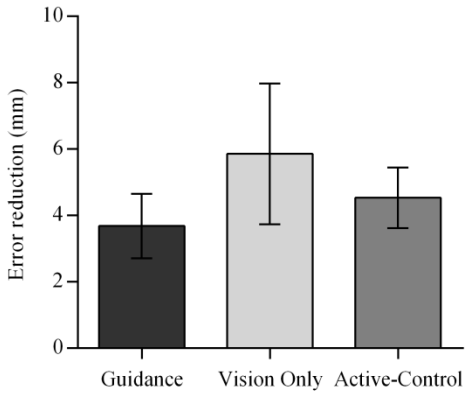

Figure 5.3. Plot of improvement in E_T from Pre-test to Post-test for all groups. Error bars indicate ± one standard error of the mean.

Further investigation of the data revealed large variations between individuals in movement strategies adopted. This suggests that some participants did not execute movements as expected: further analysis showed that some lagged the target, whilst others led ahead of it. However, participants tended to stay on the path taken by the target cursor. This was evidenced by a significantly lower variability in the Path Error (E_P) metric (the distance between the device cursor and the nearest location along the path). Whilst E_P does not directly assess performance data that is congruent with the task instructions, it is robust to the aforementioned caveats of E_T on this task. Thus, E_P was chosen as a more adequate performance measure. The previously removed outlier (from the Active-Control group) was added back into the analysis. One subject in the Guidance group obtained a Z-score greater than 2.58 for their E_P score and thus they were consequently removed from the analysis due to a failure to perform the task.

As with the E_T metric, a mixed ANOVA was used to assess baseline performance at Pre-test. Mauchly's test indicated that the assumption of sphericity had been violated, ($\chi^2(2) = 10.157$, $p = .006$), and therefore Greenhouse-Geisser corrected tests are reported ($\varepsilon = .707$). The results showed that there was a main effect of Time on E_P ($F(1.414, 28.287) = 5.673$, $p = .015$, $\eta_p^2 = .221$), a Time*Group interaction ($F(2.829, 28.287) = .422$, $p = .696$, $\eta_p^2 = .045$), and no between-subject differences ($F(2,20) = .899$, $p = .423$, $\eta_p^2 = .083$), indicating that whilst all groups were able to improve during the Pre-test trials, there were no between-groups differences in performance.

Next, the Active-Control group's Training data were examined. Figure 5.4 shows performance in E_P obtained by the Active-Control and Guidance groups at each Training trial pair (the average of paths 1 and 2). Note that like the Guidance group, the Vision group had zero error (the device autonomously followed the target) but is not plotted because there was no physical interaction between the participant and the device during the task.

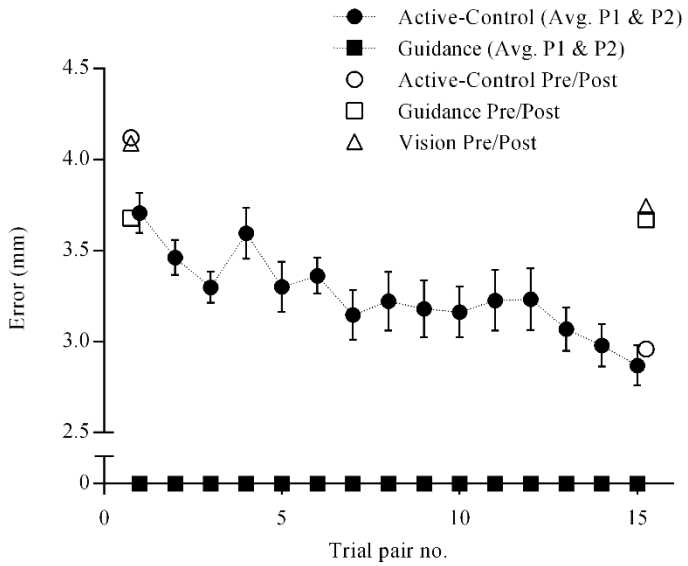

Figure 5.4. Absolute E_P obtained for the Active-control group at each Training trial for trajectories 1 and 2, indicating that the Active-Control group were able to gradually improve their performance during Training. The 'Pre/Post' values indicate the scores obtained for the final pre-test and first post-test sessions. Error bars indicate ± one standard error of the mean.

A repeated-measures ANOVA revealed a significant effect of time for Path 1 ($F(14, 98) = 4.425$, $p < .001$, $\eta_p^2 = .387$) and Path 2 ($F(14,98) = 2.289$, $p = .009$, $\eta_p^2 = .246$),

suggesting that participants were able to improve their performance during the Training trials.

In line with the a-priori hypothesis that error feedback is critical for learning, E_P data for Pre-test and Post-test were examined. A plot of each group's error reduction in E_P from Pre-test to Post-test (the 'normalised E_P') is shown in Figure 5.5. A one-way ANOVA revealed a significant between-groups difference on the normalised E_P ($F(2,20) = 5.237$, $p = .015$, $\eta_p^2 = .344$). Pairwise Bonferroni comparisons revealed a significant difference between Guidance ($M = 0.34$, $SD = 0.50$) and Active-control ($M = 1.16$, $SD = 0.78$) (mean difference = 0.817, $SE = .307$, $p = .038$) and between Vision ($M = 0.30$, $SD = 0.42$) and Active-control (mean difference = 0.858, $SE = .296$, $p = .023$). There was no significant difference between Guidance and Vision (mean difference = 0.041, $SE = .307$, $p = .99$).

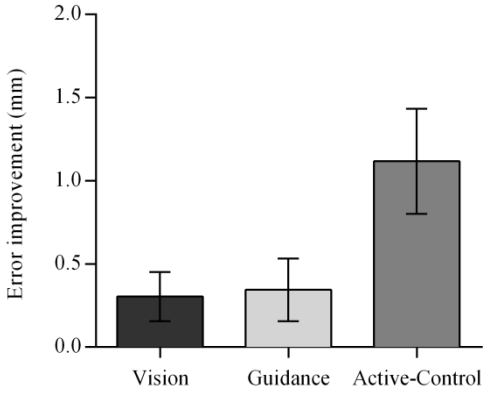

Figure 5.5. Plot of the mean difference between Post-test and Pre-test (normalised E_P) for each group. Error bars indicate ± one standard error of the mean.

These findings suggest that the Active-control group improved their performance with Training whilst the Guidance and Vision Training interventions had no effect on performance.

5.2.3 Discussion

In line with existing literature (Sigrist et al., 2012), completely passive movements do not provide error feedback and thus do not allow for the tuning of internal models of motor control (Wolpert et al., 2011). The fact that E_T captures spatial and temporal performance characteristics and that there were no group differences in this metric from pre-test to post-test disagrees with previous findings reported in the literature that haptic guidance can help to teach the temporal, yet not the spatial, aspects of a task (Lüttgen & Heuer,

2012a). It was predicted that to drive learning the motor system must be able to experience error and correct it. The results confirmed this hypothesis. Indeed, a comparison of the Guidance and Vision groups relative to the Active-control group indicated that moving a passive limb around a path is no more effective for learning than simply watching the device end-effector autonomously move around the path.

The notion that error correction aids motor learning has led researchers to consider forces that augment execution error as a possible means of increasing the rate of motor learning (Cesqui, Macri, Dario, & Micera, 2008; F. Huang et al., 2007; Reinkensmeyer & Patton, 2009; Sigrist et al., 2012). This is on the premise that such interventions may lead to a "richer experience" of task dynamics (F. Huang et al., 2007) and thus a faster internal construction of the task's structure (i.e. the parameter space). Such interventions appear to fit well within the framework of increased learning rate via increased parameter space exploration using 'structured noise' (i.e. an error augmenting force) presented here. However, a theoretical account of the mechanisms that underpin parameter space exploration, and how this leads to motor skill acquisition, is yet to be established.

Within the framework outlined above, optimal parameter space exploration is likely to lead to a maximum rate of motor learning. One way of increasing parameter space exploration could be forcing the learner to experience properties of the task that they would not necessarily experience under normal conditions. Conceptually, destabilising the environment through the use of error augmentation forces could be one way of achieving this. Experiment 2 investigates this concept by testing the effects of superimposed error augmentation ('noise') forces over a complex visuohaptic task. Based on the above and in concert with existing literature on error augmentation (Cesqui, Macri, et al., 2008; Reinkensmeyer & Patton, 2009; Sigrist et al., 2012), it is predicted that error augmentation will lead to greater workspace exploration (which acts as a proxy to parameter space exploration). Based on the model-based (MBL) and model-free (MFL) mechanisms discussed previously, error augmentation will largely modulate MBL processes, but confer no benefit on MFL.

5.3 Experiment 2: Assistance versus disturbance forces

Experiment 1 showed that error correction is integral to motor learning - a finding consistent with existing literature. These data raise the possibility that exploiting the system's innate error corrective process can be a way of increasing the rate of motor skill formation.

Experiment 2 tests the prediction that learning rate can be accelerated through the provision of disruptive forces. Training with partially assistive (Assistance group), disruptive (Disruption group) or no guidance (Active-Control group) forces was examined. All subjects completed movements in an artificial environment with a complex force field, designed to produce sufficient novelty to prevent rapid learning via a low dimensional change to an existing forward model. The force field was a heterogeneous force bias distribution that varied along the device's workspace in two dimensions (see Figure 5.6). The end result was an environment that is not readily amenable to adaptation learning. It was conjectured that participants would need to develop a relatively novel model over time (via MFL processes) in order to perform the task well.

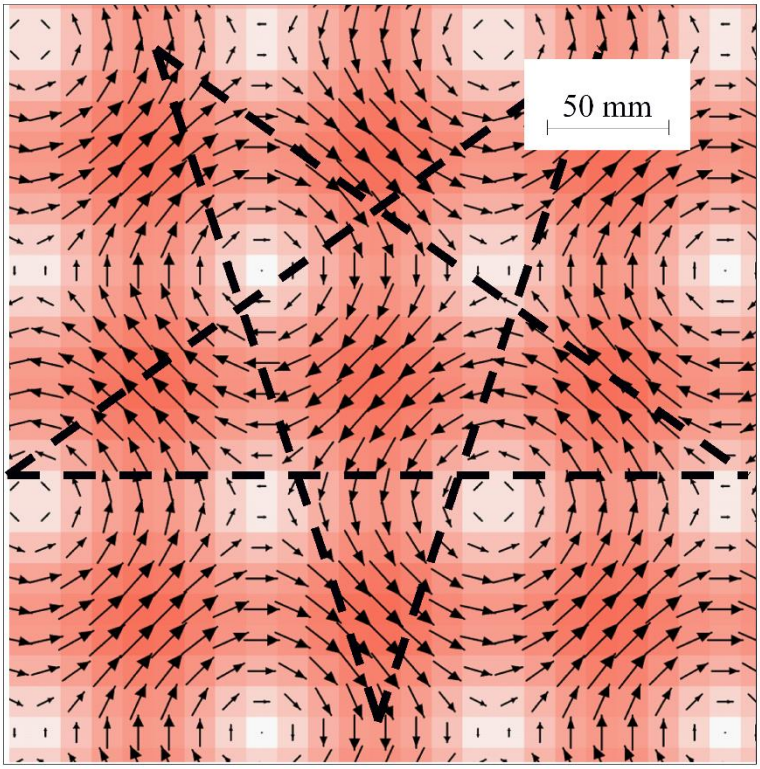

Figure 5.6. Force distribution of the novel force field, also showing the relative location of the path: feather plot of the workspace distortion force field for a section of the workspace measuring 0.16 x 0.16 m. Arrows indicate the direction and proportional magnitude of the force vector at discrete locations within the workspace. Relative magnitude is also represented using a colour map, where white = no force and dark red = high force.

The force field was implemented using a sine wave, which acted along the x and y axes of the workspace, defining a two-dimensional force vector, $F_{x,y}$, as described by equation (5.5).

$$F_{x,y} = Asin\left(\frac{1}{P_x}(x) + \theta_x\right) + Asin\left(\frac{1}{P_y}(y) + \theta_y\right) \quad (5.5)$$

Amplitude (A) of 1 N, period (P) of 0.1 m, and phase offset (θ) of 0 along both axes (X and Y) were the parameters used to implement the force field.

In the training period, the Assistive and Disruption groups were presented with an additional force vector, F_N, which acted to pull the device toward (Assistance) or push it away from (Disruption) the target position. The latter required participants to generate a compensatory force vector acting in the opposite direction to the positional error vector. This was the opposite of the Assistance condition.

Conceptually, error adjustment forces can be considered ‘noise’ if they are irrelevant to the underlying task (in our case, moving through the complex force field). For structural learning, it is critical that the CNS is able to dissociate such noise. Humans constantly operate with the presence of noise, and there is convincing evidence that humans are able to dissociate task-irrelevant noise from the environment (Todorov, 2004). The computations needed to solve this rule are relatively low dimensional and it was predicted that participants would be able to learn the compensatory force required to offset the bias relatively quickly, via MBL mechanisms.

5.3.1 Methods

5.3.1.1 Materials

The haptic assessment toolkit (HAT, described in Chapter 2) was used to build the virtual visuohaptic environment and experimental framework used for the task. The HapticMASTER was used to deliver forces and record the position and velocity of the device’s end-effector at a rate of approximately 1 kHz. The standard experimental setup described in Chapter 2 was used.

5.3.1.2 Participants

Thirty seven participants (11 female, aged: 19 to 61 years, M = 25.1, SD = 8.7) were randomly allocated to one of the three groups. All participants reported that they had normal or corrected vision and were right-handed. All participants received £15 as compensation for taking part in the study. No participants had previous experience of using the haptic device. The study was approved by The University of Leeds ethics committee and was performed in accordance with BPS ethical guidelines. All participants provided their informed consent prior to the start of the study.

5.3.1.3 Experimental design

The experiment was made up of three session types: Pre-test (always performed on a Monday), Training (always performed on a Tuesday, Wednesday and Thursday) and Post-test (Friday). Pre-test and Post-test were identical, allowing for the assessment of performance before and after training. A mixed design was used to measure any relative changes from Pre-test to Post-test, whilst a within-subjects design was used to assess the effects of Training over time.

Participants were required to attend one session per weekday for two consecutive weeks (i.e. they completed the Pre-test-Training-Post-test procedure twice). There were three groups: at training, the Assistance group were given error reduction forces that assisted them as they followed a moving target (but some active movement was still required). The Active-Control group did not receive any intervention forces. This condition used the same algorithm as that used in the Active-Control condition in Experiment 1. The Disruption group were given forces that acted to augment execution error, by pushing the hand away from the target.

5.3.1.4 Stimuli

Error adjustment forces

The Active-Control training condition was the same as the Active-Control condition in Experiment 1 and was generated using a stiffness, *k*, of 0 N/m and damping, *c*, of 1 Ns/m. Assistance was implemented using k = 100 N/m and c = 10 Ns/m. This acted to pull the device towards the target position for any non-zero error. Disruption was an error augmentation force generated using k = -100 N/m and c = 10 Ns/m. This acted to push the device away from the target position in the direction of the error (as error increased, so did the force magnitude).

Visual stimuli

Visual stimuli were set up in a similar way to those described in Experiment 1. The main difference was that paths were not displayed on-screen. This was done with the intention of increasing the participants' attention to the target cursor instead of trying to stay on the path as accurately as they could. The target cursor was a hollow red circle, and the current position cursor was a filled blue circle. A dotted black line was used to indicate the magnitude of the error between the current position and target cursors. With the aim of highlighting error, a dashed line was used to join the target and device cursors. This trajectory was based on 2D aiming tasks that have previously been used in the assessment of manual dexterity (Flatters et al., 2014).

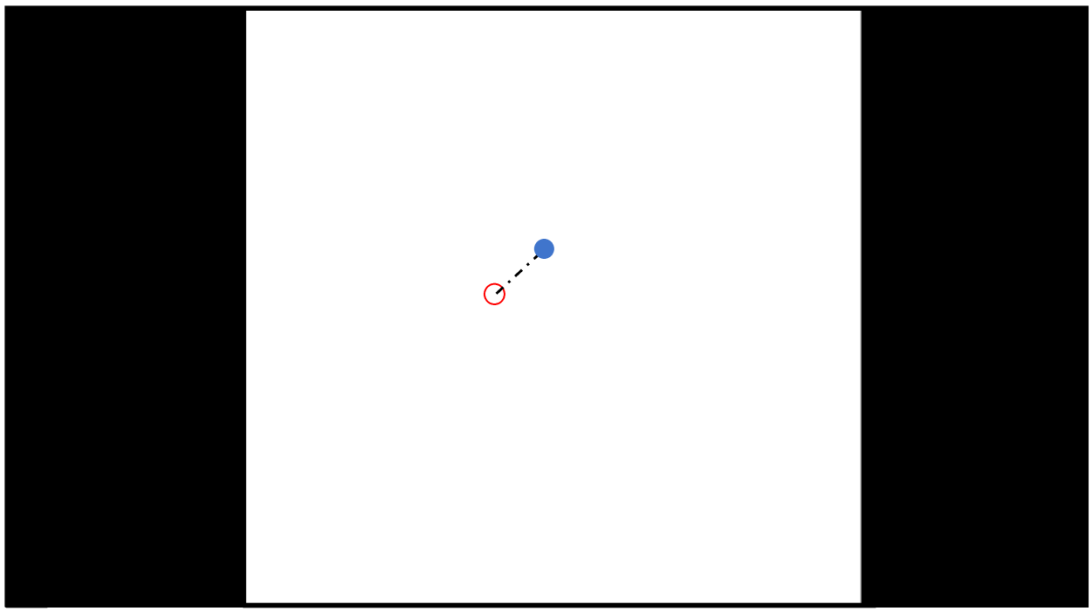

Figure 5.7. Visual display, showing the target (red) and device (blue) cursors. The dashed black line between the two cursors was designed with the aim of highlighting execution error.

An illustration of the paths used in the experiment (P1 and P2) is provided in Figure 5.8.

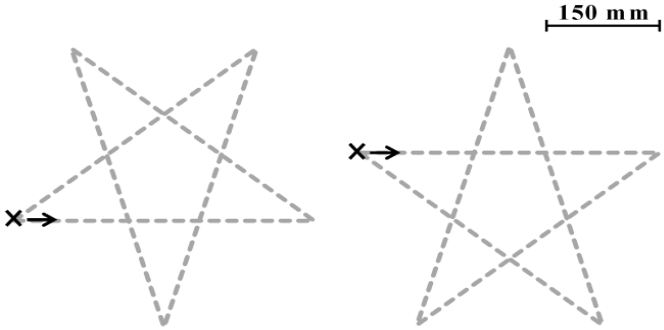

Figure 5.8. Illustration of the paths used for Pre-test and Post-test trials (left, path 1, 'P1'), and Training trials (right, path 2, 'P2'). The dotted lines indicate that the trajectories were not displayed (only the cursors indicated target and actual positions within the workspace). The crosses and arrows show the start locations and movement directions for each path.

Each component (straight line) measured 285 mm. A constant movement speed of 100 mm/s was used. In contrast to Experiment 1 (where movement speed was varied as a function of Path curvature), velocity was constant across the whole component movement. This was done with the aim of exacerbating the effects of the noise forces (i.e. the difference between the target and device cursors needed to be corrected at each change

of direction when moving from one component to another), and thus promote exploration of the parameter space.

5.3.1.5 Task and procedure

Participants were required to attend one session per week day for two weeks. Each session consisted of a number of 'blocks', where each block was a set of ten trials. Pre-test and Post-test consisted of three blocks (thirty trials) of Path 1 and had no error adjustment forces (i.e. the NG condition was used). Pre-test and Post-test were identical for all groups. Each Training session consisted of four blocks (forty trials) of Path 2.

Participants read an information sheet describing the experiment and were given general written instructions about how to complete the task. They were then able to ask the researcher any questions relating to the task. More specific instructions were displayed on-screen, showing the appearance of the target and device cursors, as well as visually indicating the requirements of the task. This was done to give more contextual information about the task. All participants performed the task with their dominant (right) hand.

After the on-screen instructions, participants were given an on-screen message to hold the end-effector of the device. Upon HAT detecting that the device was held, a message then appeared prompting them to move to the start position when ready and the target cursor appeared on-screen. After reaching the start position, the target cursor started moving immediately along the first component movement for that Path at a constant speed of 0.1 m/s. Once at the end of the component movement, the target cursor waited until the end of the component was reached by the device cursor before starting the next component. This process was repeated for each block of trials. To minimise fatigue, a compulsory 30-second rest was given after each block at which point participants were given the opportunity to sit down. The rest period could be extended indefinitely if desired by the participant. Each session lasted approximately fifteen minutes. The same process was repeated until all blocks for that session were complete.

A summary describing each stage of the experiment is shown in Table 5.1. Day 1 (a Monday) was Pre-test, which consisted of three blocks of ten trials of Path 1, under the Active-Control condition. Days 2-4 were Training, which consisted of four blocks of ten trials of Path 2, under the Assistance, Active-Control, or Disruption error adjustment

forces. Post-test was carried out on day 5 and was identical to Pre-test. Days 6-10 were identical to days 1-5.

Table 5.1. Summary of the protocol used for Experiment 2.

Day	Session ID	Group		
		ASS	ACC	DIS
1,6	Pre-test	3x10xP1,NG		
2,7	Training 1	4x10xP2, ASS	4x10xP2, ACC	4x10xP2, DIS
3,8	Training 2			
4,9	Training 3			
5,10	Post-test	3x10xP1,NG		

5.3.1.6 Data capture and analyses

Mean path error (E_P) was calculated for each component movement. Path error is defined as the absolute distance between the current position cursor and the closest point on the path (see Chapter 2). With the aim of smoothing the kinematic data, for each component, E_P scores that lied outside of 99% of the sample (i.e. those with a Z-score of 2.58) were removed from the analysis.

5.3.2 Results

A mixed ANOVA performed on the three pre-test Blocks revealed no significant effect of Time on E_P ($F(2,62) = .037$, $p = .963$, $\eta_p^2 = .001$), no Time*Group interaction ($F(4,62) = .103$, $p = .981$, $\eta_p^2 = .007$), and no significant between-group differences ($F(2,31) = 1.677$, $p = .203$, $\eta_p^2 = .098$). These findings indicate that all groups performed at the same level at baseline.

A plot of the normalised E_P (E_P at Post-test minus E_P at Pre-test) for each group is shown in Figure 5.9. A one-way ANOVA revealed a significant between-groups difference in normalised E_P scores ($F(2,27) = 6.565$, $p = .005$, $\eta_p^2 = .327$). Pairwise Tukey comparisons then revealed a significant difference between Assistance ($M = -0.253$, $SD = 0.738$) and Disruption ($M = -1.515$, $SD = 0.789$) (mean difference = 1. 263, SE = 0.364, $p = .005$) and between Active-Control ($M = -0.526$, $SD = 1.063$) and Disruption (mean difference = 0.989, SE = 0.396, $p = .048$), but no difference between Assistance and Active-Control (mean difference = 0.273, SE = 0.396, $p = .771$).

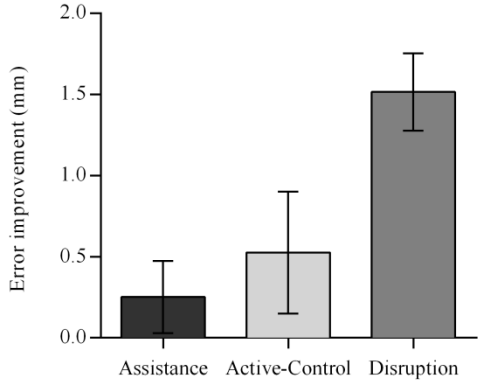

Figure 5.9. The Disruptive training group were able to generalise their learning better than those with Assistive and Active-Control. Error bars represent ± one standard error of the mean.

These results indicate that the group that received superimposed disruption forces showed the greatest improvement from Pre-test to Post-test after one week, and that assistive forces were no more or less effective than no error adjustment forces at all.

Next, the Training data were examined. A plot of E_P at each Training block is shown in Figure 5.10, also showing absolute mean performances at each pre- and post-test sessions. For the analysis, Training sessions were aggregated into blocks of four (four blocks were completed per session) to examine the effects of Training on learning. A Condition (3) x Time (6 blocks) ANOVA revealed a significant interaction (F (10, 130) = 4.829, $p < .001$, $\eta_p^2 = .271$). Decomposing the interaction revealed significant effects of block for Assistance ($F(5,45) = 10.318$, $p < .001$ $\eta_p^2 = .508$) and Active-Control ($F(5,35) = 13.304$, $p = .004$, $\eta_p^2 = .571$) and, importantly, the largest effect was found for the Disruption condition ($F(5,45) = 22.665$, $p < .001$, $\eta_p^2 = .716$) indicating a bigger difference in performance over time. These data indicate that, whilst all groups were able to improve during Training, that the Disruption group improved the most. These analyses were validated by inspection of the resulting learning curves, which confirmed faster learning rate for the Disruptive condition. This is in line with the prediction that error augmentation (noise) forces would quickly be adapted to, via a MBL process.

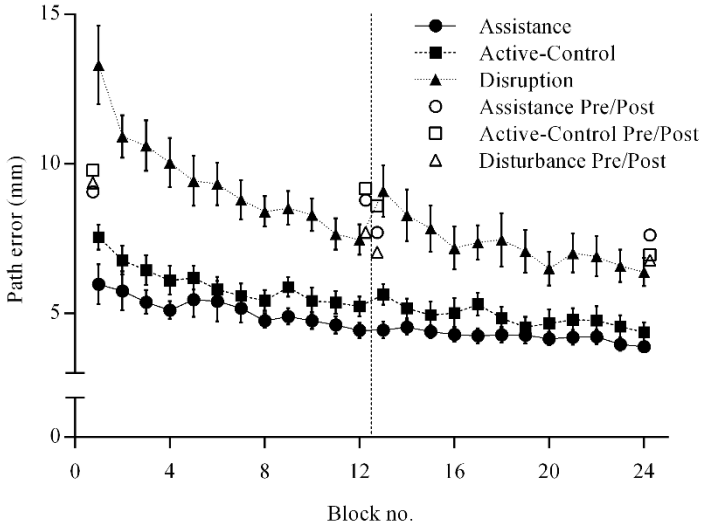

Figure 5.10. Plot of path error (E_P) during Training. The rapid initial adaptation (from Block 1 to Block 2) observed for the Disruptive group relative to the other groups is attributed to learning the Disturbance forces via predominantly MBL mechanisms. The vertical dashed line indicates the final block of the first week of testing (after which there was a two day break). The 'Pre/Post' values indicate the scores obtained for the final pre-test and first post-test sessions. Error bars represent ± one standard error of the mean.

The performance decrement after a two day break (from the last session of week 1 to the first session of week 2) was examined (see Figure 5.11). There was a significant effect ($F(2, 27) = 7.84$, $p = .0021$). The break resulted in only marginal differences for the Assistance and Active-Control conditions, but the Disruption group performed significantly worse relative to both (p's $< .03$).

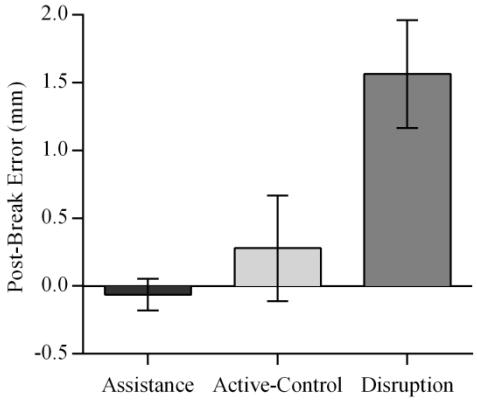

Figure 5.11. Difference in performance error before and after two day break. A lower score indicates better retention. Error bars represent ± one standard error of the mean.

To examine if these data were consistent with our a-priori prediction (that improved learning is driven by increased exploration of the parameter space), the amount of workspace exploration (as a proxy measure for parameter space exploration) was computed through analysis of path length (the total distance travelled in each condition – refer to Chapter 2; see Figure 5.12). ANOVA revealed a significant between-groups difference ($F(2,26) = 16.294$, $p < .001$, $\eta_p^2 = .556$). Pairwise Tukey comparisons then revealed significant differences between Assistance ($M = 18.09$, $SD = 0.756$) and Disturbance (M = 20.05, SD = 0.758) (mean difference = 1.959, SE = 0.343, $p < .001$), between Active-Control (M = 19.00, SD = 0.859) and Disturbance (mean difference = 1.042, SE = 0.373, $p = .025$) and difference between Assistance and Active-Control (mean difference = 0.917, SE = 0.365, $p = 047$). These findings indicate that disruptive forces may lead to greater workspace exploration.

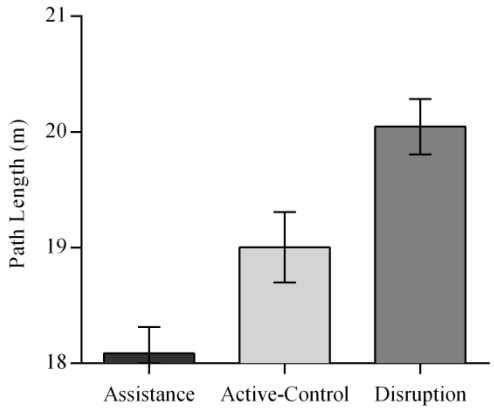

Figure 5.12. The average path length during training for each condition provides a measure of workspace exploration during the Training trials. Error bars represent ± one standard error of the mean.

5.3.3 Discussion

Experiment 2 tested the proposition that error augmentation forces can lead to faster motor learning. Subjects were required to learn task parameters in a complex novel force field (i.e. one where MBL solutions are constrained) with constant source of disruptive noise (error augmentation), no noise and assistive noise (error reduction). In line with our hypothesis, the Disruption training group performed better than the Assistance and Active-Control groups.

The two day break between week 1 and week 2 resulted in only marginal decrements for the Assistance and Active-Control conditions, but the Disruption group performed

significantly worse relative to both. An analysis of path length then revealed a pattern consistent with the a-priori hypothesis that there were significantly more movements around the workspace in the Disruption group relative to Assistance as well as Active-Control groups.

The results suggest that the system is able to distinguish error adjustment (noise) forces from the underlying task, likely via MBL mechanisms. An MFL strategy seems to be used to learn how to navigate the novel environment in absence of useful prior models. It seems also that the CNS is able to distinguish between noise and the underlying task. The data suggest that disruptive noise helps to increase the rate of sensorimotor learning, consistent with a number of previous studies (Cesqui, Aliboni, et al., 2008; F. Huang et al., 2007; Patton et al., 2013). However, it is unclear whether it is the error corrective process that directly drives learning.

One possibility that could explain this effect is that an increased amount of workspace exploration elicited by error augmentation forces could lead to more exploration of the parameter space and thus, better learning. The Disruption group exhibited the most amount of workspace exploration, suggesting that workspace exploration might help to increase learning. Whilst this experiment does not directly answer this question, there is a case to be made for learning via workspace exploration. Workspace exploration could act as a proxy to parameter space exploration, whereby moving within an unstable environment exposes the learner to task parameters which would not usually be experienced. In other words, increasing the dynamic instability of a task through provision of disruptive forces may improve learning by providing the system with more information about the underlying structure of the task (i.e. the parameter space). It is speculated that this process is likely to occur predominantly via model-free learning. To dissociate these two hypotheses and investigate which mechanism might offer the best learning rate, another experiment is required where the effect of workspace exploration is compared against a training algorithm that controls the amount of disruption given.

5.4 Experiment 3: Why does error augmentation facilitate sensorimotor learning?

The results from Experiment 2 show that disruption results in faster learning in a manner consistent with the hypothesis generated earlier from theories of structural learning and workspace exploration. Nonetheless, an alternative post-hoc account is that participants showed enhanced learning because the disruption created more errors and this improved the individual's ability to detect and correct deviations from planned trajectories. This account is less satisfactory because: (i) errors needed to be detected and corrected in all three conditions; (ii) it is not clear why feedback mechanisms would show better training with larger magnitude errors; (iii) it fails to explain how the system distinguishes the task-relevant errors related to the underlying complex force field from the task-irrelevant low dimensional noise (in contrast to an account based on structural learning theory).

Experiment 3 was constructed to test the idea that it is workspace exploration rather than error correction per se that is critical for motor learning. The experiment was also designed to pit three different algorithms against each other to determine the optimal manner of providing disruptive forces in haptic feedback devices. Exploring the impact of randomly applied assistive and disruptive forces was of particular interest as this seems to best mirror the intrinsic motor variability that predicts motor learning rates (H. G. Wu et al., 2014). The three algorithms were as follows. *Adaptive Algorithm* (AA) where the additional force varied as a function of task performance (i.e. increased disturbance when performance improved and increased assistance when performance declined). *Adaptive Disruptive* (AD) where a baseline level of performance was established before participants were exposed to ever increasing error augmentation as performance improved. *Random* (RAN) where an unpredictable force was provided (varying between high disturbance and high assistance) across trials.

5.4.1 Methods

5.4.1.1 Participants

Thirty-eight participants (24 female, aged 19 to 61 years, M = 25.1, SD = 8.7) were recruited and randomly allocated to one of three training groups. One participant withdrew voluntarily from the experiment. As with Experiment 1, all participants reported that they had normal or corrected vision, and that they were right-handed. All

participants received £15 as compensation for taking part in the study. No participants had experience of using the haptic device. The study was approved by The University of Leeds ethics committee and was performed in accordance with BPS ethical guidelines. All participants provided their informed consent prior to the commencement of the study.

5.4.1.2 Experimental design

Participants were required to attend one session per weekday (i.e. they completed the Pre-test, Training, Post-test procedure described for Experiment 2 once). The experiment consisted of one week of testing (Experiment 2 consisted of two weeks of testing). There were three groups: the Random (RAN) group were given error augmentation ('noise') forces that varied randomly between high assistance and high disturbance from one trial to the next. The Adaptive Algorithm (AA) group were given forces that varied as a function of performance: if they performed badly in trial *n*, more assistance was provided on trial *n+1*. Conversely, if they performed well, the next trial contained higher disturbance forces. The Adaptive-Disruptive (AD) group was identical to the Adaptive condition except that the bias force from one trial to the next could only either stay constant or become more disruptive (i.e. it never became more assistive, irrespective of performance).

5.4.1.3 Stimuli

Error adjustment ('noise') forces

Error adjustment forces varied between error reduction, with a maximum stiffness setting of 100 N/m (Assistance in Experiment 2), and error augmentation, with a minimum stiffness setting of -100 N/m (Disruption in Experiment 2). The noise force was updated at each trial. In the RAN condition, a random stiffness value within the operating envelope was assigned to each trial. In the AA condition, the force field magnitude changed as a function of performance in previous trials. The first trial of all conditions was always set to no guidance (0 N/m) in order to obtain a common benchmark measure of performance at the start of each session. For the AA group, the magnitude of the bias force at each trial was adjusted as a function of performance in previous trials, as described in equation (5.6). This algorithm has been used previously as a computational model of motor adaptation to predict the force required to minimise adaptation time to a viscous environment during treadmill walking tasks (Emken & Reinkensmeyer, 2005). In this case, the model was used to adjust the magnitude of a force field in the current trial as a

function of performance in previous trials. For the AD group the forces only became more disruptive (i.e. drops in performance were ignored).

$$k_{i+1} = f.k_i - g(x_i - x_d) \quad (5.6)$$

The stiffness, *k*, of the force field for the next trial is a function of the stiffness in the current trial, *i*, multiplied by a 'forgetting factor', *f*, and the difference between the demand error and actual error (x_d and x_n, respectively), multiplied by a gain value, *g*. The values of *f* and *g* dictate the relative sensitivity of the algorithm to previous performance (captured by k_i) and error. The sensitivity of the controller to performances obtained in previous trials is controlled by adjusting *f*: a larger forgetting factor will weight previous trials more heavily, whereas a smaller forgetting factor will result in more influence by the current trial's force magnitude. A value of 0.5 was used for both *f* and *g*, meaning that half of the weight was made of previous performance and the other half was made up of the current stiffness setting. This acted to give an equal balance between performance in previous trials, and that in the current trial.

For the purposes of this experiment, it was important to choose an error metric which was congruent with the instructions given to participants (i.e. to 'follow the target cursor as closely as possible'). Thus, the E_T metric was chosen as the set point variable that defined the magnitudes of *x* and x_d at each trial. To define an appropriate set point value of the error metric, some indication of the expected performance after adaptation was needed. Thus, group average performance data from experiment 2 were used. The average trajectory error (E_T) for all groups at the end of week 1 Post-test of Training was 9.1 mm. Thus, this value was used for the variable x_d. This meant that, in the AA condition, for a mean E_T greater than the set point for the current trial, the controller automatically decreased the stiffness coefficient of the error adjustment force to move closer towards the assistance realm in the next trial. Conversely, if the error was less than the set point, the force field automatically moved towards the Disturbance realm. For the AD condition, only the latter was true.

5.4.1.4 Task and procedure

A summary describing each stage of the experiment is shown in Table 5.2. Day 1 (a Monday) was Pre-test, which consisted of three blocks of ten trials of Path 1, under the Active-Control condition. Days 2-4 were Training, which consisted of four blocks of ten

trials of Path 2, under the AA, AD and RAN error adjustment forces. Post-test was carried out on day 5 and was identical to Pre-test. Days 6-10 were identical to days 1-5.

Table 5.2. Summary of the protocol used for Experiment 2.

Day	Session ID	Group		
		AA	AD	RAN
1	Pre1	3x10xP1,NG		
2	Tra1	4x10xP2, AA	4x10xP2, AD	4x10xP2, RAN
3	Tra2			
4	Tra3			
5	Pos1	3x10xP1,NG		

5.4.2 Results

Figure 5.13 shows a representative plot of the degree of assistance or disruption received by participants in all groups, along with a visual illustration of workspace exploration. For the AA and AD groups, the magnitude of the error adjustment forces varied as a function of trial-by-trial performance, whilst the Random group received a random magnitude (within the operating envelope). Importantly, the AD condition exposed the participants to more average disruption but less workspace exploration relative to the RAN condition.

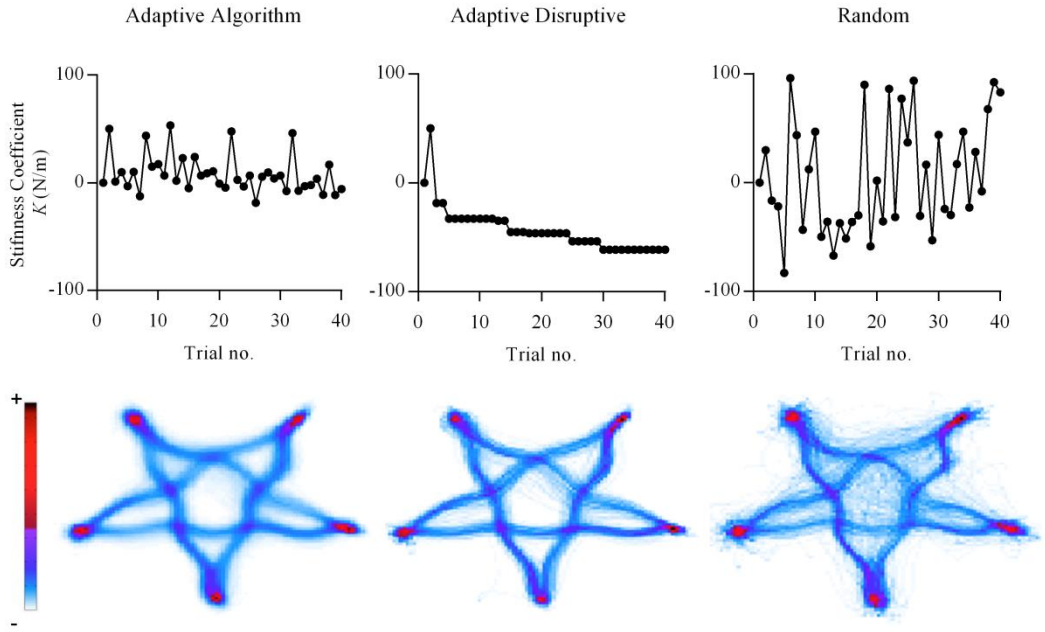

Figure 5.13. Top: The stiffness coefficient *K* (N/m) demonstrates the degree of assistance (positive values) and disruption (negative values) on a movement-by-movement basis for example subjects in the Adaptive Algorithm (AA), Adaptive Disruptive (AD) and Random (RAN) conditions. Bottom: Heat maps of movements during all training sessions for a single participant in each training group.

Next, the Training data were examined. Figure 5.14 shows the rate of error reduction for the three conditions during training. For statistical analysis, the 12 sessions were parsed into three blocks (mean average of four trials per block) and a 3 (Block) x 3 (Condition) ANOVA was conducted to examine the effects of performance over time. There was a marginally significant Block x Condition interaction ($F(4, 68) = 2.687$, $p = .054$, $\eta_p^2 = .136$). Decomposing the interaction for condition, revealed no difference in performance over time for AA ($F(2, 28) = .679$, $p = .515$, $\eta_p^2 = .046$) - this pattern was expected as task difficulty was intrinsically linked to task performance. There was a significant improvement in performance for the AA group ($F(2, 22) = 5.64$, $p = .011$, $\eta_p^2 = .339$) and the random group showed the best improvement over time ($F(2, 18) = 19.70$, $p < .001$, $\eta_p^2 = .686$).

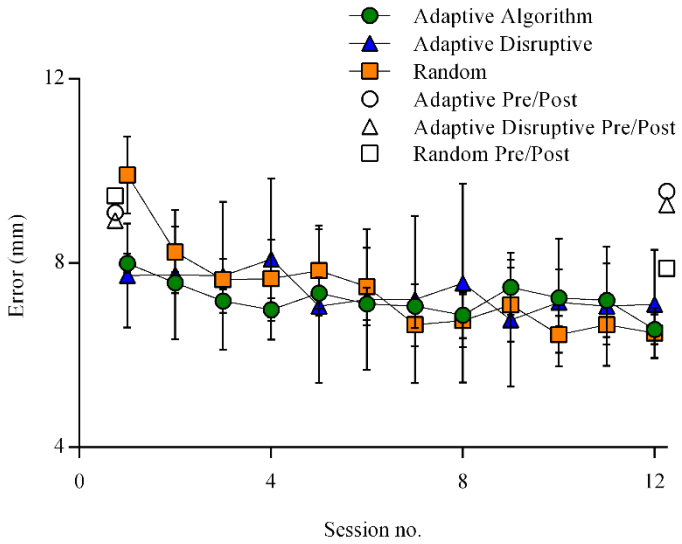

Figure 5.14. Rate of error reduction across training for all groups. The 'Pre/Post' values indicate the scores obtained for the final pre-test and first post-test sessions. Error bars represent ± one standard error of the mean.

Figure 5.15 shows the normalised path error (E_P) scores obtained for all groups. A one-way ANOVA revealed that there was a significant difference in error reduction between groups ($F(2, 34) = 3.87$, $p = .03$, $\eta_p^2 = .186$). Posthoc Tukey's comparisons showed learning in the Random condition was significantly better than AD ($p = .043$) and marginally better than AA ($p = .054$). There was no difference between AA and AD ($p = .97$).

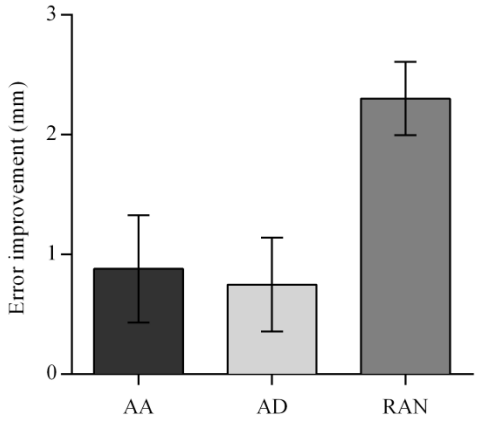

Figure 5.15. Random forces demonstrated better learning, as indexed by the amount of error reduction post-test relative to pre-test. Error bars represent ± one standard error of the mean.

5.4.3 Discussion

Experiment 3 investigated whether directly manipulating exposure to the physical workspace would result in faster learning, on the assumption that workspace exploration acts as a proxy to exploration of the parameter space. The Random group was exposed to superimposed noise forces that randomly varied between a continuum of assistance and disturbance, whilst the Adaptive Algorithm group received forces that varied as a function of their performance (on the same continuum). The Adaptive Disruptive condition was different to the AA condition in that the forces only became more disruptive.

It was found that training participants on Random forces led to better learning than algorithms that tweaked the error adjustment forces according to performance. On the assumption that workspace exploration is a proxy for parameter space exploration, it appears that actively experiencing more of the dynamics of the underlying environment (i.e. the one with the 'novel' force field) seems to allow the learner to more quickly build an internal model to skilfully navigate the novel environment, regardless of the amount of disruption provided.

5.5 General discussion

The results from the three experiments reported in this chapter support the hypothesis that the imposition of a low dimensional force can accelerate motor learning via increased physical workspace exploration. In Experiment 1, error reduction impaired MFL and thus hindered learning. In Experiment 2, the provision of disruptive forces improved learning. In Experiment 3, it was possible to accelerate MBL by increasing exposure to the workspace through delivery of a random selection of assistive and disruptive forces. These results are predicted by 'structural learning' theories that suggest that increased sampling can improve the posterior probability estimates required to learn the underlying structure of novel tasks (in this case, moving though a complex novel force field).

These findings are consistent with a number of previous results suggesting that disruptive forces might be beneficial for motor learning (Cesqui, Aliboni, et al., 2008; Emken & Reinkensmeyer, 2005; Huang & Shadmehr, 2007; Lee & Choi, 2010). The current work advances these reports by providing and testing a theoretical account of why disruption might accelerate learning. Moreover, evidence that disruption allows for generalisation beyond transient movement after-effects, rather than simple performance facilitation, is

provided (Reinkensmeyer & Patton, 2009; Reisman, Wityk, Silver, & Bastian, 2007). This work thus complements but advances previous observations about the potential benefits of disruption. For example, a previous study showed that performance on a tracking task could be improved through delivery of haptic disturbance (Lee, 2010). This finding could be explained, however, by an enhanced ability to deploy feedback control and, indeed, the authors of the study explained their results in terms of a general improvement in the 'attentional' capabilities of their participants. These mechanisms cannot explain the present experimental results where factors related to cognitive function (such as task switching; Mushtaq, Bland, & Schaefer, 2011) were controlled across the Adaptive Disruptive and Random conditions in Experiment 3. Our results also complement work showing that increased intrinsic variability predicts motor learning rates (H. G. Wu et al., 2014). The current findings demonstrate that extrinsic variability delivered through haptic disturbance can also augment learning (by accelerating the acquisition of MBL). Importantly, the system was able to adapt to the imposition of the 'low dimensional' disruptive force so that the net long-term learning outcome was beneficial – providing support for the notion that MBL mechanisms can identify and rapidly compensate to such perturbations (as predicted by structural learning theories). The general notion that increased workspace exploration can lead to faster learning is well explained by theories of structural learning and has good support from a range of empirical studies (Braun, Mehring, & Wolpert, 2010), including investigations of laparoscopic surgical training (White et al. 2014).

These findings raise the issue of the neural substrates underpinning the learning process. Previous work has indicated that the cerebellum is most likely to be responsible for the maintenance of models about task parameters (Haruno, Wolpert, & Kawato, 2001; Miall, Weir, Wolpert, & Stein, 1993; Paulin, 1993; Wolpert et al., 1998), with damage to this structure resulting in impairments in adaptation across a number of tasks (Bastian, 2011; Taylor, Klemfuss, & Ivry, 2010). One putative mechanism for MFL processing is likely to reside in the motor cortex, with systems responsible for dopaminergic neural firing in the primary motor cortex regulating trial and error learning (Hosp, Pekanovic, Rioult-Pedotti, & Luft, 2011; Huntley, Morrison, Prikhozhan, & Sealfon, 1992; Luft & Schwarz, 2009; Ziemann, Tergau, Bruns, Baudewig, & Paulus, 1997). Nevertheless, the neural processes that implement the computational algorithms exploited by the human nervous

system remain to be discovered (Franklin & Wolpert, 2011). Likewise, the underlying control mechanisms supporting skilled arm movements are poorly understood.

As discussed in Chapter 1, there are three simultaneous and interacting control mechanisms which could be present during motoric behaviour: feedback, impedance and predictive control. Perceived errors that arise during execution of a task are used to make adjustments (feedback control). By stiffening the limb, the destabilising (i.e. unpredictable) effects of a complex environment can be attenuated using impedance control. Lastly, once an inverse model is generated and 'fine-tuned', predictive control may become more dominant, resulting in a lesser reliance on feedback and impedance control and therefore improving accuracy and/or reducing effort. In these experiments it is not possible to determine how the individuals have learned to compensate for the complex force field, although the learning is likely to involve processes related to all three control mechanisms (Franklin & Wolpert, 2011). Speculatively, the high complexity of the tasks in experiments 2 and 3 are likely to require feedback and impedance control as the main contributors to the overall control strategy during the initial learning stages (i.e. before a suitable inverse model has been generated to implement predictive commands).

It is important to note that this research used neurologically intact adults as participants and whilst the force field in Experiments 2 and 3 allowed examining novel skill learning, the difficulty was tuned to a level such that all subjects could complete the task (moving between the points). It is speculated that disrupting the training of individuals with neurological deficits (e.g. cerebral palsy) might not be beneficial, and indeed that constraining errors in these populations could speed up the development of 'low-level controllers' through trial and error. Consistent with this, work with stroke survivors has shown that error amplification is useful in rehabilitation for mild impairment but error guidance is necessary for patients with more severe damage (Cesqui, Macri, et al., 2008). Likewise, haptic guidance has been found to be beneficial for people with relatively low skill levels, but error enhancement may be better for highly skilled individuals (Milot et al., 2010). The current work builds on these observations and provides a theoretical framework and empirical support for the development of optimised robotic training devices in sensorimotor skill training and rehabilitation.

CHAPTER 6

DISCUSSION AND CONCLUSIONS

6.1 Introduction

Humans display a remarkable ability to learn from and adapt to their environment. Typically developed adults have a substantial repertoire of experiences that can be utilised to make very fast predictions about the consequences of their actions when faced with subsequent similar challenges (Wolpert et al., 2011). Skilled behaviour is crucial in various clinical disciplines such as surgery and dentistry (Hamdorf & Hall, 2000). However, time restrictions and other limiting factors have recently created pressures within clinical training, and created interest in the types of training interventions that can be used to develop the sensorimotor skills of trainees within a short time period (Parsons et al., 2011). This has highlighted the need to develop novel training methodologies that can help to improve the acquisition rate of clinical sensorimotor skills.

Virtual reality (VR) training systems are becoming increasingly popular in medical training as they lever recent technological advances to produce realistic computer generated environments. However, the focus so far has been on developing 'high-fidelity' simulations of the environment which look convincing but don't necessarily provide the same perceptual cues that guide action in real world learning environments. For example, laparoscopic surgery simulators have focused on replicating the look and feel that the surgeon encounters but have paid little attention to the sensorimotor system that must interact with the simulator. Building on previous research that has investigated the role

of haptic feedback in sensorimotor learning (Patton et al., 2013; Reinkensmeyer & Patton, 2009), this thesis examines novel training paradigms that could act to improve the perceptive and active abilities of trainees through paradigms that complement existing clinical VR training systems. Chapter 2 presented two novel experimental tools that were used to carry out the experiments reported in this thesis (and which can support other researchers investigating similar questions). The experimental work was split into two general themes: the first theme (covered in chapters 3 and 4) was an investigation of compliance discrimination skills within an active-perceptive framework. The second theme (chapter 5) investigated the role of active behaviour, error augmentation forces and parameter space exploration on motor skill acquisition rate.

6.2 Review of experimental investigations

Detailed discussions have been given in each chapter. Below is an outline of the main findings.

Chapter 2 outlined the general tools developed and methodologies used throughout the experiments described in the thesis. The compliance simulation interface (CSI) was developed specifically to address the need for a method of objectively investigating (and training) the human ability to assess the compliance of virtual objects of homogeneous and heterogeneous force response distributions across their surface. This allowed for the implementation of robust experimental protocols in Chapters 3 and 4 that allowed the high-fidelity haptic simulation of compliant objects and the acquisition of kinematic data. The haptic assessment toolkit (HAT) provides a development framework which greatly simplifies the process of implementing bespoke visuohaptic environments and novel force fields. Although the current functionality of HAT is restricted to the specific requirements of the experiments discussed in Chapter 5, it is an open architecture which allows further development and scalability for increased functionality. Integration with other haptic devices and the incorporation of additional visual and haptic stimuli could help lever the full potential of HAT as a novel research tool that could prove useful to the research community as a common platform for extending this research.

In Chapter 3 the ability to assess compliance differences (the just noticeable difference, JND) with a handheld tool, and the effect of training on this skill, was assessed. This was done within the theoretical framework of ‘action-perception’, where perception’s primary

goal is to support action and where humans obtain perceptual information through active interactions with the world. The perception of compliance did, indeed, seem strongly dependent on the probing strategy employed. However, despite significant improvements in sensitivity to compliance cues, intensive training did not result in a systematic adjustment to probing strategy (over the investigated time frame). A potential application of this work could be to use a similar task to provide intensive training to trainee dentists, with the aim of increasing their compliance sensitivity thresholds when assessing tooth structures.

The aim of Chapter 4 was to investigate training strategies that can help to improve the rate of compliance discrimination skills on a simulated real world task: the detection of embedded tumours in human tissue using a handheld probe. Two different training conditions were designed to 1) increase sensitivity to compliance (this was the same as the JND training condition in the previous chapter), and 2) inform an effective probing strategy for compliance discrimination. Upon comparing to a control, the effectiveness of each training strategy was unclear due to unexpected between-group differences at baseline. In retrospect, the experiment could have been improved by adding a group which combined point (1) and (2), thus acting to train both the active and perceptive elements of the task: this way, participants would perhaps earn the benefits of both increased sensitivity to compliance, and an increase to the quality and quantity of information available to the CNS via the use of an effective probing strategy. Whilst the results from this study were inconclusive, further work on this topic could prove useful for informing novel environments for training the critical skill of compliance discrimination in a variety of medical applications.

Chapter 5 consisted of three experiments investigating the role of active behaviour, error augmentation forces, and physical workspace exploration. The broad framework of two distinct learning mechanisms was reviewed to inform the research: model-based (MBL) and model-free (MFL) learning. MBL is a fast process which adapts internal models into a ‘ball-park’ configuration of a new task. MFL is a slower process involving trial and error to sample the environment and gradually refine action. In line with the existing literature, Experiment 1 showed that the error corrective process is critical for any learning to take place. Indeed, being fully guided around a path by a device seems no more effective than simply visually observing the device as it autonomously follows the path. Experiment 2 used a novel force field to induce task novelty, allowing for an

observation of MBL and MFL processes in action when different error adjustment forces were superimposed (assistive, no intervention, and disruptive forces). It seems that humans are able to quickly (via an MBL process) dissociate noise from an underlying task, at which point MFL processes allow for gradual improvements of the underlying (novel) task. These findings have potential implications in areas such as laparoscopic surgery for increasing the learning rate of trainee laparoscopic surgeons, for example, where forces are applied about the virtual tool tip during manipulation of the laparoscopic instruments.

6.3 Overall discussion

Experienced surgeons and dentists possess a comprehensive set of knowledge-based, procedural, and sensorimotor (technical) abilities to accurately and efficiently interact with the environments that they operate in. Of these, technical skills often take the longest to learn, an issue which has presented significant challenges in modern surgical training. This thesis set out to investigate novel, virtual reality and haptics-based training methodologies for accelerating the learning rate of sensorimotor skills that are needed in clinical settings, with a particular focus on surgery and dentistry. The overall aim was to further the current understanding of the underpinning mechanisms involved in sensorimotor learning, and of how these can be exploited to maximise sensorimotor skill acquisition.

Sensorimotor skills involve perceptive (sensory) and active (motoric) abilities. These two modalities are intimately related during the execution of a task. Some sensorimotor skills that are pertinent to surgery and dentistry include using a handheld probe to proficiently navigate an environment to indent, manipulate, cut and suture human tissues. One example is the assessment of tissue health to detect physiological anomalies such as tumours or cavities.

The findings of this thesis indicate that active behaviour has a significant impact on perceptual abilities, whereby an effective behavioural strategy can inform the quality and quantity of information available to the CNS (e.g. moving at a suitable speed increases the ability to detect compliance differences – see Chapters 3 and 4). Within this dualistic framework, information obtained from the environment can also be used to inform future actions (i.e. which area of tissue to probe next, and at what speed). Optimising the

learner's exposure to the environmental parameters could be one way of maximising learning, whereby the learner is able to sample information to inform an effective control strategy. One way of achieving this could be to control exposure to the physical workspace (see Chapter 5).

This work opens up number up a number of exciting new research themes. One obvious example could be to further investigate the effectiveness of the techniques discussed here on modern virtual reality training systems. Another theme could be the assessment of the relative contributions of different control strategies during learning of a novel skill, which could further indicate the suitability of different training techniques for different applications. These are discussed in more detail below.

6.4 Future work

There are a number of limitations in this thesis that could be addressed in future work. These are described briefly below.

1. The tasks employed in Chapters 3 and 4 were relatively time consuming, whereby participants were required to attend more than one session per day. This limited the sample numbers that were possible for this work. Larger sample numbers would help validate the findings reported in Chapter 3, and also clarify the findings of Chapter 4. As discussed in the related chapter, an additional group that receives both haptic and kinematic training could provide valuable further insight into what training interventions are most effective for the provision of compliance discrimination skills.
2. The simulated compliance ranges employed in Chapters 3 and 4 were representative of the compliance of human liver. However, further validation is needed to qualify the effectiveness of different training methodologies for other environments. For instance, compliance discrimination on tooth structures would require larger forces and thus a haptic device capable of a larger force output would be needed. Fortunately, the compliance simulation interface (CSI, described in Chapter 2) is compatible with all PHANTOM devices, which makes it easy to use other haptic devices (i.e. with a larger workspace and/or force output, for instance) without the need to make any changes to the experimental protocols discussed here.

3. One limitation of the work discussed in Chapter 5 in Experiments 2 and 3 was that power output of participants was not measured as they actively controlled the haptic device to, effectively, move under a variety of force conditions which acted to disturb their movements. An analysis of arm impedance, for instance, could provide a powerful insight into the way that humans attenuate the effects of disturbing forces at various stages of training (Burdet et al., 2000).
4. Further development of the HAT could potentially create a powerful research tool which is common across research themes in this research area. This would allow researchers to configure, run and analyse experiments involving visuohaptic environments and haptic intervention strategies.
5. All of the studies presented here involved single arm movements. However, most clinical procedures involve bimanual actions, which may have implications on the control mechanisms adopted by the nervous system (Diedrichsen, 2007, p. -). Thus, further work will be needed to assess the best way of using the methodologies presented here to train bimanual sensorimotor skills.
6. A factor that was not investigated in any of the studies presented here is that of long-term skill retention (Schaverien, 2010). Whilst this was deemed to be outside the scope of this work, in moving forward it will be important to consider the impact of training interventions such as the ones discussed here on the retention of sensorimotor skills. This is critical for informing the best way of delivering haptic enhanced interventions: in the case that long-term skill retention is not possible, then 'warm-up' training (i.e. training immediately prior to performing a procedure) may be more effective (Ajemian, D'Ausilio, Moorman, & Bizzi, 2010).
7. The techniques discussed in this thesis were carried out under controlled laboratory conditions. Moving forward, it will be critical to further assess the effectiveness of such techniques in real applications.

6.5 Concluding remarks

The work presented in this thesis supports the notion that haptic feedback delivered through robotic systems can benefit sensorimotor skills training. Crucially, the techniques discussed here are highly amenable to state-of-the-art clinical VR training technologies

(i.e. many modern systems come with integrated visual and haptic feedback). For instance, the compliance just noticeable difference (JND) task employed in chapters 3 and 4 could be implemented on the Simodont (MOOG) dental trainer for trainees to undertake intensive training with the aim of increasing their ability to detect decaying teeth. The disturbance (error augmentation) algorithms used in Chapter 5 could be integrated into existing LapMentor (Simbionix) procedures to help train complex movements. However, in order to successfully transfer this research into applied clinical settings, it will be necessary to further assess the content validity of the training interventions discussed. This should be addressed in future work.

Whilst the main theme presented in this thesis has been clinical training with an emphasis on laparoscopic surgery, the research questions that have been explored are generally relevant within a number of applications. This thesis lays the groundwork for a number of potential further research themes. Physical rehabilitation, for example, has been widely discussed in the literature as one area which could benefit from haptic interventions for the re-learning of skills after neurological injury. Another potentially fruitful, but less explored area is sports training, where skilful movements are critical for success.

References

Aggarwal, R., Hance, J., & Darzi, A. (2004). Surgical education and training in the new millennium. *Surgical Endoscopy And Other Interventional Techniques*, *18*(10), 1409–1410.

Ahmed, N., Devitt, K. S., Keshet, I., Spicer, J., Imrie, K., Feldman, L., … Rutka, J. (2014). A Systematic Review of the Effects of Resident Duty Hour Restrictions in Surgery. *Annals of Surgery*, *259*(6), 1041–1053.

Ajemian, R., D'Ausilio, A., Moorman, H., & Bizzi, E. (2010). Why Professional Athletes Need a Prolonged Period of Warm-Up and Other Peculiarities of Human Motor Learning. *Journal of Motor Behavior*, *42*(6), 381–388. http://doi.org/10.1080/00222895.2010.528262

Andersen, R. A., & Cui, H. (2009). Intention, Action Planning, and Decision Making in Parietal-Frontal Circuits. *Neuron*, *63*(5), 568–583.

Bakr, M. M., Massey, W., & Alexander, H. (2013). Evaluation of Simodont® Haptic 3D virtual reality dental training simulator. *International Journal of Dental Clinics*, *5*(4). Retrieved from http://intjdc.org/index.php/intjdc/article/view/824

Bann, S. D., Datta, V. K., Khan, M. S., Ridgway, P. F., & Darzi, A. W. (2005). Attitudes towards skills examinations for basic surgical trainees. *International Journal of Clinical Practice*, *59*(1), 107–113.

Barnes, R. W. (1987). Surgical handicraft: Teaching and learning surgical skills. *The American Journal of Surgery*, *153*(5), 422–427.

Basteris, A., & Sanguineti, V. (2011). Toward "optimal" schemes of robot assistance to facilitate motor skill learning. In *2011 Annual International Conference of the*

IEEE Engineering in Medicine and Biology Society,EMBC (pp. 2355 –2358). http://doi.org/10.1109/IEMBS.2011.6090658

Bastian, A. J. (2011). Moving, sensing and learning with cerebellar damage. *Current Opinion in Neurobiology*, *21*(4), 596–601.

Botden, S. M. B. I., & Jakimowicz, J. J. (2009). What is going on in augmented reality simulation in laparoscopic surgery? *Surgical Endoscopy*, *23*(8), 1693–1700. http://doi.org/10.1007/s00464-008-0144-1

Braun, D. A., Aertsen, A., Wolpert, D. M., & Mehring, C. (2009). Motor Task Variation Induces Structural Learning. *Current Biology*, *19*(4), 352–357. http://doi.org/10.1016/j.cub.2009.01.036

Braun, D. A., Mehring, C., & Wolpert, D. M. (2010). Structure learning in action. *Behavioural Brain Research*, *206*(2), 157–165. http://doi.org/10.1016/j.bbr.2009.08.031

Braun, D. a., Mehring, C., & Wolpert, D. M. (2010). Structure learning in action. *Behavioural Brain Research*, *206*(2), 157–165. http://doi.org/10.1016/j.bbr.2009.08.031

Braun, D. A., Mehring, C., & Wolpert, D. M. (2010). Structure learning in action. *Behavioural Brain Research*, *206*(2), 157–65. http://doi.org/10.1016/j.bbr.2009.08.031

Brooks,Jr., F. P., Ouh-Young, M., Batter, J. J., & Jerome Kilpatrick, P. (1990). Project GROPEHaptic displays for scientific visualization. In *Proceedings of the 17th annual conference on Computer graphics and interactive techniques* (pp. 177–185). New York, NY, USA: ACM. http://doi.org/10.1145/97879.97899

Brouwer, I., Mora, V., & Laroche, D. (2007). A viscoelastic soft tissue model for haptic surgical simulation. In *EuroHaptics Conference, 2007 and Symposium on Haptic*

Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2007. Second Joint (pp. 593–594). http://doi.org/10.1109/WHC.2007.12

Bruner, J. S. (1973). Organization of early skilled action. *Child Development*, *44*(1), 1–11. http://doi.org/10.2307/1127671

Burdet, E., Osu, R., Franklin, D. W., Yoshioka, T., Milner, T. E., & Kawato, M. (2000). A method for measuring endpoint stiffness during multi-joint arm movements. *Journal of Biomechanics*, *33*(12), 1705–1709. http://doi.org/10.1016/S0021-9290(00)00142-1

Bush, R. R., & Mosteller, F. (1953). A Stochastic Model with Applications to Learning. *The Annals of Mathematical Statistics*, *24*(4), 559–585. http://doi.org/10.1214/aoms/1177728914

Catchpole, K., Panesar, S. S., Russell, J., Tang, V., Hibbert, P., & Cleary, K. (2009). Surgical Safety can be improved through better understanding of incidents reported to a national database. *National Reporting and Learning Service. London: National Patient Safety Agency*.

Cesqui, B., Aliboni, S., Mazzoleni, S., Carrozza, M. C., Posteraro, F., & Micera, S. (2008). On the use of divergent force fields in robot-mediated neurorehabilitation. In *2nd IEEE RAS EMBS International Conference on Biomedical Robotics and Biomechatronics, 2008. BioRob 2008* (pp. 854–861). http://doi.org/10.1109/BIOROB.2008.4762927

Cesqui, B., Macri, G., Dario, P., & Micera, S. (2008). Characterization of age-related modifications of upper limb motor control strategies in a new dynamic environment. *Journal of NeuroEngineering and Rehabilitation*, *5*, 31.

Chandler, J., Dickson, M., Jamieson, E., Mueller, T., Reid, T. (Unpublished). Computational models and haptic tools for next generation surgery.

Chen, X., & Agrawal, S. K. (2012). Assisting versus repelling force-feedback for human learning of a line following task. In *2012 4th IEEE RAS EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)* (pp. 344–349). http://doi.org/10.1109/BioRob.2012.6290678

Chen, X., & Agrawal, S. K. (2013). Assisting Versus Repelling Force-Feedback for Learning of a Line Following Task in a Wheelchair. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *21*(6), 959–968. http://doi.org/10.1109/TNSRE.2013.2245917

Choi, S., Walker, L., Tan, H. Z., Crittenden, S., & Reifenberger, R. (2005). Force constancy and its effect on haptic perception of virtual surfaces. *ACM Transactions on Applied Perception*, *2*(2), 89–105.

Conditt, M. A., Gandolfo, F., & Mussa-Ivaldi, F. A. (1997). The Motor System Does Not Learn the Dynamics of the Arm by Rote Memorization of Past Experience. *Journal of Neurophysiology*, *78*(1), 554–560.

Cotin, S., & Delingette, H. (1998). Real-time surgery simulation with haptic feedback using finite elements. In *1998 IEEE International Conference on Robotics and Automation, 1998. Proceedings* (Vol. 4, pp. 3739 –3744 vol.4). http://doi.org/10.1109/ROBOT.1998.681425

Courtine, G., Gerasimenko, Y., van den Brand, R., Yew, A., Musienko, P., Zhong, H., … Edgerton, V. R. (2009). Transformation of nonfunctional spinal circuits into functional states after the loss of brain input. *Nature Neuroscience*, *12*(10), 1333–1342. http://doi.org/10.1038/nn.2401

Crofts, T. J., Griffiths, J. M. T., Sharma, S., Wygrala, J., & Aitken, R. J. (1997). Surgical training: an objective assessment of recent changes for a single health board. *BMJ*, *314*(7084), 891. http://doi.org/10.1136/bmj.314.7084.891

Culmer, P. (2007). *Development of a Cooperative Robot System to Aid Stroke Rehabilitation*. University of Leeds. Retrieved from http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.486310

Culmer, P., Barrie, J., Hewson, R., Levesley, M., Mon-Williams, M., Jayne, D., & Neville, A. (2012). Reviewing the technological challenges associated with the development of a laparoscopic palpation device. *The International Journal of Medical Robotics and Computer Assisted Surgery*, *8*(2), 146–159.

Culmer, P., Levesley, M. C., Mon-Williams, M., & Williams, J. H. G. (2009). A new tool for assessing human movement: The Kinematic Assessment Tool. *Journal of Neuroscience Methods*, *184*(1), 184–192.

Cuschieri, A. (1995). Whither minimal access surgery: tribulations and expectations. *The American Journal of Surgery*, *169*(1), 9–19.

Darzi, A., & Mackay, S. (2002). Skills assessment of surgeons. *Surgery*, *131*(2), 121–124.

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci*, *8*(12), 1704–1711. http://doi.org/http://www.nature.com/neuro/journal/v8/n12/suppinfo/nn1560_S1.html

Dayan, P. (2009). Goal-directed control and its antipodes. *Neural Networks*, *22*(3), 213–219. http://doi.org/10.1016/j.neunet.2009.03.004

Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, *8*(4), 429–453. http://doi.org/10.3758/CABN.8.4.429

Dental Lab Home. (n.d.). Retrieved November 22, 2015, from http://www.dentsable.com

Derossis, A. M., Bothwell, J., Sigman, H. H., & Fried, G. M. (1998). The effect of practice on performance in a laparoscopic simulator. *Surgical Endoscopy*, *12*(9), 1117–1120. http://doi.org/10.1007/s004649900796

Desmurget, M., & Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Sciences*, *4*(11), 423–431. http://doi.org/10.1016/S1364-6613(00)01537-0

Diedrichsen, J. (2007). Optimal Task-Dependent Changes of Bimanual Feedback Control and Adaptation. *Current Biology*, *17*(19), 1675–1679.

Diedrichsen, J., White, O., Newman, D., & Lally, N. (2010). Use-Dependent and Error-Based Learning of Motor Behaviors. *The Journal of Neuroscience*, *30*(15), 5159–5166. http://doi.org/10.1523/JNEUROSCI.5406-09.2010

Drucker, S. D., Prieto, L. E., & Kao, D. W. K. (2012). Periodontal Probing Calibration in an Academic Setting. *Journal of Dental Education*, *76*(11), 1466–1473.

Emken, J. L., Benitez, R., Sideris, A., Bobrow, J. E., & Reinkensmeyer, D. J. (2007). Motor Adaptation as a Greedy Optimization of Error and Effort. *Journal of Neurophysiology*, *97*(6), 3997–4006. http://doi.org/10.1152/jn.01095.2006

Emken, J. L., & Reinkensmeyer, D. J. (2005). Robot-enhanced motor learning: accelerating internal model formation during locomotion by transient dynamic amplification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *13*(1), 33 –39. http://doi.org/10.1109/TNSRE.2004.843173

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.

Fermin, A., Yoshida, T., Ito, M., Yoshimoto, J., & Doya, K. (2010). Evidence for model-based action planning in a sequential finger movement task. *Journal of Motor Behavior*, *42*(6), 371–379. http://doi.org/10.1080/00222895.2010.526467

Fitts, P. M., & Posner, M. I. (1979). *Human performance*. Greenwood Press.

Flatters, I., Mushtaq, F., Hill, L. J. B., Holt, R. J., Wilkie, R. M., & Mon-Williams, M. (2014). The relationship between a child's postural stability and manual dexterity. *Experimental Brain Research*, *232*(9), 2907–2917.

Franklin, D. W., & Wolpert, D. M. (2011). Computational Mechanisms of Sensorimotor Control. *Neuron*, *72*(3), 425–442. http://doi.org/10.1016/j.neuron.2011.10.006

Franklin, D. W., & Wolpert, D. M. (2011). Computational Mechanisms of Sensorimotor Control. *Neuron*, *72*(3), 425–442. http://doi.org/DOI 10.1016/j.neuron.2011.10.006

Friedman, R. M., Hester, K. D., Green, B. G., & LaMotte, R. H. (2008). Magnitude estimation of softness. *Experimental Brain Research*, *191*(2), 133–142.

Geijtenbeek, T., Steenbrink, F., Otten, B., & Even-Zohar, O. (2011). D-flow: Immersive Virtual Reality and Real-time Feedback for Rehabilitation. In *Proceedings of the 10th International Conference on Virtual Reality Continuum and Its Applications in Industry* (pp. 201–208). New York, NY, USA: ACM. http://doi.org/10.1145/2087756.2087785

Gibson, E. (1988). Exploratory Behavior In The Development Of Perceiving, Acting, And The Acquiring Of Knowledge. *Annual Review of Psychology*, *39*(1), 1–41. http://doi.org/10.1146/annurev.psych.39.1.1

Gibson, J. J. (1986). *The ecological approach to visual perception*. Erlbaum, Hillsdale, NJ (Original work published in 1979).

Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, *66*(4), 585–595. http://doi.org/10.1016/j.neuron.2010.04.016

Gomes, P. (2011). Surgical robotics: Reviewing the past, analysing the present, imagining the future. *Robotics and Computer-Integrated Manufacturing*, *27*(2), 261–266. http://doi.org/10.1016/j.rcim.2010.06.009

Grantcharov, T. P., Bardram, L., Funch-Jensen, P., & Rosenberg, J. (2002). Assessment of Technical Surgical Skills. *The European Journal of Surgery*, *168*(3), 139–144. http://doi.org/10.3109/110241502320127739

Grantcharov, T. P., & Reznick, R. K. (2008). Teaching procedural skills. *BMJ : British Medical Journal*, *336*(7653), 1129–1131. http://doi.org/10.1136/bmj.39517.686956.47

Haith, A. M., & Krakauer, J. W. (2013). Model-Based and Model-Free Mechanisms of Human Motor Learning. In M. J. Richardson, M. A. Riley, & K. Shockley (Eds.), *Progress in Motor Control* (pp. 1–21). Springer New York. Retrieved from http://link.springer.com/chapter/10.1007/978-1-4614-5465-6_1

Hamdorf, J. M., & Hall, J. C. (2000). Acquiring surgical skills. *British Journal of Surgery*, *87*(1), 28–37. http://doi.org/10.1046/j.1365-2168.2000.01327.x

Hart, C. B., & Giszter, S. F. (2010). A Neural Basis for Motor Primitives in the Spinal Cord. *The Journal of Neuroscience*, *30*(4), 1322–1336. http://doi.org/10.1523/JNEUROSCI.5894-08.2010

Hart, R., & Karthigasu, K. (2007). The benefits of virtual reality simulator training for laparoscopic surgery: *Current Opinion in Obstetrics and Gynecology*, *19*(4), 297–302. http://doi.org/10.1097/GCO.0b013e328216f5b7

Haruno, M., Wolpert, D. M., & Kawato, M. (2001). MOSAIC model for sensorimotor learning and control. *Neural Computation*, *13*(10), 2201–2220. http://doi.org/Doi 10.1162/089976601750541778

Hesse, S., Schmidt, H., Werner, C., & Bardeleben, A. (2003). Upper and lower extremity robotic devices for rehabilitation and for studying motor control. *Current Opinion in Neurology*, *16*(6), 705–10. http://doi.org/10.1097/01.wco.0000102630.16692.38

Heuer, H., & Rapp, K. (2011). Active error corrections enhance adaptation to a visuo-motor rotation. *Experimental Brain Research*, *211*(1), 97–108. http://doi.org/10.1007/s00221-011-2656-5

Hogan, N. (1984). Impedance Control: An Approach to Manipulation. In *American Control Conference, 1984* (pp. 304–313).

Hosp, J. A., Pekanovic, A., Rioult-Pedotti, M. S., & Luft, A. R. (2011). Dopaminergic projections from midbrain to primary motor cortex mediate motor skill learning. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(7), 2481–7.

Huang, F., Patton, J. L., & Mussa-Ivaldi, F. (2007). Interactive priming enhanced by negative damping aids learning of an object manipulation task. In *29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2007. EMBS 2007* (pp. 4011 –4014). http://doi.org/10.1109/IEMBS.2007.4353213

Huang, Haith, A., Mazzoni, P., & Krakauer, J. W. (2011). Rethinking Motor Learning and Savings in Adaptation Paradigms: Model-Free Memory for Successful Actions Combines with Internal Models. *Neuron*, *70*(4), 787–801. http://doi.org/10.1016/j.neuron.2011.04.012

Huang, & Shadmehr, R. (2007). Evolution of Motor Memory During the Seconds After Observation of Motor Error. *Journal of Neurophysiology*, *97*(6), 3976–3985. http://doi.org/10.1152/jn.01281.2006

Huang, V. S., Haith, A., Mazzoni, P., & Krakauer, J. W. (2011). Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron*, *70*(4), 787–801.

Huntley, G. W., Morrison, J. H., Prikhozhan, A., & Sealfon, S. C. (1992). Localization of multiple dopamine receptor subtype mRNAs in human and monkey motor cortex and striatum. *Molecular Brain Research*, *15*(3-4), 181–188.

ISCP. (2013, October). The Intercollegiate Surgical Curriculum - Educating the surgeons of the future. Interceollegiate Surgical Curriculum Programme.

Jackson, A., Culmer, P., Makower, S., Levesley, M., Richardson, R., Cozens, A., … Bhakta, B. (2007). Initial patient testing of iPAM - a robotic system for Stroke rehabilitation. In *IEEE 10th International Conference on Rehabilitation Robotics, 2007. ICORR 2007* (pp. 250–256). http://doi.org/10.1109/ICORR.2007.4428435

JCST. (2012, June). ICSP Evaluation Report. Retrieved from http://www.jcst.org/ISCP%20Evaluation%20/full_evaluation_report

Judkins, T. N., Oleynikov, D., & Stergiou, N. (2008). Objective evaluation of expert and novice performance during robotic surgical training tasks. *Surgical Endoscopy*, *23*(3), 590–597. http://doi.org/10.1007/s00464-008-9933-9

Kaim, & Drewing, K. (2011). Exploratory strategies in haptic softness discrimination are tuned to achieve high levels of task performance. *Haptics, IEEE Transactions on*, *4*(4), 242–252.

Kaim, L. R., & Drewing, K. (2009). Finger force of exploratory movements is adapted to the compliance of deformable objects. In *World Haptics 2009.* (pp. 565–569). IEEE.

Karadogan, E., Williams, R. L., Howell, J. N., & Conatser, R. R. (2010). A Stiffness Discrimination Experiment Including Analysis of Palpation Forces and

Velocities: *Simulation in Healthcare: The Journal of the Society for Simulation in Healthcare, 5*(5), 279–288. http://doi.org/10.1097/SIH.0b013e3181e9e783

Kirk, R. M. (2002). *Basic Surgical Techniques* (5th ed.). Elsevier Health Sciences.

Kitagawa, M., Dokko, D., Okamura, A. M., & Yuh, D. D. (2005). Effect of sensory substitution on suture-manipulation forces for robotic surgical systems. *The Journal of Thoracic and Cardiovascular Surgery*, *129*(1), 151–158. http://doi.org/10.1016/j.jtcvs.2004.05.029

Klein, J., Spencer, S. J., & Reinkensmeyer, D. J. (2012). Breaking It Down Is Better: Haptic Decomposition of Complex Movements Aids in Robot-Assisted Motor Learning. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *20*(3), 268 –275. http://doi.org/10.1109/TNSRE.2012.2195202

K, L., S, G., S, T., Je, D., & M, C. (2012). Cochrane review: virtual reality for stroke rehabilitation. *European Journal of Physical and Rehabilitation Medicine*, *48*(3), 523–530.

Koçak, U., Palmerius, K., Forsell, C., Ynnerman, A., & Cooper, M. (2011). Analysis of the JND of Stiffness in Three Modes of Comparison. In E. Cooper, V. Kryssanov, H. Ogawa, & S. Brewster (Eds.), *Haptic and Audio Interaction Design* (Vol. 6851, pp. 22–31). Springer Berlin / Heidelberg.

Kopta, J. A. (1971). An approach to the evaluation of operative skills. *Surgery*, *70*(2), 297–303.

Krebs, H. I., Mernoff, S., Fasoli, S. E., Hughes, R., Stein, J., & Hogan, N. (2008). A comparison of functional and impairment-based robotic training in severe to moderate chronic stroke: a pilot study. *NeuroRehabilitation*, *23*(1), 81–87.

Kuschel, M., Di Luca, M., Buss, M., & Klatzky, R. L. (2010). Combination and Integration in the Perception of Visual-Haptic Compliance Information. *IEEE Transactions on Haptics*, *3*(4), 234–244.

LaMotte, R. H. (2000). Softness Discrimination With a Tool. *Journal of Neurophysiology*, *83*(4), 1777–1786.

Larsen, C. R., Soerensen, J. L., Grantcharov, T. P., Dalsgaard, T., Schouenborg, L., Ottosen, C., … Ottesen, B. S. (2009). Effect of virtual reality training on laparoscopic surgery: randomised controlled trial. *BMJ : British Medical Journal*, *338*. http://doi.org/10.1136/bmj.b1802

Latash, M. L., Scholz, J. P., & Schöner, G. (2002). Motor control strategies revealed in the structure of motor variability. *Exercise and Sport Sciences Reviews*, *30*(1), 26–31.

Lederman, S. J., & Klatzky, R. L. (1987). Hand movements: A window into haptic object recognition. *Cognitive Psychology*, *19*(3), 342–368.

Lederman, S. J., & Klatzky, R. L. (2004). Haptic identification of common objects: Effects of constraining the manual exploration process. *Perception & Psychophysics*, *66*(4), 618–628.

Lee, J. (2010). Effects of haptic guidance and disturbance on motor learning: Potential advantage of haptic disturbance. In *2010 IEEE Haptics Symposium* (pp. 335–342). IEEE. http://doi.org/10.1109/HAPTIC.2010.5444635

Lee, J., & Choi, S. (2010). Effects of haptic guidance and disturbance on motor learning: Potential advantage of haptic disturbance. *2010 IEEE Haptics Symposium, HAPTICS 2010*, 335–342. http://doi.org/10.1109/HAPTIC.2010.5444635

Leek, M. (2001). Adaptive procedures in psychophysical research. *Attention, Perception, & Psychophysics*, *63*(8), 1279–1292.

Lieberman, H., & Pentland, A. (1982). Microcomputer-based estimation of psychophysical thresholds: The Best PEST. *Behavior Research Methods*, *14*(1), 21–25.

Luft, A. R., & Schwarz, S. (2009). Dopaminergic signals in primary motor cortex. *International Journal of Developmental Neuroscience : The Official Journal of the International Society for Developmental Neuroscience*, *27*(5), 415–21.

Lüttgen, J., & Heuer, H. (2012a). Robotic guidance benefits the learning of dynamic, but not of spatial movement characteristics. *Experimental Brain Research*, *222*(1-2), 1–9. http://doi.org/10.1007/s00221-012-3190-9

Lüttgen, J., & Heuer, H. (2012b). The influence of haptic guidance on the production of spatio-temporal patterns. *Human Movement Science*, *31*(3), 519–528. http://doi.org/10.1016/j.humov.2011.07.002

Martin, J. A., Regehr, G., Reznick, R., Macrae, H., Murnaghan, J., Hutchison, C., & Brown, M. (1997). Objective structured assessment of technical skill (OSATS) for surgical residents. *British Journal of Surgery*, *84*(2), 273–278. http://doi.org/10.1046/j.1365-2168.1997.02502.x

McCaskie, A. W., Kenny, D. T., & Deshmukh, S. (2011). How can surgical training benefit from theories of skilled motor development, musical skill acquisition and performance psychology? The“ father of microsurgery” was a skilled pianist who developed the instruments for microsurgical operations using his“ pianist’s dexterity”. *Medical Journal of Australia*, *194*(9), 463.

McCloy, R., & Stone, R. (2001). Virtual reality in surgery. *BMJ*, *323*(7318), 912–915. http://doi.org/10.1136/bmj.323.7318.912

Melendez-Calderon, A., Masia, L., Gassert, R., Sandini, G., & Burdet, E. (2011). Force Field Adaptation Can Be Learned Using Vision in the Absence of Proprioceptive

Error. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *19*(3), 298 –306. http://doi.org/10.1109/TNSRE.2011.2125990

Miall, R. C., Weir, D. J., Wolpert, D. M., & Stein, J. F. (1993). Is the Cerebellum a Smith Predictor. *J Mot Behav*, *25*(3), 203–216.

Miall, R. C., & Wolpert, D. M. (1996). Forward Models for Physiological Motor Control. *Neural Networks*, *9*(8), 1265–1279. http://doi.org/10.1016/S0893-6080(96)00035-4

Milner, T. E., & Franklin, D. W. (2005). Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *The Journal of Physiology*, *567*(2), 651–664. http://doi.org/10.1113/jphysiol.2005.090449

Milot, M.-H., Marchal-Crespo, L., Green, C., Cramer, S., & Reinkensmeyer, D. (2010). Comparison of error-amplification and haptic-guidance training techniques for learning of a timing-based motor task by healthy individuals. *Experimental Brain Research*, *201*(2), 119–131. http://doi.org/10.1007/s00221-009-2014-z

Mon-Williams, M., Wann, J. P., Jenkinson, M., & Rushton, K. (1997). Synaesthesia in the normal limb. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *264*(1384), 1007–1010.

MOOG. (2011, January). HapticMASTER Programming Manual.

Moorthy, K., Munz, Y., Sarker, S. K., & Darzi, A. (2003). Objective assessment of technical skills in surgery. *BMJ : British Medical Journal*, *327*(7422), 1032–1037.

Mueller, S., & Sandrin, L. (2010). Liver stiffness: a novel parameter for the diagnosis of liver disease. *Hepatic Medicine: Evidence and Research*, *2*, 49–67.

Mushtaq, F., Bland, A. R., & Schaefer, A. (2011). Uncertainty and cognitive control. *Frontiers in Psychology*, *2*(October), 249. http://doi.org/10.3389/fpsyg.2011.00249

Mussa-Ivaldi, F. A., Giszter, S. F., & Bizzi, E. (1994). Linear combinations of primitives in vertebrate motor control. *Proceedings of the National Academy of Sciences, 91*(16), 7534–7538.

Nagasaki, H. (1989). Asymmetric velocity and acceleration profiles of human arm movements. *Experimental Brain Research, 74*(2), 319–326.

Oropesa, I., Sánchez-González, P., Lamata, P., Chmarra, M. K., Pagador, J. B., Sánchez-Margallo, J. A., … Gómez, E. J. (2011). Methods and Tools for Objective Assessment of Psychomotor Skills in Laparoscopic Surgery. *Journal of Surgical Research, 171*(1), e81–e95. http://doi.org/10.1016/j.jss.2011.06.034

Otaduy, M. A., & Lin, M. C. (2006). High Fidelity Haptic Rendering. *Synthesis Lectures on Computer Graphics*, *1*(1), 1–112. http://doi.org/10.2200/S00045ED1V01Y200609CGR002

Paisley, A. M., Baldwin, P. J., & Paterson-Brown, S. (2001). Validity of surgical simulation for the assessment of operative skill. *British Journal of Surgery, 88*(11), 1525–1532. http://doi.org/10.1046/j.0007-1323.2001.01880.x

Parsons, B. A., Blencowe, N. S., Hollowood, A. D., & Grant, J. R. (2011). Surgical Training: The Impact of Changes in Curriculum and Experience. *Journal of Surgical Education, 68*(1), 44–51. http://doi.org/10.1016/j.jsurg.2010.08.004

Patton, J. L., & Mussa-Ivaldi, F. A. (2004). Robot-assisted adaptive training: custom force fields for teaching movement patterns. *IEEE Transactions on Biomedical Engineering, 51*(4), 636 –646. http://doi.org/10.1109/TBME.2003.821035

Patton, J. L., Mussa-Ivaldi, F. A., & Rymer, W. Z. (2001). Altering movement patterns in healthy and brain-injured subjects via custom designed robotic forces. In *Proceedings of the 23rd Annual International Conference of the IEEE*

Engineering in Medicine and Biology Society (Vol. 2, pp. 1356 – 1359 vol.2). http://doi.org/10.1109/IEMBS.2001.1020448

Patton, J. L., Stoykov, M. E., Kovic, M., & Mussa-Ivaldi, F. A. (2006). Evaluation of robotic training forces that either enhance or reduce error in chronic hemiparetic stroke survivors. *Experimental Brain Research*, *168*(3), 368–383. http://doi.org/10.1007/s00221-005-0097-8

Patton, J. L., Wei, Y. J., Bajaj, P., & Scheidt, R. A. (2013). Visuomotor Learning Enhanced by Augmenting Instantaneous Trajectory Error Feedback during Reaching. *PLoS ONE*, *8*(1), e46466. http://doi.org/10.1371/journal.pone.0046466

Paulin, M. G. (1993). The Role of the Cerebellum in Motor Control and Perception. *Brain, Behavior and Evolution*, *41*(1), 39–50.

Philips, H. (2003). *The European working time directive: interim report and guidance from the Royal College of Surgeons of England Working Party*. Royal College of Surgeons of England.

Prasad, S. K., Kitagawa, M., Fischer, G. S., Zand, J., Talamini, M. a, Taylor, R. H., & Okamura, A. M. (2003). A Modular 2-DOF Force-Sensing Instrument for Laparoscopic Surgery. *Proc. 6th Int. Conf. Medical Image Comput. Computer-Assisted Intervention*, 279–286. http://doi.org/10.1007/978-3-540-39899-8_35

Reinkensmeyer, D. J., & Patton, J. L. (2009). Can Robots Help the Learning of Skilled Actions? *Exercise and Sport Sciences Reviews*, *37*(1), 43–51. http://doi.org/10.1097/JES.0b013e3181912108

Reisman, D. S., Wityk, R., Silver, K., & Bastian, A. J. (2007). Locomotor adaptation on a split-belt treadmill can improve walking symmetry post-stroke. *Brain*, *130*(7), 1861–1872.

Rescorla, R. a, & Wagner, a R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II Current Research and Theory*, *21*(6), 64–99. http://doi.org/10.1101/gr.110528.110

Reznick, R. K., & MacRae, H. (2006). Teaching Surgical Skills — Changes in the Wind. *New England Journal of Medicine*, *355*(25), 2664–2669. http://doi.org/10.1056/NEJMra054785

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

Salmoni, A. W., Schmidt, R. A., & Walter, C. B. (1984). Knowledge of results and motor learning: A review and critical reappraisal. *Psychological Bulletin*, *95*(3), 355–386. http://doi.org/10.1037/0033-2909.95.3.355

Schaverien, M. V. (2010). Development of Expertise in Surgical Training. *Journal of Surgical Education*, *67*(1), 37–43. http://doi.org/10.1016/j.jsurg.2009.11.002

Schmidt, R. a., & Bjork, R. a. (1992). New Conceptualizations of Practice: Common Principles in Three Paradigms Suggest New Concepts for Training. *Psychological Science*, *3*(4), 207–217. http://doi.org/10.1111/j.1467-9280.1992.tb00029.x

Sedef, M., Samur, E., & Basdogan, C. (2006). Real-Time Finite-Element Simulation of Linear Viscoelastic Tissue Behavior Based on Experimental Data. *IEEE Computer Graphics and Applications*, *26*(6), 58–68. http://doi.org/10.1109/MCG.2006.135

Selen, L. P. J., Franklin, D. W., & Wolpert, D. M. (2009). Impedance Control Reduces Instability That Arises from Motor Noise. *The Journal of Neuroscience*, *29*(40), 12606–12616. http://doi.org/10.1523/JNEUROSCI.2826-09.2009

Selwitz, R. H., Ismail, A. I., & Pitts, N. B. (2007). Dental caries. *The Lancet*, *369*(9555), 51–59.

Selzer, M., Clarke, S., Cohen, L., Duncan, P., & Gage, F. (Eds.). (2006). *Textbook of Neural Repair and Rehabilitation 2 Volume Hardback Set* (1st ed.). Cambridge University Press.

Seymour, N. E. (2008). VR to OR: A Review of the Evidence that Virtual Reality Simulation Improves Operating Room Performance. *World Journal of Surgery*, *32*(2), 182–188. http://doi.org/10.1007/s00268-007-9307-9

Seymour, N. E., Gallagher, A. G., Roman, S. A., O'Brien, M. K., Bansal, V. K., Andersen, D. K., & Satava, R. M. (2002). Virtual Reality Training Improves Operating Room Performance. *Annals of Surgery*, *236*(4), 458–464.

Shadmehr, R., & Krakauer, J. W. (2008). A computational neuroanatomy for motor control. *Experimental Brain Research*, *185*(3), 359–81. http://doi.org/10.1007/s00221-008-1280-5

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error Correction, Sensory Prediction, and Adaptation in Motor Control. *Annual Review of Neuroscience*, *33*(1), 89–108. http://doi.org/10.1146/annurev-neuro-060909-153135

Sigrist, R., Rauter, G., Riener, R., & Wolf, P. (2012). Augmented visual, auditory, haptic, and multimodal feedback in motor learning: A review. *Psychonomic Bulletin & Review*, 1–33.

Squeri, V., Basteris, A., & Sanguineti, V. (2011). Adaptive regulation of assistance "as needed" in robot-assisted motor skill learning and neuro-rehabilitation. In *2011 IEEE International Conference on Rehabilitation Robotics (ICORR)* (pp. 1 –6). http://doi.org/10.1109/ICORR.2011.5975375

Srinivasan, & LaMotte, R. H. (1995). Tactual discrimination of softness. *Journal of Neurophysiology*, *73*(1), 88–101.

Srinivasan, M. A., Beauregard, G. L., & Brock, D. L. (1996). The impact of visual information on the haptic perception of stiffness in virtual environments. In *ASME Winter Annual Meeting* (Vol. 58, pp. 555–559).

Sturm, L. P., Windsor, J. A., Cosman, P. H., Cregan, P., Hewett, P. J., & Maddern, G. J. (2008). A Systematic Review of Skills Transfer After Surgical Simulation Training: *Annals of Surgery*, *248*(2), 166–179. http://doi.org/10.1097/SLA.0b013e318176bf24

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.

Takahashi, C. D., Scheidt, R. A., & Reinkensmeyer, D. J. (2001). Impedance Control and Internal Model Formation When Reaching in a Randomly Varying Dynamical Environment. *Journal of Neurophysiology*, *86*(2), 1047–1051.

Tan, Durlach, N. I., Beauregard, G. L., & Srinivasan, M. A. (1995). Manual discrimination of compliance using active pinch grasp: The roles of force and work cues. *Perception & Psychophysics*, *57*(4), 495–510.

Tan, Durlach, N. I., Shao, Y., & Wei, M. (1993). Manual resolution of compliance when work and force cues are minimized. *Advances in Robotics, Mechatronics and Haptic Interfaces*, *49*, 99–104.

Tang, Hanna, Joice, & Cuschieri. (2004). IDentification and categorization of technical errors by observational clinical human reliability assessment (ochra) during laparoscopic cholecystectomy. *Archives of Surgery*, *139*(11), 1215–1220.

Tan, Pang, X. D., & Durlach, N. I. (1992). Manual resolution of length, force, and compliance. *Advances in Robotics*, *42*, 13–18.

Taylor, J. A., Klemfuss, N. M., & Ivry, R. B. (2010). An explicit strategy prevails when the cerebellum fails to compute movement errors. *Cerebellum (London, England)*, *9*(4), 580–6.

Teodorescu, K., Bouchigny, S., & Korman, M. (2013). Training Haptic Stiffness Discrimination Time Course of Learning With or Without Visual Information and Knowledge of Results. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *55*(4), 830–840.

Thelen, E. (1989). The (re)discovery of motor development: Learning new things from an old field. *Developmental Psychology*, *25*(6), 946–949. http://doi.org/10.1037/0012-1649.25.6.946

Thorndike, E. L. . W. R. S. (1901). The influence of improvement in one mental function upon the efficiency of other functions. *Psychological Review*, *8*, 247–261.

Tiest, & Kappers, A. M. (2008). Kinaesthetic and cutaneous contributions to the perception of compressibility. In *Haptics: Perception, Devices and Scenarios* (pp. 255–264). Springer.

Tiest, W. M. B., & Kappers, A. M. L. (2009). Cues for Haptic Perception of Compliance. *IEEE Transactions on Haptics*, *2*(4), 189–199.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*(9), 907–915. http://doi.org/10.1038/nn1309

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*(11), 1226–1235. http://doi.org/10.1038/nn963

Tresilian, J. (2012). *Sensorimotor Control and Learning: An introduction to the behavioral neuroscience of action*. London: Palgrave Macmillan.

Tsuda, S., Scott, D., Doyle, J., & Jones, D. B. (2009). Surgical Skills Training and Simulation. *Current Problems in Surgery*, *46*(4), 271–370. http://doi.org/10.1067/j.cpsurg.2008.12.003

Turchetti, G., Palla, I., Pierotti, F., & Cuschieri, A. (2011). Economic evaluation of da Vinci-assisted robotic surgery: a systematic review. *Surgical Endoscopy*, *26*(3), 598–606. http://doi.org/10.1007/s00464-011-1936-2

van Beers, R. J. (2009). Motor Learning Is Optimally Tuned to the Properties of Motor Noise. *Neuron*, *63*(3), 406–417. http://doi.org/10.1016/j.neuron.2009.06.025

van der Linde, R. Q., Lammertse, P., Frederiksen, E., & Ruiter, B. (2002). The HapticMaster, a new high-performance haptic interface. In *Proc. Eurohaptics* (pp. 1–5).

van der Meijden, O. A. J., & Schijven, M. P. (2009). The value of haptic feedback in conventional and robot-assisted minimal invasive surgery and virtual reality training: a current review. *Surgical Endoscopy*, *23*(6), 1180–1190. http://doi.org/10.1007/s00464-008-0298-x

van Hove, P. D., Tuijthof, G. J. M., Verdaasdonk, E. G. G., Stassen, L. P. S., & Dankelman, J. (2010). Objective assessment of technical surgical skills. *British Journal of Surgery*, *97*(7), 972–987. http://doi.org/10.1002/bjs.7115

Verstynen, T., & Sabes, P. N. (2011). How Each Movement Changes the Next: An Experimental and Theoretical Study of Fast Adaptive Priors in Reaching. *The Journal of Neuroscience*, *31*(27), 10050–10059. http://doi.org/10.1523/JNEUROSCI.6525-10.2011

von Websky, M. W., Vitz, M., Raptis, D. A., Rosenthal, R., Clavien, P. A., & Hahnloser, D. (2012). Basic Laparoscopic Training Using the Simbionix LAP Mentor:

Setting the Standards in the Novice Group. *Journal of Surgical Education*, *69*(4), 459–467. http://doi.org/10.1016/j.jsurg.2011.12.006

VR Laboratory - University of Twente. (n.d.). Retrieved November 22, 2015, from http://www.vrlab.ctw.utwente.nl/

Wattiez, A., Soriano, D., Cohen, S. B., Nervo, P., Canis, M., Botchorishvili, R., … Bruhat, M. A. (2002). The Learning Curve of Total Laparoscopic Hysterectomy: Comparative Analysis of 1647 Cases. *The Journal of the American Association of Gynecologic Laparoscopists*, *9*(3), 339–345. http://doi.org/10.1016/S1074-3804(05)60414-8

Wilkie, R. M., Wann, J. P., & Allison, R. S. (2008). Active gaze, visual look-ahead, and locomotor control. *Journal of Experimental Psychology: Human Perception and Performance*, *34*(5), 1150–1164.

Wilson, A. D., Collins, D. R., & Bingham, G. P. (2005). Human movement coordination implicates relative direction as the information for relative phase. *Experimental Brain Research*, *165*(3), 351–361.

Winstein, C. J., Pohl, P. S., & Lewthwaite, R. (1994). Effects of Physical Guidance and Knowledge of Results on Motor Learning: Support for the Guidance Hypothesis. *Research Quarterly for Exercise and Sport*, *65*(4), 316–323. http://doi.org/10.1080/02701367.1994.10607635

Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences*, *1*(6), 209–216.

Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature Reviews Neuroscience*, *12*(12), 739–751. http://doi.org/10.1038/nrn3112

Wolpert, D. M., & Flanagan, J. R. (2010). Motor learning. *Current Biology*, *20*(11), R467–R472. http://doi.org/10.1016/j.cub.2010.04.035

Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, *2*(9), 338–347. http://doi.org/10.1016/S1364-6613(98)01221-2

Wu, H. G., Miyamoto, Y. R., Gonzalez Castro, L. N., Ölveczky, B. P., & Smith, M. A. (2014). Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nature Neuroscience*, *17*(2), 312–21.

Wu, W. C., Basdogan, C., & Srinivasan, M. A. (1999). Visual, haptic, and bimodal perception of size and stiffness in virtual environments. *ASME Dynamic Systems and Control Division*, *67*, 19–26.

Yiannakopoulou, E., Nikiteas, N., Perrea, D., & Tsigris, C. (2015). Virtual reality simulators and training in laparoscopic surgery. *International Journal of Surgery*, *13*, 60–64. http://doi.org/10.1016/j.ijsu.2014.11.014

Ziemann, U., Tergau, F., Bruns, D., Baudewig, J., & Paulus, W. (1997). Changes in human motor cortex excitability induced by dopaminergic and anti-dopaminergic drugs. *Electroencephalography and Clinical Neurophysiology/Electromyography and Motor Control*, *105*(6), 430–437.

Appendix 1

“Computational Models and Haptic Tools for Next Generation Surgery” (MEng team report)

COMPUTATIONAL MODELS AND HAPTIC TOOLS FOR NEXT GENERATION SURGERY

James H. Chandler, Matthew J. Dickson, Earle S. Jamieson, Thomas Mueller, Thomas Reid

ABSTRACT

Recent years have seen a transfer of surgical procedures from traditional open surgery to Minimally Invasive Surgery (MIS), and more recently, to Robot-Assisted Laparoscopic Surgery (RALS). These have shown significant benefits over open surgery, but have resulted in the reduction or complete loss of haptic (force and touch) feedback. This has decreased the perception of applied force and ability to discriminate tissue features. Despite previous attempts to resolve this issue using haptic technology, a clinically available solution has not been realised. This study presents an approach for the detection of tumours in human liver through haptic palpation. Human liver with embedded tumours has been modelled for Finite Element Analysis and physical experiments. The produced results form the basis of a virtual haptic surgical system that allows palpation with integrated visual feedback. Results of a human factors study on the effects of size and depth of inclusions have shown statistical significance in the mean accuracy and search time using the haptic system (p<0.05). An inverse method that allows the characterisation of tumour parameters using input experimental data has also been implemented. The system has shown potential for effective tumour detection and sets a framework for future development.

Keywords — Haptics, Surgical Robotics, Palpation, Biomechanical modelling.

1 INTRODUCTION

In the UK, more than one in three people will develop some form of cancer in their lifetime [1] and one in four of all deaths in the UK are caused by cancer [2]. Cancer commonly manifests itself as hard abnormal masses (tumours) embedded within softer tissue (organs) [3]. Tumours emerge as a result of neoplasia (irregular tissue growth) and invade and destroy surrounding tissue. In the case of malignant tumours, early detection and removal decreases the chances of the disease spreading, hence increasing the patient's likelihood of survival [4].Currently, the most effective method for curing cancer is surgery, however, efforts to improve the current techniques are low in relation to other cancer-related research activities [5].

This paper focuses on surgical systems used in Minimally Invasive Surgery (MIS) for tumour resection (removal of tissue), specifically the lack of haptic (force and touch) feedback they provide to surgeons during non-open surgery. In an effort to improve surgical equipment, work has been carried out using various techniques to model a diseased human liver and initial findings show promising results.

1.1. Overview of Current Surgical Techniques

MIS is currently the preferred method of surgery for many procedures. It consists of inserting specialised instruments through small incisions, or natural orifices, to perform surgery. In relation to open surgery, MIS reduces blood loss, tissue trauma, pain experienced by the patient, recovery time and the risk of post-operative infection [6, 7].

Conventional laparoscopic surgery (laparoscopy) is a particular form of MIS whereby the surgical tools are inserted into the patient's abdominal cavity through small ports (trocars). This method, however, has inherent limitations: the point at which the laparoscopic tool is inserted into the abdomen acts as a pivot, which requires the surgeon to cope with the reversal of hand movements at the instrument tip. This results in loss of intuitiveness which consequently leads to incorrect movements, fatigue and premature tremor for the surgeon [6]. In comparison to open surgery, further disadvantages of this approach include: reduced depth perception, dexterity and hand-eye coordination [7]. Tactile and kinaesthetic feedback received by the surgeon are reduced or eliminated. This is primarily due to the lack of direct contact between finger mechanoreceptors and the tissue, and is exacerbated with issues such as friction between the shaft and the trocar, and force scaling created around the tool pivot [6].

In order to solve the motion constraint problems experienced in MIS, Robot-Assisted Laparoscopic Surgery (RALS) systems such as the ZEUS and da Vinci (Intuitive Surgical, Sunnyvale, CA, USA) have been developed. These master/slave devices allow the surgeon to remotely control the movement and actions of instrumented tools using a

console with flexible graphical user interfaces and, in the case of the da Vinci system, three-dimensional visualisations. RALS allows the surgeon to perform a variety of minimally invasive operations more effectively than laparoscopy. This is primarily due to increased dexterity, motion scaling and the reduction of physiological tremor [8]. A major limitation of RALS, however, is the incapability of providing force feedback to the surgeon, causing them to rely on visual feedback alone [9]. This makes the diagnostic technique of palpation (assessment of tissue properties through physical manipulation) unachievable and can lead to excessive tissue trauma during manipulations [6] [7].

1.2. Potential Benefits of Palpation in MIS

Vision alone is not enough for the effective detection of tumours [10]. In contrast to using palpation, as practiced in open surgery, surgeons are forced to estimate the location of tumours by comparing pre-operative imaging scans to the observable operative region [11]. Errors in such estimations can result in involved resections (edge of extraction containing tumour tissue), putting the patient at risk of increased local tumour recurrence [12]. More accurate definition of the tumour position using palpation during MIS could potentially aid the surgeon in defining the location of embedded tumours more accurately. In order to achieve effective palpation during MIS, force feedback must be achieved to translate measured reaction forces through to the surgeon's kinaesthetic senses. This would additionally allow the surgeon to perceive the tension or hardness of tissue and be able to measure the variation in their properties, all of which are possible through the sense of touch.

Coupling the qualities of bi-directional human haptic perception with a next generation surgical device could potentially increase the ability for the surgeon to make improved operative assessments through remote palpation. For assessing the feasibility of such a system, the human liver is an appropriate case study as the literature contains relevant information, such as Young's Modulii [13], for both the liver and the tumours it may contain [14]. Additionally, its relatively large size and regular shape are beneficial for modelling its structure [15].

1.3. Overview of Haptics

The term *haptic* is used to describe interactions relating to or based on the sense of touch. The haptic sensory system employs cutaneous and kinaesthetic receptors to interact with objects during active (dynamic) procedures. Haptic feedback has been used in virtual environments for over five decades, starting with flight simulators and master-slave teleoperated robotic devices [16]. Since then, the range and quality of haptic applications has improved tremendously within both kinaesthetic and tactile feedback. A range of commercial devices have been developed, hence increasing the accessibility of haptics to industrial and research developers [17]. With the aim of exploiting the benefits of haptics in other industries, the surgical field has seen the development of applications for training and teleoperation becoming more frequent, and is currently an area of active research.

1.3.1. Current State of Haptics in Surgery

To date, haptic technologies have been implemented within many virtual environments specifically aimed at surgical training applications [18-20]. In efforts to expand the benefits of haptics into real surgery, extensive research has been carried out with the aim to augment surgeons' abilities in the operating theatre [21]. Due to the need for a master/slave device for the implementation of haptics, however, it is only possible to integrate haptic feedback into RALS. Currently, the only commercial surgical device with integrated haptics is the neuroArm (Calgary University, Canada), in which various systems have been integrated to recreate the sight, audio and feel of brain surgery. NeuroArm has proven to be a less invasive, more accurate method than previous procedures, partly due to the employment of haptics [22]. Although this system is only applicable to neurosurgery, it shows potential for the commercialisation of haptics in the operating theatre.

1.3.2. System Requirements for Palpation in MIS

Mahvash *et al* [23] found, during a palpation exercise on a physical tissue model, that using direct force feedback resulted in the highest percentage of correct inclusion locations by their test participants over other feedback methods including: Graphical force feedback, both direct and graphical force feedback, and no force feedback at all. Furthermore, the use of virtual haptic interfaces in combination with visual feedback has shown significant advantages within research and surgical training environments [19] [24].

Jungsik *et al.* [21] present a framework for a non-invasive, real-time haptic palpation system (Fig. 1). The surface force response of a physical tissue model was measured through telemanipulation using a

slave manipulator and force sensor, and transmitted to the user via a haptic device. The study does not, however, justify the characteristics of the physical model as representative mechanical properties of biological tissue. Furthermore, it does not assess human perceptibility of measured forces and the effects of varying tissue and inclusion parameters.

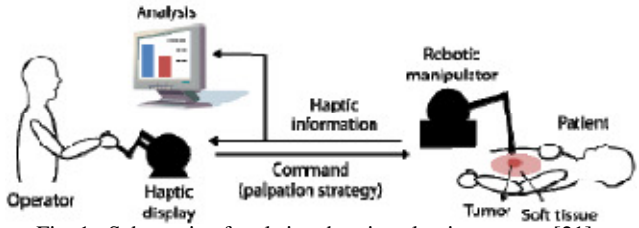

Fig. 1. Schematic of real-time haptic palpation system [21]

Previous studies [19] [25] have shown that size, depth and dissimilarity in stiffness between healthy and diseased tissues are factors that affect the ability to detect tumours during palpation. This indicates that these are important considerations for the validation of a haptic system.

This paper discusses the application of Finite Element Analysis (FEA) and experimental data in the development and validation of a novel haptic simulation tool for the detection of tumours in artificial human liver through virtual palpation. The results of human trials for the assessment of detection rates of embedded tumours with different parameters are also discussed.

2 MODELLING & EXPERIMENTATION

In order to allow the user to "feel" the representative stiffness of a human liver with an embedded tumour, representative Response Force Surfaces (RFS) were imported into a bespoke haptic system. These were generated using an FEA model and compared with results from physical experiments. A human factors study enabled both quantitative and qualitative assessment of human performance for the detection of tumours. An inverse method allows for the determination of tissue and tumour parameters with given input data from physical or computational models.

2.1. Research Overview

A set of parameters common to all work areas of the project were produced. This allowed the exchange of data between the different modelling methods, as illustrated in Fig. 2.

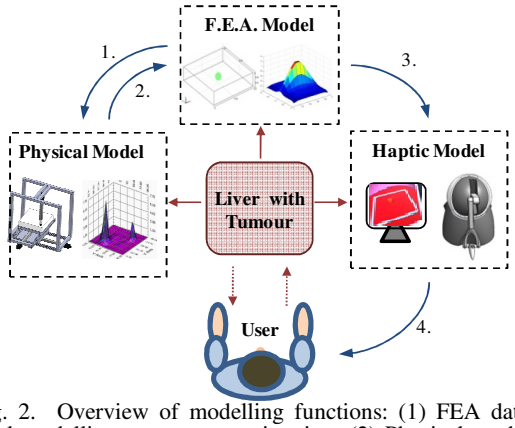

Fig. 2. Overview of modelling functions: (1) FEA data for physical modelling parameters estimation; (2) Physical modelling data for inverse methods; (3) Use of FEA data for implementation of haptic feedback; (4) Virtual haptic system presented to user.

The mean length of the male human liver is 105 mm [15], and liver tumour sizes range from a few millimetres up to 20 cm [19]. Furthermore, a study [25] found notable differences in human and force sensor abilities to detect inclusions of 6.5, 9.5 and 12.5 mm diameter. Based on information from the literature and available resources, the common parameters across the FEA, physical and haptic models were defined as:

- Liver size – 100 x 100 x 25 mm
- Inclusion sizes – 6, 9 and 12 mm diameter
- Embedded inclusion depth – One inclusion flush with top surface, another flush with bottom surface, and three evenly spaced in between.
- Indenter – Cylindrical, with a 10mm diameter semi-spherical contacting end.
- Indentation depths - 5, 10 and 15 mm into the tissue, measured from the surface.

The RFSs generated using FEA and imported into the haptic system used Young's Moduli of 1 and 75 kPa for the liver and tumour components, respectively. These are representative properties of real human tissue. The RFSs generated through experimental methods, however, were obtained using silicone models and steel ball bearings. The Young's Moduli of these are approximately 6.5 kPa and 200 GPa, respectively. The parameters of the FEA model were adjusted to match the physical experiments to allow a direct comparison, and hence assess the validity of the computational and physical results.

A common testing protocol has been adopted for all testing methods to define the size and depth of inclusions within the tissue. Table 1 shows the testing protocol for each test permutation, T_{SD} (where S and D

are the inclusion size and depth indices, respectively), and indicates which experiments were used in each method.

Table 1. Test protocol. All tests were run in F.E.A., black dots indicate tests run in physical experiments, red dots indicate tests run in haptic models

		Inclusion Depth				
		Surface	T_{S1}-T_{S3} midpoint	Centre	T_{S3}-T_{S5} midpoint	Base
Inclusion Size (mm)	6	T_{11} ••	T_{12}	T_{13} ••	T_{14}	T_{15} •
	9	T_{21} ••	T_{22}	T_{23} ••	T_{24}	T_{25} •
	12	T_{31} ••	T_{32}	T_{33} ••	T_{34}	T_{35} •

2.2. Physical Testing

Models of liver with embedded inclusions were created; an automated testing system was developed to measure their RFSs using a sensory probe.

2.2.1. Physical Models

A physical model of a section of human liver was created using a 2-part (A and B) silicone formula (Platsil Gel 10). A ratio of 1:1:3 of A and B components and Deadener (a stiffness reduction additive) respectively, was used in the mixing process. This produced a physical tissue model with a measured Young's Modulus of approximately 6.5kPa; which is in the upper stiffness region of a healthy liver [13]. Steel ball bearings (Young's Modulus ≈ 200 GPa) were used to represent embedded tumours. Although these are much stiffer than real tumours, they were chosen as a proof-of-concept method for assessing sensing capabilities. The inclusion stiffness is seven orders of magnitude greater than the silicone, and hence a large force response difference was expected between them. Two inclusions (each at a different depth) were integrated into each model to decrease testing time. Data from an FEA model was used to determine the minimum required separation between the inclusions to avoid influence between them and hence treat them as decoupled entities.

2.2.2. Physical Testing System

During palpation, forces do not usually exceed 5N [26]. A Force Sensing Resistor (FSR) (*FlexiForce A201*), capable of measuring this force range, was chosen for measuring RFSs of the physical tissue models. An amplifier circuit is used to produce a voltage proportional to the sensor's change in resistance under varied load. The sensor was incorporated into a mechanical housing to form a sensory probe for unidirectional indentation normal to the tissue surface. Sensitivity of low forces was increased by compressively pre-loading the sensor, and hence bypassing an inherent low response region. The sensory probe was calibrated by applying known, incremental loads through its full sensing range and recording the output voltage at each load increment. A fifth-order polynomial fit was implemented to characterise its non-linear force-voltage relationship, enabling the measurement of forces in Newtons. The effects of friction between the indenter and its enclosure have been neglected.

A testing system was developed to enable the measurement of RFSs of physical tissue models using the sensory probe. It consists of three linear actuators with position feedback, mounted in the x, y and z axes. This allows automated tri-axial movement of the sensory probe relative to the physical tissue model. Control of the actuator movements is made possible using data acquisition hardware (compact DAQ, NI) and software (LabVIEW, NI). Response force measurements were taken and logged at pre-defined indentation depths and positions at a specified in-plane resolution. Each test was repeated to observe the reproducibility of the RFS. Fig. 3 shows the system and components.

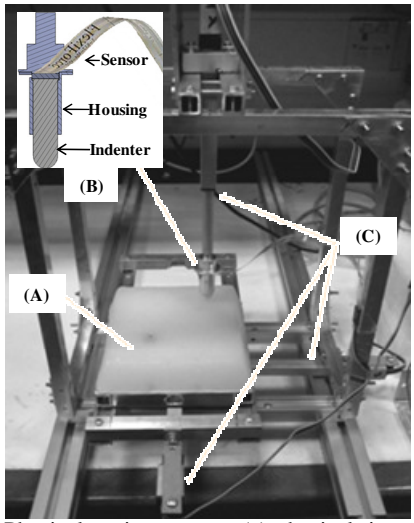

Fig. 3. Physical testing system: (a) physical tissue model, (b) sensory probe with integrated FSR and (c) linear actuators.

To quantitatively determine if the embedded inclusions were detected using the physical testing system, statistical analysis was conducted. For each RFS the mean and standard deviation (σ) were calculated and used to determine a Gaussian probability density distribution. Response forces were considered highly likely ($p < 0.03$) or almost certainly ($p < 6\text{x}10^{-7}$) caused by the presence of an inclusion if

they were between 3σ and 5σ, or greater than 5σ, respectively, from the corresponding RFS mean.

2.3. Finite Element Analysis

FEA was employed to model samples of liver with embedded inclusions to enable the prediction of their behaviour under various loading regimes. The data obtained was used to predict RFSs and had its main application in prescribing the haptic force response.

2.3.1. Finite Element Analysis Tissue Model

The model representing the liver sample was created using a FEA software package (COMSOL Multiphysics 3.5). The inclusion is defined as a function of position within the domain's material properties. Respective Young's Moduli are applied for the liver and inclusion regions. To mimic palpation, a surface deformation in the shape of the generic spherical probe geometry was applied, acting in one discrete location. The integration of the deformed surface yields the force required to prescribe the surface deformation.

The FEA model computes one instance only. In order to iterate the computation of the response force, the model was exported into a programming environment (MATLAB, The MathWorks). This allows the generation of two- and three-dimensional response force profiles that depict the variation of response force for a range of positions of variable resolution. Following the establishment of mesh independence for the model, response force profiles were determined for each scenario laid out in the test series.

2.3.2. Inverse Method

The inverse method enables the determination of variables in the model that had previously been pre-defined inputs. The method was tested by determining the in-plane surface coordinates of the inclusions. For this, results from the physical testing system and an arithmetic function were used as input data. The arithmetic function replaced the FEA model to reduce computational expense in the process of establishing proof of concept.

A Root Mean Square Error (RMSE) between the two datasets evaluates their agreement. The inverse method minimises for this error by varying the X and Y position of the inclusion in the computational dataset. This is done using the *fmincon* function from MATLAB's optimisation toolbox. The solution that the function converges to represents the suspected location of the inclusion in the experimental data set. The accuracy of the method is assessed by measuring the radial error.

2.4. Haptic System

A custom Dynamic Link Library (DLL) has been developed to enable communication between a 6 degree-of-freedom haptic device (PHANToM Omni, SensAble Technologies) and a bespoke User Interface (UI) developed in a graphical development environment (LabVIEW, NI). This was implemented using a 3 GHz single core Pentium processor. The DLL enables data transfer between the UI and the OpenHaptics Haptic Device Application Programming Interface (HDAPI), allowing low-level control over the device. The UI applies three-dimensional (3D) position and orientation data to a virtual end-effector within a simulated visual surgical scene and allows for high-level control of the haptic system. A schematic of the full system is shown in Fig. 4.

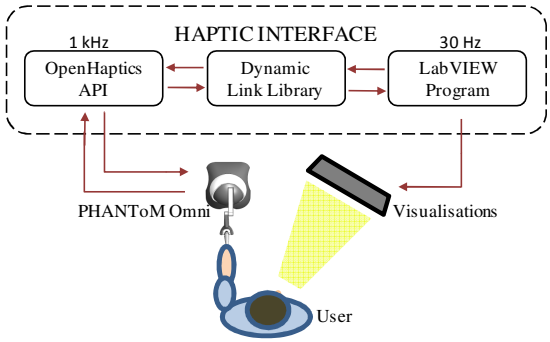

Fig. 4. Overview of haptic system illustrating data flow and loop rates

The parallel architecture enables real-time, high fidelity force rendering and visualisations, whilst providing a flexibile platform with potential of future scalability with other SensAble haptic devices.

In order to implement high-fidelity force rendering and effective visualisations, update rates of approximately 1 kHz [16] and 30 Hz [17], respectively, are recommended. Whilst the visual update rate is heavily dependent on the complexity of the visual scene, the HDAPI enables consistent force rendering at 1 kHz.

2.4.1. Real-time Force Rendering

Discrete FEA RFS data is approximated using a Gaussian function (described in Equation 1) and implemented in the haptic system.

$$F_G = Ae^{-\frac{x_r^2}{2\sigma^2}} + B \quad (1)$$

Where F_G is the output force (N), A is the peak force (N) relative to the baseline, B, σ is the function width variable and x_r is the radial distance to the inclusion centre (m).

For any input RFS dataset, the σ value of the Gaussian function is varied until the overall RMSE is minimal. Fig. 5 shows a representative Gaussian fit to real FE data, where edge effects have been omitted.

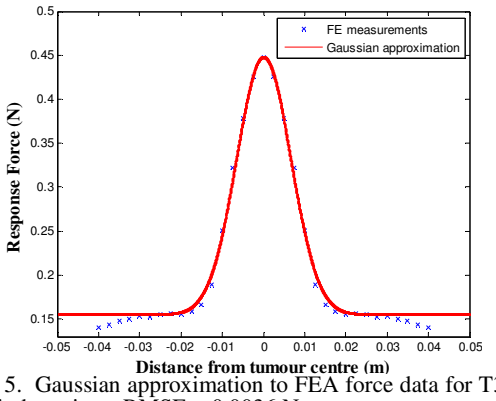

Fig. 5. Gaussian approximation to FEA force data for T33with 10 mm indentations. RMSE = 0.0036 N

This approximation represents an inclusion within an infinite domain of healthy tissue. The variables for the Gaussian function are imported into the haptic interface and a force is computed as a function of position at a frequency of 1 kHz.

The FEA data shows that the force-indentation relationship for indentations from 5 to 20 mm is almost linear and hence a linear assumption has been made by implementing a spring stiffness model according to Hooke's Law, F=kx. Upon collision of the virtual tool with a virtual object, forces are generated as a function of indentation depth.

In order to increase the discrimination between inclusions and soft tissue within the kinematic constraints of the haptic device, response forces have been augmented by a factor of 5. This increases the magnitude of the perceived forces without adjusting the peak to baseline stiffness ratio.

2.4.2. Visualisation

A graphical display has been developed and integrated into the UI to provide the user with a simulation of a representative visual scene of MIS procedures. Following insertion into the virtual abdomen, the user is presented with a surgical probe and a deformable liver surface. Forward kinematic analysis of the Phantom Omni was undertaken to define the available workspace and validate the position and orientation data obtained from the HDAPI. A variable Gaussian function is used to insert a deformation profile into the liver surface as a function of the 3D position of the indenter. The function parameters have been adjusted to match the deformations predicted using the FEA model. Fig. 6 shows the complete haptic system during operation.

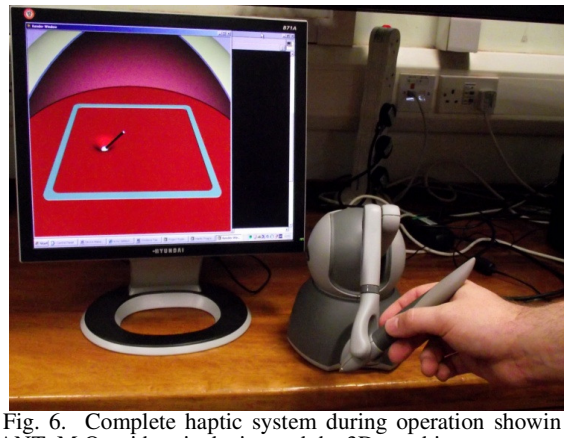

Fig. 6. Complete haptic system during operation showin the PHANToM Omni haptic device and the 3D graphics

2.5. Human Factors Study

To evaluate the current haptic system, a human factors study was conducted using 20 participants with no surgical training and little or no previous experience with virtual simulations. The aim of the study was to (1) assess the user's ability to discriminate between healthy tissue and inclusions, (2) to evaluate the effects of tumour size and depth on required search time and the accuracy of location, and (3) test the haptic system under operational conditions.

A test group of 20 participants (16 males and 4 females) with a mean age ± SD of 22.7 ± 2.7 years took part in the study. Approval for the trials was gained from the Research Ethics Committee at the University of Leeds. Participants were given a full introduction to the research area and asked to give qualitative feedback regarding the haptic system, in addition to the quantitative measurements made.

2.5.1. Testing Protocol

Each participant was presented with a total of 14 virtual surfaces; these contained the FEA test data for the three inclusion sizes (6, 9 & 12 mm), each at two depth levels (surface and mid-depth) and the control surface (no inclusion). Each of these was presented twice throughout the trial, with the order and inclusion position randomised independently for each participant. For each surface, a 3 minute time limit was allocated as appropriate palpating time based on

feedback from surgeons. Users were tasked with palpating the surface and deciding whether one or zero inclusions were present. For a positive selection, the user was required to mark the location of the estimated centre of the inclusion; the radial error between selected and actual tumour position was taken as one variable for assessment. Negative selections and timed out rounds were recorded. The user was asked to give their level of confidence for each of their selections. Two-way Analysis of variance (ANOVA) was used to assess the significance of differences in the mean values of radial error and search time for the factors tested (tumour size and depth).

3 RESULTS

3.1. Experimental Validation

Results from the physical tissue models were obtained using the physical testing system, and FEA results were obtained from the computational model.

3.1.1. Physical Experiments

The sensory probe indented the physical tissue models with an in-plane resolution of 5mm, producing a 19 x 19 array of force measurements for each one. RFSs were generated for indentation depths of 5, 10 and 15mm. Fig. 7 shows a physical tissue model with its corresponding RFS at 10 mm indentations.

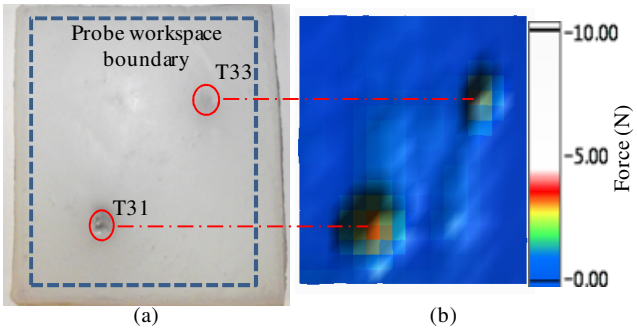

Fig. 7. (a) Physical tissue model and (b) RFS for T31/T33 (marked) obtained using physical testing system. Higher force responses are observed at the tumour positions.

The statistical method used to determine the likelihood of inclusions within the T31 physical tissue model for 10mm indentations is shown graphically within Fig. 8. The histogram illustrates the range of response force measurements with the statistically significant response forces indicated. For the tested physical models (Table 1) RFSs at all indentation depths, significant response forces were found in: T31, T21, T11, T33 (5mm) and T31, T21, T11, T33, T23, T35 (10 and 15 mm).

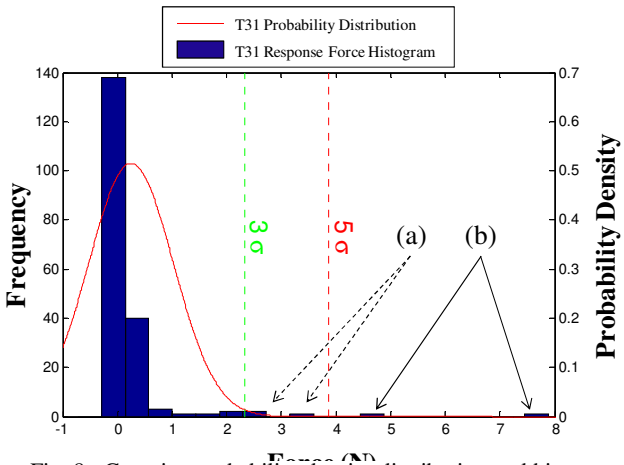

Fig. 8. Gaussian probability density distribution and histogram for T31 RFS data for 10mm indentations, where (a) indicates measured values between 3 and 5σ from the mean, and (b) indicates measured values greater than 5σ. These show measurements of high statistical significance ($p < 0.03$ and $p < 6\text{x}10^{-7}$, respectively).

3.1.2. Comparison of Physical and FEA results

A visual comparison of RFSs from FEA and physical experimentation is shown in Fig. 9. Inclusions that are large or located near the sample surface are visually identifiable, whereas inclusions that are small or near the bottom of the sample yield no significant response force. Although the obtained force values vary between the methods, a clear correlation can be observed.

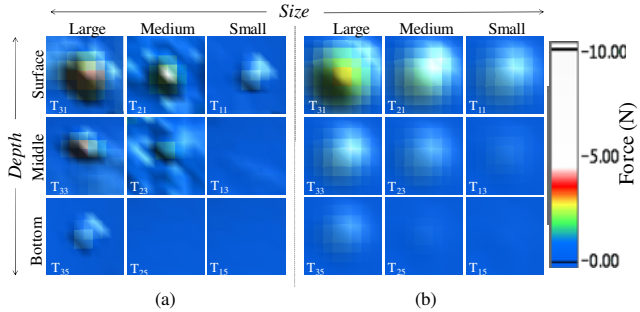

Fig. 9. (a) Experimental vs. (b) FEA data force intensity maps showing close agreement

3.2. Inverse Method (IM)

The performance of the inverse method was tested using the same data sets as for the human trials. The results obtained are shown in Table 2 and graphically in Fig. 11.

Table 2. Inverse Method radial error for different test series, where '*' indicates failure to identify inclusion location

Test Series	T31	T21	T11*	T33	T23*	T13*
Radial error [mm]	7.8	16.3	22.7	11.2	32.4	33.0

The results show that large inclusions and inclusions near the surface can be identified more accurately. Deviations from these conditions yield greater errors. The IM failed to identify the locations of T11, T23 and T13 and inherently identified their suspected locations as lying on the domain boundary.

3.3. Human Factors Study

Data collected as part of the human factors study are presented in Fig. 10. The selected positions for each trial are shown along with indication of 'timed out' and 'no tumour' selections. Participants were given randomised tumour positions, and hence the selected positions have been centralised to allow direct visual comparison.

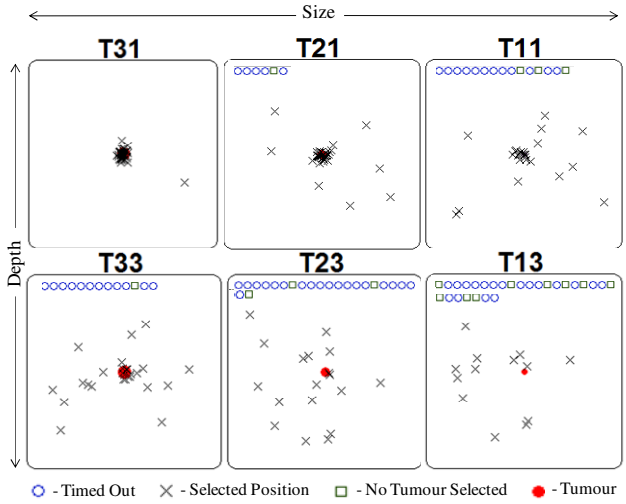

Fig. 10. Visual representation of human studies results, showing an increase in error and "timed out" and "no tumour" selections proportional to increase in inclusion size and/or depth.

Statistical analysis for the measured variables of radial error and trial time of each selection is described below. The mean and standard error of the results for the radial error are shown in Fig. 11. The results obtained from the inverse method assessment (Table 2) of the physical tissue models are superimposed. It should be noted, however, that these are *not* directly comparable to the human factors study results as they were obtained using different models. They have been included for the sole purpose of illustrating similar trends across the different techniques.

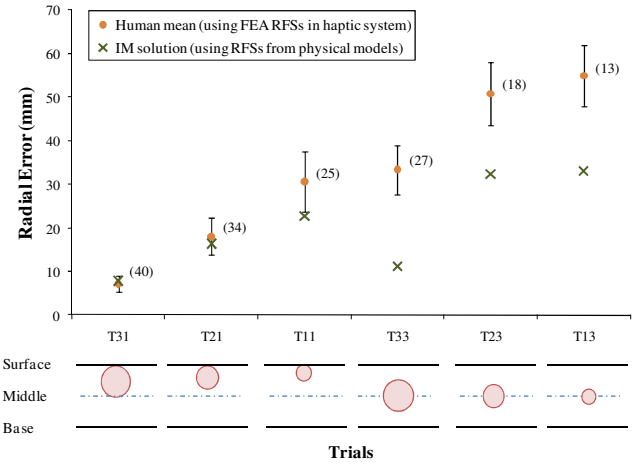

Fig. 11. Mean error and standard error of results, showing the effects of varying inclusion size and depth. Sample size is shown in brackets. The results from the inverse method have been superimposed to illustrate similar trends in detection.

Radial Error - The recorded radial errors were assessed only when 'tumour present' selections were made leading to disproportionate sample sizes. As such the standard error has been plotted alongside the mean for each of the trials, as shown in Fig. 11. The mean radial errors for the trials ranged from 7 mm to 55 mm. Variation in the mean radial error due to the size of the tumour present was found to be statistically significant (p = 0.0001). Also, the effect of tumour depth showed a significant effect on the mean radial error (p < 0.00005) with no significance shown for interaction of size and depth (p = 0.7228).

Trial Time - Time taken for each round was recorded. Average trial times ranged from 67 s up to 149 s, where an upper limit of 180 s resulted in a 'timed out' result. Statistical analysis carried out on the data (two-way ANOVA) used the data from every trial, giving an equal sample size of n = 40. The factors of tumour size and tumour depth both showed significant effect on the mean search time (p < 0.00005). Additionally, the interaction of size and depth on the mean time taken to make a decision for each trial was found to also be significant (p=0.0028). Each of the factors is therefore shown to be dependent on the specific value (level) of the other factor.

4 DISCUSSION

Physical testing and FEA – In the majority of the 10 mm indentation tests, it was possible to confidently identify the inclusions both qualitatively through visual inspection, and quantitatively through statistical assessment. At a 5mm indentation, inclusions located deep within the model could not be detected.

Additionally, at a 15mm indentation, inclusions located at the surface of the model were dislodged by the indenter during palpation and hence were not easily identified under visual inspection of the RFS, although statistical analysis showed a positive detection . Results for indentations of 10 mm were comparative to those obtained using the FEA model (Fig. 9). This indicates confidence within the modelling processes used, although larger discrepancies were shown at different depths. These have been accredited to some of the limitations of the testing environment such as friction between the indenter and probe housing and the sensor's poor low range sensitivity. Furthermore, issues of friction, calibration error within the linear actuators as well as limitations of the motor controllers restrict the abilities of the control system. Assuming linear structures and neglecting effects of friction in the FEA model may have further contributed to the discrepancies. It should be noted that the palpation depths used have not been verified for use within clinical practice, due to the absence of relevant literature. This issue would require clarification prior to commissioning of a complete system.

Inverse method - The IM achieved the identification of inclusions in clearly defined datasets i.e. smooth baselines and large peaks in response force. The error in detection increases with increasing implantation depth and decreasing inclusion size; a trend also observed in the human trials. The sensor's poor low range sensitivity directly affects how well response force peaks distinguish against baselines and subsequently the IM's ability to detect inclusions. Improved sensing hardware coupled with improved algorithms would improve the IM's performance and subsequently pose a potentially powerful tool in aiding medical staff to identify inclusions in organs during MIS or teleoperated diagnostics.

Haptic System – The interface allows high-fidelity haptic rendering with robust performance, and enables variation to the stiffness properties of the liver and tumour components. The system is limited due to the quality of the haptic device, including low force output, unwanted mechanical forces generated by joint movements and overheating after long operating periods. Improved graphics would be achieved using a higher specification PC with a dual core processor.

Human Trials - It is clear to see that the inclusions close to the surface are far easier to detect than those deeper into the tissue. In addition, the data presented in Fig. 10 and Fig. 11, along with subsequent statistical analysis, indicate that the positional accuracy and precision of selected tumours is affected by tumour size and depth. Indeed for a successful clinical haptic system, capable of effectively determining the location of tumours, it should promote both accuracy and precision. For this reason, force augmentation would almost certainly be required to increase the difference in force between the peak and baseline values. This would increase human perceptibility to any measurable size-depth combinations. The trade-off to this, however, is that rendered forces would not be truly representative of the tissue characteristics.

A non-direct comparison between the results of the inverse methods and human trials (Fig. 11) shows similar trends in the ability to detect embedded inclusions. A more direct study would involve using the same modelling data, specifically RFS data from the same physical tissue models, to observe the capabilities of the two techniques. Directly comparing the capabilities of the two techniques would yield information relating to the feasibility of teleoperational palpation and tissue inverse methods

5 CONCLUSIONS

- Representative response force data of human liver tissue with embedded tumours were generated through computational FEA models and palpation of physical models using a custom, automated testing system with integrated sensory probe.
- The statistical analysis carried out on the physical test results agrees with the FEA data and the results obtained from the human trials. These show similar detection rates for tumours of varied size and depth.
- Discrete response force surfaces produced using the FEA models were approximated using a continuous function for efficient implementation into the haptic system.
- The developed haptic virtual simulation system allows the user to palpate the surface of a simulated liver and feel its stiffness properties. The system is capable of high fidelity haptic rendering allowing users to detect tumours.
- A human factors study showed that the accuracy and precision of tumour detection depends greatly on tumour size and depth. Statistically significant variation was found in the mean values of accuracy of detection and required search time.
- An inverse method was successfully implemented to predict the location of a tumour within a tissue

sample, showing similar trends in detection accuracy to human trials. A potential application in MIS is the automatic detection of tumour position.

- Discrepancies between the FEA and physical results suggest the presence of errors in one or both of the methods. Potential factors contributing to these have been discussed. Identifying the major contributing factors would allow for refinement of the method(s).
- Augmentation of measured forces prior to haptic rendering could potentially increase accuracy in tumour location and decrease variability.

6 FUTURE WORK RECOMMENDATIONS

Recommendations for further work in this area include:

- Improvements to the force rendering functions used within the haptic system, to account for non-linear stiffness effects.
- Inclusion of time-dependent properties within the computational models to represent human liver more closely, e.g. poroelasticity.
- Improvements to the force sensing accuracy and precision to increase the validity of the response force surfaces from the physical tissue models.
- Development of a teleoperational system with haptic feedback from remote force sensing, allowing clinical feasibility to be assessed.
- Further human trials, testing the effects of factors such as force augmentation and user training within haptic tumour detection.

ACKNOWLEDGEMENTS

The team would like to thank our supervisors Peter Culmer and Rob Hewson, for their continued support and guidance throughout the project. We would also like to thank our industrial mentor, National Instruments, for supplying hardware and software to make completion of the project possible.

REFERENCES

1. Sasieni P, S.J., Ormiston-Smith N, Thomson C,Silcocks P, *Lifetime risk of cancer*, [Submitted].
2. Research, C. *All cancers combined statistics - Key Facts*. 2011 15/04/11 [cited 2011 19/04/2011]; Available from: http://info.cancerresearchuk.org/cancerstats/keyfacts/Allcancerscombined/#How3.
3. R. C. Bast , D.W.K., R. E. Pollock , R. R. Weichselbaum , J. F. Holland and E. Frei *Holland-Frei Cancer Medicine*2000: BC Decker
4. Andrew W. ElBardissi, M., MPH; Joseph A. Dearani, MD; Richard C. Daly, MD; Charles J. Mullany, MD; Thomas A. Orszulak, MD; Francisco J. Puga, MD; Hartzell V. Schaff, MD, *Survival After Resection of Primary Cardiac Tumors - A 48-Year Experience* 2008.
5. UK, C.R., *Cancer Reseach UK's Research Strategy*, 2008. p. 48.
6. Puangmali, P., et al., *State-of-the-Art in Force and Tactile Sensing for Minimally Invasive Surgery.* Sensors Journal, IEEE, 2008. **8**(4): p. 371-381.
7. Tholey, G., J.P. Desai, and A.E. Castellanos, *Force Feedback Plays a Significant Role in Minimally Invasive Surgery: Results and Analysis*, 2005.
8. Gomes, P., *Surgical robotics: Reviewing the past, analysing the present, imagining the future.* Robotics and Computer-Integrated Manufacturing, 2011. **27**(2): p. 261-266.
9. Meijden, S., *The value of haptic feedback in conventional and robot-assisted minimal invasive surgery and virtual reality training: a current review.* Surgical Endoscopy, 2009. **23**: p. 1180-1190.
10. Stoianovici, D., *Multi-imager compatible actuation principles in surgical robotics.* The International Journal of Medical Robotics and Computer Assisted Surgery, 2005. **1**(2): p. 86-100.
11. Current Medicine, I., *Current review of minimally invasive surgery*, ed. D.C. Brooks1998, Philadelphia: Springer.
12. NHS, *Laparoscopic liver resection - guidance*, in *IPG135*2005.
13. Mueller, S., *Liver stiffness: a novel parameter for the diagnosis of liver disease*, 2010, Dove Press.
14. Masuzaki, R., et al., *Assessing liver tumor stiffness by transient elastography.* Hepatology International, 2007. **1**(3): p. 394-397.
15. Walker, H.K., W.D. Hall, and J.W. Hurst, *Clinical methods: the history, physical, and laboratory examinations*1990: Butterworths.
16. Otaduy, M.A. and M.C. Lin, *High Fidelity Haptic Rendering.* Synthesis Lectures on Computer Graphics, 2006. **1**(1): p. 1-112.
17. Salisbury, K., F. Conti, and F. Barbagli, *Haptic rendering: introductory concepts.* Computer Graphics and Applications, IEEE, 2004. **24**(2): p. 24-32.
18. Alhalabi, M.O., et al. *Medical training simulation for palpation of subsurface tumor using HIRO.* in *Eurohaptics Conference, 2005 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2005. World Haptics 2005. First Joint.* 2005.
19. Langrana, N., et al., *Human performance using virtual reality tumor palpation simulation.* Computers & Graphics, 1997. **21**(4): p. 451-458.
20. Wang, D., et al., *Haptic rendering for dental training system.* Science in China Series F: Information Sciences, 2009. **52**(3): p. 529-546.
21. Jungsik, K., et al. *Inclusion detection with haptic-palpation system for medical telediagnosis*. in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*. 2009.
22. Greer, A.D., P.M. Newhook, and G.R. Sutherland, *Human-Machine Interface for Robotic Surgery and Stereotaxy.* Mechatronics, IEEE/ASME Transactions on, 2008. **13**(3): p. 355-361.
23. Mahvash, M., et al., *Force-Feedback Surgical Teleoperator: Controller Design and Palpation Experiments*, in *Proceedings of the 2008 Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*2008, IEEE Computer Society. p. 465-471.
24. Frederick P. Brooks, J., et al., *Project GROPEHaptic displays for scientific visualization.* SIGGRAPH Comput. Graph., 1990. **24**(4): p. 177-185.
25. Gwilliam, J.C., et al. *Human vs. robotic tactile sensing: Detecting lumps in soft tissue*. in *Haptics Symposium, 2010 IEEE*. 2010.
26. Richards, C., et al., *Skills evaluation in minimally invasive surgery using force/torque signatures.* Surgical Endoscopy, 2000. **14**(9): p. 791-798.