UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING, SCIENCE & MATHEMATICS

School of Engineering Sciences

Enhanced Pre-Clinical Assessment of Total Knee Replacement Using Computational Modelling with Experimental Corroboration & Probabilistic Applications

by

Anthony Michael Strickland

Thesis for the degree of Doctor of Philosophy

April 2009

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE & MATHEMATICS

SCHOOL OF ENGINEERING SCIENCES

Doctor of Philosophy

ENHANCED PRE-CLINICAL ASSESSMENT OF TOTAL KNEE REPLACEMENT USING COMPUTATIONAL MODELLING WITH EXPERIMENTAL CORROBORATION & PROBABILISTIC APPLICATIONS

by Anthony Michael Strickland

Demand for Total Knee Replacement (TKR) surgery is high and rising; not just in numbers of procedures, but in the diversity of patient demographics and increase of expectations. Accordingly, greater efforts are being invested into the pre-clinical analysis of TKR designs, to improve their performance *in-vivo*. A wide range of experimental and computational methods are used to analyse TKR performance pre-clinically. However, direct validation of these methods and models is invariably limited by the restrictions and challenges of clinical assessment, and confounded by the high variability of results seen *in-vivo.*

Consequently, the need exists to achieve greater synergy between different pre-clinical analysis methods. By demonstrating robust corroboration between *in-silico* and *in-vitro* testing, and both identifying & quantifying the key sources of uncertainty, greater confidence can be placed in these assessment tools. This thesis charts the development of a new generation of fast computational models for TKR test platforms, with closer collaboration with *in-vitro* test experts (and consequently more rigorous corroboration with experimental methods) than previously.

Beginning with basic tibiofemoral simulations, the complexity of the models was progressively increased, to include *in-silico* wear prediction, patellofemoral & full lower limb models, rig controller-emulation, and accurate system dynamics. At each stage, the models were compared extensively with data from the literature and experimental tests results generated specifically for corroboration purposes.

It is demonstrated that when used in conjunction with, and complementary to, the corresponding experimental work, these higher-integrity *in-silico* platforms can greatly enrich the range and quality of pre-clinical data available for decision-making in the design process, as well as understanding of the experimental platform dynamics. Further, these models are employed within a probabilistic framework to provide a statistically-quantified assessment of the input factors most influential to variability in the mechanical outcomes of TKR testing. This gives designers a much richer holistic visibility of the true system behaviour than extant 'deterministic' simulation approaches (both computational and experimental).

By demonstrating the value of better corroboration and the benefit of stochastic approaches, the methods used here lay the groundwork for future advances in pre-clinical assessment of TKR. These fast, inexpensive models can complement existing approaches, and augment the information available for making better design decisions prior to clinical trials, accelerating the design process, and ultimately leading to improved TKR delivery *in-vivo* to meet future demands.

LISTING OF CONTENTS

DECLARATION OF AUTHORSHIP

I, **Anthony Michael Strickland**, declare that the thesis entitled:

"Enhanced Pre-Clinical Assessment of Total Knee Replacement Using Computational Modelling with Experimental Corroboration & Probabilistic Applications"

and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published or submitted for publication as:

  "Could Passive Knee Laxity be Related to Active Gait Mechanics? An Exploratory Computational Biomechanical Study Using Probabilistic Methods". M. Strickland, M. Browne and M. Taylor, CMBBE *(In Press)*.

  "In-silico Wear Prediction for Knee Replacements - Methodology and Corroboration". M.Strickland and M. Taylor J.Biomech *(In Press)*.

Signed: ……

Date: ……

ACKNOWLEDGEMENTS

Recognition must first be given to those who have contributed directly by their advice and experience to the work in this thesis, in particular my supervisors, Professor Mark Taylor and Doctor Martin Browne, as well as the postdoctoral researchers and my fellow postgraduates within the Bioengineering Sciences Research Group at Southampton. Similarly, I would like to thank my secondment hosts at the University of Kansas: Professor Lorin Maletsky and his research team in the Experimental Joint Biomechanics Research Group. More broadly, thanks to those from other institutions whose information and suggestions have been so helpful, especially the researchers at the Universities of Leeds and Denver.

We have benefited much from developing closer knowledge-sharing ties with industrial researchers, and particular thanks are due to the researchers at DePuy Orthopaedics, in particular to Matt Dressler, Craig Ernsberger and Todd Render who have been a valuable source not only of experimental data, but of ideas, insights, inquiries and inspiration too.

Producing this thesis has been an educational experience in many ways; not least in negotiating the throes of those academic 'doldrums'! My warmest thanks are due to a fantastic group of friends and colleagues in the form of my fellow researchers in the bioengineering research team; whose grace, good humour and camaraderie helped me through some of the darker hours! So (in no particular order): Johnny, Adam, Ian, Suk, Lucy, Sara, Mav, Olly, Carolina, Alex, Catherine, Becky, Pramod, Hamid, Pete, Hatice, Meadhbh, Margaret & Francis, I hope I made it a bit more fun for you guys, too – and my sincere best wishes to each of you, wherever your future lives may lead – thank-you!

Finally, my warmest thanks to those people outside of the academic world who've been there for support and encouragement over the last few years, in particular my dearly-loved family, and a wonderful group of friends from Highfield church – along the way you've helped me to remember what's really important in life – and for that I am deeply grateful.

Funding for this research was provided by the Engineering & Physical Sciences Research Council (EPSRC), as well as DePuy, A Johnson & Johnson Company. Some additional materials (data sources & CAD geometry) were also provided by DePuy.

GLOSSARY OF TERMS

ACL	- Anterior Cruciate Ligament
ADL	- Activity of Daily Living (any regular subject activity; e.g. walking)
ADAMS	- Automatic Dynamic Analysis of Mechanical Systems (MBD Software)
AMTI	- Advanced Mechanical Technologies, Inc (company name)
A-P	- Anterior-Posterior (translational direction)
Ad-Ab	- Adduction-Abduction (rotational axis of knee; see also V-V)
Arthroplasty	- lit. '*arthron*-' (joint) + '*-plastia*' (moulding); hence 'joint surgery'
BW	- Body Weight; normalising term for joint forces (typically ~800N)
C-D	- Compression-Distraction (knee motion: see also I-S)
CDF	- Cumulative Density Function (in probability modelling)
Co-Cr	- Cobalt-Chromium (metal alloy common in orthopaedics)
COV	- Coefficient of Variation (measure of statistical accuracy)
CR	- (PCL) Cruciate-Retaining (design option for knee implants)
CS	- Cross-Shear (crossing motions for knee contact mechanics)
DD	- Displacement-Driven (knee simulator test method)
DOE	- Design of Experiment
DOF	- Degree (or Degrees) Of Freedom
EF	- Elastic Foundation (computational method for modelling contact)
FB	- Fixed Bearing (design option for knee implants)
FD	- Force-Driven (knee simulator test method)
F-E	- Flexion-Extension (primary rotational axis of knee)
FE(A)	- Finite Element (Analysis); computational method
FPI	- Fast Probability Integration (stochastic techniques)
HIKIN	- High-Kinematics (knee simulator DD gait profile)
HS	- Heel-strike (beginning of stance phase in gait cycle)
I-E	- Internal-External (secondary rotational axis of knee)
in-vitro	- Experimental (lit. 'in glass', used as synonym for ex-vivo)
in-vivo	- Clinical (lit. 'in body', antonym ex-vivo – out-of-body)
in-silico	- Computational (permutation of *in-silicio*; lit. 'in silicon')
I-S	- Inferior-Superior (translational direction)
ISM	- Importance Sampling Method (statistical modelling method)
JCF	- Joint *Contact* Force – *internal* force experienced at joint surface
JRF	- Joint *Reaction* Force – *external* force transmitted by joint
KKS	- Kansas Knee Simulator (servo-hydraulic knee rig)
KU	- University of Kansas
LCL	- Lateral Collateral Ligament (also Fibular Collateral Ligament)
LHS	- Latin Hypercube Sampling (statistical modelling method)
LUT	- Look-Up Table
MBD	- Multi-Body Dynamics (software modelling methods)

MCL	- Medial Collateral Ligament (also Tibial Collateral Ligament)
MCS(T)	- Monte-Carlo Simulation (Technique)
MCycle	- Mega-Cycle (1 million cycles; used in long-term wear tests)
MIS	- Minimally Invasive Surgery
M-L	- Medial-Lateral (translational direction)
MoP	- Metal on Polyethylene (articulation for knee implants)
MPP	- Most Probable Point (for FPI methods)
(A)MV(+)	- (Advanced) Mean Value (+) (family of FPI Methods)
NESSUS	- Numerical Evaluation of Stochastic Structures Under Stress
OA	- Osteoarthritis (synonyms: Osteoarthrosis, Arthrosis)
(F/S)ORM	- (First/Second) Order Reliability Method; FPI methods
PCL	- Posterior Cruciate Ligament
PDF	- Probability Density function (also *Distribution* function)
PE	- Polyethylene (of which UHMWPE is a particular form)
PFC	- Press-Fit Condylar (J&J TKR design)
PID	- Proportional-Integral-Derivative (control scheme)
PKS	- Perdue Knee Simulator (servo-hydraulic knee rig; basis for KKS)
PL	- Patellar Ligament (occasionally referred to as patellar tendon)
PMMA	- Polymethylmethacrylate (principal constituent of bone cement)
PS	- PCL Substituting/Sacrificing (design option for knee implants)
QT	- Quadriceps Tendon
RA	- Rheumatoid Arthritis (systemic form of arthritis)
RMS	- Root-Mean-Square (convenient measure of vector signals)
(A/P)ROM	- (Active / Passive) Range of Motion
RP	- Rotating Platform (design option for knee implants)
RSE	- Response Surface Equation
RSM	- Response Surface Method (FPI Method)
S/C	- Semi-Constrained (specific tibial insert design)
SA	- Sensitivity Analysis (in statistical modelling)
SD	- Standard Deviation (statistical measure; also denoted by σ)
SKS	- Stanmore Knee Simulator (wear simulator rig)
T(J/K)A	- Total (Joint / Knee) Arthroplasty
T(J/K)R	- Total (Joint / Knee) Replacement
TO	- Toe-Off (end of stance phase in gait cycle)
V-V	- Varus-Valgus (secondary rotational axis of knee)
U/C	- Un-Constrained (specific tibial insert design)
UHMWPE	- Ultra-High Molecular Weight Polyethylene
UKR	- Unicompartmental Knee Replacement
UMKC	- University of Missouri (Kansas City)

SUMMARY

The purpose of this thesis is to demonstrate that computational and experimental methods can be used together more effectively to provide an enriched pre-clinical analysis toolset for design of knee replacements, and further to show that accounting for variability (by using probabilistic methods) plays an essential role in developing a more holistic understanding of knee mechanics.

With rising life expectancy, joint problems are increasingly common in the developed world. In recent decades, hip and knee surgery has become commonplace, with millions of procedures now performed annually worldwide. However, several factors limit the effectiveness of joint replacement surgery, especially for the knee: high inter-patient variability, rising expectations from surgery and imperfect understanding of the mechanics, biology & tribology of the joint. ***Chapter One*** begins by discussing the biology, pathology and intervention options related to knee replacement in more detail.

In consequence of these challenges, considerable effort goes into analysis and design of new implants. To avoid expensive and risky clinical (*in-vivo*) trials, there is an increasing emphasis on pre-clinical testing, using experimental (*in-vitro*) and computational (*in-silico*) methods. A review of the relevant literature and historical developments in this area is presented in ***Chapter Two***. Important examples in the field of knee testing are experimental knee wear simulators, lower limb simulators, and finite element stress analysis. However, these are often used as either isolated computational studies, with inadequate experimental corroboration, or isolated experimental studies, lacking the enriched analysis and visualisation which computational modelling can provide. The need exists to use *in-vitro* and *in-silico* methods together more effectively, providing a holistic, data-rich means of assessing the many variables and uncertainties in knee biomechanics.

A major aim of the thesis was to demonstrate that 'single-shot' models alone are not adequate to fully account for the results observed *in-vivo* and *in-vitro*. Historically, many biomechanical studies have failed to account for the wide range of variability seen in such factors as patient activity or surgical positioning accuracy. Models which do simultaneously take account of the multiple different variable input factors are termed 'probabilistic' models. ***Chapter Three*** describes some of the existing probabilistic techniques that can be used to obtain a more

'holistic' overview of system performance, and considers how these methods might be applied to model the influence of variability for biomechanical systems.

Capturing the influence of variability in knee analysis is a key aim of the thesis. To facilitate the large number of simulations necessary to explore this variability, a faster modelling approach is needed. ***Chapter Four*** describes the early modelling efforts, exploring the use of multi-body dynamics models as a low-cost surrogate to replace deformable finite-element methods. These models can potentially deliver faster performance (simulation times of minutes rather than hours) with only a limited loss in accuracy. Early results demonstrated satisfactorily that MBD could be used as a fast surrogate for finite element methods, able to reproduce the kinetics & kinematics for force- and displacement-driven tests, and also approximate the finite-element predictions for contact pressure and wear. We used these fast rigid-body models to simulate multiple *in-vitro* wear tests, and so assess the performance of existing *in-silico* wear algorithms.

With these baseline deterministic models in place, ***Chapter Five*** describes how the models were coupled with probabilistic methods to begin exploring the influence of variability on knee biomechanics, firstly reproducing the results of existing studies in the literature, and then progressing on to novel investigations. However, none of the existing studies corroborated the probabilistic methods with corresponding probabilistic experimental data; therefore it was determined that better corroboration against specific experimental test platforms was needed, to serve as the basis for a corroborated probabilistic study.

The first targeted corroboration was for the Kansas lower-limb simulator, described extensively in ***Chapter Six***. This focused model design & verification exercise revealed the importance of accurately capturing the 'dynamics' of the simulator (e.g. damping effects and inertia), and also the influence the controller can have upon the overall system performance. As a result of this, it was recognised that in order to achieve accurate corroboration, a knee-wear simulator model needed to include the control system, and had to accurately represent all the mechanical dynamics of a specific rig design.

Therefore, the AMTI knee simulator was used as the basis of a much more robust set of *in-silico* models, as described in ***Chapter Seven***. Using this higher-fidelity model, forces and kinematics could be accurately predicted using experimental feedback data. With this deterministic model performing well, the final aim was to build a probabilistic study around this model, to investigate whether the variability predicted by *in-silico* models matched *in-vitro* test results (i.e. a first-of-

kind 'probabilistic corroboration'). Data from over one-hundred experimental trials was compared with a probabilistic computational study, to discover if the resulting variations in knee mechanics and wear corresponded for *in-vitro* and *in-silico* analysis methods.

The key outcomes of the thesis may be summarised as follows:

- Rigid 'multi-body-dynamics' based models were shown to perform acceptably as fast surrogates in place of finite-element models.
- Models of specific real-world experimental platforms were developed. Good agreement with *in-vitro* results was achieved (to within 5% averaged RMS errors), and the experimental data could be augmented with additional *in-silico* data (e.g. wear 'decomposition' to visualise contact pressure or cross-shear).
- The 'probabilistic' models including variability reveal that small variations in input test conditions can considerably alter the resulting outcomes; e.g. it is possible to achieve very high wear rates (2 to 3 times the normal level) even with small component misalignments of only a few degrees.
- The focus on computational wear modelling has led to developments in our understanding of wear; it has been robustly demonstrated that *in-silico* models can qualitatively rank designs in terms of wear performance, but there is still too much uncertainty to have quantitative confidence in the results to a high degree of accuracy.
- By combining probabilistic experimental and computational results for the first time, this thesis reveals important limitations of computational models when predicting the variation in real-world wear rates: *in-silico* methods seem to be under-estimating the true amount of observed variation by at least a factor of four (more for some older theoretical wear approaches).

This work has demonstrated that much faster *in-silico* models can still deliver acceptable accuracy. This speed increase can be harnessed to investigate variability, to perform multiple analyses, and so extend the domain of testing beyond what is possible with purely experimental means. But computational models must always be grounded in reality by robust corroboration with experimental methods, so close collaboration is essential. These studies demonstrate that close collaboration between computational and experimental specialists can yield benefits for both; the computational models are more

realistic and more accurate, and the experimental data can be augmented with enriched *in-silico* visualisation methods, and complemented with additional data from probabilistic studies (which would be too time-consuming to run experimentally).

In this thesis the combination of fast computational analysis with experimental data was used to investigate the current generation of wear concepts, and it is shown that whilst these existing wear formulae are useful and beneficial analytical tools, they are not perfect, and there is room for considerable improvement in the current theories of wear. The combination of experimental and computational methods, taking account of the important role of variability in the tests, has been the key to advancing our understanding of the capabilities and limitations of existing wear theories.

To progress our understanding of knee biomechanics further in the future, it will increasingly be necessary for different research specialists to collaborate together, in order to corroborate their methods and so build confidence in our results collectively, and enhance the total quantity and quality of pre-clinical data available to researchers, implant designers and clinicians. Equally, as a research community, emphasis must shift from 'single-shot' experimental and computational models of knee biomechanics to more 'probabilistic' approaches capturing the substantial variability observed in the knee *in-vivo*. The methods that have been used and the models that have been developed demonstrate what is possible when such a holistic approach to pre-clinical analysis is adopted.

CHAPTER ONE - REVIEW OF THE HUMAN KNEE

Anatomy, Physiology, Pathology & Clinical Treatments

1.1. Introduction: Motivation for Understanding the Human Knee

Advances in medicine, diet & living conditions have led to increases in life expectancy across the developed world. As a result, UN estimates predict that the percentage of the population aged 60 or over will have risen from about 10% in 2000 to over 20% in 2050 worldwide, with the European rate rising from 20% in 2000 to almost 35% in 2050 [1].

As a result, many of the physiological problems associated with older age are more prevalent. Joint problems are particularly common. Over decades of normal daily activity, the articulating surfaces experience damage leading to pain and hence reduction of mobility. Less frequently, joint problems can also be the result of progressive or congenital diseases affecting younger individuals as well as older subjects.

In order to improve patient quality of life, it is desirable to remove, or at least reduce, the pain and loss of mobility caused by joint degeneration. Various remedial procedures such as knee surgery (called 'knee arthroplasty') have been devised for this purpose. Historically, arthroplasty procedures have traditionally been associated with the lower limb joints (specifically, the hip and knee), which bear the largest loads and hence are most susceptible to osteoarthritis. The focus of the present body of work will be the knee joint, which now undergoes more procedures than any other single joint nationally [2]. A necessary pre-requisite for this study is therefore a preliminary review of the fundamental anatomy, physiology & relevant pathology of the knee joint, and the surgical options for intervention. This review is presented in this first chapter.

1.2. The Human Knee: an Anatomical Review [1]

The knee is the largest 'synovial' joint in the body (i.e. the joint is enclosed in a fibrous capsule, containing synovial fluid). Although often referred to as a 'ginglymus' (simple hinge) joint, it is in fact a complex multi-condylar joint, with secondary motions including considerable anterior-posterior (A-P) translation, and internal-external (I-E) rotation. Technically, it is not one single joint; there is a patellofemoral articulation, and two distinct tibiofemoral articulations (both

[1] *Extensive use is made in this document of the anatomical frames of reference. These are defined in Appendix A. The material presented in this section is derived from standard human anatomy texts [3, 4].*

medial and lateral condyles). It may therefore be considered as three 'compartments', and in this sense a 'total' knee replacement may be referred to as a 'tri-compartmental' replacement.

The knee experiences very demanding mechanical loads, since most of the body weight acts through the joint and large torques are present due to the large thigh & shank moment arms. Furthermore, unlike the hip, the knee is inherently unstable, so additional soft tissue forces (from surrounding muscles and ligaments) are required to stabilise the joint. The following anatomic description progresses from the skeletal components of the knee to the more superficial soft tissues:

Skeletal

The ***femur*** is the principle bone of the 'thigh' (or upper leg); the 'shank' (or lower leg) consists of both ***fibula*** and ***tibia***, although the tibia is the principle load-bearing structure. The femur & tibia articulate together directly (two convex condyles on the distal epiphysis of the femur articulate with the superior surface of the proximal tibial condyles), thus forming the ***tibiofemoral*** joint. Additionally, the anterior 'groove' of the distal femur also articulates with the ***patella***, (the 'kneecap') a 'sesamoid' (intra-tendonous) bone providing attachment for, and improved leverage to, the quadriceps muscles (Figure 1, left) – this forms the ***patellofemoral*** joint. The area of the bones where contact occurs is covered with a thin layer of ***articular cartilage***, a collagen-based soft-tissue which provides impact-damping and reduces joint friction. It is the deterioration of this cartilage protection which often leads to joint failure (see Section 1.4).

Menisci

In both of the tibiofemoral condyles, a cartilage ***meniscus*** is present, which reduces joint friction, distributes loads to reduce local contact stresses, and provides further impact protection to the joint. The menisci are located over the lateral & medial condyles of the tibia, connected posteriorly by a transverse ligament, and to both the femur & tibia by additional ligamentous attachments.

Synovial Membrane

The articulating region is enclosed by a synovial membrane, containing the synovial fluid which assists in lowering joint friction and providing fluid ingress for nutrient supply to the cartilage.

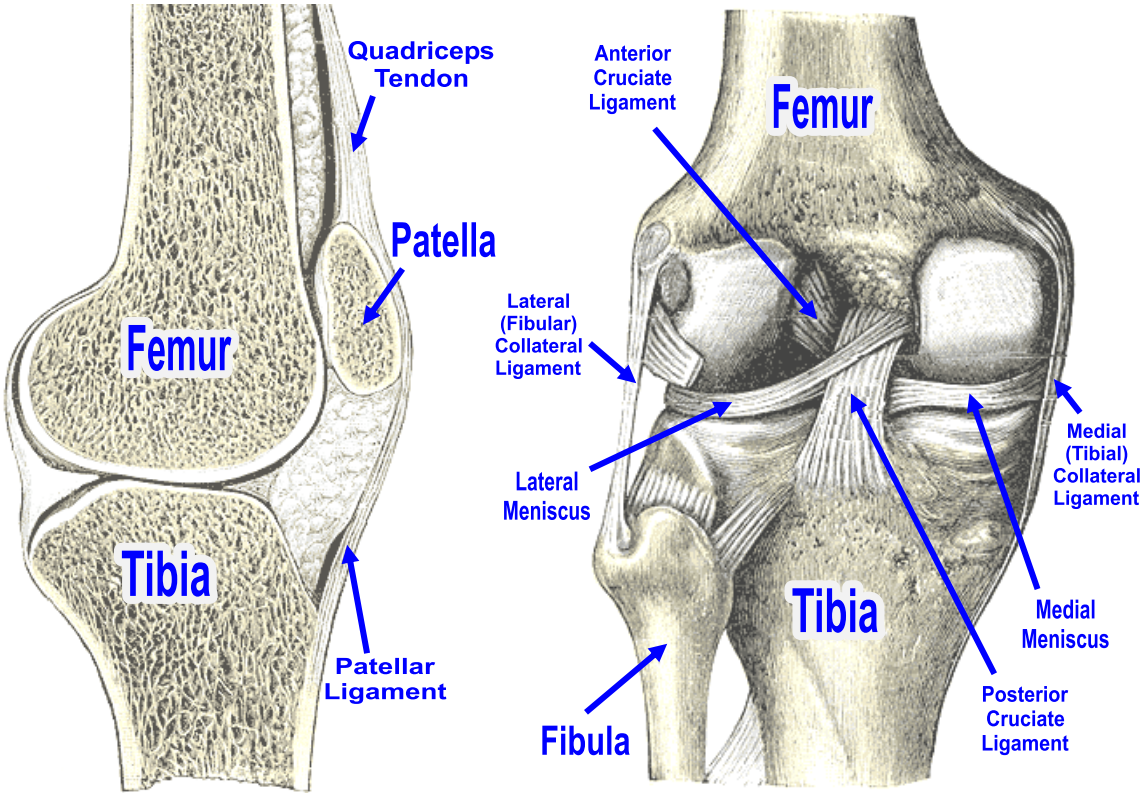

Figure 1: Sagittal cross-section (left) & posterior view (right) of the knee (from [3]).

Fibrous Capsule

An extensive fibrous capsule surrounds the entire joint, blending with the surrounding tendons and ligaments, providing additional protection and soft-tissue restraint.

Ligaments & Tendons

The patella is embedded within a tendonous link between the ***tibial tuberosity*** (on the anterior aspect of the proximal tibia), and the different muscles which form the ***quadriceps*** group. The (inferior) tendonous link between the tibia and patella is called the ***patellar ligament*** (PL), while the (superior) link between the patella and the quadriceps muscles is the ***quadriceps tendon*** (QT). Embedded within this tendonous link, the patella provides increased leverage for the quadriceps muscles; in deeper flexion angles the quadriceps wraps over the anterior surface of the distal femur (quadriceps 'wrapping'). The patella, articulating in the patellar groove on the anterior aspect of the distal femur, controls the line of action of the quad muscle forces, and by increasing the moment arm, increases the magnitude of the extension moment which the quadriceps can generate at the knee.

Ligaments form an essential part of the tibiofemoral joint. The knee is stabilised by four main ligaments: two ***cruciates*** (anterior & posterior) and two ***collaterals*** (medial & lateral), abbreviated ACL, PCL, MCL & LCL respectively (Figure 1, right).

The MCL is also called the ***tibial collateral*** ligament (as it inserts distally to the tibia), while the LCL is the ***fibular collateral*** ligament (as it inserts to the fibula).

The function of the ligaments is to constrain the kinematics of the knee, increasing the stability of what is an inherently unstable joint (it should be noted that the fibrous capsule and the surrounding muscle & tendon tissues also provide additional soft-tissue constraint). The collateral ligaments are recruited when resisting larger rotational motions; e.g. I-E and varus-valgus (V-V), since they have a larger moment arm against such torques. The cruciate ligaments are particularly important for their role in guiding A-P translation through different flexion angles. In reality, ligaments are complex multi-bundle structures, with different origins & insertions, different mechanical properties (stiffness & tensile strength) between bundles [5], and a non-linear behaviour (due to the fibrous structure 'crimping' when relaxed), including differing levels of 'pre-strain' (the degree of pre-tensioning the ligament experiences in the 'neutral' stance position). A typical ligament load-extension response is shown in Figure 2. Various studies have demonstrated that ligament properties vary considerably between different subjects [6], and that the precise configuration of ligament bundles is important in determining the overall ligament behaviour [7].

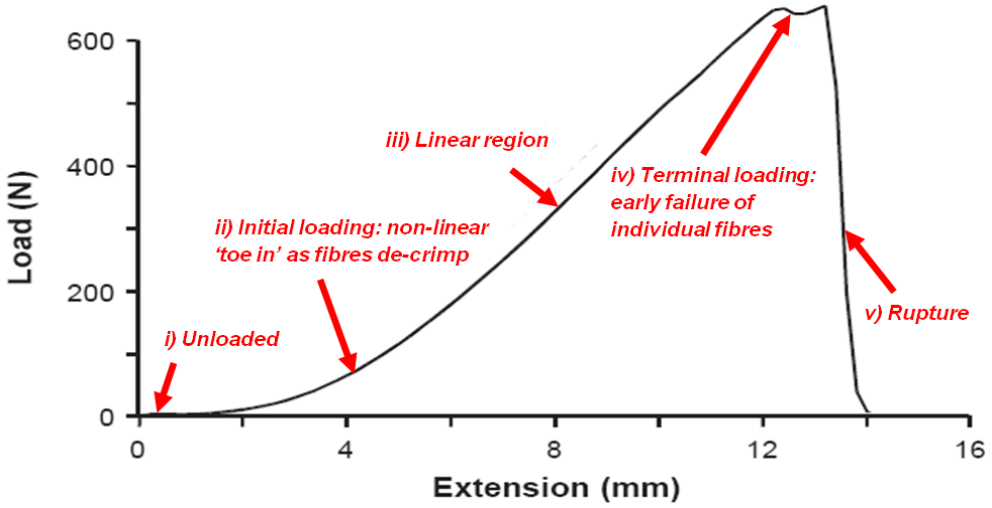

Figure 2: Typical load-extension response (here: MCL bundle, adapted from [5]).

Muscle Groups

The most notable muscles are those responsible for sagittal-plane knee flexion (the ***hamstrings***: ***biceps femoris***, ***semimembranosus*** & ***semitendinosus***) and extension (the ***quadriceps***: ***rectus femoris*** and the ***vastus*** muscle group: ***v.mediales v.intermedius & v.laterales***). Other muscles also play a part in flexion-extension (F-E), as well as I-E rotation (see Table 1). The muscles included in this table are also depicted graphically in Figure 3. Note: for clarity, this figure only includes those muscles directly surrounding the knee joint; as such, some

other major lower-extremity muscle groups (e.g. ***gluteus*** & ***soleus***) are excluded; in reality, there is of course always an interdependence between the role of different muscle groups during different activities, and the full musculature of the lower limb must be considered as a single system for dynamic analysis.

Muscle Group	Muscle	Function
'Hamstrings'	*Biceps Femoris*	Flexion, External Rotation
	Semimembranosus	Flexion, Internal Rotation
	Semitendinosus	Flexion, Internal Rotation
'Quadriceps'	*Rectus Femoris*	Extension
	Vastus Intermedius	Extension
	Vastus Laterales	Extension
	Vastus Mediales	Extension
Plantar-flexors	*Plantaris*	(Secondary) Flexion
	Gastrocnemius	(Secondary) Flexion
Secondary flexors	*Gracilis*	Flexion, Tibial Internal Rot.
	Sartorius	Flexion
	Tensor Fascia Lata	Stability (extra tension)
	Popliteus	Flexion, Internal Rotation

Table 1: List of Functional muscles associated with the knee joint.

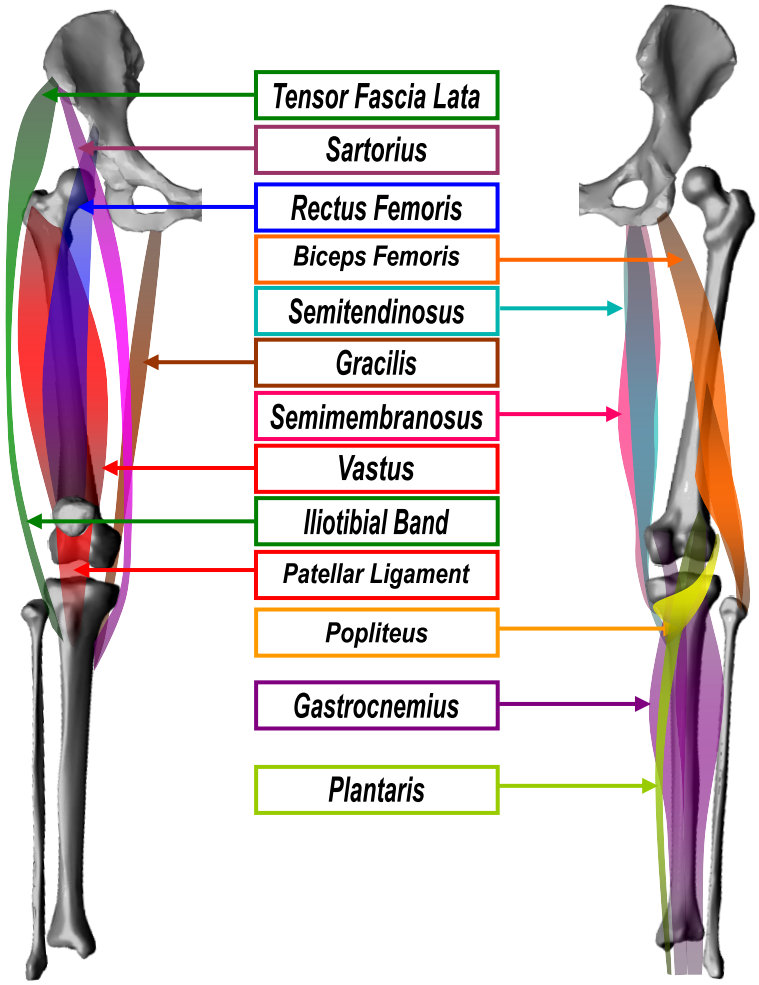

Figure 3: Muscles around the knee joint (left: anterior view, right: posterior view).

1.3. Knee Mechanics: Kinematics & Kinetics

1.3.1. Motions of the Knee

It is possible for the knee to move to a limited extent in any of the six possible degrees of freedom (3 translations, 3 rotations). F-E is by far the most visually apparent rotational action; however considerable I-E and V-V rotation is possible. The translational motions are less apparent, although several millimetres of A-P and medial-lateral (M-L) displacement is possible, and condylar 'lift-off' may result in slight compression-distraction (C-D) displacements. A number of specific issues related to knee kinematics are briefly outlined below:

Range of Motion at the Knee

It is difficult to define a 'typical' range of motion (ROM) for the knee joint, for two reasons; firstly, inter-patient variability means that this envelope would not be the same for any two subjects. Secondly, the degree of motion achieved is dependent upon the loads applied to the knee. This has lead to a distinction being made between the 'active' and 'passive' ROM (abbreviated AROM and PROM respectively) - i.e. whether the motion is made under the subject's own muscle action, or whether external manipulation is used to achieve the motion. Clinically, AROM is reported to average ~130°, decreasing with age. PROM is higher (~160°, again decreasing with age) [8, 9]. Terminologically, flexion angles over 90°, and especially those beyond 120°, are often referred to as 'deep flexion' (not required for general ambulatory activities, but required for some kneeling & squatting everyday activities, such as gardening, domestic cleaning or kneeling prayer). Facilitating this 'deep flexion' ROM is a key goal for next-generation TKR designs.

Knee 'Locking' and Screw-Home

Although the knee is unstable, several effects combine to increase stability in full extension. The distal radius of the femoral condyle is larger than the posterior radius, thus increasing conformity in full flexion. For normal subjects, the line of action of body-weight is slightly anterior to the tibiofemoral contact when in full knee extension, tending to maintain the knee in extension. This is accompanied by an internal rotation of the femur relative to the tibia, causing the surrounding soft tissues to tighten, resulting in a higher degree of stability. This 'locked' stance state is released when the popliteal muscle contracts, causing the femur to rotate externally relative to the tibia and so reducing the soft tissue constraint prior to the knee flexing (see Figure 4). This mechanism for increasing stability in full extension is often referred to as the 'screw home' effect [10].

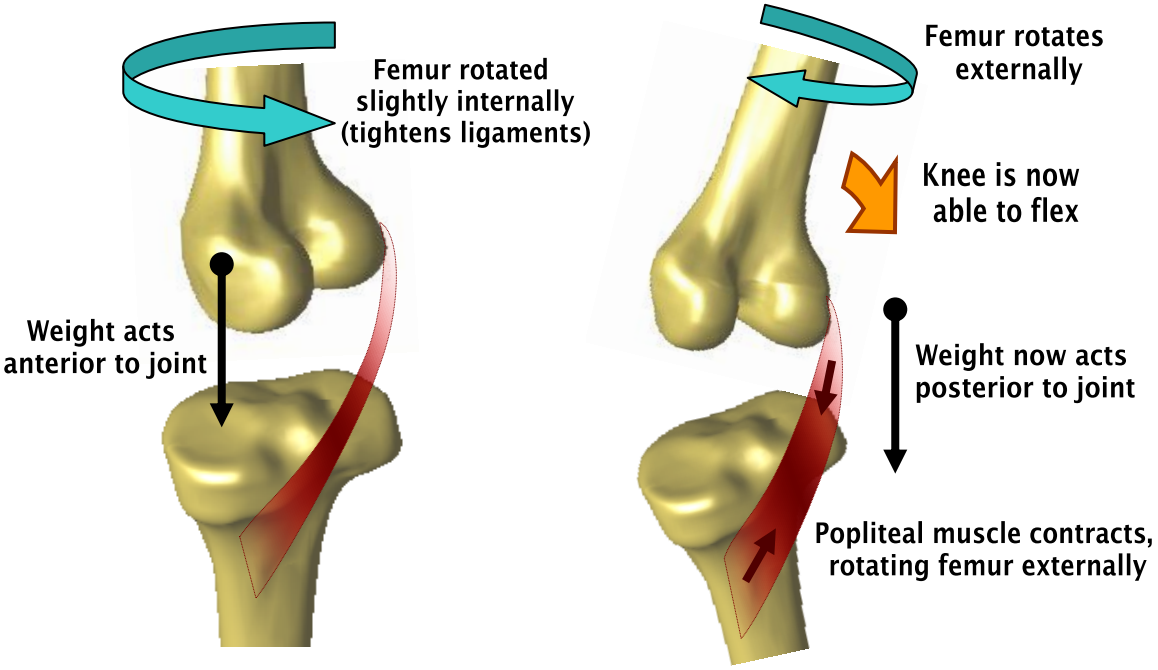

Figure 4: Knee 'locked' in screw-home position (left) and released in flexion (right).

Femoral Rollback

The term 'femoral rollback' refers to a posterior movement ('roll-back') of the femur (or conversely, an anterior movement of the tibia) as the knee flexes (and therefore vice-versa, as the knee extends). This concept first gained currency at the beginning of the 20th century [11], and was subsequently analysed using rudimentary four-bar-linkage models of the knee (Figure 5); both the femoral axis of rotation and the tibiofemoral contact point are predicted to move posteriorly as flexion increases, according to these simple rigid-linkage predictions. The concept became the subject of some debate within the orthopaedic research community, with studies both confirming and refuting the femoral rollback phenomenon. Consensus is building that femoral rollback is not apparent for the medial condyle, but is often observed at the lateral condyle under passive loading [12]. Recent fluoroscopy studies (e.g. [13]) reveal the situation during active loading (i.e. when the knee is subject to large muscle loads during daily activities) to generally be much more variable [14]. Further, there are differences between natural and implanted knees, with the latter more likely to exhibit 'paradoxical' anterior femoral motion with flexion (particularly in deep flexion) [15]. Finally, it is important to distinguish between the movement of the two bones (defined by hard anatomical landmarks), and the movement of the *contact point* between the bones; it is possible to have 'paradoxical' motion of the contact point relative to the motion of the two bones.

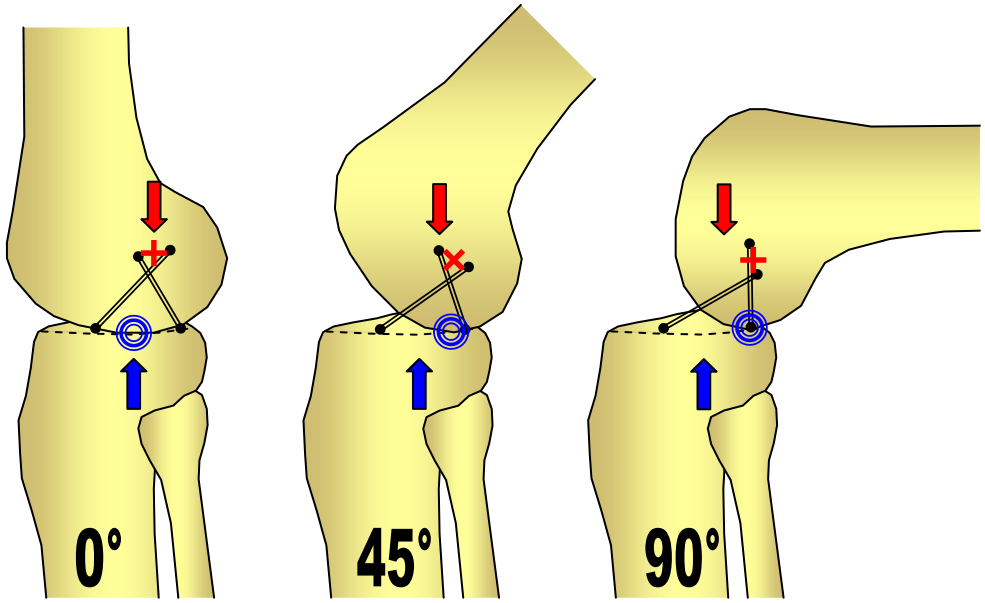

Figure 5: Simple 2-D (sagittal) linkage model, showing the 'femoral rollback' concept.

Medial Pivot

The 'medial pivot' concept is related to the concept of femoral rollback. It is widely reported that the femur tends to rotate externally as the knee flexes (i.e. the tibia rotates internally relative to the femur). This, coupled with the hypothesised posterior motion of the femur during femoral rollback, would result in a combination of rotation and translation about the long axis of the bones, which could equivalently be represented by a single rotation (with no corresponding translation) about a 'virtual' pivot point shifted towards the medial condyle (see Figure 6). Note that the 'medial pivot' concept is dependent upon the 'femoral rollback' assumption, and so the caveats associated with that concept apply equally to the medial pivot hypothesis. If paradoxical motion occurs, the virtual pivot will not be medially-shifted. Once again, inter-subject variability is considerable, and there is no single 'correct' description of the medial pivot effect; however it is widely reported within the literature [16].

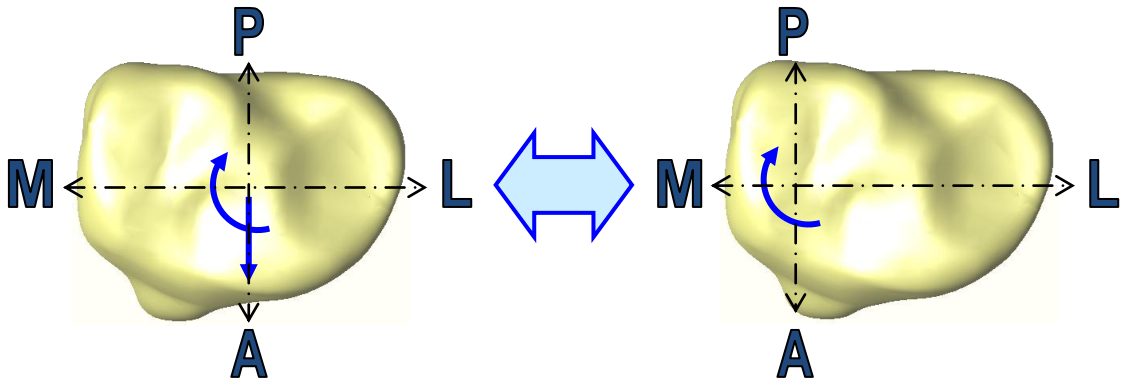

Figure 6: The 'medial pivot' concept (illustrated on the tibia): rotation plus translation (left) is equivalent to rotation about medial condyle (right).

Describing Knee Motions: The Grood & Suntay System

The multiple degrees of freedom and complex motions at the knee mean that kinematics can be complex, so kinematics must be defined clearly and reported consistently to avoid ambiguity or confusion. An important and widely-adopted method was proposed by Grood & Suntay [17]. In this cylindrical-axis co-ordinate system, the sequence in which the different rotations and translations are applied does not alter the final position & orientation (i.e. the system is sequence-independent; this is an important advantage over e.g. the Euler co-ordinate system); see Figure 7. Although intended for natural knee motions, the Grood & Suntay system is equally applicable for *in-vitro* lower-limb simulators.

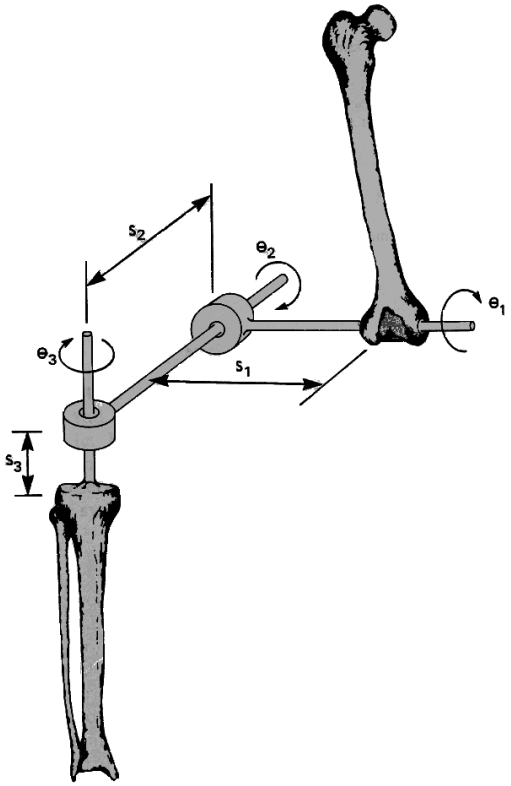

Figure 7: Grood & Suntay co-ordinate system: graphical illustration (from [17]).

1.3.2. Kinetics of the Knee

For the human knee, loading varies from subject to subject depending upon the activity or mix of activities of daily living (ADL), and there is considerable inter-patient variability. Experimental studies have demonstrated that the most important common activities considering both loading *and* frequency are walking ('gait') and stair usage [18]. Other activities (e.g. sitting & lying down) may be more prominent in terms of duration or frequency of occurrence but place limited dynamic loading demands on the knee; conversely some highly demanding activities (e.g. 'shock' loads due to tripping or stumbling) may result in greater loads, but occur only very rarely.

It is important to characterise both the *variety* (range) and *volume* (frequency) of activities the joint will be subjected to. Basic data like step rate measurements can be obtained using pedometers or foot-switches, whereas the relative mix of activities can be extrapolated based on observations across a short time period. Inter-patient variability means that generic assumptions are rarely applicable; for example, 'typical' subject step rate is often approximated as a million steps per year (when a 'standard' is required for testing), but for a diverse sample of healthy subjects, Seedhom & Wallbridge reported an average of some 1.8 million steps/year per joint [19]. A study by Schmalzried et al focussed on arthroplasty patients, and found an average of 0.9 million steps/year per joint, but this average masked a wide *variability*, with outliers ranging from just over 70,000 to as high as 3.2 million steps/year per joint [20]. Clearly a single 'representative' figure has only limited practical meaning.

For common ADL types, knee mechanics can be recorded or estimated by various methods, including clinical motion analysis using video recording (or, more recently, fluoroscopy studies – e.g. [21]) & force plates (for external joint reaction forces), coupled with optimization algorithms (based on inverse dynamics methods) and/or EMG data (for internal joint contact forces). Rarely, more 'invasive' assessment methods have been used; e.g. Lafortune et al used markers fixed with traction pins directly into the bone [22]. The data collected by these studies can be used as the basis for input waveforms into a simulated knee model, and 'standardised' waveforms have been devised for comparative testing of implant designs (e.g. ISO testing standards [23, 24]).

Often-cited examples of these studies are the early work by Morrison [25] for ambulatory gait, and Andriacchi et al [26] for stair climbing. Other studies have reported for more demanding ADLs; e.g. 'deep flexion' squats (Nagura et al [27]). More recent studies have included a larger number of subjects, giving some indication of inter-patient variability; e.g. McFadyen & Winter, who recorded both stair ascent and descent [28], and the studies for gait and stair activates by Costigan et al [29]. More recently, telemetric measurements using prosthetics with embedded sensors have provided direct *in-vivo* data to compare with the theoretical results of earlier investigators; first for the hip joint (as pioneered in the early 1990's by Bergmann et al [30]), and subsequently for the knee, since the late 1990's (notably studies by Taylor et al [31-33] for a distal femoral implant, Kaufman et al [34] and most recently D'Lima et al [35-37] for an instrumented tibial tray). Together, this large body of work provides a picture of the kinematics and kinetics of the knee joint for a range of ADLs – this information is invaluable

for subsequently modelling the knee, as it provides the 'raw data' to drive knee simulations in a physiologically-representative manner.

It is important to make a clear distinction between the *internal* forces acting between the contacting joint condylar surfaces (often termed joint contact force, or JCF), and the *external* resultant forces experienced by the whole limb segments (termed joint reaction force, or JRF). By necessity of Newtonian mechanics, the static magnitude of the external JRF will be of the same order as the subject's bodyweight (BW), (although dynamic external forces can exceed 1BW due to accelerating/decelerating forces in locomotion). The internal JCF can be much higher however even under static conditions (often several times BW), since antagonistic muscular co-contraction (necessary to stabilise the joint) are considerable.

At the knee, forces are not loaded equally between condyles; the medial condyle will typically carry a larger load. It is also larger, however, to balance contact pressures. Surface contact pressures reflect the distribution of compressive joint forces across the surface of the femoral and tibial condyles. Generally, in a natural knee, the combination of low-stiffness articular cartilage on the condyles with the load-distributing effects of the meniscal cartilage results in low contact pressures when compared to the more rigid materials used in artificial implants. Although measuring knee contact pressures *in-situ* is challenging, several *in-vitro* cadaveric studies have attempted this using pressure-sensitive dyes or transducers. Depending on activity, typical *mean* condylar surface contact pressures have been found to be below 2MPa, with *peak* values around 6MPa [38] (a simple order-of-magnitude consideration would anticipate this, since forces of a few thousand Newtons are acting on an area of between 10-15cm²). Naturally, contact pressures rise with higher loads, or when the contact area is reduced (e.g. after a meniscectomy). Unfortunately the relatively stiff synthetic materials used for artificial prostheses result in much smaller contact areas, and are known to result in contact pressures several times higher than this (often approaching 20MPa [39]).

1.3.3. Mechanics of Normal Gait

Whilst the knee is used in many different ADLs, it is the most mechanically demanding which are of interest here, as these contribute most to mechanical failure. Of particular importance are the conditions during normal walking (i.e. 'active gait'). This is because, although gait does not result in the most extreme forces or kinematics, it is by far the most prevalent daily 'active' ADL for most typical arthroplasty patients. For example, a study by Morlock et al [18] found

that for hip patients, walking accounted for over 10% of the monitoring time – although this is low compared to some of the 'passive' activities (e.g. sitting, lying down), it was much higher than other high-loading activities (e.g. more than 25 times more frequent than stair climbing). As such, analysis of gait receives considerable attention in the literature; therefore the mechanics of 'normal' gait will be reviewed in further detail in this section.

Knee Flexion (Kinematics)

The flexion of the knee in gait is the most apparent kinematic feature, with a very clear & intuitive functional basis. Flexion of the knee serves two primary purposes: To provide 'shock absorption' damping as the limb is loaded immediately after heel strike (HS), and to ensure adequate clearance of the foot above the ground during swing phase. These two requirements lead to a characteristic *biphasic* waveform, as shown in Figure 8.

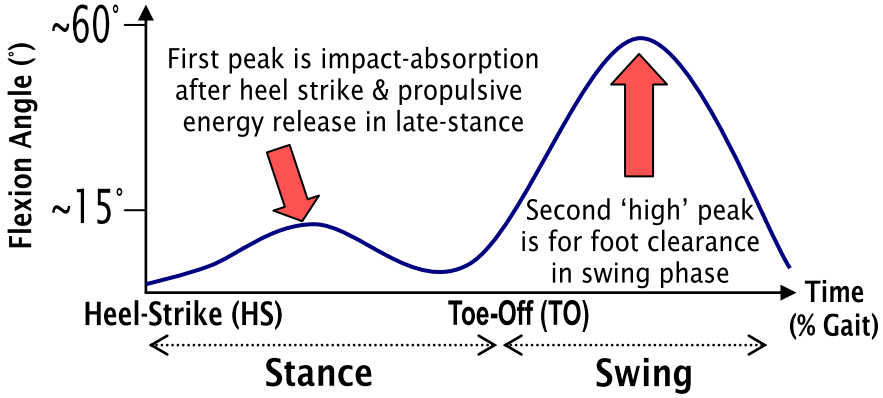

Figure 8: Typical knee flexion for normal gait.

The first peak is smaller, and quite variable between subjects. A typical knee flexion angle for this first peak is around 10-20°. The second peak is much larger, typically around 60° of flexion.

Note that the knee will not necessarily achieve full extension (i.e. 'flexion' = 0°) at any point in the gait cycle. For some subjects, the knee will remain in slight 'positive' flexion, even throughout the stance phase. This is subject-specific; e.g. the subject telemeterised by D'Lima et al [35] never exhibited under ~13° flexion in gait. Depending on the subject, the knee may remain always in slight (positive) flexion, or conversely achieve slight hyperextension (e.g. [22]), in gait.

Axial Compressive Force Loading (Kinetics)

During bipedal motion the lower limbs alternately support the weight of the body while in contact with the ground; therefore, to a first approximation, the knee bears a 'high' load in stance, and a 'low' load in swing. However, the actual

loading is more complex (Figure 9). Inverse-dynamics analyses predict a 'double-peak' during stance phase (corresponding to vertical acceleration of the trunk); this is sometimes observed *in-vivo* (e.g. [40]), although not consistently. Loading in swing phase is not 'zero', due to the passive restraint provided by soft tissues, and antagonistic muscle action. For the purposes of *in-vitro* tests it is often assumed that the swing phase load is constant (a few hundred Newtons), but telemeterised data reveal considerable variation of the load even within swing phase, with the lowest loads around mid-swing [35]. Antagonistic co-contraction of the muscles around the knee means that JCFs are higher than corresponding JRFs. However, whereas the early muscle-optimisation algorithms (e.g. Morrison [25]) anticipated JCFs of 3-4BW during gait, latest *in-vivo* measures suggest actual values may not be much higher than 2-2.5BW [35, 40] (see telemeterised waveform in Figure 10).

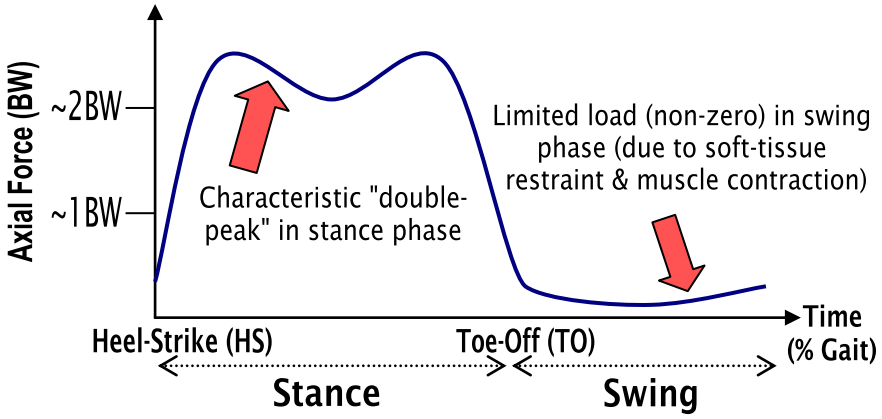

Figure 9: Typical theoretical axial JCF during gait.

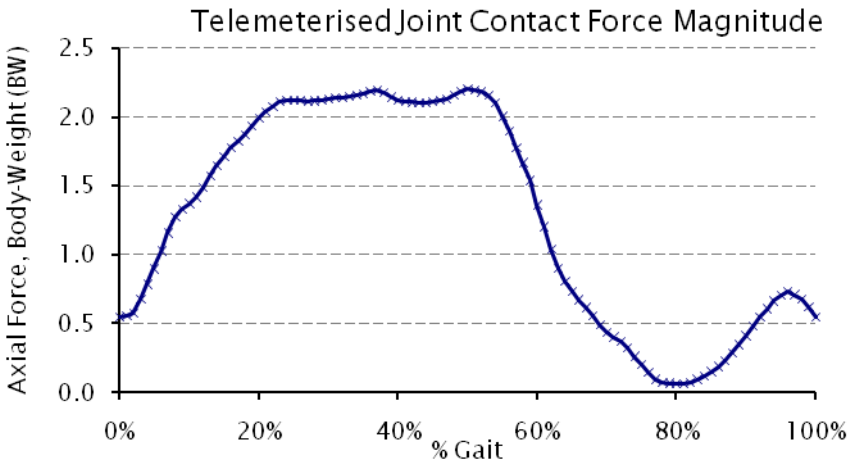

Figure 10: Telemetry data for axial JCF (adapted from D'Lima et al [35]).

Internal-External Torque & Rotation (Kinetics & Kinematics)

Early experiences with hinged knee prostheses demonstrated that the I-E kinematics and kinetics are an important characteristic of normal bipedal gait, and cannot be neglected; (early 'fixed-hinge' prostheses often failed because of

the high rotational torques occurring [41]). This I-E action helps to establish favourable trunk orientation for the proceeding step. Since the 'stance' foot is fixed on the ground, the I-E moment to twist the trunk must be generated across the lower limb. Given that the moment on the trunk is *external*, the reaction moment must be an *internal* moment; hence at the knee the proximal side (femur) experiences an external moment, whereas the distal (tibia) experiences an internal moment. This is seen in clinical gait assessment; a large torque peak (typically several Newton-metres) is seen in late stance phase (see Figure 11). The torsional effect acts to cause an internal rotation of the tibia relative to the femur (or, conversely, an external rotation of the femur relative to the tibia).

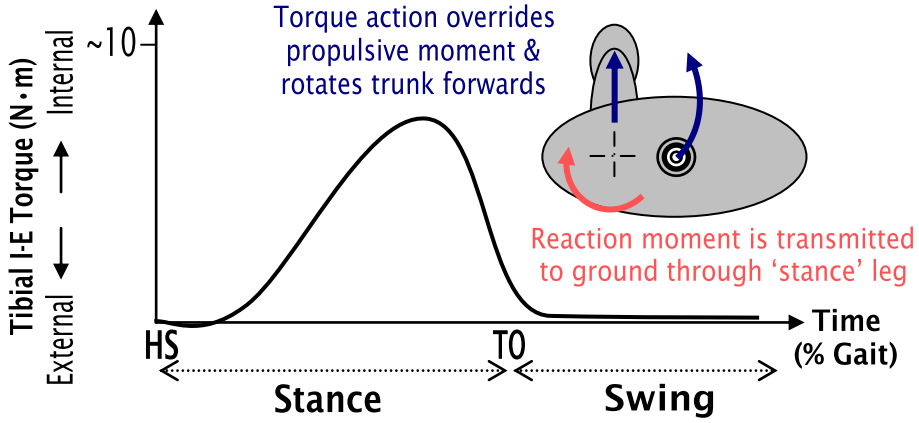

Figure 11: Typical I-E torque acting at the knee. The torque is principally 'external' on the femur, and hence 'internal' on the tibia.

Anterior-Posterior Force & Translation (Kinetics & Kinematics)

A-P forces and motions are important for TKR performance (and for defining the input conditions for force- and displacement-driven knee simulators – see Chapter Two). However, there is little consensus on the A-P forces or motions at the knee, owing to inter-subject variability (as discussed in Section 1.3.1 regarding the femoral rollback and medial pivot concepts). The A-P shear forces are known to be of considerable magnitude (Taylor et al reported peak A-P loads of 0.5BW [40]; D'Lima et al reported loads of over 0.3BW [35, 36]). However, there is very little consensus on the 'shape' of this shear-force profile. Even with the 'standardised' ISO waveform [23], the polarity is inverted by some testers (e.g. [42]).

For A-P kinematics, the ISO-standard adopts a predominantly anterior motion of the femur on the tibia [24], (in line with the findings of Lafortune et al [22]), but other groups [42, 43] have adopted a predominantly posterior femoral motion, in keeping with the medial-pivot hypothesis (Figure 12). Clearly, in light of this lack of consensus, further *in-vivo* fluoroscopic & telemetric studies are required for larger cohorts, to better understand how A-P motions vary between subjects.

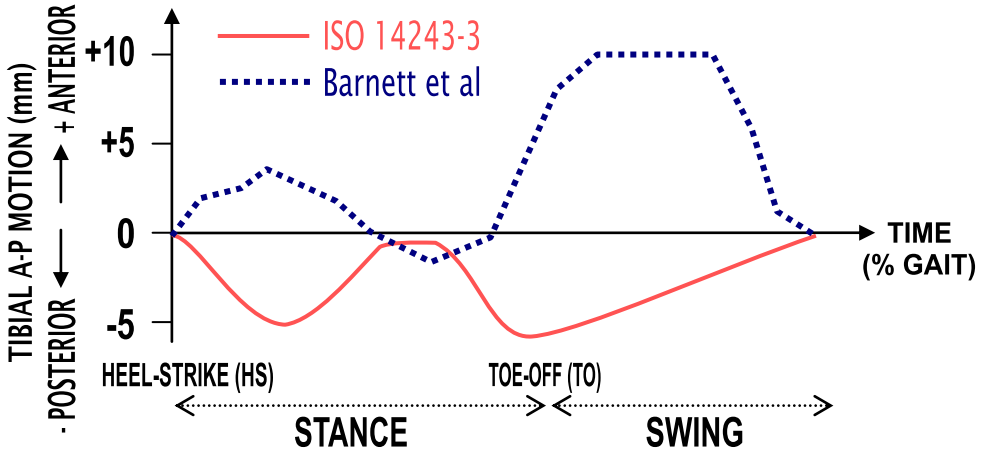

Figure 12: Different A-P motion profiles: ISO-standard [24] with 'paradoxical' anterior femoral rollback (solid), and from Barnett et al [44] with posterior rollback (dashed).

In the above sections, the anatomy and mechanics of a normal *healthy* knee have been discussed. However, TKR is only required when the knee ceases to function correctly. It is next necessary to consider how the knee joint can 'fail' & hence come to require intervention.

1.4. Pathology & Failure of the Knee Joint

For the overwhelming majority of cases where some form of clinical intervention is required, the cause is some form of *arthritis* (literally meaning 'joint-inflammation' in the Greek). Note that arthritis is not a *causal* diagnosis; the definition is based on the symptoms rather than any specific cause. Generally, the cause of this pain and inflammation of the joints is damage to (or total wear-out of) the cartilage at the joint.

The most common form of arthritis is *osteoarthritis* (OA). This is a localised degenerative condition associated with old age and overuse of the joint – essentially, natural 'wear and tear'. Hence the eventual onset of 'primary' OA is simply an inherent consequence of a long and active life. Something of a trade-off exists, since it is in every patient's health interests to remain active in later life, and whilst regular physical activity can help control joint swelling and pain [45], excessive activity levels can increase the incidences of joint complaints [46]. However, other 'secondary' causes can advance the onset of OA, such as injury, obesity or diabetes [46-48]. The effect of OA is that moving or loading the joint results in considerable pain; this in turn makes the subject reluctant to engage in activity, effectively causing loss of joint functionality and impairing quality of life.

The second most common cause is rheumatoid arthritis (RA). This is a progressive disease in which the immune system triggers inflammation of the synovial fluid, causing destruction of the joint soft tissues. RA generally begins to cause problems at an earlier age than OA, and is systemic, often affecting

multiple joints. Because the condition can continue to progress post-operatively, there is an increased risk of revision surgery being required (see Section 1.5.3).

Several countries maintain national databases to register TJR patients, recording reasons for surgery, implant design, revision history, and other relevant statistics [49]. A review of recent registry reports demonstrates that OA is easily the most prevalent indicator for surgery, with RA consistently in second place - statistics for several registries are listed in Table 2 (note: the data is for primary TKR, and excludes revision cases). There are other possible reasons why an implant might be needed, e.g. *osteonecrosis* (damage and death of bone tissue) or serious bone or soft tissue damage (e.g. due to severe trauma). In mild or early cases (e.g. unicompartmental OA), a full TKR may not be used; in extreme cases (e.g. limb reconstruction following *osteosarcoma*) a more extensive prosthetic than a standard TKR would be required.

Registry Report	1st	2nd	Other
Australia, 2008 [50]	OA (96.8%)	RA (2.0%)	1.2%
Canada, 2007 [51]	OA (93%)	RA (4%)	3%
Denmark, 2007 [52]	OA (90.9%)	RA (5.4%)	5.7%
England & Wales, 2007 [2]	OA (97%)	RA (2%)	1%
Finland, 2006 [53]	OA (92%)	RA (4%)	4%
New Zealand, 2006 [54]	OA (92.2%)	RA (3.5%)	4.3%
Norway, 2008 [55]	OA (77.9%)	RA (7.8%)	14.3%
Scotland, 2008 [56]	OA (93.8%)	RA (4.1%)	2.1%
Sweden, 2008 [57]	OA (93%)	RA (3%)	3%

Table 2: Top reasons (with %) for primary TKR.

Note: derived from most recent available registry data (non-concurrent).

1.5. Surgical Options, Techniques & Limitations

1.5.1. A Review of Joint Replacement Technologies

As with many problems associated with old age, treatment for joint complaints aims to alleviate undesirable symptoms rather than reverse the causal underlying aging process. Partial or total joint replacement is generally the last resort when other less drastic measures to alleviate pain and/or restore function via lifestyle changes, physiotherapy or medication are unsuccessful. There are a range of possible surgical options, depending on the degree of joint deterioration. Although TKR is the focus of this study, the other options are briefly outlined below (in order of progression from most to least conservative). Whilst some of the more conservative options may be less robust or long-lasting, they should not

be overlooked; for example, uni-compartmental knee replacements may not perform as well as total knee replacements, but they can be used earlier without the risk of damaging bone stock, precisely because they are more conservative. Therefore, it is often desirable to use such methods, to forestall the need for a full TKR as long as possible. Therefore, the bullet-list below should be seen as a scale of intervention options, with the earlier options being most conservative, and therefore being preferable, where possible.

- Tissue resection: For younger patients, it may not be appropriate to use an implant at first, instead resecting the natural knee tissues, e.g. *meniscectomy*, where the damaged meniscal cartilage is partially or totally removed, and *osteotomy*, where a portion of bone is removed to better distribute loads across the knee.

- Interpositional spacers: where only the meniscus is damaged, a conservative option is an interpositional spacer, to replace the worn cartilage (so preventing bone-on-bone articulation) without any resection of bone stock.

- Hemiarthroplasty: hemiarthroplasty replaces only the articulating surface of *one* bone, e.g. a tibial hemiarthroplasty may replace only one of the tibial condyles, with an anatomically representative resurfacing implant.

- Unicompartmental & bi-lateral arthroplasty: When damage is limited to one condyle a popular option is to use a *unicompartmental* knee replacement (UKR) – this does require limited resection of both the femur and tibia, but leaves sufficient bone stock for subsequent revision to a full TKR if needed. In some cases, separate UKR implants can be used for the medial and lateral condyles (called *bi-lateral* arthroplasty), allowing the intercondylar region and associated cruciate ligaments to be entirely retained. Early clinical data shows UKR has a higher revision rate than TKR [57], and some concerns remain over whether UKR can accelerate contra-lateral condyle degradation [58]; however this is based on early experiences, and results will potentially improve as the technique is more widely practised. Nonetheless UKR is an attractive option, since despite any shortcomings in longevity it is generally easier to revise from a UKR to a TKR, than to revise a TKR.

- **Primary TKR ('tri-compartmental' knee arthroplasty)**: TKR involves resection of considerable bone stock, including at least part of the intercondylar region of both the femur and tibia. Many design variants exist, for example fixed-base, rotating platform & mobile bearing; PCL-retaining, substituting & sacrificing (see Section 1.5.2). Compared to other surgical approaches, TKR is

well-established with a proven clinical record. TKR may or may not include a *patellar resurfacing*; this will depend upon the condition of the patient's own patellofemoral joint, the design of implant and the surgical team's practice.

- Revision: Revision normally requires a more extensive implant than the original TKR – e.g. a long-stemmed hinged joint replacement might be used if considerable bone has been lost. However, if the primary procedure used a UKR, then a standard TKR might be appropriate for the revision. Generally, every new revision procedure carries further risk of infection or complications and also further reduces bone stock. Therefore, driving down revision rates by increasing implant longevity is highly desirable.

Rarely, in severe cases it may no longer be possible to provide a joint replacement, due to severe infection, or serious loss of bone stock. In these cases, the only options available may be *arthrodesis* (the 'fusing' of the joint into a fixed position, with the associated mobility impairment), or amputation.

Naturally, on this 'scale' of intervention options, it is desirable to forestall progression to more extensive solutions for as long as possible, since it is not presently possible to reverse the increasing damage to natural tissues and structures caused by the more aggressive procedures. Of all the listed options, TKR is currently the most common, and many recipients of conservative implants will eventually have these revised to TKR. As an established and widely adopted technique, it also has lower revision rates than some of the less-established alternatives, representing the best opportunity to halt the spiral of implant failure and revision. Consequently TKR is a natural focus for any design efforts to improve longevity and function. The following section outlines the design philosophies of TKR.

1.5.2. TKR Design Characteristics

In this study, the term 'TKR' is taken to refer to the design family of *endoprostheses* which resurface the entire distal surface of the femoral condyles, and resect the proximal tibial condyles, such that two artificial surfaces articulate together to form the new tibiofemoral joint. This design format became popular in the 1970's, in the wake of the success of new materials applied to total hip replacement designs [59]. Early designs either sought to mimic the geometric anatomy of the natural knee, or else to work from mechanical principles to accommodate the functionality of the natural knee. Both approaches (*anatomic* & *functional*) resulted in some common features; note the geometric similarity of the implants in Figure 13. However there are several important design aspects

where different designs follow different philosophies; the most major variations are outlined below.

Figure 13: Commercial TKR designs, from left to right: PFC Sigma (DePuy), NexGen (Zimmer), Advance MP (Wright Medical) and LCS (DePuy).

Materials: The tibiofemoral joint must be low-friction (to minimise tribological damage to the implants); metal-on-polyethylene (MoP) articulation had proven to be successful for hip implants, and so was adopted for TKR. Whereas modern hip implants are now migrating to more advanced technologies, such as all-metal or ceramic bearings, two features make this less appropriate for the knee. First, the geometry of hip bearings is simple (spherical); this is not true for knee components, which must generally be hand-finished to achieve a suitable surface. Second, the hip is constantly held in compression by the musculature; this is not true for the less stable knee joint, where tibiofemoral 'lift-off' can occur, resulting in impact loading. Therefore MoP remains the state of the art for TKR. The femoral component is generally manufactured from cobalt-chromium (Co-Cr), providing high strength, good biocompatibility and excellent corrosion resistance. The tibial articulating insert is a medical grade ultra-high molecular weight polyethylene (UHMWPE), e.g. GUR-1020, GUR-1050 or GUR-4150; however experiences with early designs demonstrated that the lower stiffness of UHMWPE against cancellous bone could lead to failure [60], and it soon became standard for the tibial polyethylene insert to be mounted in a metal tray (often Co-Cr or titanium) for stiffer backing. The use of polyethylene leads to potential concerns over the effects of wear debris (see Section 1.5.3). To counter these problems, a range of refinements have been made to the production processes for UHMWPE (e.g. gamma-ray vacuum sterilisation is used to encourage polymer cross-linking, which can greatly reduce wear susceptibility [61]).

Patellar replacement: The patella may or may not be separately re-surfaced; if it is resurfaced the implant is generally all-polyethylene. The shape of the patellar implant may be anatomically representative, or simply oval or dome-shaped. Similarly, the anterior groove on the femoral implant may or may not be

symmetrical: in the natural knee, the patellar groove slopes laterally by several degrees as it progresses proximally. However if this feature is adopted, separate femoral components must be manufactured for left and right knees. To avoid this, some designs use a straight vertical patellar groove. The disadvantage is that this may change the line of action of the quadriceps force, diverging from the normal loading of the natural knee.

Tibial bearing design: The use of a metal-backed insert is now widespread, and many designs now also introduce an additional degree of freedom between this tray and the polyethylene insert. The theory is that allowing the tibial insert this extra freedom of motion can split the tibial motion across two different bearings – this means both that the tibial insert can rotate to a more conformal position against the femoral component (increasing surface contact area and so decreasing pressure), and also that the kinematic motions associated with producing wear can be reduced [42] (see discussion on wear and cross-shear motions in Chapter Four).

One design concept is to use a central peg, permitting only I-E rotation between the tray and insert; i.e. a *rotating platform* (Figure 14, centre). Another concept is a slotted peg permitting both rotation and translation; i.e. *mobile bearings* (Figure 14, right). Nonetheless, *fixed* designs with no tibial bearing (Figure 14, left) are still common; although theoretically rotating & mobile bearings offer advantages, currently these benefits do not clearly translate to improved clinical results [62, 63].

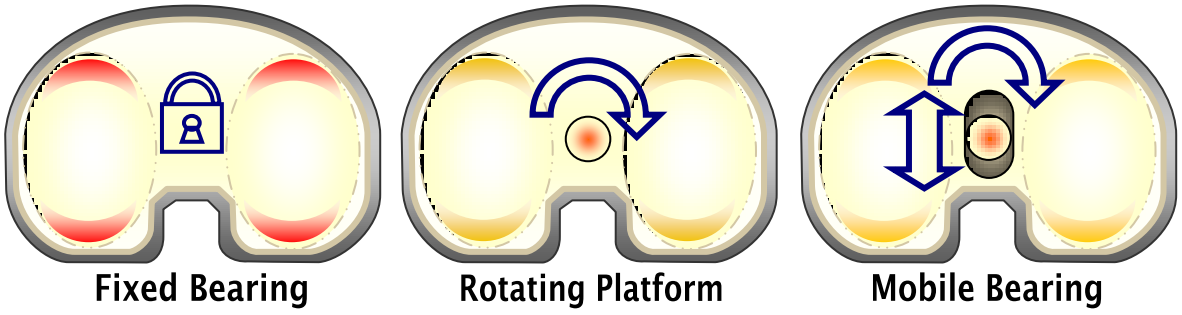

Figure 14: Comparison of tibial bearing designs.

Cruciate retention/resection: surgical treatment of the cruciate ligaments is an important decision in the choice of knee implant design. Almost every design must resect the ACL, as a necessity in order to install a full resurfacing implant (this has been one impetus for bi-lateral use of UKR, where the ACL and PCL can both be preserved), and there is no need to resect the collateral ligaments (which do not obstruct the intercondylar region). However there is no clear consensus on whether to resect or retain the PCL; the final decision rests with the practising

surgeon, and the key factors are typically the physical condition of the patient's ligaments, and individual surgical preference.

Some implants are designed to leave the PCL intact, taking advantage of the stability it provides for large flexion angles. These designs are referred to as PCL-retaining, or simply cruciate-retaining (CR), and generally have less conformal sagittal geometry, since the PCL helps restrict A-P motion (see Figure 15). However this requires the PCL to be in good condition, and correctly tensioned when the implant is fitted. Such low conformity-surfaces can also result in higher tibial contact pressures.

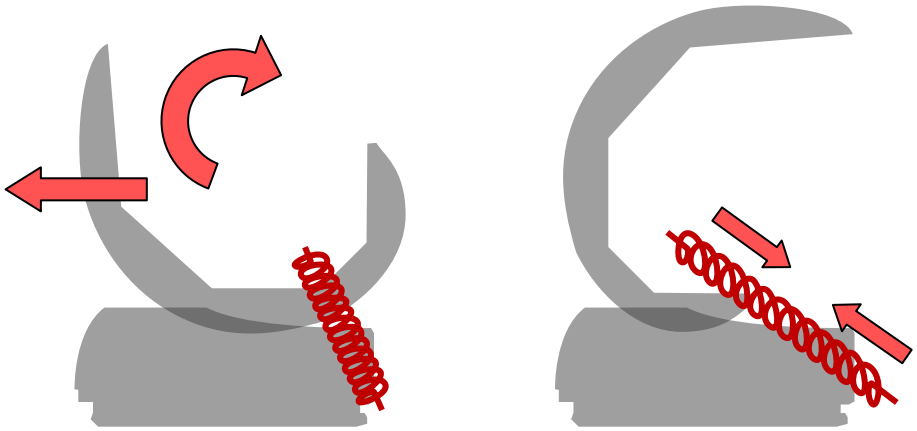

Figure 15: 'CR' design. As the femur flexes and experiences anterior force, the PCL acts as a 'spring' to constrain the anterior motion.

The alternative is to resect the PCL; devices which do this are PCL-substituting or PCL-sacrificing (PS). The key design feature is either a more conformal geometry, or else a distinct motion-constraining feature, e.g. a cam system in the intercondylar region (see Figure 16). The implant must provide the constraint which the PCL would otherwise offer, so larger restraint forces must be supported. If features such as camming systems are used this can lead to large shear stresses within the cam, whereas using a more conformal surface will distribute contact forces better, reducing surface contact pressure. The decision must be based on the condition of the PCL in the patient. Many designs are modular (so that the surgeon can choose an alternative if intra-operative inspection of the PCL indicates that additional constraint is needed). Reported clinical results for both PCL retaining and resecting approaches are mixed; although retaining the PCL is considered preferable when the surgical team is experienced [64].

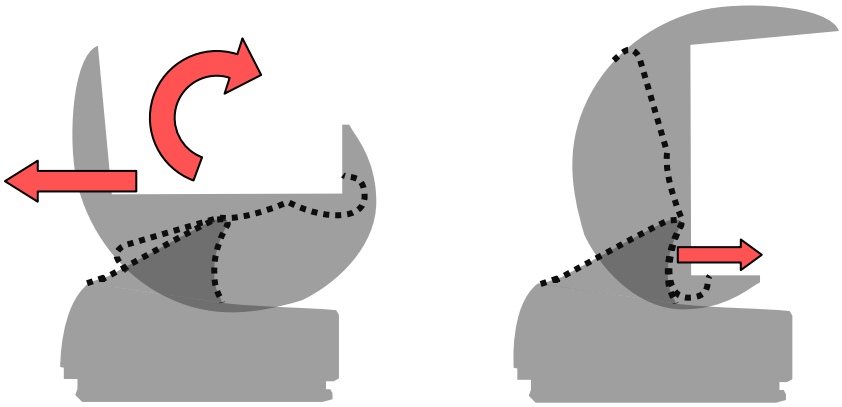

Figure 16: 'PS' design. As the femur flexes and experiences anterior force, the inter-condylar 'cam' system engages to resist anterior motion.

Fixation: As with hip replacements, the option exists to use either cemented or cementless fixation. Cemented designs typically use cement to fix the tibial tray and the femoral resurfacing implant (the tibial insert is normally held in place by mechanical interlock with the tray). Direct cementing of an all-polyethylene tibial component may be used, reinforcing the low stiffness of the UHMWPE with the higher-modulus polymethylmethacrylate (PMMA) bone cement. Components may be designed for a 'press-fit' (e.g. the femoral resurfacing implants), although due to the high rotational torques experienced by the knee (see Section 1.3), fixation pegs are often included. For cementless designs, a further possibility is to provide a coating for improved fixation; this may take the form of a porous coating for better mechanical interlock, and/or *osteoconductive* coatings such as hydroxylapatite to encourage bone in-growth. Whereas for hip prosthesis, cementless designs have risen considerably in popularity in recent years (with 30% of recent UK hip procedures using cementless fixation vs. 48% cemented), for the knee implants, only 7% are cementless vs. 83% cemented (the remainder in both cases being hybrid or conservative implants – data from NJR 2007 [2]).

1.5.3. TKR Failure Mechanisms

Assuming the initial arthroplasty surgery is successful, there are still many risks of failure post-operatively. Some of the most common are listed below. Note that the factors are not independent or exclusive (e.g. wear-induced osteolysis may lead to loosening, or increase the risk of direct mechanical failure of the component). Some of these factors are unrelated to the mechanical environment; others depend strongly on joint mechanics and it is these factors (designated by an asterisk*) that are most relevant to the results of the subsequent computational mechanical modelling; other factors are recorded here for completeness but will not be attended to further in this thesis. Statistics for different national joint registers are compared in Table 3, showing the frequency

with which these failure types occur. It is noteworthy that in the table, objective 'mechanical' failure criteria such as loosening do dominate; however other more subjective criteria such as pain also feature for some national registries.

- **Wear*:** Although any surfaces in moving contact will experience wear, for TKR the most prevalent occurrence is wear of the polyethylene used in the tibial and patellar components. Although the small degree of wear might not compromise the implants directly (i.e. the structural integrity of the implant itself may remain satisfactory), the build-up of wear debris in the surrounding biological tissue can lead to problems; in particular *osteolysis* [65]. This is a process whereby *macrophages* (biological agents which are part of the body's immune response) attempt to remove the foreign wear debris from the body, but in the process also reabsorb the surrounding natural bone stock. This leads to decreased bone density and hence ultimately can compromise the mechanical integrity of surrounding bone tissue (which in turn can cause bone fracture and/or component loosening). Macrophage activity is part of the normal bone remodelling process, but can be increased by the presence of wear debris particles. Use of appropriate low-wear materials and controlling the articulating motion can both help reduce the volume of wear debris created. For a given implant design and material type, wear is hypothesised to be a function of the sliding path motions and the contact pressure; these factors can be readily investigated with mechanical models, and so will be a key metric for assessment in the analyses described in this thesis. (For more on the mechanics of wear, refer to Chapter Four).

- **Loosening***: the failure of a mechanical fixation interface (for cemented implants, either the bone-cement or implant-cement interface; for cementless designs, the bone-implant interface, or if coatings are used, at the coating interface). This is commonly reported (e.g. [66]), and can have many causes; e.g. wear-induced osteolysis or migration (due to poor initial fixation, or cementing quality leading to excessive micro-motion). Although loosening will not be investigated directly, it is the most frequently reported mode of failure, and is often causally related to other failure mechanisms which will be explored (e.g. wear).

- **Dislocation / Subluxation***: this can potentially occur for the patellofemoral or tibiofemoral joint if soft tissue constraint (from the ligaments or fibrous capsule) is inadequate, or if the component is severely malpositioned during surgery. Although rare, such incidents are reported in the literature [67]. Since

these events are a direct function of the knee kinematics, they can be easily assessed using mechanical modelling.

- **Instability***: The knee is inherently unstable, depending on the extensive surrounding soft tissue (ligaments and muscles) to provide stability. Consequently, if this soft tissue support is damaged (e.g. trauma from surgical incision, or resection / improper balancing of the ligaments) the internal kinematics of the knee would be under-constrained. This can cause the patient to feel the sensation of instability, reducing their confidence when walking or moving, and hence impairing mobility [68]. Note that *instability* is a result of underlying mechanical issues, but also psychological patient-perceptions; the latter are very difficult to analyse; however the actual degree of mechanical constraint can be assessed readily as an indicator of instability.

- **Patellar complications***: patients may report pain specifically around the patella; poor patellar tracking is a common problem. In cases where the patella was not initially resurfaced, a revision may be indicated to include a patella implant if the original results are not satisfactory. Again, there is a degree of subjective patient perception involved, but the mechanics (patellofemoral kinematics and contact pressures) are more objective and can be measured directly.

- Disease Progression: with TKR for RA, the joint may continue to degenerate after surgery, leading to further problems. Good survival rates can be achieved for RA patients [69], but rates are generally lower than for OA [2, 50-57]).

- Mechanical Fracture: Implants can potentially suffer structural failure *in-vivo*. Historically there were problems in particular for the tibial tray [70] (due to poor design, or osteolysis undermining the bone supporting the tray); standard tests are now used to reduce this risk [71].

- Limited function: if components are misplaced or ligaments are not balanced, knee ROM may be reduced, impeding some ADLs (e.g. kneeling, stair climbing). Rarely, this may even be an indication for revision [72].

- Infection: deep infection may necessitate the removal of the implant. Good hygiene practice, implant irradiation, including antibiotics in cement and similar measures can mitigate this risk [73], but infection rates of a few percent remain typical [2, 50-57].

- Pain: may be related other failures, (e.g. where 'progression' is reported, the patient will also be in pain). Post-operative pain is common [74], and even

without clear causes may necessitate revision. Unlike mechanical modes of failure, pain is a very subjective metric to assess.

Registry	1st	2nd	3rd
Australia, 2008 [50]	Loosening (37%)	Infection (15%)	Wear (8%)
Canada, 2007 [51]	Loosening (33%)	Wear (30%)	Instability (17%)
Denmark. 2007 [52]	Loosening (35%)	Pain (21%)	Instability (18%)
England & Wales, 2007 [2]	Loosening (46%)	Pain (16%)	Osteolysis (16%)
Finland, 2006 [53]	Infection (25%)	Misalignment (12%)	Patella (8%)
New Zealand, 2006 [54]	Pain (33%)	Loosening (33%)	Infection (26%)
Norway, 2008 [55]	Loosening (25%)	Pain (22%)	Infection (10%)
Sweden, 2008 [57]	Loosening (25%)	Patella (21%)	Instability (11%)

Table 3: Top specific reasons for revision (%). Most recent data (non-concurrent). Percentage values averaged for all revision types.

1.5.4. Success Rates with TKR

At first consideration, TKR is a very successful procedure, with most registries and studies typically indicating survivorship of 90%+ at 10 years, and 80%+ at 20 years [75]. However, the fact that an implant has 'survived' (i.e. not been removed) does not automatically make it faultless (a patient may still be in pain, or suffer from lack of function). A single discrete 'success/failure' metric does not capture these other problems and limitations; a more graded, multi-factor scale is needed to identify underlying issues.

Subjective scores such as the Oxford Knee Score [76], or Knee Society Score [77] can be used to gauge the implant's success on a more continuous scale, and these generally show that the majority of patients experience *some* pain and/or loss of function post-operatively. For example, in the England & Wales National Joint Registry, less than 10% of knee patient respondents reported 'no' or 'hardly any' problems – implying that over 90% of respondents had some problems with their new implant [2]. Compare this to the hip patients in the same registry report: around 30% had 'no' or 'hardly any' problems; over three times the equivalent knee rate. The proportions with moderate to severe problems were also about twice as high for the knee patients as for hip patients: 11% versus 6.1%. This illustrates a considerable disparity between the patient-perceived outcomes of hip and knee replacements. Studies which have used 'pain' as a failure endpoint (rather than revision) see a much higher 'failure' rate (e.g. around 30% [74]). Similarly, although very few knees are actually revised due to inadequate flexion range [72], this masks the higher rate of patients with imperfect knee functional scores post-operatively (e.g. inadequate flexion range,

which is not serious enough to indicate revision, but may impede participation in certain activities). The clear conclusion is that, across the population, knee replacements are not performing as well as hips.

Further, pre-operative patient selection masks the true effectiveness of the treatment; younger patients are contra-indicated for TKR, based on the increased demands of a longer potential lifespan and higher activity levels. Therefore the true number of people whose needs and expectations are not fully met by current TKR procedures is much larger than the headline revision rates alone suggest.

Even laying these caveats aside and reviewing revision rates alone, percentages should still be considered in terms of the underlying real numbers. Data from those nations with national joint registries accounts for over 120,000 knee replacements per year [2, 50-57], and with an estimated 300,000 knee replacements annually in the US alone [75], the annual figure worldwide is substantially over half a million procedures. Consequently, even a few percent represents tens of thousands of patients every year for whom TKR has been unsuccessful. As the number of patients continues to increase, the case for driving down the percentage of failures is strengthened. This should include addressing all aspects of sub-optimal performance (i.e. improving longevity, reducing pain, and also increasing functionality.)

However, a review of long-term registry data demonstrates that the *rate* of improvement is decreasing (for example, in the longest-running Swedish registry, it is reported that whereas TKR revision rate at 5 years dropped from over 12% to ~5% between 1980-1990, the corresponding drop between 1990-2000 was only ~2%, from ~5% to ~3% [57]). This is a classic example of the 'Pareto Principle'; the drive for improvement becomes progressively more challenging as improvements are made: as the most obvious and effective improvements are implemented, the remaining outlier cases are more difficult to address, requiring more detailed understanding of the system and more effort to engineer appropriate solutions.

1.6. Summary

This chapter has presented an overview of knee anatomy, demonstrating that the knee is a complex joint facing demanding loading conditions. The challenges this presents for TKR designs has been discussed, along with a review of existing design solutions, and a discussion of how these designs are performing *in-vivo*. It has been shown that there are important shortcomings which still need to be addressed, and the challenge facing orthopaedic researchers is increasing as the

drive for continuous improvement requires an ever great understanding of the knee, and ever more design effort.

However, the time and resources that can be dedicated to orthopaedic research cannot correspondingly increase indefinitely; therefore future research efforts must become more carefully focused, to achieve the most benefit with the finite resources available. To this end, the role of pre-clinical assessment tools has grown more prominent over recent decades; developers cannot afford the costs associated with development, prototyping and clinical trials before detecting problems with a design. Increasingly, efforts focus on predicting likely outcomes whilst the implant is still in the early design stages; i.e. *pre-clinical* analysis.

Note that the subjective nature of many 'failures' presents a particular design problem; for instance, if a patient reports severe pain, this may indicate revision, even if no causal explanation can be found for the pain. It is very challenging for an implant designer to address such subjective and poorly-understood issues in the pre-clinical design phase; and indeed how this may be done is beyond the scope of the present thesis. Rather, we will attempt to demonstrate how the *general methodology* of pre-clinical design might be improved, and demonstrate these improvements within the domain of some of the more 'established' pre-clinical testing that is routinely performed (i.e. mechanical phenomena such as kinematics and wear). It is hoped that in the future, some of the lessons learnt will be more generally applicable to other forms of pre-clinical analysis for TKR implant designs.

Chapter Two will review some of the techniques developed to assist in this pre-clinical assessment, and demonstrate the need for more intelligent assessment tools as the research field continues to develop.

CHAPTER TWO - PRE-CLINICAL ANALYSIS METHODS

Experimental and theoretical methods used for pre-clinical analysis, and the benefits of cross-corroboration between alternative approaches

2.1. The Motivation for Pre-clinical Modelling & Analysis

As has been discussed in Chapter One, the human knee is a complex system, and any surgical intervention or implant design requires a robust understanding of this system behaviour to achieve the optimal outcome. The earlier in the design process that change decisions can be made, the lower the subsequent development costs; hence there is a strong incentive to have an effective set of analysis methods available pre-clinically. This requires a representative model (offering some advantage in terms of time, risk, cost or ethics compared to a clinical trial) to enable practical predictions of the likely performance (and hence suitability) of a design proposal, modification or feature. If this model is of sufficient robustness & integrity, it may then be used to predict behaviour under 'perturbed' conditions, and hence ultimately be used for broad-based 'probabilistic' studies of the full range of varying factors.

Quantitative pre-clinical assessment tools are well established, and have flourished with the improvement over recent decades in electronics, sensors and computing performance. Such models can be broadly split into two disciplines: theoretical/analytic models, which virtually model the expected behaviour of the system to make predictions, and the experimental/empirical models, which directly test 'real-world' models under representative physical conditions. The analytical approach is generally modelled computationally, to handle the complexity of the models, and may broadly be referred to under the label *in-silico* modelling; the empirical approach is equivalently termed *in-vitro* modelling. Both classes of analysis can be contrasted to clinical trials within a living subject, termed *in-vivo* tests. It is important to appreciate that *in-silico*, *in-vitro* and *in-vivo* analysis can be complementary rather than competitive, since their functions are fundamentally different:

- ***In-silico*** models generally either explicitly model the underlying 'physics' of the system, or else fit some response to an extant 'training' set of known results; this can result in fast, low-cost models which are readily parameterised and monitored to study the effect of perturbations or configuration changes within the domain of the 'known' operational behaviour. However, they cannot address the phenomenologically 'novel'; if the system is operating under conditions where the underlying laws of physics are not properly understood, purely analytic models are not applicable. (In

other words, *in-silico* models may interpolate between known conditions, but when extrapolated to operate under unknown conditions or additional confounding effects, less confidence can be placed in the results).

- ***In-vitro*** models by contrast implicitly invoke the laws of physics in order to operate; therefore the underlying behaviours do not need to be understood to undertake testing – however this does limit the ability of the models as predictive tools. Experimental tests can be more expensive and time consuming than computational tests, and can be limited by the practical achievability of the proposed test conditions (e.g. hardware limitations).

- ***In-vivo*** tests may be considered the 'gold standard', since they directly test under *in-situ* conditions; however, the associated expense, timescales and ethical issues often make such testing highly challenging. Further, *in-vivo* tests tend to be very specific and narrow in focus, which is particularly problematic given the very high variability associated with biological systems. For example, tests on a single patient with *in-vivo* telemeterised sensors may provide 'real' results, but for the equivalent cost and effort, it might be possible to perform gait analysis using inverse dynamics on a large cohort of patients, giving some indication of the statistical distributions observed in the variability of force magnitudes.

Both analytical and empirical approaches have been used extensively for analysis of the human knee; the following sections will consider some of this extant published research.

2.2. Theoretical/Analytical Methods

Mathematically-based models are valuable, because they use the underlying physics of the system to predict outcomes without the cost and risks of a physical simulation. In reality, the knee is a complex system, and cannot be accurately defined without extensive and complex mathematical formulations. Whilst some of the constituent mechanical behaviours can be represented satisfactorily using parametric equations (e.g. ligament stiffness/strain relationships), others (e.g. contact mechanics) can only be loosely approximated by such simplistic functions. Again, this complexity has resulted in a range of different approaches to modelling, depending on the objectives, and the available resources. Fundamentally, a distinction may be made between models which remain purely analytic, using parametric and differential equations to describe the knee holistically at the system-level, and those models which adopt a discretised 'numerical integration method', e.g. the 'finite element' approach (see Figure 17).

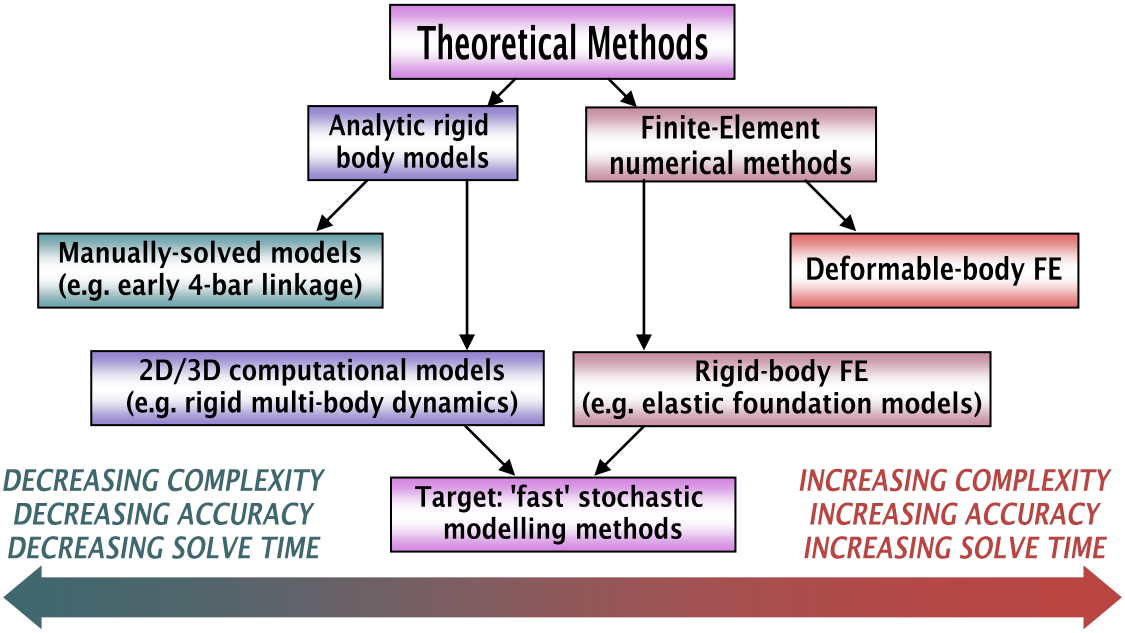

Figure 17: A broad classification of 'theoretical' knee models.

Rigid-Body Modelling

The fact that a simplified analytic model could be solved with less numerical effort made these ideal candidates for very early models of the knee, before the advent of affordable and accessible computing. Simple rigid-body models in the sagittal plane can be dated to the early part of the 20th century, e.g. the 'four-bar-linkage' model, treating the cruciate ligaments as rigid restraints during flexion-extension F-E motion (employed by Strasser as early as 1917 [78]). This model was progressively developed by subsequent research, for example including non-linear elastic spring elements for ligaments [79], and performing sensitivity studies on the model [80]. These four-bar linkage models have been widely used to develop understanding of the mechanics of the knee (e.g. the work of O'Connor et al [81]).

Initially, when computational power was limited, this sagittal-plane modelling approach proved popular, as it captured the single most obvious motion (i.e. F-E), but could also describe some secondary motions (e.g. A-P 'rollback' as the knee flexes). However, sagittal-plane-only models cannot include I-E, V-V or M-L motions. 3-D analytic models did begin to emerge in the 1980's [82], taking advantage of computers to assist with the calculations. However, these could still be classed as analytic models, since the system was still defined globally using analytic equations (not discretised numerical integration methods); the computer purely assisted with evaluation of the mathematics. Today, purely analytic models are still used for some studies of the knee, notably in the work of Hefzy et al [83], which has advanced to include detailed analytic representation of both patellofemoral and tibiofemoral articulations, with nonlinear ligaments and

quadriceps wrapping (Figure 18). These models may very readily be formulated with *dynamic* equations, whereas FE models are often static or quasi-static.

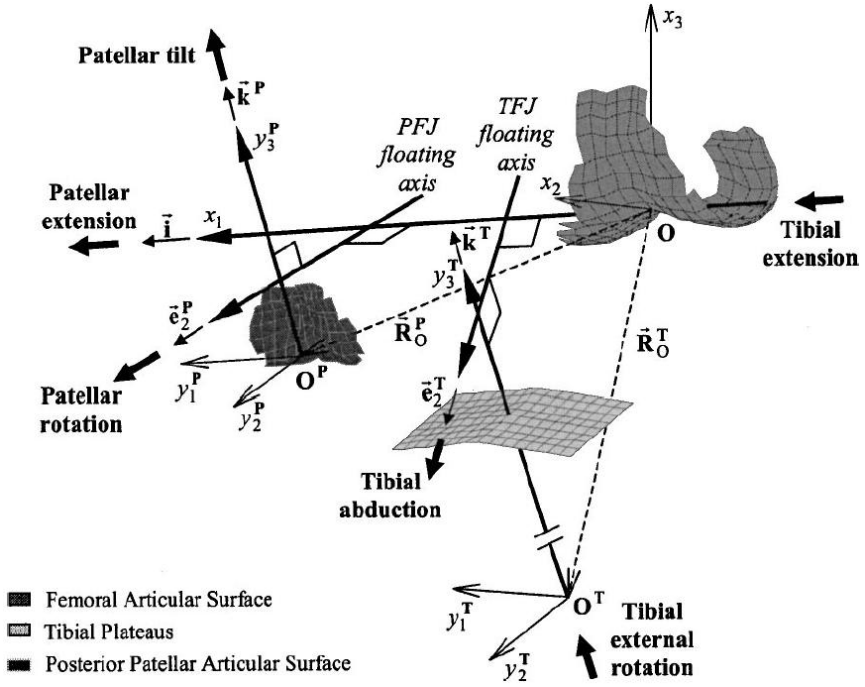

Figure 18: Analytic modelling of the knee joint (from [83]).

Although much research has been dedicated to FE-based models, analytic solutions remain a very useful tool, especially for stochastic simulations, since they offer the low computational cost essential for large volumes of trials. Multi-body dynamics (MBD) simulations fall within this category, and are still widely used (e.g. Bei et al [84] demonstrating the combination of dynamic simulation with multi-body deformable contact). This approximation to the true deformable behaviour of the material uses a pre-defined relationship between pressure and 'penetration distance', or 'overclosure'. Typically, the penetration will be estimated at a number of points, forming a rudimentary discretised 'point cloud', reminiscent of the finite element approach discussed below. Despite this use of discretisation, MBD models are effectively a class of analytic model, since they still seek to apply analytic equations to describe the system dynamics, and do not fundamentally have to depend upon discrete numerical integration methods. They may be distinguished from 'pure' analytic models in that the geometries involved are not represented in analytic form; rather, they are typically obtained from CAD data, making the analysis too complex to evaluate without a computer. MBD models are widely used for models of both natural knees & TKR implants (Figure 19).

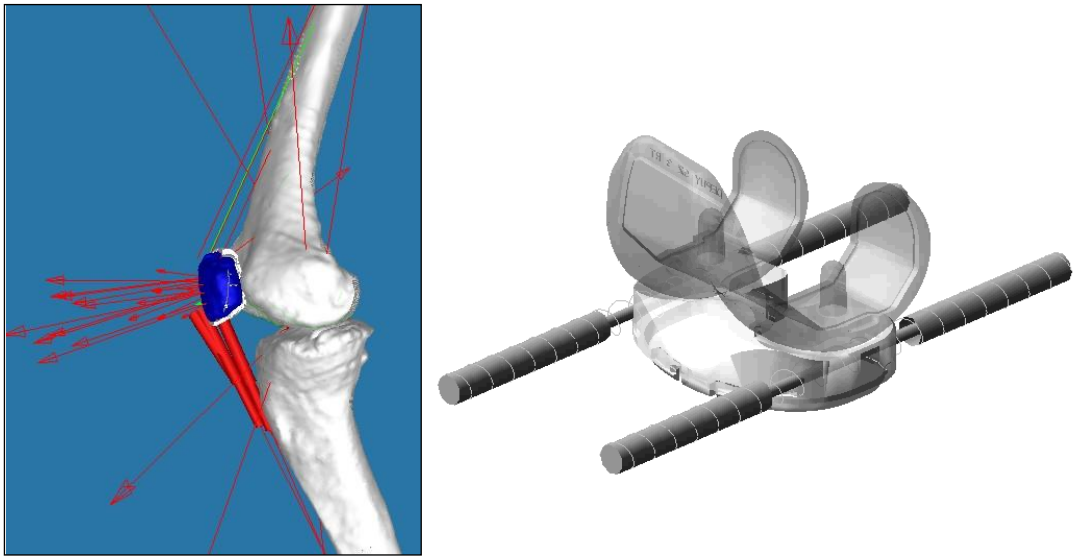

Figure 19: MBD for the natural knee, left (courtesy Guess et al, UMKC), and for artificial implants, right.

Finite Element Methods:

Whilst computers may be used to ease the calculation of the simple purely analytic models described above, they have also enabled a fundamentally different approach to the analysis of biomechanics – finite element analysis (FE, or FEA). Unlike analytic models, the basis of numerical integration methods is not to provide 'exact' solutions, but instead approximate the true result by a 'brute force' approach to the solution, applying fundamental physical equations discretised across small spatial and/or temporal intervals. As these intervals become smaller, the approximation becomes better, but computational effort also increases as the number of separate discretised equations increases.

As with other analytic *in-silico* models, an advantage of computational numerical techniques for stochastic studies is that the process of parameterising the model is simplified; a numerical value representing the input parameter can be changed instantaneously. Conversely, the individual trials are now much slower, such that a high-fidelity FE model of the knee gait cycle might require hours to simulate an event lasting around one second. So, in contrast to experimental simulations, computational models are very easy to *re*-configure (after initial pre-processing), but currently take much longer to simulate. As such they are better suited to multi-variable parametric studies, but less well suited if a highly adaptive model is required to run many successive cycles. (*In-silico* models may still be used for adaptive studies, e.g. in adaptive long-term wear damage studies [85]; typically a large 'step size' must be used between adaptive updates, to limit the number of analysis runs required). A stochastic analysis requires more trials (potentially thousands, but even if 'fast' stochastic methods are used, dozens of trials will still be needed). This is not problematic if there are sufficient computational

resources available for simulation; however, in the present case, the work in this thesis is being done as a 'proof-of-concept' with (relatively) limited processing power available. Therefore, for these studies a high-resolution model using deformable FE would not be appropriate.

The Application of FE to Biomechanics

FE was first applied to knee bioengineering in the 1970s, with simple 2-D static simulations (e.g. Askew et al [86], who demonstrated the effect of bone anisotropy on the fixation of the tibial component). Full 3-D static models soon followed; in the early 1980s Lewis et al used 3-D FE to evaluate different tibial component designs [87]. Most early studies focussed on bone stresses; later studies began to investigate the stresses within the polymer components [88], with studies focusing on both the tibia and patella (e.g. the work of Bartel et al [89, 90], using FE to differentiate between designs, for example to demonstrate the effect of increased conformity or insert thickness). For non-static assessments (i.e. implant kinematics), explicit quasi-static models are used (as demonstrated by Godest et al for TKR gait [91]). Combined with stress predictions, this allows FE-based models to predict wear performance (e.g. Knight et al [92]), and to be used for design optimisation algorithms (e.g. Willing et al [93]). However, these sophisticated studies come at considerable computational cost.

Modern FE simulations can be elaborate, including fully non-rigid deformable bodies, membranes for modelling ligaments, and complex contact friction effects. This can result in simulations requiring several hours to achieve a full solution across a single gait cycle (e.g. an explicit deformable-FE gait cycle can require 6-7 hours [91]) – making large multi-cycle analyses very laborious. It is possible to simplify the FE model to achieve much faster solve times (using rigid-body contacts essentially similar to the MBD models); for example work by Halloran et al using the elastic foundation (EF) method [94] has been used as the basis for the first stochastic studies of TKR mechanics [95].

2.3. Experimental Methods

Empirical measurements, using experimental physical rigs or simulators, are often considered superior to purely theoretical simulations, since the physical laws controlling the system are implicitly invoked, rather than explicitly modelled (meaning the model can simulate un-investigated physical conditions). However, the simulation must still represent the *in-vivo* conditions as closely as possible, and must be carefully designed to ensure that all influential behaviours are captured. For TKR pre-clinical analysis, a top-level distinction can be made between two classes of experimental methods (Figure 20). General materials

testing screens for material properties (e.g. strength, wear or biocompatibility), whereas mechanical knee testing directly uses 'true' natural or implanted knee geometry, capturing the combined effects of materials *and* design form.

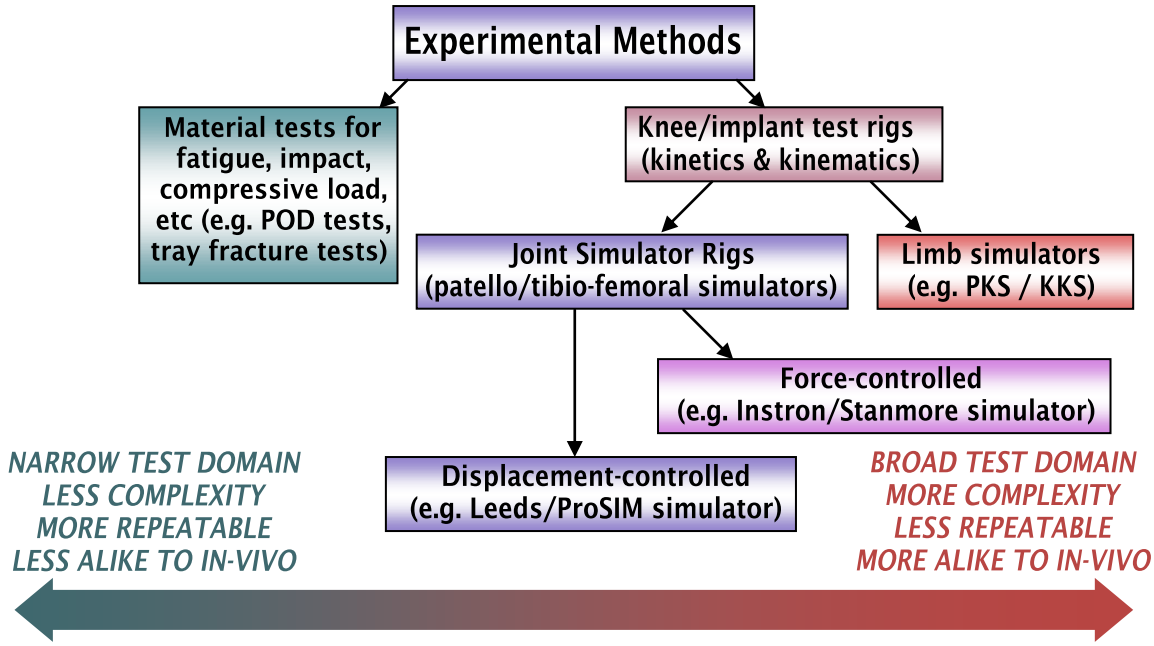

Figure 20: Representative classification of biomechanical simulators.

Materials Testing

A wide range of different tests will be performed on any new material prior to clinical adoption, including various forms of biological and mechanical screening (for example, impact testing on ceramic components, fatigue testing on tibial trays, or wear testing on metal-alloy and polymer components). A study of the literature on *in-vitro* testing (e.g. amongst many others [96-99]), the specific ISO or ASTM standards for materials testing (e.g. [23, 71, 100]) or any of the relevant textbooks in the field (e.g. [101, 102]) will provide more detail on some of the different forms of testing undertaken. In this particular thesis, the only testing method which will be described in further detail (owing to its direct relevance to the present body of work) is the widely-used pin-on-disc (POD) test, also known as pin-on-plate (POP). This is a tribological assessment of the material, to gauge the likely wear-rate that might be seen *in-vitro*, and ultimately *in-vivo*. POD tests do not use implant geometry, which would otherwise become a confounding factor. Instead, the same geometry is used for each candidate material (a simple flat-headed 'pin' of the material articulating against a flat disc of the opposing material), and a simple, repeatable motion is driven under compressive loading (see Figure 21). The motions and loads are normally not physiologically-derived (the sliding path normally follows a simplified path profile, and the POD machine

cannot distinctly model rolling as opposed to sliding contact). TKR tests and POD tests tend to give different values for normalised wear rates; this is not fully understood as the precise mechanisms of wear are not fully characterised. Nonetheless, data from POD wear testing provides a qualitative ranking of relative wear rates compared to existing established clinical materials, and can be useful for providing a first estimate of the expected wear for TKR implants. Following the advent of modern orthopaedic implants, POD tests have been used as a baseline wear test of materials since the early 1970s and through subsequent decades (e.g. [97, 103, 104]).

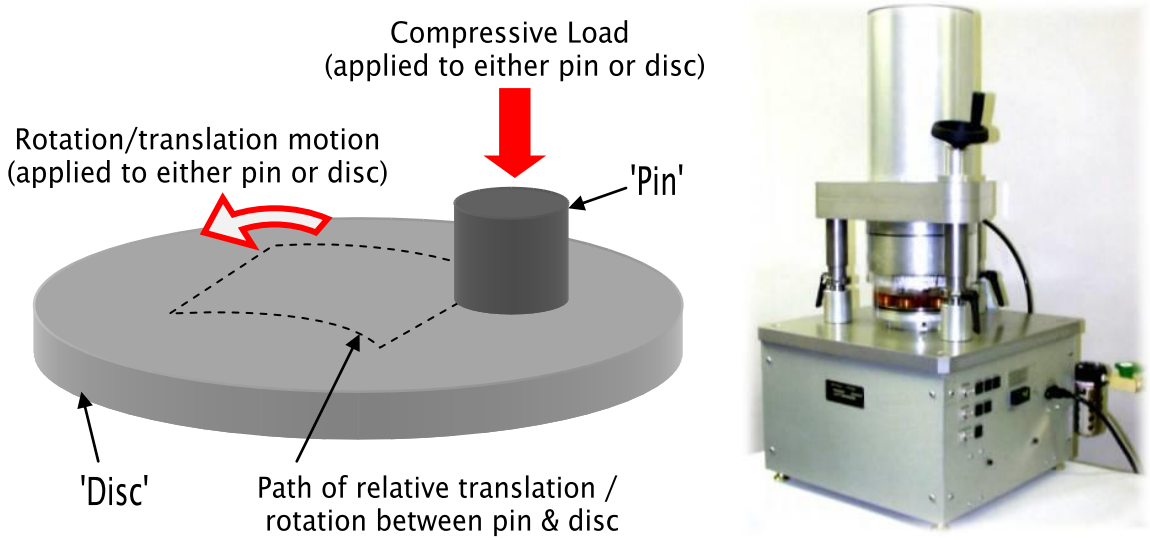

Figure 21: Typical POD wear test configuration (left) & commercial POD tester: AMTI 6-station 'Orthopod' (right, image: Advanced Mechanical Technology, Inc.)

Today POD testing is still widely used to explore the relationship between the factors hypothesised to influence wear performance (e.g. the relative effects of cross-shear [105] and correlation to contact pressure [106, 107]), as it provides a much more controllable environment to define specific motion paths and contact pressures. POD testing is used extensively to explore the fundamentals of wear, in order to devise theoretical algorithms for wear prediction (see Chapter Four).

Mechanical Knee Testing

Whilst materials screening is necessary, it is still essential to perform tests which are more representative of the full *in-vivo* conditions. For this reason, mechanical knee test simulators were developed, which directly perform static or dynamic mechanical tests on the implants themselves. In only a few decades these *in-vitro* models have developed from rudimentary 2-D rigid-linkage models to sophisticated representations of the complete lower limb (Figure 22).

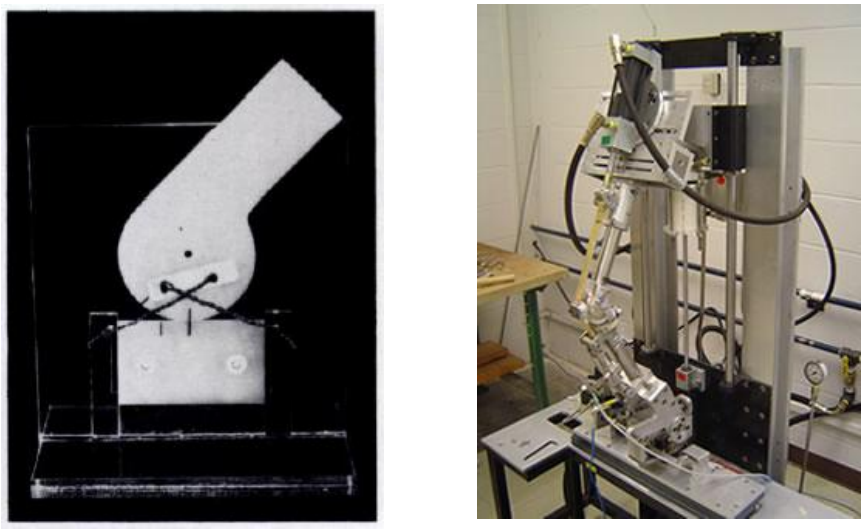

Figure 22: Early (left) & modern (right) *in-vitro* knee models (from [108] & [109]).

There are a few examples in the literature of early rudimentary biomechanical simulators (e.g. the 2-D sagittal model reported in [108]); however the first noteworthy efforts came in the late 1970's and early 1980's. For example, Werner et al performed limited testing of I-E torques and rotations for different prostheses [110], and later Thatcher et al developed a more comprehensive rig capable of applying axial compressive loads, shear loads and torques, to monitor both A-P translation, and rotations [111].

In reality it is difficult for a simulator to accurately reproduce the mechanical environment of the knee. The knee is an unstable joint with complex geometry and kinematics, driven by multiple muscle forces and restrained by a complex arrangement of active and passive soft tissue constraint. Further, there is no 'standard' human knee (due to the degree of patient variability) – geometries, forces and tissue properties all vary considerably. In response to this complexity, different modelling approaches have been used for *in-vitro* test designs.

A fundamental difference in the 'extent' of the modelling scope emerged (see Figure 20): in some cases, the model focussed only on the internal kinematics and kinetics of the implanted prostheses, often focusing on a single articulation (e.g. tibiofemoral only) – these may be termed 'joint' simulators. In other cases, the entire lower limb would be modelled (to include the effects of muscle forces & lines of action) – these may be termed 'limb' simulators. Both classes of rig are discussed in more detail below.

Joint Simulators

The knee is part of a complex system of muscles, joints, ligaments and bones which together form the functional lower limb. To avoid modelling the full complexity of this limb-level system, many tests are devised to model solely the

loads at the specific joint interface. These 'joint' simulators can then use aggregate loading and restraint conditions to mimic the effect of the surrounding muscles and soft tissues. The various knee-wear simulators (as described in ISO-14243 [23, 24]) are an example of this design ethos.

Two rival approaches to control may be adopted: *force*-driven or *displacement*-driven (Figure 20). Whereas the biological knee is inherently *force*-driven (due to body weight, muscle loads and ground reaction forces), many simulators directly drive the relative *displacement* of the femur and tibia (to by-pass the complex interplay between limb lever arms, muscle forces & moments, articular surface geometry and restraining soft tissue).

To speak of 'force' or displacement' control is slightly misleading, since either strategy is generally a hybrid; e.g. axial compression is universally applied as a force rather than an inferior-superior (I-S) displacement, conversely F-E is generally applied as an angular displacement input regardless of the other control inputs. The difference emerges for 'secondary' effects such as A-P shear force vs. A-P displacement, and I-E torsion torque vs. I-E angle.

Note that no simulator operates in true 'displacement' control, since the 'real-world' actuators are inherently force-driven (typically pneumatic or hydraulic). Rather, the control feedback loop will use displacement (measured with LVDTs or potentiometers) as the target control signal. This does mean that achieving accurate tracking is important in displacement-controlled simulator design, and it is possible for the achieved 'true' displacement-driven kinematics to be very different to the intended 'target' input waveforms (e.g. see [112], where with a more conforming implant under test, the displacement-driven simulator could not scarcely achieve 50% of the desired displacement during stance phase).

One of the first tibiofemoral joint simulator designs was by Walker & Hsieh [113], which simply oscillated the femoral component whilst applying a constant stance-phase load on a multi-station machine. Many variants followed; notable recent examples are the Leeds/ProSIM simulator (Simulation Solutions, Stockport, UK), the MTS Bionix knee wear simulator (MTS, Eden Prairie, MN, USA) and the AMTI-Boston simulator (Advanced Mechanical Technology, Watertown, MA, USA), all of which can be displacement- or force- driven, and the Instron/Stanmore knee simulator (SKS) [114], a force-controlled tibiofemoral simulator (Figure 23). These knee wear simulators are widely used for commercial testing of TKR designs.

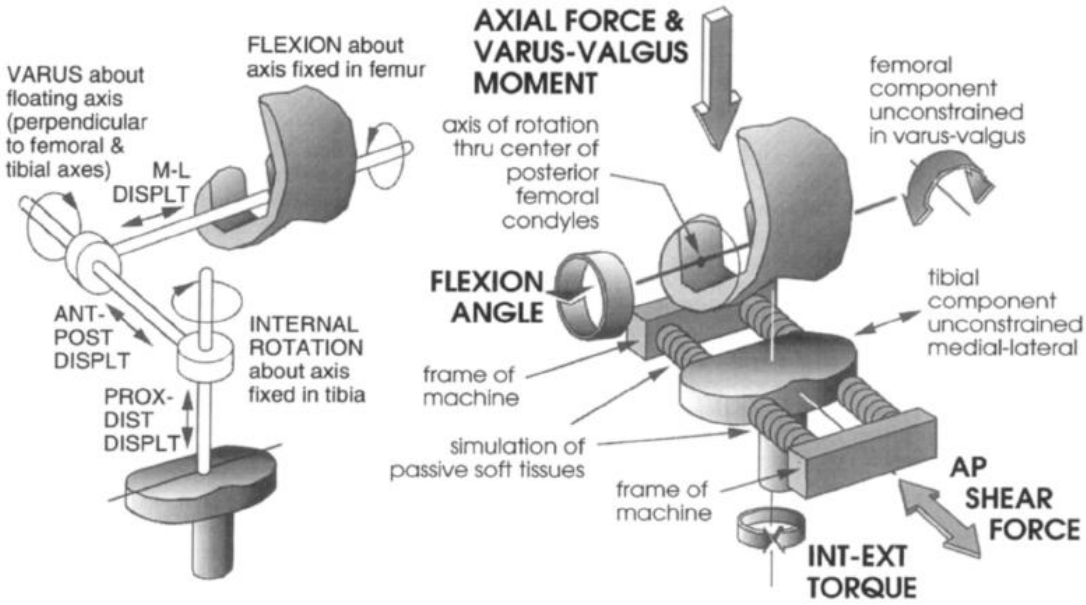

Figure 23: SKS Mechanical Configuration (from [114]).

There are limitations to this strategy of simulating conditions directly at the joint. Since the various contributing forces from muscles & ligaments are modelled as an 'aggregate' load (and not individually incorporated), the loading and restraint applied in the simulator are not truly physiologically representative. Consequently, operating outside of the intended *in-vitro* conditions can result in non-physiological mechanics; for example cadaver knees loaded in the SKS (using the knee's natural ligaments in place of the standard horizontal springs) have exhibited very exaggerated kinematics [115]. Nonetheless, knee implant wear simulators have become established as the *de-facto* standard for pre-clinical implant testing.

Limb Simulators

Full lower-limb simulator rigs have a more extensive modelling scope; by applying loads at the hip and the foot/ankle, they can more realistically account for such factors as variations in muscle forces or component positioning, since the actual knee joint is not directly (artificially) constrained, and has all six potential degrees of freedom of motion. However, these simulators are inherently more complex, since more of the lower limb is modelled, requiring more components (to represent the full thigh, shank, hip and ankle) and more actuators (to provide representative loading) to be physically incorporated. This requires more comprehensive understanding of the behaviour of the equivalent biological elements to model them correctly. It also places demands on the physical engineering of the rig, which should aim to match the inertia, strength, speed and power of the natural lower-limb (a difficult challenge, owing to the high performance of the target biological system).

An important benefit of the full lower-limb test rigs is that the simulator can be driven directly with clinical data. Ground plate reaction forces & torques can be applied at the 'ankle', and motion-capture video recordings and inverse-dynamics then used to determine corresponding kinematics or kinetics at the more proximal joints. Any study intending to vary the muscle forces and limb moment arms would be much simpler to implement on such a rig. However, since the kinematics at the knee are not directly driven (but are determined by loads and motions at the hip and ankle), it can be difficult to reproduce specific implant kinematics & kinetics at the knee. Hence it is more difficult to directly match the input waveform profile of a knee-wear 'joint' simulator to a corresponding 'limb' simulator; so the two cannot easily be directly compared. (A limb simulator must of course include both the tibiofemoral and patellofemoral joints, so in theory it should be possible to compare results with either a patellofemoral or tibiofemoral knee joint simulator).

Early examples of these whole-limb simulators emerged in the 1970's (e.g. Shaw & Murray demonstrated a single-axis, manually-operated mechanical rig for quadriceps-driven F-E as early as 1973 [116]), with subsequent designs becoming more sophisticated. The configuration adopted by Perry in 1975 [117] fixed the 'ankle' with a sagittal-plane hinge and allowed the 'hip' to translate vertically (see Figure 24, left). The Oxford knee rig [118] was also based on this same configuration, but including additional degrees of freedom for out-of-plane loading (see Figure 24, right). These rigs were designed to be used for quasi-static analysis; i.e. the rig did not have the capability to be dynamically 'actuated' by loading the knee with representative muscle forces.

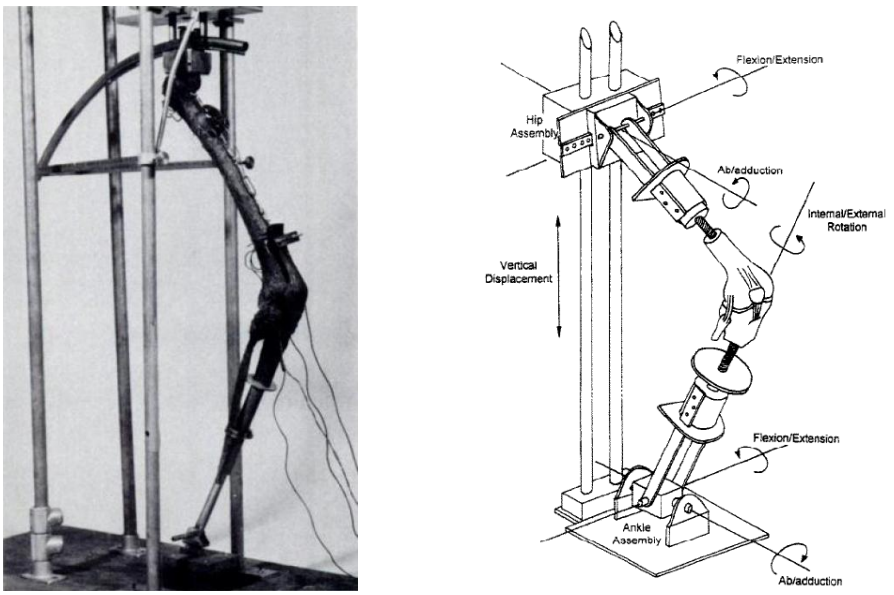

Figure 24: Left: Perry's knee testing fixture [117]; Right: Oxford knee rig [119].

One of the more advanced rigs of this type is the Perdue/Kansas design, which is a servo-hydraulic powered *dynamic* five-axis simulator, driven by the forces at the ankle and hip. The original Perdue Knee Simulator (PKS) was designed as a next-generation wear testing station in the late 1970's (see early work by Zachman et al [120]). The conceptual design of this rig is shown in Figure 25. The 'shank' & 'thigh' are fixed to sliding 'hip' & 'ankle' sleds. Four actuators drive the rig: in the sagittal plane, a vertical force emulates bodyweight, and a 'quad' actuator replicates the quadriceps muscles. At the ankle, actuators drive vertical rotation (equivalent to I-E for small flexion angles), and also M-L load, to produce adduction-abduction (Ad-Ab) moments, hence allowing out-of-plane loading to be applied.

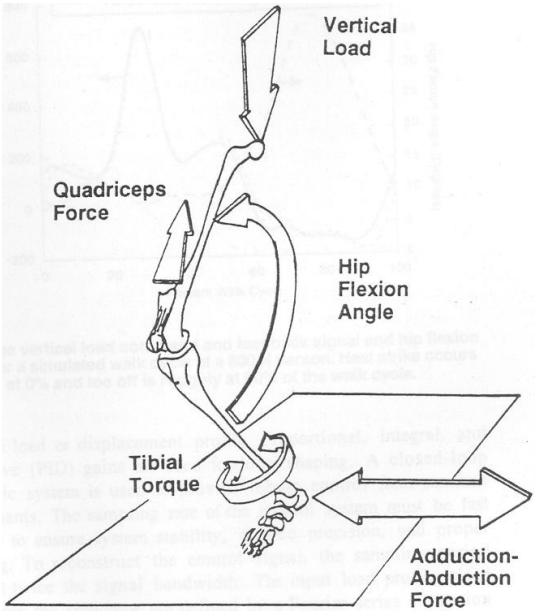

Figure 25: Mechanical configuration of the original PKS (From Zachman [120]).

Although the rig has gone through various re-builds subsequently (to decrease weight and increase flexion range) the only fundamental design alteration to the simulator has been to include an additional sagittal-plane actuator to provide ankle F-E moments. This actuator (mounted anterior of the distal tibia) is not physiologically representative of specific muscle groups (*in-vivo*, ankle plantar-flexion is provided by the *gastrocnemius* & *soleus* muscles situated posteriorly in the shank). However, it does compensate for the lack of an antagonistic 'hamstring' force to counter the quad actuator, and provides a means to apply a strong flexion moment when the knee is close to full extension (normally the 'body weight' applied by the vertical actuator can produce a strong flexion moment, but not at full extension). This five-axis version of the simulator has been further developed by Maletsky et al [121]. The latest build now features more feedback data by the inclusion of a six-axis load cell to directly measure loads at the tibia [122]. This current Kansas Knee Simulator (KKS) design

represents a versatile platform for *in-vitro* testing, and has also been the subject of derivative computational models [109, 123].

2.4. Corroboration & Validation, and the Case for Stochastic Analysis Approaches

The previous sections have illustrated the wide range of extant pre-clinical analysis models in use, both computational and experimental. However, no single form of testing is sufficient in isolation. *In-silico* studies in isolation are subject to suspicion as long as there is no consensus on the precise causal mechanics of wear. But *in-vitro* studies alone cannot provide the range and volume of information which can be quickly and efficiently evaluated computationally. Rather, the combination of *in-vitro* and *in-silico* wear prediction methods corroborated together provides a better, more extensive toolset for pre-clinical analysis of TKR wear.

It is important to make a clear distinction between 'corroboration' and 'validation' of a model. A model is only truly 'validated' when it matches 'true' reality (i.e. in the case of knee assessment, post-clinical *in-vivo* outcomes) – this may be considered the 'gold standard'. However, when it is not possible to directly validate a model (due to difficulties of capturing data in the real world), two independent models may be 'corroborated'. This means they are in '*relative'* agreement with each other; which does not necessarily mean that they are both correct (compared to the '*absolute'* reality), but it does allow greater confidence to be placed in the models, as corroboration would reveal any obvious errors, mistakes or serious differences between the two (they would both have to be wrong in exactly the same way for the error to go undetected). As such, corroboration may be considered more of a 'silver' standard, which has its place earlier in the design process, or when validation is not practically possible.

Clearly there is a natural inter-dependence between the different forms of analysis. *In-vitro* studies can be used to begin investigating the underlying phenomena. Using this data, *in-silico* models can then be developed to match this behaviour, and then extrapolated to predict the system response under a wider range of conditions. A good agreement between *in-vitro* & *in-silico* models may be considered to 'corroborate' the models. Finally, a few selected *in-vivo* tests may be used when the understanding of the system is more mature, to provide this important 'validation' with the *in-situ* real world application of the system.

Figure 26 illustrates this in the case of knee assessment. Corroboration is possible early in the pre-clinical analysis stages between *in-vitro* and *in-silico* models, provided they model comparable conditions (i.e. based on the same mechanical conditions, using the same components under test). Comparing data between *in-vitro* platforms is not straightforward due to the many confounding factors (e.g. comparing POD tests to TKR wear tests), and experimental models are not close enough to the *in-vivo* reality to truly 'validate' them directly. Hence these interlinks are shown with dashed arrows. By contrast, *in-silico* simulations provide a parallel, complementary modelling 'domain' where cross-communication & data transfer between different stages is much easier. The 'silver standard' of corroboration between *in-silico* and *in-vitro* models gives greater confidence in the understanding of the test mechanics & modelling domain, however ultimately the aim is the 'gold standard' of validation with clinical performance results. Accurately reproducing the complex holistic *'in-vivo'* environment is beyond the capability of experimental methods; however *in-silico* musculoskeletal models may ultimately provide the best means to achieve this. In order to produce computational models of sufficient quality, the 'early stage' corroboration with *in-vitro* methods is essential however. *In-vitro* tests continue to be necessary, to provide the real-world grounding for the explicitly defined mechanics of the *in-silico* models.

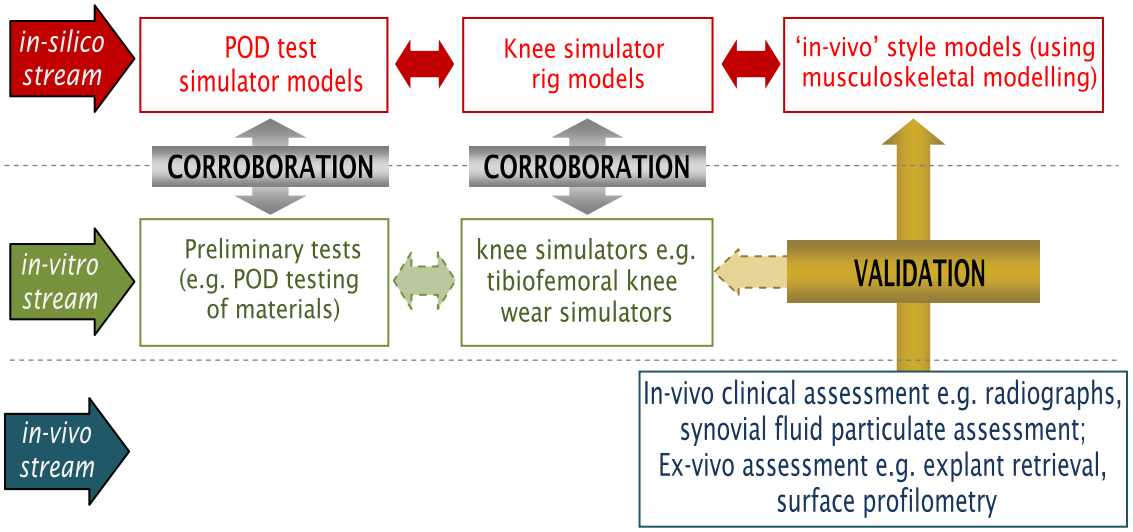

Figure 26: The role of corroboration and validation - conceptual diagram.

The value of corroborative studies has been recognised by a number of researchers who have published *in-silico* models based on the experimental results of *in-vitro* simulations. For POD models, this work has been performed by

a number of teams, including Hamilton et al [124] who corroborated the work of Turell et al [105], and Kang et al [125], who used their own proprietary data. For TKR simulations, computational-experimental corroboration has been performed by amongst others Godest et al and Halloran et al for kinematics, [91, 94], Knight for displacement driven wear [126] and long term wear [92], and Willing & Kim, Hamilton et al and others for force-driven wear [124, 127]. Generally, the aim is to demonstrate that the *in-silico* models & methods produce comparable results to the experimentation, so that the computational model can then be extended to be used for further investigation; in other words, the aim is not just duplication, but to produce a quality *in-silico* analysis tool which can be used to augment the experimental capability

However, in all of this discussion a key consideration has been excluded up to this point; namely, the *in-vivo* domain is inherently highly variable; there can be no single 'validation' with clinical data, since the data in any two cases would be different (different patients, different mechanical loads, different surgical outcomes, different activity levels). This is a major confounding factor when attempting to 'validate' a model; a single-run on an *in-vitro* test platform cannot possibly reflect this range of outcomes. *In-silico* models, used as a complementary analysis tool, have the speed and power to run multiple cycles (e.g. with probabilistics to explore variations in alignment or loading conditions), but it is essential that they are well-corroborated in order for this data to be meaningful. Therefore, the *in-silico* modelling domain has the potential to act as a 'bridge' between the *in-vitro* and *in-vivo* domains, introducing stochastic analysis approaches to the models derived from the experimental laboratory.

2.5. Summary

The use of experimental and theoretical methods for *in-vitro* and *in-silico* pre-clinical analysis is well-established. These various models have been demonstrated to provide useful analysis and predictions of knee behaviour whilst operating under normal knee conditions. However, as has been stated, high variability is an inherent feature of biological systems such as the human knee, and in order to provide a truly complete picture of knee performance, analysis should not be limited to 'normal' conditions only, but should include more complete variation in factors such as loading, geometry and alignment.

Most of the studies cited above are either deterministic (operating for a single case with specific known inputs), or parametric 'one at a time' studies (sweeping across a range of values for one variable or a small group of related variables). Only limited work has been done in recent times to extend these models across the entire domain of uncertainty associated with TKR. But without this 'holistic' perspective, it is impossible to be confident that the system is fully characterised by the current body of literature. Therefore, a more complete stochastic analysis of TKR is called for, to map the areas of variability and cross-coupling effects not explored by existing studies. This requires further extensions to the current body of knowledge; to understand the sources and levels of variability within the system, and also to demonstrate the application of stochastic techniques using knee mechanical models & statistical data. This stochastic framework will be explored in the next chapter.

CHAPTER THREE – PROBABILISTIC METHODS [2]

Sources of variability & review of numerical probability integration methods

3.1. Concepts of Probabilistics

To facilitate a discussion of the case for stochastic modelling, it is necessary to begin with a few definitions of relevant concepts. Any given model may be reduced to an "input-system-output" paradigm; certain influences on the system will influence the resulting output states; these input values may be termed input *factors*. Input factors may have known 'fixed' values which are accurately measurable and controllable. If all the factors are of this nature, there is no uncertainty and the model may be described as *deterministic*. However, if there is a degree of uncertainty in one or more of the factors, this input variability makes the system *indeterminate*; in which case stochastic (or 'probabilistic') modelling may be applied. The variable factors may be denoted individually as X_1, X_2, X_3 ... X_N, or collectively as a single vector value, $\mathbf{X}$ (or $\vec{\mathbf{X}}$).

If a factor does not have a single fixed value, it may take a continuous or discrete range of values. For a model with *N* variable factors, the range of possible combinations may be represented in an *N*-dimensional region of space, with each factor forming a separate orthogonal axis – the resulting representation is the *design-space*. Figure 27 illustrates this concept for a 2-dimensional example, with two variable factors, X_1 and X_2. X_1 can take values from -3 to 3; X_2 can take values from 0 to 5; the resulting design space is a 2-D surface, where every point on the surface represents a different unique combination of the two variables.

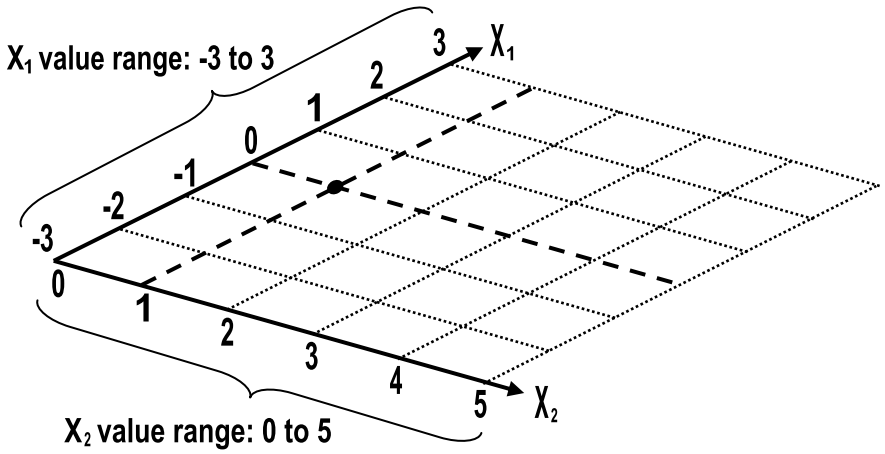

Figure 27: Illustration of the design space concept in 2-D.

[2] *The concepts presented in the following section are derived from various standard texts on the subject of probabilistic modelling & reliability theory [128-133]. Any of these may be consulted for further information on the techniques of stochastic analysis.*

However, the design space only indicates whether a combination of values is *possible*; it does not contain any information about whether that combination is *probable*. For this, it is first necessary to know the probability density function (PDF) associated with the different factors. A PDF is a function spanning the range of possible values for X, the magnitude of the PDF (p_x) indicates the probability associated with a given value of X. A PDF may take any form, however a number of 'standard' types are commonly encountered (e.g. Normal or Gaussian, lognormal, Poisson, binomial, Weibull, Rayleigh, etc). Figure 28 assigns two different PDFs to the variables in the present 2-D example; X_1 takes a Normal distribution with mean of 0 and standard deviation of 1; X_2 takes a lognormal distribution with mean 1 and standard deviation ½. (Note the PDFs have been clipped slightly to fit within the bounds; in reality they extend to ∞). Note that, for a multi-factor problem, it is possible that the PDF for one factor could change depending on the value of the other factors; in this case this coupling of the factors must be accounted for. However, in the absence of better data, it is often assumed that factors are *independent*; i.e. that variations in one have no effect on the others, and hence the PDF does not change.

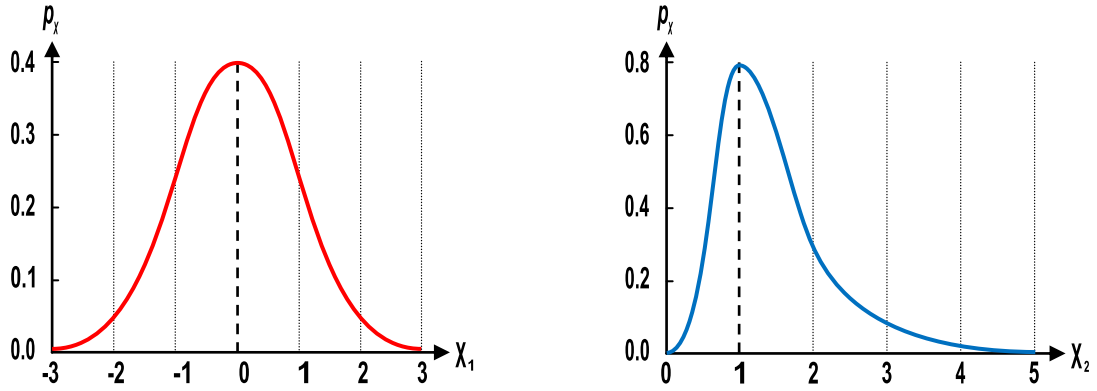

Figure 28: Two typical PDFs; left – Gaussian, right – lognormal. Note: the area integral under any PDF is always unity.

Combining the geometric mapping of the design space with the information in these PDFs allows the construction of the *possibility space* – this is again shown for the 2-D example in Figure 29. The advantage over the design space is that it is now apparent which combinations of variables are more or less likely; so for example, events around the region $X_1 = 3$, $X_2 = 5$ have a very low associated probability of occurrence (as a result it might not be too relevant how the system performs under these conditions).

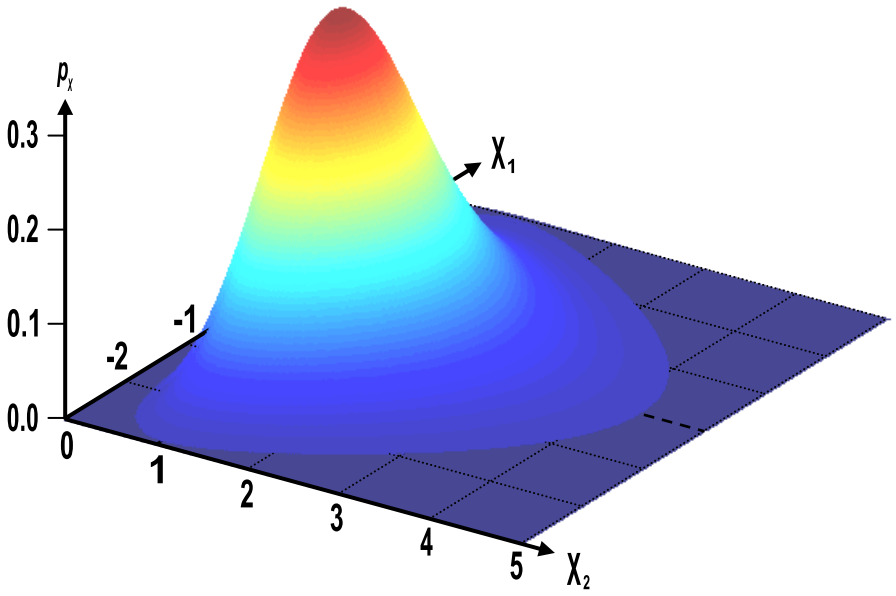

Figure 29: Illustration of the 'possibility space' in two dimensions. Note that the volume-integral of the PDF across the possibility space is always unity.

3.2. The Case for Probabilistic Analysis

The studies described in Chapter Two included a wide range of deterministic studies (i.e. considering only the 'neutral' case without regarding any perturbations of input factors). Historically, when *in-vitro* & *in-silico* models were first developed, the aim was simply to 'validate' these simulations with a single 'target' output, considered to represent 'typical' real-world conditions (for example, [91, 92, 112, 114]). Of course, such studies are an essential first step, but beyond this, they provide no information about the effect of any variability.

As the science of orthopaedics matured, it became desirable to better understand the influence of various factors identified as important. Studies began using parametric 'one at a time' sweep methods (varying one factor, or a small number of factors, across a range of values), for example, [134, 135]. This provides a valuable first indication of the factor's influence. However there are two limitations of such methods.

Firstly, 'one-at-a-time' studies are *decoupled* from statistical information about the PDF for the input factor. Trials evenly-sampled across an input range do not give information about the probability of a given outcome; for this, information about the probability of the input conditions is also required. (It is of course possible to perform a one-dimensional 'sweep' study with better selection of input values based on measured statistical distributions of the input parameter, but in the field of orthopaedic research, historically many studies have used regular step sizes, taking no account of the true input PDF, e.g. [136, 137]).

Secondly, parametric 'one-at-a-time' studies fail to map out the entire possibility space. This can be readily visualised in 2-D for two variables (Figure 30).

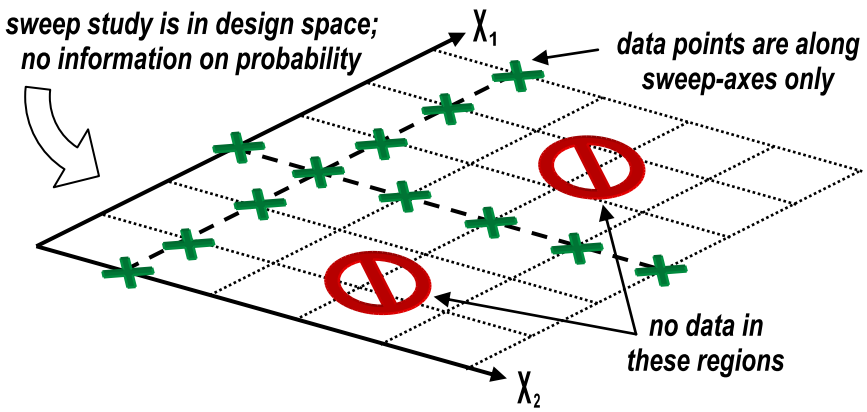

Figure 30: 'One-at-a-time' studies (above) provide no information about probability distributions, or the correlations between factors across multiple dimensions.

This introduces the need for 'stochastic' or 'probabilistic' studies: the entire possibility space can be investigated, and the output data will be related to input variability, giving corresponding statistical data (see Figure 31). However, as is also clear from the illustration, this greatly increases the space that must be explored, and this increases according to a power law of the number of input variables; 3 variables gives 3-D space, 4 variables gives 4-D space, etc. This rapid increase in the scale of the task is often referred to as 'the curse of dimensionality', and presents a serious challenge to stochastic study design. To address this challenge, a number of different methods exist for implementing probabilistic studies. These will be discussed in the next section.

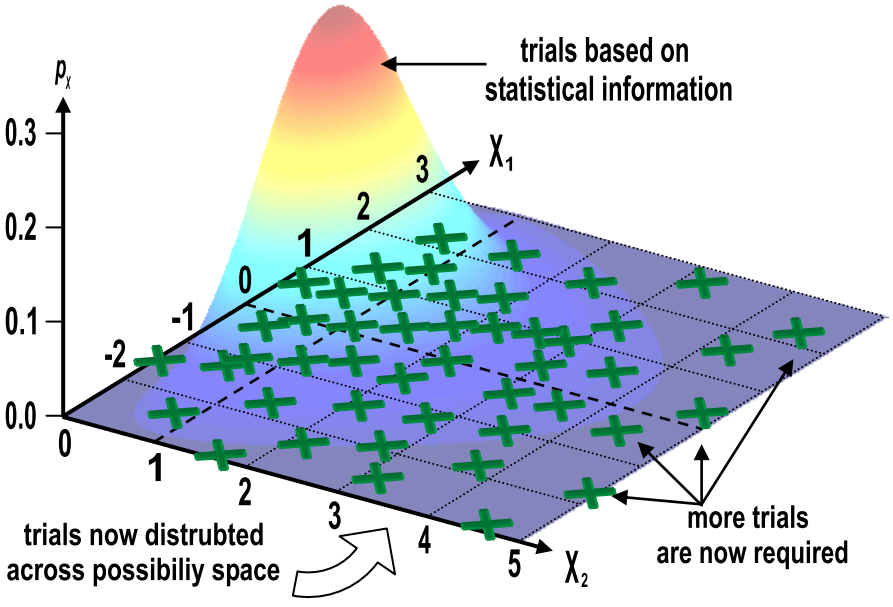

Figure 31: Example 'probabilistic' study; samples are distributed across the possibility space, based on the PDFs of the input factors. Note that this requires more trials than the deterministic or 'one-at-a-time' studies.

3.3. Numerical Approaches to Modelling Probability [3]

Most computer-based statistical modelling relies on the concept of 'numerical probability integration' – in other words, the summing (integration) of individual numerical trials, to approximate the true (analytical) probability. The method of probability approximation by numerical integration of statistical samples has only been applied to the specialist field of orthopaedic biomechanics relatively recently, as the computational resources it requires have become more readily available. The most established method is the 'Monte-Carlo' simulation technique (MCST), which uses purely random samples across the possibility space. However, this is computationally intensive, and consequently other methods have been devised, which can broadly be split into two categories. Importance sampling methods (ISM) fundamentally use the MCST approach, but improve efficiency by selectively reducing the sample-space based on knowledge of the system. Fast probability integration (FPI) methods are alternative approaches which are more approximate, but more efficient.

3.3.1. Monte-Carlo Simulation Technique

The Monte-Carlo technique is essentially a 'gamble'; the approach uses brute force rather than careful selection of trials to achieve a good result, relying on a very large number of trials to achieve high-fidelity. Random (or pseudo-random) samples are created, based on the known (or assumed) PDF associated with each input variable. These values are then used to generate associated output values.

As with other approximate numerical-integration methods, the 'resolution' of the integration determines the accuracy of the calculations; typically many thousands of trials are required to obtain useful results, and the number of trials required will increase when the probabilities involved are small. Once the trials are completed, output distributions (mean, standard deviation or specific probability levels) can be determined readily, e.g. if the measure of interest is a probability of failure (p_f), this can be estimated by taking the ratio of failures (N_f) to total trials (N); see Figure 32. Alternatively, the outcome associated with a particular probability range can be calculated, by taking the p^{th} percentile of the trial results. Intuitively, accuracy increases as the number of trials increases (with reference to the figure, a larger number of trials gives a higher 'resolution' image of the possibility space).

[3] *Figures in this section adapted from the conference paper: "Probabilistic Computer-Aided Analysis of Variables Affecting the Performance of Total Knee Replacement". Strickland et al, 2006, Biomedical Futures 2006 – Musculoskeletal Biomechanics (Durham, UK)*

Note that the figure illustrates how MCST reveals the path of the *'limit state'* function, sometimes denoted g(**X**). This is the boundary between 'success' and 'failure' trials (although the term may be used for the threshold between any two distinguishable system outcome states; e.g. achieving a particular performance-level or not). Only an approximation to the 'true' limit state is obtained with numerical integration methods.

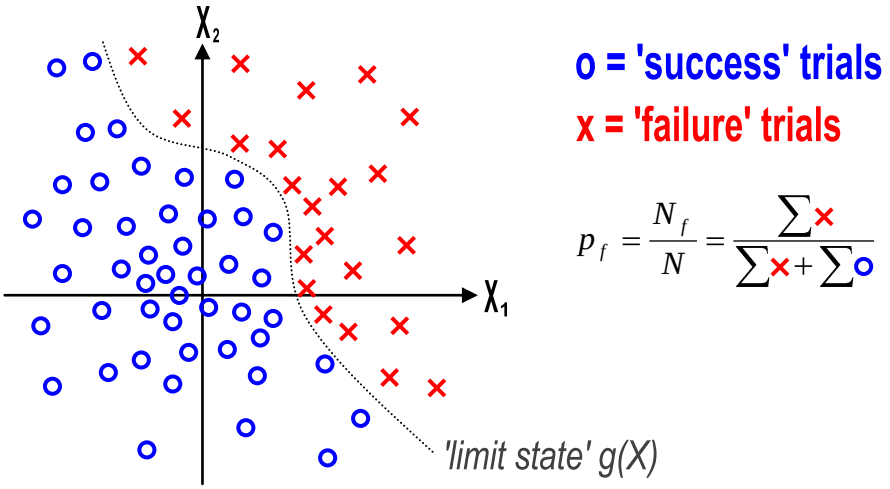

Figure 32: Estimating Probability of Failure (p_f) via MCST.

It is often important to know the accuracy of these probability estimates, and this can be ascertained approximately by calculating the 'coefficient of variation' (COV). This uses a binomial approximation to model the maximum possible error in MCST, for a given number of trials, N, and probability of failure, p_f:

$$COV = \frac{\sqrt{\frac{(1-p_f)p_f}{N}}}{p_f} \qquad (1)$$

Note that it is clear from this equation that COV will approach zero as N tends to infinity; however small values of p_f will also yield a larger COV. A smaller COV means the results of the MCST are more accurate. Hence the accuracy of MCST suffers if the number of trials is low, but also if the probability being estimated is very small. A typical numerical relationship between N and COV is shown in Figure 33.

Note that the when applying the COV measure in practice, p_f is not known, and so the estimate generated by the MCST trials must be used. The danger is that this value may not be accurate, especially for low values of N; therefore the apparent relationship does not match the 'smooth' theoretical relationship (in fact the COV may seem quite low for some very low values of N) – however this is erroneous; COV must be used with care for small values of N.

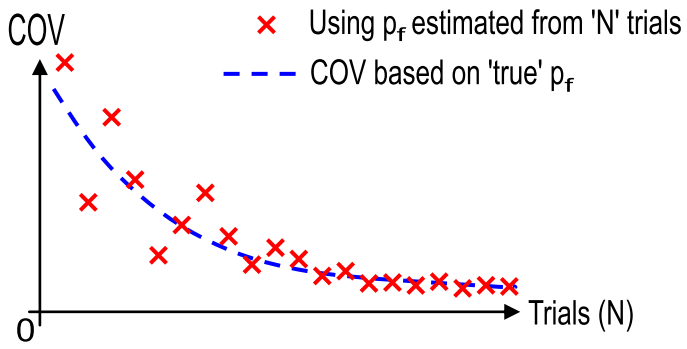

Figure 33: Typical relationship between number of trials and coefficient of variation (actual & estimated). Note the estimate may be above or below the true COV, but converges towards the true value (as N → ∞ and true COV → 0).

3.3.2. Latin Hypercube Sampling and Orthogonal Sampling

A variation of MCST which is sometimes used is Latin Hypercube Sampling (LHS). Rather than distribute the trials entirely at random across the possibility space, LHS attempts to distribute them to ensure an even coverage of the possibility space. Figure 34 illustrates this principle with a 'worst case' example; whereas for a small number of trials it is possible for the trials to 'cluster' with MCST, LHS constrains the trials to be evenly spread over the possibility space in distinct partitions, so reducing this risk. This can potentially reduce the error in probability estimates.

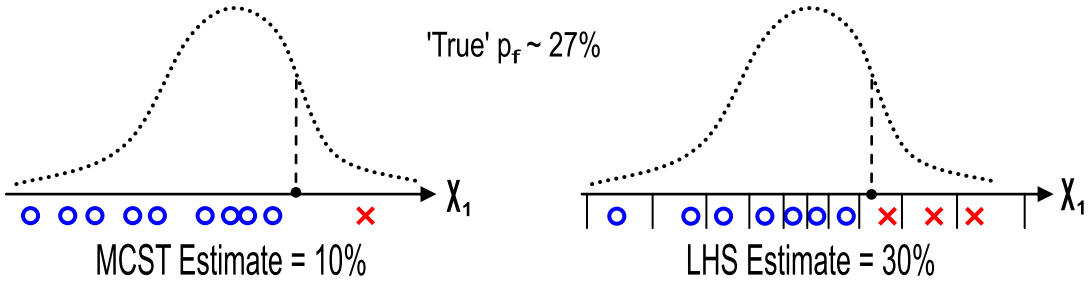

Figure 34: 'Worst-case' comparison of MCST (left) with LHS (right). Note in this example, the MCST error in p_f is several times larger than with LHS. This is an extreme case; the differences would generally be much less pronounced.

Within each partition, the sample may be taken at a random point, or using the mean or median value within the partition. For best results the LHS sampling is *weighted* such that the partitions are not of equal width, but rather of equal *area integral* beneath the PDF, i.e. the associated *probability* of each partition is equal. Consequently, partitions are smallest closest to the mean value 'peak' of the PDF, as in Figure 34. For problems with multiple dimensions, samples are selected to give a good statistical spread by ensuring that each sample falls into a unique row and column. The 2-D case is called the 'Latin Square'; the more general N-dimensional case is the eponymous 'Latin Hypercube'. To achieve higher sample

rates, this procedure can be applied with smaller bins, or else repeated multiple times with different LHS arrangements.

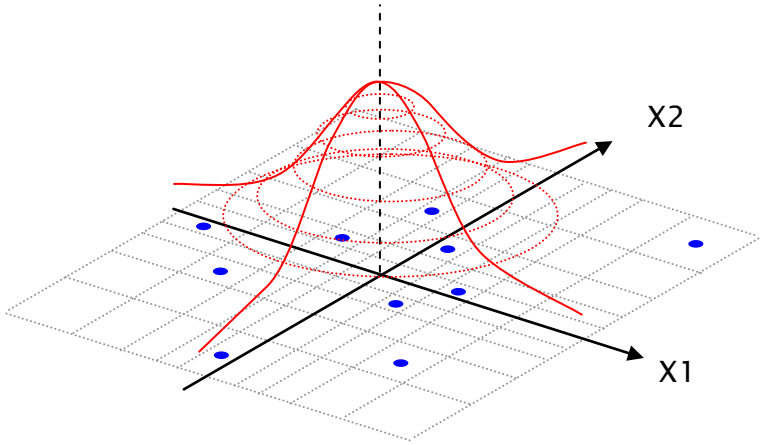

Figure 35: Example of 2-D LHS (or 'Latin Square'), showing variation of partition area, and also unique sampling in each row/column.

However, for problems with multiple dimensions, the LHS method can still result in 'clustering' (because the sampling between dimensions is independent). A further refinement of LHS is 'orthogonal' sampling (for problems with more than one dimension). Here, the possibility space is partitioned into smaller segments *across dimensions*, and the additional constraint imposed that an equal number of samples must be selected from each segment. Figure 36 again illustrates a 'worst case' example, comparing MCST, LHS and the orthogonal LHS method. Note that the orthogonal case should yield the most representative distribution of samples, although LHS in turn is generally more evenly distributed than MCST trials. However, this figure depicts an exaggerated case, since there are a small number of samples. The benefits of LHS & orthogonal sampling are greatest with a small number of trials. As the number of trials increases, the possibility of such a 'clustering' scenario with MCST decreases; all methods converge towards the true solution, so the difference becomes negligible.

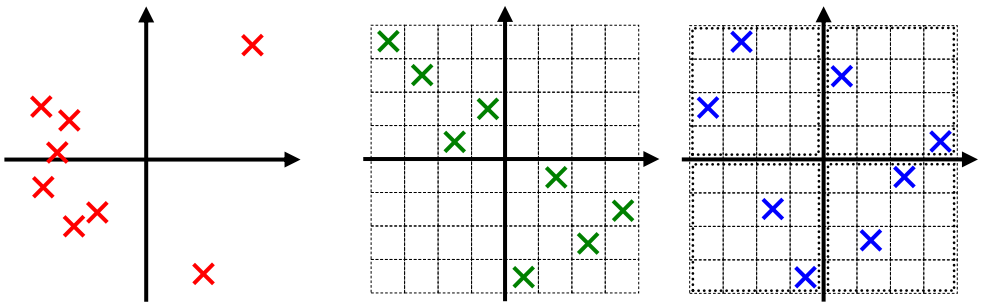

Figure 36: A 'worst case' comparison of MCST (left), LHS (centre) & orthogonal (right) sample sets; LHS & orthogonal methods are more robust against clustering, which is most apparent for small numbers of trials, as in this example.

Clearly, in the limiting case as N tends to infinity, the MCST family of methods offers an excellent solution, and may be considered the 'gold standard'. However, coming close to this solution may require many thousands of trials (depending on the problem & the required accuracy, the number of trials is routinely of the order 10^3, 10^4 or 10^5). This is often not feasible where the individual trials are computationally expensive; in such cases, techniques are sought to reduce the number of trials required for a given level of accuracy. Two sets of methods will be reviewed in the following sections: importance sampling methods, and fast probability integration methods.

3.3.2. Importance Sampling Methods (ISM)

An adaptation of MCST is ISM. There are a number of methods within this category; the common feature is that the possibility space is *not* fully explored; instead, trials focus on areas of interest, e.g. only assigning trials to areas on the 'fringes' of the possibility space. The effect is to multiply the accuracy; for example, if it is known beforehand *with confidence* that ¾ of the possibility space will not be associated with failure then samples can be focused in the remaining ¼, such that the same accuracy is achieved 4 times faster. An example of one of the simplest forms of ISM, the radius-based method, is shown in Figure 37. The failures all lie beyond a given radius from the mean value – therefore samples are not needed from this inner region, and can be concentrated on the outer region. The result is then scaled by the probability of the entire outer domain (p_d).

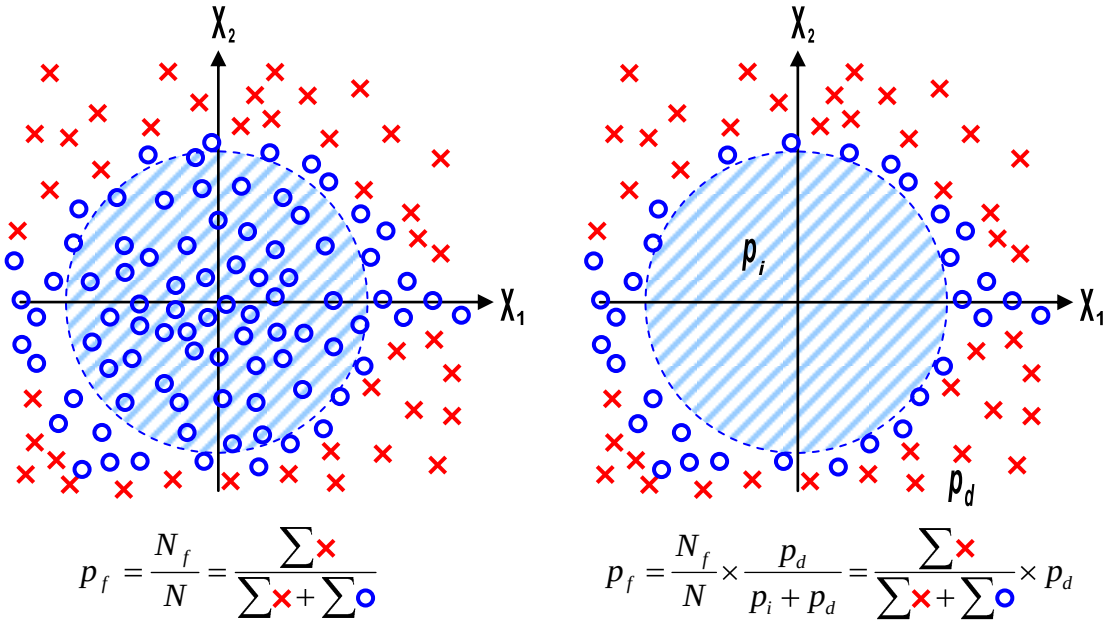

Figure 37: Comparison of MCST (left) with radius-based ISM (right). ISM achieves the same result as MCST with fewer trials; however, it is important to be confident that no failures would occur within the inner radius. Note that by definition, $p_i + p_d = 1$.

The risk is that any failures within the region presumed 'safe' will not be detected, so there must be sufficient confidence in the assumptions made. Note ISM depends on some additional knowledge of the system to reduce the sample space, so if the system is completely unknown, ISM cannot be used directly (since no region can be considered 'safe'). One possibility is to apply a low-resolution MCST to 'screen' the possibility space, before using ISM. Another, more sophisticated approach is to use 'adaptive' importance sampling, which gradually refines the sample space, based on new information obtained as the sampling progresses.

3.3.3. Fast Probability Integration (FPI) Methods

3.3.3.1. Response Surface Methods (RSM)

The response-surface modelling approach was first described in the literature by Box and Wilson [138], and in its most basic form is a none-adaptive, DOE-based FPI method (although more sophisticated variations of RSM are also now used). The concept of RSM is to fit a simple analytic function of the input variables to approximate the output parameter, across the full range of the sample space. Typically, this will be a low-order polynomial (called the 'response surface equation', RSE), and regression techniques will be used to select the term coefficients. However, it is possible to derive alternative forms of RSM, not based on simple polynomials but based on non-linear models with the outcome of the earliest trials being used to adaptively select the subsequent trials. (For more details on these methods, the reader is referred to the texts referenced in the footnote at the beginning of this chapter). These more sophisticated alternatives could be used for biomechanical problems in future, but will not be considered further for the concept studies discussed in this thesis; simple polynomial-based RSE methods will be demonstrated in the first instance.

Once a simple RSE is derived, this can be used as the basis for a MCST, since the RSE can be evaluated much faster than the true model. This method works best when the true output can be well-represented by an analytic function, e.g. very linear models can easily be fitted; highly non-linear systems are not well-represented. Figure 38 illustrates the method used to approximate a simple limit-state.

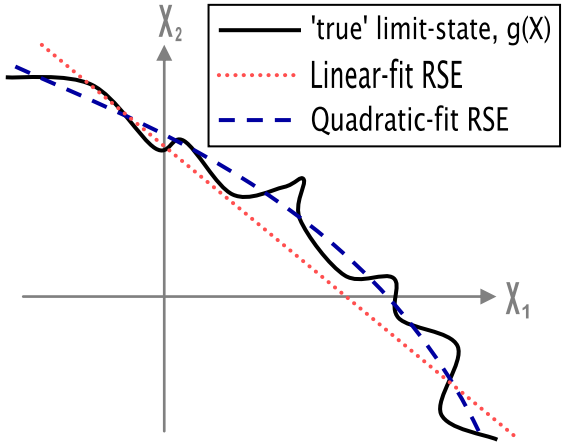

Figure 38: Fitting an RSE to a system response. The higher the order of the equation, the better the potential fit.

Trials could be random, but a better result is achieved by distributing the trials regularly across the sample space (e.g. using LHS or orthogonal sampling). The higher the order of the RSE used, the more terms that will be included; hence the more samples needed to achieve a good fit with the regression. For an N-dimensional model, the number of terms required for up to a cubic RSE is given in Table 4, in combinatorial and polynomial form (for quantitative comparison, the number of trials needed for a 10-factor system is also listed in each case). The actual number of trials used to achieve the RSE fit must in turn be several times this number, to achieve a reliable fit:

RSE Order	Combinatorial Expression	Polynomial Expression	Example (N = 10)
Constant	1	1	1
Linear	1+N	1+N	11
Linear with cross-terms	$1+N+{}^{N}_{2}C$	$1+\frac{1}{2}N+\frac{1}{2}N^2$	56
Full quadratic	$1+N+{}^{N}_{2}C+N$	$1+\frac{3}{2}N+\frac{1}{2}N^2$	66
Full cubic	$1+N+{}^{N}_{2}C+N+{}^{N}_{3}C+2\,{}^{N}_{2}C+N$	$1+\frac{11}{6}N+N^2+\frac{1}{6}N^3$	286

Table 4: Number of terms required for different RSE models (with N factors).

Clearly, higher order RSEs require more runs according to the highest-power polynomial term. Beyond cubic terms this becomes impractical for most models (the number of trials required is scarcely less than a low-resolution MCST approach). Results are generally not highly accurate, because the RSE is a *global* model; the same analytic function must approximate the output across the *entire* sample space. If accuracy is only required about one region of interest, a better

result may be achieved by fitting a *local* model at that point; this is the approach taken by a number of the FPI models, as discussed in the following section.

3.3.3.2. First and Second Order Reliability Methods (FORM & SORM)

FORM and SORM are based on the underlying assumption that, somewhere along the limit state function g(**X**), there is a region of most statistical significance (i.e. the conditions which are most likely to be responsible for failure) – this point is variously termed the '*design point*', or '*most probable point*' (MPP). The aim of F/SORM is to fit an analytic model at and around this specific localised point, to achieve higher accuracy.

To better visualise the concept of the MPP, consider a 1-D system, which 'fails' if the single variable (X_1) exceeds a certain 'limit state' value, g(**X**) (Figure 39, left). This can be generalised to 2-D and higher models. In the 2-D case, the limit state is no longer a point, but a boundary line across the 2-D space (Figure 39, right). In the general N-dimensional case, the limit state is a hyper-dimensional surface, generally termed the *limit-state surface*. All locations in the possibility space that lie beyond this surface constitute a 'failure'.

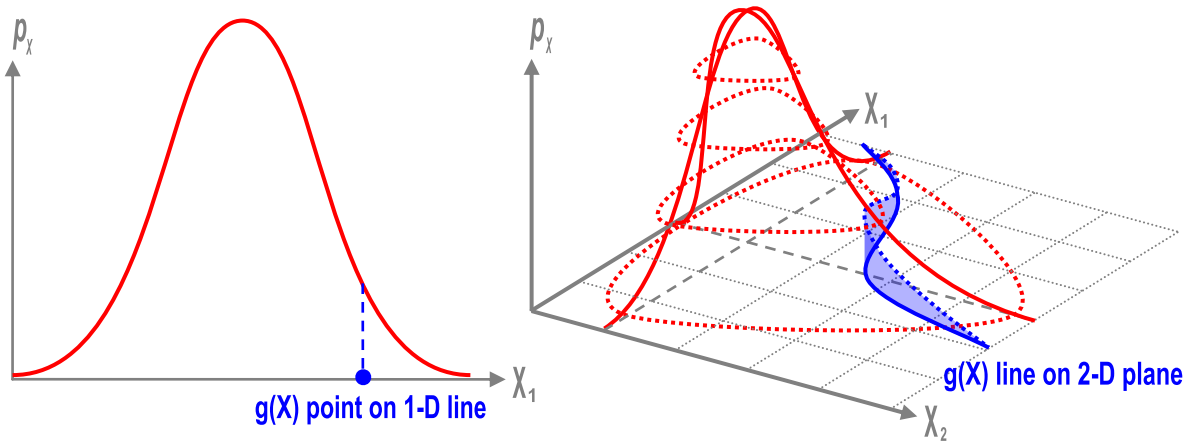

Figure 39: Limit state concept for 1-D (left) and 2-D (right) systems.

Now, in the 2-D case (and for higher dimensions), a PDF can be 'mapped' along this limit state (see the shaded region in Figure 39, right); this represents the limit-state probability of occurrence. There will exist a point at which the PDF reaches a global maximum; by definition this point is the 'MPP' of the limit-state.

Finding this point is complex if the input factors all have differing distributions; therefore the possibility space may be *normalised*, to re-map all the factors as normal distributions, with mean of 1 and standard deviation of 0. (The normalised input factors are then designated using '**U**' rather than '**X**'). Now, conveniently, the MPP on the limit-state PDF will be the point of closest approach to the origin, and can be located using geometric methods (Figure 40).

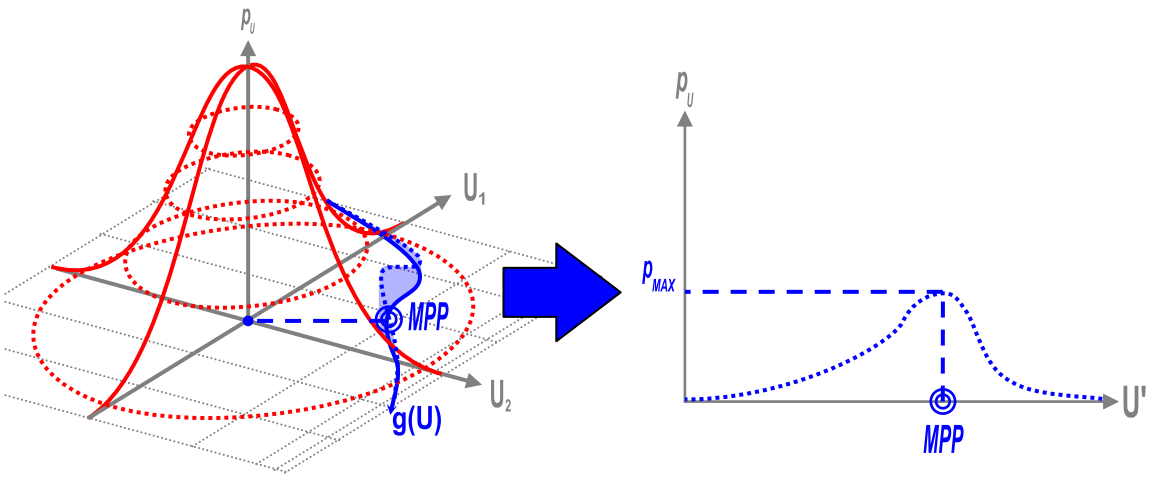

Figure 40: Locating the MPP in normalised (U) space, based on geometric proximity to the origin.

Typically an iterative approach is required to locate the MPP, which requires multiple trials. Once the design point is located, further trials are required to fit an appropriate analytical model. FORM fits a linear model, which is less accurate and converges to the design point more slowly, but requires fewer trials per iteration. SORM fits a higher-order model, requiring more trials to fit per iteration, but offers faster convergence and more accuracy.

In both cases, there are risks associated with the method. As with MCST, the input variables must be correctly characterised. Further, it is possible that the area of high probability along the limit state g(**X**) may be broadly distributed, such that no singular region represents the 'majority' probability of failure. (For instance, consider a broad, low PDF; many areas have moderate probability of occurrence, but no area is significantly the *most* probable; In this case the MPP/design point concept is less applicable). The MPP search algorithms also have limitation; for example, they may converge to a sub-optimal 'local' MPP which is not the true 'global' MPP. However, for well-conditioned problems, FORM & SORM can provide a much better localised approximation around the design point than is achieved by the global RSM approach.

3.3.3.3. Mean Value (MV) & Advanced Mean Value (AMV) Methods

The MV family of methods (MV, AMV, AMV+), again begin with an approximation of the function; in this case the approximation is made about the mean value of the input functions (i.e. the 'origin' of the possibility space). For MV, the expansion of the function takes the form of a first order (i.e. linear) Taylor-series expansion (requiring N + 1 runs for an N-dimensional problem), comparable to a first-order RSM. Note that here the linear model is applied directly to the factors in **X**-space, without converting to normalised **U**-space. The MV model is suitable

for fairly linear problems, but is not accurate for non-linear behaviour; its main practical use is as the basis for the subsequent AMV method.

The AMV method takes the linear model derived by the MV method, and attempts to include corrective terms to approximate the higher-order effects. However, unlike RSM, FORM and SORM, the AMV formulation does *not* provide a parametric function that can be applied elsewhere in the possibility space. Instead, it takes the MV prediction and (using data from the calculated MPP of interest), corrects this value for a *single point* in the possibility space. The higher-order approximation achieved by AMV cannot be applied at any point other than that for which it was derived; hence estimating probabilities for additional points requires additional applications of the AMV method.

Figure 41 illustrates the application of MV and AMV, for a 2-D domain. In illustration a), the mean centre-value is evaluated, along with small perturbations in the two variables X_1 and X_2 (1+2=3 evaluations). This gives the MV model of the output objective (the sloped plane). In illustration b), a probability level has been chosen, and the approximate location of the limit state g(**X**) estimated with the linear projection from the MV model. The MPP for this limit state is found, and another evaluation performed at this point. This new value allows for a corrected AMV estimate of the output at the probability level. This has given a good estimate at the point in question with only four evaluations; however, the disadvantage is that no similar information is available for the rest of the possibility space or limit state function.

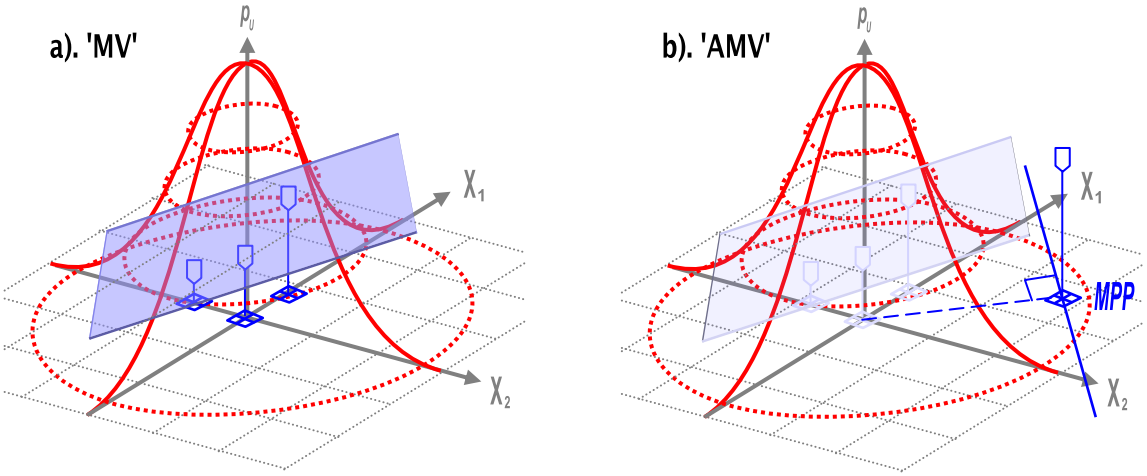

Figure 41: a) Deriving a linear MV model; b) Using AMV for a given probability level.

AMV+ is an extension of this method; essentially it is the AMV method applied iteratively, so each new estimate is further corrected with additional trials, to a specified error level – although this can quickly become computationally expensive. Due to the good convergence of the AMV method, the AMV+ method

is rarely required, except for very non-linear or non-monotonic functions, or when particularly high accuracy is required.

The MV method is fast, but not accurate for non-linear models. The AMV methods are efficient, but as with FORM/SORM, they rely on the MPP concept, so again can be confounded by local minima in the limit-state evaluation. The AMV method's impressively low computational cost comes at a price: this method only provides information for a single point, so in order to construct a full PDF the method must be applied repeatedly for every point of interest.

For any of the practical probabilistic methods discussed above, from full MCST to FPI methods, care must be applied when implementing a stochastic study; the results will *only* be accurate if the parameters of the various input factors, as well as the physical model, are representative of reality. Statistical properties such as mean, standard deviation, distribution type and inter-variable correlations must all be accurately characterised for every factor under study.

3.4. Visualising Probabilistic Results

A common feature of all probabilistic methods is that they generate a large volume of output data; instead of a single 'deterministic' value, a full PDF of possible values can potentially be constructed. For a system with multiple output objectives of interest (and especially if these outputs are vector values, e.g. time-varying metrics), this can potentially produce an overload of information, making the results & important observations less accessible. It is important to ensure that the stochastic study can still deliver simple, clear results or it will not be a useful tool for designers or clinicians. Therefore ways must be found to clearly present key results.

3.4.1. Performance Envelopes

For time-varying output objectives, a useful visualisation tool is the 'performance envelope'. Effectively, every time instant represents a unique 'system', with a unique stochastic behaviour and hence a unique PDF for each of the outputs, **Y**. This is too much data to present to the user; consequently, an effective simplification is to only display specific 'levels' on each PDF; e.g. the mean value, and an 'upper' and 'lower' bound (1 or 2 standard deviations, or a fixed percentile value). This is illustrated in Figure 42. The 3-D data is reformatted into a 2-D time-plot which is easy to interpret.

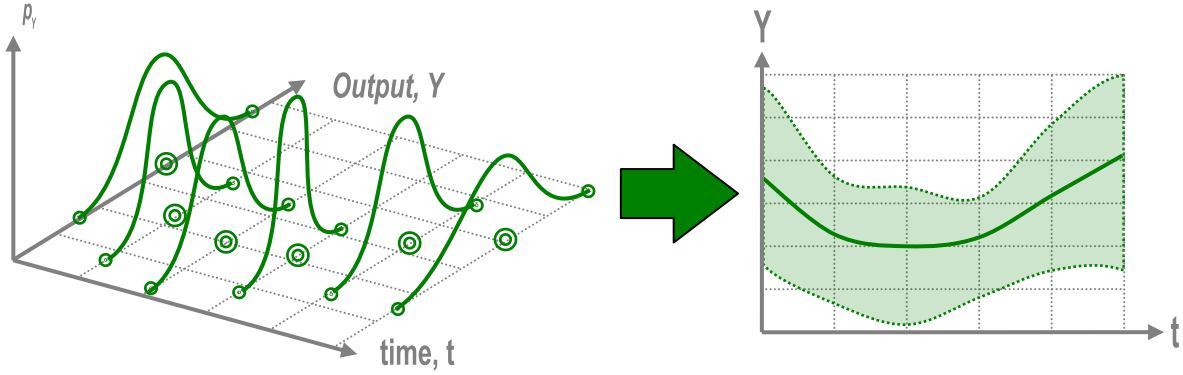

Figure 42: Concept of 'performance envelope' illustrations. Displaying only selected points from the PDF at each time-instant allows a simpler, more readily interpretable visualisation to be produced.

Note that when 'performance envelopes' are used, care must be taken in choosing the limits to display. Clearly 0% and 100% are not appropriate, as for many distributions these would extend to ±∞ and so provide no useful information. Values very close to these extremes are unlikely to be appropriate either; for any numerical probability approximation, generally the 'tails' of the distributions (i.e. furthest from the mean value, where the PDF is very low) are the most poorly approximated area. Hence larger errors are likely in these outlying regions. So if a model is based on a thousand MCST trials, a 1% to 99% envelope might be justifiable, but if only a hundred trials were used, the accuracy would be lower, and a 5% to 95% interval would be more appropriate. (The COV discussed earlier provides a useful means to quantify the error that would be associated with a given choice of envelope interval range). Conversely, if the bands are too narrow (e.g. ±1σ), they may not reveal the range of outlier cases (e.g. for a Gaussian distribution the first standard deviation either side of the mean excludes some 32% of all possible results).

3.4.2. Sensitivity Analysis (SA)

Another highly valuable technique is to report 'sensitivity factors'. The aim is to clearly illustrate which of the input factors are having most effect on a given system output Y, so that designers or clinicians can quickly see which of the factors is having most impact, and focus only on these factors. (Similarly, it may allow stochastic-study designers to determine which factors to preserve or omit for future probabilistic studies, if computational resources are limited).

Even for a single, scalar-value output objective (i.e. space- & time-invariant), providing a single 'bottom line' value for sensitivity is not straightforward. The value of the 'sensitivity' can vary, depending on the location in the possibility space (at any point, the *local* sensitivity is the partial derivative of the output-

function $\partial Y/\partial X_N$; this value can of course vary across the design space, depending on the shape of the response function). One may quote the sensitivity at a specific *local* point of particular relevance (e.g. the MPP from an FPI approach), or alternatively fit a *global* approximation (e.g. RSM), and quote sensitivity based on this global function. Obviously, the global sensitivity is less accurate, but more broadly applicable.

When a signal is varying across time or space, (i.e. vector-outputs), the task becomes still more challenging; the same factor may have very different effects at different times, or different locations in the system (e.g. in knee mechanics, a particular malpositioning of the components may increase pressure locally on one condyle, but decrease pressure on the contra-lateral condyle). It may not be appropriate to attempt to express a single-value of 'sensitivity' for such a case. Alternatives are to make a distinction between different 'regions' (spatial or temporal) and quote separate sensitivity values for each region; or to report only the averaged-magnitude of the sensitivity, to give a general indication of the overall influence of the factor. Nonetheless, sensitivity analysis should be applied with care in uncertainty analysis, to avoid providing misleading or overly-simplistic data. Note also that for systems where factors are heavily interdependent, quoting individual sensitivities is again misleading. Sensitivity factors are most meaningful therefore, for scalar-outputs of relatively linear systems with independent factors.

It is often desirable to display sensitivities for different factors alongside each other, to give a quantitative indication of the relative ranking of factors. In this case, it is important that the sensitivities are *normalised*, to compare like-for-like. Consider two factors, X_1 and X_2. If at some point of interest the system output Y is twice as sensitive to X_2 as X_1 (i.e. $\partial Y/\partial X_2 = 2\partial Y/\partial X_1$), it may seem that X_2 is the more sensitive factor. However, if the input variations in X_1 are actually ten times larger than X_2 (e.g. X_2 has a standard deviation ten times lower), then it is apparent that in reality the actual system response will be more affected by X_1 owing to its much greater variability. Hence normalised sensitivity is a function of both the raw sensitivity value and also the input variance for each factor. This requires possibility-space information (not just design-space), since the statistical properties of the input factors are required. Figure 43 shows a typical SA plot; note that if the sensitivity factors are based on a linearised fit then, by trigonometry, the squares of the normalised values (or β-values) will sum to unity.

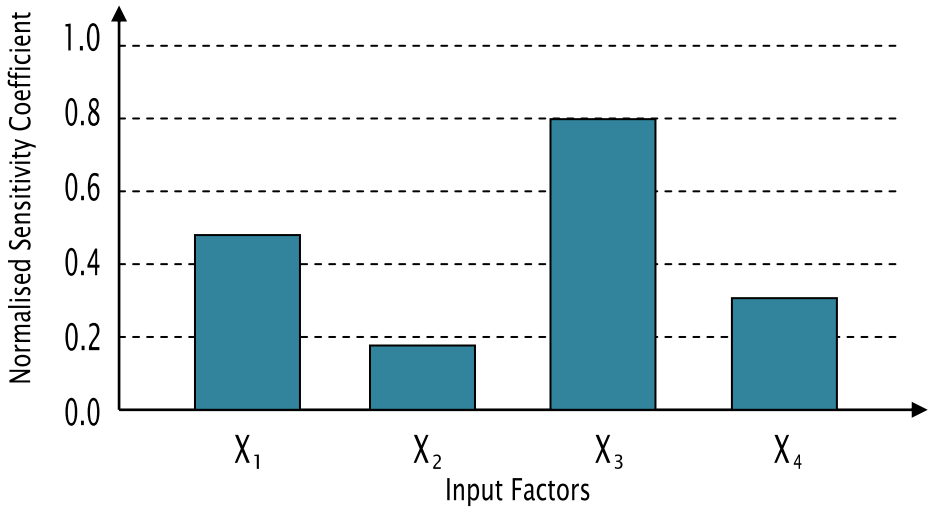

Figure 43: Typical sensitivity analysis results, displayed as a bar graph. SA can provide a very clear, accessible visualisation tool for stochastic analysis; e.g. it is immediately clear in this example which factor is most influential.

3.5. Considerations for Correlating Multiple Outputs

When there are multiple output objectives, designers or clinicians may wish to explore the correlation between the different outputs. This is particularly useful when it is difficult to directly assess the values of the uncertain input factors. For example, in the case of the human knee, directly assessing the individual stiffness, pre-strain or insertion sites of the ligaments is not readily achievable *in-vivo*, but these input factors will influence the behaviour of the knee in active gait (one set of 'output objectives'). Equally however, they will affect the passive laxity of the knee (a different set of outputs). It may be hypothesised that some correlations exist between these two different sets of outputs (active & passive); if this were proved to be the case, then by understanding these correlations it may be possible to infer the *probable* results for one output objective based on the known results of other outputs measured. (For example, assessing the passive laxity for some particular case may indicate that the knee is more likely to experience exaggerated kinematics post-operatively in normal gait; the surgeon may then be better informed to adjust the ligament balance intra-operatively).

In order to analyse such correlations, the outputs must be evaluated and compared under corresponding input conditions (so all factors are controlled, to be compared 'like-for-like'). There are two possible approaches; either the individual trials must be sampled at exactly the same input levels for the different output objectives, or alternatively a continuous functional expression must be built (e.g. an RSE model), so that after the trials are completed, the functions for the different objectives can be evaluated by re-sampling at co-incident points (see Figure 44). In either approach the final outcome is a series of coincident points to compare; however re-sampling with a functional model is much faster than

running original trials (as only the fitted function needs to be evaluated), so provided less trials are required to build the function, then this may be computationally more efficient.

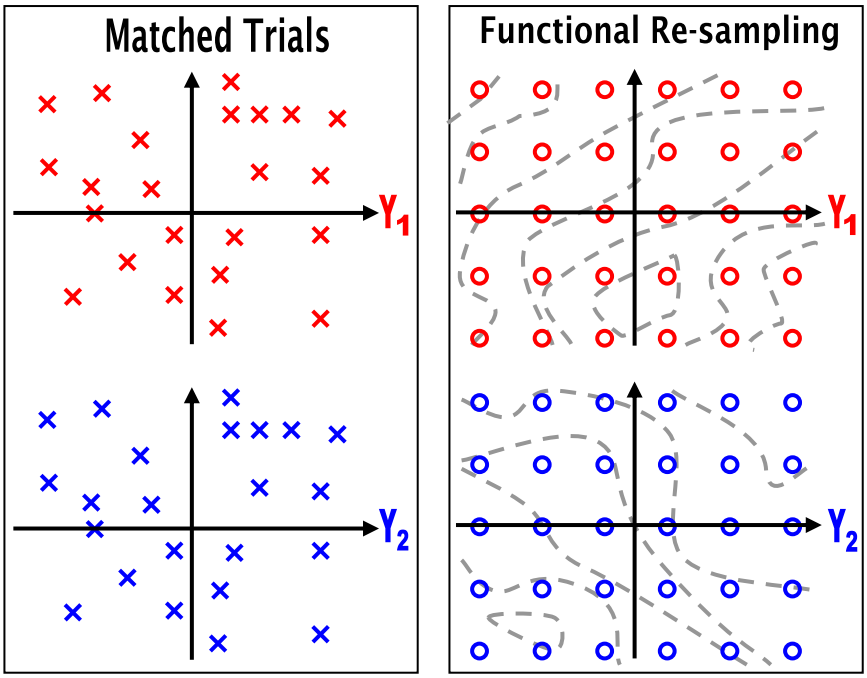

Figure 44: Alternatives for correlating outputs. Left: matched (coincident) trials; Right: function-fitting & re-sampling.

If the stochastic method uses pre-selected trials, and does not feature a recursive 'search' approach, then the first method with matched trials can be used (obvious examples are the Monte-Carlo or LHS approaches, where the trials are all specified before the evaluation begins). However, for the search-based methods such as FORM/SORM and AMV, (which use an iterative approach whereby the inputs for new trials are decided based on the value of the previous trials), this cannot be done, as each different output objective will result in a different iterative search path. In these cases, a functional model across the possibility space is necessary. Some models inherently accommodate this; for example, the essence of the FORM/SORM approaches is to build a first-or-second order regression model of the system; this can then be used directly as the function for correlating the two output objectives. However, the same is not true of the AMV family; the AMV approach only evaluates the model at a single input level and does not give any functional description across the possibility space. AMV is therefore unsuitable for determining multi-output correlations.

3.6. Performance Issues for Large-scale Stochastic Studies

The fundamental determining influence when choosing a numerical probabilistic method is always the performance-accuracy trade-off. Mechanical models of the knee are inevitably quite complex and multi-factorial; this means that there is

always a relatively high 'cost' associated with the probabilistic studies, and makes the trade-off between required solve-time and accuracy of the results more challenging. In practice, how this trade-off will be made depends on the computational resources available; for a full industrial deployment, with many thousands of processor-hours available, complex fully deformable models and intensive MCST methods may be a realistic option. However, for 'proof of concept' exploratory studies (such as the present work), much less computational power is available, so the trade-off must be more in favour of lower simulation cost, at the expensive of accuracy.

As has been discussed in this and previous chapters, there are several possible complementary strategies available to reduce the cost of the evaluation. One approach is to make appropriate simplifications to the mechanical 'physics' of the model (i.e. the actual TKR simulation). Examples are the use of rigid bodies, elastic-foundation contact algorithms or linear material properties (as discussed in Chapter Two). Alongside this trade-off in the mechanical domain, similar performance-accuracy tradeoffs can be made in the statistical domain too, as discussed in this chapter (e.g. using FPI methods).

However, other methods are also available which blur the distinction between a 'mechanical' and 'statistical' model in conjunction – two important classes are *surrogate models* and *statistical emulators*. A surrogate model is essentially still a mechanical model, but no longer necessarily modelling the causal physics. This is distinct from simplified mechanical models; the models in MBD or rigid-body FE are still based on underlying physics; the physics are just simplified for faster performance. In a surrogate model, there is no physically causal link between the input conditions and the output; often, it is reduced to a simple analytic function dissociated from any physical meaning (e.g. a response-surface style function). The use of such surrogate models has been explored for TKR mechanics [139], although it has not been applied in any published stochastic studies.

A statistical emulator is fundamentally different, in that it does not represent a mechanical model of the system at all; it is a purely statistical description of the mechanical simulation (see Figure 45). The emulator must be trained with a data set from the simulator to be emulated (this could be *in-vivo*, *in-vitro* or *in-silico* data), so a mechanical model is still required; however, subsequently the emulator can be used *in lieu* of any mechanical model. An advantage of this approach is that, because it is a purely statistical model, it is possible to associate a statistical error level with the prediction returned by the emulator – i.e. the emulator can predict its own accuracy. A disadvantage of these more abstract

modelling approaches is that, the further divorced they become from the underlying physics, the more difficult it is to verify their behaviour (e.g. they do not have mechanically meaningful intermediate outputs or states that can be corroborated with physical reality); similarly they are not well-suited to extrapolating outside of the 'known' physics into novel operating conditions (whereas a mechanical model may be able to extrapolate, within reason).

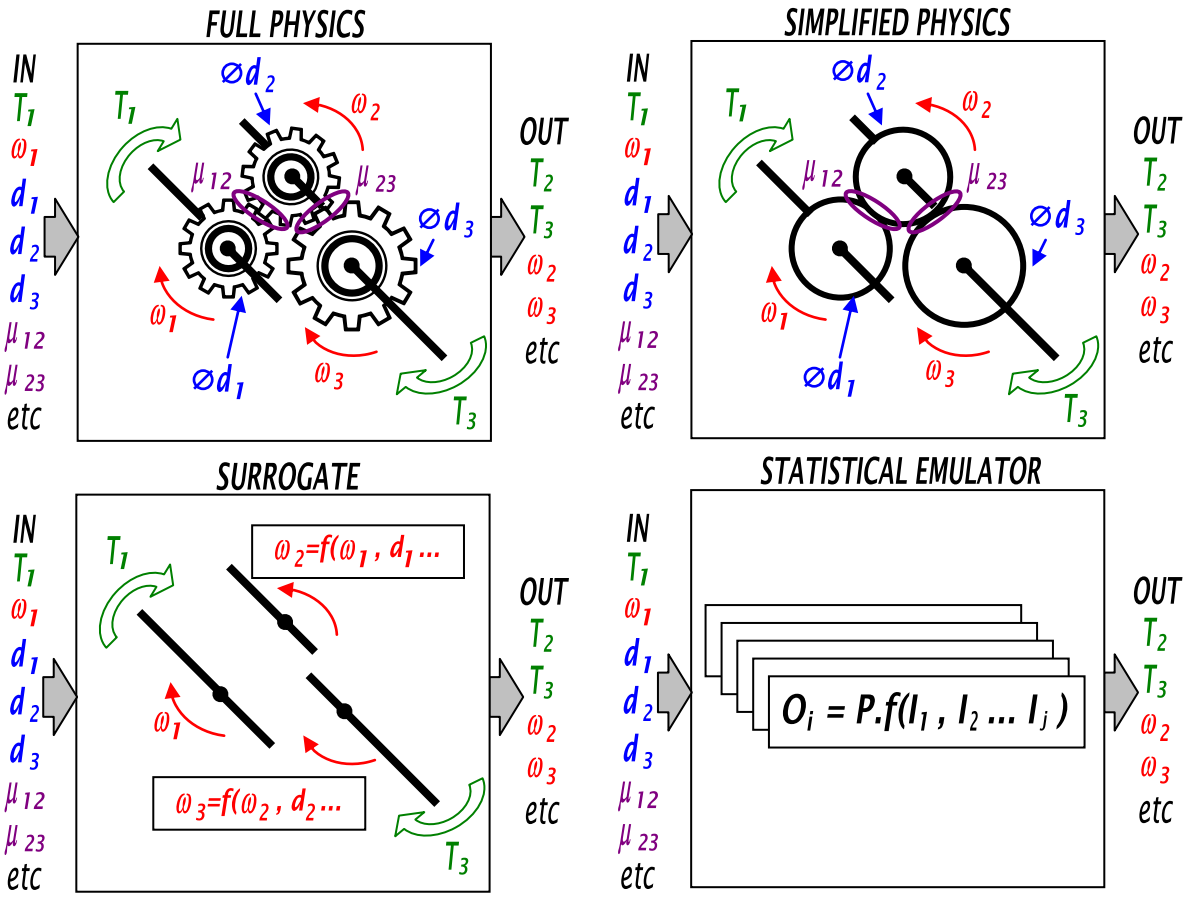

Figure 45: Rationale for surrogate & emulation techniques; emulation uses a purely statistical model of the system.

For initial conceptual work, as in the present studies, the use of a 'true' physical model (albeit simplified for computational efficiency) is desirable to simplify analysis, troubleshooting and debugging of technical development issues. However, for any subsequent high-volume work (e.g. a professional/commercial high-performance highly-automated probabilistic framework), entirely eliminating the computational overhead of mechanical modelling by using a statistical emulator may be an advantageous approach to consider.

3.7. Probabilistics Applied to Knee Biomechanics

In these studies, the application of interest for probabilistic methods will be knee biomechanics. Historically, probabilistic methods were first applied in the fields of more 'conventional' engineering (civil & structural engineering applications;

e.g. [140, 141]). Only relatively recently were probabilistic methods first applied to the field of bioengineering; initially in relation to structural mechanics of knee replacements (by Browne et al [142]). Other studies began applying probabilistic methods not just to the structural strength, but to the kinematics and kinetics of implants. These studies were first applied to the comparatively simple domain of THR mechanics (e.g. [143, 144]). Most recently, they have been applied to TKR models, as reported by Laz, Pal et al [95, 145]. These most recent studies represent the 'state-of-the-art', and will be the starting point for the development work in the present project.

Ideally, every knee, every patient, & every TKR would be identical. Then, the same remedial procedure would always result in a fixed outcome ('success' or 'failure'), and the design of implants and techniques could easily be adapted accordingly. The tremendous challenge of TKR is the amount of variability, in terms of implant design options & rationale, surgical procedure and inter-patient variations. Whereas a heavily automated process working exclusively with synthetic components can achieve a very high repeatability and very low tolerances, TKR is a specialist highly manual procedure, operating upon biological systems which can exhibit high levels of physiological and pathological variability. Examples of variables within these categories are listed in Table 5; despite this list being extensive, it is not exhaustive. Some of the variables represent *discrete* choices (e.g. to retain or resect the PCL); others represent *continuous* 'distributions' (e.g. subject weight). Clearly a very large number of factors can be influential, and where studies concentrate on a few 'key' variables, any number of uncontrolled or unexplored secondary factors can confound results, making predictions and recommendations less reliable.

Implant Design	Surgical Procedure	Patient Factors
Geometry (e.g. degree of articular conformity)	Surgical experience	Weight *(& weight changes post-operatively)*
Size *(standard sizing; component size mismatching)*	Malpositioning	Anthropometry (e.g. limb length)
Material *(polyethylene grade / cross-linking; stiffnesses, friction coefficients)*	Bone resection accuracy *(cutting errors affect malpositioning, component fit)*	Soft tissue conditions *e.g. Muscle forces, Ligament/Capsule quality, Ligament/Muscle Insertion sites*
Assistive surgical tool design *(affects surgical accuracy)*	Cement mantle quality (*porosity, thickness, coverage)*	Clinical / Pathological; *e.g. effect of RA Progression, Physiotherapy regime*
Fixation method	Ligament balancing	Patient recovery times
Tibial Bearing *(fixed vs. mobile)*	Soft tissue trauma due to incision (*surgical approach; MIS vs. conventional surgery)*	Post-operative lifestyle *(diet, exercise, activity level)*
PCL retention/resection	Patellar treatment *(preserved, resurfaced?)*	Range and frequency of daily activities
Manufacturing tolerances	Surgical Approach *(conventional vs. computer assisted)*	Post-operative gait adaptation

Table 5: A sample of factors influencing the outcome of TKR.

For accurate results, all of these must be accurately characterised, so that the statistical distributions (PDF shape, mean & standard deviation), and correlations between factors are known. This requires data on these sources of variability to be collected. A wide range of possible sources are available; for the pre-operative factors, information on design geometry and sizes is easily captured with CAD data sources; the variability in geometry is specified by manufacturing tolerances, which are generally well-documented. Variations in material properties can be acquired from materials testing standards (e.g. NISTS standards for UHMWPE [146, 147] indicate variations in material strength & stiffness), or from other experimental testing (e.g. POD tests can give an indication of variability in friction co-efficient). Intra-operative variability factors have been recognised as an area of concern, and as such several studies have published relevant data measuring the variation in different malpositioning outcomes for TKR (e.g. [148-150]). Some patient-variability factors (e.g. bodyweight and limb anthropometry) are well-reported in large-cohort population surveys (e.g. the US NHANES [151]). Further, specialist studies have used cadaveric mechanical testing to determine the properties of internal biological structures such as ligaments & tendons; the ranges they report provide some indication of variability e.g. [152-155]. For *in-vivo* mechanics, some gait analysis studies with larger cohorts have included an 'envelope' of variability on the data (e.g. see [28, 29]), and occasionally other studies include variability effects (e.g. the telemeterised data from Taylor et al [33] or the step-rate data from Schmalzried et al [20]); however the available data is limited, and incomplete. For these conceptual studies, it will sometimes be necessary to estimate variability based on other comparable sources, and it will generally be necessary to assume independence between factors and normal-distribution of variation, in the absence of better data. An important conclusion for further work is that, ultimately, better data on variability will have to be sourced (or else directly measured) to achieve greater accuracy in future probabilistic studies.

Although the available input data for the present studies is not ideal, an important outcome for this study will be to demonstrate a methodology allowing engineers and surgeons to have visibility of the complete scope of input variability and its effect on kinetics and kinematics, in particular in terms of sensitivity. This may help to identify areas to focus research efforts or procedural guidelines; for example, if a particular implant is highly susceptible to variations in one or a handful of factors, the designer could focus analysis efforts on these factors and their influence.

3.8. Conclusions

Clearly there is considerable scope for a study of variability in TKR. The number of variables involved means that a study will require considerable computational effort; this can be lessened by choosing appropriate probability integration techniques.

It has been seen that there are a number of standard techniques for numerical probability integration. While MCST and its derivatives represent the most reliable method, this is a computationally expensive strategy to pursue. For any complex numerical engineering problem, the faster alternatives are worth investigation. The results of different methods will be compared in subsequent studies, but as has been discussed there are risks and disadvantages with every method; therefore for the early investigations, these methods will always be validated with an MCST analysis.

The following chapter will describe how early work has applied some of these stochastic techniques to simplified numerical rigid-body models of the implanted knee.

CHAPTER FOUR – DETERMINISTIC MODEL DEVELOPMENT & CORROBORATION

4.1. Defining the Model

4.1.1. Study Scope

As with any project, time and resources for these studies are limited. Because they are being approached from a mechanical engineering perspective, emphasis will be placed on the causal link between the input variability and the resulting kinematics & kinetics. In some areas, the further causal relationship between kinematics/kinetics and mechanical failure modes will be partially explored (e.g. wear prediction). However, to fully characterise the different 'failure' modes (as discussed in Chapter One) would require a much more extensive model including the causal links between mechanical performance and failure (e.g. interface de-bonding and bone re-modelling), besides other non-mechanical failure influences (e.g. infection, inflammation and necrosis). A systematic representation of the study scope is illustrated in Figure 46.

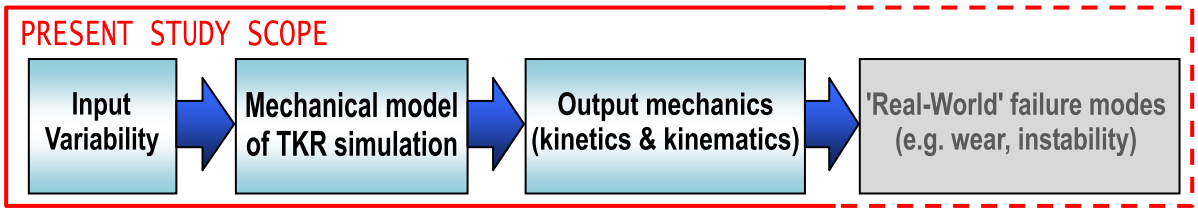

Figure 46: Scope of present study; dashed partition indicates 'partial' inclusion.

4.1.2. Modelling Environment (Software)

In Chapter Two, fast *in-silico* MBD models were identified as an appropriate platform for the stochastic study methods described in Chapter Three. A number of software applications support MBD modelling; the software used in these studies is MSC.ADAMS (MSC Software Corp) - "Automatic Dynamic Analysis of Mechanical Systems". It is widely-used for MBD, and is dedicated to the solving of specifically *dynamic* problems (whereas FE is traditionally associated with static or quasi-static solutions). The software allows CAD geometry for the implant design under test to be imported directly and used within the model, and allows co-simulation with MATLAB/Simulink for control-plant modelling. Probabilistic studies can be managed internally by the ADAMS/'Insight' module (for design-of-experiment (DOE) studies), or externally by any hand-coded or 3rd party statistical software, e.g. NESSUS (South-West Research Institute).

4.1.3. Target Simulator Configuration

For any model of TKR mechanics, there are many potential factors to include, and the complexity of the model can very quickly expand drastically. This brings with it associated risks (unexpected behaviours, errors and simulation failures). To mitigate these risks, the earliest developmental models should *not* be fully-featured stochastic studies. Instead, the various sources of variability should be introduced sequentially in phases (i.e. a 'crawl-walk-run' approach). Therefore, the first objective is to demonstrate that a 'baseline' *deterministic* model (similar to those used in existing FE models) can equivalently be implemented using the alternative MBD software environment. This has three purposes: to gain experience with the software, to provide a platform for further modelling, and to corroborate results with existing computational & experimental data. The target for early corroboration efforts was the study by Halloran et al [94, 156], who used FE methods to simulate the SKS and compared experimental and computational results. In his thesis, Halloran performed extensive FE-based modelling [157]; he also explored rigid (non-deformable) FE, using the 'elastic foundation' bed-of-springs approach to model contact, based on interpenetration of the geometries and a pressure-overclosure relationship to determine the resulting contact forces. These rigid-FE linear-elastic-foundation models are particularly suitable targets for a 'silver standard' *in-silico* versus *in-silico* corroboration, since their use of rigid body models makes them comparable to MBD-based methods. Therefore, the following section describes a specific, targeted corroboration against Halloran's baseline SKS model.

4.2. Initial (Deterministic) Corroboration

4.2.1. Implant Geometry

This deterministic study tested a standard, widely-used CR TKR design, with a Co-Cr femoral component on a FB tibial UHMWPE insert. For this TKR model, two alternative designs for the insert are available – a '*semi-constrained*' (S/C) design with more conformal articular geometry in the sagittal-plane, and a less conformal '*unconstrained*' (U/C) design (see Figure 47). Both alternative tibial inserts use the same femoral component. The 'parasolid' format CAD models of two designs were acquired from the manufacturer for use in these studies. Unless otherwise stated, results are always presented for a *right* knee.

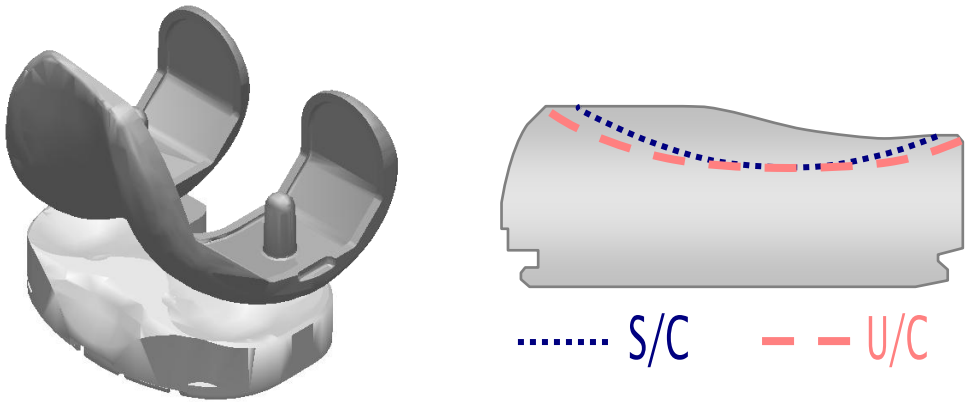

Figure 47: CAD geometry for femoral & tibial components (left), & sagittal-plane comparison between S/C & U/C tibial inserts (right).

4.2.2. Driving Inputs

The experiment was designed to simulate a normal gait cycle, loosely based on ISO-14243-1 force-driven gait [23]. As discussed in Chapter Two, the SKS is a hybrid force/displacement-driven simulator, so the four input waveforms required are A-P force, I-E torque, axial compressive force & flexion angle. In this case, the inputs were not identical to the ISO standard, so the 'feedback' data retrieved from the experimental rig is used instead (see Figure 48). Note that the illustrated gait cycle begins at heel strike; the stance phase is then the first ~60% of the cycle, followed by toe-off, then the remaining ~40% represents the swing phase through to the next heel strike event. The cycle is intended to be driven close to real-time speeds of around 1Hz.

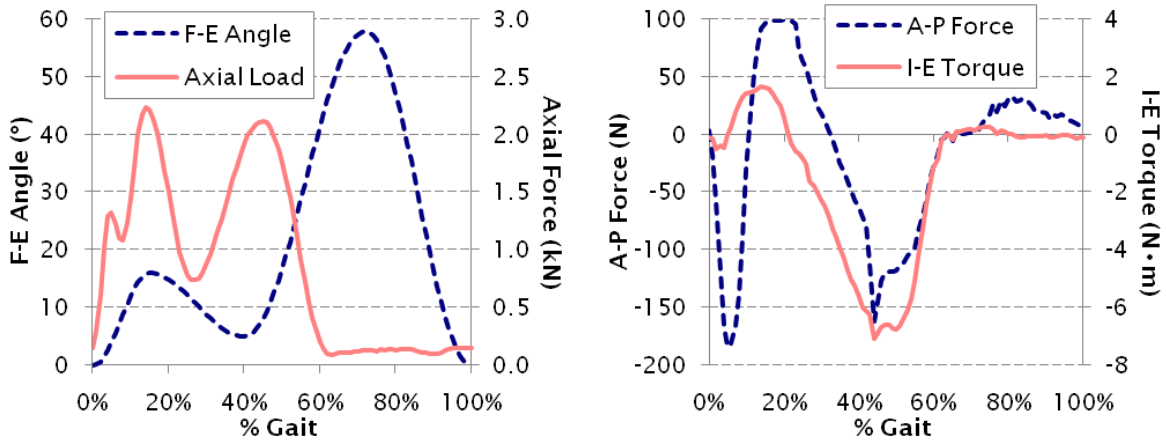

Figure 48: Input waveforms for force-driven SKS (adapted from Halloran et al [94]).

4.2.3. Spring Restraint Model

The SKS uses a transverse-plane spring restraint model. This is not directly anatomically representative of ligaments (since this is not the aim of the original wear simulator); however it does appropriately restrain the implant kinematics when used in conjunction with the input waveforms depicted above (i.e. it

provides an *aggregate* restraint at the knee, broadly equivalent to the effect of natural soft tissues in the transverse plane). The system consists of four spring buffers in the A-P orientation, with a fixed M-L separation providing the moment arm for I-E rotational torque restraint (see Figure 49). Various values of different spring stiffness have been proposed [158], sometimes featuring a short 'dead zone' permitting a few mm of unrestrained motion. For this model, Halloran et al adopted a spring configuration used by DesJardins [159]; this configuration is accordingly reproduced here. Each spring was treated as linear, (with the dead zone neglected) and with a stiffness of 5.21N/mm, with the spring M-L separation, $\Delta_{M\text{-}L}$, set at ±28.7mm. (It may be verified that the corresponding total transverse stiffness is ±20.84N/mm, with angular torsional stiffness of ± 0.30 N·m/°).

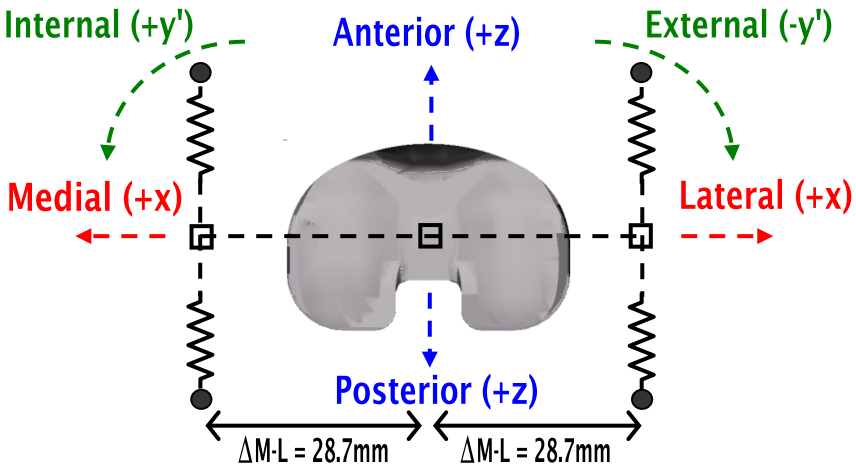

Figure 49: Spring restraint in the transverse plane.

4.2.4. Mechanical Configuration

The natural tibiofemoral joint has no fixed axes, but in practice most simulators constrain the femoral and tibial components, typically reducing the system to 6 (out of a possible 12) degrees of freedom (DOF), since it is generally only the *relative* orientations of the femoral and tibial components with respect to one another that are of interest. The SKS model follows this convention, as shown in Figure 50 (compare to Figure 23 in Chapter Two). Note that because each DOF is applied sequentially, the order of application does affect the kinematics (unlike e.g. Grood & Suntay co-ordinates [17]). However, for small angles (< ~10°) this has negligible influence; hence the F-E rotation, which is not limited to small angles, must be applied last. The configuration is summarised in Table 6.

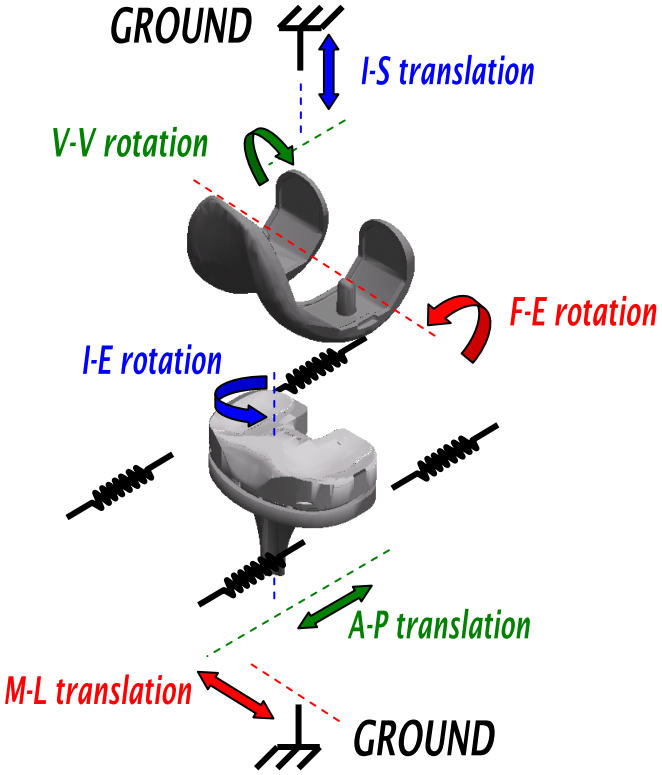

Figure 50: Mechanical DOF Arrangement.

Femoral Component	Tibial Component
I-S translation	M-L translation
V-V rotation	A-P translation
F-E rotation (driven as SKS input)	I-E rotation

Table 6: SKS configuration for the six degrees of freedom.

Dynamic terms had to be assigned for the various properties of the resulting bodies. The target study by Halloran et al was based on limited experimental data, and so the inertia, friction and damping for the model were all estimated based on engineering judgement. For this present corroboration study, the values used are based on the target study, in order to match the computational model; as such these values do not accurately represent the *in-vitro* SKS rig. Representative inertia was assigned to the femoral component (5kg). A higher inertia was assigned to the tibial component (25kg). The tibiofemoral friction coefficient was constant at 0.04, and the transverse-plane damping was set to 1% of the spring stiffnesses (50N.s/m per spring).

4.2.5. Component Positioning

Technically, the ISO-standard defines specific guidelines for the position of a fixed femoral axis of rotation [23, 24]. However, many studies do not precisely observe this standard, and as the aim of this exercise is corroboration rather than rigorous adherence to standards, the proprietary axis positioning within the original experimental studies will be adhered to when differences exist.

4.2.6. Contact Algorithm

An important metric for knee performance is contact pressure (CP). However, whereas *kinematics* (i.e. translations & rotations) can be readily determined from MBD analysis, the *kinetics* (forces, torques) are only reported as resultant values for the whole rigid body. Therefore in order to model the force *distribution* (and hence CP), a macro was written to discretise the surface of the tibial insert into multiple elements (see Figure 51). The resultant contact force for each of these elements (and the element area) could then be used to estimate CP. (Essentially this is comparable to the elastic foundation approach [94, 160]). For this study, it was determined that a resolution of 1mm² would be adequate for CP visualisation (based on sensitivity studies and comparisons to previous FE models [91, 94]).

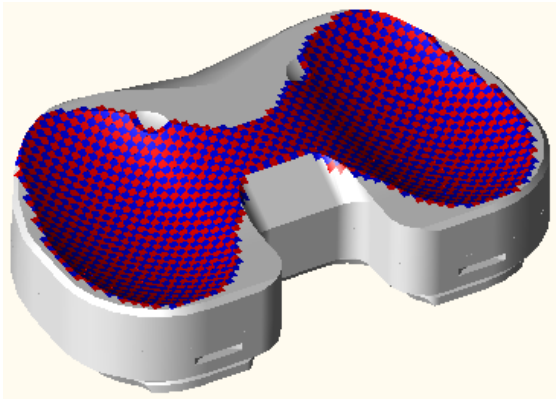

Figure 51: Tibial insert, showing 1mm² grid.

ADAMS features an internal 'impact' function [161]; this allows the two rigid bodies to partially 'interpenetrate', and then approximates deformable contact by relating contact normal force (F_N) to the interpenetration depth (g), using an exponential relationship, where k is the stiffness co-efficient; e is the exponent:

$$F_N = k \times g^e \qquad (2)$$

With suitable values for k and e, this can be used to fit an EF-style contact model. For a simple linear model, with material thickness ~10mm and cell area 1mm², it may be shown that $e = 1$ and $k \approx 10^5$ gives a reasonable first approximation, which can be further 'tuned' to experimental data (for more, see Appendix C).

4.2.7. Output Measures

The output measures reported here are A-P translation & I-E rotation (reported relative to the 'settled' reference positions of the components, as defined above), contact area and peak CP (the highest surface pressure recorded for any of the mesh elements on the tibial insert at each point in time).

4.2.8. Results

In Figure 52, results for A-P Translation, I-E Rotation, peak CP and contact area are shown. In the graphs, the rigid-body result from MBD (ADAMS) is compared with FE and experimental data for the same implant design presented by Halloran et al. Maximum A-P range was ~4mm, with I-E rotation varying by ~7˚, and CP closely following the axial force waveform, with a maximum value of ~ 17MPa.

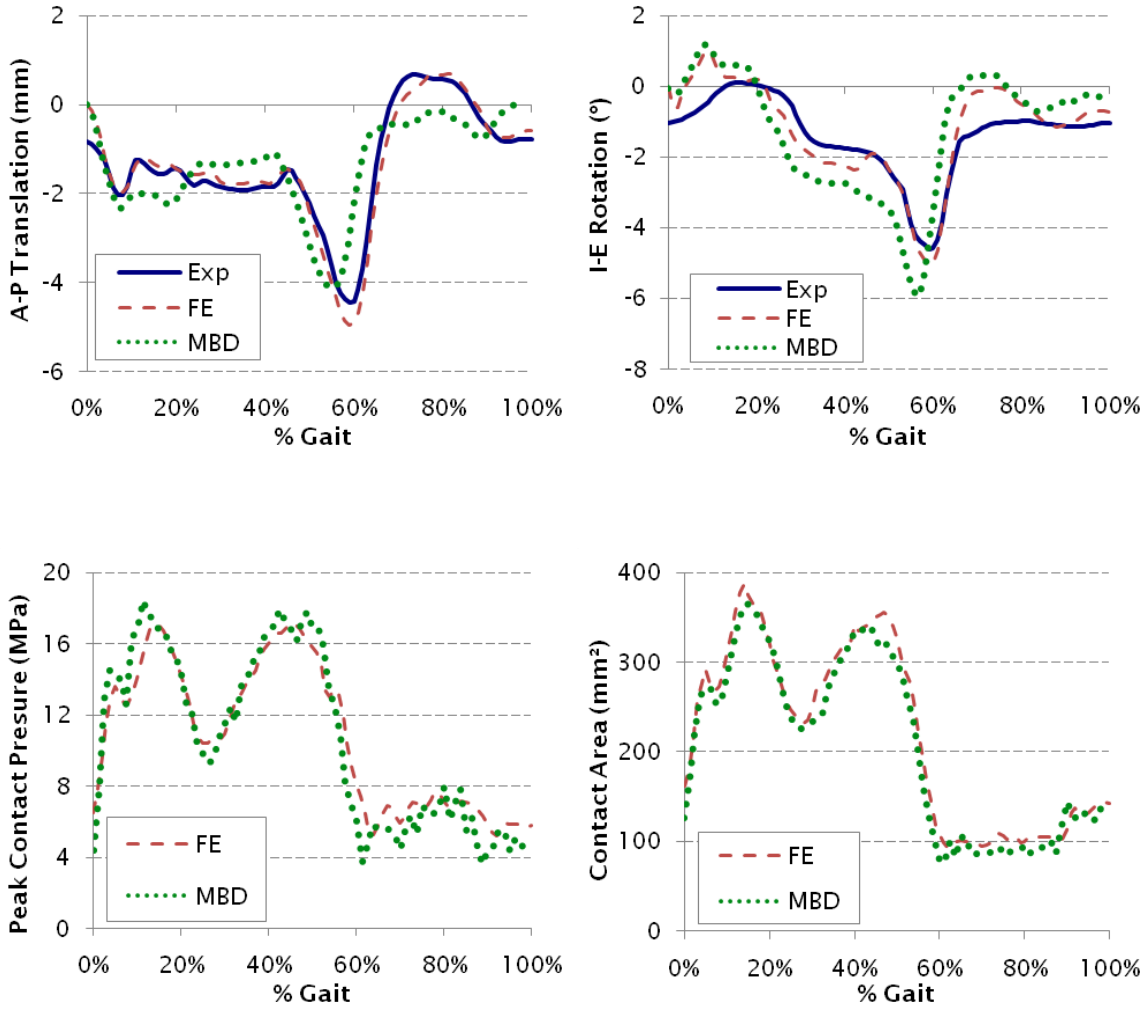

Figure 52: Tibial A-P translation (top left), I-E rotation (top right), maximum CP, (bottom left), & contact area (bottom right). MBD - ADAMS (dotted trace) vs. FE - Halloran [94] (dashed trace), with experimental kinematics (solid trace)

Figure 53 shows an example contour plot for CP, at a single point in the gait cycle, comparing both the FE and MBD methods – note the two are very similar. This MBD-based model solved in under 10 minutes (Intel P4 3GHz, 2Gb RAM); a favourable computational cost compared to more complex deformable FE models.

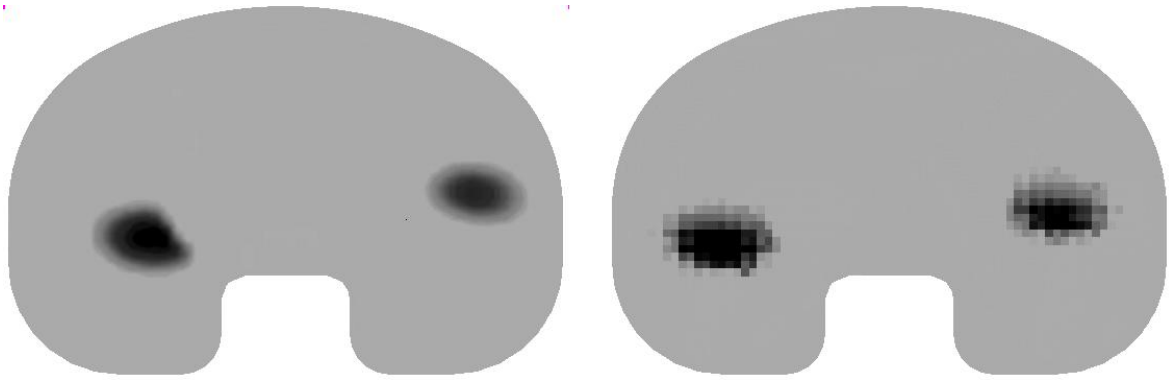

Figure 53: Contour maps for CP Distribution. Left: FE (Halloran [94]). Right: MBD (ADAMS). (Taken at 10% gait; scale maximum 10MPa).

4.2.9. Discussion

The ADAMS model results corresponded well with both the magnitudes and trends reported in the corroborated publication study, although there are some small differences due to various differences in the modelling approach taken. The 'contact' and 'friction' models are slightly different in the present study, and the system of discretising the tibial component is proprietary, introducing small differences. Also, although FE is fully capable of dynamic analysis, historically biomechanical FE studies have often been 'quasi-static', meaning each time instant is evaluated in isolation. This means that inertial effects (the mass of the components) or dynamic effects (damping in the polymer contact or spring restraint) were not considered. ADAMS is fully dynamic so does incorporate such influences, introducing further differences.

An important observation is that whilst the FE models *appear* to 'validate' well with the experimental data, by applying quasi-static conditions (and hence neglecting dynamic effects), they are failing to capture the full mechanics of the *in-vitro* test. The SKS configuration is quite forgiving, owing to the 'hard' restraint provided by the physical springs; nonetheless dynamic terms can play a very important role in dictating the kinetics and kinematics. Unfortunately, dynamic properties are not widely reported as historically they have not been widely modelled. Subsequently, these properties *must* be better understood for more accurate fully-dynamic modelling.

Nevertheless, the results are sufficiently similar to published studies to give further confidence in the choice of modelling environment, and the model itself. The solve time is comparable to published rigid body studies (e.g. the rigid EF model variant reported by Halloran et al [94]) and sufficiently low to allow multiple-trials in a larger-scale stochastic study. With a well-corroborated deterministic model as a baseline, the analysis capabilities of this model could now be extended, by incorporating *in-silico* wear prediction.

4.3. Implementing *In-silico* Wear Prediction

The ability to predict wear *in-silico* is valuable, as wear and wear-related osteolysis & loosening are leading causes of TKR failure (see Chapter One). However, the precise physical mechanisms of wear are not fully quantitatively understood, and 'wear' is a catch-all term which includes a number of distinct tribological 'mechanisms'; for example:

- Surface fatigue wear – caused by contact between two bodies, where there is limited or no sliding motion (e.g. rolling contact). This mechanism generally produces minimal wear damage.
- Adhesive – caused by two bodies sliding together with a compressive load applied between them; the degree of wear varies based on the motions and compressive loads.
- Abrasive – caused by hard particles mechanically abrading against a softer material. The particles may be embedded with a composite material; e.g. barium sulphate particles in bone cement.
- Three-body wear – a form of abrasive wear where loose ('third body') particles become located between the two contacting surfaces, causing accelerated wear.
- Corrosive – caused by the degradation of the material properties, rather than purely mechanical effects. Corrosive wear can however be exacerbated by the mechanical environment.

Any or all of these processes may be occurring in a given TKR *in-vivo*; it is therefore important to differentiate and understand what is being modelled by the predictive wear algorithm. For a well-implanted modern TKR, the mechanism believed to be most important is adhesive/abrasive wear; this is the focus of the models described henceforth.

4.3.1. Adhesive / Abrasive Wear Theory

Fundamentally, the same concept underlies all forms of mechanical wear. Energy is transferred from the *kinetic* energy of the moving surfaces, to the increased *surface* energy of the wear particles generated. This energy transfer is facilitated via surface friction during motion; therefore the amount of energy lost to friction dictates the *maximum* energy available for tribological processes (this theoretical relationship has been described in detail by Wang [162], who illustrates a strong relationship between wear rate and friction coefficient). Despite attempts at a theoretical basis such as [162], wear prediction remains largely an empirical

science. The foundational work in quantifying wear predictions was the relationship reported experimentally by Archard in the early 1950s, in relation to wear of bearings in electrical machines [163]. The relationship has been widely used in the tribological study of prosthetics, as the Archard/Lancaster equation:

$$W = k.CP.s \qquad (3)$$

Where W is the experimentally-measured wear depth, k is a scaling constant, CP is the contact pressure, and s is the sliding distance (the product of contact pressure & sliding distance is sometimes termed 'tribological intensity'). Note this is a *localised* expression for the *localised linear wear depth*; to estimate a wear volume, this must be integrated across the contact area, taking account of variations in contact pressure and cumulative sliding distance (in other words, calculating wear computationally requires piecewise numerical integration).

Beginning in the mid-1990s this wear formulation has been applied to computational biomechanics, notably in the FE-based work of Maxian et al [85, 164-166]. Essentially, the simplified form of the equation above is applied to each finite element in contact, at each time step. Summing together these discrete contributions gives the total wear *volume*. This can be converted to an equivalent wear *mass* based on the density of UHMWPE (~0.93 mg/mm³ [167]).

Although the Archard equation has *some* theoretical basis, it is *not* an analytic formulation. In practice, it does not account for variations in material properties across the surface (inhomogenities, varying surface roughness, varying degrees of cross-linking on the molecular level), or variations in the size & surface energy of wear particles. Rather, it provides an empirical approximation, designed to provide an aggregate estimate of wear as a macroscopic-level phenomenon.

4.3.2. Modelling the 'Cross-Shear' Effect

The Archard/Lancaster wear formulation assumes that the wear constant 'k' is a fixed value. This means that, regardless of variations in contact pressure, sliding distance or any other parameter, the same proportion of frictional energy is assumed to be producing the same volume of wear debris. However, it has been demonstrated that the wear 'constant' varies considerably. One particular influence appears to be the orientation of the sliding contact between the two bodies, relative to the alignment of the polymer fibres on the contact surface. However, the surface alignment in turn depends upon the time-history of the previous sliding motions. As a result, the overall wear rate can increase if the orientation of sliding motions *at any fixed point on the surface* changes considerably over time through the activity cycle. Generally, relatively 'linear'

motions (unidirectional or bi-directional sliding) result in much lower wear, whereas a path featuring considerable lateral shearing motion (i.e. sliding motions tangential to the principle sliding direction) results in higher wear than the Archard/Lancaster prediction would suggest. This has resulted in alternative formulations for wear based on the crossing-motion, or cross-shear (CS) theory, where essentially the wear constant '*k*' becomes variable, as a function of the CS:

$$W = k(CS).CP.s \qquad (4)$$

The effect of different crossing-motion paths has been demonstrated empirically by Turell et al [105]. Causally, it is postulated that the CS increases wear by causing fibrillar de-bonding. For linear sliding paths, the UHMWPE fibrils are found to be aligned with the sliding direction; for high-CS sliding paths, the fibril alignment is more multi-directional [168]. There are several proposed metrics for defining the 'degree' of cross-shear. The most rudimentary is the M-L/A-P ratio, assuming the principle sliding direction to be along the A-P axis, (Figure 54, left condyle). However, this formulation produces singularities if the motion is entirely in the M-L direction. Therefore, an alternative is to use a 'bounded' form (Where the denominator is not 'A-P' but the sum of 'M-L+A-P'). A more sophisticated approach is to determine a specific principle direction vector based on the actual path data. This *principle sliding direction* is designated 'B', with the transverse (i.e. 'cross-shear') direction designated 'A' (see Figure 54, right condyle).

Whereas the M-L/A-P model breaks down if the principal sliding direction becomes predominantly lateral, the A/B model can account for any sliding direction, so is more robust. For the A/B formulation, both the unbounded & bounded (A/A+B) formulations are again possible (see Table 7).

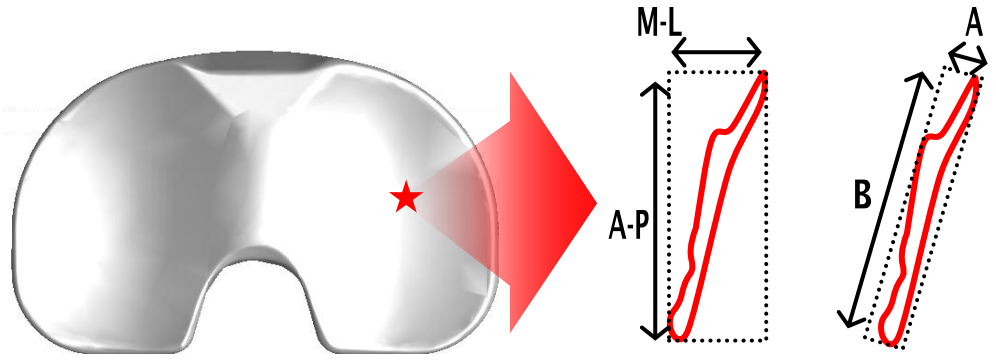

Figure 54: Defining cross-shear for M-L/A-P ratio (left) and A/B ratio (right). These simple models use a 'cycle-averaged' measure of the crossing motions seen at any point on the insert surface. For 'skewed' path orientations, M-L/A-P tends to over-predict CS (as shown by the lower-aspect ratio of the enclosing rectangle).

Wear model	Unbounded	Bounded
M-L/A-P	$k(CS) = k_0 . \frac{\sum \lvert ML \rvert}{\sum \lvert AP \rvert}$	$k(CS) = k_0 . \frac{\sum \lvert ML \rvert}{\sum \lvert ML \rvert + \sum \lvert AP \rvert}$
A/B	$k(CS) = k_0 . \frac{\sum \lvert A \rvert}{\sum \lvert B \rvert}$	$k(CS) = k_0 . \frac{\sum \lvert A \rvert}{\sum \lvert A \rvert + \sum \lvert B \rvert}$

Table 7: Four alternative formulations for cycle-averaged cross-shear.

These proposed CS models are purely empirical, and have no physically-based analytical derivation; they are based on limited observations from *in-vitro* experimentation. Unfortunately limited data exists for such *in-vitro* experiments, owing to the costs and timescales involved. Further reported studies often vary different factors (e.g. implant design, material) from test to test, which means the results are not directly comparable. To compound these challenges, wear tests exhibit a high degree of variability even when repeating the same test on multiple stations (e.g. Fisher et al reported variations of >±30% [169]).

In consequence of this paucity of good data, there is no consensus on a definitive 'correct' model for wear available, and alternative formulations continue to be suggested. One example is the 'crossing intensity' formulation proposed by Fregly et al [124]. Here, cross-shear is estimated as the 'spread' of different sliding directions, weighted by both the sliding distance and contact pressure. This effectively gives a 'standard deviation' of the spread of sliding directions; normalising this by the 'worst case' sliding path (circular rotation) gives the crossing-intensity value, which can then be used as a CS term (see Figure 55). Yet another proposal by Willing et al [170] considers a 'closed-path' of sliding vectors, and expresses CS as the ratio between the perimeter length and enclosed area of this shape (see Figure 56).

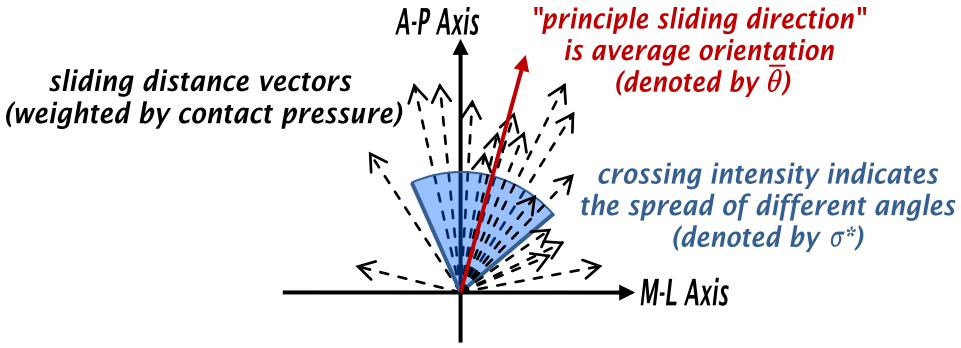

Figure 55: Alternative CS concepts: Hamilton's statistically-based "crossing intensity". Wear increases as the 'spread' of sliding vectors increases.

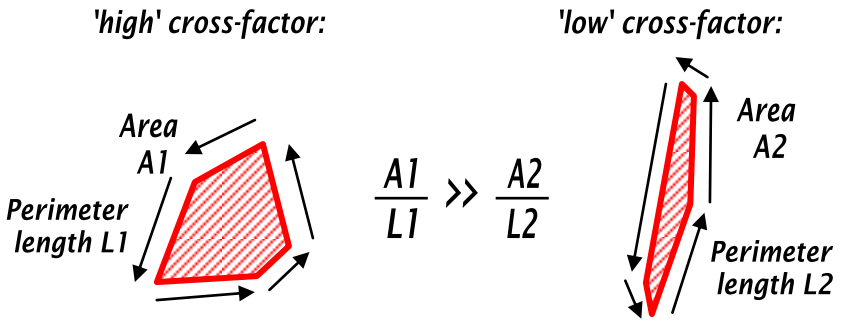

Figure 56: Alternative CS concepts: Willing's geometrically-based "cross-factor" method. Wear increases for paths where "Area : Perimeter" ratio is higher.

4.3.3. The Role of Contact Pressure in Wear

Historically, ever since Archard's original hypothesis it has been assumed that linear wear depth rate is directly proportional to contact pressure (as discussed above). More recently however, this assumption has been challenged in studies by a number of authors. Using POD testing, Mazzucco et al [107] and Ernsberger et al [106] have both argued that there is no apparent relationship between wear depth rate and contact pressure (i.e. it is independent); work by Kang et al goes further, arguing that there is an *inverse* relationship between wear and contact pressure (i.e. wear factor *decreases* as contact-pressure increases, at least within the range of CP tested [125, 171]).

However, these studies were all performed in the simpler domain of POD tests, where geometry is not a confounding factor, and contact pressure is (ideally) constant across the articulating surface. How applicable these conclusions are for more the complex geometries, kinetics and kinematics of TKR wear is a matter of ongoing debate. A major obstacle in comparing and testing these different proposals for wear algorithms is that there is often limited experimental data to base the formula on, and small numbers of trials (often in the limited domain of POD tests) cannot provide sufficient grounds to explore the differences between the various algorithms proposed. Therefore, the need exists to apply these algorithms across a wider range of experimental TKR tests to corroborate their performance on a larger scale.

Clearly, until the precise details of adhesive/abrasive wear are better quantified experimentally, debate will remain as to which mathematical model gives the most accurate results. In light of this, it was decided to include a range of *different* wear algorithms within the models in the present studies.

Note that there is a considerable workload of post-processing involved in evaluating wear – a typical activity might include several hundred time-frames; in any frame, as many as a thousand elements may be in contact – this results in a

very large number of individual pressure & sliding calculations to evaluate. The computational cost of post-processing may be an important factor when choosing a wear formulation to use with a large-volume DOE or probabilistic study. The A/B formulation and the crossing-intensity formulation are both *recursive*; it is necessary to scan through the data once to determine the 'principal' sliding direction, and then again to apply the actual wear summation. (As such, for high-speed stochastic studies it may sometimes be preferable to use the less accurate but faster M-L/A-P formulation).

For the exploratory studies here and in subsequent chapters, a number of the algorithms described above are incorporated and used in parallel, in order to compare their performance in different situations.

4.3.4. Adaptive Wear Modelling

The process of wear is inherently *dynamically adaptive*; localised high wear can result in faster deformation in certain locations, thus altering the surfaces of the articulating geometries, and thus altering the kinematics and contact pressure distribution for subsequent cycles. This is most clearly seen in the 'bedding in' phenomenon, where concentrations of high pressure will tend to result in localised high wear, such that the surfaces become more conforming and the contact is more evenly distributed, thus reducing pressure concentrations.

It is possible to simulate this adaptation, by re-modelling the contacting surface between consecutive simulations, as first demonstrated by Maxian et al [85]. In reality, the surface adaptation is a continuous process. Numerically, however, this would be very impractical; the geometry would need to be minutely modified during every individual cycle. For a 5- or 10- million cycle test requiring ~5-10 minutes of computation time per cycle, this would result in simulation times of thousands of years (i.e. clearly not feasible).

Instead, the surface adaptation is applied in discrete blocks of cycles; it is assumed that for 'short' periods the wear rate and wear depth are approximately linear (convergence tests suggest a maximum step size of 500,000 cycles to 1Mcycle [85, 92], corresponding to a few months of *in-vivo* use). An entire long-term simulation of several million cycles may then be completed in just a handful of iterations. Generally, a sensitivity test will be used, testing decreasing step sizes to verify that the overall wear rate converges towards the rate observed under 'continuous' conditions. The flow diagram for this algorithm is illustrated in Figure 57.

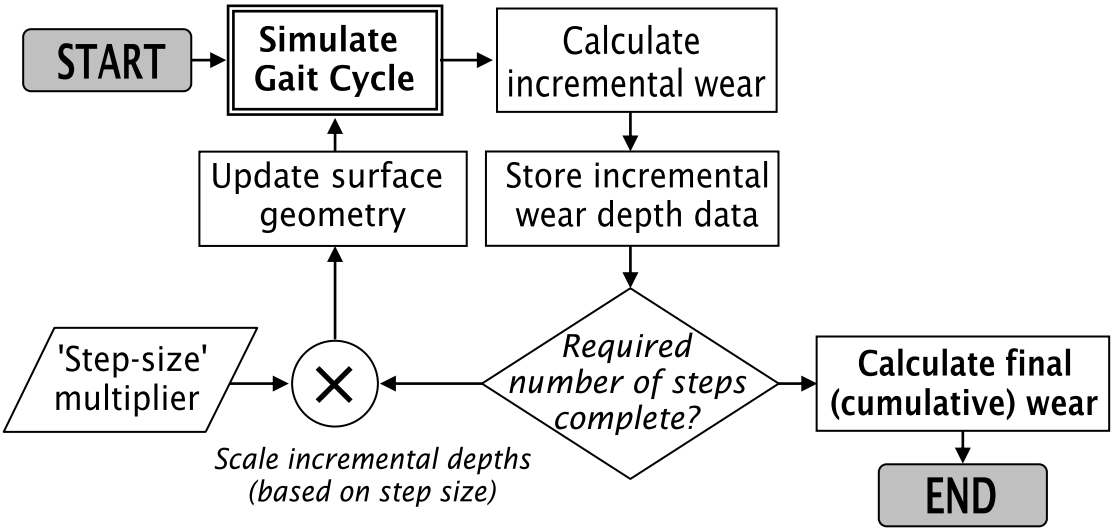

Figure 57: Computational adaptive wear prediction: process flowchart.

4.3.5. Implementing the Wear Algorithms with MBD

In order to calculate wear within ADAMS, the simulation results must be post-processed. This could potentially be done using a number of software tools; however using an external third party application would require data to be imported and exported between programs. Instead, the command-scripting capabilities of ADAMS were used, to perform the post-processing *internally* within the ADAMS environment. Computationally, this is considerably slower, since the macro scripting language is not compiled; however, the convenience of keeping all processing within ADAMS is considered to outweigh this disadvantage.

The macros operate by interrogating the results database (containing kinematic & kinetic information from the previous analysis). The 'output' is a series of numerical arrays (indexed to reference the discretised elements of the tibial surface) containing information such as total sliding distance, contact pressure, cross-shear, or wear depth, for each cell. The ADAMS GUI has been adapted to display this data in the form of colour plots (similar to the contact pressure plots illustrated in the earlier deterministic corroboration, see Figure 53). For example, see Figure 58, showing colourised 'contour maps' for linear wear depth. This facilitates both numerical & graphical visualisation of the predictive wear results.

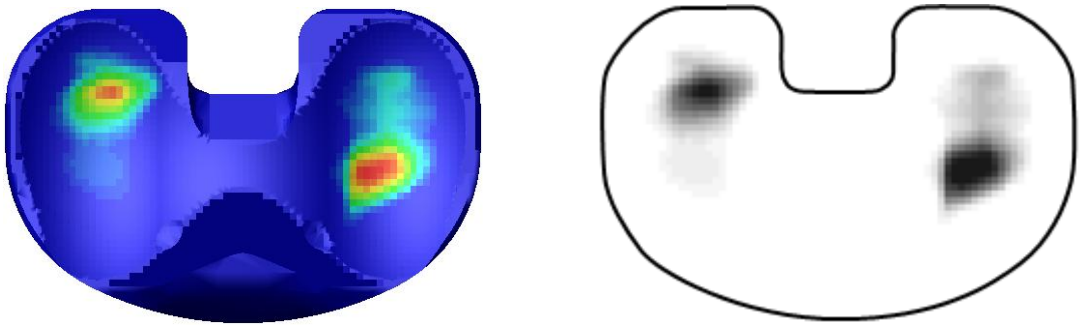

Figure 58: Linear wear depth contour plots in hue (left) & gray-scale (right).

The capability to model adaptive wear has also been incorporated. In FE models, this is achieved by altering the entire surface mesh; the surface must remain congruent; therefore the individual *nodes* are displaced based on the linear wear depth, forming the new 'smooth' surface mesh. In ADAMS, the surface may be allowed to become incongruent, since the contact algorithm is based purely on the interpenetration depth for the entire surface element; therefore individual cells may be entirely displaced vertically (see Figure 59). This is less physically representative, but an acceptable compromise within the MBD environment.

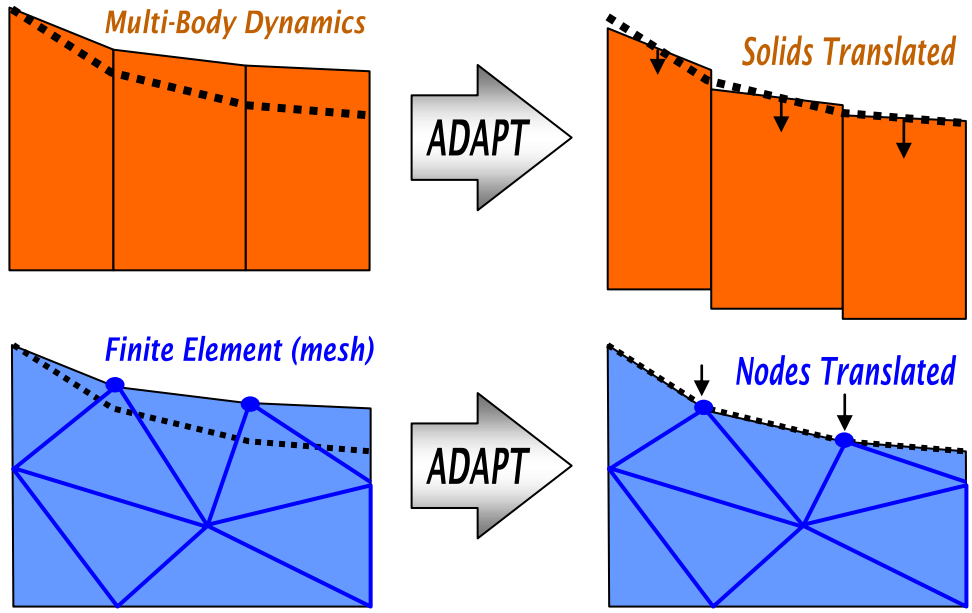

Figure 59: Comparing surface adaptation methodology for MBD (above) & FE (below).

4.4. Corroborating *In-silico* Wear Models

It is recognised that *in-silico* computational wear predictions are not analytically-based, and so cannot ever provide 'exact' predictions of *in-vitro* wear. In light of this, it is important that the results from MBD-based wear modelling do not introduce any additional error or variation – they should corroborate closely with other *in-silico* FE-based methods, as well as *in-vitro* studies. With the capabilities of the wear model developed, the model was corroborated with published results using established FE-based models & wear simulator rigs.

A first step is to corroborate a single 'deterministic' case in detail. The long-term adaptive wear study by Knight et al [92] was selected, because it includes both FE and experimental results, and is information-rich, presenting data for adaptive wear steps, showing wear depths, volumes, and contour plots. An MBD-based model was constructed to replicate the test conditions (using proprietary force-driven inputs with soft/hard springs in a SKS configuration, and a CR FB commercial knee implant). Adaptive wear was simulated for 10 'steps' of 500,000

cycles each, out to 5MCycles. The wear model used was the simple 'Archard' formulation (no cross-shear), with a wear constant, k, of 2.64×10^{-7} $mm^3/N{\cdot}m$.

In Figure 60, the MBD-based model is compared to the results reported by Knight et al. Wear depth and volume are very similar; there are differences in the precise wear contours, but the trends are similar (greater & more concentrated wear on the medial side, with a more dispersed wear scar shifted posteriorly on the lateral condyle). Note that the wear post-processing adds an additional overhead to computation times; hence whereas the baseline mechanical model solves in 5-10 minutes, simulations with a full wear analysis (for multiple wear methods) can take 15-20 minutes (computing the wear rates for all the different alternative wear algorithms incorporated into the ADAMS model). There is therefore a solve-time trade-off which must be considered when incorporating wear algorithms.

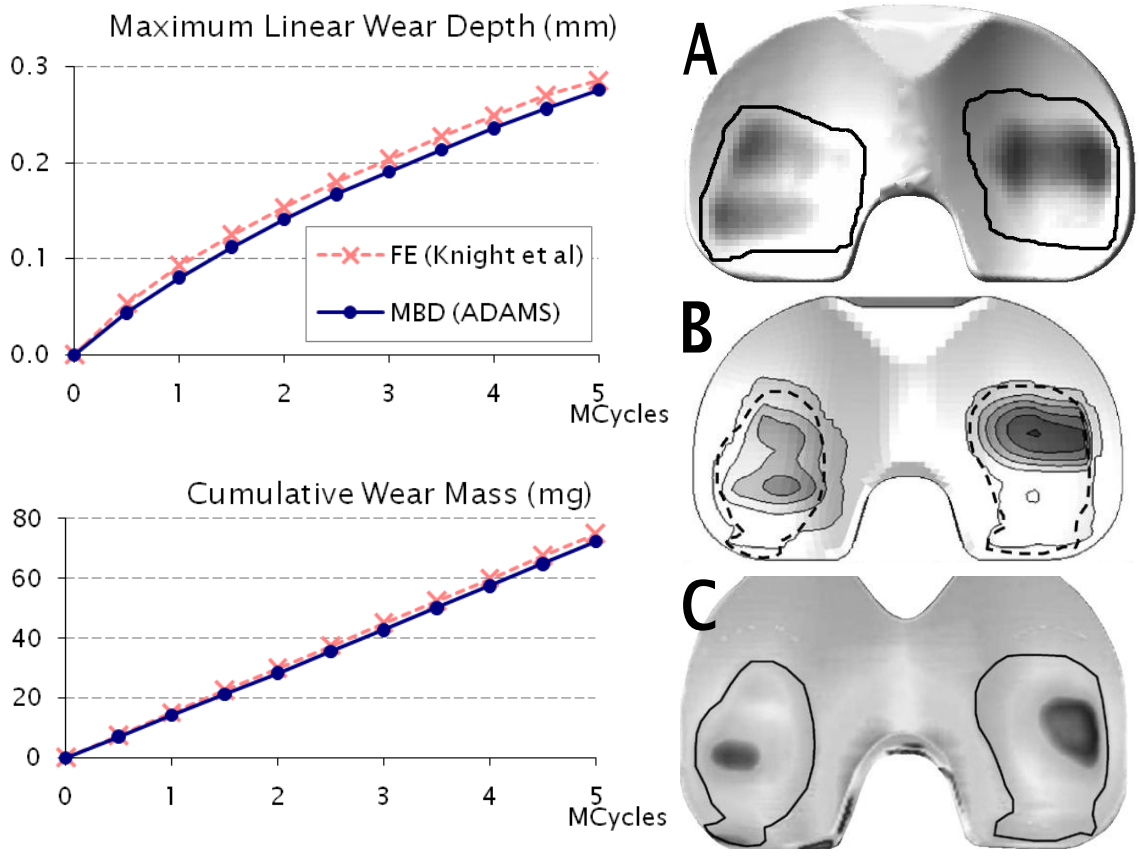

Figure 60: Comparison of FE & MBD adaptive wear. Left: cumulative wear depth (above) & mass (below). Right: wear depth contours for (A) ADAMS MBD, (B) FE – Knight et al [92], and (C) experimental.

Alongside this comparison test, a number of other corroboration studies were performed, comparing the ADAMS model to existing FE-based wear predictions and experimental wear data. For the sake of brevity, these are not reported in detail in this thesis, although some of the results from these studies were utilised as part of the work described in the following section.

4.5. Investigating *In-silico* Wear Theories [4]

4.5.1. Methods

The MBD-based model has now been shown to provide a fast and accurate alternative modelling approach to FE, matching deterministic results for kinematics, kinetics, and also for existing wear prediction methods. However, the power of these wear prediction methods has not been effectively demonstrated; to-date, *in-silico* wear models have been 'tuned' to and compared with only small experimental datasets, either using published pin-on-disc (POD) data, e.g. in [105, 124], or else TKR wear simulator results, e.g. [92, 127]. Whilst these studies demonstrate the value of *in-silico* methods in individual cases, they cannot broadly corroborate across a range of test conditions.

The initial wear predictions used with this model are based on standard algorithms discussed above; the baseline Archard/Lancaster sliding-distance model [163] (without CS), and other algorithms including CS (e.g. M-L/A-P [126], A/A+B [105], and σ^* 'crossing intensity' [124]). Alongside these existing formulations, alternative arrangements have been included to explore the effect of excluding CP from the wear model [106, 107, 125].

Twenty-two different experimental tests were selected, sourced from the public literature and proprietary test data, where 'conventional' polyethylene was tested (with minimal or no cross-linking), to ensure that the tests would be broadly comparable. Implant geometry was acquired from manufacturers or reverse-engineered. Results for a range of kinematics under displacement-control for the PFC sigma (fixed and mobile bearing designs) and LCS were sourced from [172, 173]. These implants were also tested under ISO 14243-1 force-control [174]. Results for the NexGen CR implant were corroborated under force-control [92, 175] and displacement control [43]. Additional implants included were the Vanguard PS under ISO force-control [176], and Triathlon CR under displacement control [177]. Proprietary unpublished test data was also used to corroborate semi-constrained & unconstrained design variants of the PFC sigma under displacement-controlled conditions. Finally, tests of femoral components against 'flat' polyethylene surfaces using displacement control [178] were included to corroborate the wear algorithms across a wider range of contact pressures & areas *in-vitro*. The full list of test-cases is summarised in Table 8. Note that because of the number of tests, it is not possible to include the full set of test-

[4] This section is adapted from the journal article: "In-silico Wear Prediction for Knee Replacements - Methodology and Corroboration". Strickland et al, J.Biomech (*In Press*).

conditions in this document for every case. In each model, the same procedure was followed; component positioning, allowed motions, spring constraint (where applicable), input loading profiles and any other relevant factors were matched to the reported test conditions in the literature. Where these conditions were not stated, and where the original investigators could not be successfully contacted for further clarification, 'generic' test conditions were imposed (e.g. assuming a 60-40 M-L load split [23], using a representative friction co-efficient of 0.04 [91], and adjusting the model configuration according to a typical set-up for the test machine being used; i.e. replicating the standard mechanical configurations for Instron, ProSIM, or AMTI simulator rigs, as available from the manufacturers). The original papers may be referred to for more details on individual test cases.

Source(s)	***Implant(s) (PE derivative)***	***Inputs (forces & kinematics)***
McEwen et al [173]	*Sigma FB & RP; LCS (GUR1020 & 1050)*	*Displacement (various kinematics) & ISO 14243-1 (Force) Gait*
Galvin et al [178]	*Sigma femoral on flat PE (GUR1020)*	*Displacement-driven Gait (various levels of kinematics)*
Knight et al [92]	*NexGen CR (GUR1050)*	*ISO-derivative Gait*
Cottrell et al [175]	*NexGen CR (GUR1050)*	*ISO 14243-1 (Force) Gait*
Muratoglu et al [43]	*NexGen CR (GUR1050)*	*ISO-derivative Gait*
Williams et al [177]	*Triathlon (GUR1020)*	*ISO-derivative Gait*
Haider et al [174]	*Sigma FB & RP (GUR1020)*	*ISO 14243-1 (Force) Gait*
Haider et al [176]	*Vanguard PS (GUR1050)*	*ISO 14243-1 (Force) Gait*
Proprietary unpublished data	*Sigma FB S/C & U/C (GUR1020)*	*Displacement-driven ISO-derivative & high-kinematics gait*
Proprietary unpublished data	*Sigma femoral on flat PE (GUR1020)*	*ISO-derivative; High & low levels of axial load & I-E rotation*

Table 8: Listing of test-cases used for corroboration, with references where applicable.

Wear rates reported in mg were converted to mm^3 using a density of $0.93mg/mm^3$. Although the model is capable of adaptive wear, to limit computational times for this exploratory study, volumetric wear rate for each case was calculated based on a single-cycle. Published experimental and computational long-term studies demonstrate that whilst linear wear depth rates may vary over time (e.g. [92]: Figure 7, [179]: Figure 3a), volumetric wear is reasonably linear within the first few million cycles, (e.g. [92]: Figure 6, [179]: Figure 3b, [42]: Figure 2, [44]: Figure 2, [180]: Figure 10). (Although the precise mechanics are not quantitatively understood, the increase in contact area due to 'bedding-in' seems to offset the gradual decrease in linear wear depth rate).

Once all the necessary experimental configuration data had been obtained for these tests (e.g. implant geometry, loading input waveforms, spring restraint setup and available degrees of freedom), the tests were simulated *in-silico* using the fast rigid-body model, and predicted wear was evaluated for each of the proposed wear formulations included in the model. The computationally-derived rates were then compared to the reported experimental wear rate (with error levels, where available). This allowed the predictive power of different wear algorithms to be compared directly.

4.5.2. Results

All of the test-cases were simulated successfully and were post-processed to evaluate predicted wear using the different algorithms. The volume of data generated is considerable, so wear contour maps are not compared here; only the baseline volumetric wear rate for each model using each algorithm is reported.

The following figures show correlation plots for a few of the selected models. Note that in every plot, there is considerable 'scatter', and the uncertainty (shown by error bars) in the experimental results is very large. The results very clearly confirm the current general consensus that the baseline Archard model has very limited predictive power to assess wear (Figure 61) – this is equally applicable for the knee as for the hip, despite the typically lower degree of cross-shear.

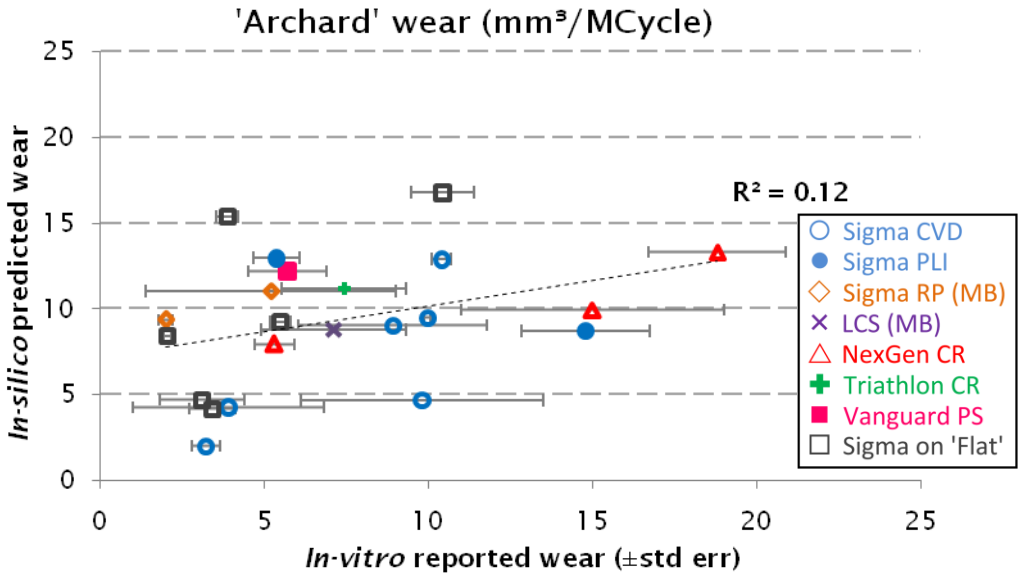

Figure 61: Experimental wear vs. wear predicted using the 'Archard' algorithm.

By comparison, *every* variation of wear algorithm which includes some representation of CS has a much greater predictive power (typically R^2 of 0.5 to 0.6; e.g. see A/A+B model in Figure 62). Considering these CS models, there are several important observations. First, the inclusion or exclusion of contact pressure as a proportional term within the algorithm does not consistently or

considerably alter the predictive power of the model for this particular set of test-cases. Second, the precise 'definition' (i.e. mathematical formulation) of CS used is of secondary importance compared to the decision to include or exclude a CS metric – the relative difference between alternative CS-based models is less than the difference between models with and without CS (compare Figure 62 and Figure 63).

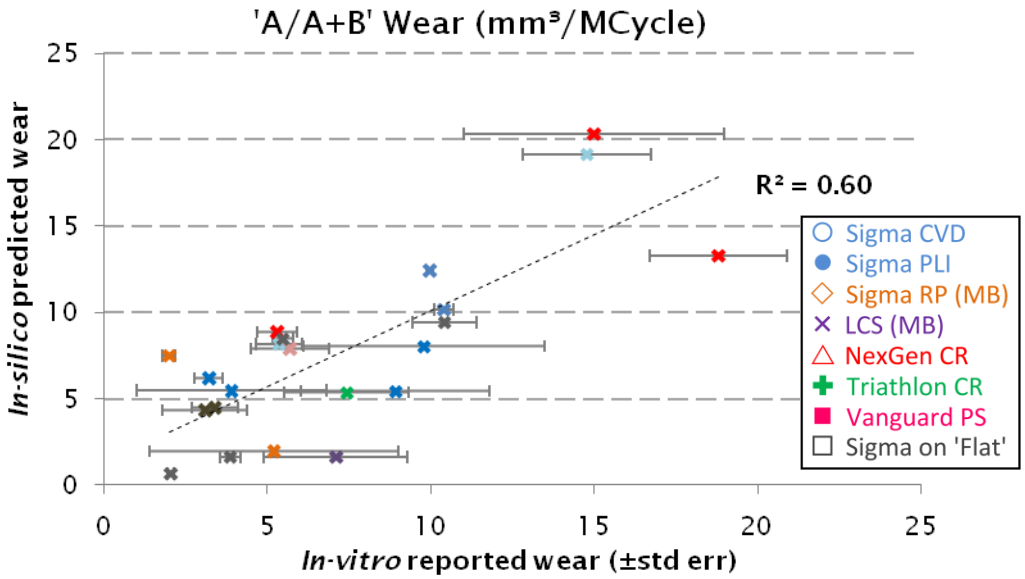

Figure 62: Experimental wear vs. wear predicted using 'A/A+B' algorithm.

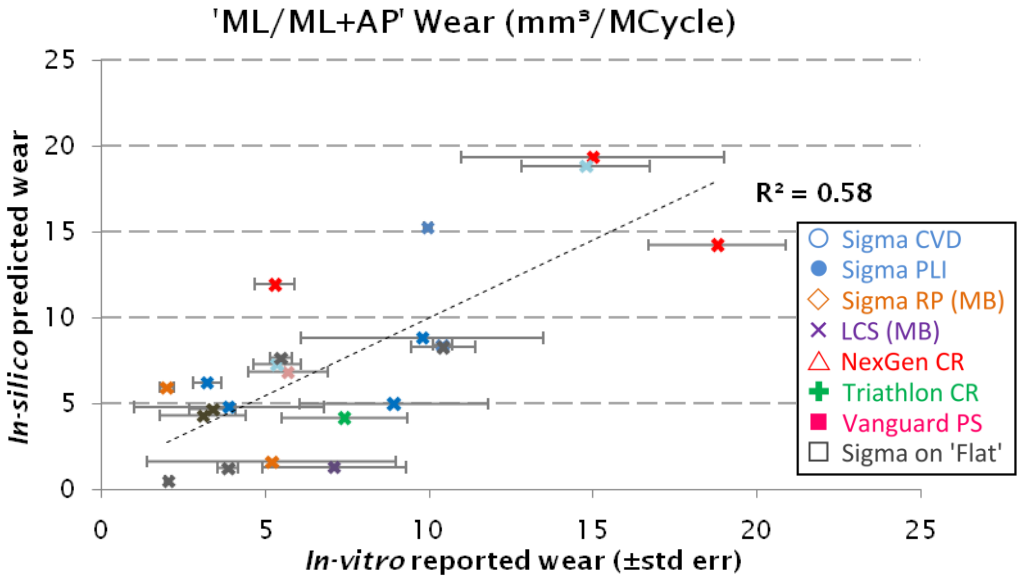

Figure 63: Experimental wear vs. wear predicted using 'M-L/M-L+A-P' algorithm.

Again, the treatment of CP within the algorithm also appears to be of secondary importance; both models with a proportional-CP term, and with no CP term, have similar predictive power for this set of test cases, provided that a CS metric is included (compare Figure 63 and Figure 64); the models including a proportional CP term appear slightly stronger, however the role of contact-pressure in wear

mechanics remains unclear – a plot of wear rate vs. cycle-averaged CP reveals no noteworthy correlations (see Figure 65).

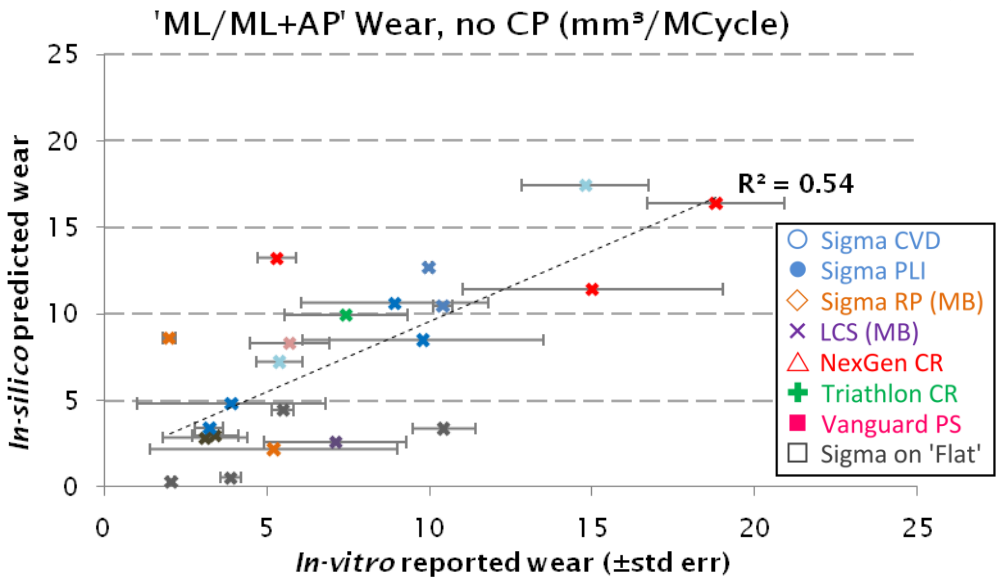

Figure 64: Experimental wear vs. wear predicted using 'M-L/M-L+A-P' (without CP).

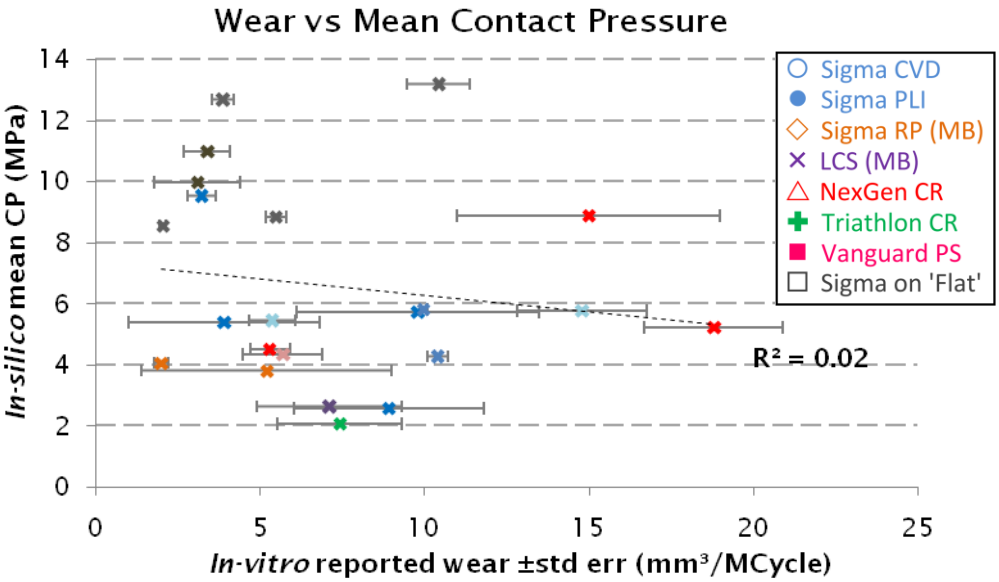

Figure 65: *In-vitro* wear vs. cycle-averaged contact pressure (no strong correlation).

Despite these uncertainties, it is possible to 'rank' the performance of the different CS algorithms for this particular test-case set. Based on this set of test-cases, the A/A+B wear model proposed by Turell [105] appears to be marginally the strongest predictor of *in-vitro* wear (Figure 62).

Previously, the reported empirical wear constants used in mathematical models of wear have been based on limited data-sets (e.g. a small sample of POD test results [105]). Based on this study, regression-fitting techniques were used to provide a set of wear constants for the different models tuned to this group of test-cases, for future use by other researchers to improve their TKR wear predictions. This has two advantages; the constants are directly based on TKR

tests, rather than derived from POD or THR tests (removing a potential confounding factor) and the values have been assigned based on this larger 'training' data set. The values suggested for the different models are listed in Table 9. Note that, although these values are more robust for general use than values derived from a smaller test set, they are still only approximations; using a larger data set, or including a wider range of activities, or considering different materials, could all result in different wear constants. Further, for any specific subset of tests (within a single research centre where test conditions are more repeatable & comparable), a better 'specific' constant may be selected; however this would have less applicability to test results from other research centres. Ultimately, with better experimental data, factors currently included under the 'constant' term may have to be recognised as distinct variables within the wear algorithm.

Wear Depth Formulation	Historical (Legacy) Constant, K_w	Revised Constant, K_w (based on test-cases)	Model predictive power with new constant (R^2)
Archard $H = K_w.p.S$	2.64×10^{-7} mm³/N.m	2.0×10^{-7} mm³/N.m	.12
Sliding distance $H = K_w.S$	-	1×10^{-6} mm/m	.04
M-L/M-L+A-P $H = K_w.CS.p.S$	3×10^{-6} mm³/N.m	2.7×10^{-6} mm³/N.m	.58
A/A+B $H = K_w.CS.p.S$	3×10^{-6} mm³/N.m	3.3×10^{-6} mm³/N.m	.60
σ^* $H = K_w.(\sigma^*)^2$	-	1.1×10^{-5} mm³/N.m	.29
M-L/M-L+A-P (no CP) $H = K_w.CS.S$	-	1.43×10^{-5} mm/m	.54
A/A+B (no CP) $H = K_w.CS.S$	-	1.8×10^{-5} mm/m	.49

Table 9: Summary of current and suggested wear constants for different algorithms.

4.5.3. Discussion

It is not possible to speak of an empirically-defined model as being 'correct', since it has no direct analytic derivation. Therefore, the relevant question is: "which model appears to offer the greatest predictive power?" Previously, published studies have only corroborated with individual experimental tests, and so the performance of these models is not well-understood. Undertaking a more comprehensive corroboration requires multiple simulations from different sources, which necessitates faster *in-silico* modelling methods (e.g. rigid body FE- or MBD- based models). The combination of *in-vitro* & *in-silico* wear prediction

methods corroborated together provides the fullest, most powerful toolset for pre-clinical analysis of TKR wear. *In-silico* studies in isolation are subject to suspicion as long as there is no consensus on the precise causal mechanics of wear. But *in-vitro* studies alone cannot provide the same range and volume of information as can be quickly and efficiently evaluated computationally.

Of course, there are important limitations to these studies; the simulation can only perform well if the underlying behaviours are modelled correctly, so the actual mechanical conditions must be accurately captured to set a 'benchmark' for corroboration. A pertinent observation from the multiple test-case corroboration is that there is considerable variability in the experimental results reported in the literature (both within, and especially between, different research centres). This could be due to variations in standard experimental procedure (e.g. whether wear is reported for the counter-face or not, or whether secondary axes such as M-L translation or V-V rotation are fixed or free), or simply due to unintentional errors (e.g. component malpositioning and measurement tolerances). This is a serious confounding factor in attempting to provide a more exhaustive corroboration; the 'noise' due to experimental variability masks the finer influence of the choice of wear algorithm. This can be mitigated to some extent if all the particulars of the experimental procedure are fully reported (and so can be recreated in the computational model), and if tolerances on *in-vitro* uncertainty are reduced to a minimum. Only by corroborating with a 'tighter' set of experimental test results will it be possible to determine with greater confidence which is the most appropriate empirical algorithm for wear prediction (i.e. the best formulation for CS, the true influence of contact pressure & area, etc). To re-iterate: a central conclusion of this study is that it will not be possible to further refine our theoretical models of TKR wear prediction, until more and better experimental data is available to differentiate clearly between proposed algorithms.

Nonetheless, this study clearly has some selective power, e.g. in discounting the Archard/Lancaster sliding distance models (as has been advanced elsewhere in the literature [44, 126]), and supporting cross-shear models. However, the quality of the available data is not adequate to preferentially select individual wear algorithms within this sub-set of cross-shear based models.

The present study compared models with and without a proportional term for contact pressure, in light of current debates about the role of CP in polyethylene wear. The results are not conclusive; both families of models had comparable predictive power; with neither showing a clear advantage. This may indicate that

the range of contact pressures encountered in standard TKR wear tests does not vary sufficiently for the influence to become apparent, or that there are antagonistic factors which have a confounding influence (e.g. increased articular conformity will reduce CP, but may also be influencing lubrication and debris transport). Again, ultimately the best way to resolve this issue is with a greater number of well-defined, targeted corroborations between *in-vitro* and *in-silico* wear analysis platforms.

There are many possible improvements and extensions to the models presented here; besides the challenge of accurately capturing experimental conditions, adaptive models could be used to account for variations in PE depth over time and so investigate long-term wear for each test case (as in [92, 179]), and more detailed deformable FE models could be used to better predict contact pressures, so achieving higher accuracy. Probabilistic methods could be used to attempt to capture the experimental uncertainty *in-silico*. However, whilst all of these are desirable goals, they also all entail a considerable increase in the computational modelling workload, which is not currently justified by the quality of published experimental data.

As understanding of wear mechanics improves, the wear algorithms could be customised to different combinations of articulating materials (e.g. different UHMWPE grades). All these tests are for gait-simulation (mostly based on a derivative of the ISO standard); it would be beneficial and informative to extend this to include a much wider range of activities with more varied loading. However, this would of course require extensive corresponding experimental test data. Corroborating within a single framework for a wider range of implant designs, simulator configurations, lubrication conditions, materials and loading regimes will all ultimately play a part in augmenting our holistic understanding of TKR wear.

This study illustrates the valuable role *in-silico* models can play in exploring and refining fundamental concepts concerning TKR polyethylene wear. It demonstrates that the current generation of CS-based empirical wear models have useful predictive power when corroborated with *in-vitro* experiments and are able to qualitatively rank the wear performance of different designs under different loading regimes; however there is room for further refinement in our current understanding and predictive modelling of wear. The best way to advance our understanding of wear is through greater corroboration between both computational and experimental approaches, to exploit the unique strengths of both domains. By doing so, future pre-clinical analysis tools used for wear

prediction will offer designers a richer, faster and more accurate insight into the causes of TKR wear.

4.6. Conclusions

The ADAMS MBD model now features the capability to model wear based on sliding distance, contact pressure and optionally cross-shear. Wear depths, volumes and contour plots can be reported, as well as cross-shear pressure & sliding distance surface maps. The wear can also be applied adaptively to simulate the 'bedding in' effect associated with long-term wear studies. Further, this functionality has been corroborated with both *in-silico* FE and *in-vitro* wear simulator results, and used for a detailed assessment of the 'state-of-the-art' in *in-silico* wear prediction. This gives good confidence for integrating this additional output metric reporting capability into the framework of subsequent probabilistic studies.

The use of *in-silico* wear prediction within deterministic models has been demonstrated by (amongst others) Knight, using FE-based methods [181], and also Bei [137]. This earlier work included many of the capabilities implemented within this chapter (e.g. adaptive wear, contour-map visualisation, etc), albeit using fewer alternative algorithms and fewer comparisons to experimental data. However, these previous studies were purely *deterministic*, aiming to demonstrate 'proof of concept' *in-silico* of the extant theoretical wear models. They did not attempt to incorporate the effect of uncertainty, and the consequent variation of possible resulting wear outcomes. The ability to include wear within a *probabilistic* study (e.g. to report the typical distribution of wear rates for typical component malpositioning variability) is a powerful additional tool for supporting TKR implant design. All of the necessary components are now in place for probabilistic methods to be applied to these MBD models of TKR mechanics – this will be the objective in the following chapter.

CHAPTER FIVE – PROBABILISTIC MODELLING: CORROBORATION & APPLICATIONS

5.1. Prerequisites: Stochastic Modelling Methodology

With a deterministic rigid-body model generating comparable results to existing FE models, the next objective was to demonstrate and corroborate the proposed stochastic tools and methods in conjunction with this new mechanical model.

Essentially, the statistical modelling environment will sit as a 'wrapper' around the inner mechanical model. This mechanical model is fundamentally the same deterministic model as developed in the previous chapter, although with added parameterisation of various input factors, and additional measurements of output characteristics). As such, there are a range of options for statistical modelling; any software able to 'interface' with the mechanical model to write input variables and read output measures is a potential candidate for use in probabilistic modelling.

For mechanical models in ADAMS, three options were explored for statistical modelling:

- Using the native 'ADAMS Insight' software. This has the advantage that it can very easily interface directly with ADAMS, so no user-coding is required. The main disadvantage is that it exclusively uses *design-of-experiment* (DOE) methods. This means the input values for all trials must be fixed at the start of simulation – the subsequent trial values cannot be dynamically adapted based on the results of earlier trials. This excludes any of the adaptive FPI methods – e.g. FORM/SORM, the AMV family of methods, and adaptive ISM (See Chapter Three for more details). However, other standard methods (RSM or MCST & LHS) are supported by 'Insight'.

- Using third-party software: commercial packages are available for statistical analysis, designed to interface with other modelling environments. One of the most well-established is NESSUS; this software does not natively support interfacing with ADAMS, but by using the custom application support it can be tailored to interface. This is a cumbersome procedure compared to using 'Insight', but has the advantage that NESSUS *does* support adaptive trial sampling, so fully supports all FPI methods.

- Using proprietary coding: the above options are convenient for stand-alone models in ADAMS, however it is sometimes necessary to run ADAMS as a co-simulation with MATLAB/Simulink to facilitate more sophisticated control

plant modelling. In these instances, 'Insight' cannot be used, and using NESSUS becomes even more convoluted. For these co-simulation models, directly encoding a statistical wrapper by hand within MATLAB is a relatively straightforward alternative.

For simplicity of implementation in these first studies, it was decided to use the 'Insight' module. The simulations will use non-adaptive sampling methods: Monte-Carlo and RSM, so 'Insight' is adequate for purpose. However, some later models used the third approach mentioned above (see Chapter Six).

5.2. Probabilistic Corroboration Study: Knee Wear Simulator Mechanics [5]

5.2.1. Background

A recent probabilistic study by Laz et al (2006) [95] performed a preliminary investigation into the use of stochastic methods to measure knee simulator variability. The study used a baseline rigid-body FE simulation of a standard *in-vitro* knee wear simulator setup (the SKS). The aim was to demonstrate the use of probabilistic methods, and to compare the conventional MCST approach with a computationally more expedient AMV approach.

It was decided to attempt to re-create this study with the MBD-environment models, in order to develop the stochastic capabilities needed for later studies; in this way the results generated could be corroborated with existing data in the literature.

5.2.2. Methods

The deterministic model already developed in Chapter Four was used as the basis for this study. It was parameterised to allow user-specified variations for 12 input 'factors'. These input parameters were chosen to duplicate the published study; the twelve parameters are listed in Table 10, along with the values for mean and standard deviation. The published study compared two levels of variability (nominally titled level 'A' - low, and level 'B' - high).

[5] Results in this section are adapted from the conference paper: "Comparison of two methods for probabilistic finite element analysis of total knee replacement" C.Arsene, M.A.Strickland, P.J.Laz and M.Taylor. In: 8th International Symposium on Computer Methods in Biomechanics and Biomedical Engineering 2008: Porto, Portugal. The present author contributed the MBD-based models used in this study.

Factor (Abbreviation)	Mean value	σ (Level 'A')	σ (Level 'B')
F-E Axis A-P position (FEax_AP)	0mm	0.25mm	0.5mm
F-E Axis I-S position (FEax_IS)	25.4mm		
I-E Axis A-P position (IEax_AP)	7.62mm		
I-E Axis M-L position (IEax_ML)	0mm		
Initial F-E angle (Init_Fem_FE)	0˚	0.5˚	1˚
Tibial tilt malrotation (Insert_Tilt)			
Femoral I-E malrotation (Fem_IE)			
Tibial V-V malrotation (Insert_VV)			
Spring M-L separation (Δ_{M-L})	28.7mm	0.5mm	
Spring stiffness (K)	5.21N/mm	0.09N/mm	
M-L Load Split % (ML_Load)	60%	2.5%	
Friction Coefficient (μ)	0.04	0.01	

Table 10: Input factors for probabilistic study (from [95]).

These factors are based on *in-vitro* wear simulators, not *in-vivo* knee replacements. For example, the concept of a 'fixed axis' for F-E and I-E rotation is not applicable to the natural knee; nor do the transverse-plane spring factors directly represent any equivalent *in-vivo* property. Similarly, the levels of variability are based on estimated *in-vitro* simulator setup errors, rather than *surgical positioning* errors (*in-vitro* simulators can be configured more accurately; intra-operative *in-vivo* positioning is generally more variable [148-150]). All factors were assumed to be independent with Gaussian distributions (values were bounded to lie within ±3σ to avoid extreme outliers). Whilst this assumption is tolerable for a 'proof of concept' study, it may have considerable implications. Note also that in some cases, this assumption of Gaussian distribution is clearly inappropriate; for example, such a distribution can always potentially have negative values; this is not suitable for a friction co-efficient which should never be less than zero (a lognormal distribution would be more appropriate). However, for consistency with the published study the same set of conditions will be adopted at present. (Before the method is deployed for practical real-world problems, it would be beneficial to justify these assumptions using experimentally-collected data).

5.2.3. FPI Methods

Laz et al used the AMV method (corroborated with a 1000-trial MCST run) for the analysis, interfacing to NESSUS from the Abaqus FE modelling software (Abaqus, Inc). If 'insight' is used initially then the AMV method cannot be corroborated in this study. However, it is desirable to demonstrate the capability for both

'standard' MCST methods, and also 'fast' (low-cost) alternatives. Therefore a 1000-trial MCST will also be the baseline for this corroboration. In addition, lower-cost response surface methods (rather than AMV methods) will be benchmarked against the MCST results, to explore how suitable these 'fast' models are for subsequent studies. RSM models based on 50, 100 and 1000 samples were generated for comparison to the MCST data. In this first instance, a 'linear' model was selected, for computational efficiency and to provide a first indication of how much non-linearity might be evident in the TKR system.

5.2.4. Corroborating Results with Published Data

For corroboration, figures are presented directly comparing the Laz et al results (finite element) with the present study (ADAMS MBD); in each case, the comparison is for the MCST results. The 1% - 99% envelopes for both levels of variability ('A' & 'B') are shown for A-P translation (Figure 66), I-E Rotation (Figure 67) and Peak CP (Figure 68). Figure 69 compares the sensitivity factors (for A-P translation). The results demonstrate a good correlation with the 'performance envelope' kinematics in the literature data. Similar trends in the envelope size and shape are seen; for example both sets of results show a clear decrease in kinematic variability range in the swing phase (>60% gait) compared to stance phase. Some differences are notable; particularly in swing phase. These differences are partly attributable to differences in the model mechanical set-up (e.g. initial component positioning, and the 'neutral' point for the fixed femoral and tibial axes), and partly to internal model parameters (e.g. dynamic terms; inertia, friction & damping).

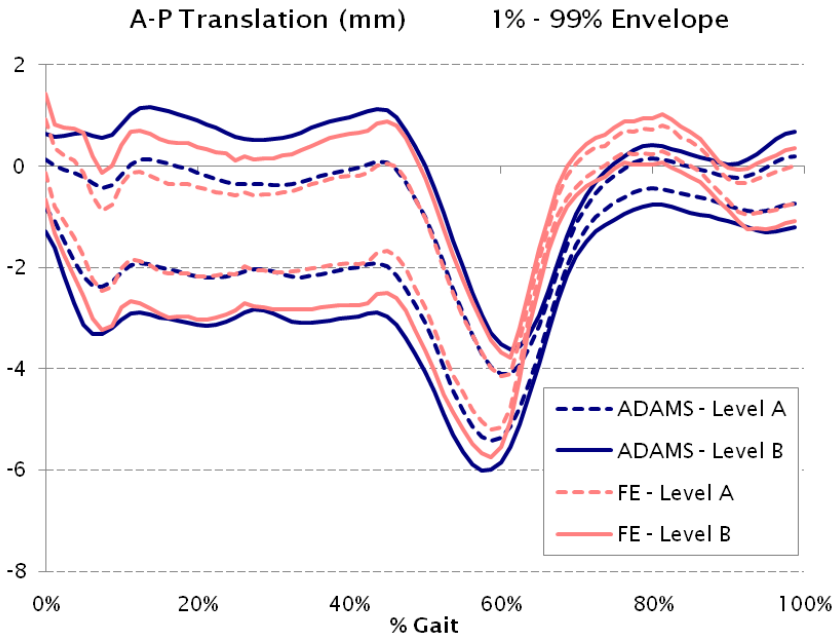

Figure 66: Comparison of MCST A-P envelopes with published data.

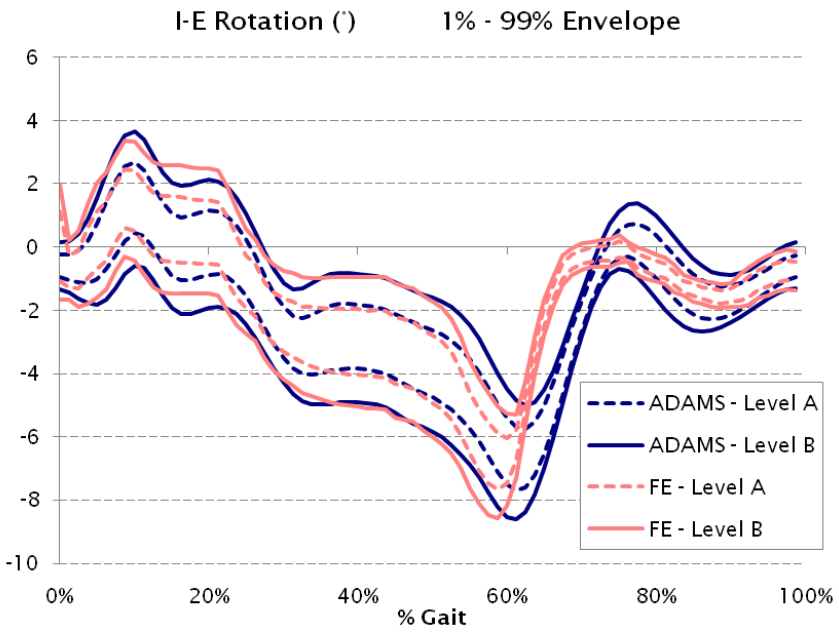

Figure 67: Comparison of MCST I-E envelopes with published data.

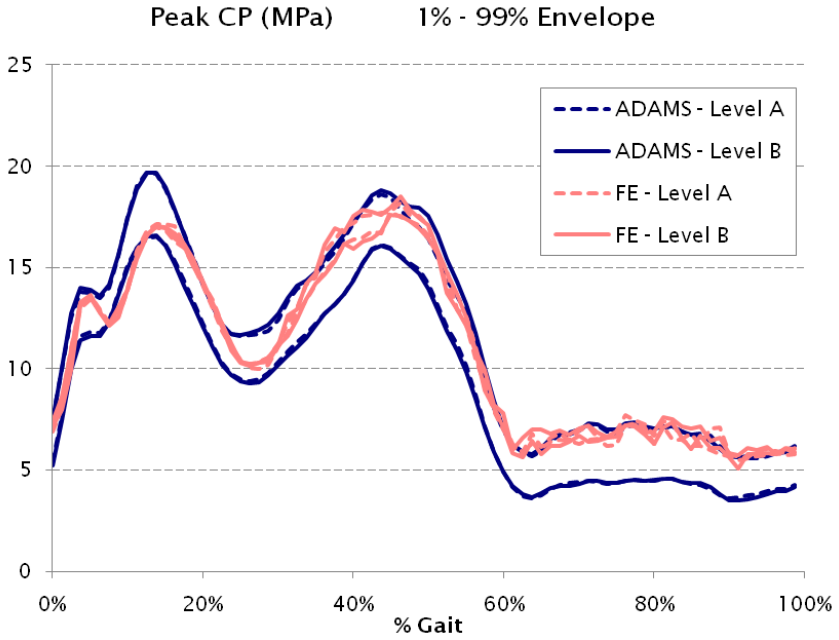

Figure 68: Comparison of MCST peak CP envelopes with published data.

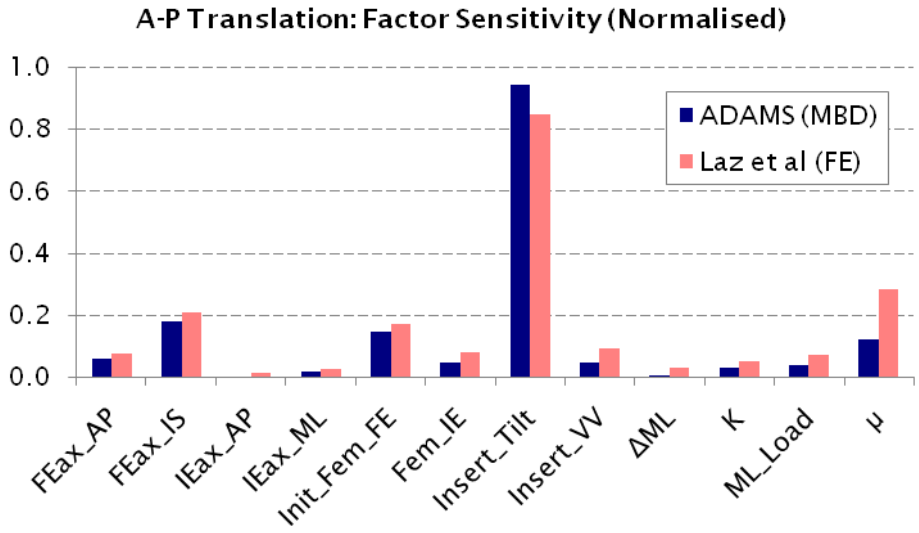

Figure 69: Comparison of MCST A-P sensitivity factors (based on Level 'A' results).

In terms of sensitivity, the same factors were found to be most dominant, with the ranking identical for the top factors. For A-P and I-E kinematics, the corroboration was particularly strong, with sensitivity differences well under 0.1 on the normalised scale in every case. The differences were somewhat larger for the contact pressure data (up to 0.2), due to the different formulation of contact model used.

5.2.5. Comparing 'Fast' Methods (RSM) with MCST

Results were compared for 'level A' variability using RSM & MCST (Figure 70 and Figure 71). For kinematics (A-P & I-E) in particular, RSM closely matches MCST; 50 or 100 samples are adequate to approximate the envelopes; there is no benefit in having more trials (see Figure 70 as an example). This suggests that the kinematics are not particularly non-linear.

However, peak contact pressure is more non-linear (partly due to numerical solver 'noise'), resulting in poorer response-surface fitting, especially if fewer samples are used – although regardless of the number of trials, the fit is not perfect (see Figure 71).

Nonetheless, RSM provides a fair first-approximation to the peak CP envelope, which may be adequate for exploratory studies where high accuracy can be sacrificed for speed of evaluation.

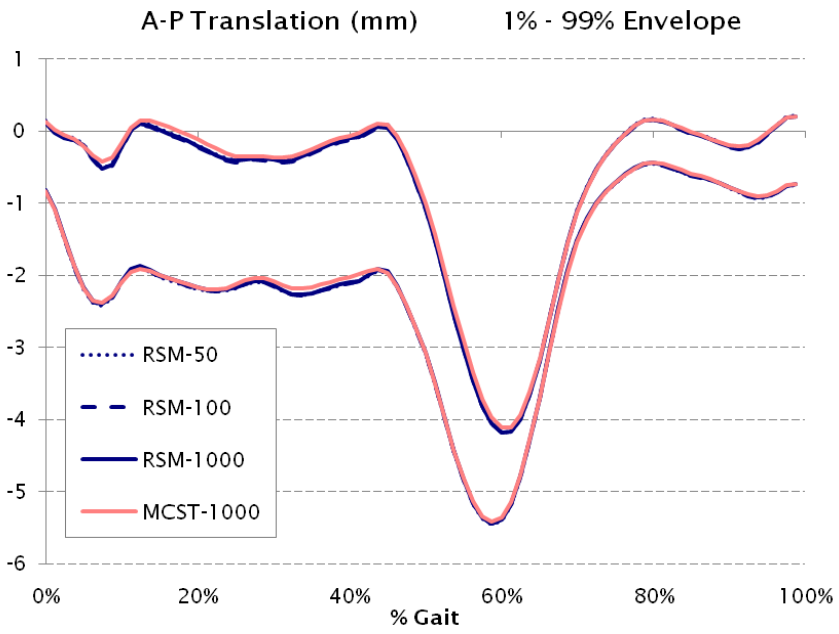

Figure 70: Level 'A' MCST vs. RSM (based on 50, 100 or 1000 trials): A-P translation.

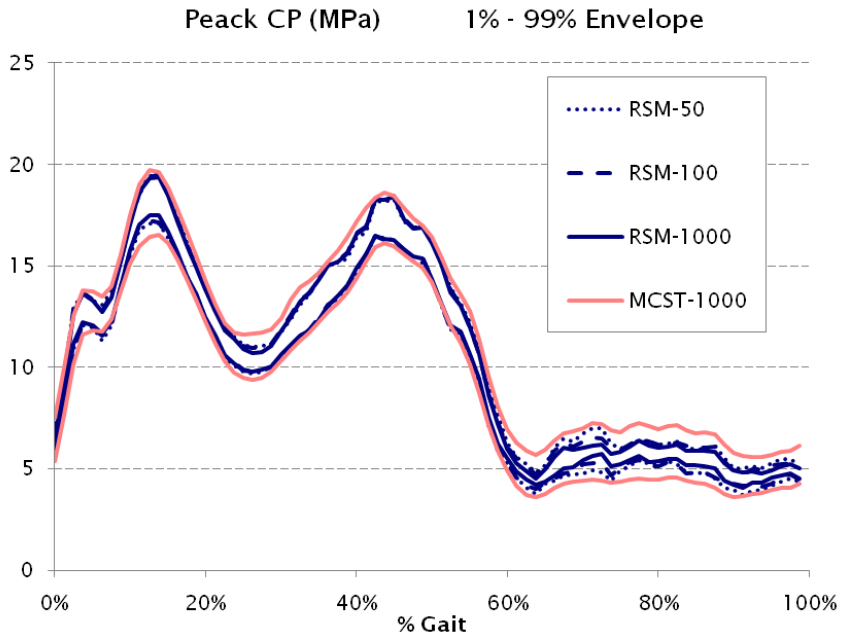

Figure 71: Level 'A' MCST vs. RSM (based on 50, 100 or 1000 trials): Peak CP.

Discussion: Despite some evident differences, the results presented are broadly comparable to the study by Laz et al. This is encouraging, since it corroborates the proposed software & methods with data from the literature. Differences can be accounted for by the slightly different modelling approach used in the published study. There, non-deformable FE was used with elastic-foundation contact, and the statistical method used to produce the envelope points was AMV. As discussed in Chapter Three, there are potential issues with AMV for a very irregular or non-linear system behaviour – the contact pressure in particular is prone to numerical 'spikes' and oscillations; these can cause particular difficulty with AMV, and differences between AMV and MCST become more apparent [179].

The comparison between RSM and MCST is promising; for the system studied here, RSM is able to produce very comparable results to MCST with 10-20 times less computational workload. (In terms of solution time, the MCST required around 5 days for 1000 trials; by comparison RSM with 50 trials required only 6 hours). Note that the quality of this match depends on the variability levels considered in the 'performance envelope'. For lower ranges (e.g. 10%/90% or 5%/95%) the match is exceptionally good; however for more extreme ranges (e.g. 1%/99% or broader) the accuracy deteriorates. This is because these values sit on the 'tails' of the distributions, and as discussed in Chapter Three, obtaining an accurate prediction of distribution tail values is difficult for FPI methods. Therefore 'fast' probabilistic methods are more suited to reporting performance envelopes within a narrower range (2 standard deviations or lower).

It is always prudent with any new model to first validate the FPI method using a more exhaustive MCST simulation; however the close agreement here demonstrates that there are only limited differences in results due to the FPI techniques.

Note one very important observation: although the FE and MBD approaches have been *corroborated* here, it is not possible to *validate* either approach with *in-vitro* data because comparable probabilistic experimental data for this virtual test is not available (the study is purely theoretical). This is a major potential limitation; until probabilistic computational methods have been corroborated with probabilistic experimental data (and not just deterministic experimental data), questions will remain about the integrity of the *in-silico* tests in isolation.

5.2.6. Extension Work: Further Results & Discussion

A very important point to recall is that the AMV method used by Laz et al has a key limitation: every individual sample point of interest requires its own ***unique*** evaluation using AMV. This means that monitoring multiple objectives for the same experiment *actually increases the computational overhead*. To appreciate the impact this has, consider the case in this example:

If the model has 12 input factors, then 12+1 = 13 trials are required for the initial 'MV' analysis. Now, A-P, I-E and CP are monitored throughout the gait cycle, at a sample-rate of 80 samples per cycle. For each output, two levels are monitored (the 'low' 1%, and the 'high' 99% levels). This means that for these 3 output measures, 2x3x80 = 480 trials are required; in addition to the original 13 – this is now close to 500 trials. *But* if it was desirable to monitor a new output (say V-V rotation, or contact area) at the same sample rate, this would add another 2x80 = 160 trials for each output. Hence with only a few more time-varying outputs the number of trials required could exceed the 1000-trials needed for Monte-Carlo.

Compare to the RSM model; only 100 trials are used; this is a fixed overhead. The RSM results will be less accurate than AMV at any targeted point. But *any number* of output measures can be retrieved from the model, as it is not 'tuned' to local points in the possibility space, but is a broad 'global' fit (as discussed in Chapter Three). This means that AMV is well-suited for a small number of outputs when accuracy is important, but RSM may be a better choice for an investigatory study such as this, where many outputs may be of interest.

To illustrate this advantage, a number of additional metrics are reported below which were not included in the original publication by Laz et al. Together, they

provide additional insights into the influence of variability on knee simulator mechanics.

Contact Area

Contact area can easily be retrieved for the model. An envelope showing contact area variability (1%-99%) is shown for the two levels in Figure 72. Variations are limited compared to the contact pressure envelopes; however contact area variations may be important, as some studies suggest total wear volume correlates better to contact area than contact pressure [106].

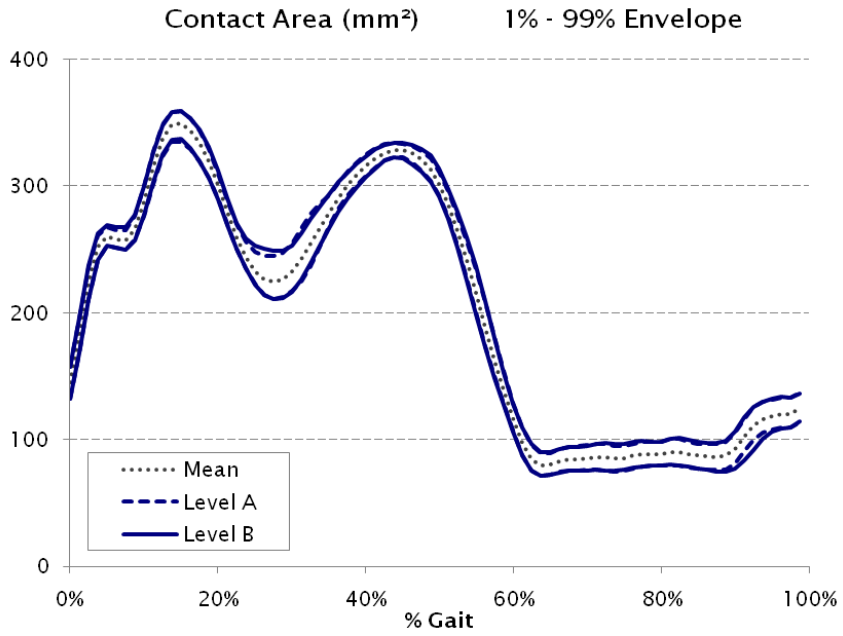

Figure 72: Contact Area envelopes (1-99%).

M-L Load Split

Medial-Lateral load split is an important metric as it defines the amount of loading on each condyle, which in turn affects the kinematics of the knee as well as polyethylene stresses [136]. The ISO-standards go to considerable lengths to define a 'controlled' application of M-L load split [23, 24] (based on the *assumption* that in a 'normal' knee the loading is split approximately 60% on the medial condyle, 40% on the lateral condyle). However this load-split is achieved by specifying a 'fixed' translation along the M-L axis for the application of the axial force. As such it cannot account for dynamic changes in component positioning throughout the gait cycle. Only recently has the actual M-L load split been measured *in-vitro* by Zhao et al [182], this revealed considerable variation from the 'target' 60-40 Medial-biased loading. However, this was for one subject only and did not take account of positioning variability. Here, we have post-processed M-L load split in the present probabilistic study; the results are shown in Figure 73. There is very considerable variability in the load split, which is not apparent

from monitoring only the peak CP (as in Figure 68). Note that extreme values (close to 0% or 100%) would indicate uni-condylar loading occurring. This would have implications beyond the modelling scope of this study, since condylar "lift-off" is known to be associated with a considerable increase in wear rates [172] which cannot be readily explained by existing theoretical wear models. Here, although the peak values approach 80% (in swing-phase, corresponding with higher flexion), the loading is never entirely uni-condylar.

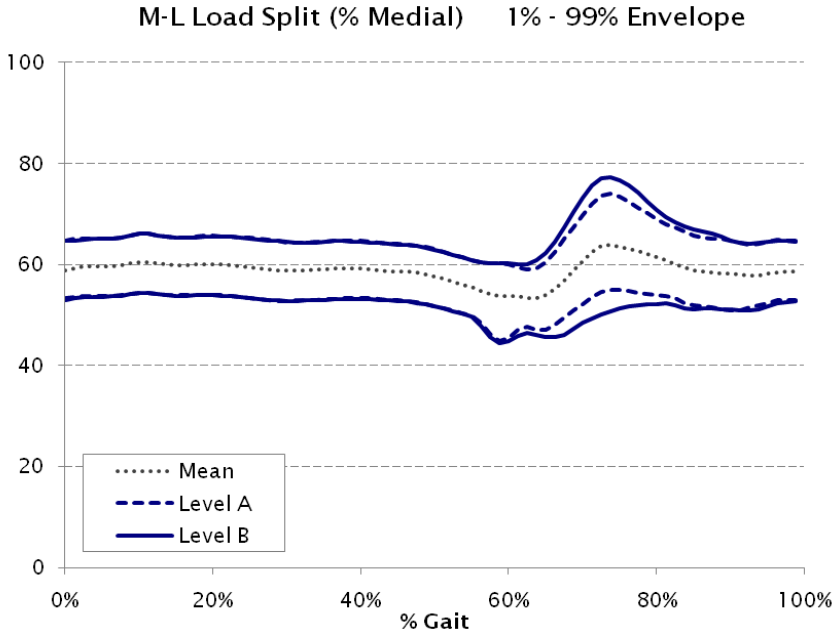

Figure 73: Medial-Lateral Load Split (1-99% Envelopes).

'Relative' Kinematics (A-P & I-E)

In the original published study, A-P and I-E kinematics were reported as 'absolute' values. This means that the 'zero' position was based on the settled component positions for the *deterministic (unperturbed)* case, *for every trial*. So, if a trial had a 3° shift of femoral I-E rotation, and then traced a similar motion to the unperturbed case, then the 'envelope' would show an offset 'variation' of 3 degrees. The 'absolute' values are useful for some purposes (e.g. laxity/ROM assessment), but not for others. It is possible that the *relative* motions could actually be completely unchanged (merely offset), but the variability envelope would still appear to be very large. The relative motions are of particular importance for understanding contact paths for wear prediction. Therefore it is useful to report the *relative* kinematic envelopes alongside the *absolute* kinematics (i.e. for each trial, the 'zero' position is based on the average component positions *specifically for that trial*).

The 'relative' envelopes are reported here for A-P translation (Figure 74) and I-E rotation (Figure 75). It is apparent that these offset envelopes are quite different

to the 'absolute' envelopes (Figure 66 and Figure 67 earlier). These 'relative' envelopes are narrower than the envelopes for 'absolute' values, showing that while the precise location of the contact may vary considerably, the actual kinematics relative to the starting point are not as variable. (This is important, because wear rate will be related to this degree of relative motion; i.e. the total sliding distance).

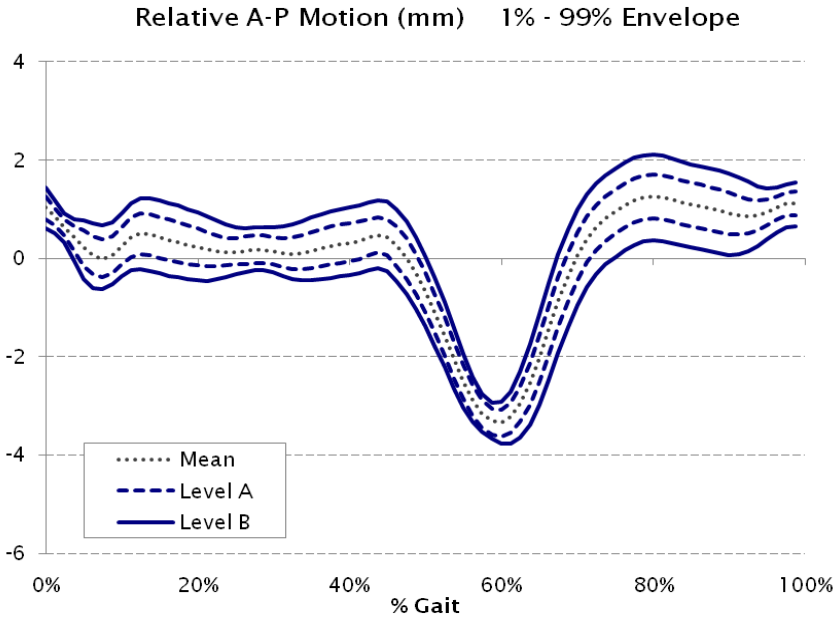

Figure 74: Envelope of relative kinematics for A-P translation (1% - 99%).

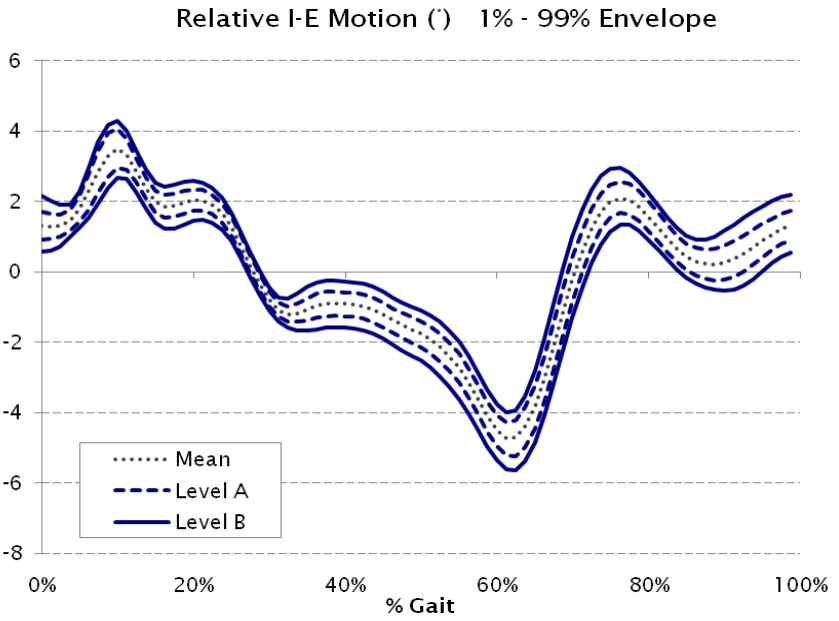

Figure 75: Envelope of relative kinematics for I-E rotation (1% - 99%).

These further output metrics further demonstrate the range of data that can be extracted from a probabilistic study. They also illustrate that the choice of FPI method has implications for how many metrics can be reported; using AMV introduces an overhead for each new output monitored, so may not be the ideal choice for a 'data-rich' investigation where many metrics are of interest.

5.3. Concept-Study: Can Passive Laxity Predict Gait Mechanics? [6]

5.3.1. Background

The probabilistic framework based on ADAMS-MBD has now been corroborated with published FE-based methods. This gives greater confidence for using the MBD models for other probabilistic studies.

The study in the previous section considered only one *activity*, (although multiple output metrics were monitored for this activity). An interesting feature of probabilistic modelling is that *multiple* activities can be compared together to explore potential correlations (see Chapter Three).

In this section, the model is extended to form an original study demonstrating conceptually how probabilistic studies can provide a framework to explore relationships not just *within* but *between* different activities. In this study, two distinct classes of activity are compared: the *'passive'* laxity motions of the knee, and the kinematics & peak contact pressures experienced in an *'active'* gait cycle.

Simple passive laxity drawer loading can readily be performed intra-operatively, but the question of whether these tests can yield information about the likely post-operative 'active' performance of the knee has yet to be rigorously addressed. Currently, this is a subjective judgement based on the expertise of the clinical professional.

A comparison using simulation methods may allow more quantitative statements to be made about the predictive power (and hence practical value) of such passive laxity tests. The present conceptual study will demonstrate how such an investigation might be structured, using simplified computational simulations of simulated gait and laxity drawer loading.

5.3.2. Methods

This study is based upon an adaptation of the probabilistic setup corroborated in the previous section [95], and incorporates several factors included in that study (misalignment, friction and M-L load split). However, various developments are introduced in order to explore the passive/active performance correlations.

[6] This section is adapted from the journal article: "Could Passive Knee Laxity be Related to Active Gait Mechanics? An Exploratory Computational Biomechanical Study Using Probabilistic Methods". Strickland et al, CMBBE (*In Press*).

Mechanical Modelling:

The study will compare the two variant knee implants described in Chapter Four (S/C and U/C), in order to compare designs and consider the influence of design-specific variations (in this case, sagittal-plane conformity).

The standard configuration for the SKS (as used in the earlier models) features springs in the transverse plane to provide A-P and I-E restraint only. However, to conceptually explore correlations with laxity range tests (which include out-of-plane forces and moments), here a simplified 3-D restraint model was used instead. This uses nonlinear spring-elements to approximate the combined restraint provided by the knee ligaments, within the range-of-motion of interest (0˚ to 60˚ flexion).

Since every new spring element increases the number of factors involved (which greatly increases the number of simulations required), the restraint model was kept to a rudimentary minimum of three elements: together providing the necessary I-E, V-V and A-P restraint.

Note that this model is *not* representative of the complexity of true ligament restraint *in-vivo* (with multiple bundles and insertion/origin sites [7]). Further, in an intact knee *in-vivo*, the degree of laxity in full extension would be reduced to a minimum by the 'locking' mechanisms discussed in Chapter One; however in this simplified computational model there is still some laxity, since some of the important contributors to restraint (the capsule and surrounding tissues, and the patellar restraint) are omitted.

Nonetheless this model provides a comparable 'aggregate' restraint force across the gait envelope, when compared to 'physiological' ligament-based models, as described in the literature [183, 184]. Although not identical, the envelopes are sufficiently similar for conceptualisation purposes (Figure 76 shows the envelopes for the S/C insert throughout the flexion range, for both A-P translation and I-E rotation).

Unified 'mean value' properties for the spring-restraint elements (i.e. the stiffnesses and pre-strain levels selected for each spring-element) are listed in Table 11.

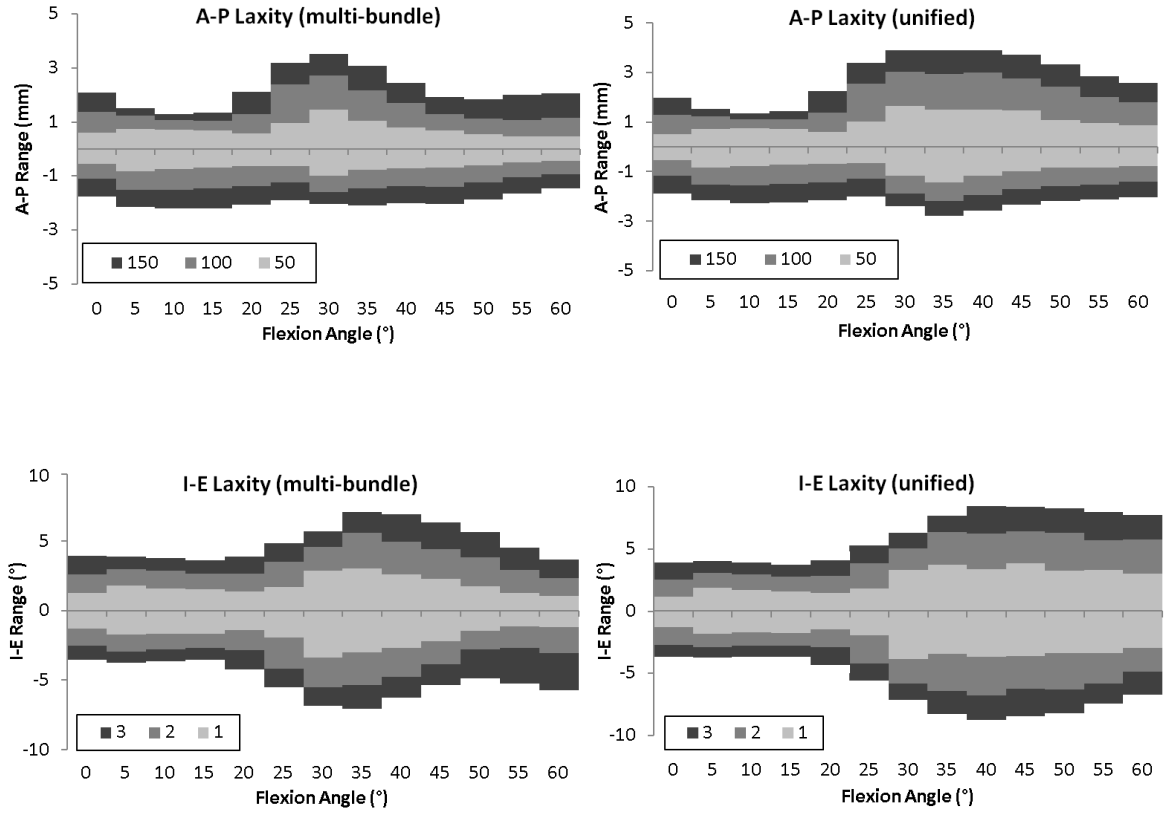

Figure 76: Comparison of laxity envelopes for multi-bundle and unified spring-restraint models, at three levels of drawer force (±50, 100 & 150N) and torque (±1, 2 & 3N·m). The envelope 'width' shows the degree of laxity at a given flexion angle.

The nonlinear force (F) / strain (ε) relationship for the spring elements was adopted from previous analytic studies [185]:

$$F = \begin{cases} 0 & \varepsilon \le 0 \\ \left(\frac{k}{4\varepsilon_1}\right)\varepsilon^2 & 0 < \varepsilon \le 2\varepsilon_1 \\ k(\varepsilon - \varepsilon_1) & 2\varepsilon_1 < \varepsilon \end{cases} \qquad (5)$$

Where the instantaneous strain (ε), is defined relative to the initial spring-element length (L_0), the instantaneous length (L), and the 'pre-strain' (ε_p) which determines whether the element is considered to be under tension at its initial length:

$$\varepsilon = \frac{L(\varepsilon_p + 1)}{L_0} - 1 \qquad (6)$$

(Note the similarity between this analytic expression, and the experimental curves presented in Chapter One - Figure 2). There are three controllable parameters within this non-linear model: linear-region stiffness (k), toe-in (ε_1) (which determines the strain level at which the linear region begins), and pre-strain (ε_p), as a percentage of the natural length.

Variability levels for these factors were derived from the literature [6], with the three parameters (ε_1, ε_p and k) for each of the three spring-elements giving a total of nine variable factors. This is more appropriate for a demonstration study to maintain a moderate number of factors in total; however a fully-featured musculoskeletal model would need many more factors to be accurately representative of *in-vivo* dynamics – this in turn would require many times more simulation trials.

Combined with the factors adopted from published studies, the complete set of input variables are listed in Table 11; all variables are assumed to be independent, following a Gaussian distribution bounded at ±3σ.

Factor	Mean (μ)	S.D. (σ)
F-E axis I-S	25.4 mm	0.5 mm
F-E axis A-P	0 mm	
Tibial axis M-L	0 mm	
Tibial axis A-P	7.62 mm	
Initial F-E angle	0°	1°
Initial I-E angle		
Initial tilt angle		
Initial V-V angle		
Friction (μ)	0.04	0.01
M-L load split	60M-40L	2.5%

Factor	Mean (μ)	S.D. (σ)
LSpr k	70 N/mm	20%
MSpr k	100 N/mm	
PSpr k	130 N/mm	
LSpr ε_p	+5%	1%
MSpr ε_p	0%	
PSpr ε_p	+2%	
LSpr ε_1	+3%	1%
MSpr ε_1		
PSpr ε_1		

Table 11: Input factors, with mean & SD. The new terms (right) are related to the spring-restraint ('M'edial, 'L'ateral and 'P'osterior 'Spr'ings); ('*k*' is the linear stiffness; 'ε_P' is the pre-strain, & 'ε_1' is the toe-in, as a percentage of natural length).

As in previous studies, output kinematics and peak contact pressures were analysed through a standard 1-second gait cycle, based on ISO-derived force/displacement input waveforms [91] (Figure 77).

For gait, the selected output measures were A-P translation & I-E rotation, and peak contact pressure, sampled throughout the cycle. Kinematics are reported in terms of 'offset' values; i.e. normalised relative to the initial equilibrium position.

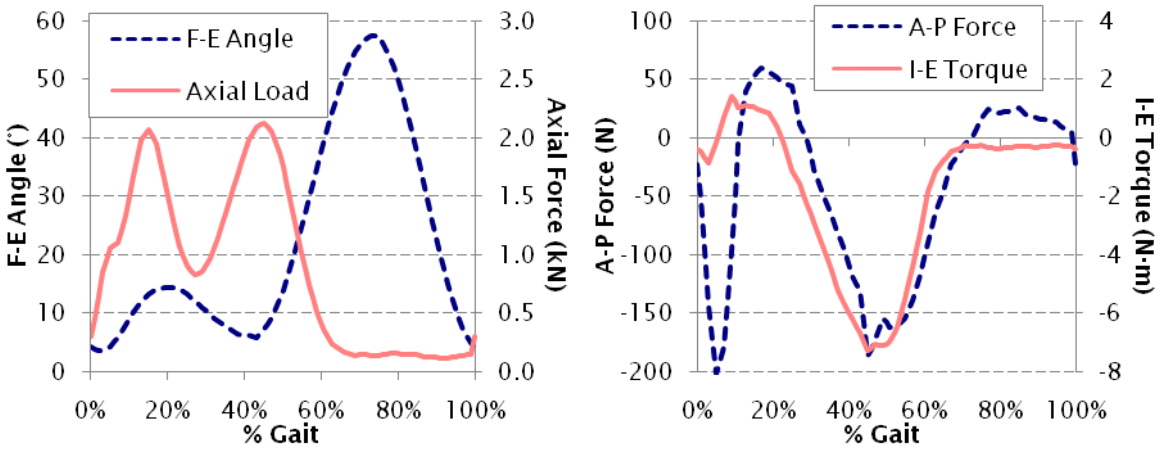

Figure 77: Input waveforms for force-driven gait simulation (Adapted from [91]).

Additionally, three paired tests of passive laxity drawer loading were simulated with typical clinical passive loading levels [186]: anterior-posterior (A-P) draw (±100 N), internal-external (I-E) torsion (±5 N·m), and varus-valgus (V-V) torsion (±10 N·m). Laxity loading was simulated both in full extension and at 20° flexion (reflecting the clinical practice of testing at high-laxity flexion angles associated with stance), with compressive axial loading limited to 300N for 'passive' restraint [187]. For these output measures, the displacement (translation or rotation) was reported relative to the initial 'offset' reference position (when unloaded). A 'positive' value indicates displacement in the direction of the applied force or torque.

Statistical Modelling:

The key concept of this study is that by analysing the output of multiple different tests, it may be possible to identify correlations *between* them; this could potentially allow a test with one activity to be a predictor for the probable outcome of a different activity, e.g. allow passive laxity to be a predictor for active gait.

For this study, in order to provide a matched set of trials (to directly compare correlations as discussed in Chapter Three), a randomised 1000-trial matched MCST analysis was performed, with the same matrix of input factor settings used for both the gait cycle and laxity draw simulations.

Because this simplified model did not include capsule or musculature contributions to joint restraint, the range and levels of variability studied meant a handful of statistically outlying trials resulted in subluxation under high laxity drawer loading. This was particularly the case for those trials where the spring stiffnesses were lower (especially the 'MCL' spring-element, which averaged below 50% of its mean stiffness in the subluxation outlier cases). Without the

additional restraint that would normally be provided by other sources (e.g. joint capsule, patellar mechanics) a value of stiffness significantly below the 'mean' level in these spring elements alone cannot provide adequate restraint to realistically constrain the tibiofemoral mechanics. These outlier trials were therefore excluded from the subsequent correlation analysis.

Results from the simulations were used to determine 1%-99% performance envelopes for gait cycle kinetics & kinematics, and to determine the statistical distributions for laxity drawer displacements for both designs.

To identify correlations between active gait and passive laxity, scalar statistical metrics for the time-varying gait waveforms were required for each trial. The waveform *minimum*, *maximum*, *mean*, *range*, and *standard deviation* were chosen for this purpose. Each of these five values was calculated for the three gait cycle output measures, and the results from all trials were correlated with the three pairs of laxity drawer displacements, giving a 3×15 correlation matrix. This matrix was generated for both the S/C and U/C implants to allow comparison between designs.

5.3.3. Results

Active Gait Simulation Characteristics

Probabilistic performance envelopes for the simulated ISO-wear gait cycle were calculated for comparison with the previous studies. The gait kinematics and peak contact pressure are shown in Figure 78. Both S/C and U/C designs are included on the same axes for comparison; the kinematics are reported as 'offset' values to more clearly illustrate design-specific differences between S/C and U/C designs.

For the S/C design, it is apparent that the offset A-P and I-E motions of the knee are more closely constrained during stance phase, whereas the U/C design permits more variability of motion. For both designs, the variability envelope expands to its widest during swing phase (due to lower articular surface conformity and lower compressive forces). The differences between the S/C and U/C designs are also apparent in general envelope trends, with larger envelopes for the U/C design (reflecting the larger mean-value kinematics and contact pressures).

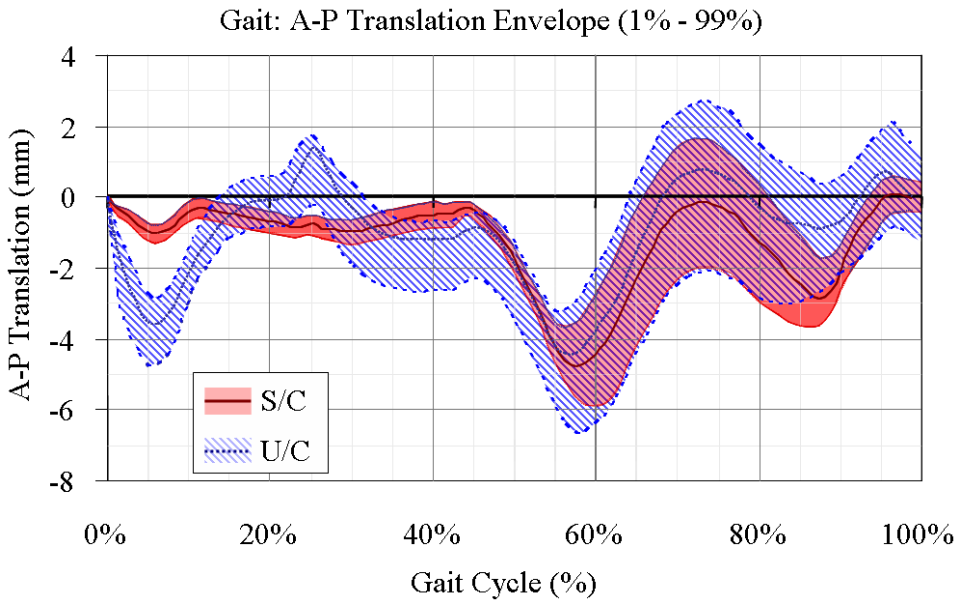

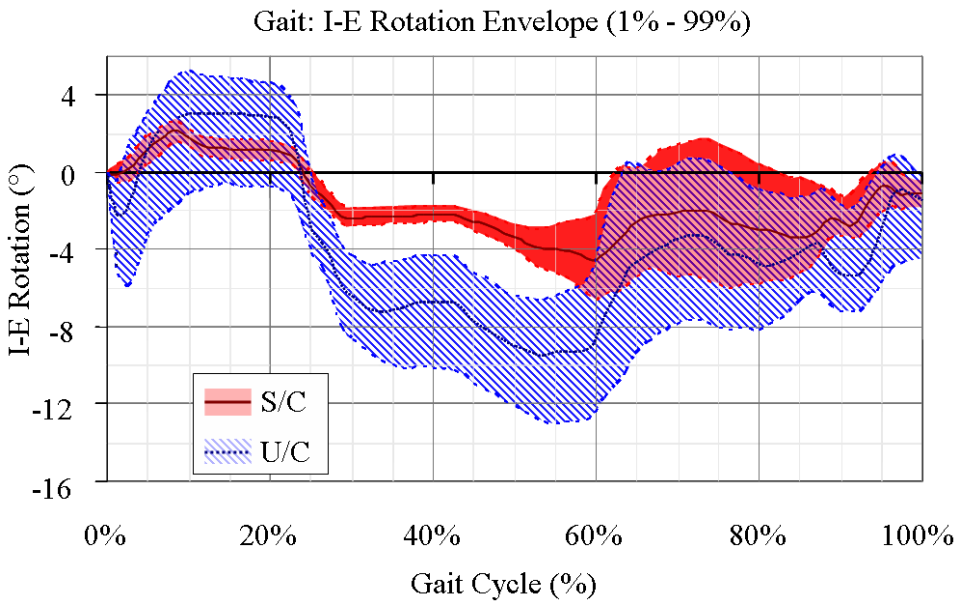

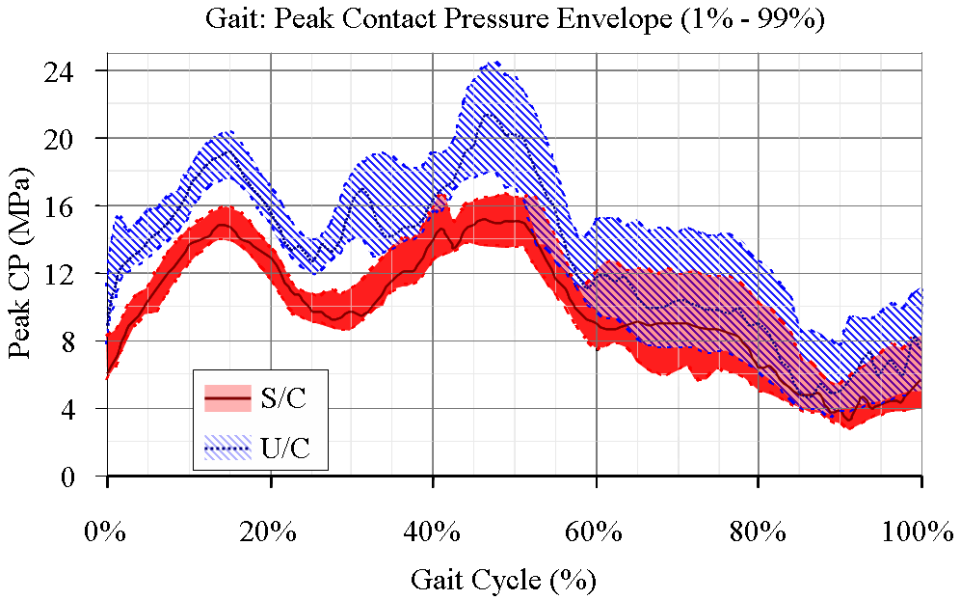

Figure 78: Kinematics and peak contact pressures for gait simulation, with variability envelopes (1% - 99%). Solid fill: S/C, hatched fill: U/C.

Passive Laxity Drawer-Loading Distributions

The distributions of laxity draw range for A-P, I-E and V-V are shown in Figure 79. The 'range' is the total difference in displacement between the two opposite draws; e.g. between the displacement for posterior draw of -100N, and the displacement for anterior draw of +100N.

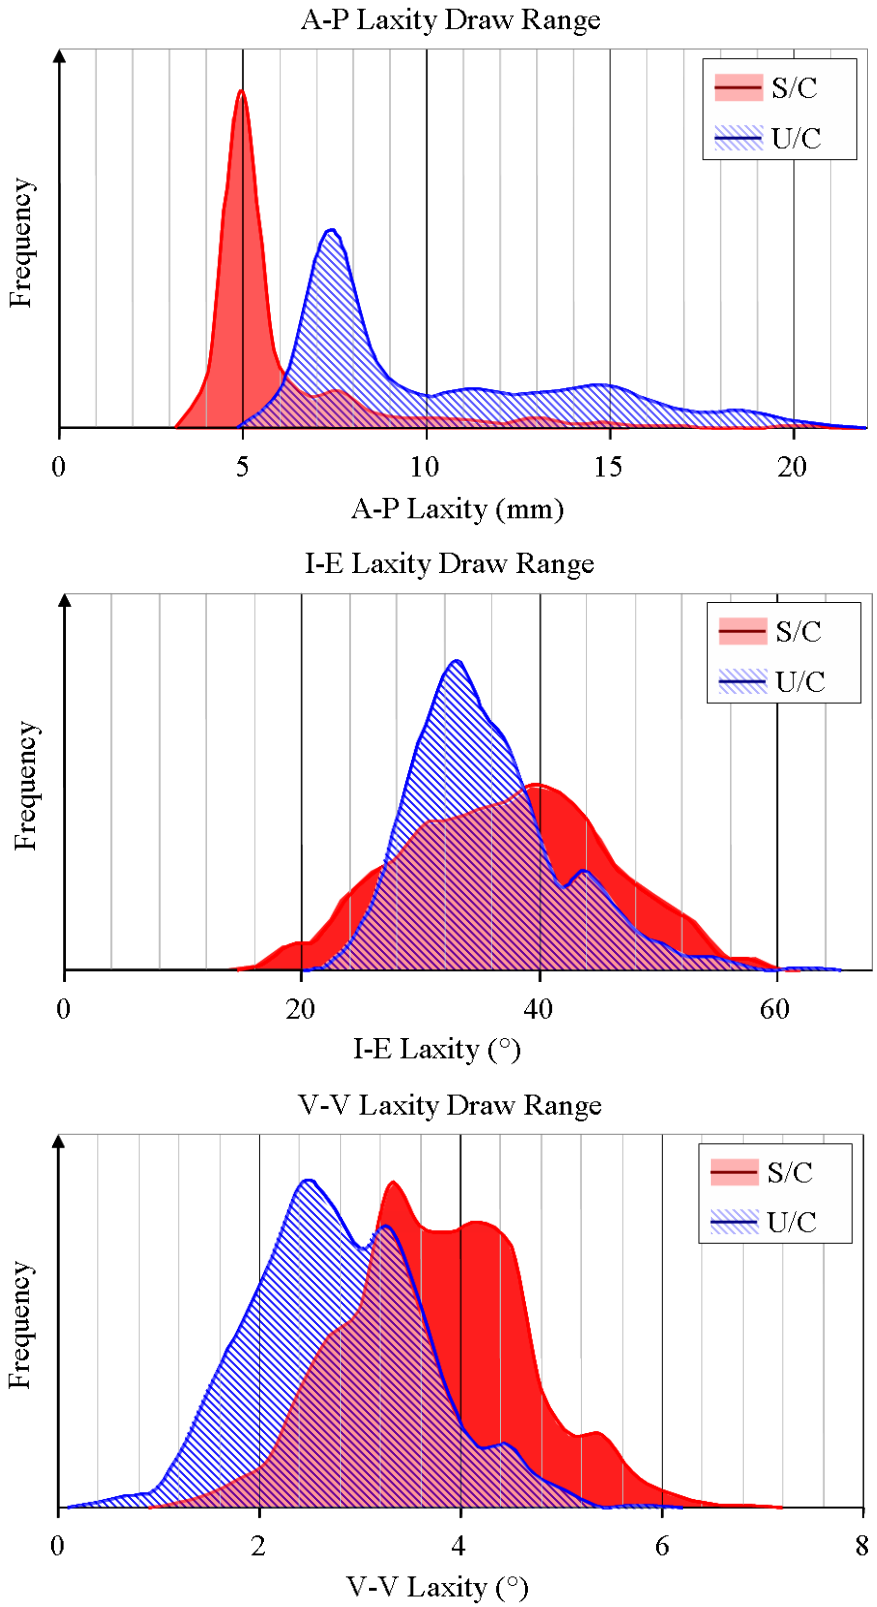

Figure 79: Distribution in laxity draw ranges due to input variability. Solid fill: S/C, hatched fill: U/C.

It is clear that the greatest difference between designs is for the A-P drawer loading, where the lower sagittal conformity of the U/C design allows higher draw ranges. The distributions of I-E laxity for the two designs lie within a similar range, whilst for V-V laxity, the distributions are very similar in shape, with higher laxity for the S/C design.

Passive-Active Correlations

The correlations are reported in terms of Pearson-squared (R^2) values in Table 12 with the strongest correlations highlighted. Note that these values indicate the strength of the correlation only (not whether a correlation is positive or negative).

		S/C			U/C		
		A-P	I-E	V-V	A-P	I-E	V-V
GAIT A-P	MIN	0.05	0.03	0.01	0.12	0.04	0.00
		0.01	0.00	0.03	0.02	0.00	0.05
	RANGE	0.05	0.08	0.12	0.10	0.06	0.07
	MEAN	0.02	0.00	0.01	0.07	0.01	0.01
	ST.DEV	0.07	0.04	0.02	0.06	0.04	0.08
GAIT I-E	MIN	0.01	**0.12**	0.05	**0.15**	0.08	0.00
	MAX	0.01	0.00	0.09	0.02	0.02	**0.21**
	RANGE	0.01	**0.15**	0.10	**0.20**	**0.36**	**0.34**
	MEAN	0.05	0.09	0.04	**0.16**	0.07	0.00
	ST.DEV	0.03	**0.12**	0.09	**0.25**	**0.38**	**0.33**
GAIT CP	MIN	0.08	**0.20**	**0.31**	**0.23**	**0.31**	**0.42**
	MAX	0.00	0.01	0.00	0.09	0.02	0.02
	RANGE	0.03	0.10	0.10	**0.24**	**0.13**	0.03
	MEAN	0.04	**0.36**	**0.36**	0.06	**0.27**	**0.33**
	ST.DEV	0.06	**0.43**	**0.40**	**0.41**	**0.52**	**0.33**

Table 12: Correlation matrix: active gait parameters (rows, headings left) versus passive laxity draw ranges (columns, headings top) for S/C & U/C designs.

The correlation coefficients are mostly low; this is to be expected of a complex mechanical system with multiple influential factors. Nonetheless, some of the correlations are sufficient to provide some degree of predictive power (R^2 up to 0.5). Notable trends are apparent for both the S/C and U/C designs, especially under the I-E and V-V torsional loading. The largest difference between the two designs was for A-P drawer loading; this would be expected, since the design difference between the two is in the sagittal plane, and would most directly affect A-P laxity. The more conforming S/C design showed limited predictive power between active & passive mechanics (with laxity motion restricted), whereas the U/C design had higher correlation coefficients. As would be expected, some of

these metrics are inter-related; e.g. a good correlation for 'range' tends to be accompanied by a good correlation for 'standard deviation' (although note that this is not necessarily the case, depending on the shape of the time-varying waveforms).

For active gait parameters, the strongest correlations occurred for contact pressures, with moderate correlations for I-E rotation and very little correlation for A-P translation. The tests for laxity drawer loading had greater predictive power for variations in minimum (i.e. swing phase) contact pressures, suggesting that the influence of the modelled restraint force is causing this correlation. I-E rotation was not well-correlated for the S/C design; only small I-E rotations occur for this design during gait, so any correlations are likely to be less evident. For illustrative purposes, Figure 80 provides representative example correlation scatter-plots for the weakest and strongest correlations observed. Note that in the 'weaker' correlation plot (A-P range for the S/C design) the limited correlation in the main 'clustering' of trials is masked by the number of outlier trials with greater A-P laxity. (In this particular case, the high-laxity outliers are due to lower pre-strain of the 'PCL' spring-element permitting greater motion).

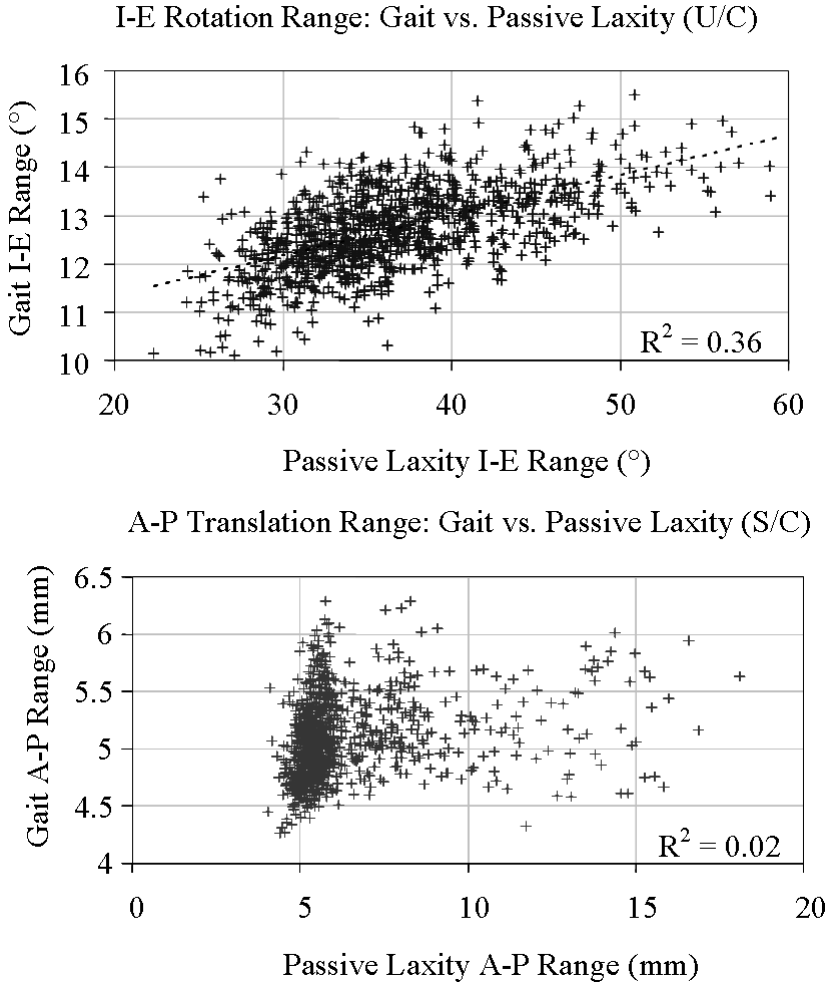

Figure 80: Examples of observed correlations: stronger (I-E laxity range for U/C), and weaker (A-P laxity range for S/C).

5.3.4. Discussion

The performance envelopes predicted for the normal gait cycle are comparable to the previous corroboration study. A larger degree of output variability is evident in the present study, due to the additional input variability factors related to the spring restraint model. This is most apparent in the peak contact pressure envelopes, where a variability range of up to 6MPa is seen for both designs. The effect of compressive forces due to the spring element restraint will increase contact pressure ranges; this is further compounded by malpositioning leading to exaggerated gait kinematics with lower tibiofemoral conformity at contact, and hence higher contact pressures.

When the envelopes for S/C and U/C designs are compared, these results do suggest that insert design *can* play a role in controlling the influence of variability on gait mechanics. However, this may be specific to the simplified mechanical configuration being demonstrated in this conceptual study, and further studies using more extensive models would be required to confirm this observation.

The laxity loading reveals a high degree of torsional laxity in these simulations at both flexion angles (0˚ and 20˚); this must be interpreted in light of the reduced transverse-plane restraint provided by the spring elements used in the model (comparable cadaver experiments using this tibiofemoral test configuration have also yielded unusually high levels of rotation [115, 188]). Trends were similar at both flexion angles simulated, with magnitudes of laxity range generally greater for the 20˚ position, as would be expected. As anticipated, the most apparent differences between designs were for the A-P drawer loading; it is notable that when the full distribution of variability is considered, higher V-V laxity is evident for the more constrained S/C design than the U/C (this is associated with greater I-E rotation of the tibial component under V-V torques for the S/C compared to U/C designs, suggesting that this rotation may facilitate the higher V-V laxity without requiring condylar lift-off).

Although the range of passive laxity motion is very low for V-V rotation, the correlations are generally strongest; this is most likely because much of the correlation is due to variability in the 3-D restraint model: Whereas A-P and I-E loads act in the transverse plane, V-V loads are out-of-plane and can directly distract the joint, resulting in increased restraint forces and so giving more ready indication of variability in the soft-tissue constraint. Note that many *in-vitro* simulators model spring-restraint only in the transverse plane, and so cannot simulate this V-V restraint. Again, this illustrates the influence that the choice of mechanical model may have upon the study outcomes.

As was noted in the results, some correlations may be masked by outlier effects, due to one particular input variable (e.g. the PCL 'pre-strain' term for A-P range with the S/C design). This has an important implication; it is possible that there are other important input variables that have not been included in this study, and the inclusion of these further additional inputs might serve to weaken, or conversely strengthen, the observed correlations.

This conceptual study has explored the relationship between the influence of variability on passive laxity and gait kinematics & kinetics, for two specific TKR design variations, using a simplified mechanical model of the tibiofemoral joint with rudimentary 'soft tissue restraint' representation. Correlations were demonstrated for certain parameters: in some cases, with predictive powers up to $R^2 = 0.5$. This *may* allow design-specific predictions about gait mechanics to be made based on tests of laxity drawer loading; for example, high V-V laxity means it is more probable that a knee with the U/C insert will experience greater I-E rotation in gait; the same trend is less probable for the S/C insert. This becomes clinically relevant when these mechanical observations are related to modes of failure; for example, studies have associated more pronounced I-E rotational kinematics with higher component wear [173, 189]. It could therefore be hypothesised that the U/C implanted knee exhibiting higher passive V-V laxity *might* be more susceptible to greater wear damage from prolonged active gait; however, this relationship would be less apparent for the S/C implanted knee, where the correlations are weaker.

This study is intended only to illustrate the use *in principle* of statistical correlations to link the characteristic mechanics of different active and passive daily activities, and there are some important limitations to the models which should not be overlooked. The simplified restraint model and rigid-body contact formulations will result in reduced fidelity. Further the set of input variables studied is limited, with generalised assumptions made about distributions and correlations; these would need to be better modelled. The opportunity exists to develop this methodology with more complete and accurate anatomical models to explore (for specific designs) whether passive laxity can be a predictor for active gait mechanics. Note however that, the more complex the model, the more variable factors that must be accounted for, and hence the more trials that are needed to obtain a sound statistical model of the system. In consequence, a more accurate model could also require far more computational time than the 7-days needed for the 2×1000 trials in this proof-of-concept. The results of this exploratory study suggest this *may* be the case, for a limited sub-set of gait characteristics, and subject to design-dependency. It remains to be investigated

whether other activities (e.g. stair usage or deep flexion manoeuvres) would exhibit similar correlations to passive laxity motion, for a range of different flexion angles.

It would require further investigation with a wider range of variability factors and implant designs to determine how much of this correlation is universal (i.e. related directly to the variable input factors themselves, such as soft-tissue effects), and how much is controlled or constrained by the implant design (such as A-P motion for the S/C design in this study). However, it seems apparent that certain design features (e.g. lower constraint) can improve predictive power, and that some tests of laxity drawer loading (e.g. V-V) correlate better than others to gait characteristics. Measurement errors are known to be associated with assessments of passive laxity [190]; this would erode the strength of these correlations, so a more exhaustive study would also need to account for uncertainty in the laxity ranges.

This study illustrates conceptually another potential application of probabilistics, demonstrating the design-dependent correlations between passive laxity and active gait mechanics, and suggesting that for some gait characteristics these correlations potentially offer useful predictive power as a decision-support tool.

5.4. Probabilistic Wear Assessment: Multi-Design Comparison [7]

5.4.1. Background

Chapter Four demonstrated the valuable role *in-silico* wear assessments can play in pre-clinical analysis. These wear methods are easily incorporated to probabilistic studies, giving a more holistic perspective on the influences and variability of predicted wear.

Wear is known to be highly variable both *in-vitro* and *in-vivo*, but it is difficult to collect large enough data-sets clinically or experimentally to explore this variability (due to time and cost). By contrast, *in-silico* models can use large numbers of trials with low associated time & cost. Therefore using probabilistic computational methods it is possible to explore whether input variability (e.g. component malpositioning) can account for the high degree of wear variability observed.

[7] This section is adapted from the conference proceedings: "*In-silico* Predictions of TKR Robustness to Wear Variability: A Probabilistic Cross-Design Comparison". 2009, M.A. Strickland and M. Taylor. In: Transactions, ORS 55th Annual Meeting (Las Vegas, NV).

In this study, we will combine wear prediction with probabilistic methods to compare the predictions of multiple different wear algorithms, and to compare multiple TKR designs (to observe whether some are more robust to wear variability than others). The only existing published study for probabilistic wear of TKR [179] was for a single implant design, and failed to include any cross-s (using Archard wear only). This study will provide a more complete overview, including multiple designs and multiple wear theories.

5.4.2. Methods

Existing TKR designs were incorporated from CAD geometry or reverse-engineering, including 6 fixed-bearing (CR) and 2 rotating-platform designs. For each one, an *in-silico* simulation of an *in-vitro* wear test was used. Once again, the mechanical configuration was based upon the force-driven SKS [114] with a soft/hard spring combination as recommended by Haider et al [158] (7.24N/mm anterior & 33.8N/mm posterior). The inputs used were 'true' ISO-standard gait [23] (*not* experimental feedback data this time, as different designs were under test so no single feedback dataset would be appropriate).

Wear was evaluated using the standard algorithms discussed in Chapter Four, including variants without contact pressure terms. Distributions were plotted to form a PDF of wear rate for each design with each of the different wear algorithms. These PDFs could then be compared to evaluate the different TKR designs and wear algorithms. Due to the number of trials required for a multi-design probabilistic study, it was not feasible to use adaptive wear methods; therefore wear estimates were based on single-cycle analyses.

A 'streamlined' probabilistic analysis was used with higher levels of variability than previously, somewhat closer to *in-vivo* levels of variability. 7 Factors were included: six component malpositioning angles (with SD of 2°) and M-L load split (with SD of 12.5%). Having previously demonstrated the relative linearity of the SKS system, RSM-100 was selected instead of a more expensive MCST – this required 800 trials in total (approximate simulation time: 10 days).

5.4.3. Results

The choice of wear algorithm has a major influence on the degree of variability observed; see Figure 81. Algorithms excluding cross-shear (e.g. the Archard model) grossly *under-predict* wear variability. When CS is included, the SD of the resulting wear PDF is typically 3 to 5 times greater. Algorithms ignoring contact-pressure predict a moderate probability of wear levels below the 'neutral' (unperturbed) wear rate.

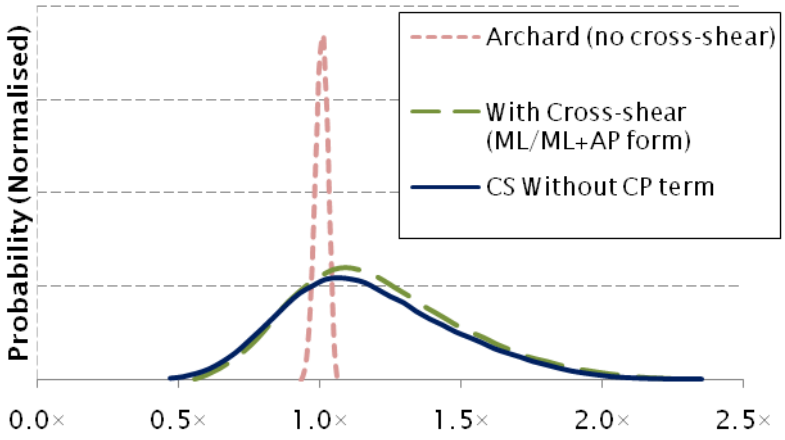

Figure 81: Typical PDFs for different wear algorithms, normalised relative to the deterministic wear rate (in this figure: for S/C fixed-bearing design).

The comparison between designs reveals that there are clearly design-specific differences (Figure 82). The deterministic (unperturbed) wear rate for designs varies, as has been reported in many *in-vitro* studies. However, this probabilistic study reveals that the *spread* of wear rates due to variability is also different.

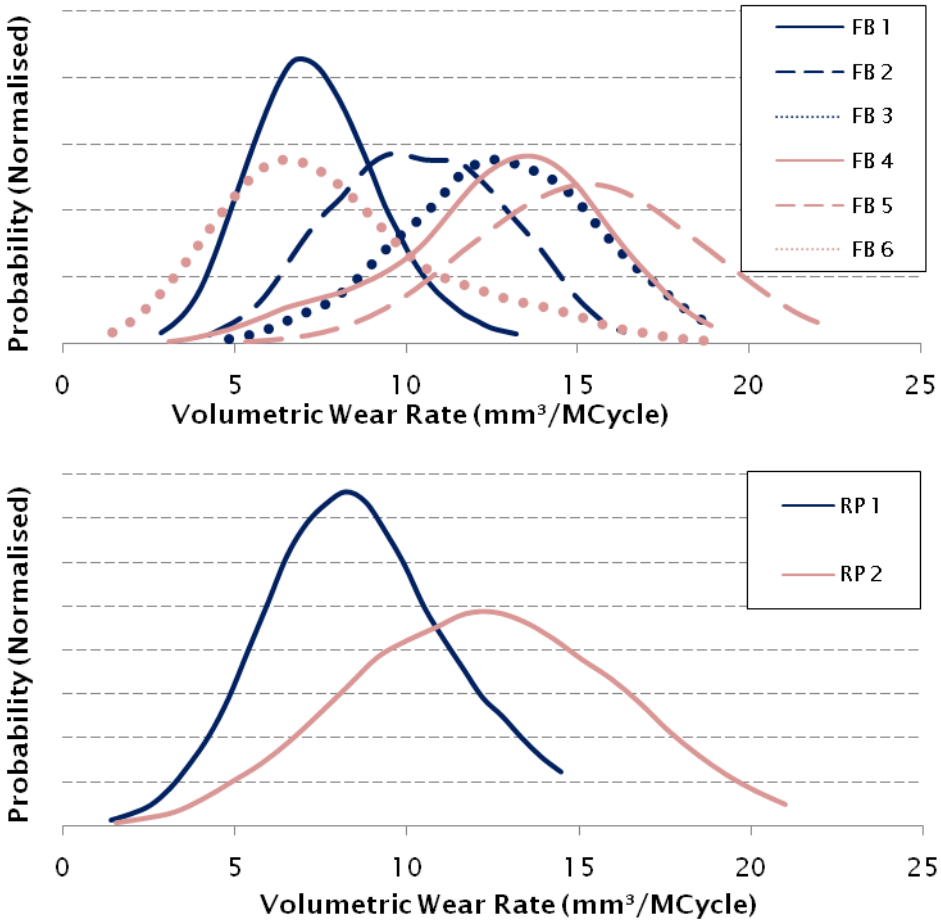

Figure 82: Comparison of PDFs for multiple designs, based on M-L/M-L+A-P wear model (6 × fixed-bearing, top; 2 × rotating-platform, below). Note: commercial TKR design brands have been anonymised.

Some designs appear more resilient to malpositioning and do not exhibit such a high spread of wear rates. For example, consider designs FB1 and FB6; looking only at the 'normal' (mean) wear rate, FB6 appears to marginally outperform FB1. However, looking at the full *distribution,* it is clear that FB6 is considerably more likely to have a high-wear rate in the event of malpositioning. Note that it is in no way possible to extract this information from individual deterministic models (either *in-vitro* or *in-silico*); a probabilistic approach is essential.

Note that wear rates of 3 or more times the neutral level have a significant (>5%) probability of occurrence for some of the designs studied; again, this is an important result which would be overlooked by a simple one-off deterministic analysis.

5.4.4. Discussion:

This probabilistic application of *in-silico* wear prediction once again reinforces the observation that wear models without CS do not predict the variations reported by *in-vitro* wear tests. Typically, experimental results show a large spread of results for any given design (even for small values of 'N'); the Archard formula does not predict this, CS must be included to capture this degree of variability.

Probabilistic studies provide a more challenging validation test for wear theories: a complete PDF of wear results is generated, providing a more complete data set to corroborate with (rather than an individual wear-rate value). If this computational probabilistic approach could be compared with a similar 'probabilistic' data set from *in-vitro* testing, it may help to identify the most accurate wear models under a wider range of test conditions.

The multi-design comparison reveals two very important observations:

- Firstly, wear rates can be much higher (greater than three times) the 'neutral' wear rates seen in correctly-aligned *in-vitro* simulators. This implies that those *in-vitro* results may also under-predict clinical *in-vivo* wear with malpositioning; further work would be needed to explore this.
- Secondly, wear distributions appear to be design-dependent. This implies that the TKR designer does have some ability to 'design-in' a degree of robustness to reduce the 'spread' of wear rates.

Once again, limitations to this study must be noted; the models used represent *in-vitro* (not *in-vivo*) conditions, so could not be expected to reproduce *in-vivo* wear variability (this would require a musculoskeletal modeling approach). As has been noted, the wear algorithms used are historical empirical models since

UHMWPE wear is not fully quantitatively understood. Better wear models could yield different distribution shapes. It is also important to note that the tests are abstract; this data is not being compared to any real set of experimental data. Without this corroboration, it is not possible to be certain that the factors included are the factors which would be relevant in a real experiment.

Nonetheless, it appears that alignment variability results in much higher top-end wear rates, and that this is a design-specific effect. These observations justify further investigation with better data, and better-corroborated models.

5.5. Conclusions

This chapter has demonstrated that the models & methods introduced in Chapter Four are well-suited to probabilistic analysis approaches. Published results in the literature have been corroborated, and the same modelling framework has been extended to demonstrate potential correlations between active and passive gait.

An important issue highlighted is that many of these models & methods have not been adequately compared to real-world experimental data. In many cases, published studies are 'validated' using a single 'feedback' dataset (e.g. kinematics for one isolated gait cycle), or a single wear-rate value. This does not give a complete proof of the model's performance. As such it is difficult to know with confidence what impact if any the limitations and assumptions of the model are having on the results. This is especially true for the probabilistic approaches in this chapter: to-date, *no* probabilistic computational study has been validated with true probabilistic experimental data (only isolated deterministic data has been used). The aim in subsequent chapters will be to apply these *in-silico* probabilistic methods to much richer data from specific real-world test platforms, to achieve a much higher level of corroboration between computational and experimental models. This requires test-platforms which are highly controllable, well-understood and well-documented. It is also valuable to have good collaborative links with the experimental test specialists, to obtain access to high-quality data, information on operating procedures and technical expertise.

As such, the studies in the final chapters describe models constructed and corroborated in much closer collaboration with experimental researchers. Chapter Six outlines the development of a new MBD-based model of the 'Kansas' Knee Simulator (KKS), whilst Chapter Seven demonstrates how the lessons learnt have been re-applied to the AMTI knee-wear simulator.

CHAPTER SIX – LOWER-LIMB MODELLING

Development of holistic lower limb models, based on the KKS

6.1. Background: Motivation for KKS Modelling

The strengths and limitations of different *in-vitro* knee simulator configurations were discussed in detail in Chapter Two. The full lower-limb simulator configuration is potentially more powerful (since the scope of the model is greater), but also more complex. Whereas tibiofemoral knee simulators are in commercial production, full limb simulators remain bespoke one-off investigation platforms, in the domain of academia rather than industry. An advantage of this is that, by close collaboration with the experimental KKS research team, much more full and detailed specifications and data are available for this platform than for the 'black box' commercial systems. Therefore, a more specific, targeted corroboration is possible with this platform than with the previous work – this can then be used as the basis for probabilistic analyses.

6.1.1. The KKS: Technical Description

The Purdue/Kansas knee simulator is one of the most well-established & technically advanced knee simulators available, and has been widely used for peer-reviewed research and industrial TKR design & development. Originally conceived as a next-generation knee wear simulator, the current KKS design is now used to support research on knee kinematics, loading, laxity & stability (gravimetric wear assessment is not supported on the current rig for *in-vitro* testing, but this can be estimated using coupled *in-silico* methods). The KKS is a highly versatile platform, able to operate using artificial implant test-pieces, implanted cadavers, or natural-knee cadavers, with the capability to track and record force-feedback (via load cells) and kinematics (via an 'Optotrak' motion tracking system) in real-time during testing. For further reference, the capabilities of the platform are described in more detail in a series of papers by Zachman, Hillberry, Rullkoetter & Maletsky [120, 121, 191-196].

Alongside the mechanics of the rig, the control system is a very important feature of any lower-limb model. For the tibiofemoral knee wear simulators discussed previously, the system is essentially stable (the tibial orientation is fixed, and compressive loads will tend to stabilise the femur in the conformal condyles) so accurate control is less critical. By contrast, the natural knee (and by extension any *in-vitro* lower limb simulator) is unstable. When the knee is in flexion, vertical loads at the hip & ankle will tend to increase flexion – this in turn increases the moment arm of the vertical loads, and so increases the flexion moment, creating

a ‘positive feedback’ loop. In the natural knee, stability is achieved through the complex holistic system-level operation of the entire neuro-musculoskeletal system (i.e. soft tissue restraint, sensory feedback, antagonistic muscle action, etc). It is therefore to be expected that a sophisticated control system is also required for the KKS.

The KKS control system is a five-channel PID controller with full cross-compensation, and the ability to operate all five axes in force- or displacement-control, based on a number of uni-axial load cells and linear & rotary displacement transducers mounted on the simulator. Consequently, control of the KKS is challenging: a range of activities are simulated, requiring different loading profiles, and the same control scheme may not be appropriate under all conditions. Inappropriate control commands could potentially damage cadaveric tissue under test, or even the rig itself. Consequently, there is a strong incentive for augmenting the *in-vitro* test rig with *in-silico* modelling, to devise and test profiles before they are used on the simulator rig itself as a ‘risk reduction’ exercise. This would also allow more unconventional profiles or control schemes to be investigated without the additional sensors, actuators, or reconfiguration time that would be required on the *in-vitro* rig.

For these reasons, the KKS has previously been modelled using MBD, by Guess et al [109, 123] (Figure 83). For more technical detail on this original model, the reader is referred to the doctoral thesis of Guess [197]. This earlier model (dating from 2003 [198]) was based on an older configuration of the KKS, and whilst it conceptually demonstrated MBD modelling of the KKS and its controller, there were several key limitations:

- The model itself used *Hertzian* contact only for the patellofemoral and tibiofemoral articulations. This simplified contact model is fast, but assumes spherical contact surfaces, and so is not very accurate. It also cannot provide any information about contact pressure distribution at the contact surface – (this information is needed for wear prediction).
- The model did not feature any form of ‘wrapping’ for the quadriceps; this meant that the *in-silico* and *in-vitro* results diverged after around ~80˚ knee flexion, when the Kevlar strap representing the quadriceps tendon (QT) begins to wrap across the patellar groove on the femoral component. Without this wrapping, the model is effectively limited to shallow flexion activities. There is particular interest in ‘deep’ flexion performance for many new knee designs, so it would be highly desirable for the model to reflect the flexion range of the rig. Theoretically, this is approximately 0˚ to 135˚; beyond this mechanical

interference between the femur and tibia restricts any further flexion. In practice, a range of 120˚ is closer to the realistically achievable limit.

- The controller used was based on the internal ADAMS controls toolkit, and was limited to a single axis under feedback control.
- The model did not account for dynamic joint friction or actuator damping.

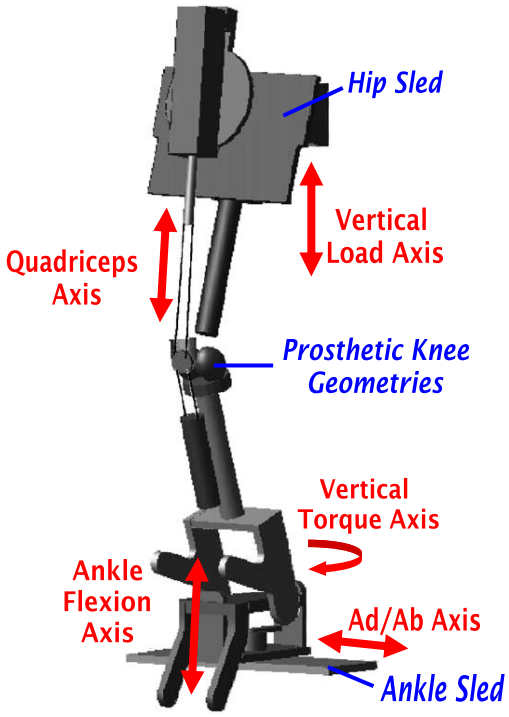

Figure 83: Original KKS ADAMS MBD-based model (adapted from [123]).

In light of this, the need was identified for a new computational model to overcome these limitations and so provide a more robust and complete comparison to the KKS rig. In particular, it was desired that the new model:

- Better reflect the re-designed KKS geometry & inertia
- Incorporate the new tri-axial load cell and inclusion of 'collateral ligaments'
- Allow the model to be easily re-positioned & parameterised
- Include 'deep flexion' capability (i.e. quadriceps wrapping, and modelling of any other relevant mechanical interference)
- Achieve greater accuracy corroborating with the KKS.

The development of the new model involved collaboration between Kansas University (KU), the University of Missouri (UMKC) and the University of Southampton. Each research group had specific requirements for the new model; KU required a more robust *in-silico* model for generating deep flexion testing profiles to use on the KKS rig; the UMKC & Southampton research groups required a more extensive & capable baseline model for purely *in-silico* studies. The following section describes modelling undertaken as part of this work-plan by the author whilst on secondment with KU.

6.2. Model Development

6.2.1. Geometry

Extant CAD geometry from Pro/Engineer (Parametric Technology Corporation) was available for the new (redesigned) KKS configuration. This geometry was used as the basis for the new ADAMS model. Inertial properties were derived for each part based on assigning known material densities to the solid volume. Where necessary (e.g. assemblies comprised of multiple materials) certain parts were weighed directly to validate this estimated inertia. Note that although the KKS can accommodate implanted 'cadaver' samples, the initial ADAMS model is based only on the artificial jig used for direct component testing. The implant described in the following tests is a standard-size, fixed-bearing, posterior-stabilised (PS) variant of the cruciate-retaining implant used in Chapters Four and Five.

The complete simulator model includes over 30 parts, compared to the 14 in the original model by Guess et al. This reflects the additional components required to capture all inertial and dynamic effects; for example the new model explicitly includes the actuators, with the moving actuator rods and linkages individually modelled, in order to include the additional damping effects, and friction at the linkages. Figure 84 compares the two models side-by-side, indicating the main areas where the scope of the model has been revised.

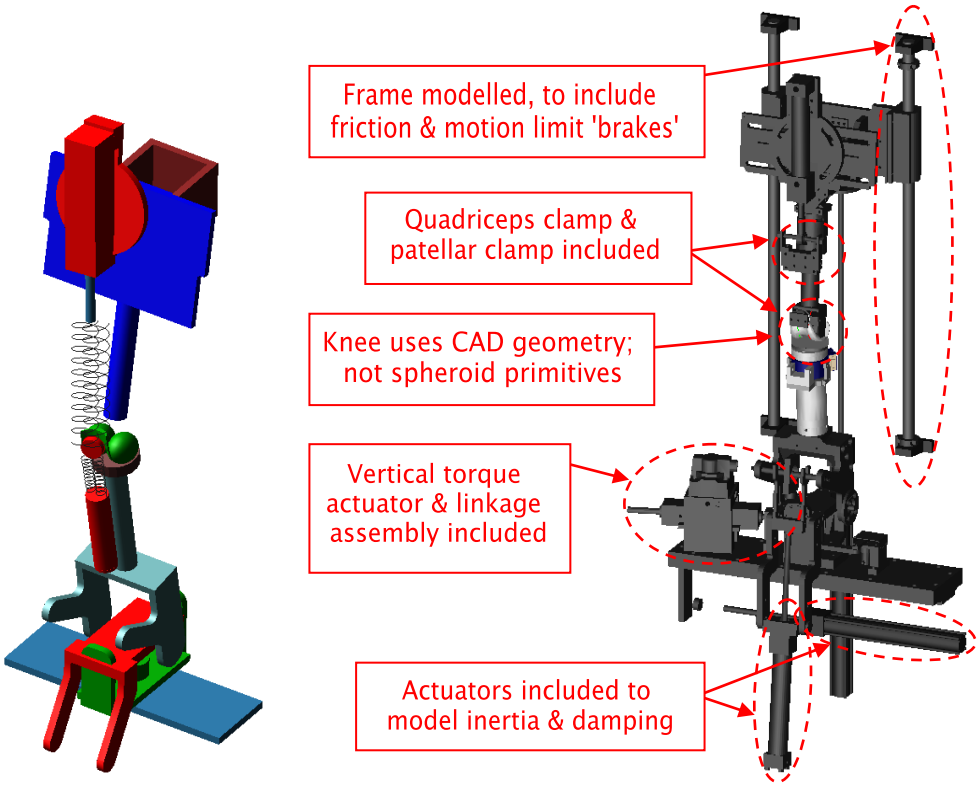

Figure 84: Comparing the original (left) and revised (right) KKS models.

The parts were assembled in ADAMS using appropriate 'joints' to limit the degrees of freedom. Initially, components were coupled via 'rigid' hinged or sliding linkages. However, it was subsequently discovered that 'soft' bushing elements (with a finite 'stiffness' on all axes) were far better suited to simulate the pliancy & damping effects within the simulator. The model was revised to use physically representative joint types (which are more intuitive for design work). This does considerably increase the number of degrees of freedom in the model (hence increasing solution time). The original model featured ~15 DOF, whereas the new model has almost 60. This increase in complexity is not desirable in terms of computational cost, but is *necessary* to correctly model the system dynamics (something overlooked by most *in-silico* models).

With the underlying mechanical assembly complete, it was possible to begin including some of the new features. These were introduced sequentially, so that each new extension could be individually tested and 'debugged' in turn.

6.2.2. Instrumented Tibial Assembly

The KKS was recently upgraded to include a tri-axial load-cell for measuring tibial loads & moments. The load-cell mounts directly below the tibial insert, to measure loads as close to the proximal tibia as is possible (similar in ethos to *in-vivo* instrumented tibial inserts). This requires a modified tibial assembly to accommodate the large and quite heavy (~1kg) load-cell. The new tibia is slightly longer than the original, and is designed to mount the tibial insert with a posterior slope of 5˚. The differences are illustrated in Figure 85.

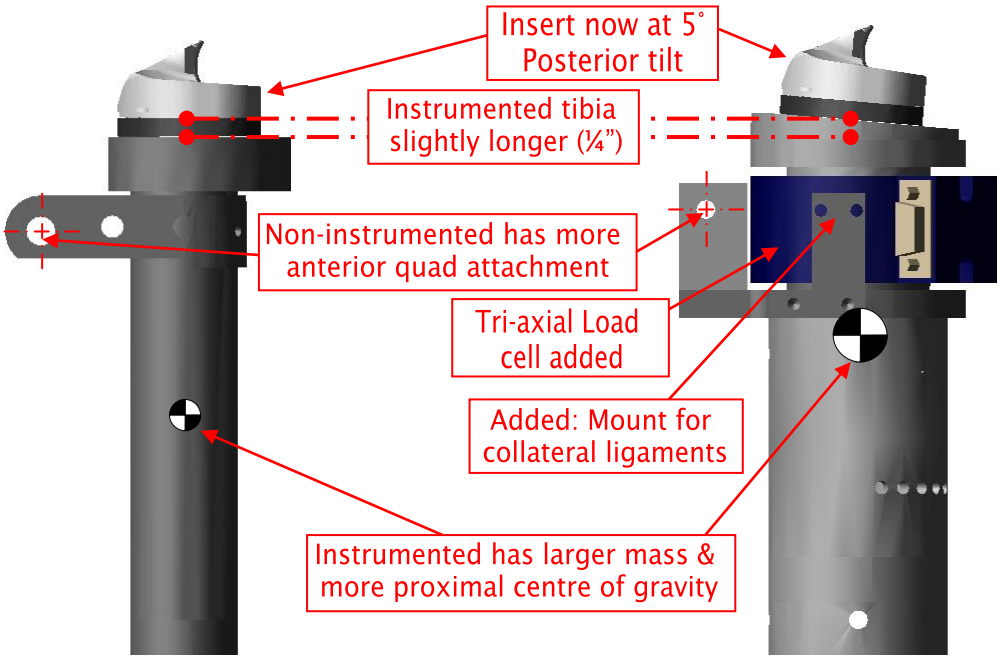

Figure 85: Comparison of non-instrumented (left) & instrumented (right) tibial assemblies.

This new part was modelled in ADAMS based on parasolid geometry for the old assembly, and also the new manufactured parts. Rather than removing the original non-instrumented tibia, the model was configured with a custom script to activate/deactivate and appropriately reposition components, such that either the instrumented or non-instrumented variants could be selected using a single command-line instruction. The two different variants have different inertia and a different centre-of-mass; also the attachment point for the quadriceps is altered, and the tibial insert is sloped. All these factors can influence the knee kinematics, so the appropriate selection must be made when corroborating with *in-vitro* data (depending which part was used to collect the experimental data).

6.2.3. Quadriceps 'Wrapping'

On the KKS, the quadriceps load is applied by a servo-hydraulic actuator mounted on the proximal 'femur', anterior to the hip joint. The force is transmitted to the proximal tibia via a Kevlar strap, intended to represent both the QT & PL. The patella is mounted on the strap via a specially designed clamp (such that the initial I-S location of the patella in extension can be freely adjusted). In the original model by Guess et al, this was represented by two pairs of spring-damper elements, providing a line-of-sight restraint force between the quad actuator and the patella on the proximal side, and the patella and tibial 'tuberosity' (a mounting point on the proximal anterior tibia) on the distal side. This line-of-sight spring model did not detect interference with the femoral component, so as the knee flexion increased, the proximal springs in particular would penetrate the femoral component without being deflected; this reduced the moment arm of the quad actuator, such that it could not correctly resist the flexion moment. This in turn would lead to further knee flexion, further exacerbating the limitation of the non-wrapping model. This 'positive feedback' effect made the model unable to operate beyond ~80˚.

Therefore, a more realistic model able to account for wrapping was required. Essentially, this introduces an additional set of contacts between the strap and the femoral component. This obviously has adverse implications for solution time. However it is particularly problematic, because whereas the tibial and patellar contacts are generally close to perpendicular, the strap wrapping contact is acting antagonistically, directly *against* the tibiofemoral contact (see Figure 86). This makes numerical iterative convergence for the two contacts much more challenging, and can greatly increase the required solve time.

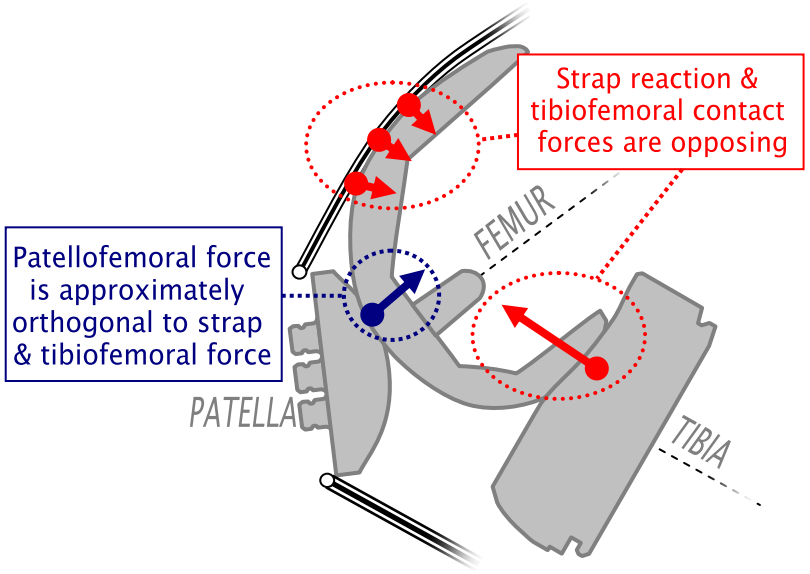

Figure 86: Antagonistic action between the tibiofemoral joint and strap wrapping contact.

Because of the envisaged impact on solution times, two different methods for strap wrapping were proposed, and implemented for a comparison study. The methods are outlined below:

'Fast' Point-on-Curve Method

This method does not use the femoral component geometry directly; instead, the sagittal-plane geometry of the femoral component is traced with a simple vector-path. Contact between the femoral surface and the strap is then based on this vector-path representation. For the strap, several point-nodes are then embedded along the region where wrapping occurs, to provide contact reaction points (i.e. the strap becomes a series of chained spring elements, with point-contact occurring only at the linkage points). The advantage of this approach is that no 3-D solid geometry calculation is involved, thus greatly reducing solve-time. The obvious disadvantage is accuracy; the model is only accurate as long as the patellar is tracking 'normally', (i.e. close to the sagittal plane). In addition, a specific path must be created for each different femoral implant design under test.

Sphere-on-Solid Method

Using this method, a series of 3-D ellipsoid primitives (spheres) are embedded into the strap. A contact is defined between these spheres and the actual CAD geometry of the femoral resurfacing component. This means that the contact *will* conform correctly to the articular surface of the femoral patellar groove. However, the spheres still only provide a *discretised* contact, rather than the continuous contact on the physical strap – hence this method is still only an approximation.

The use of 3-D solid contact greatly increases the computational complexity, which coupled with the antagonistic action between the tibiofemoral reaction force & the strap wrapping, can drastically increase solve time. Therefore, this approach is not suited for 'fast' modelling, (e.g. during development or debugging).

As with the switching script for the tibia, both methods are embedded in the model, with a custom macro to allow the user to toggle quickly between the two alterative configurations. In either case, a decision must be made as to how many distinct sections the strap should be discretised into – this is a standard performance/accuracy trade-off, and sensitivity studies by UMKC Demonstrated that even a small number of discrete wrapping contact points (2 or 3) gives acceptable accuracy, provided that the points are appropriately located and spaced along the strap (wrapping tends to occur only at the distal end of the QT, for the flexion angles of interest – see Figure 87).

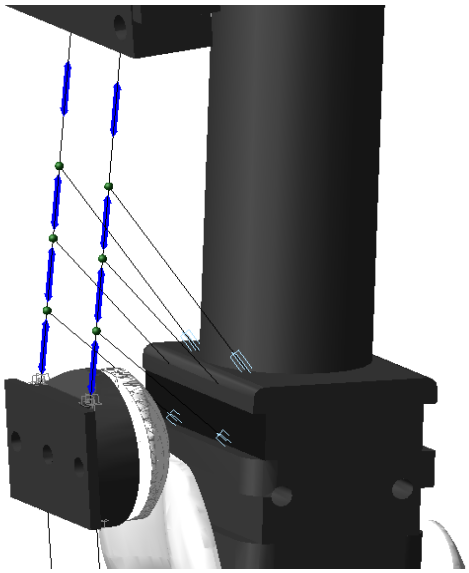

Figure 87: Strap wrapping. Three pairs of discrete contacts emulate continuous wrapping, in the contact region of interest between the QT & proximal patellar groove on the femoral component.

6.2.4. Collateral Ligaments

The original KKS did not include *any* ligament restraint at the knee – it was designed for principally sagittal-plane loading (such that collateral ligaments would have limited effect), and for use with PCL-sacrificing TKR designs (such that neither cruciate ligament needed to be included). In order to provide more physiological restraint, and to allow greater out-of-plane loads & motions, it was decided to incorporate springs onto the KKS rig to represent the collateral ligaments. For simplicity in the first iteration of implementing collaterals, the two ligaments were modelled as single-line of action force ('SFORCE') elements in

ADAMS, with zero stiffness under compression, and a constant stiffness of 30N/mm under tension. This stiffness was dictated by the availability of suitable small springs for the physical rig; it is recognised that this is less restraint than the values reported for experimental tests on the collateral ligaments [5]. (The intention is to revise these springs on the KKS rig to specialist fittings of higher stiffness in the future).

Since the springs have a single line of action, they cannot reflect the true physiological behaviour of the multi-bundle collateral ligaments. Instead, the *in-silico* model was used to determine insert & origin locations such that the ligaments provided additional restraint through the operational flexion range into deep-flexion. Based on these recommendations, the collateral springs were then installed, mounted on brackets attached to the femoral and tibial assemblies. The ADAMS model was updated accordingly to reflect these changes. Note that the tibial insertion is *distal* to the load cell, so that the load cell reports the compressive load experienced at the tibiofemoral joint *inclusive* of ligament forces. A complicating factor is that the KKS can be used for either left- or right-knee components; therefore the asymmetry present in the natural LCL & MCL could not be modelled; instead, a 'mean' line of action was chosen for the ligaments, with two alternative attachment points included on both the left and right brackets, to accommodate both positions on both sides (Figure 88).

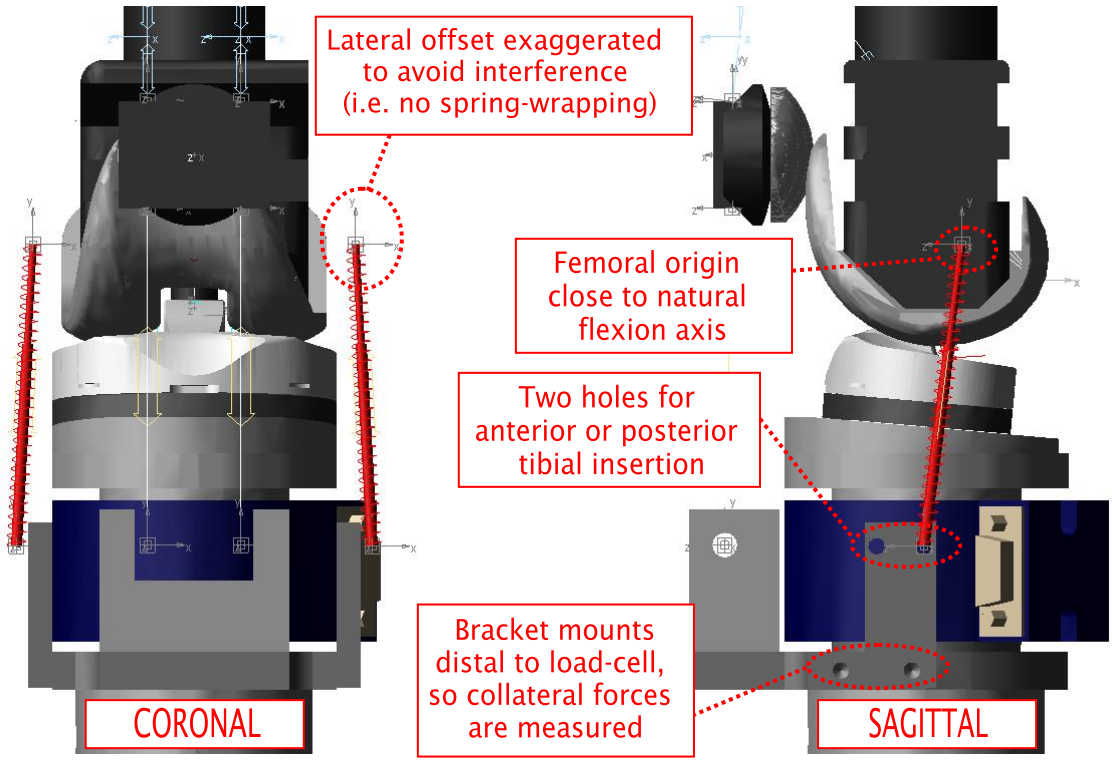

Figure 88: Proposed attachments for the collateral 'ligaments': coronal view (left) & sagittal view (right), showing posterior tibial attachment point for collateral springs.

6.2.5. Contact Switching

In the model developed by Guess et al, the contact method used was Hertzian contact only, and contact pressures were not considered. In the new model, contact pressure on the tibial and patellar UHMWPE components may be required as output metrics in stochastic studies. Therefore, the surfaces of both the tibial and patellar inserts were discretised, according to the method described in Chapter Four. However, this increased complexity inevitably results in slower computational solve-times. As an alternative, a contact 'switching' macro script was devised. Similar to the switching macro for the strap wrapping, this allows the user to quickly toggle between a simplified model based on Hertzian contact (for 'fast' purely kinematic analyses or profile generation) and a more complex model based on the discretised contact, for obtaining contact pressure or wear results when needed. The difference in performance between the two alternatives is considerable, and may be an important factor in deciding the best approach for further stochastic studies.

Hertzian contact is an analytic approximation, specific to the simple case of linear-elastic bodies with simple geometry. Under a given load, a penetration is predicted, based on the separation between the two surfaces, and the material properties (modulus and Poisson ratio) – see Figure 89. This produces an elliptical load distribution, from which the overall contact area (and hence peak pressure) can be estimated. To apply this method to TKR, the surface of the femoral and tibial condyles must be approximated as elliptical spheroids of appropriate radius. Appropriate material properties are applied (femoral: CoCr; $E_F \approx 2.3\times10^{11}$Pa, $v_F \approx 0.3$. tibial: UHMWPE; $E_T \approx 1$GPa, $v_T \approx 0.45$), and used to determine parameters for the contact, using the standard Hertzian contact formulae. First, an equivalent 'contact' radius, R_C, is evaluated as the reciprocal sum of the tibial radius, R_T, and femoral radius, R_F:

$$\frac{1}{R_C} = \frac{1}{R_T} + \frac{1}{R_F} \qquad (7)$$

Next, an equivalent 'contact' stiffness, K_C, is evaluated, using the modulus of the tibial & femoral materials (E_T & E_F) and the Poisson ratio of the materials (v_T & v_F):

$$K_C = \frac{4}{3}\left(\frac{1-{v_T}^2}{E_T} + \frac{1-{v_F}^2}{E_F}\right)^{-1} \qquad (8)$$

Finally, the Hertzian contact equation can be re-expressed in terms of an IMPACT-style force-interpenetration relationship (see Chapter Four, Equation 2) for use in

ADAMS (where h_C, the contact separation distance, is equivalent to g, the interpenetration depth). This gives the IMPACT stiffness co-efficient k and exponent e for the normal force F_N:

$$F_N = K_C R_C^{0.5} h_C^{1.5} \Leftrightarrow F_N = k \times g^e \quad (9)$$

$$\therefore k = K_C R_C^{0.5} \qquad e = 1.5 \quad (10)$$

Note that in the Hertzian model, the value of e is always 1.5 by definition; however the value of k will vary depending on the material properties and implant geometry, and so must be evaluated on a case-by-case basis. (Note: this is a whole-body aggregate contact equation, so these values of k and e cannot be compared to the 'discretised' contact methods described in Chapter Four and Appendix C).

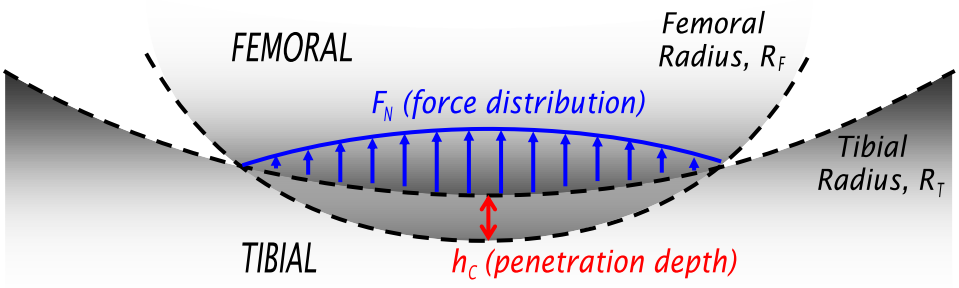

Figure 89: Hertzian contact illustration.

6.2.6. Standardising Polarity within the KKS Model

An important lesson learnt during the development of the model and *in-vitro* corroboration testing is the ease with which errors may be introduced in the polarity of inputs and measures. In the case of primary motions (e.g. flexion), these errors would be obvious and easily detected. However, in the case of secondary motions (e.g. small amounts of I-E rotation or Ad-Ab rotation) these errors are less obvious and can confound subsequent corroboration testing. To mitigate against this, the polarity of the different input/output forces & displacements was standardised, based on the current polarities used on the KKS rig. The subsequent figures illustrate the reference (positive) direction for the polarity of the different *linear* translations & forces (Figure 90) and *angular* rotations & moments (Figure 91). Whilst this may sound an obvious issue, it is nonetheless a frequently-overlooked source of mistakes in many *in-silico* & *in-vitro* tests.

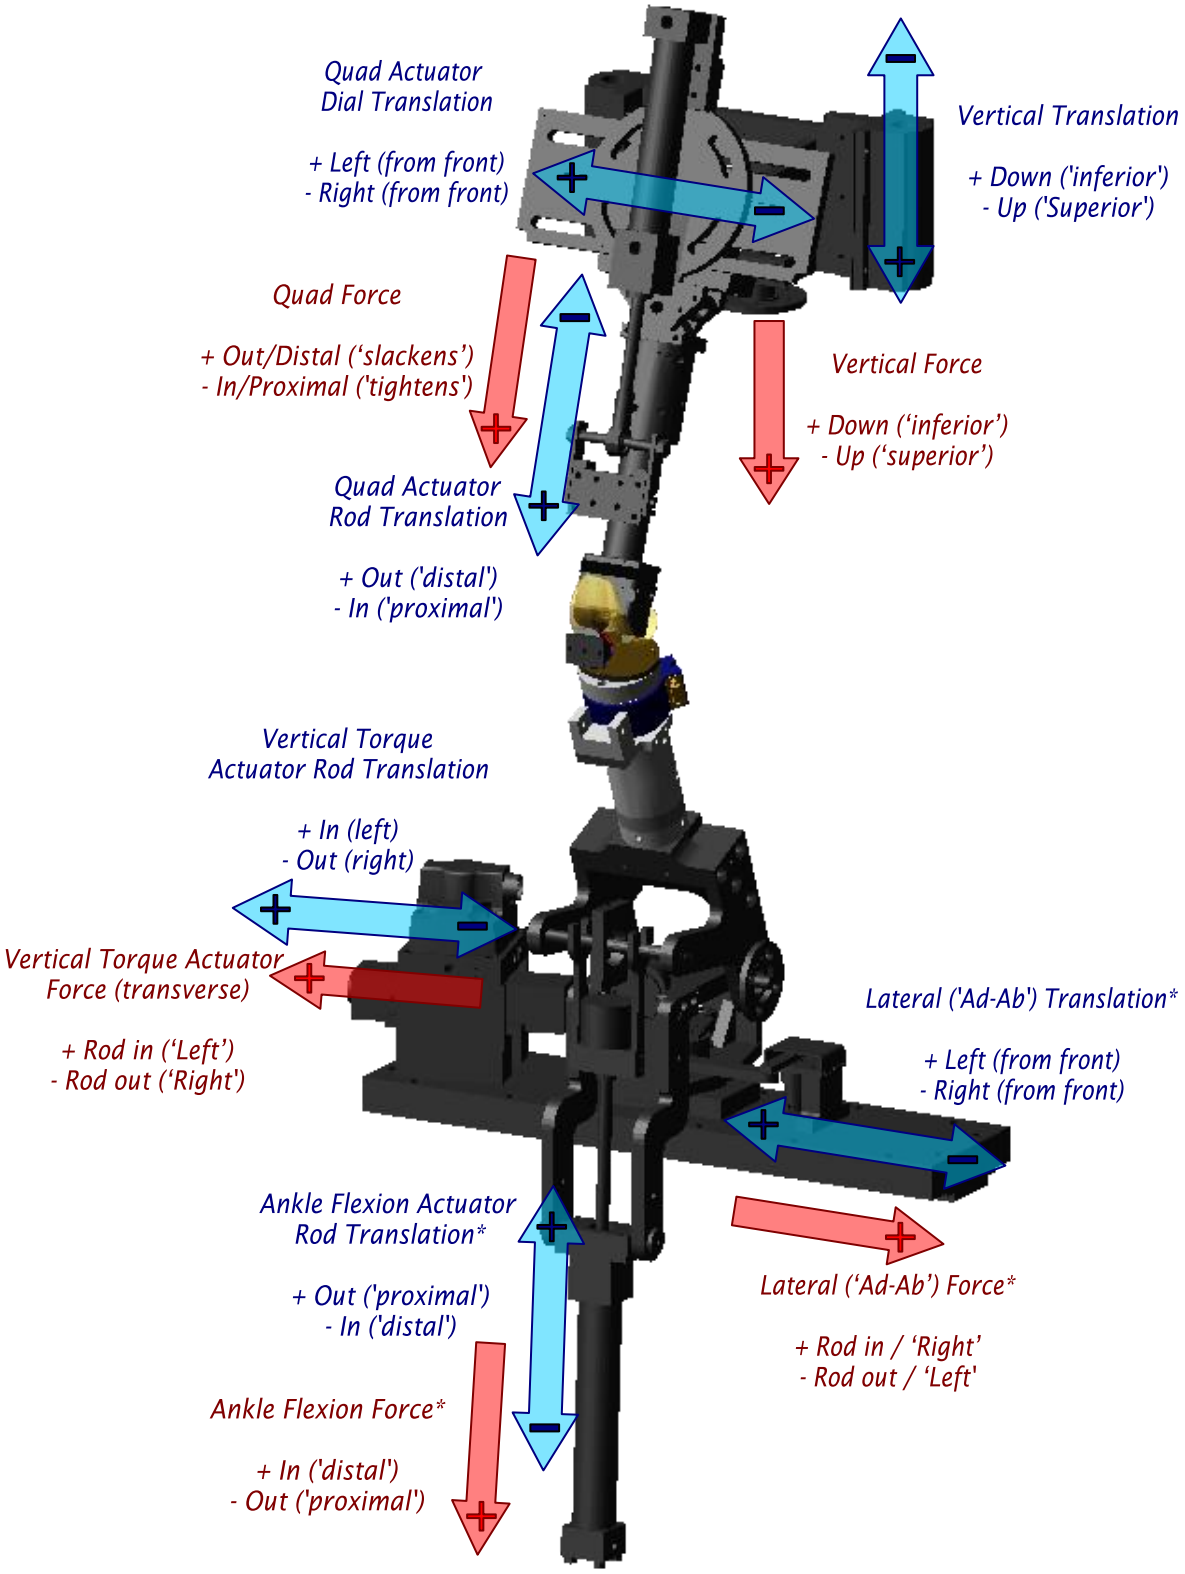

Figure 90: Polarity of *translational* forces & displacements on the ADAMS model (Note: for axes denoted by *, polarities of force and displacement are inverted).

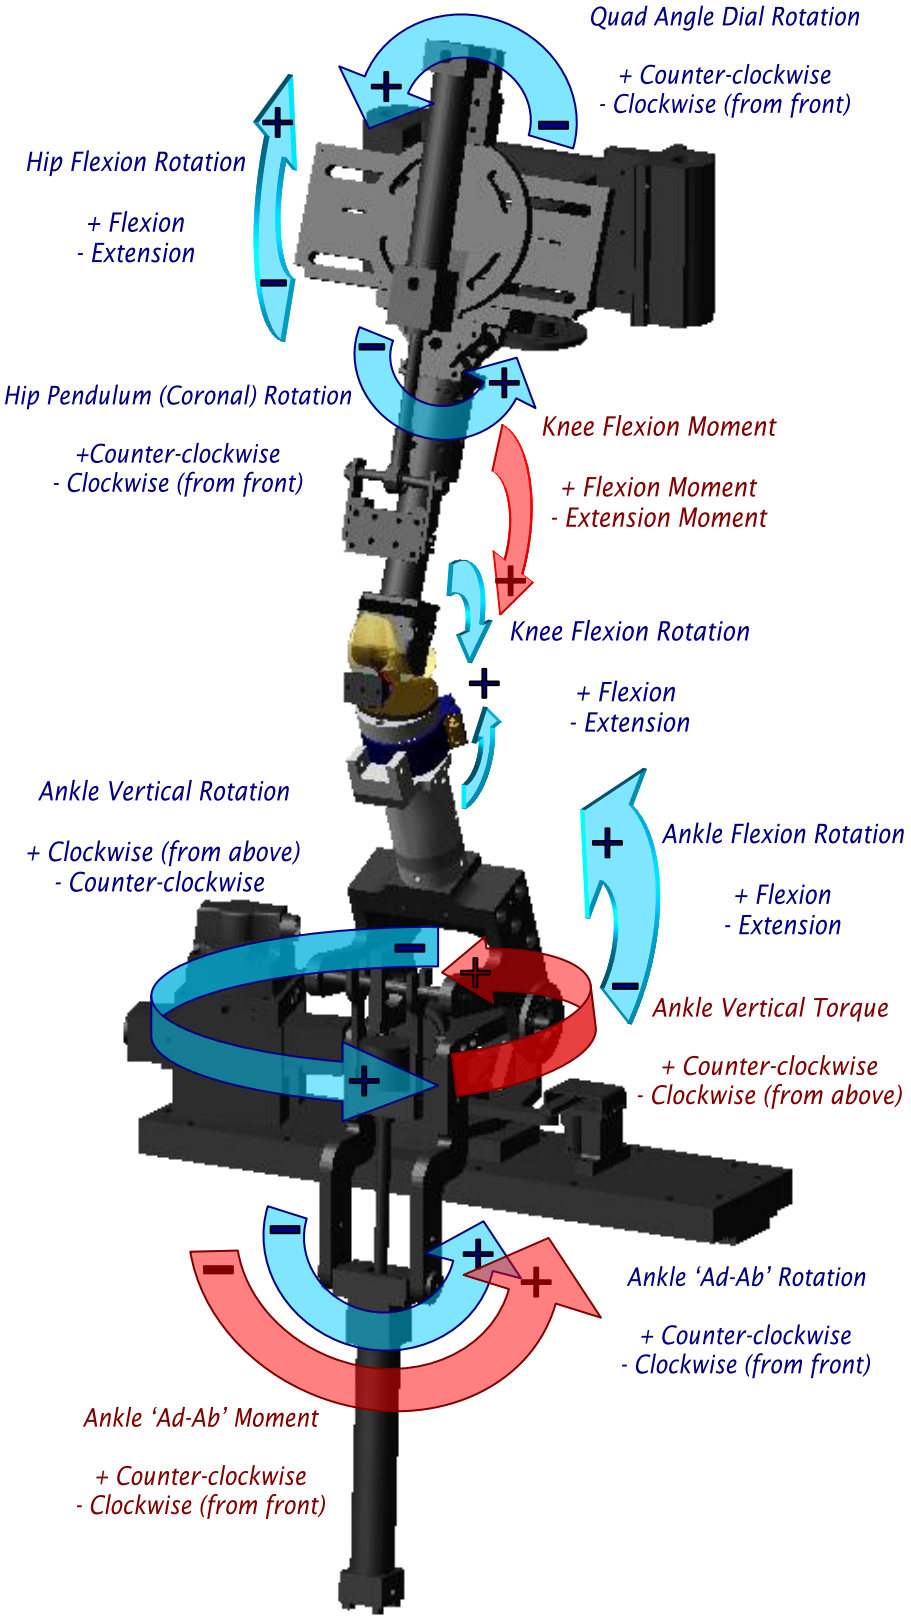

Figure 91: Polarity of *rotational* torques & angles on the ADAMS model.

6.2.7. Model Parameterisation

The model has been extensively parameterised; almost every DOF can be set by the use of numeric 'design variables' (30 factors in total). Considerable effort has been taken to ensure that the parameters are correctly inter-related; e.g. adjusting the quadriceps 'q-angle' will also appropriately shift the M-L dial position, and setting the initial vertical (I-E) rotation will also reposition the various linkages and actuator heads associated with vertical rotation. Table 13 lists the factors which have been parameterised within the model. The advantage of performing this parameterisation is that these factors could now be used as inputs for subsequent probabilistic studies (obviously, further specific factors for material properties could be added subsequently). Note however that the mechanical configuration of the KKS imposes certain additional limitations on the variables; for example although the 'shank' and 'thigh' length may be varied in the computational model, in reality they must be closely matched, or else the knee will not articulate correctly (the KKS does not permit A-P position adjustment at the hip or ankle; the design intent is that both hip & ankle flexion angles should be approximately equal in order to mate correctly at the knee, and the two segment lengths therefore should also be approximately equal).

Component malpositioning is based on the Grood & Suntay system [17]. Although this system is intended for joint motions, it can also be used for the static malpositioning of a component relative to the bone. This results in 3 independent sets of terms, describing the position of the femoral component relative to the femur, the tibial insert relative to the tibia, and the patellar button relative to the patella. Each has 6 potential degrees of freedom (3 translations: S1, S2 & S3; 3 rotations: E1, E2 & E3), resulting in 18 malpositioning parameters (Figure 92).

Values Fixed Throughout Simulation	
Tibial (shank) length adjustment	Femoral (thigh) length adjustment
Fixed femoral I-E rotation offset	Fixed femoral V-V rotation offset
Fixed quad coronal plane Q-angle	Fixed quad actuator lateral offset
Initial Values at Start of Simulation	
Initial M-L position of ankle sled	Initial I-S position of hip sled
Initial ankle flexion angle	Initial hip flexion angle
Initial ankle vertical rotation angle	Initial ankle Ad/Ab angle
Implant Malpositioning	
Femoral: 6×Grood & Suntay cylindrical axis system positioning variables	
Tibial Insert: 6×Grood & Suntay cylindrical axis system positioning variables	
Patella: 6×Grood & Suntay cylindrical axis system positioning variables	

Table 13: Factors parameterised on the new KKS model.

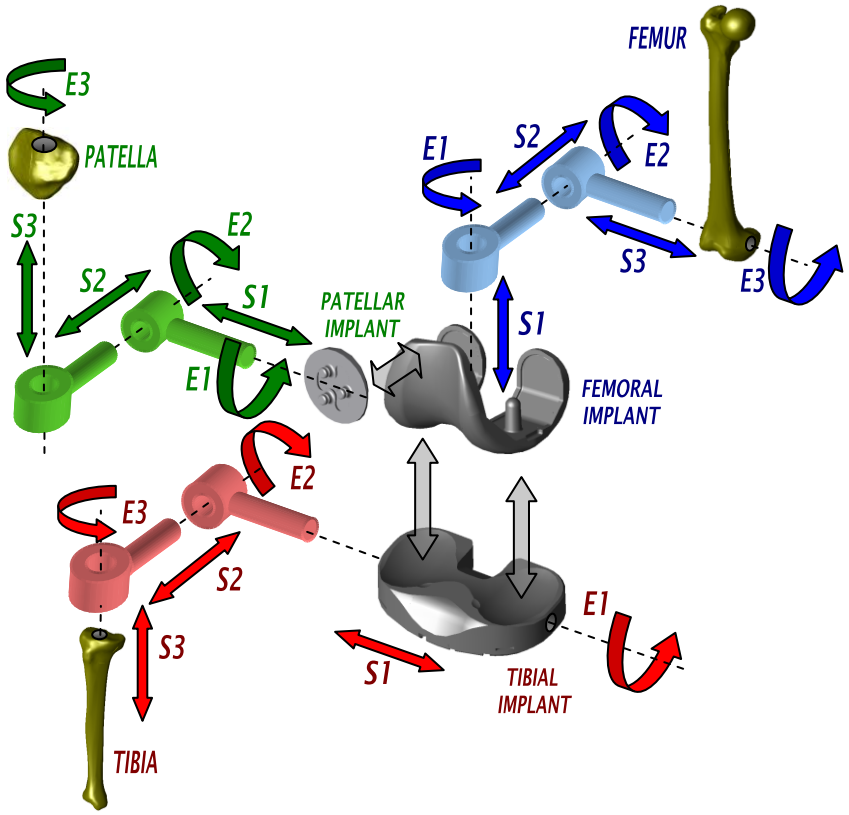

Figure 92: Grood & Suntay malpositioning for components on the KKS model.

6.2.8. Configuring Dynamic Properties of the KKS Model

Early testing demonstrated that the computational rig was less stable than the physical rig, due to the lack of any dynamic resistance effects in the model. On the *in-vitro* KKS rig there were two important classes of behaviour overlooked in the initial modelling. Firstly, the assumption of 'rigid' fittings and joints was not precisely correct; the parts have limited stiffness and so flex slightly, and the joints also exhibit a degree of pliancy. These effects essentially 'soften' the system. Secondly, damping & friction effects were found to be substantial; these tended to attenuate high-frequency motions, and limit the rate of movement. This combination of 'softening' and 'dampening' intrinsically improves the stability of the system – with the (unintentional) benefit of making system control easier: the system has less of the high-frequency characteristics which can lead to oscillatory behaviour and instability under PID control.

It was apparent that the *in-silico* model could not mimic the behaviour of the KKS unless these effects were included. To reflect these effects, various spring-damper elements (e.g. on the vertical load axis) were deliberately added to the physical simulator to mimic the *in-vitro* conditions and increase system stability. However, the number of potential sources of damping, pliancy and friction meant that it would not be possible to experimentally determine every term without stripping the rig, individually performing a sweep of tests on every component, and then testing the assemblies at each joint. This was beyond the scope of the

time available for modelling & development work during the secondment at Kansas University. Instead, some coefficients could be estimated based on engineering data sheets (e.g. manufacturer's friction coefficients for the branded roller-bearings & servo actuators). Selected other components or assemblies (those which could be readily removed & tested in isolation, or those which were considered particularly influential) *were* removed for testing. This testing included dead-weight loading for some spring stiffnesses, MTS tensile testing for the Kevlar strap stiffness & damping terms (e.g. see Figure 93), and testing of the upper hip assembly.

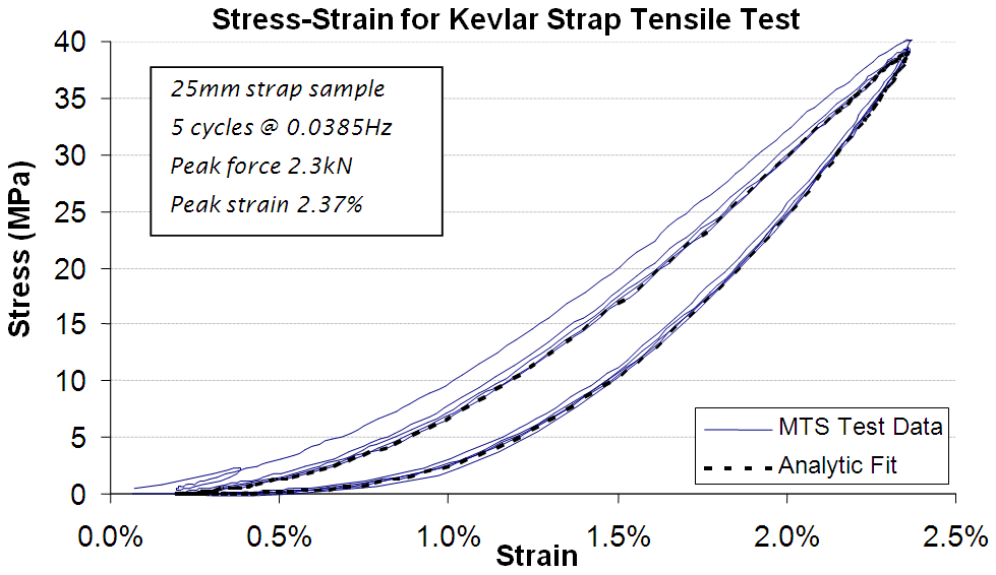

Figure 93: MTS tensile testing to determine analytic fit for strap stiffness/damping.

To capture any outstanding terms which had been neglected, the model included global damping terms which could be 'tuned' to experimental data. To do this, a series of triangular ramp-up/ramp-down waveforms were applied to each axis and the system response was measured. This comparison revealed the considerable influence of dynamic resistive terms. For example, see Figure 94; in this test, the hip was driven through a triangle-wave for flexion angle, between 5˚ - 30˚ at a lower rate, and between 5˚ - 17.5˚ at a higher rate. If there were no dynamic losses (friction or damping), the resulting plot of quad angle versus flexion angle would be a single curve (with no difference between quad force for flexing & extending). However, there *is* dynamic resistance, so the resulting plot exhibits '*dynamic hysteresis'*. (Note: the term 'hysteresis' is used here in its etymologically-correct sense of referring to *any* system which has path-dependant behaviours; not just the elastic hysteresis effects most commonly associated with the term. For *dynamic* hysteresis, the behaviour depends not only on the path taken, but the rate of progress along that path).

This dynamic resistance increases as the flexion rate increases (i.e. the hysteresis effect is greater for the 5˚ – 17.5˚ sweep). The control-system tracking for this test was well-tuned, so it may be deduced that the degree of hysteresis for the quadriceps positional tracking is almost entirely due to mechanical resistances. It is then possible to empirically or analytically fit terms to these results. Note: similar hysteresis plots to those observed have been reported by other testers for friction and damping effects on other *in-vitro* platforms [199, 200].

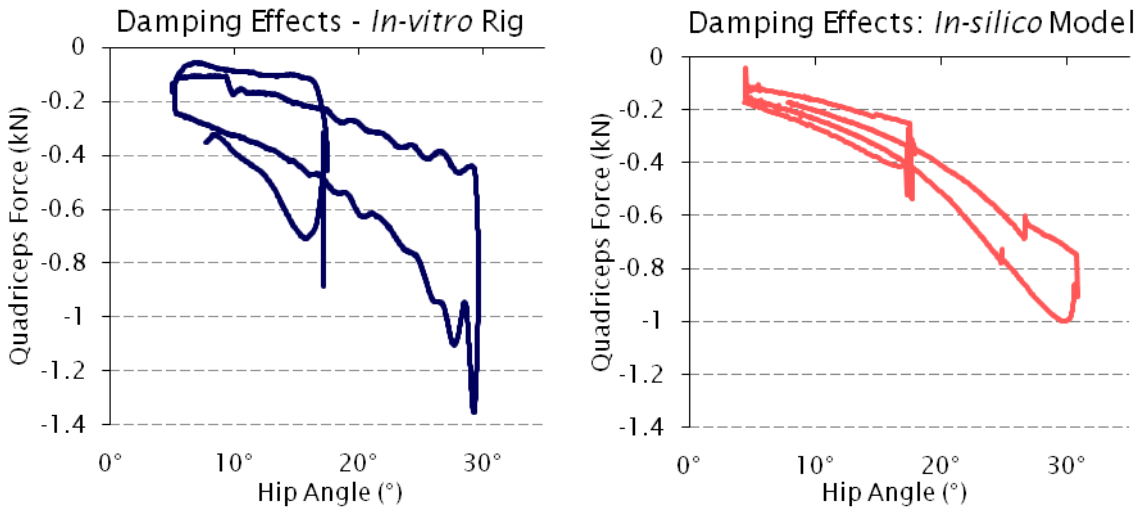

Figure 94: Damping comparison tests. Zero dynamic resistance would result in no hysteresis. Both plots show more hysteresis for the shorter path (5˚-17.5˚) with the higher rate-of-change, but note the greater hysteresis *in-vitro* (left).

6.2.9. Controller Development

The KKS consists of both a mechanical system (the rig itself) but also importantly a *control system*; both must be modelled correctly if the physical and virtual simulations are to be comparable. There are two possible approaches to modelling control feedback within ADAMS:

Internal: ADAMS features a basic 'controls toolkit' capable of constructing the control system via a series of interlinked equations (for summing, gains, PID controllers, etc). The advantage of this control system is that it is entirely internal to ADAMS, providing performance benefits. However, there are a number of important disadvantages.

- The GUI is based on a series of database objects with interlinked equations, and cannot be readily visualised; this makes tracing connectivity and debugging more difficult.
- The range of features and functionality within this toolkit is quite limited; more sophisticated operations such as cross-compensating and signal slewing cannot easily be implemented directly.

- The controller is inherently 'continuous'; although this gives a 'smoother' control response, it is not ideal when modelling the effect of a discrete digital controller, and also results in slower performance, since iterative convergence must account for the continuous controller response (this is not necessary for discrete control).

External: It is also possible to interface ADAMS to external third-party software applications, such as EASY5 or MATLAB/Simulink. This works by a process of 'co-simulation'; the control system determines the inputs for a given sample time step, and invokes the ADAMS solver to solve the system mechanics for that single step. The results are then output back to the control software, to evaluate the required inputs for the next step. The obvious disadvantage of this method is that extra software is required, so data must be transferred between both programs and additional system memory is required to run both concurrently. However, the advantages are substantial. Dedicated control software such as Simulink provides a far more powerful toolkit with a more effective GUI for visualisation during design & debugging. The ability to operate in discrete time-sampled mode can actually improve performance considerably, even if the number of individual solve steps is increased, simply because the iterative convergence process is not coupled to the controller.

For these reasons, although the controls toolkit was explored, it will not be used for these studies. It was decided to investigate the option of a controller system in MATLAB/Simulink, based on availability of software & licenses. The ADAMS model was configured such that both the mechanical actuator inputs & force/displacement measures could be used as inputs and outputs for a 'plant' subsystem within a hand-coded Simulink-based control system. In theory, this control system could be augmented to include capabilities beyond the current KKS controller (e.g. tracking 'virtual' measures such as M-L load split or contact pressure). However, for these corroboration tests the control scheme was designed to mimic the extant PID feedback control and channel cross-coupling for the controllable 5 axes of the *in-vitro* rig. In the computational model as for the real controller, force or displacement control can be toggled by switching between different feedback channels. (Note that different PID values are required depending on the feedback input used). Some other minor features, such as signal filtering, load limiting, and display output, were also included. The complete Simulink controller is shown in Figure 95. The GUI layout makes extensive use of colour-coding, masked sub-systems & signal routing to simplify and organise the controller layout – this is necessary because of the control

system complexity. For illustration of the 'true' controller path layouts, a single control path (for quadriceps-driven knee angle control) is shown in Figure 96.

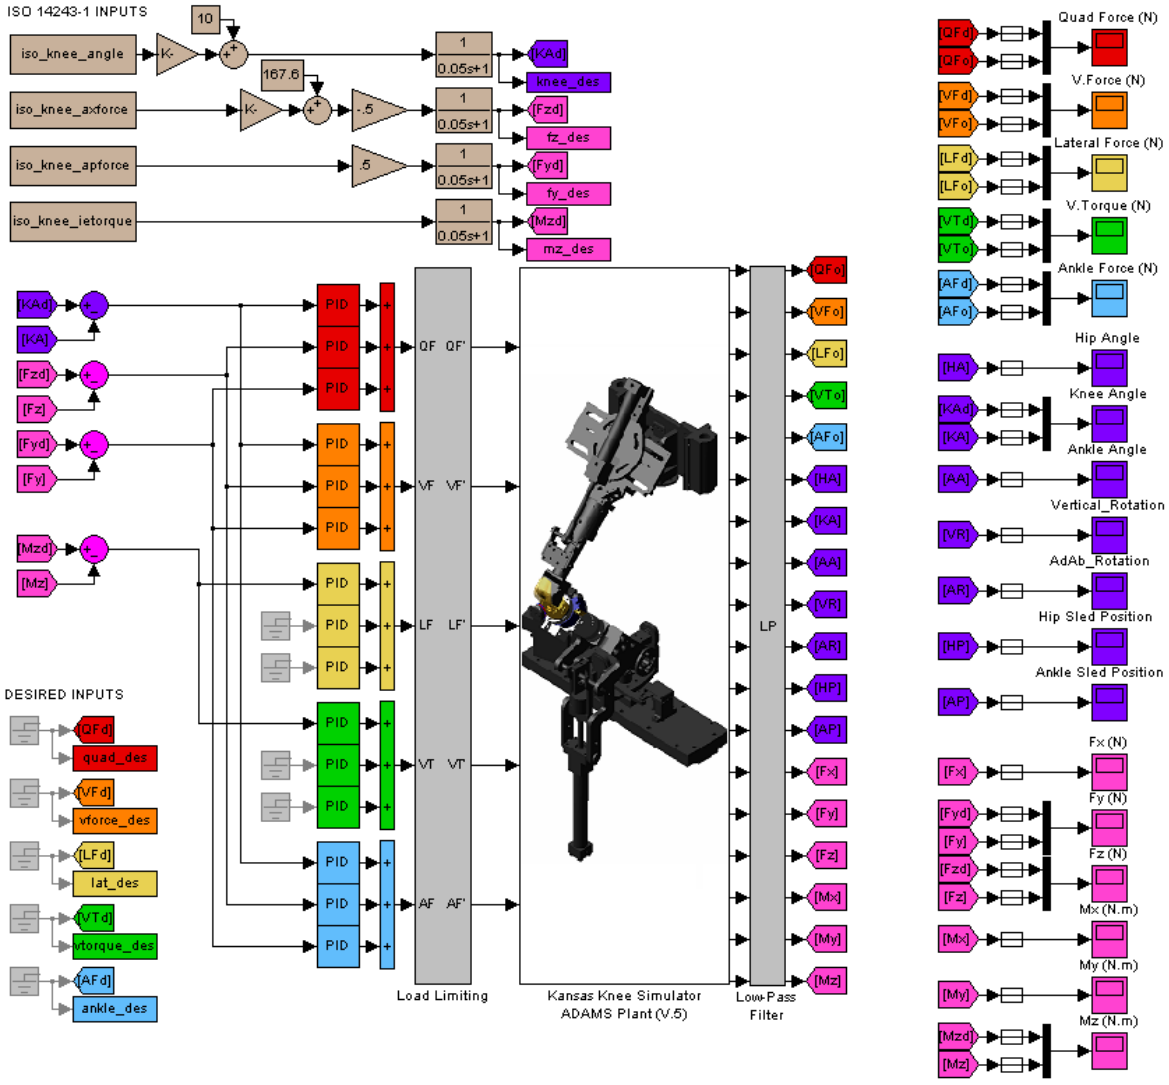

Figure 95: Simulink controller for the KKS model – top level view.

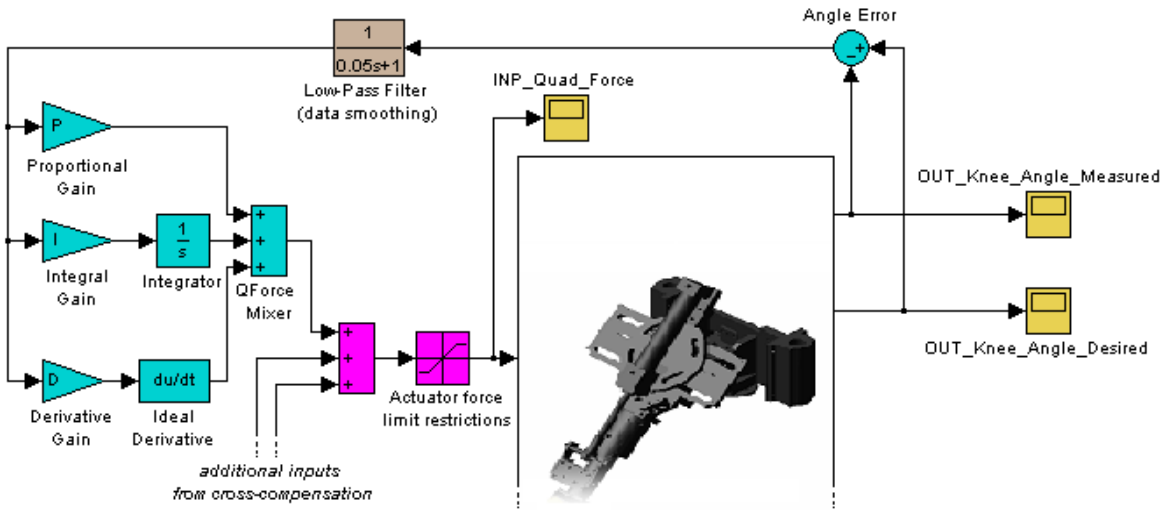

Figure 96: Simulink controller for the KKS model – partial expanded view, showing quadriceps force / knee angle control loop.

Originally, it was hoped that PID controller settings could be imported directly from the KKS Instron 8x00 6-axis servo-hydraulic controller. Unfortunately, the internal PID values for the controller software were based on a proprietary system with logarithmic-scale sample-based units. By comparison, the ADAMS/Simulink controller was designed to use true SI units for its PID controls, on a linear scale. Additionally, the polarity of some inputs and feedback channels are inverted. Hence the extant controller settings could not be used, and the model had to be re-tuned from scratch. Unfortunately, this means that the controller tuning on the *in-vitro* & *in-silico* models does not correspond exactly, introducing an additional potential source of discrepancies during correlation testing. In practice, experimental controller settings are changed from test-to-test, and so to accurately corroborate this *in-silico*, the controller values are also test-specific. However, as a general indication of the sign magnitude of the different PID terms, tuned PID values for a 'typical' gait cycle are listed in Table 14.

	Quad Load	Vertical Load	Ankle Force	Lateral Force	Vertical Torque
Proportional (P)	100	-4	-4	*Not used*	50
Integral (I)	500	-2	-16		200
Derivative (D)	0	-1	0		5

Table 14: Typical PID controller values for the KKS (gait cycle).

6.3. Deterministic Corroboration Testing [8]

6.3.1. Validation Test-Cases

With the model development complete, testing and corroboration were performed. A series of simple test profiles were devised, to be run in parallel on both the physical rig and the computational model. The test would first be run on the KKS rig (which did not track the desired inputs precisely). The feedback data was then collected and used to drive the computational model. This meant the *in-silico* case accounted for the tracking errors in the *in-vitro* controller, so the focus was on the only mechanical behaviour of the system.

[8] The validation testing in this section is reported in the conference proceedings: "Verification of a dynamic knee simulator computational model". 2008, A.N. Reeve, M.A. Strickland, L.P. Maletsky and M. Taylor. In: Proceedings, ASME SBC 2008 (Florida, USA). The present author was responsible for much of the computational model development and testing, and co-assisted with the test case verifications.

Only feedback data for position & uni-axial load cells was used at this stage; the tri-axial load cell data was kept 'blinded'. After the *in-silico* model had been simulated, the experimental data from the tri-axial load cell could be compared to the computational predictions, to give an independent comparison between the *in-vitro* and *in-silico* models (i.e. the system had not been 'tuned' to this data). This procedure could then be repeated *vice-versa*, using the ADAMS model first to obtain feedback data to drive the physical KKS.

The initial family of test-case profiles devised focused mainly on sagittal-plane kinetics. These initial tests were intended to be based on purely sinusoidal waveforms, to limit the complexity of the system response. (Subsequently, the tests would be extended to include more complex loading; e.g. applying constant loads to the M-L sled or vertical-torque axis).

An initial set of test cases were devised and run - however, these profiles used *rectified* sine-waves, i.e. $|\sin(t)|$. This results in a non-smooth inflexion at the end of each cycle (as the polarity of the sine-wave inverts) – see Figure 97. The inclusion of these tracking errors made corroboration between the *in-vitro* & *in-silico* cases more uncertain. Note however that the *in-silico* model was able to track better than the *in-vitro* rig under these conditions. This is a reflection of how much faster and easier it is to tune the controller computationally, thanks to the intuitive GUI-based controller and fast solve times.

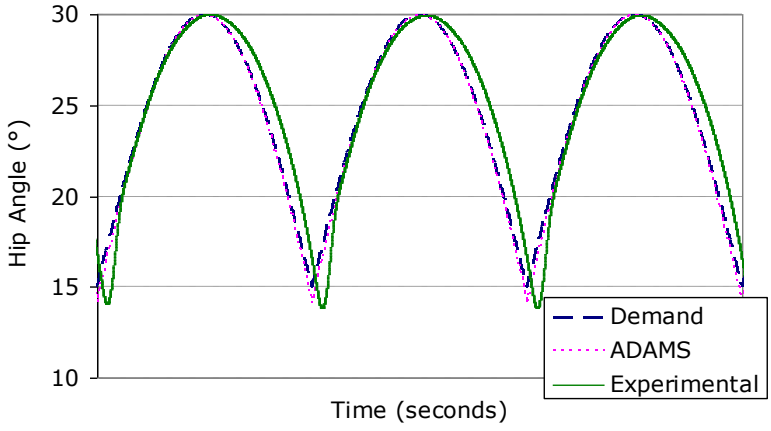

Figure 97: Initial 'rectified' sine-wave profiles. *In-vitro* tracking for the rectified sine-waves was inaccurate due to the inflexion at 15° hip flexion. These profiles were abandoned in favour of 'true' sine-waves.

Subsequently, a second family of profiles were devised, using pure sinusoids to avoid the inflexion effect; these profiles are outlined in Table 15. The quad force was used to vary position control of the hip to track a 'true' sine-wave, whilst the force-controlled loads applied to the other axes were held constant. Tracking for

these profiles was much better on both the rig and *in-silico* model, simplifying the corroboration process.

Test Case	Quadriceps Force	Vertical Load	Lateral Force	Vertical Torque	Ankle Force
TC1	Sine-wave, range ~10° - 40°, period 10-seconds	75N	-	-	-
TC2		75N	30N	-	-
TC3		75N	-30N	-	-
TC4		75N	-	3N·m	-
TC5		75N	-	-3N·m	-
TC6		75N	-	-	75N

Table 15: Revised test-cases, with pure sinusoid 'position' waveform (quad axis).

Initially, the root-mean-square (RMS) error was still quite large even with the new profiles (as high as ~15% for the sagittal plane kinematics). This was investigated further, and it was found that considerable error was induced due to *in-vitro* malpositioning. In the ADAMS model, the components were assumed to be in perfect alignment. However, on the KKS simulator, the components had been cemented to their fittings with slight inadvertent misalignment (the femoral component was placed a few degrees in varus, and slightly externally rotated).

Since the *in-silico* model had earlier been extensively parameterised, it could easily be re-positioned to match the *in-vitro* rig misalignment. This was done accordingly and the simulations re-run. Once this was accounted for, sagittal plane average RMS error dropped to around 5%; see the comparison in Figure 98. This means that the differences between the experimental and computational results are on the same order as variations due to control system tracking and sensor measurement errors. The cycle-averaged RMS errors for the six revised test cases are summarised in Table 16 (both before and after accounting for misalignment).

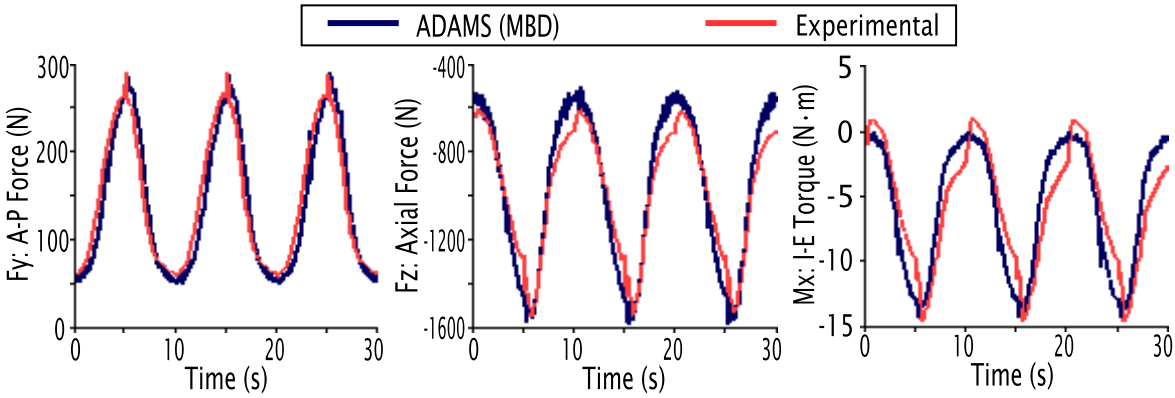

Figure 98: Sagittal-plane comparison (here shown for test-case TC6); cycle-averaged error is ~5%. Left: A-P force (F_y), Centre: axial force (F_z), Right: I-E torque (M_x).

Test-case	Initial	Re-aligned
TC1	8.6%	4.1%
TC2	12.4%	5%
TC3	13.1%	4.8%
TC4	10.5%	4.5%
TC5	10.2%	4.7%
TC6	14.9%	5.3%

Table 16: Averaged sagittal RMS errors across the profile cycle, comparing initial (perfect alignment) & re-aligned (to reflect the misalignment *in-vitro*).

These test-cases demonstrated the ability of the *in-silico* model to match the output of the *in-vitro* rig. Next, simple test profiles were devised on the *in-silico* model to drive the rig (i.e. corroborating in the reverse direction). Figure 99 shows a 'constant compressive load' profile; in which one axis drives flexion, whilst the other actuators compensate to maintain a constant axial load at the knee. Two test cases were created, with hip flexion angle (position control) driven by quad-force in one and vertical load and in the other. Note that the achieved tracking is very good for both force- and position-control (after the first 10-second cycle, which is a 'transient' and is discarded).

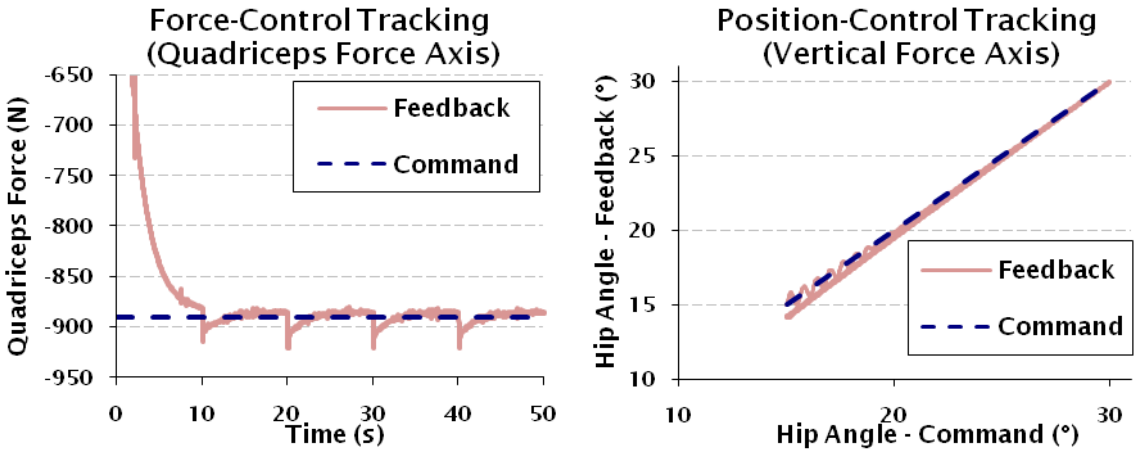

Figure 99: Example validation test (sagittal plane loading only). Constant 200lb quadriceps load (left), with 15°-30° sine-wave position-control on vertical force actuator (right). Transient 1st cycle (0 – 10 seconds) is discarded.

This profile was generated on the model, and then used to drive the simulator. Figure 100 shows the 'achieved' versus 'desired' tracking for the position control axis and the vertical load axis; the close agreement indicates the *in-vitro* profile has quite successfully been devised by the *in-silico* model. Once again, the tracking is poorest when the simulator encounters non-smooth inflexions in the profile; Position tracking (left) exhibits almost no error; vertical load (right) has

some high-frequency oscillation 'spiking' at mid-cycle (at maximum flexion), and noticeable deviation from the desired profile at the end-of-cycle inflexion.

Despite these minor differences, the sagittal tracking corroborates to within a few percent (note that this is less than the errors due to malpositioning shown in the previous test-cases). This demonstrates the concept of using the *in-silico* model to generate profiles for use *in-vitro*.

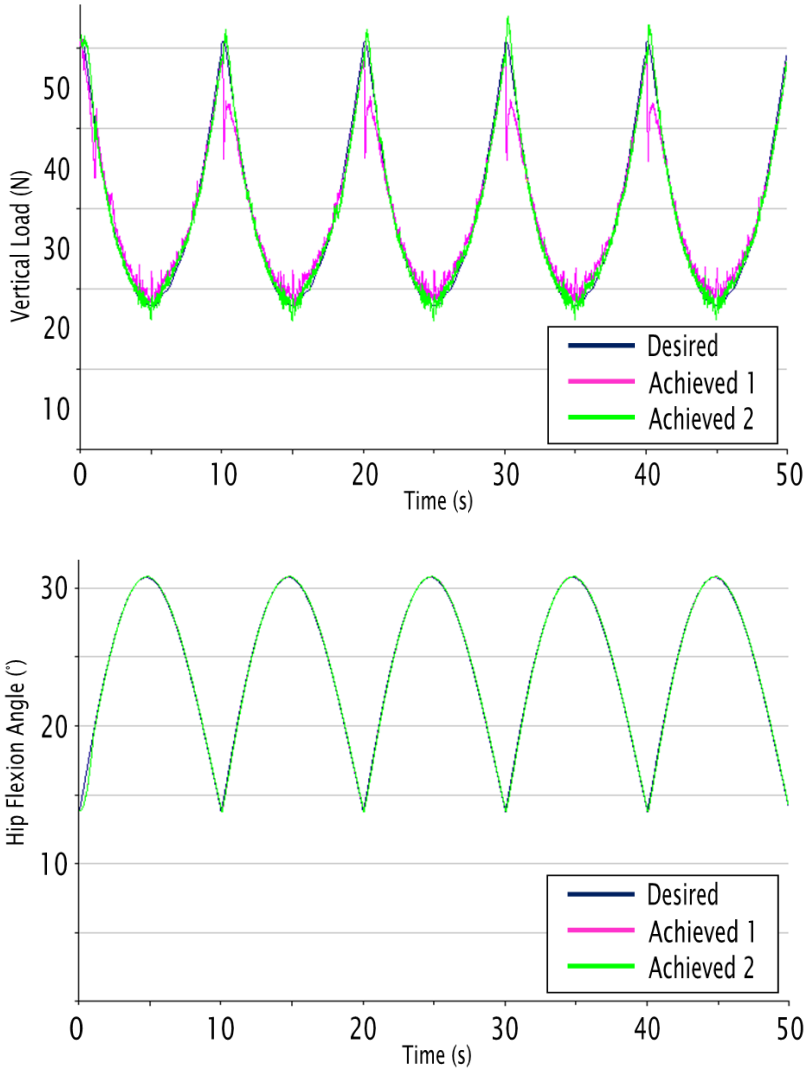

Figure 100: *In-vitro* KKS rig feedback for constant-load profile (using input from *in-silico* model). Top: force-feedback comparison for vertical load. Bottom: position-feedback comparison for hip angle.

6.3.2. Profile Generation & Testing

The computational and experimental simulators had been corroborated together to within an acceptable tolerance level; it was now possible to proceed with more physiologically-representative test conditions. Data for a wide range of activity profiles was collated from the literature, including profiles for gait [21-24, 35, 43, 44, 91, 92], stair ascent [21, 26, 28, 29, 35, 40, 175], stair descent [28, 201] and deep-flexion activities such as squat & chair rise [27, 35]. These data profiles

were then used as the input for the *in-silico* model. This did require some adaptation of the Simulink controller on a profile-by-profile basis (to account for different starting positions, or to change terms in the controller). Figure 101 illustrates the tracking achieved for one particular profile (gait data from D'Lima et al [35]). The feedback from the *in-silico* model could then be used to drive the KKS rig. Some of these profiles generated using the new ADAMS model have subsequently been used for cadaveric TKR tests by KU as part of ongoing industrial research. This gives good confidence going forward for the *in-silico* model to be used in a 'stand-alone' capacity for probabilistic studies.

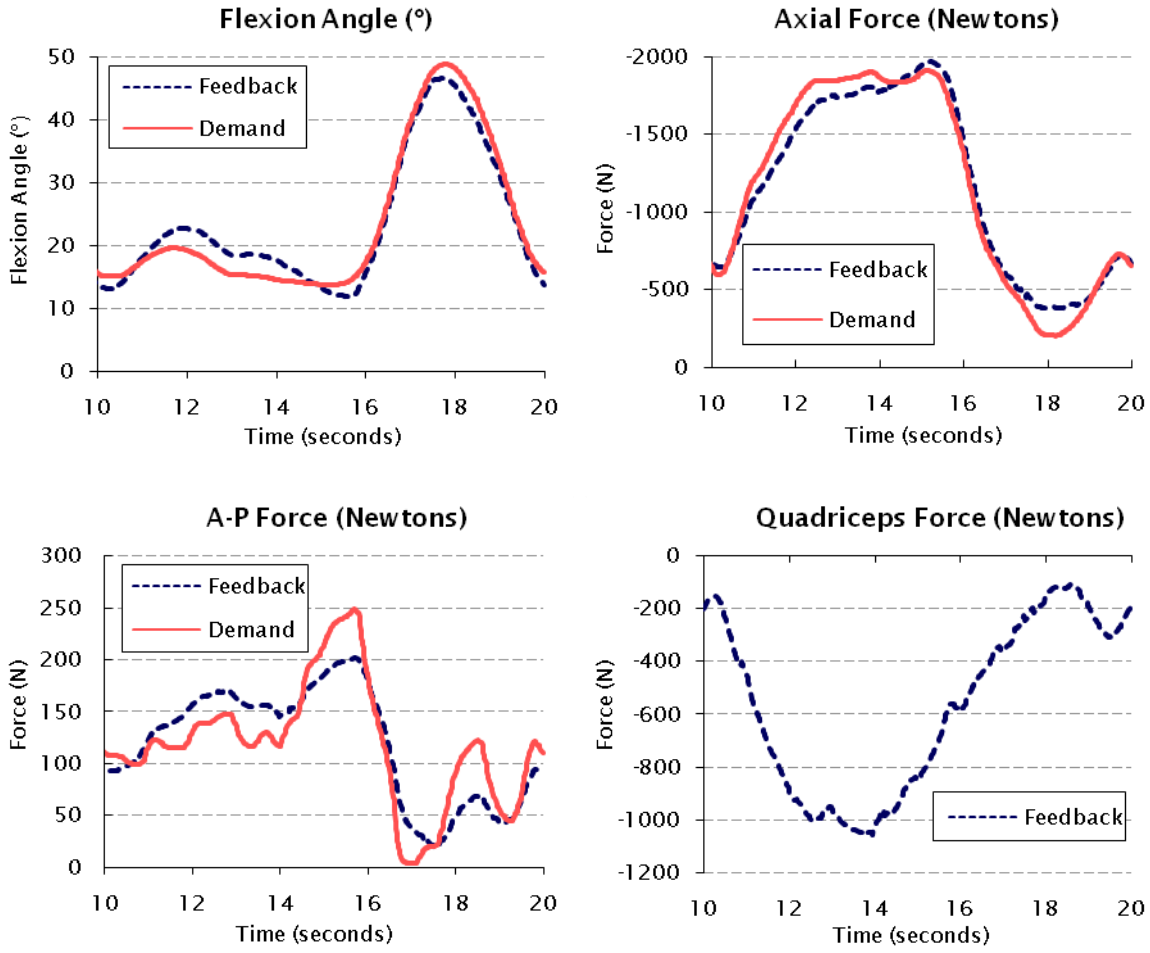

Figure 101: Profile tracking for the ADAMS model (based on *in-vivo* gait data [35]); feedback vs. demand for flexion angle and axial & A-P shear force, with an example actuator feedback waveform (quadriceps force).

6.4. Probabilistic KKS Modelling

6.4.1. Methodology

The *in-silico* KKS model is implemented differently to the previous MBD-based models; therefore before including a probabilistic 'wrapper' for, it is necessary to re-evaluate the most appropriate statistical software and models to use.

The most important difference is the use of co-simulation between MATLAB/Simulink and ADAMS; this means that ADAMS is now run 'externally' from within the MATLAB environment. As such, 'internal' applications (such as the native 'Insight' module) cannot be used. Instead, as discussed in Chapter Five, the simplest option is to directly encode a DOE 'wrapper' using the native m-file scripting language in MATLAB. This approach was used with the KKS model, as illustrated in Figure 102.

The input variable perturbations are read in as a raw 'matrix' (this can be copy/pasted from another application such as 'Insight', or generated directly). These values are then used to individually run co-simulations with Simulink/ADAMS; at the end of each simulation, the data is retrieved and stored (again in matrix-form) using comma-separated-value (*.csv) files. The data can subsequently be transferred into MATLAB, 'Insight', MS-Excel or any other suitable application for post-processing. This approach is somewhat more cumbersome than the earlier 'Insight'-based models, but is necessary to accommodate the co-simulation of the control-plant.

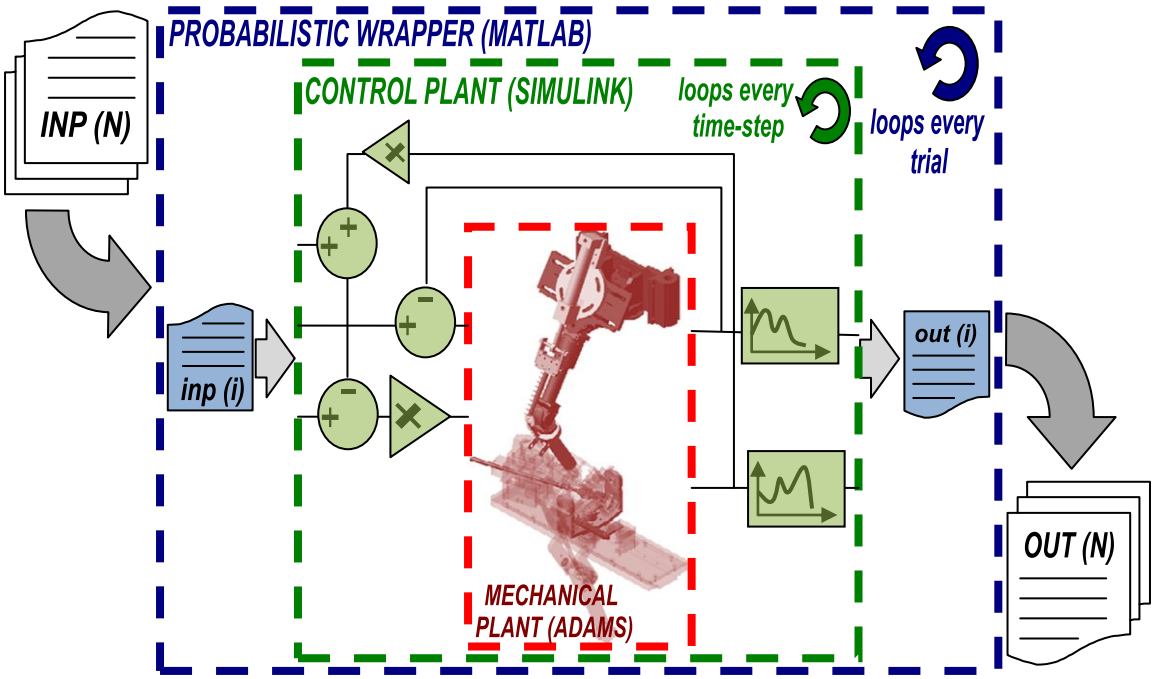

Figure 102: Concept structure for probabilistic study with KKS model. The MBD-based mechanical 'plant' model is nested within a controller 'wrapper', which in turn is nested within a probabilistic wrapper managing the multiple trials.

Note that this basic model is not capable of adaptive sampling (as discussed in Chapter Three); however, it is adequate for a conceptual study. Also, the KKS-based trials are considerably slower than the simple SKS-based models (~30 minutes versus ~5 minutes); coupled with the additional doubled overhead of wear-post-processing for *both* the tibiofemoral and patellofemoral joints, a full

trial and post-processing can require almost an hour. As such, full MCST studies are more computationally expensive, and fast RSM-based methods are a more attractive option. However, it is necessary for this first analysis to use MCST, to demonstrate the linearity of the system (in order for subsequent studies to use RSM with greater confidence).

Therefore, a 1000-trial MCST will serve as the 'baseline' probabilistic study. The study is based on the gait profile by D'Lima et al [35], as shown in Figure 101. For this first study, eight variables were selected: malrotation of the three angles for positioning of both the femoral and tibial components (as in earlier studies), as well as two new variables reflecting the full lower-limb scope of the model: The initial 'height' of the patellar implant (I-S position) which is known to affect the quadriceps extension-moment, and the quadriceps actuator 'dial' position (which controls the Q-angle on the *in-vitro* rig. These last two variables alter the effect of the quadriceps force actuator, which is a step towards more physiologically-representative loading compared to the tibiofemoral knee-wear simulators. The input factors are summarised in Table 17; for all variables, a Gaussian distribution was used, cropped at ±3σ.

Factor	Abbreviation	Mean	S.D.
Femoral F-E Rotation	Fem_FE	0°	2°
Femoral I-E Rotation	Fem_IE	0°	2°
Femoral V-V Rotation	Fem_VV	0°	2°
Tibial 'Tilt' (F-E Rotation)	Tib_FE	0°	2°
Tibial I-E Rotation	Tib_IE	0°	2°
Tibial V-V Rotation	Tib_VV	0°	2°
Patellar Clamp 'Height' (I-S)	Pat_IS	0mm	2mm
Quad Dial Angle	Q_Dial	0°	2°

Table 17: Input factors for initial KKS probabilistic study.

A wider range of output measures were monitored for this study. The standard metrics introduced for the tibiofemoral knee wear simulators are preserved here (e.g. A-P translation, I-E rotation and contact pressure). In addition, for the patella kinematics (tilt, rotation) and contact pressure were also monitored. Wear results were evaluated for both the tibial insert and the patellar insert, using the different algorithms introduced in Chapter Four. Besides this, the different force-feedback (uni-axial and tri-axial sensors) and displacement-feedback sensors on the rig were also monitored for each trial.

This represents a very large total data set and only selected results of interest are presented in the following section (since this concept study is only intended as a

demonstration of probabilistic methods for the validated *in-silico* KKS model). An important issue for probabilistic studies in general is condensing the volume of data produced, to make it concise, relevant and accessible to designers and clinicians.

6.4.2 Results & Discussion

The volume of data generated by a probabilistic study on the KKS is considerable; it is possible to retrieve force feedback and displacement feedback for the entire rig as a whole, as well as the tibiofemoral and patellofemoral joints in isolation. Here, a number of pertinent observations are made regarding selected results:

KKS Rig Feedback: Tracking Responses

The probabilistic study was run in conjunction with a controller; for some specific channels (flexion angle, axial load), this means that the controller is working to reduce any variability for those axes. This is apparent in Figure 103; for this simulation, quadriceps force is used to control flexion angle. Consequently, there is almost no flexion-angle variability, whereas the quadriceps actuator effectively compensates for the variability in flexion angle, and hence shows very high levels of variability itself. This is important, conceptually, when devising the control system for a probabilistic study (as will be discussed further subsequently).

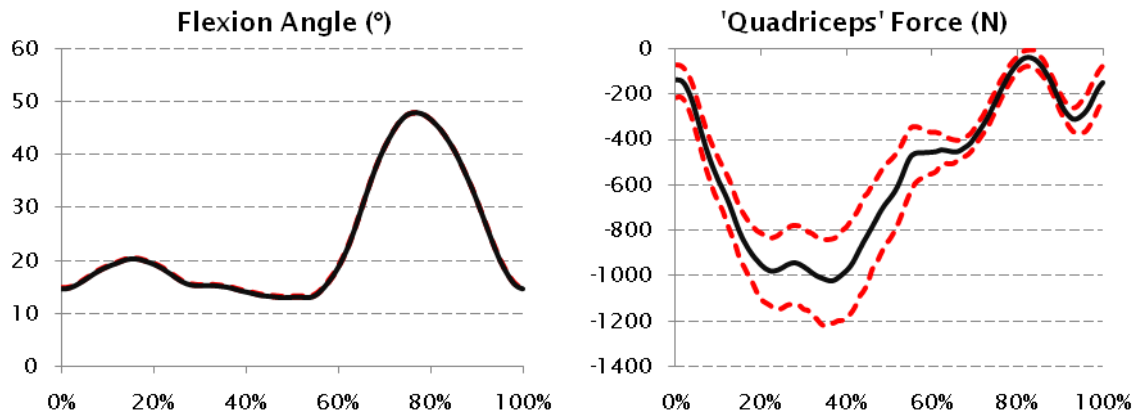

Figure 103: Comparing flexion angle (controlled) and quad force (driving); 5%-95% envelopes.

Tri-axial Load-cell Feedback

The six channels from the load-cell (Figure 104) re-emphasise the above observation; three of the channels are 'controlled': A-P force (F_Y), axial force (F_Z) and I-E torque (M_Z). The other three are uncontrolled. The differences are immediately apparent. Putting certain axes under tight closed-loop control effectively constrains other available degrees of freedom to compensate. Ultimately, this leads to higher variability in the uncontrolled channels.

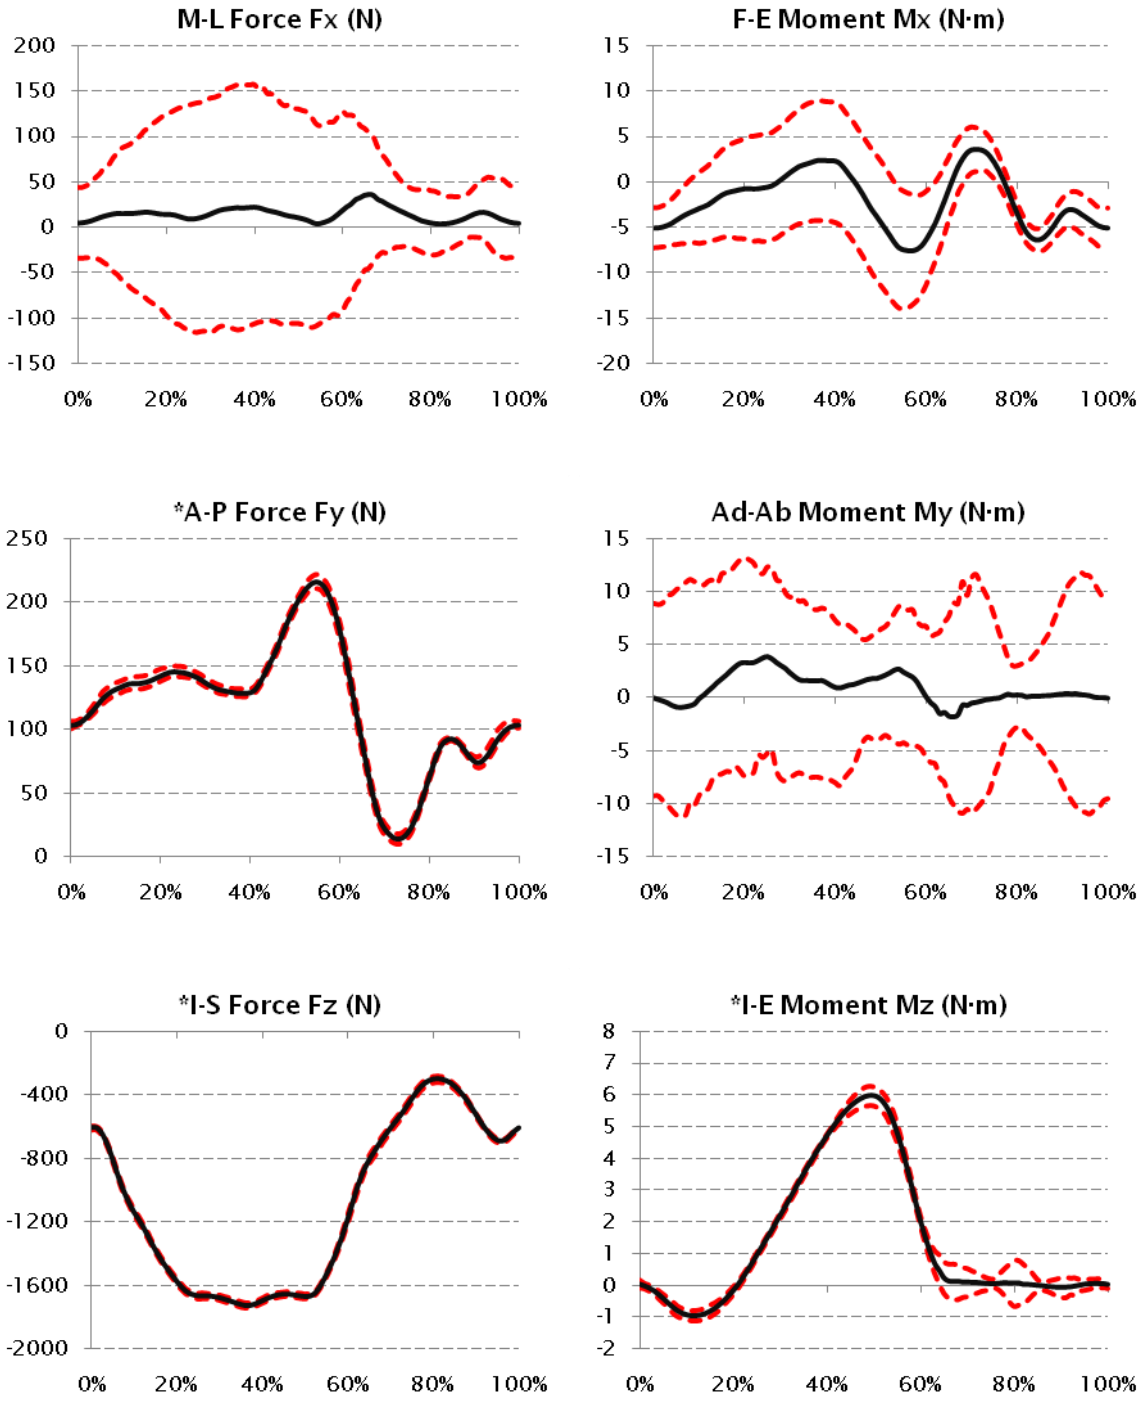

Figure 104: Response envelopes for the load-cell forces & moments (5%-95%). Note: axes pre-scripted with an asterisk (*) are under direct force-control.

This raises an important question for the pre-clinical test designer: what is the aim in incorporating variability into a study? In this case, as a baseline analysis, the controller was commanded to track the same input loads and displacements, regardless of the implant mal-alignment. However, in reality, severely misaligned components would probably lead to adaptations in the nature of the gait cycle kinematics and kinetics; in other words, a single 'target' profile would not be applicable. In this case, constraining such a tightly-controlled profile leads to a

wide imbalance between a number of control channels with very limited variability, and a much higher level of variability on the other uncontrolled axes. It may be important for future modelling efforts to re-evaluate this approach, to achieve more physiologically meaningful variability studies.

Tibiofemoral Mechanics

The isolated feedback for the tibiofemoral joint might be compared to the earlier tibiofemoral probabilistic simulations in Chapter Five. The profiles are not strictly comparable (besides differences in the *in-vitro* platform, the input profiles are also different: the earlier studies were based on ISO-prescribed gait, whereas this study was based on telemeterised data), nonetheless, comparisons can be drawn.

Figure 105 shows the kinematic envelopes. The A-P envelope shows similar trends to the earlier models. The input variability (standard deviation of misalignment angles) is lower; however, more factors were included in the earlier studies, which would increase the envelope size. The envelope for I-E rotation is quite different; it reveals a very high level of variability for the KKS in swing phase. (This swing-phase variability is also apparent for some of the force-feedback and load-cell data).

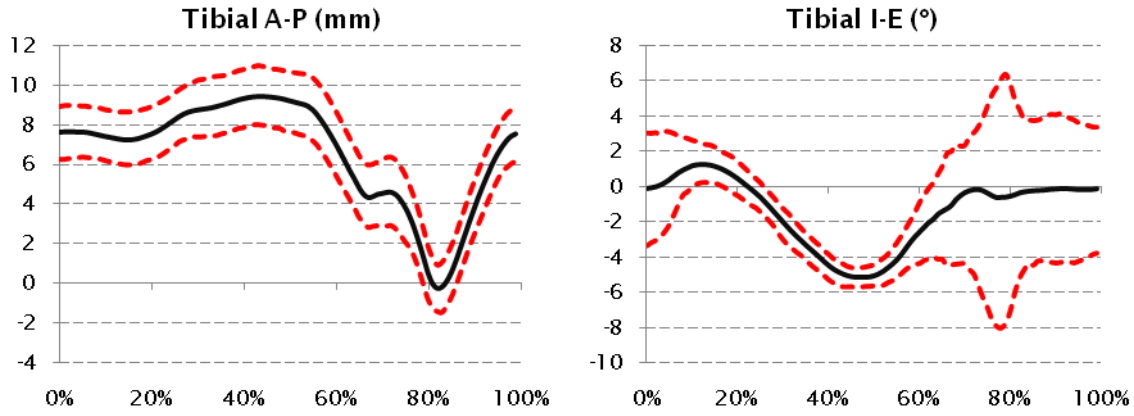

Figure 105: 5% - 95% Envelopes for tibiofemoral kinematics: A-P (left) and I-E (right).

The peak CP (Figure 106) is higher for the KKS data, *despite* the fact that the input profile specifies lower axial forces than the ISO standard (~1800N compared to 2600N). CP variability is also high; given that the contact forces are well-controlled this suggests that the contact area is quite variable, depending on component mal-positioning.

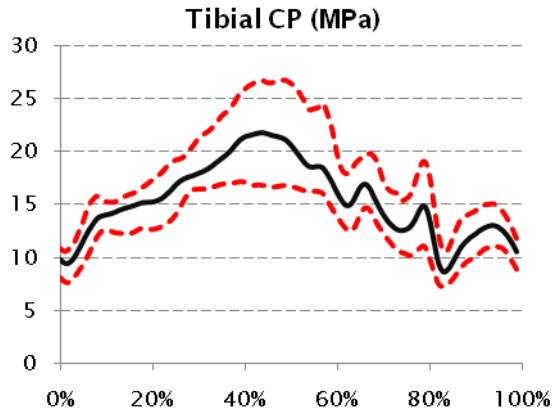

Figure 106: Peak CP for tibial insert (5%-95% envelope).

Patellofemoral Mechanics

The KKS patellar kinematics are not strictly representative of *in-vivo* kinematics; the KKS 'patella' assembly has no lateral constraint, so higher levels of M-L translation and patellar tilt are possible. Figure 107 illustrates this; I-S translation and patellar rotation are both relatively well-constrained. However, the M-L translation and patellar tilt show very high variations (translations up to 30mm and rotations up to 30°).

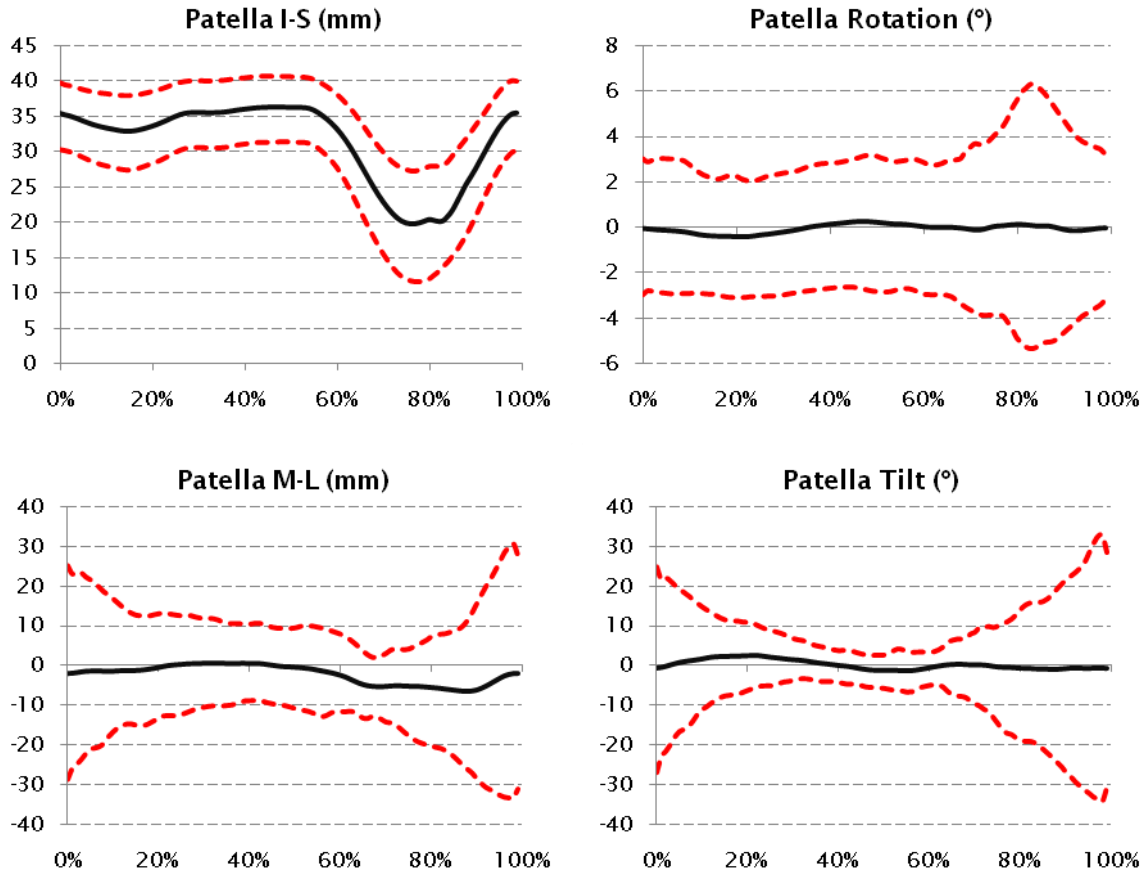

Figure 107: Patellar kinematics: I-S translation (top, left) and patellar rotation (top, right), are well-constrained. M-L translation (bottom, left) and patellar tilt (bottom, right) exhibit excessive variability.

These extreme values are *not* physiologically representative; this illustrates the important role of the para-patellar retinaculum and medial patellofemoral ligament in controlling patellar kinematics. Note that in the *in-silico* model, the degree of tilt and M-L translation was found to be very sensitive to the friction and damping at the patellofemoral joint – once again demonstrating the importance of accurately characterising system dynamics between the *in-vitro* and *in-silico* models.

Patellar peak CP correlates closely with the quadriceps-actuator force (Figure 108). This would be expected, since this actuator provides the constraining force which is principally responsible for the compressive load on the patella. Note that the periods of lower contact pressure correspond to the greater variability in M-L translation and patellar tilt – the patellar motion is more erratic when the constraining load is reduced. These pressures reported for the patellar are well-beyond the range of linear-elasticity for UHMWPE; in practice rigid-body modelling will not yield meaningful values for contact pressure under such extreme conditions (inspection of the polyethylene components on the *in-vitro* rig demonstrates that visible plastic deformation of the components does occur under normal use). Therefore, these values must be interpreted with some caution, especially for the highest reported pressures.

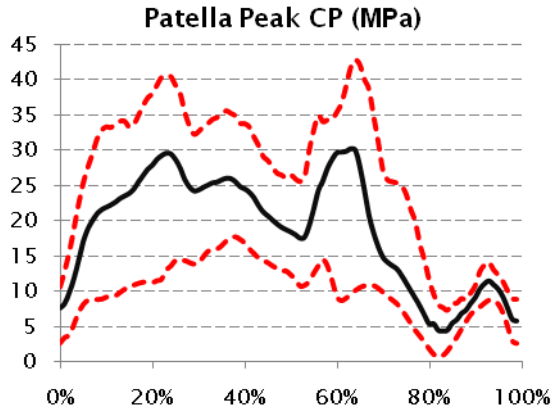

Figure 108: Patella insert peak contact pressure, MPa (5%-95% envelope).

In-silico Wear Prediction

The tibiofemoral wear results compare well to previous studies with tibiofemoral knee wear simulators. The different wear models all predict a 'mean' wear rate in the region of ~7–9 mm^3/MCycle (Figure 109). Once again, Archard wear is the least variable, with a standard deviation of only ~0.25mm^3/MCycle; for the CS-based models the wear variability is many times higher. Once again, a characteristic asymmetry is apparent in the wear PDFs; the 'tail' of the

distribution for top-end wear rates predicts a considerable proportion of high-wear outcomes, given this level of input variability.

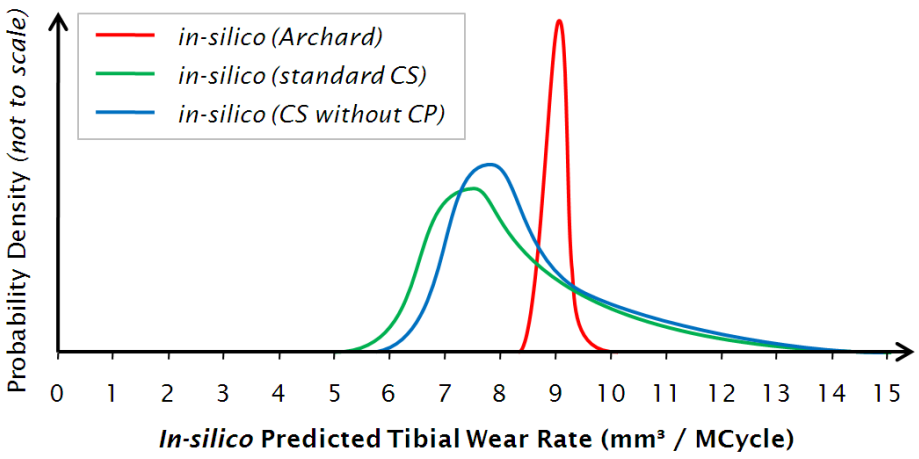

Figure 109: Predicted KKS tibiofemoral wear rates for different wear models (PDF magnitudes scaled for clear comparison in figure).

The patellofemoral wear results are less reliable, in light of the under-constrained kinematics for tilt and M-L translation (Figure 110). The 'top-end' wear rates (highest values in the PDF) are unrealistically high, owing to the inflated sliding distances from the exaggerated kinematics. However, the mean-value predictions for the patellofemoral wear (~2-5 mm³/MCycle) are comparable to the limited available data for patellofemoral wear *in-vitro* (e.g. Ellison et al reported rates of 2.2±1.2 mm³/MCycle [202]).

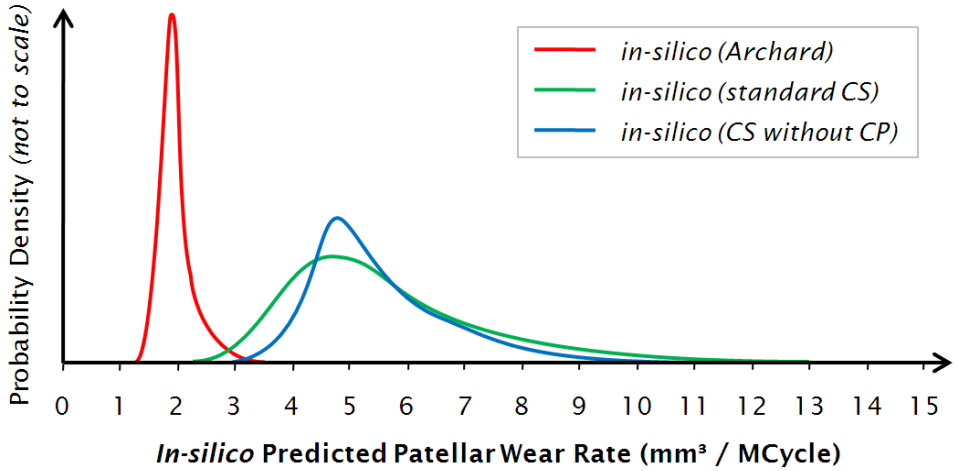

Figure 110: Predicted patellofemoral wear rates for different wear models (PDF magnitudes scaled for clear comparison in figure).

It is possible to use SA to determine the sensitivity of this variation in wear to the different input factors. Figure 111 shows the linearised sensitivity factors for the tibial and patellar wear rate, based on the "M-L/M-L+A-P" wear model. The results reveal some similarities, but also notable differences between the factors contributing to wear of the patellar and tibial inserts.

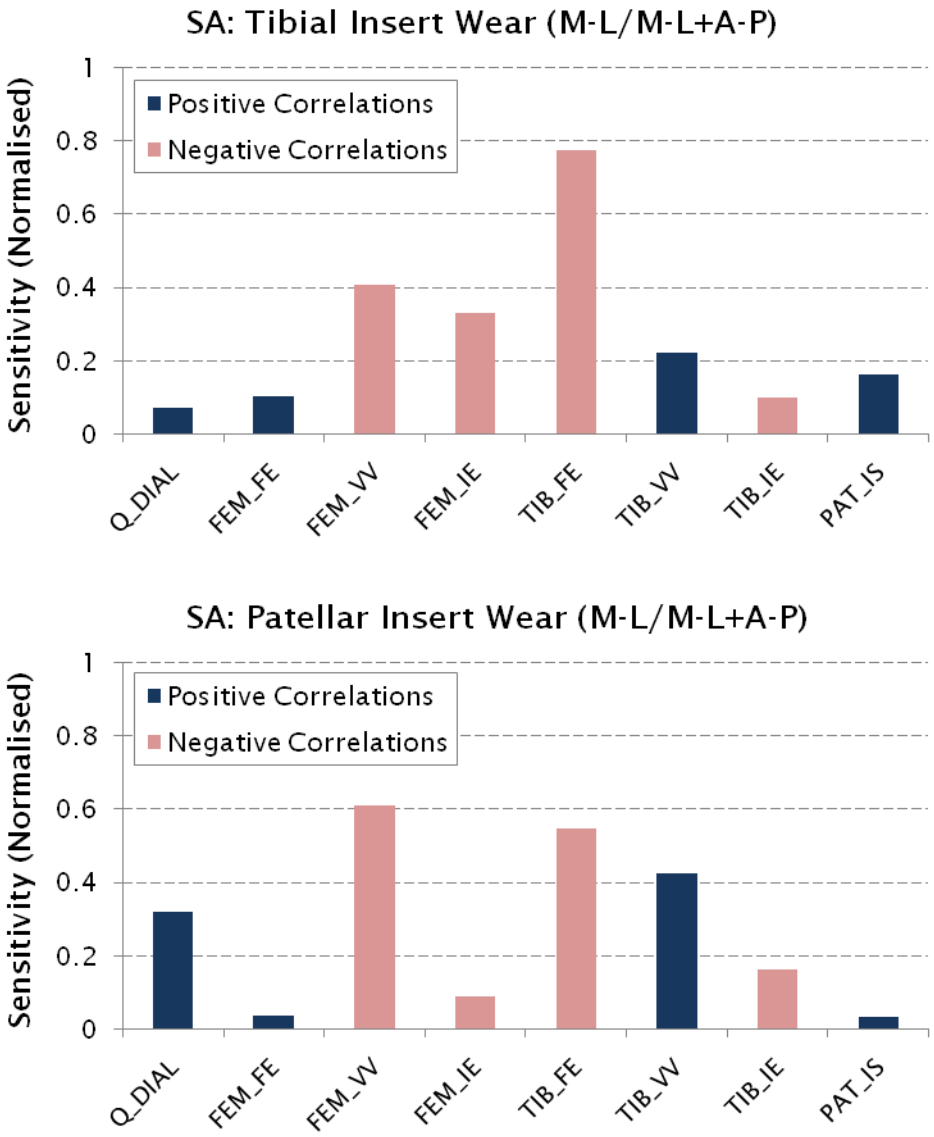

Figure 111: Sensitivity analysis results for tibial wear (above) and patellar wear (below). Wear model analysed is the 'M-L/M-L+A-P' formulation.

For the tibial insert, wear is dominated by the tibial tilt (Tib_FE), with the other mal-rotations also moderately strong. This is comparable (but not identical) to the findings of Pal et al [179] - note that their study was based on the SKS knee wear simulator, so results would not be identical (friction, which was the main factor in that published study, was not varied in this initial KKS probabilistic study). For the patellar insert, the sensitivities are more distributed; some of the mal-positioning factors are still important, but the quadriceps dial angle (which directly affects patellar tracking) is now also more influential. In both cases, the F-E offset of the femoral component has minimal influence; the clamping height of the patellar insert on the Kevlar 'QT' strap is also relatively unimportant. Note that these sensitivity factors are design-dependent and activity-dependent, so general rules cannot be inferred from this one data-set as to which factors are influencing wear under other test conditions. A similar probabilistic approach would be needed to investigate on a case-by-case basis.

The KKS is not used as a wear-assessment platform *in-vitro* (historically, it was decided by the original designers after early work on the rig not to pursue this line of development, for various practical reasons). Consequently, it is not possible to corroborate these wear results. This is a disadvantage of the KKS model; it can be used to corroborate the kinematics and mechanics of the platform, but cannot be used to further explore *in-silico* wear prediction algorithms.

The study demonstrates conceptually the application of probabilistics to the KKS *in-silico* model, and raises some important questions about the control-philosophy for studies of variability on the KKS. This work could easily be extended to include a wider range of factors, different ADL activity profiles, or indeed control-system modifications.

6.5. Discussion

The studies in this chapter represent a considerable degree of experience gained through modelling and corroborating the KKS platform. Although the target of the models is a different platform to the work in previous chapters, many of the lessons learnt are more broadly applicable.

It is clear that having better experimental data available (in terms of quantity and quality) permits much better corroboration of the computational model. For the KKS, the model is much more *specific* to a particular hardware configuration; input from CAD data to validation test feedback has been used to ensure the *in-silico* and *in-vitro* models are well aligned. This gives much greater confidence in results, since the test results can be matched directly to real-world data. A two-way collaboration between experimental and computational researchers means that the validation process has been pro-actively *designed* (i.e. choice/number of tests, degree of complexity etc); not merely attempted reactively *post-hoc*. This resulted in a more systematic, more comprehensive validation process, where complexity was progressively phased in and so the design & test process is more structured.

Note that although the initial model is quite specific, it is of course possible to generalise or customise the model, beginning with the corroborated version as a 'baseline'; the penalty trade-off is that progressive modifications make the model more flexible to new studies, at the cost of diverging further from the corroborated benchmark.

The *in-silico* model of the KKS developed represents a highly useful model, with the potential for further usage and development. The model has been used

during the course of these studies as a test-bed for developing new activity profiles, and also to explore potential rig modifications. These and other uses (e.g. predicting behaviour changes for cadaveric knee specimens) could be developed further, or combined with probabilistic methods to better-understand uncertainty in the KKS testing.

Lessons can be applied from this modelling approach to the world of *in-vitro* knee wear testing. A more specific, targeted corroboration would result in a better-defined, more accurate *in-silico* model; this could then be used as a baseline for further study. However, as has been demonstrated, this requires sound collaborative links as the prerequisite for better cross-disciplinary co-operation. It is apparent that the complexity of the target platform need not be an issue; the KKS is far more complex, mechanically and in terms of control systems, than the knee wear simulators. However in spite of this complexity and the number of unknown parameters affecting the dynamic behaviour of the KKS, corroboration to within good accuracy (>95%) was possible. In light of this, it should clearly be possible to achieve better corroboration with the simpler knee wear simulators. The final chapter describes attempts to develop such an advanced model for one specific knee wear test rig.

CHAPTER SEVEN – ADVANCED KNEE-SIMULATOR CORROBORATIVE MODELLING

Demonstrating computationally-enriched pre-clinical analysis methods for the AMTI Knee Wear Simulator

7.1. Introduction

The studies in Chapter Four demonstrated that the MBD environment is well-suited to high-speed studies of TKR mechanics *in-silico* wear assessment. Chapter Five further demonstrated that probabilistic methods can be used in conjunction with these baseline models for a more holistic picture of TKR performance. Chapter Six showed that close corroboration of *in-silico* models with *in-vitro* data makes the models much more robust and gives much greater confidence in the results.

In this final chapter, the capabilities, methodologies, and lessons learnt in the previous work are integrated into a highly-robust, extensively corroborated validation model for a specific knee-wear simulator design, including computational wear modelling and, ultimately, a probabilistic study demonstrating the corroboration between *in-silico* and *in-vitro* stochastic data sets.

7.1.1. The AMTI Knee Simulator

The modelling in this chapter is targeted specifically at a commercial knee simulator design by AMTI. This simulator is used widely in industry, and through industrial collaboration links, it was possible to access high-quality data and research expertise for this platform, which is essential for robustly corroborating any computational model to a high standard.

The AMTI-Boston KS2-6-1000 (Figure 112) is a 6-station servo-hydraulic knee simulator, conceptually similar to the other commercial rigs available for tibiofemoral knee wear testing (e.g. the SKS, the MTS-Bionix or Leeds/ProSIM designs introduced in Chapter Two). The six stations are divided into two 'banks' of three (left and right), to compare different designs under test. Note that the stations are not truly independent; feedback is based on the first station only, and common inputs must be applied to each station in the bank.

Figure 112: AMTI-Boston KS2-6-1000 Knee Simulator (Image: Advanced Mechanical Technology, Inc.)

There are some important design configuration differences between the AMTI rig and other commercial designs, as described below:

- The rig is capable of both force-driven and displacement-driven operation; however when running under force control, unlike the SKS, it does not use 'physical' spring buffers to mimic soft-tissue restraint. Instead, it uses a proprietary 'virtual' spring-restraint system (using software-based compensation of the driving inputs). This has the advantage that the soft-tissue effects can be re-programmed and customised by the user; however it does introduce an additional degree of complexity to the control system.

- The physical configuration (i.e. how the different degrees of freedom are modelled) is quite different to the SKS; for instance there is no single 'hinge' for varus-valgus; both V-V and M-L motions are combined with a roller-bearing system. Figure 113 illustrates the configuration for a single-station of the rig. Unlike other rigs, the AMTI simulator applies A-P force and displacement to the *femoral* component, with all the stations in that bank linked together and driven by a single central actuator (i.e. the stations are not independent). Note however that the force-feedback is measured beneath the tibial platen; this means that inertial effects between the points of application and measurement need to be considered.

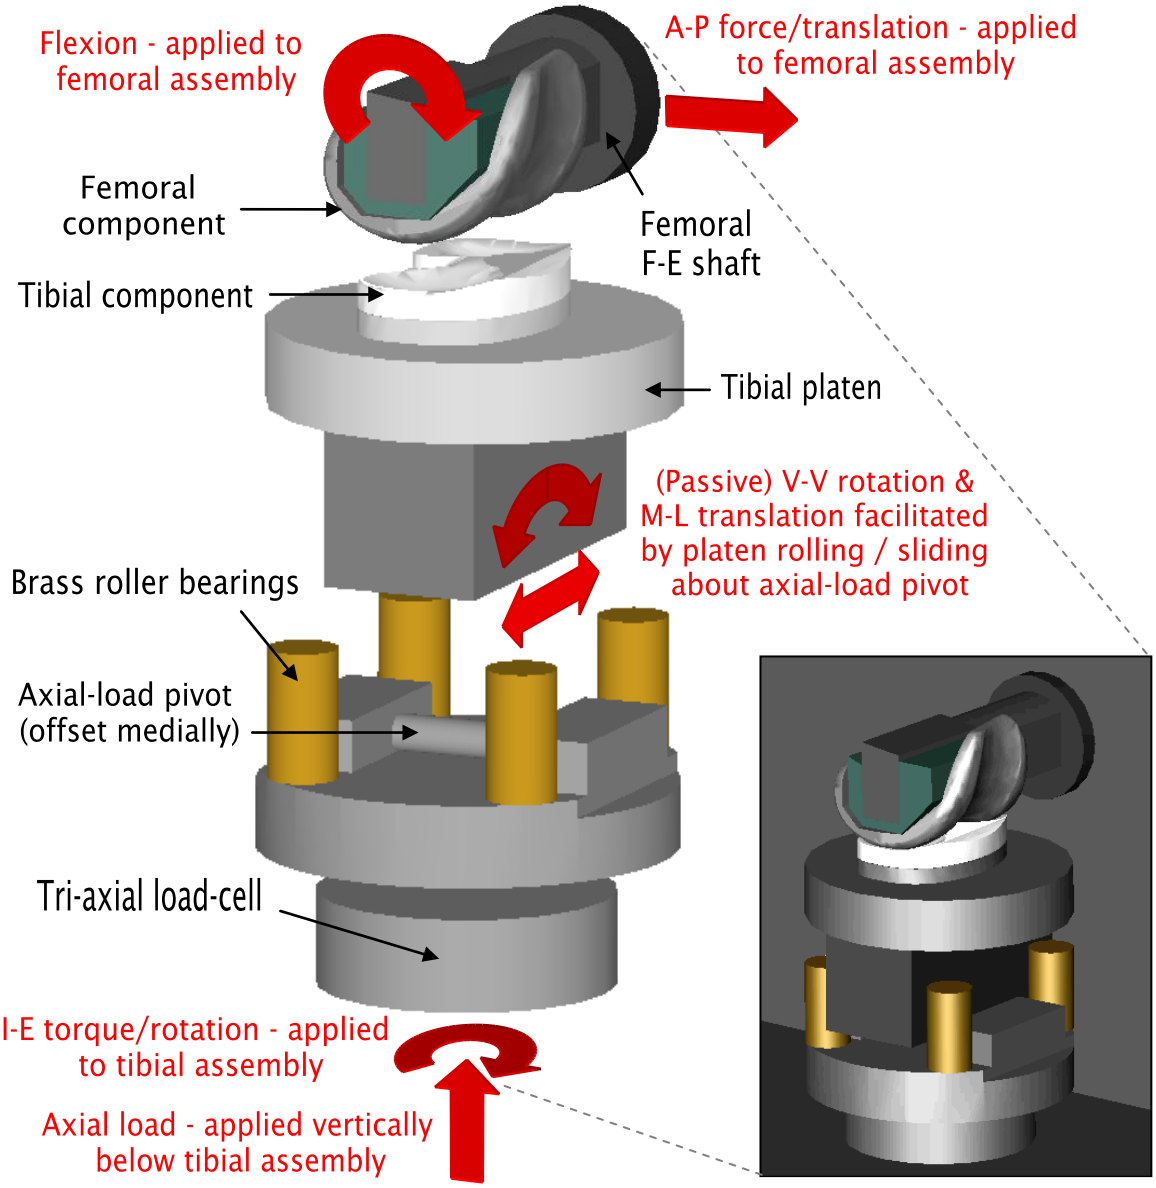

Figure 113: Schematic for single station of AMTI simulator: exploded view and (inset) *in-situ*.

7.1.2. Modelling Strategy

Because construction materials, dimensions and other details were known or could be measured directly for the parts in the rig, it was possible to construct a faithful representation of this configuration *in-silico*. Further, through industrial research links, it was also possible to run certain tests specifically to measure dynamic characteristics of the system (besides using other test data to validate the model).

Note that the machine operates under both displacement- and force-control. Therefore there are effectively two quite distinct modes of operation, which need to be characterised and corroborated separately. Displacement control is conceptually simpler, and so was addressed first. A 'crawl-walk-run' incremental approach was taken to modelling. Initially a very basic mechanical model was

constructed; subsequently the various properties (inertia, friction and damping) were measured, estimated, or 'tuned' to experimental data. Finally, more sophisticated features (controller modelling for force-driven operation and *in-silico* wear prediction & visualisation metrics) were incorporated once the baseline mechanical model was validated. Ultimately, the integrated model could be used for a probabilistic analysis.

7.2. Displacement-Driven Modelling: Corroboration

7.2.1. Modelling Details (Methodology)

The first target was a baseline mechanical model of the AMTI rig, to be operated in displacement-control. This rig has been the target of previous *in-silico* modelling. Zhao et al used MBD methods to model the AMTI simulator [203], but this model was mechanically very simplistic and neglected the friction, damping and other details.

Lanovaz et al [200, 204] made a more robust effort to corroborate the rig using FE-methods, exploring the effects of inertia, friction and pliancy; however this model failed to address dynamic damping terms so was less accurate for force control, and did not include *in-silico* wear prediction; it also used deformable FE, and so was very slow (7½ hours even with 8 processors) compared to MBD or rigid FE methods.

For the present work, the extant SKS model was used as the baseline for the new model, and overhauled to reflect the AMTI configuration. The principle changes are highlighted below:

- Model domain scope: the SKS model included only the tibial and femoral components; the actual mechanics of the rig construction were ignored. Here, the rig fittings have been explicitly modelled; the model includes the tibial platen and roller-bearing assembly down to the load-cell. This allows the model to include additional sources of inertia, friction and damping; e.g. the bearing friction between the platen base and the brass roller bearings is very variable and can sometimes be quite high; this would be neglected if the assembly was not modelled.

- Tri-axial load-cell: integral to a more robust corroboration is accurately predicting the forces measured by the load-cell. To facilitate this, a series of 6 measures (3 forces and 3 torques) were included at the same relative location as the *in-vitro* sensor, and with polarities to match the experimental data (Figure 114).

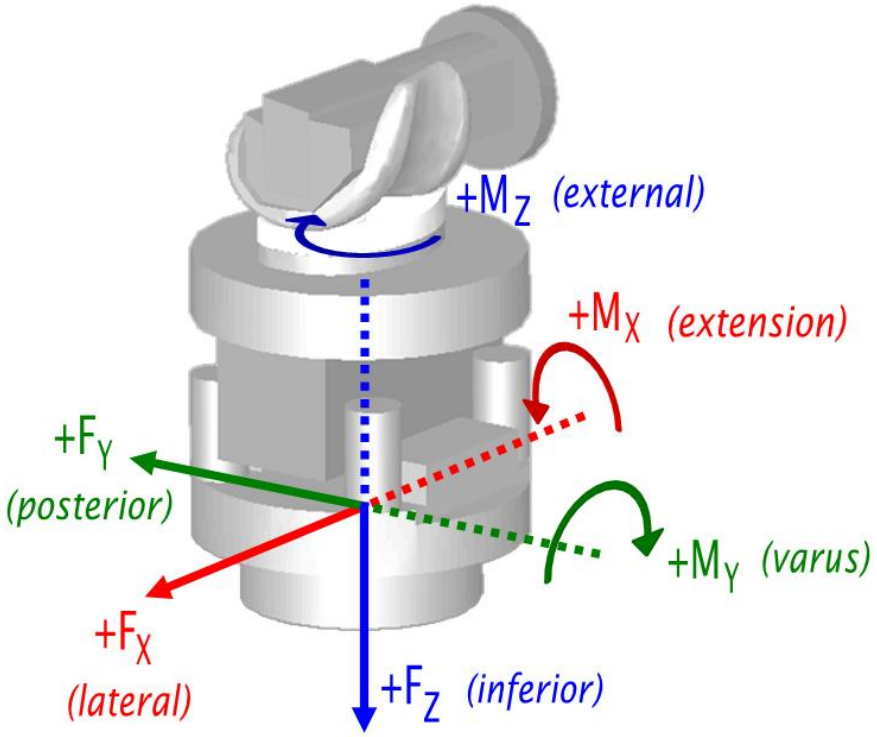

Figure 114: Orientation and polarity for load-cell measurements.

- Inertia: the inertia of the tibial assembly is particularly important, as this component forms part of the control-feedback loop (see later section on force-driven controller modelling). The mass and moments of inertia for the platen were calculated based on geometry and density, with the mass verified by direct measurement. For the femoral assembly, mass is less critical; approximate values were assigned for the components based on density estimates (on the order of ~3kg); sensitivity tests demonstrated the model is insensitive to changes in femoral assembly mass under displacement control.

- Observation of the rig suggested that there was some pliancy on the femoral axis (especially in the A-P direction). Lanovaz et al modelled this as a linear-elastic deflection of the F-E shaft [200], however investigation of the force-displacement relationship suggests a nonlinear 'backlash' effect. Whatever the source, there is a noticeable hysteresis effect between force and displacement, which must be accounted for in the computational model. The result can be up to a ~0.5mm difference between the reported and true A-P displacements for force-driven gait profiles. Note that this is still an issue for displacement-driven tests, since the system can only track the measured (not actual) displacement. The discrepancies in displacement may be small, but for a conformal implant under test, this corresponds to large differences in the A-P shear force; e.g. see Figure 115, where the predicted A-P force almost doubles, if the flexion arm pliancy is not included in the model.

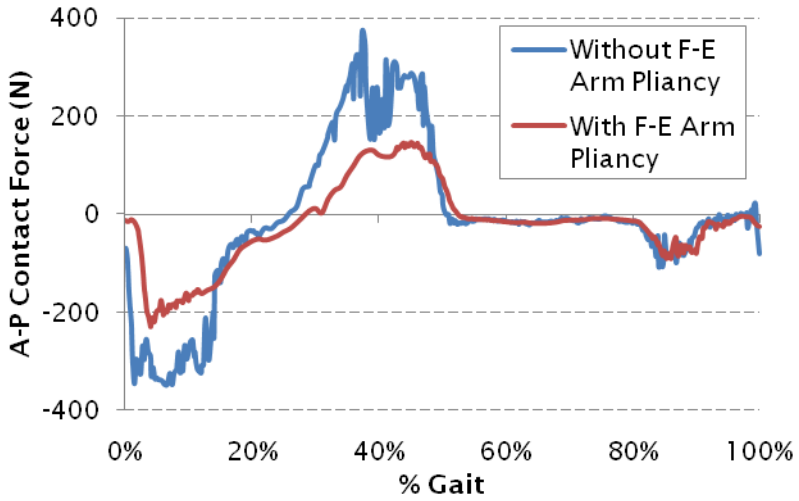

Figure 115: Effect of including F-E arm pliancy on A-P contact forces.

7.2.2. DD Corroboration Test 1: 'High-Kinematics' Gait Test

With a baseline model developed, corroboration testing was performed. *In-vitro* data was available for the two CR FB knee variants from Chapter Four (S/C and U/C). Two standard profiles had been tested under displacement-control; ISO-gait and a 'high-kinematics' (HIKIN) alternative (based on the profile in [44] – see Figure 116).

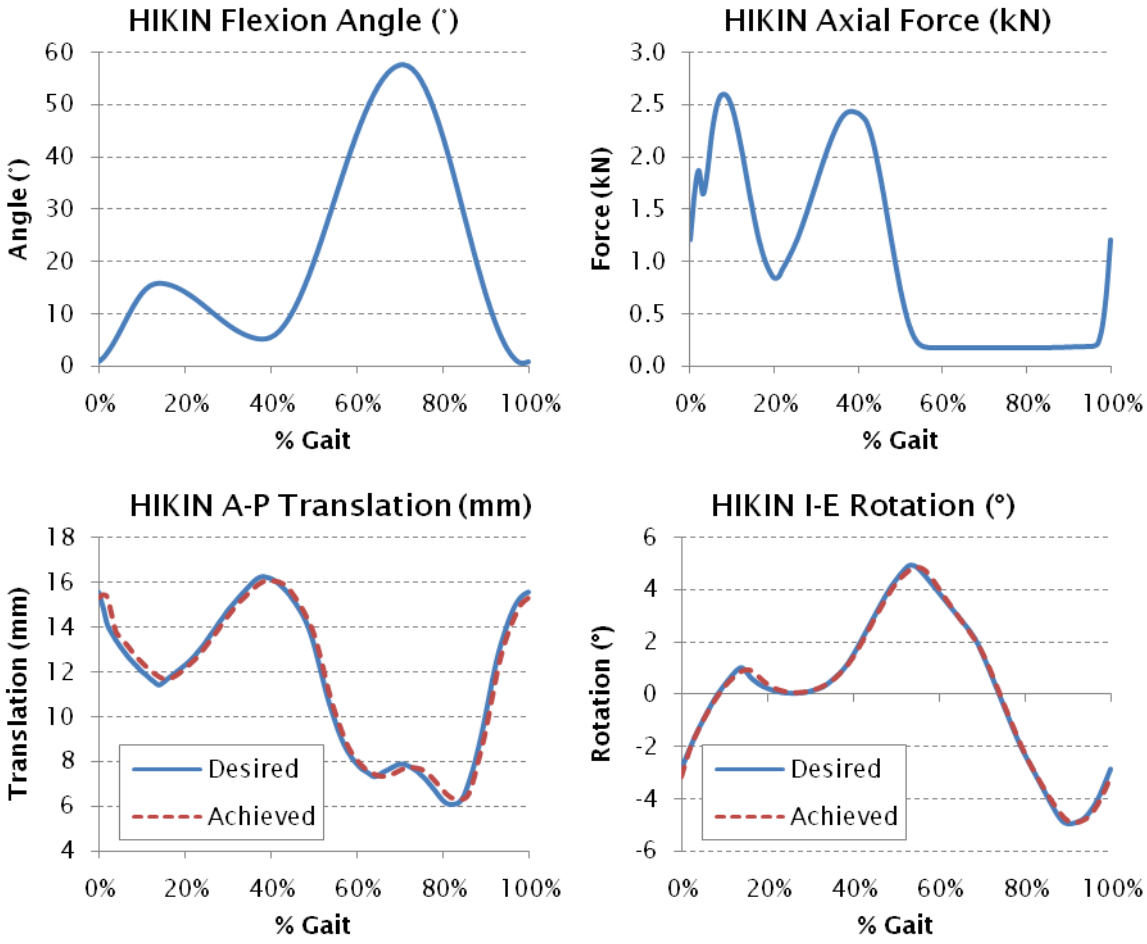

Figure 116: Input Waveforms for 'HIKIN' profile. Note 'achieved' kinematics are slightly smoother around sudden inflexions.

Various tests were corroborated, but the S/C insert with HIKIN will be focused on here. The AMTI simulator does not perfectly track these profiles (which require very sharp inflexions in the A-P and I-E profiles); the precise tracking varies depending on the implant under test, but there is always some deviation between the desired and achieved profiles. In Figure 116 the 'desired' and 'achieved' waveforms are compared. The main differences are around the sharp inflexions; the differences are small, but the impact on force-feedback (due to the higher-order derivatives of these inputs) is considerable. For corroboration testing, the 'achieved' (feedback) waveform was used (when available).

Note: for the AMTI rig, A-P translation is defined as the distance from the centreline of the fixed F-E axis to the centreline of the tibial platen; a positive value indicates the 'femur' is more posterior relative to the 'tibia'. The initial value (~15mm) is the 'dwell point'; i.e. the point where the components are in their 'settled' position under nominal compressive load at full-extension. This is a design-specific value (e.g. the value is generally lower for the U/C design).

The profile was simulated *in-silico*, and the resulting tri-axial force predictions compared to the *in-vitro* feedback data. Initial studies revealed that certain factors in particular were very influential:

- 'Dwell point': although the dwell point is theoretically prescribed by the profile offset, in practice small errors in the simulator setup, component positioning and the axis 'zero' positions can result in variability (generally < 1mm) in the exact dwell point. This is small, but sufficient to make a large difference (several hundred Newtons) to the sagittal-plane force-feedback (especially F_Y, the A-P force component). For any individual test results, the dwell-point value can be 'tuned' to match; more generally the variability is better included as a factor within a probabilistic study framework.

- Friction: tibiofemoral friction is known to be important (e.g. see the results of the probabilistic studies in Chapter Five), and for the AMTI simulator the friction coefficient makes a particular difference to the F_Y force (and hence M_X moment also). However, the friction coefficient is test-specific; POD studies have shown friction is much higher for more complex motion paths [205]; therefore tuning a 'global' friction constant for any model is less than ideal. Nonetheless, early exploratory studies comparing the *in-silico* model with *in-vitro* data suggest that values at the lower end of the reported range seem to best match the experimental results ($0.01 < \mu < 0.02$).

- Friction from the roller bearings is also important; the M-L force (F_x) and Ad-Ab torque (M_y) are most sensitive to this factor. Experimentally, the bearing friction (ideally zero) is often surprisingly high (μ up to 0.1) and can considerably reduce the freedom of motion about the axial-load pivot.

- The 'pliancy' of femoral arm was also important; this seems to vary from station to station; values of 1kN/mm are typical, but variations in the range ±50% are needed to account for the experimental feedback data.

Because of the high experimental variability observed with all of these parameters, 'tuning' is necessary for any *specific* data-set. For the present data set, values were determined based on an iterative tuning process (using the localised sensitivity to gauge the correction factor required); the final values are listed in Table 18.

Parameter	Value
A-P dwell point	12.25 mm
Tibiofemoral friction coefficient	0.01
Roller-bearing friction (M-L)	0.06
Roller-bearing friction (V-V)	0.03
Femoral axis pliancy	1000 N/mm

Table 18: Values used for the S/C HIKIN experimental corroboration.

Using these values, good corroboration was achieved for all 6 axes of the load-cell feedback. Results are shown in Figure 117. Note that the experimental data is presented for all 6 stations running the same test – this immediately shows the high degree of experimental variability from station to station. This indicates that it would never be possible to achieve an 'exact' match with deterministic studies alone; a probabilistic study is the only way to corroborate the system given the variability present.

If the 'average' of these experimental values is taken, it is possible to report quantitative error levels (see Table 19). Note that although the percentages are high on some axes, (e.g. F_x) the actual absolute errors are low. This table must be interpreted with caution, given the inherent experimental variability – if the individual experimental feedback traces were compared to the averaged mean trace, many of them would appear to exhibit worse 'errors' than the computational model. Quoting error levels is of limited value when the system includes a high degree of uncertainty; once again, a more 'probabilistic' approach to corroboration is fundamentally necessary, given this variability.

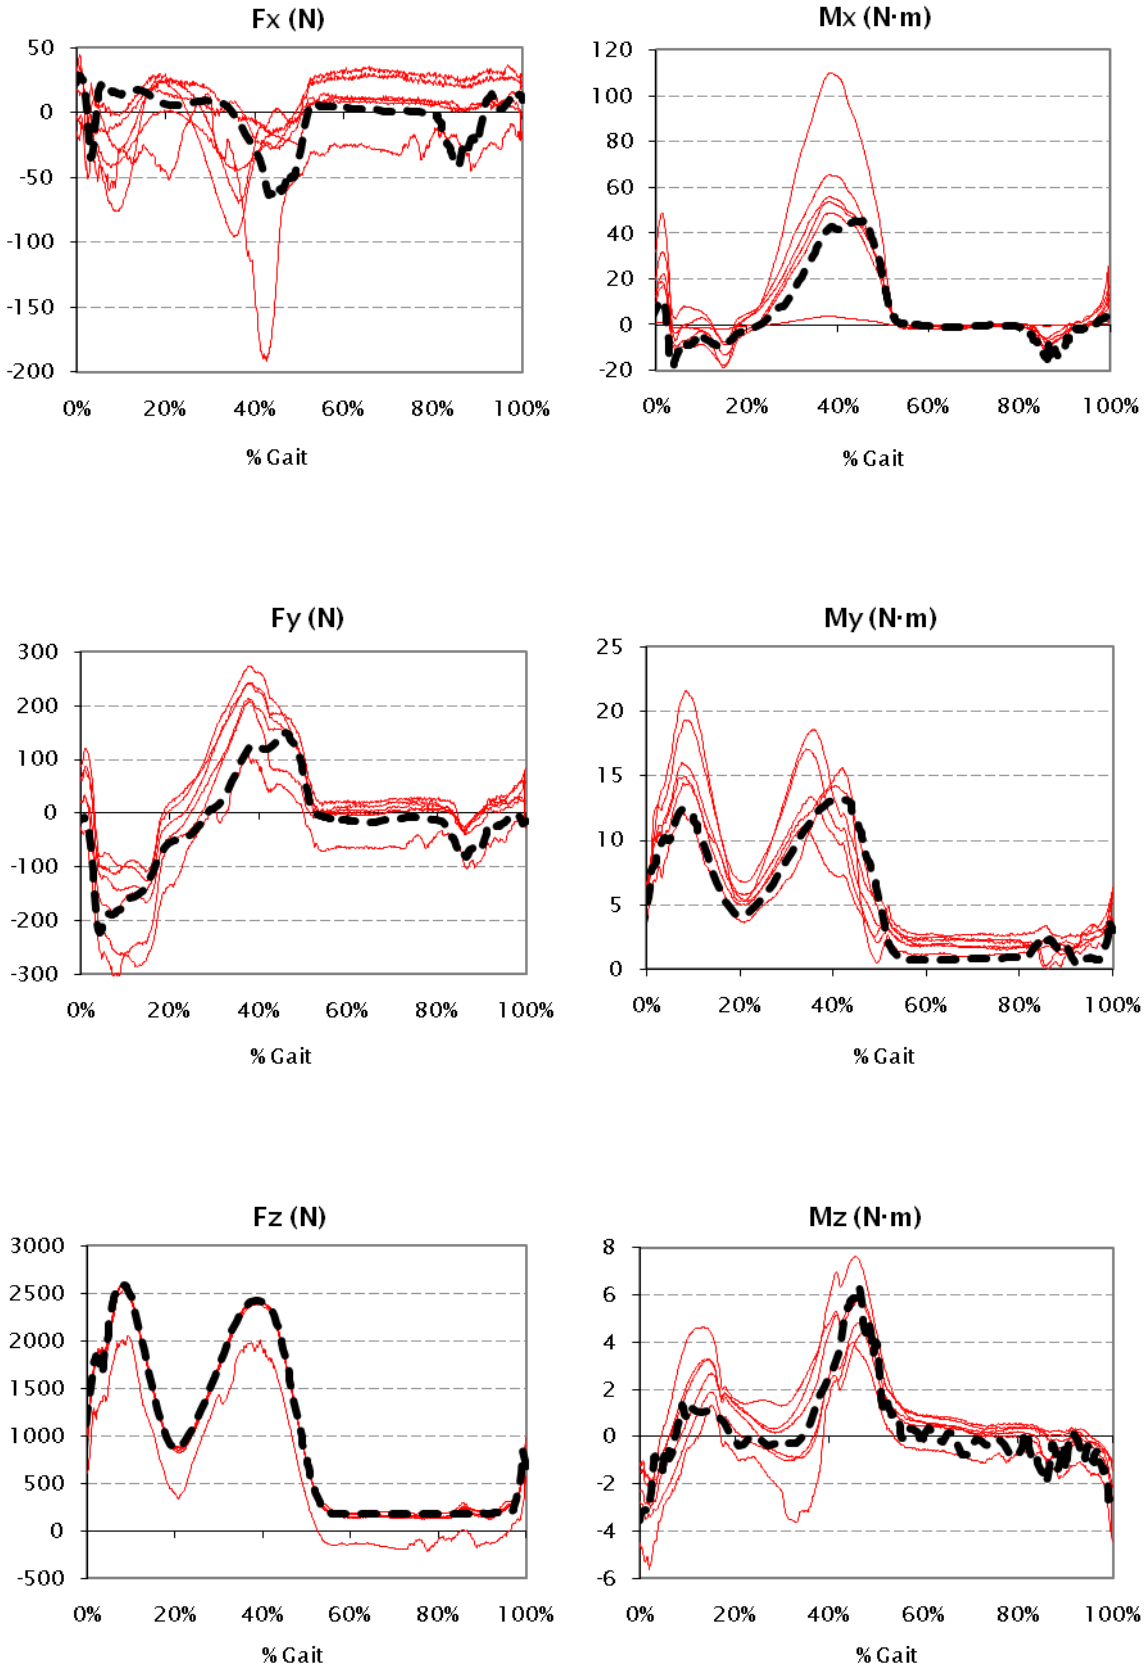

Figure 117: Corroboration for HIKIN profile with S/C insert – *in-vitro* (solid, N=6) versus *in-silico* (dashed) - forces (left) and moments (right).

Measurement	Cycle-averaged Absolute Error	Error as Percentage of Max. Absolute Value
F_X (M-L)	16.6 (N)	35.0%
F_Y (A-P)	30.8 (N)	14.6%
F_Z (I-S)	79.1 (N)	3.2%
M_X (about M-L)	5.14 (N·m)	9.2%
M_Y (about A-P)	1.68 (N·m)	10.3%
M_Z (about I-S)	0.48 (N·m)	8.9%

Table 19: Error levels in the first deterministic corroboration (S/C HIKIN).

Nevertheless, this deterministic corroboration represents an important step forward from the tests in Chapters Four and Five; the inclusion of instrumented force-feedback to compare across *in-silico* and *in-vitro* tests provides a robust extra degree of corroboration. This study has shown that the ADAMS-based computational model of the AMTI simulator is in good agreement with the spread of experimental results.

Once again, it is possible to use the new computational model to greatly enrich the data-set available from the rig alone. Figure 118 illustrates some of the additional data that can be retrieved *in-silico* which is not available directly *in-vitro*. Plots of contact area and M-L load split through the gait cycle, sliding distance distributions, sliding paths at individual nodes, contour maps for cross-shear and wear depth, and intra-cycle wear rate plots all provide an enhanced perspective on the test. Of course, the value of this *in-silico* dataset depends entirely on how representative it is of *in-vitro* conditions – hence, the importance of the more rigorous computational-experimental corroboration described above.

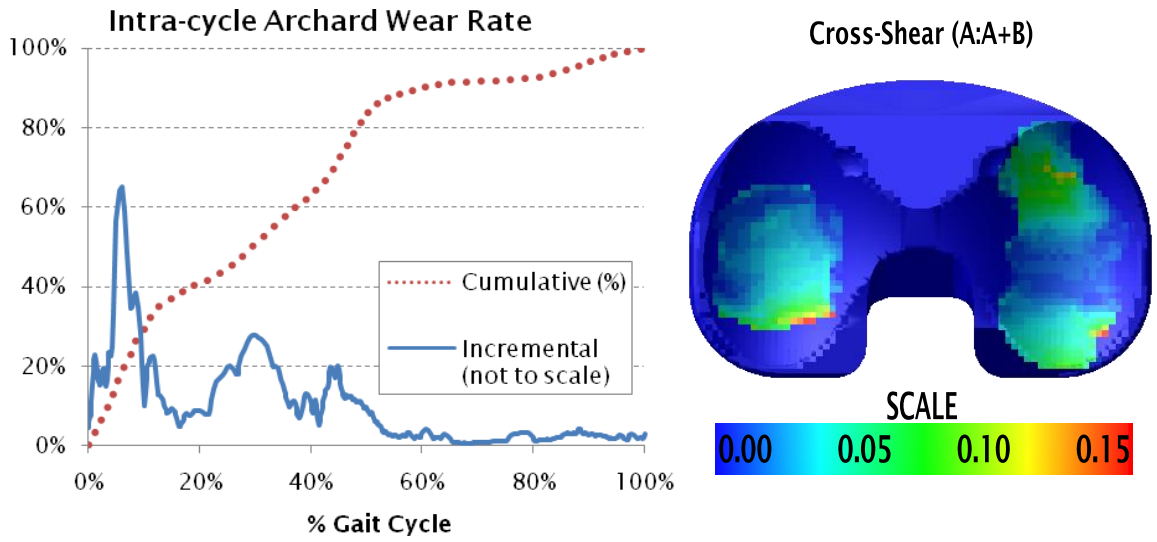

Figure 118: Additional visualisation metrics are available *in-silico* to enrich the overall pre-clinical analysis process. Here: intra-cycle wear rate (left) and cross-shear contour map (right).

7.2.3. DD Corroboration Test 2: 'Femur-On-Flat' Gait Test

The S/C HIKIN test demonstrates that the model can be tuned to a specific test-case, and the results achieve good corroboration. However, for a different test, it is recognised that variations in the implant design and procedure might lead to variations in component positioning. Therefore, the same precise 'tuned' values would not be applicable. Instead, a 'femur-on-flat' experimental test was selected for a second corroboration study. This consists of a regular S/C femoral component, articulating against a 'flat' polyethylene insert (i.e. with no condylar 'cups', such that there is no geometric conformity). Femur-on-flat studies are currently being used to investigate wear behaviour under extreme ranges of contact pressure, to better understand the mechanics of wear (e.g. see [178]). The advantage for present purposes is that theoretically there is no 'dwell' point, as the flat surface means there is zero conformity. This removes one of the most influential variables, making 'tuning' of the model less critical.

Experimentally, the inputs for this test are identical to the 'HIKIN' profile, except that axial loading is scaled up from a peak value of 2,600N (~600lb) to 3,600N (~800lb). The high loading, coupled with low conformity of the flat 'insert', results in extremely high contact pressures concentrated on a small area of polyethylene (this makes the test of interest to wear theorists). The test was run as previously, but replacing the S/C insert with a flat alternative, (once again discretised into 1mm^2 cells as described in Chapter Four). The simulation was run with a 0.1 second 'ramp-up' into the profile, followed by the 1.0 second profile itself, using the 'tuned' values described in the first corroboration test. Initial results did not corroborate as well as hoped; on further investigation a number of issues were identified:

- The high pressures in this test made the assumption of linear elasticity (inherent in the rigid-body model) inadequate: considerable plastic deformation was apparent on the *in-vitro* implants after testing. However, this could be accounted for; modified implants were generated, based on surface profiles of the actual experimental samples, which featured this deformation effectively 'built in'. These were used for subsequent modelling. Note that this now meant A-P dwell position was an important factor once again, since the sagittal profile of the insert was no longer 'flat', but deformed.

- The feedback data had considerable 'noise' which was more apparent for the small-magnitude forces & moments in this test; pre-conditioning to smooth the higher derivatives of the input waveforms helped to reduce this artefact.

- It was evident that, even with these refinements, the corroboration was not perfect; there was a minimal degree of malpositioning (under 1° on the I-E and V-V axes) which was having a small, but noticeable influence on the force-feedback. If this was accounted for (by correspondingly aligning the components in the *in-silico* model) slightly better accuracy could be achieved (although this effect was minor compared to the surface deformation and 'dwell point' issues).

In Figure 119, results are shown for the load-cell feedback channels, for all six axes. Note that more channels of *in-vitro* data were available for this test.

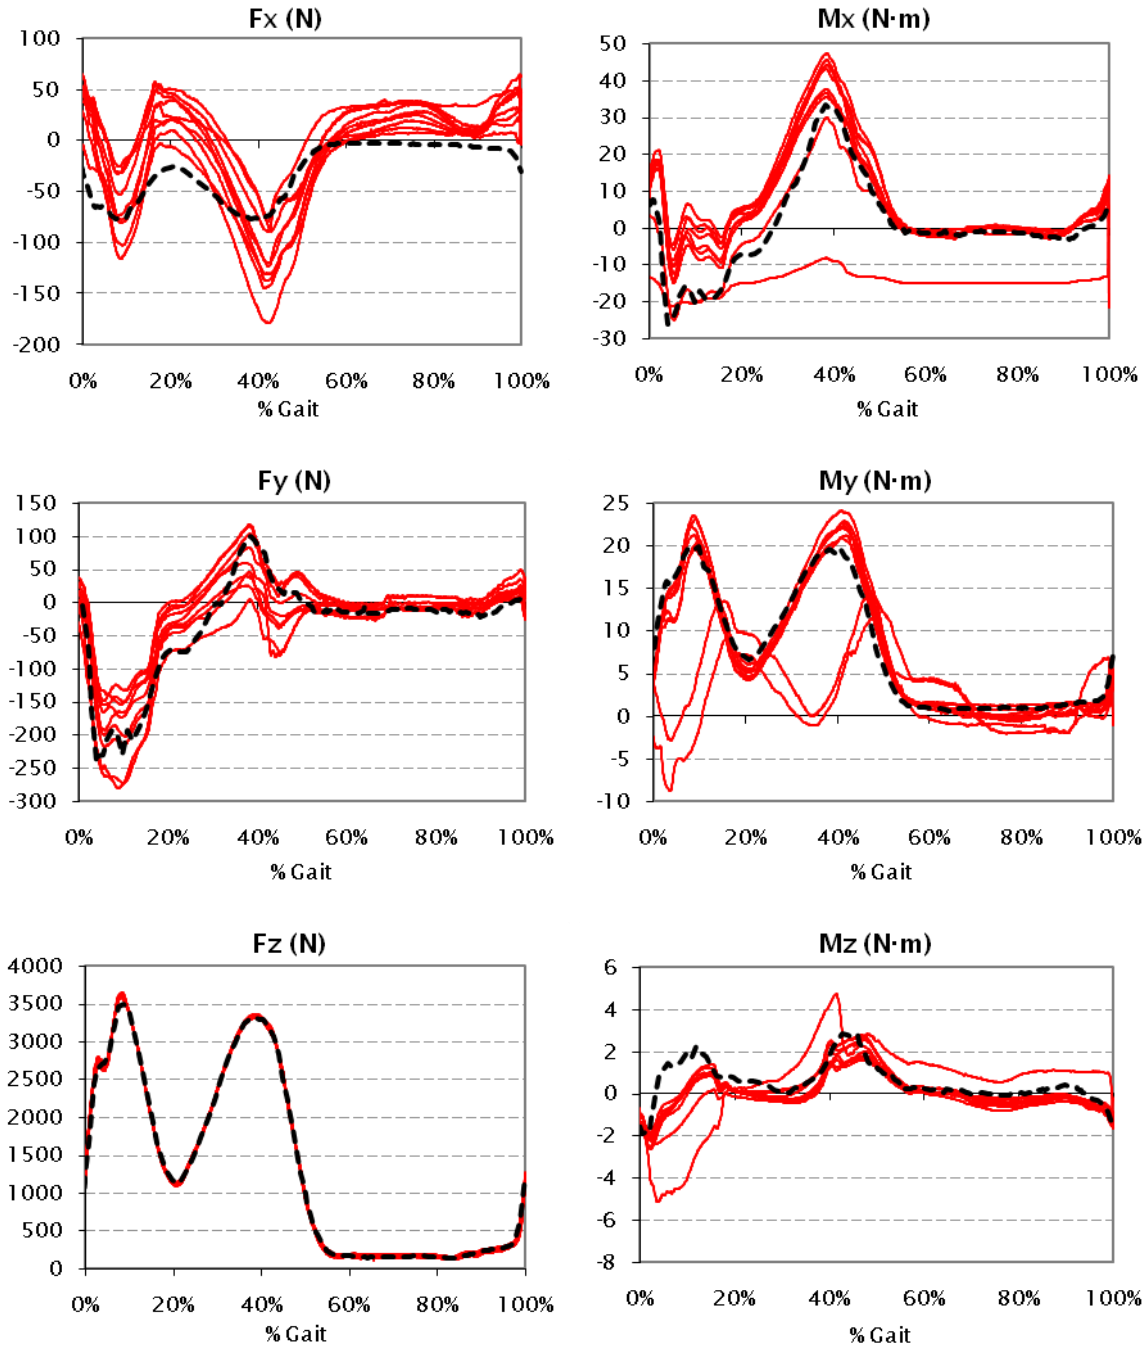

Figure 119: Corroboration for 'femur-on-flat' test – load-cell forces (left) & moments (right) *in-vitro* (solid, N=9) versus *in-silico* (dashed).

Note that (contrary to initial expectations) the predicted and observed shear forces are actually substantial (several hundred Newtons for F_X and F_Y); this is in large part due to the increased degree of conformity induced by plastic deformation experimentally (and the corresponding conformity introduced when this plastic deformation was modelled approximately in the *in-silico* simulation).

The corroboration is not perfect, but again for the experimental data, there is some variability between stations (and it is apparent that some stations are 'outlying' very noticeably from the others), so once again, no single-shot 'deterministic' simulation could match this spread of experimental results.

Once again, a table of 'error levels' was compiled, based on comparison to the experimental 'averaged' values (Table 20). Error levels are comparable to the first S/C HIKIN corroboration test; given the additional challenges presented by the femur-on-flat test, this is a positive result. Note that once again, the highest errors are in F_x, the M-L shear force (which is however uncontrolled, and so less critical).

Measurement	Cycle-averaged Absolute Error	Error as Percentage of Max. Absolute Value
F_X (M-L)	31.4 (N)	26.3%
F_Y (A-P)	20.9 (N)	10.8%
F_Z (I-S)	18.4 (N)	0.5%
M_X (about M-L)	4.15 (N·m)	12.1%
M_Y (about A-P)	1.84 (N·m)	9.8%
M_Z (about I-S)	0.53 (N·m)	23.6%

Table 20: Error levels in the second deterministic corroboration (fem-on-flat).

Note that further corroboration studies were performed, which are not reported in this thesis for brevity (including further femoral-on-flat tests, and various gait tests using the S/C and U/C inserts).

The most pertinent conclusion of these different studies is that, in every case, experimental variability means that various factors must always be 'tuned' to achieve the best possible corroboration. As such, it is always possible to raise the question whether this tuning is legitimately accounting for *in-vitro* experimental variability, or in fact compensating for deficiencies in the *in-silico* model. The best way to address this is with a full probabilistic corroboration; this will be addressed subsequently in the present chapter.

7.3. Force-Driven Modelling: Corroboration

7.3.1. AMTI Control-System Modelling

The most important difference between DD and FD operation is that the force-driven method requires a more sophisticated control system, to mimic the effect of the virtual spring restraint. The controller works by measuring A-P translation / I-E rotation, calculating the restraint force/torque at this level via a spline-based 'look-up table' (LUT), and superimposing this force/torque onto the input waveforms for A-P force / I-E torque. This is illustrated conceptually in Figure 120. The advantage of this configuration is that the 'springs' are defined only in software; therefore it is necessary to change only the data splines to alter the spring characteristics. This is both faster than physically replacing spring buffers on the *in-vitro* rig, and also allows any particular non-linear spring relationships to be defined (rather than needing to source physical springs with appropriate stiffnesses).

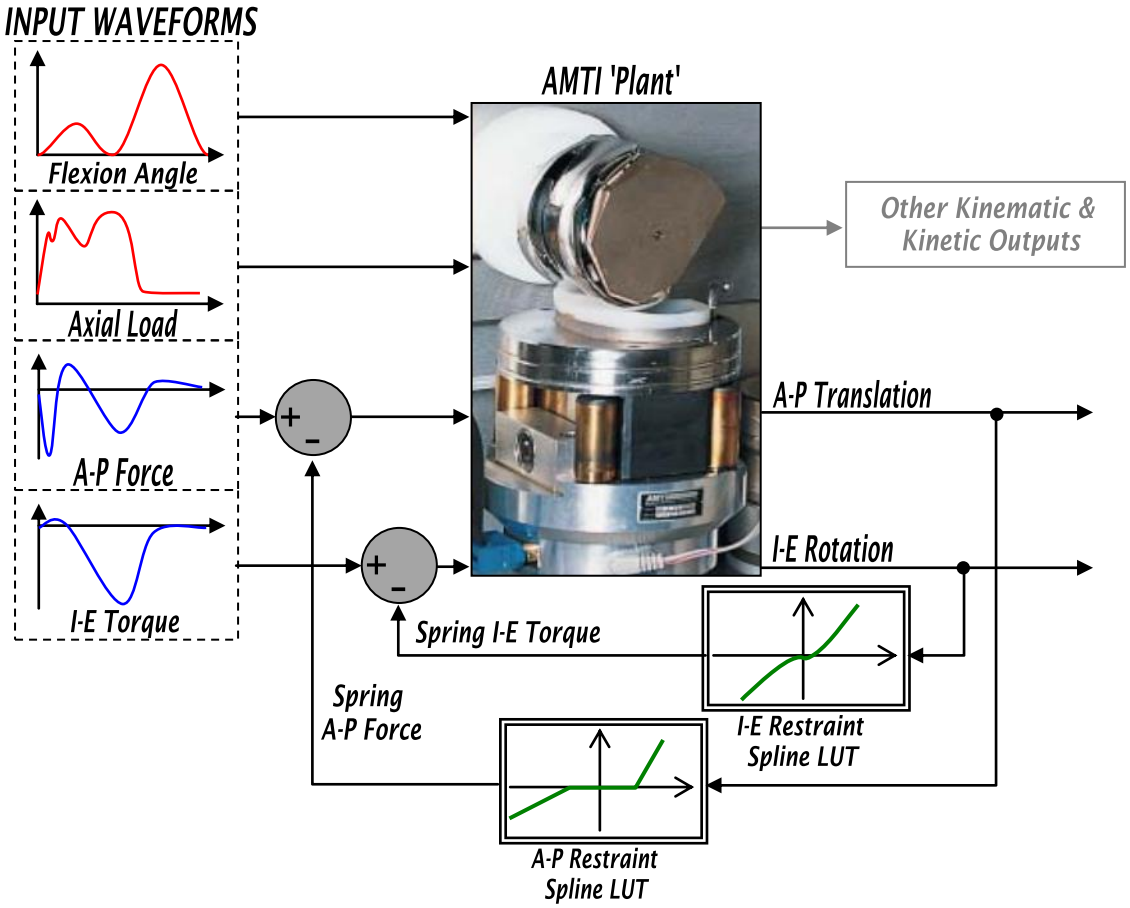

Figure 120: Conceptual AMTI FD 'virtual spring restraint' operation schematic.

For the specific model being corroborated in this study, the springs were based on the work of Haider et al [158] (who proposed a combination of 'soft' and 'hard' springs to better represent the *in-vivo* effect of resecting the ACL but retaining the PCL). The *idealised* load-displacement relationship is shown in

Figure 121. (Note that in practice the model does not precisely mimic this relationship – see later discussion and Figure 127).

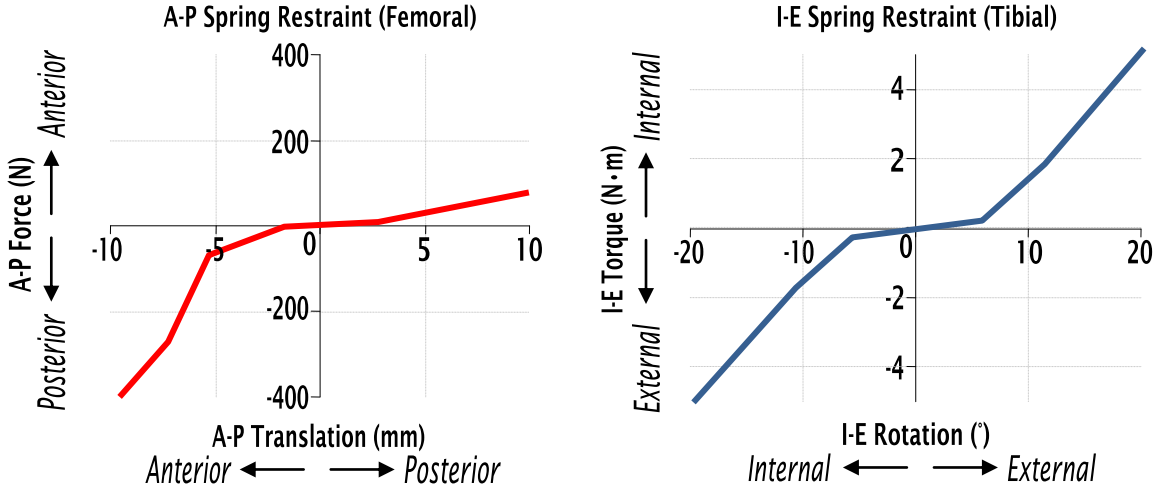

Figure 121: Spring restraint splines for the AMTI model (derived from [158]). I-E restraint is symmetric, with ±5° 'dead zone'; A-P restraint is asymmetric (to represent resected ACL & retained PCL) with ±2.5mm 'dead zone');

As already stated, on the AMTI rig the A-P motion is applied to the femoral component. This introduces a slight complication, since the force-feedback load cell is mounted beneath the tibial insert. Consequently it is important to accurately model the dynamics of the system, to capture the influence of inertia between the applied force (femoral side) and feedback (tibial side). Figure 122 illustrates the sagittal mechanics (considering A-P force components only, and neglecting angular or non-sagittal components of loads and displacements).

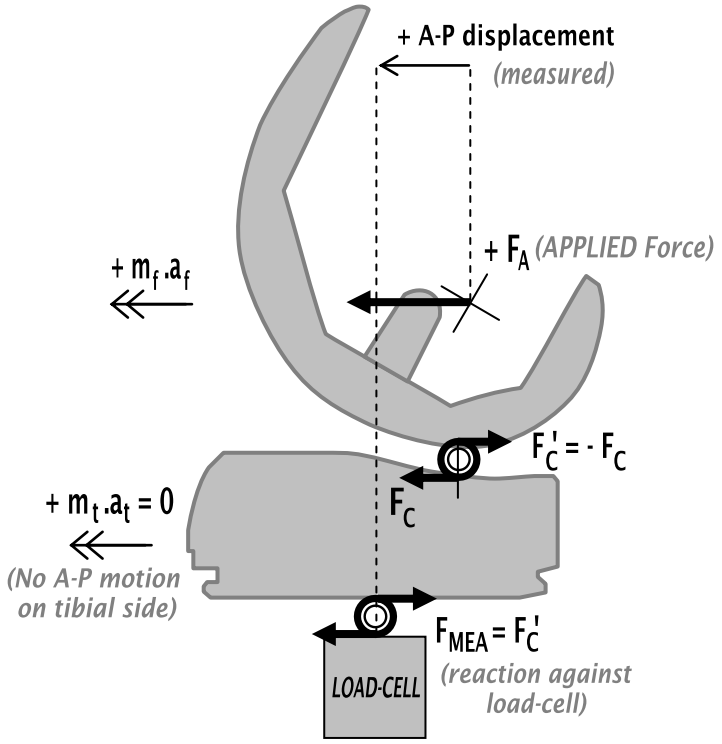

Figure 122: Sagittal plane A-P forces for the AMTI simulator.

Resolving by applying Newton's 2nd Law for the femur & tibia gives, respectively:

$$\text{R}(\leftarrow) \quad F_A - F_C' = m_f.a_f \qquad (11)$$

$$F_C - F_{MEA} = 0 \qquad (12)$$

Where F_A is the applied force, F_C is the A-P component of the contact force acting on the tibia, and F_C' is the reaction force to F_C, acting on the femur. F_{MEA} is the 'measured' force; i.e. the force 'fed back' by the load-cell beneath the tibia. The terms m_f & m_t designate the mass of the femoral and tibial components & their associated mounting jigs; the terms a_f & a_t are the corresponding accelerations (note that a_t is zero, since the tibial component has no unconstrained A-P DOF). The desired force on the tibia, denoted F_D, is the combined sum of the input ISO-derivative waveform (denoted F_{ISO}), and the superimposed spring restraint force, F_{SPR}, which is a function of the measured A-P translation, Δ_{A-P}

$$F_D = F_{ISO} + F_{SPR}(\Delta_{A-P}) \qquad (13)$$

If the control system was based only on displacement-feedback (i.e. the Δ_{A-P} measure), the only option would be to set the applied force, F_A equal to the desired force, F_D. However, it is clear that the measured force at the tibia, F_{MEA} would not then be equal to the desired force:

$$F_A = F_D \qquad (14)$$

$$F_D - F_C' = m_f.a_f \qquad \therefore F_C' = F_D - m_f.a_f \qquad (15)$$

$$F_C - F_{MEA} = 0 \qquad (16)$$

$$\therefore F_{MEA} = -F_C = F_C' = F_D - m_f.a_f \qquad (17)$$

The inertial effect of the femoral component A-P acceleration would result in a discrepancy between the achieved (measured) and desired A-P force profile. Instead, the controller must include closed-loop feedback of the measured force, in order to achieve $F_{MEA} = F_D$. (Note: the above analysis neglects pliancy of the F-E shaft also; this would introduce a further complication in reality)

As with the control systems described in Chapter Six, external co-simulation using MATLAB/Simulink was chosen for control-plant modelling. A custom controller was constructed, based upon the same design concept as the *in-vitro* AMTI controller, but including a number of additional features and visualisation

tools for enhanced diagnostic functionality. A screen-shot of the GUI designed for this purpose is shown in Figure 123.

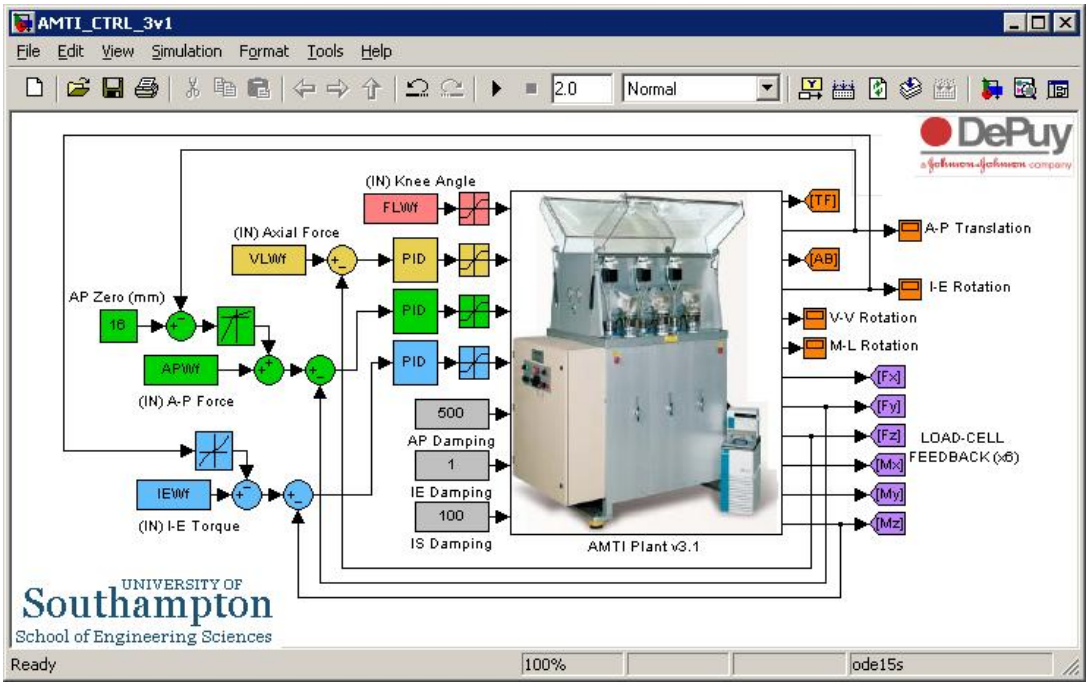

Figure 123: Custom AMTI Controller modelled in Simulink.

The precise details of the experimental control system are commercially confidential, but it is known that the *in-vitro* system uses a proprietary adaptive control system. The PID-based model used here is therefore only an approximation to the real AMTI controller; since the actual *in-vitro* system has comparatively good accuracy, compared to other commercial rig designs, the *in-silico* values are tuned to achieve the best possible tracking for the conditions under test. Typical values are given in Table 21 (note that the precise settings can vary from case to case, so the values given are only a representative 'starting point' for subsequent refinement in any specific test-case; generally, the more conformal the implant design, the stronger the controller settings must be).

	Vertical Load	A-P Force	I-E Torque
Proportional (P)	10	0.1	1
Integral (I)	3	0.1	1
Derivative (D)	1	-0.01	0

Table 21: Typical PID values for a force-driven AMTI gait test.

7.3.2. FD Corroboration Test 1: Isolated Axis Tests

With the model and controller developed, the first corroboration tests were attempted. Early efforts to corroborate full force-driven ISO gait were

unsuccessful (kinematics for A-P and I-E did not match the experimental results). Therefore, a more incremental approach was taken.

A series of tests were devised, which 'deactivated' selected inputs of a standard force-driven gait test, to simplify the ensuing mechanics. Axial compressive load was always included for stability, but the other three input channels (flexion angle, I-E torque and A-P force), were each analysed in isolation from the others (Table 22). These three 'isolation tests' allowed the behaviour of each axis to be studied without the confounding effect of influences from the others.

These tests were run by our industrial collaborators for the purpose of comparing the computational model with *in-vitro* data (note: tests were run 'dry', instead of running under lubrication for many millions of cycles as required for wear assessment: obtaining kinematics is much easier, and requires only a handful of cycles to remove initial transients. However this does of course alter the operating friction). For all the test-cases the virtual springs provided a simple linear restraint (30N/mm for A-P and 0.6N·m/°).

Test Case	Vertical Load	F-E Angle	A-P Force	I-E Torque
1. 'VL-FE'	ISO14243-1	ISO14243-1	constant 0N	constant 0N·m
2. 'VL-AP'		constant 0°	ISO14243-1	constant 0N·m
3. 'VL-IE'		constant 0°	constant 0N	ISO14243-1

Table 22: Inputs for force-driven 'isolation tests'.

The corroboration revealed the importance of dynamic effects for the AMTI rig under FD control. A large degree of damping was necessary (especially on the A-P axis) to accurately match the experimental kinematics. This experimental damping is believed to be due in part to the construction materials (e.g. the tibial platen is constructed of a relatively 'soft' polymer), in part due to damping in the system hydraulics, and in part due to other sources of pliancy within the mechanics of the rig.

Based on 'tuning' to these cases, a reasonable match was obtained between the computational and experimental results, for both kinematics (A-P & I-E) and kinetics (the load-cell feedback). Figure 124 shows the kinematics for all three tests (note, *in-vitro* data was only collected for one station, so unfortunately no indication of experimental variability is available for these isolation tests).

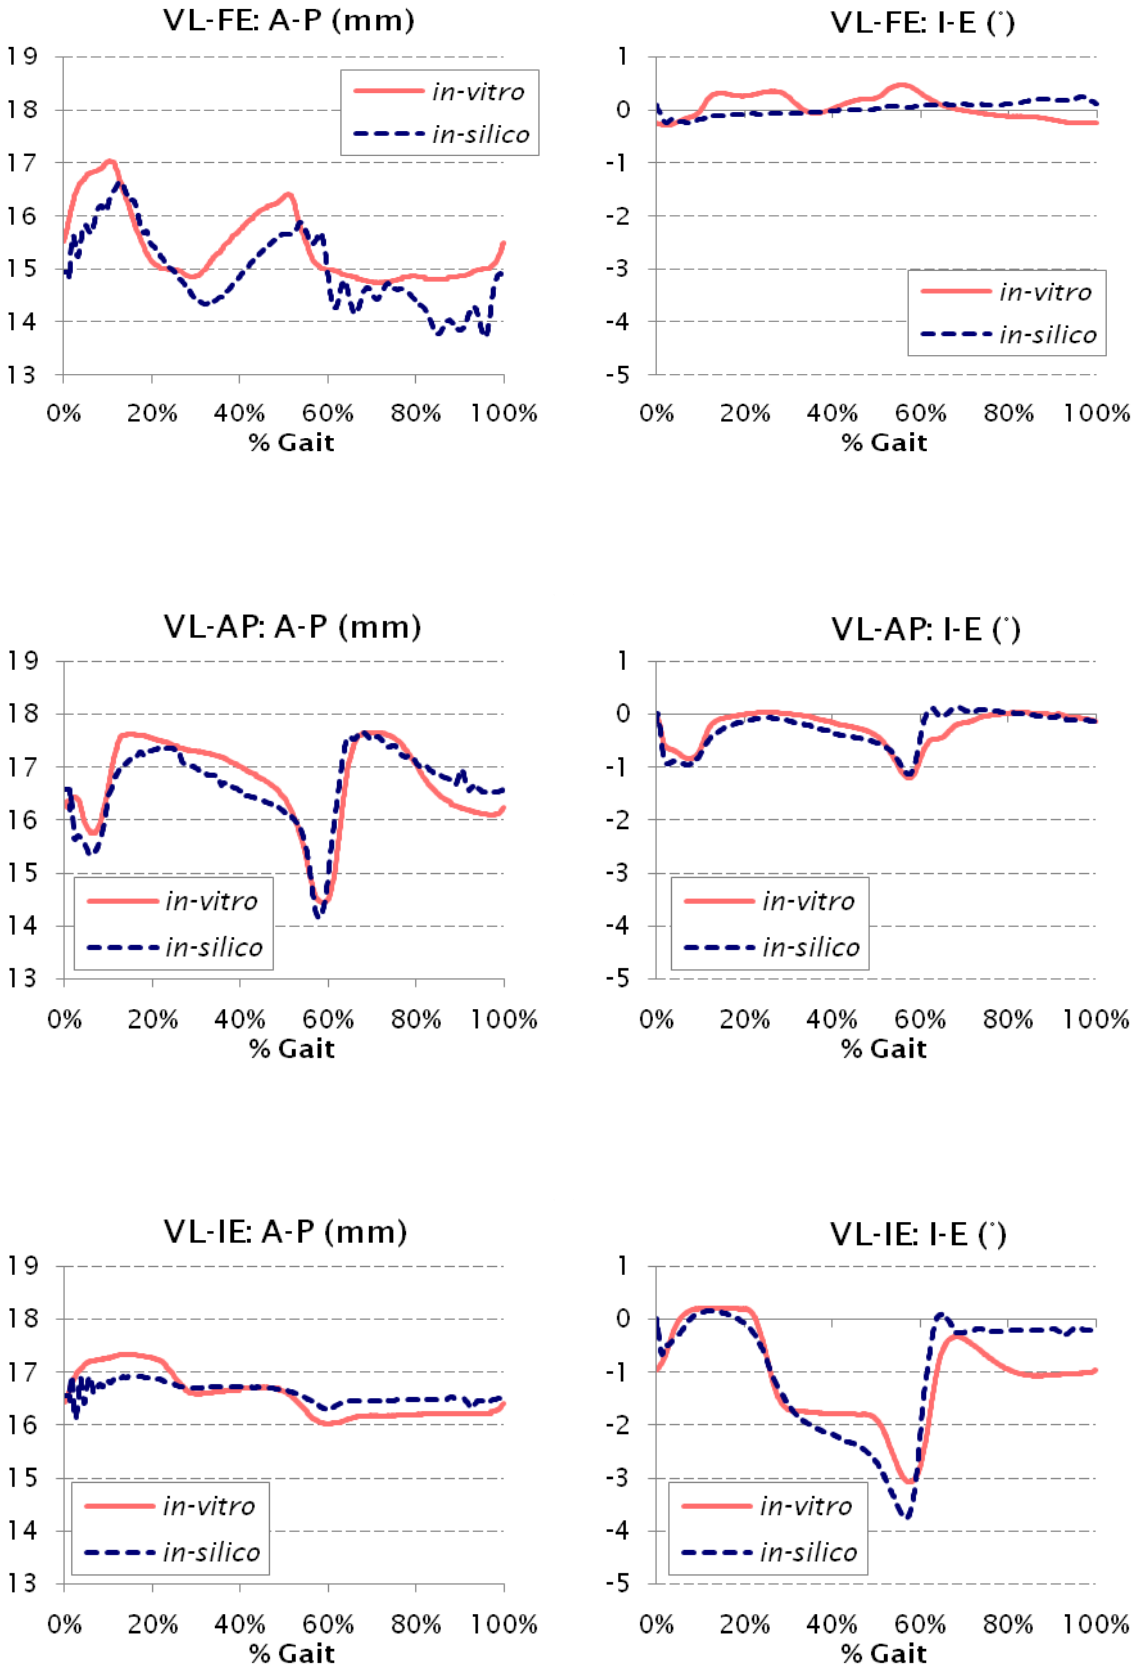

Figure 124: Kinematic feedback for all three isolation tests.

Figure 125 shows the tri-axial load cell force-feedback for just the final test case (VL-IE); again data is unfortunately only available for the first station in the bank, so it is not possible to determine how representative this single experimental data-set is, or if it was in fact an unrepresentative 'outlier'.

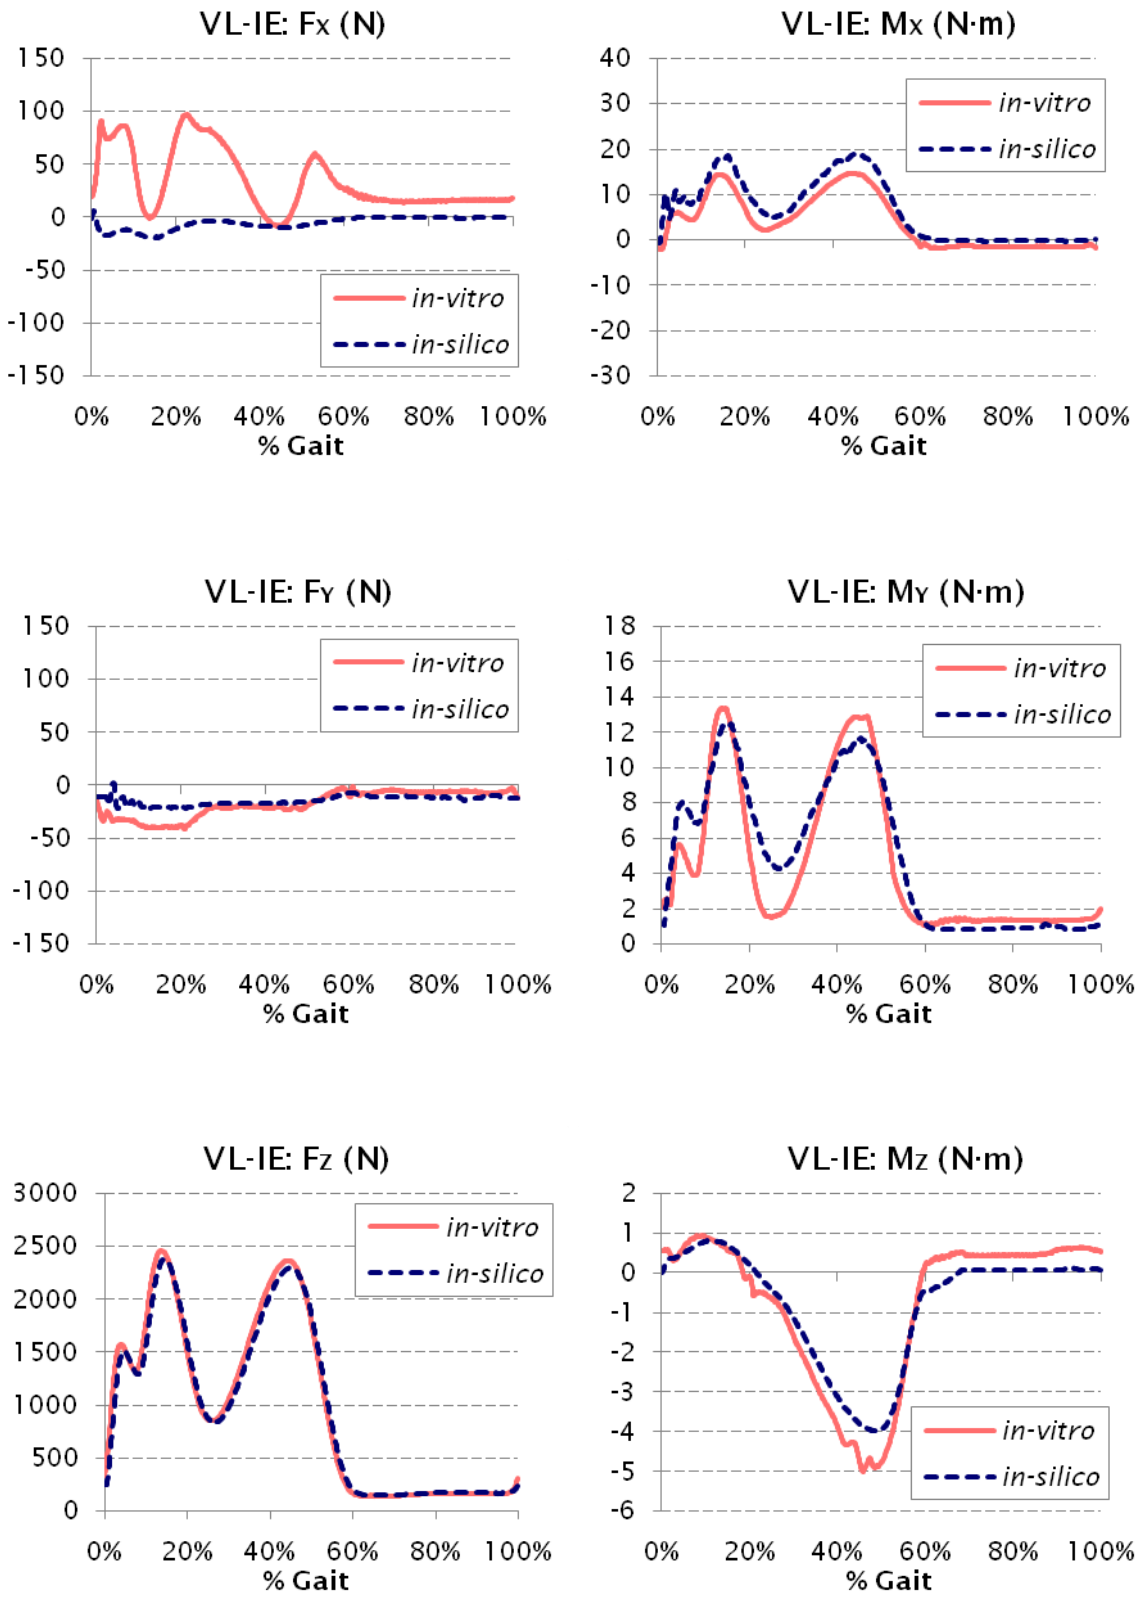

Figure 125: Load-cell feedback for isolation test 3 (VL-IE).

In the event, corroborating the force-driven model proved considerably more challenging that corroborating the displacement-driven tests. There were a number of reasons for this; firstly, limited data was available. Force-driven tests were relatively new for the AMTI platform, so there was not a large volume of historical data available. In these isolation tests, data was only available for a single station, so it is impossible to know if these individual waveforms were

representative (close to the average) or in fact outliers; previous tests have shown that experimental ranges of uncertainty are quite large, and without any indication of the experimental variability in this case the quality of the corroboration cannot be assured. Additionally, the control-system tracking of the rig itself was not as good under force control as under displacement control – this is apparent from analysis of the original *in-vitro* data. The Simulink model was designed to track the 'ideal' inputs so does not account for the experimental tracking errors in forces. (It is not possible to easily factor in tracking errors, due to the confounding influence of the 'virtual spring' system). Further, in this particular case, some of the forces and moments for the isolation tests are inherently quite small (e.g. in the VL-IE test, there is very little A-P shear force); this makes the signal-noise ratio unfavourable, so once again corroboration is difficult. Despite these challenges, progress has been made in identifying influential factors for force-controlled testing.

7.3.3. Corroboration Test 2: Full ISO-derivative Gait

Having tuned the dynamics of the individual axes using the isolation tests, a full FD gait test was next modelled. Unfortunately, the only available experimental data suitable for corroboration was a relatively early data-set – an initial 'benchmarking' test of the S/C design following the commissioning of the force-driven rig upgrade. Feedback data for this test was therefore used to undertake a computational simulation. The inputs (Figure 126) were similar to the ISO-standard [23], but were slightly phase-shifted relative to one another (e.g. the flexion waveform seems to be phase-delayed by ~10% cycle).

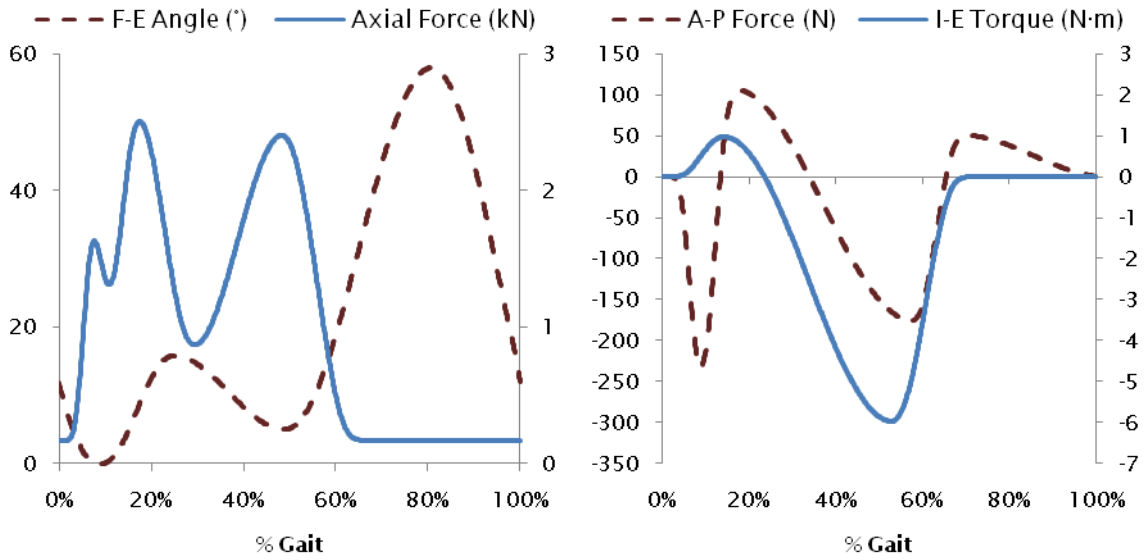

Figure 126: Inputs for ISO-derivative FD gait corroboration test. Note there are different phase shifts for the different waveforms; e.g. F-E is delayed by ~10% cycle, whereas axial force is delayed by only ~3%.

This test was conducted under lubrication *in-vitro,* so the high 'dry' friction values used for the isolation tests (0.07 to 0.1), were scaled down to 'wet' values (0.01 to 0.02). The virtual springs were configured to simulate soft/hard springs with a spring-gap (the actual feedback splines are shown in Figure 127). Note that the achieved feedback does not match the ideal relationship shown earlier in Figure 121; there is a degree of hysteresis, and the spline-interpolation sometimes leads to 'positive feedback' where the restraint force acts to increase the kinematic offset. This will cause differences in performance compared to the ideal spring relationship – therefore the *in-silico* model was based on the experimental (rather than theoretical) splines, to better match the resulting *in-vitro* data. The kinematics (A-P and I-E) and force-feedback were compared to the limited available experimental data (N=3 channels).

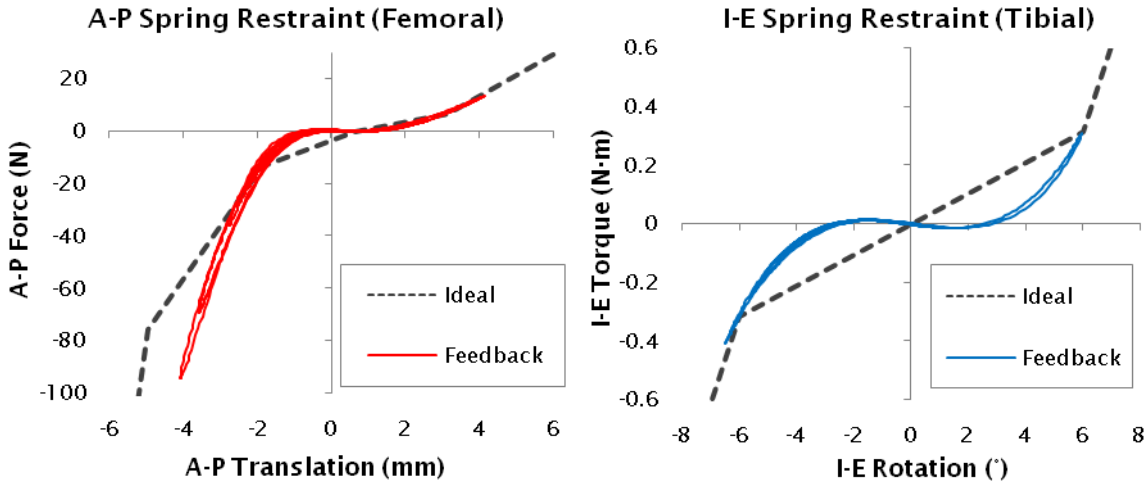

Figure 127: Actual (feedback) splines for virtual spring restraint: A-P (left) & I-E (right). For comparison, the 'ideal' relationship is also shown (c.f. Figure 121).

The results of this test did not corroborate as well as hoped. Whilst reasonable agreement was achieved for the A-P translation, the *in-vitro* I-E rotation could not be reproduced *in-silico* (Figure 128).

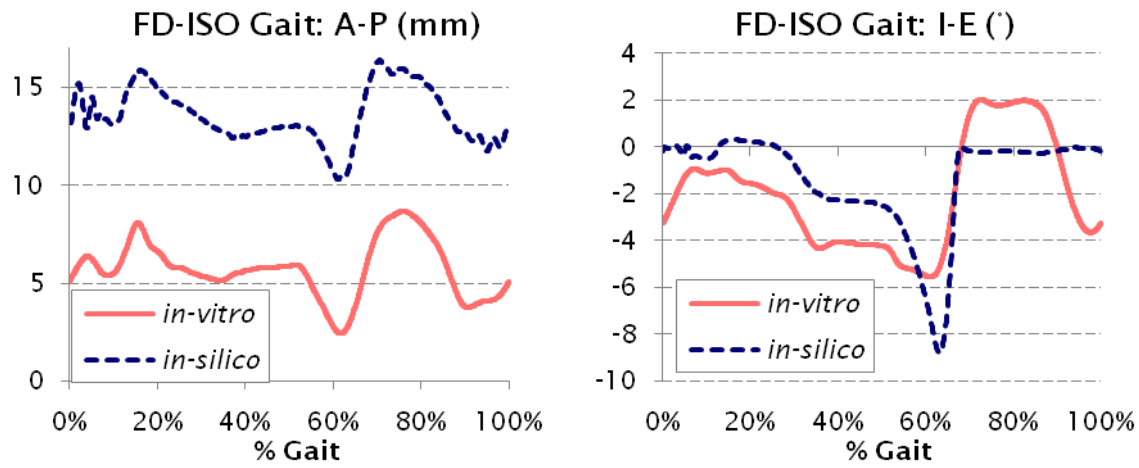

Figure 128: FD ISO-gait kinematics – *in-vitro (solid)* versus *in-silico (dashed).*

For the experimental load-cell data, considerable variability was observed (even within just 3 stations of available data); this made it difficult to determine how representative these samples were. The computational waveforms matched reasonably, except for the F_Y and M_X channels (Figure 129). The fact that M_Z (the I-E torque) is in good agreement with the experimental data, but the actual I-E rotation is so different, suggested that there was some considerable pliancy or motion between the load-cell and the insert itself.

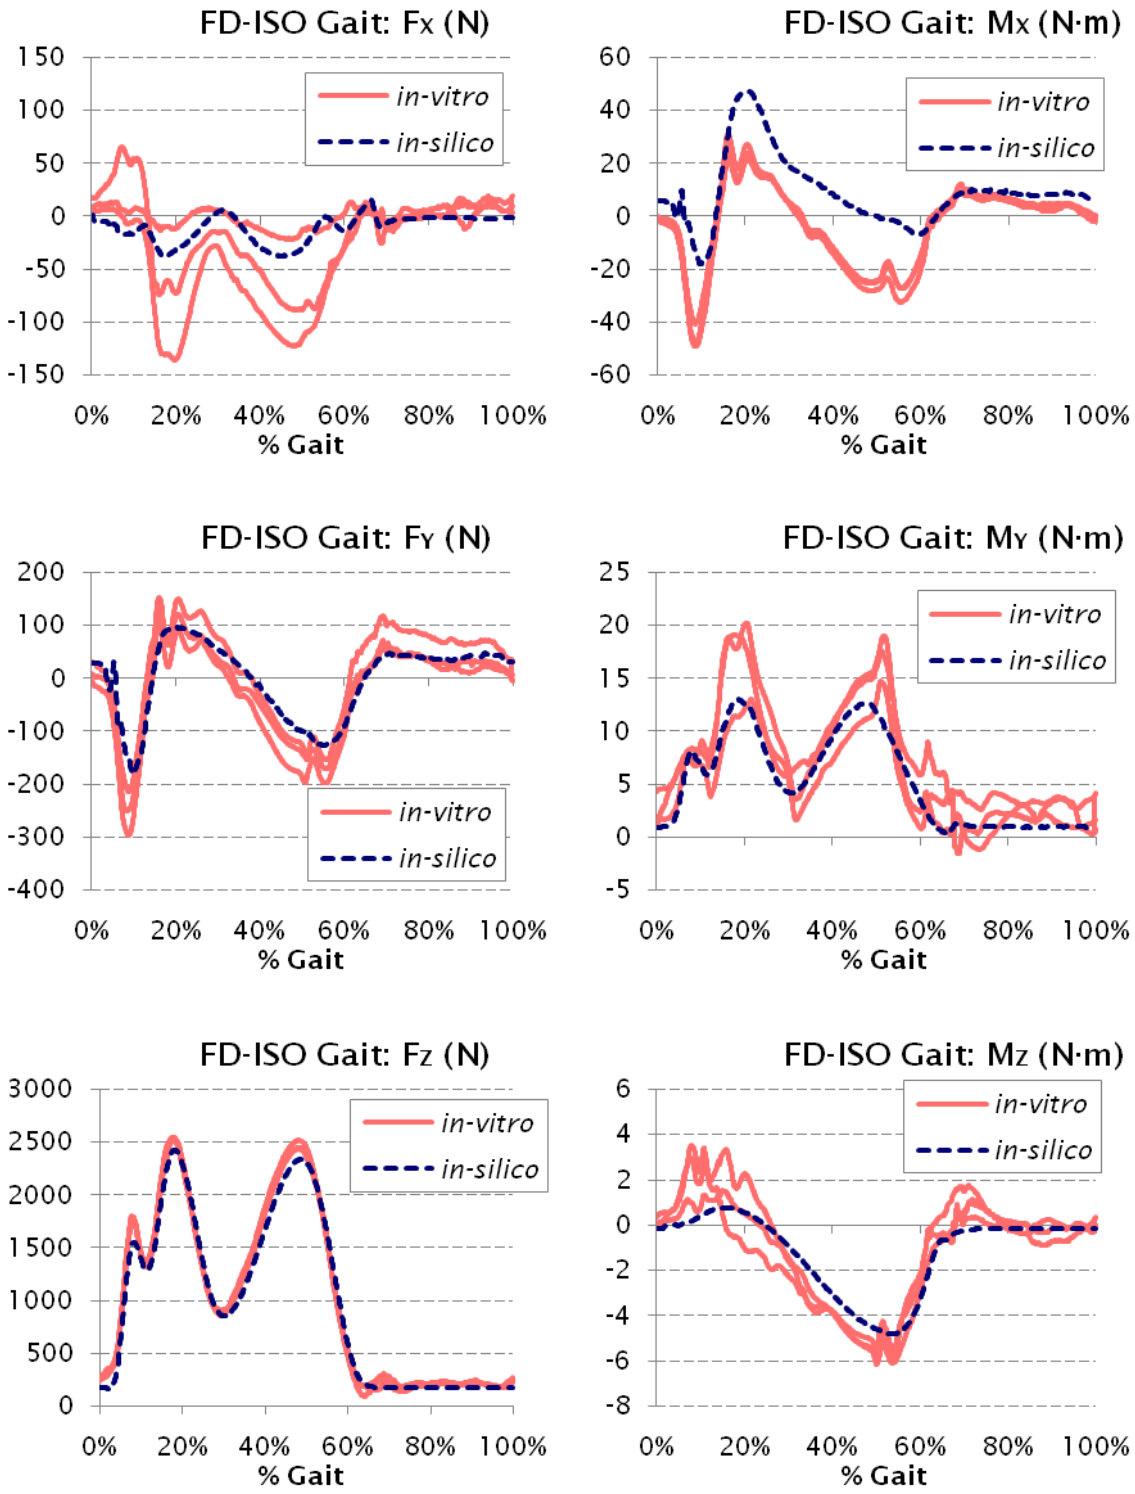

Figure 129: FD gait load-cell feedback; *in-vitro (solid, N=3)* vs. *in-silico (dashed).*

Several factors were hypothesised to be playing a role in the differences observed; the fact that the different inputs were out of synchronisation with one another, the imperfect tracking of the *in-vitro* controller, unaccountably high damping, or backlash/pliancy effects still not correctly modelled computationally. It was apparent further investigation would be required to better corroborate this study. Unfortunately, the data was quite old, and details of the precise experimental procedure have been lost. Further, some of the hardware on the rig had been changed since the original test (e.g. the tibial platen was entirely re-designed). It is therefore not possible to precisely re-create this test, to determine if changes in methodology (e.g. component positioning, or fixed-axis location) accounted for some of the differences in results.

Recently, the test conditions were re-created on the AMTI rig, using the new hardware configuration. The results of this comparison were very different to the earlier test (see the plots of A-P and I-E kinematics in Figure 130). Given that there is such variation in the *in-vitro* data, it was apparent that an accurate corroboration *in-silico* would not be possible at this stage; further experimental data will be required. There are, however, some important observations from this attempt: firstly, any *in-silico* corroboration study depends on quality, consistent *in-vitro* data to be robust and effective. Secondly, variations in methodology can have a major role in experimental results, and the corroboration effort must attempt to include consideration of these 'hidden' effects.

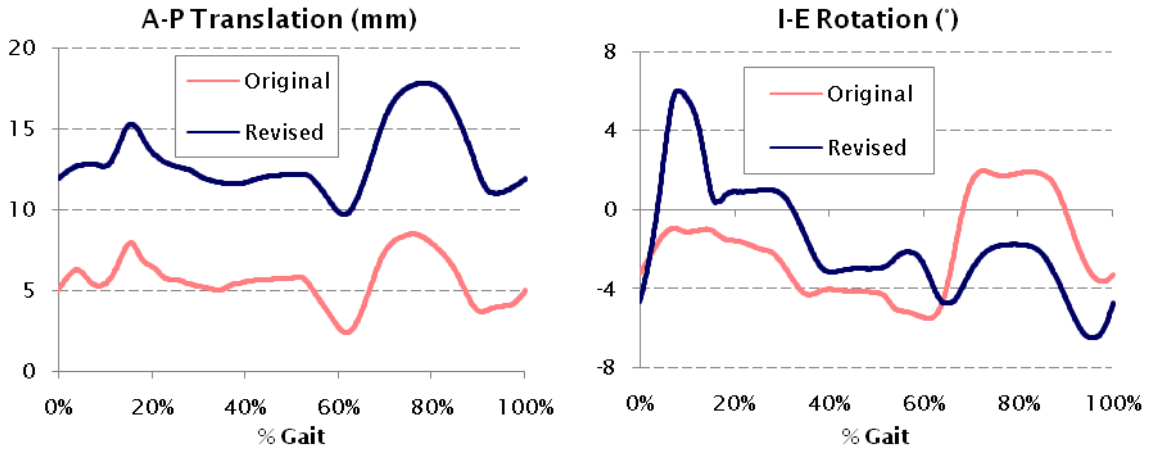

Figure 130: Revised data for the FD S/C ISO-gait test, compared to the original data. Left: A-P translation; right: I-E rotation.

For both the displacement- and force-driven models, variability has been revealed to be a key concern. The aim of the final study was to address this more exhaustively, by attempting the first probabilistic corroboration between *in-vitro* and *in-silico* experimental results of TKR testing.

7.4. Corroboration of Probabilistic Methods [9]

7.4.1. Study Structure (Methods)

The limiting factor in corroborating probabilistic methods is of course the availability of a large enough body of *in-vitro* data. An analysis of historical data available from previous industrial testing revealed that the most suitable data-set was for displacement-driven testing of the S/C design with high-kinematics (for conventional PE). There were still only a handful of comparable tests in total, but given the number of stations running simultaneously (typically 3 – 6 per test), and the number of intervals the test was run on for (typically 10 – 12 per test), it was possible to source >100 data points for force-feedback and interval wear rates.

The *in-silico* study was tailored to match this data. The experimental set-up is as per Section 7.4.2 (the deterministic corroboration). The model was parameterised with several factors previously identified as influential during the 'tuning' phase of the corroboration studies, together with additional malpositioning factors. In some cases statistical properties (mean, SD) could be based on available data (e.g. for the A-P dwell position, the *in-vitro* feedback data available indicated the degree of variation in the initial offset of the A-P waveform). In other cases, variability was assigned based on engineering judgement, or else estimated from the variability already observed in the earlier corroboration study (some indication of the variability could be obtained based on the N=6 samples from the deterministic study in Section 7.2.2 - whereas now a total of N=128 samples are available). The variables, and their assigned values, are listed in Table 23. The Normal distributions are cropped at ±3σ; the Lognormal distributions (which cannot be less than zero by definition) are cropped only at +3σ.

Factor	Description	Dist. Type	Mean	S.D.
AP_Dwell	Initial A-P Dwell offset	Normal	12.5mm	0.5mm
Fem_FE	Femoral F-E malrotation	Normal	0°	0.5°
Fem_IE	Femoral I-E malrotation	Normal	0°	0.5°
Fem_VV	Femoral V-V malrotation	Normal	0°	1°
Tib_ML	M-L offset of insert on platen	Normal	0mm	0.5mm
TF_μ	Tibiofemoral Contact Friction	Lognormal	0.01	0.02
Roll_μ	Roller-bearing Friction	Lognormal	0.02	0.01

Table 23: Input factors for the AMTI probabilistic corroboration study.

[9] This section is adapted from the conference proceedings: "Holistic Approaches to Pre-clinical TKR Analysis: Computationally-Enriched Experimental Testing". 2009, Strickland et al. In: Knee Arthroplasty 2009 IMechE MED (London, UK).

Various output metrics were monitored for this study. A-P and I-E are driven so need not be measured; instead load-cell measurements were recorded, along with peak CP, M-L load split and the various wear metrics (sliding distance, cross-shear and linear wear rate for the different wear models described in Chapter Four).

7.4.2. Results

Figure 131 shows the 6 load-cell channels, with envelopes at ±1SD from the mean value (this is quite a limited range, but represents the only available *in-vitro* data).

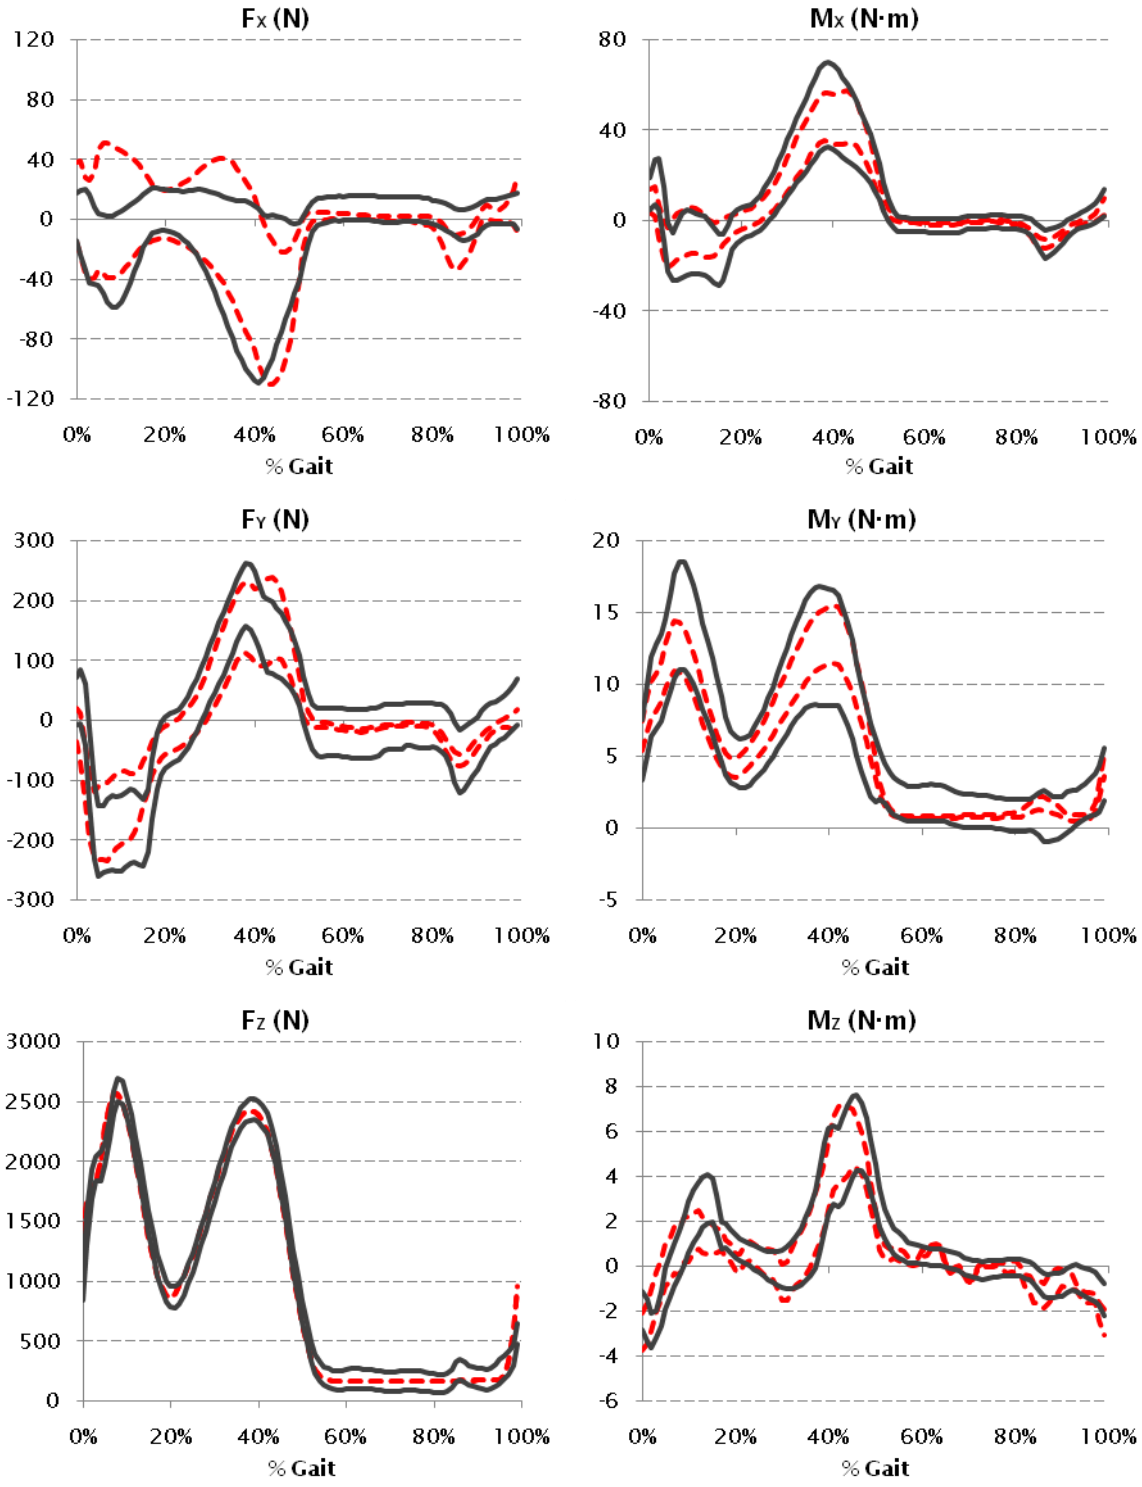

Figure 131: Comparison of response envelopes for load-cell measures: *in-vitro* (solid) versus *in-silico* (dashed). Envelopes are ±1SD (to match available *in-vitro* data).

In this first probabilistic comparison of *in-vitro* and *in-silico* data, the response envelopes show promising agreement between the two, for a 'proof of concept' study. In every case, the computational envelope tracks with similar trends and magnitudes to the experimental data. There are some clear differences; most notably in swing phase where the experimental data consistently shows more variability than is predicted *in-silico*. This is believed to be related to measurement errors within the load-cell itself. This is especially clear for the F_z channel; the computational model predicts almost no variability (as intuitively would be expected in the vertical direction, since this axis is under direct force-driven control), whereas the experimental data reveals a near-constant-width envelope of variation. The fact the width is so constant suggests strongly that these are offsets in the load-cell sensor calibration. This is an important point: in the computational domain, 'measurement' is an error-free process; however, experimentally the process of measurement can inherently introduce further error. In this case, it appears that the load-cells in stations 2 and 3 (which are not used for the control-system feedback) can carry offset or calibration errors, hence introducing further variability into the experimental results. This also may explain why the experimental envelopes are sometimes considerably larger than the computational envelopes on some of the other axes. Future models may need to account for this additional error by including a model of measurement variability – clearly, there is room for improvement. It is also possible that other important input factors have been missed in this demonstration study – including such as-yet-unidentified additional input variables may account further for some of the differences between the computational and experimental models.

The envelopes for contact pressure and M-L load split (Figure 132) reveal a higher degree of variability than was seen in the earlier theoretical/idealised probabilistic studies of Chapter Five (recall that in those earlier studies the envelopes were for 1%-99%, or about ±2½SD; here the envelopes are only ±1SD, but are still substantial). Note the considerable lateral load-shift and CP 'spike' in late swing phase (~85% gait); this is also evident in the load-cell data and is a result of the sharp A-P & I-E inflexions – this is a danger with displacement-driven testing. The results suggest that both test kinematics and system variability can have a considerable influence on whether the intended 60-40 load-split is achieved or not.

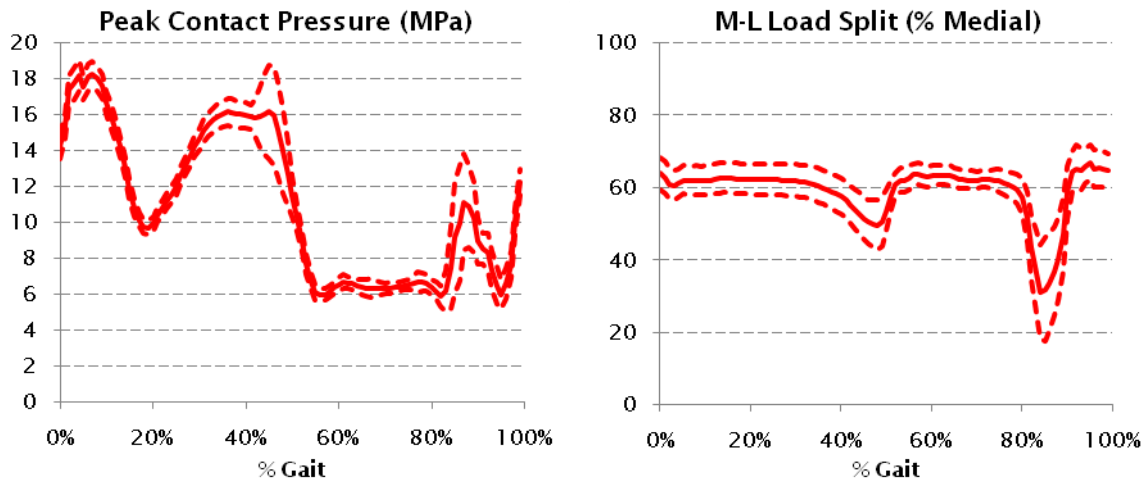

Figure 132: Response envelopes for peak contact pressure and M-L load split. Envelopes shows mean (solid line) ±1SD (dashed).

Sensitivity results (Figure 133) reveal that A-P dwell position plays a dominant role for many of the metrics under study – this is an interesting result, since Laz et al reported translational misalignment factors as being less significant than angular malrotation factors. However, that study did not consider the influence of variability in the A-P dwell position. By using real experimental data as the basis for the current study, A-P dwell was identified as a key factor with comparatively high levels of variability, which in fact dominates the malrotation terms. Similarly, the roller-bearing friction was also a strong factor – which again would be neglected by a less detailed model.

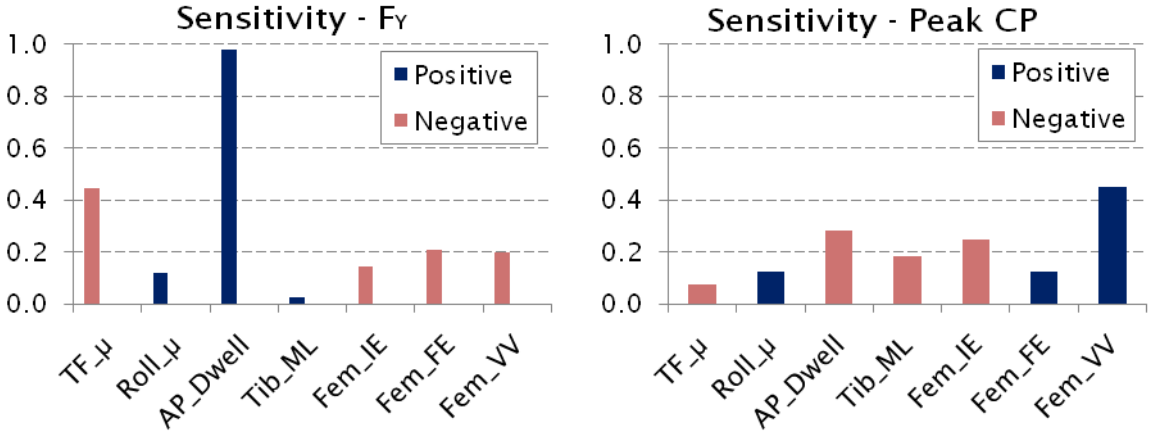

Figure 133: Selected sensitivity plots (normalised cycle-averaged values) for A-P shear force F_Y (left) and peak contact pressure (right).

The wear results provide the most pertinent observations about the current state of *in-silico / in-vitro* corroboration. PDFs for wear rate were compiled for each of the theoretical models and compared to the spread of interval wear rates recorded experimentally. Selected results are shown in Figure 134. Note that, even laying aside the differences in the deterministic 'mean' wear rates (which

have already been explored and discussed in Chapter Four), none of the *in-silico* PDFs come close to matching the level of experimental variability. The probabilistic wear study in Chapter Five revealed an approximately four-fold increase in wear between the 'Archard' wear model, and models with cross-shear. The results here show a further four-fold increase from those CS based models to the *in-vitro* results. At present, the mechanical model is still imperfect, and so some of the discrepancy could be due to the mechanical model, as well as the wear algorithm. However, the differences in the mechanical model are considerably less than the differences observed in these wear results.

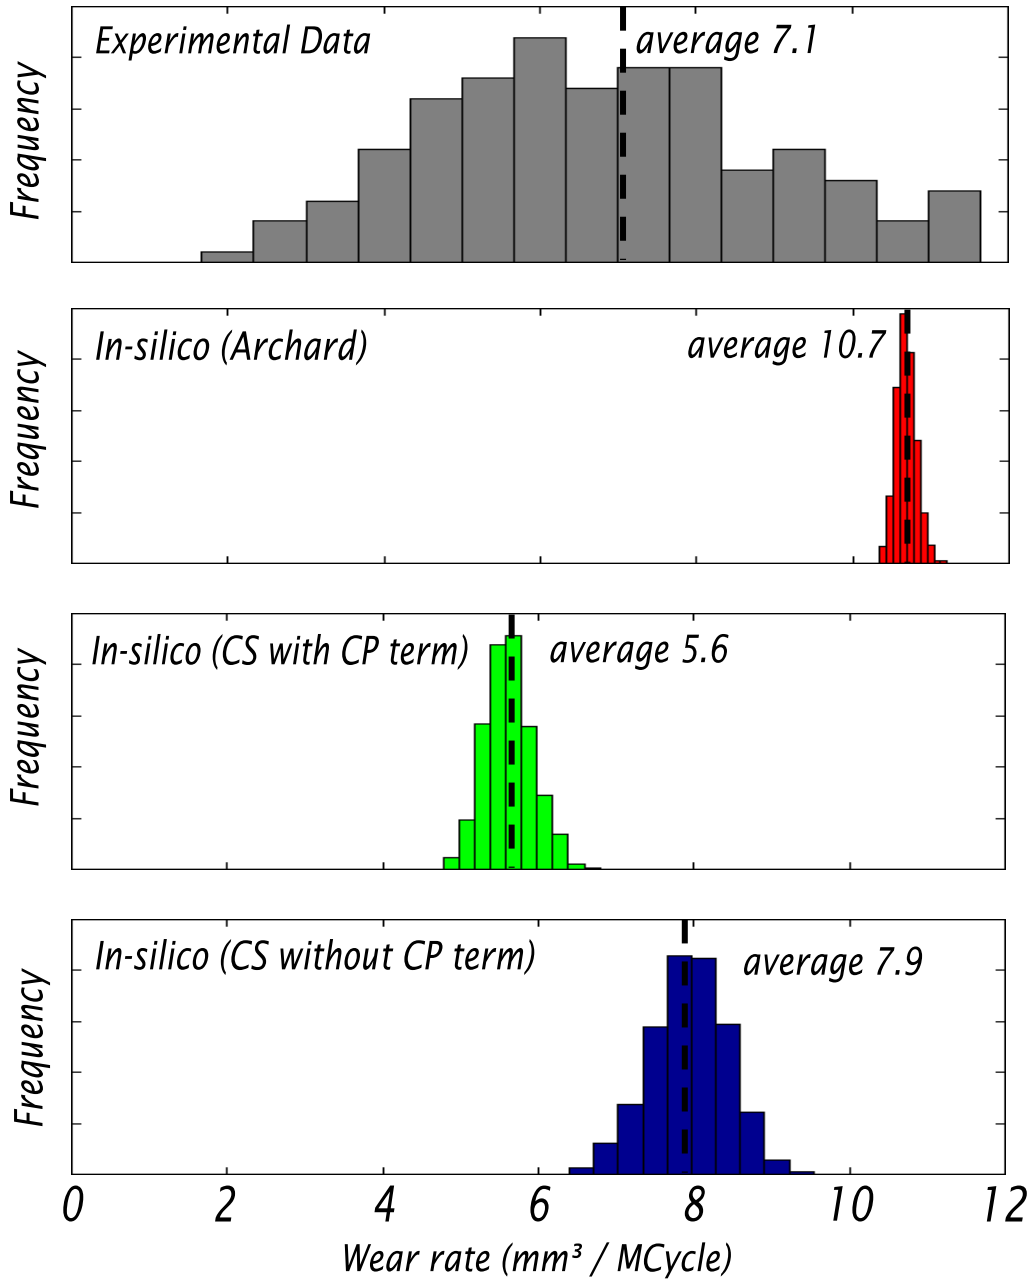

Figure 134: PDF of wear rates for experimental and selected computational results (note: the vertical axes are scaled individually, for visual clarity).

This raises important questions for computational wear modelling: principally, what is the cause of this discrepancy? Does it represent variability in the true wear rate, which is not currently detected by current wear theories? Or is it an artefact of experimental measurement procedures? Further, if it *is* an artefact of experimental methods, how reliable is *in-vitro* data as a basis for constructing theoretical wear models around? Ultimately, it should be possible to identify the best wear models by matching their PDFs to the PDF of *in-vitro* wear; if experimental methodology has such a confounding influence, this could limit efforts to better-understand the fundamental mechanics of wear.

7.4.3. Discussion

This study represents the first time a true corroboration has been attempted for a probabilistic analysis. The results are imperfect; this is to be expected, since there is no direct way to ascertain the uncertainty of 'inputs' to the system. However, the fact that magnitudes and trends are so demonstrably similar is very encouraging, and demonstrates convincingly that the fundamental methodology behind the probabilistic approach is sound.

Of particular interest is the higher degree of variability in the experimental wear data compared to the computational results. This reveals the importance of *measurement error* within the *in-vitro* assessment process. This has previously been neglected by probabilistic studies. However, as the results here clearly show, there is evidence that measurement errors (for both the load-cell results and for wear assessment) are playing a strong (if not dominant) role in the observed variability.

This discrepancy clearly needs to be addressed if *in-vitro* and *in-silico* models are to be corroborated more accurately. However, the question is, should *in-silico* models attempt to model measurement variability, or is it the role of experimental testers to reduce this variability? In either case, the key to successful *corroboration* is better *collaboration*. If *in-silico* models are to reflect the true variability inherent in experimental procedure, this requires a better understanding of those experimental procedures; spending time working alongside experimentalists to understand the methods and processes being used, in order to identify (and quantify) where variations and uncertainty are introduced. If, on the other hand, *in-vitro* variability is to be reduced, this can be greatly assisted by co-operating with computational modellers; for example, in the present study the computational results may be used to identify the key sensitivity factors; experimental procedures could then be focused on better-controlling these factors. In this case, the procedure for assigning A-P dwell

position has subsequently been revised to be more repeatable; closer attention is also now paid to the friction of the roller bearings.

There are important limitations to the study described here. The sample set remains relatively small at just over 100 samples; ideally much more data would be needed for a robust corroboration. This is the reason for presenting data only to ±1SD; presenting data for 5-95% or 1-99% would require sample sizes an order of magnitude larger. The damping and friction terms have been tuned based on limited experimental testing. Ideally, a more robust set of tests would be needed to fully characterise the dynamics of the rig. The *in-silico* controller is not identical to the *in-vitro* version; more information would be required from the manufacturer to construct a more accurate model, and so limit differences due to the control system. The wear models are based on existing theoretical concepts which are only empirical and approximate. Further influential input factors may exist, and the current factors could be more accurately characterised statistically. The wear was evaluated based upon a single-cycle analysis; this cannot account for adaptive wear effects (however, a fully adaptive probabilistic wear assessment was beyond the scope of this exploratory study). Additionally, creep and plastic deformation of the polymer were neglected, further limiting accuracy.

Nonetheless, the study very clearly shows the benefits of better corroboration. The results raise interesting questions about the underlying experimental data and the mechanics of wear. The fact that much variability is unaccounted for shows that there is still considerable scope to progress this work in the future.

7.5. Summary

The various displacement- and force- driven models of the AMTI simulator in this chapter have built on many of the lessons learnt in the earlier modelling activities. The availability of better experimental data (especially the combination of force and displacement feedback from the tri-axial load-cell as well as displacement transducers) means that there are more means by which the model can be corroborated. This gives greater confidence in the model when good corroboration is achieved, but equally provides a much richer diagnostic resource when discrepancies arise.

The displacement-driven modelling in particular was very successful. The model performs well, with good accuracy, but solution times far faster than the FE-based methods employed by Lanovaz et al [200]. This reduction in computational cost is a key enabling pre-requisite for probabilistic studies, such as that described in Section 7.4. This ‘probabilistic’ perspective is very valuable, revealing that even

for a relatively controlled gait-test, the actual kinetics of the contact are inherently variable (with implications for kinematics and wear rates). The comparison of *in-vitro* and *in-silico* probabilistic results is also highly informative; it is apparent that while the variability in mechanics can be replicated *in-silico*, the corresponding variability in wear rates cannot currently be accounted for purely based on existing wear algorithms. Probabilistic assessments of wear may prove to be a key future tool in furthering the understanding of wear mechanisms.

The challenges of accurate and repeatable force-driven simulation are illustrated by the difficulties encountered in reproducing both the kinematics and kinetics of force-driven gait. This is obviously a challenge for future computational modelling efforts. However, it is equally a challenge for experimentalists (who benefit from a sound, quantifiable understanding of the mechanics of their test simulations through *in-silico* corroboration) – since good experimental data is the basis for any effective *in-silico* modelling. The influence of bearing friction, and pliancy in the F-E axis assembly, clearly demonstrate that the results are being altered by unintentional artefacts from the experimental set-up. Equally, the 'dynamic' properties (e.g. inertia and damping), which are not tailored to represent *in-vivo* dynamics, are also influential. The compounded effect of these different factors is that the test outcomes are variable and susceptible to subtle changes in the experimental hardware or procedures. This is not ideal, as such variations confound the important aim of the tests: to understand the effect of TKR design and materials on kinematics and wear-performance. By working together, computational and experimental researchers may be able to better identify and hence mitigate some of these other confounding influences.

CONCLUSIONS & FURTHER WORK

This thesis set out to demonstrate that computational and experimental methods could be used together more effectively to provide an enriched pre-clinical analysis toolset, and further to show that accounting for variability using probabilistic methods is an essential part of any study of knee biomechanics.

The need for improved pre-clinical analysis methods, driven by rising demand for TKR, is presenting new challenges to orthopaedic designers and researchers. Established computational and experimental methods have a venerable pedigree in building the body of current scientific knowledge and providing guidance for current TKR designs. However, these studies have often been isolated, poorly corroborated and limited in scope, failing to consider the high levels of variability inherent in TKR performance.

It is clear that there are deficiencies and limitations in the existing experimental studies. The lack of standardisation on 'normal' gait profiles for wear testing (compare the profiles proposed in [44] and [24]) reflects an imperfect understanding of the true *in-vivo* mechanics. The large differences in wear rates reported between very similar tests (e.g. compare [206] and [177] where tests from the same research centre on comparable TKR designs exhibited a tenfold difference in wear) demonstrate that experimental procedures and sources of variability are also not fully controlled or understood. The results between different research centres are still less consistent, (for example wear-tests including stair activities have contradicted each other, reporting both higher [201] and lower [175] wear rates compared to normal gait) – clearly showing the degree of variability and uncertainty in current *in-vitro* methods.

There is, then, a need to better-understand the outcomes of experimental research, and fast computational models can augment experimental tests to improve understanding and provide better data for pre-clinical research and development.

Central to making progress in this field is the need for better collaboration between *in-silico* and *in-vitro* testers. By working to corroborate results across multiple test platforms, researchers can gain a more complete picture of the test mechanics, and subsequently have access to a more powerful database from which to extract and visualise the results of the test *in-silico*.

Probabilistic methods can provide the framework for understanding the influence of variability; multiple factors can be combined in a single model and explored in

a statistically robust manner. The work of Browne et al [142] in introducing these methods to the field of orthopaedics, and Laz, Pal et al [95, 179] in developing their application to TKR mechanics, has established the foundation for the broader uptake of probabilistic methods. However, probabilistic approaches require many more trials, and therefore necessitate faster modelling methods than the deformable FE models preferred historically. Rigid-body modelling (for example using MBD software) can provide this speed increase, and has been successfully adopted by various research groups, including Bei, Lin et al at the University of Florida [84, 139] as well as in the various studies included in this thesis [122, 207-213]. There is of course an accuracy trade-off; however if sufficient attention is paid to *in-vitro* corroboration efforts, it is apparent that the accuracy is still acceptable for many investigatory studies. (In fact, the errors resulting from poor or inadequate corroboration can be larger than any errors from assumptions of rigid-body mechanics).

In this thesis, the development of a new generation of MBD-based knee simulations has been charted from conceptualisation and early demonstration, through further studies incorporating probabilistic methods and *in-silico* wear prediction, up to highly-focused corroboration studies against specific data-rich *in-vitro* testing platforms (the KKS and AMTI knee simulators).

The work in these final chapters represents the 'state-of-the-art' in computational modelling of TKR *in-vitro* simulation, and *in-silico*/*in-vitro* corroboration. By combining fast rigid-body modelling techniques, contemporary theoretical wear models, and probabilistic methods, and by actively engaging in a deeper level of collaboration between computational and experimental researchers, an excellent foundation has been laid for future pre-clinical analysis efforts.

Whilst it is important to recognise that this work represents a step-change from the basic deterministic FE-based models of only a few years previously, it is equally important to appreciate that there remains a great deal of work to be done, if *in-vitro* knee testing is to be better understood. The efforts to corroborate dynamic models have demonstrated that many of these dynamic effects (friction, damping, inertia) and their influence on test outcomes are not rigorously understood, even within the experimental community. Working with *in-silico* modellers gives experimentalists an excellent opportunity to ascertain and improve their own understanding of their test-platforms; every discrepancy encountered between the computational and experimental results represents an opportunity to investigate, diagnose, and ultimately build a more sound understanding of the real-world physical mechanics.

This is particularly true in the domain of wear theories. The work in this thesis has shown that existing wear models are a valuable tool, and do have real and useful predictive power, demonstrating this more robustly and conclusively than any other studies previously. However, it is also apparent that they are not perfect; the mid-range correlations observed in Chapter Four and the comparison of wear distributions in Chapter Seven demonstrate this. The resulting challenge involves both experimental and computational researchers; experimentalists must identify and reduce the variability and uncertainty in their tests (both within and especially between different research centres) if the data they provide is to be most effective. Computational modellers must then revise the theoretical algorithms, to better reflect the observations revealed through POD and TKR testing. This is best deployed as an iterative learning process; new *in-silico* models should be based on *in-vitro* results, but the predictions and anomalies they highlight then need to be rigorously screened experimentally.

The work in this thesis, and associated modelling efforts working with experimental researchers in industry, has helped to identify key limitations in our current wear theories – it has been demonstrated robustly that the present models are not perfect, and that better data is needed in order to advance further [211]. Existing studies have already begun to challenge the assumptions about the role of contact pressure (e.g. [106, 107]), and recent POD investigations supported by the authors using MBD-based modelling have also challenged the assumption that wear is simply proportional to sliding distance (Dressler et al, [214]). The newest wear-modelling algorithms involve the concept of incorporating a 'memory' into the polymer, so that wear is a function of the time-history of sliding directions (not simply the total sliding distance or even cycle-averaged cross-shear) [215]. However, even these models neglect other known important factors; for example the choice of material type (e.g. [173]), and the phenomenon of 'lift-off' (e.g. [172]) are known to influence wear; clearly, there is considerable scope to progress the theory of wear modelling. *In-silico* simulation has a key role to play in this, because it is flexible, adaptable and can provide a rich source of supplementary data. Computational models can serve as the medium by which different experimental test platforms communicate and interface. For example, the kinematics and kinetics of lower limb simulators such as the KKS rig can be analysed computationally, and re-framed to serve as new activity profiles for wear-simulators such as the AMTI rig. In turn, these physiological profiles and loads can be post-processed from wear-simulator based TKR testing, and used to provide more appropriate sliding paths and load-profiles for POD tests. Wear models can then be developed and refined in the POD-testing

domain, and ultimately re-exported to the knee-simulator platforms such as the AMTI, SKS, or KKS rigs. Note that the computational models are not in any way supplanting the experimental tests; rather they complement, enhance and interface the experimental tests, producing a more holistic, more robust, better synchronised and integrated environment for pre-clinical analysis of TKR designs, materials and technologies.

The technical contributions made by this thesis, and the contribution to knowledge in the field, may therefore be summarised as follows:

- Rigid-body MBD models have been developed and robustly corroborated against various existing FE-based simulations, demonstrating that the performance-accuracy trade-off with MBD can be acceptable, if used appropriately. This confirms the findings of e.g. Bei [137] and Guess [197], who used such rigid-body modelling approaches extensively.

- Specifically, improved models of the KKS and AMTI platforms have been delivered, building on the previous modelling of these platforms by (amongst others) Guess [109] and Lanovaz [200]. These new models can now serve as the basis for further ongoing research.

- In the process, specific lessons have been learned about the mechanical behaviour of these experimental platforms. For example, the KKS modelling identified the considerable losses induced by frictional effects on the quad actuator in particular; subsequently the rig has been accordingly re-designed with new degrees of freedom to accommodate small amounts of misalignment. Similarly, the AMTI models highlighted the effect of friction in the roller-bearings; in light of this, much closer attention is now paid to this friction influence during experimental testing. These are two examples of how computational modelling can 'feedback' into the physical domain of *in-vitro* testing (i.e. bi-directional sharing of information between *in-vitro* & *in-silico* platforms), hence the collaboration is mutually beneficial to computational and experimental modellers.

- The combination of probabilistic methods with wear prediction has revealed that with many theoretical wear models, wear rate are quite sensitive to relatively small variations in the 'input' conditions under test. For example, the study in Chapter Five revealed that misaligning the components to a relatively small degree (with a standard deviation of only 2°) is sufficient to produce a two- or three-fold increase in wear rates. Further, this sensitivity

appears to be somewhat design-specific. This would not be observed with purely 'deterministic' assessments of wear (whether *in-silico* or *in-vitro*).

- More generally, the focus on *in-silico* wear prediction in this thesis has led to an improved appreciation of the capabilities and limitations of existing models. We are now able to compare the predictive power of many of the current wear algorithms [216], and have identified a key weakness in their inability to account for the true observed experimental variability (as shown in the final chapter).

- The application of probabilistic methods has been demonstrated, building on the foundation laid by Laz [95] and Pal [217]. We have performed a 'first-of-kind' probabilistic corroboration in the field of knee biomechanics [212], with a proof-of-concept study on the AMTI simulator providing promising initial results. This study shows that it should be possible to corroborate experimental knee test platforms probabilistically, and in the process identify which input factors are affecting the system performance. This work also sets a benchmark for future studies to work towards, in terms of combining computational-experimental corroboration with probabilistic methods.

The assumptions and limitations within the present models should not be overlooked. Fundamentally, rigid-body modelling is inherently inaccurate for TKR contact mechanics; the contact pressures for almost any knee design or ADL profile *will* exceed the elastic limit of polyethylene; permanent plastic deformation will occur; over the long-term testing timescales of wear simulations, creep will alter the surface profile. Whilst these effects are generally relatively small, in certain cases (e.g. edge-loading of the insert, or loading of the cam-post in PS designs) the nonlinear behaviour may be considerable – then the assumptions behind rigid-body modelling begin to break down. (Nonetheless, the large performance gain for this small loss of accuracy must be considered) [94].

The choice of parameters for contact and friction modelling is based upon experimental tuning, and this introduces difficulties. Contact properties may vary from material to material; these differences are not characterised. Friction co-efficient is known to be related to wear rate (Wang, [162]) and to vary across the polymer surface depending on the local motion paths at any given point (as shown by Dunn et al [205]) – again, the assumption of a single co-efficient value held constant across the surface and throughout the cycle is an over-simplification of reality. To the author's knowledge, no force-driven simulation has yet been reported which varies friction across the contact area; despite the fact that, as reported by Godest et al, friction is known to be influential for force-

driven modelling, and the reported coefficients used vary by as much as an order of magnitude [91]. There is clearly an opportunity for better accuracy in this area in future.

The statistical modelling introduces the need for further assumptions; as discussed in Chapter Five, there is only limited and sporadic experimental data available describing the variation observed in many of the relevant factors; without better data, assumptions must be made about factor interdependence, distribution shapes, mean values, levels of standard deviation, and range limits on variables. All of these have the potential to introduce errors. It is hoped that with the wider adoption of probabilistic methods, researchers in the experimental and computational community will begin to appreciate more the value of better statistical data, and consequently more effort will be made to collate and report this information. A cursory analysis of the data reported by Mahaluxmivala et al [149] suggests that component misalignments may well be relatively independent of one-another, and have distributions close to Gaussian (as assumed in this thesis and other published studies); however this must be investigated further.

In light of these limitations, and the other various obstacles and challenges identified during the various studies described, there is considerable scope for further work. Experimentally, there remains much work to be done in better understanding the mechanics of wear; *in-silico* models have a role to play in supporting this investigation, and better POD-test models will help to corroborate the most fundamental investigations into wear behaviour. There are challenges in translating the work done using POD investigations to the domain of TKR testing; again, the ability to decompose and analyse the mechanics in detail using computational methods is valuable in bridging this transition between different testing platforms.

Nonetheless, the progress made with the current work is valuable; the concepts and methods of corroborated probabilistic analysis methods have been demonstrated and applied for a range of different platforms, and de-risked by extensive comparison to existing published work. The foundation has been laid for these models and modelling approaches to be used to support commercial/industrial TKR research and design efforts.

The models created, especially the KKS and AMTI simulators, have potential to be used for further studies. The KKS model has been robustly validated, but only applied for a concept-demonstration study. The model could be used for a wide range of purposes besides profile generation and testing, probabilistic studies of misalignment. The flexibility of the *in-silico* model makes it an ideal test-bed to

explore future modifications to the rig (e.g. the possibility of a mobile A-P axis, or of including a 'hamstring' actuator). Further, computational tools would allow better cross-platform modelling; it is possible to isolate and extract the kinematics for any given profile on the ADAMS KKS model, and use these as the basis for new knee-wear simulator profiles (e.g. on the AMTI rig). The central theme once again emerges: computational corroborative modelling has the potential to bridge and interconnect the various distinct experimental testing domains, providing a more holistic perspective, and enriching existing pre-clinical analysis capabilities. But this depends on high-quality, well-corroborated models.

Note that all the work in this thesis has focused exclusively on corroboration with *in-vitro* testing; the entire domain of *in-vivo* validation has not been directly considered. The author would argue that in fact, until rigorous corroboration is possible between *in-silico* and *in-vitro* results, any application to *in-vivo* performance will always be open to question, hence fundamentally undermining the confidence of the broader healthcare community in any results presented. By first demonstrating good computational-experimental corroboration in the domain of pre-clinical *in-vitro* analysis, researchers and designers can demonstrate that they have a sound, robust, and quantifiable understanding of the physics of the systems they are working with; this in itself does not demonstrate that the results are applicable to *in-vivo* outcomes, but it is an essential foundation towards that goal.

Ultimately, the delivery of new TKR designs is a large-scale, multidisciplinary effort, involving specialists from the clinical, industrial and academic community, encompassing backgrounds as varied as surgeons, mechanical engineers, physiotherapists, materials scientists, anaesthetists, computational modellers, manufacturing engineers, and many others. The tools and methods discussed in this thesis represent a small but essential part of this larger process-chain. Pre-clinical analysis represents the enabling technology to deliver the next generation of knee replacements, in order to drive down revision rates and improve functional performance. Ensuring the very best tools are available to designers supports them in making the very best design decisions. By giving designers the confidence that their tests are reproducible and fully characterised, by presenting the 'holistic' perspective offered with probabilistic methods, and by integrating the widest possible suite of tools for assessing knee kinematics, kinetics, laxity & wear, computationally enriched pre-clinical analysis methods can help to make those design decisions better, ultimately contributing to real improvements in patients' experience of TKR and subsequent quality of life.

APPENDIX A – HUMAN ANATOMIC REFERENCE FRAMES

1. Reference Planes

The human body can be adequately decomposed into three orthogonal reference frames. A plane may be defined to cut the body at any point, although planes cutting through the midpoint of the body are sometimes termed distinctly; e.g. the mid-way sagittal plane is termed the 'median' or 'mid-sagittal' plane.

Sagittal *(lateral)*:	Plane normal to the M-L axis, formed by the A-P & I-S axes
Coronal *(frontal, dorsal)*:	Plane normal to the A-P axis, formed by the M-L & I-S axes
Transverse *(horizontal)*:	Plane normal to the I-S axis, formed by the M-L & A-P axes

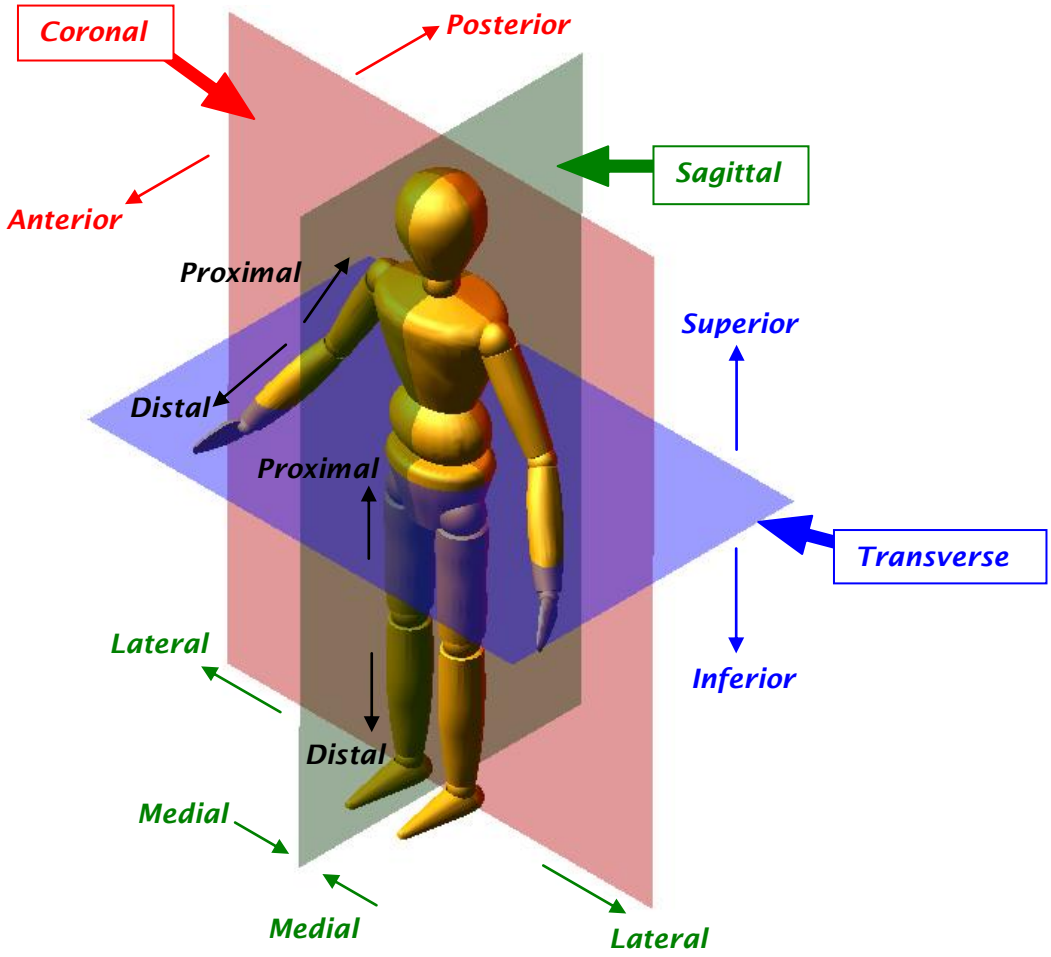

2. Directional Terms (Translations)

Many of these terms are used to indicate the *relative* position of features (or sometimes, to describe a relative motion), and do not have any 'absolute' positional meaning. The most common are summarised in the following table (Note there is more than one term for some of these directions; the preferred term is shown in bold; the alternative in parentheses).

***Inferior** (caudal):*	'Lower' (closer to the base, caudal literally refers to 'tail')
***Superior** (cranial):*	'Higher' (closer to the head)
***Medial**:*	Towards the median (mid-sagittal) plane; 'inner'
***Lateral**:*	Away from the median (mid-sagittal) plane; 'outer'
***Anterior** (ventral):*	'Forwards' (towards the front surface)
***Posterior** (dorsal):*	'Rearwards' (towards the rear surface)
***Proximal** (central):*	Closer to the centre of the body (torso)
***Distal** (peripheral):*	Further away from the centre of the body

3. Directional Terms (Rotations)

Adduction:	An active motion towards the median plane
Abduction:	An active motion away from the median plane
Internal rotation:	Rotation inwards (towards the body)
External rotation:	Rotation outwards (away from the body)
Varus:	An inward twisting of the distal limb (for the knee, 'bow-legged')
Valgus:	an outward twisting of the distal limb (for the knee, 'knock-kneed')
Flexion:	Motion that decreases the joint angle, or the state of being 'flexed'
Extension:	Motion that increases the joint angle, or the state of being 'extended'
Hyperextension:	Extension beyond the 'normal' joint range

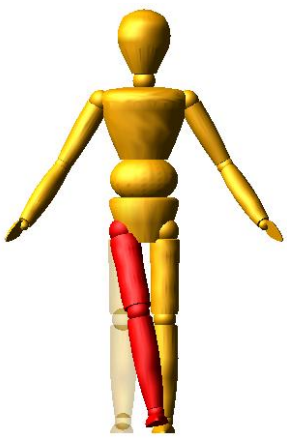

Adduction

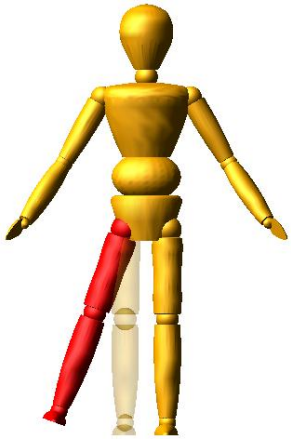

Abduction

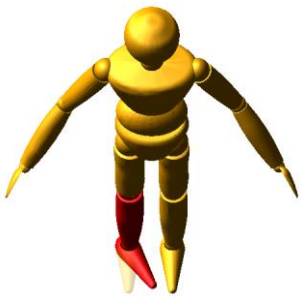

Internal rotation

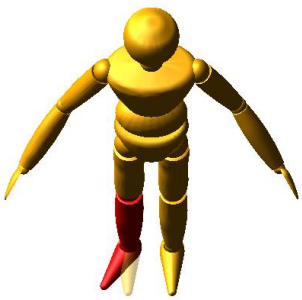

External rotation

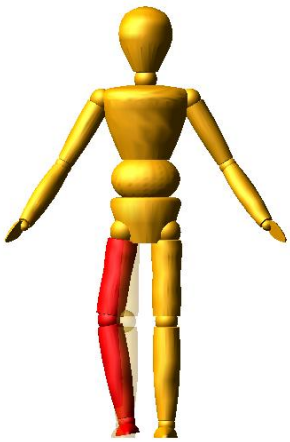

Varus

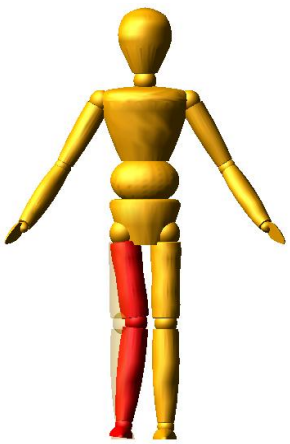

Valgus

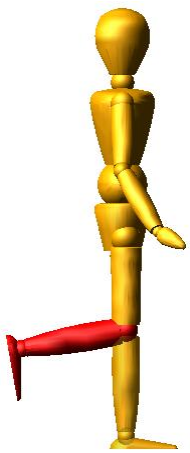

Flexion

Extension

Hyperextension

4. Motions at the Tibiofemoral & Patellofemoral Joints

It is important to use anatomical frames of reference at the knee joint with caution. No segment of the knee is stationary; the tibia, femur and patella all change position and orientation in normal gait and other activities. Therefore, use of a term such as 'inferior' or 'superior' is misleading; the terms 'proximal' and 'distal' are better suited. (For example, with the hip extended the 'distal' direction along the tibia is 'inferior' when the knee is fully extended, but 'posterior' if the knee is at 90° flexion). It is important that the frame of reference should always be reported (i.e. motion should be 'with respect to' the femur, tibia or patella), e.g. an 'anterior' motion of the tibia could equally be presented as a 'posterior' motion of the femur (and vice-versa).

However, since these various terms are widely used to describe knee kinematics, the following figures illustrate the conventional use of the terms for the tibiofemoral and patellofemoral joints. (Note: the illustrations show a 'right' knee; for a 'left' knee, the direction of medial-lateral, varus-valgus, internal-external, and patellar tilt & rotation would be mirrored.

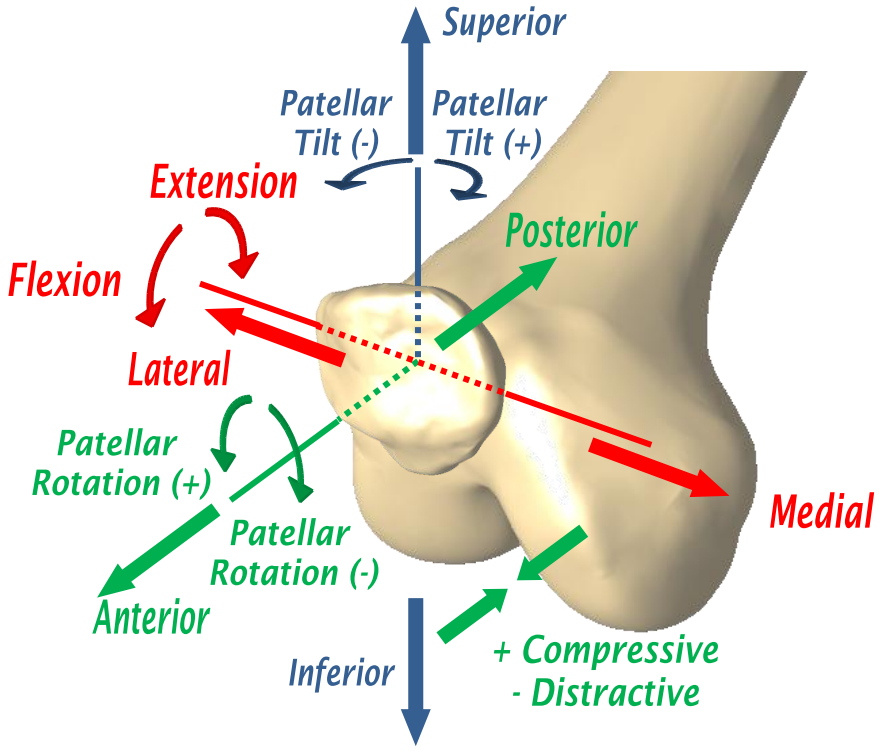

Figure A1: The motions of the tibiofemoral joint.

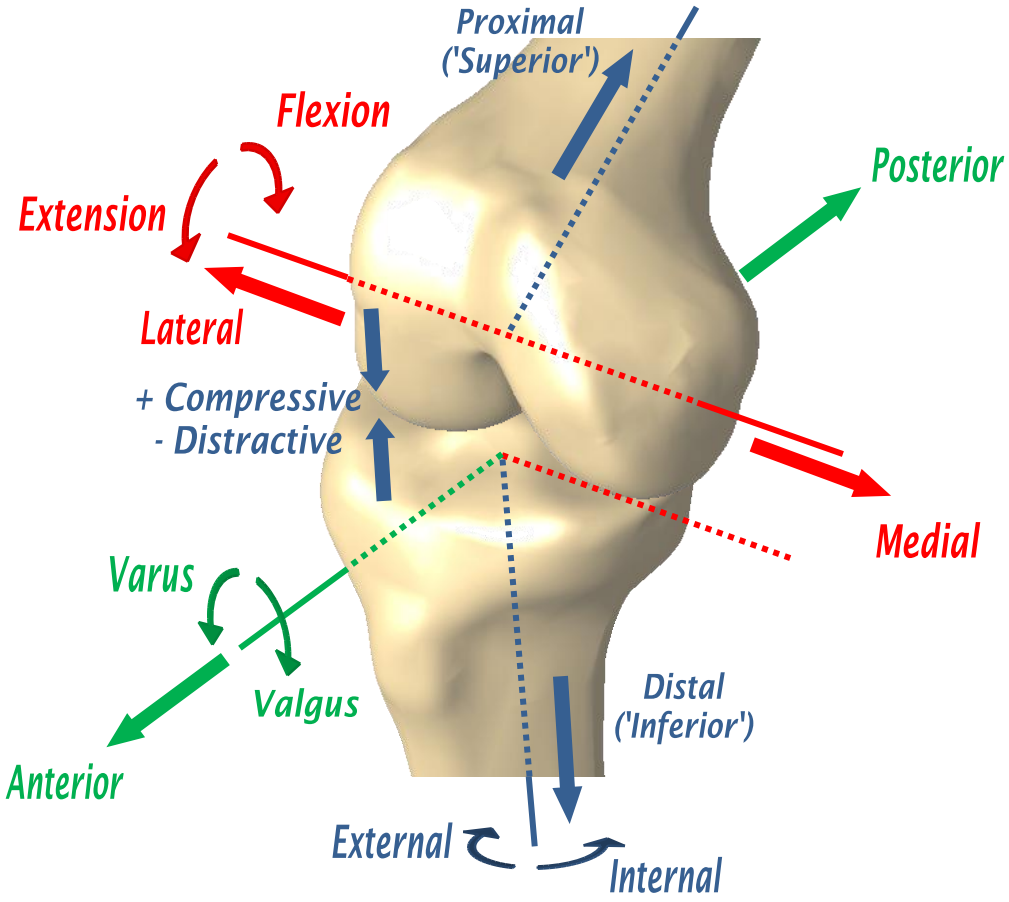

Figure A2: The motions of the patellofemoral joint.

APPENDIX B – INTERFACING MSC.ADAMS WITH NESSUS

Note: the information is presented in the context of an interface to NESSUS; however a similar technique may be used to link to any 3rd party software via an ASCII-based conduit

ADAMS includes its own proprietary stochastic analysis module (ADAMS/Insight) and does not directly interface to NESSUS. However it is useful to establish an interface to exploit some of the advanced analysis methods not supported within Insight. The following is a step-by-step guide:

1. Create the rigid-body model in ADAMS/View. Use design variables for the parameters that will be varied as input factors. Use requests or measures for the parameters to be measured as output responses. The model must contain a script written using ADAMS solver commands ('acf' syntax). The model must also contain at least one design objective; however this is a 'dummy objective' and will not be referenced by NESSUS. Save the model as a binary file (*.bin), for speed of access (this will use more disk space, but the model itself will not be duplicated)

2. Use a text editor such as notepad to create a template ADAMS/View command file (*.cmd) to run a DOE trial with the following commands:

```
file binary read file_name = "«PATH/FILENAME.BIN»"

simulation multi_run doe &
   model_name = «MODELNAME» &
   sim_script_name = «SCRIPTNAME» &
   variable_names = «FACTOR#1», «FACTOR#2» &
   objective_names = «OBJECTIVE» &
   rows_in_table = 1 &
   table_of_values = XXXXX, XXXXX
```

   Where «RED» text is replaced with the required names. The input file must include the full path. Note the string of X's for the table_of_values parameter is deliberate; this string should be long enough for the numerical precision required by the model. More variable names & values can be added using comma separators; only two are shown for brevity.

3. Beneath the above commands, add additional commands to export the output data to ASCII text file(s). The precise syntax depends upon the output data required. If a REQUEST is to be used from ADAMS, use the following syntax:

```
file request write &
        analysis_name = Last_Run &
                file_name = "«OUTPUT.REQ»"
```

If numeric data is to be exported (e.g. from a MEASURE), use the following syntax:

```
numeric_results write &
        result_set_component_name = «MEASURE.CMPT» &
        file_name = "«OUTPUT.TXT»"
```

Create as many outputs files as needed, although note that the specified filenames must not include a full path. Save the file as DOE_TEMPLATE.CMD. It will serve as the master template file to be edited by NESSUS.

4. Load NESSUS, and start a new project. Define a suitable problem statement, e.g. "output1=f(x1, x2)", and under response model define the model type as "Numerical" and application as "USER_DEFINED". Choose Interactive mode, and enter the following execution command (note this is ADAMS version-specific; replace the text 'adams05r2' according to your version of ADAMS):

   call adams05r2 aview ru-s b DOE_MODIFIED.cmd

   Underneath, for input files select your template file (DOE_TEMPLATE.CMD). Specify the destination as DOE_MODIFIED.CMD. The output files should refer to your «OUTPUT.REQ» or «OUTPUT.TXT» file(s).

5. Under Create Mappings, create a mapping for each variable, ensuring the target is DOE_MODIFIED.CMD. For each variable, select the relevant line/column index for replacement, highlighting the correct 'XXXXX' string. Ensure variables are in the correct order.

6. Under 'Select Responses', for each output make a similar mapping to the corresponding output file (*.res or *.txt). Note ADAMS can format and sort the output data using additional command line parameters which can be included in the *.cmd file - experimentation will reveal the most suitable file formatting to use for a given requirement.

NESSUS is now able to interface to ADAMS, to run stochastic analyses. Note that the output files will be put into subfolders underneath the NESSUS project directory. If the input file path is not specified in the *.cmd file, ADAMS will overwrite the NESSUS path setting, and all output files will be created in the base directory, resulting in a 'file not found' error in NESSUS. Consult the ADAMS and NESSUS help files for further reference on specific commands and options.

APPENDIX C – CHOSING CONTACT PARAMETERS IN ADAMS

This information is a guide to help the new ADAMS user chose appropriate parameters for the 'IMPACT' function in ADAMS, to create an 'elastic foundation' (EF) contact model for tibiofemoral or patellofemoral knee mechanics

The ADAMS 'IMPACT' function [161] relates contact normal force (F_N) directly to geometric penetration depth (g), using an exponential relationship of the form:

$$F_N = k \times g^e$$

Where k is the stiffness co-efficient and e is the force exponent. To derive suitable values for a rudimentary first approximation, several assumptions are made. For the discretised tibial or patellar insert, every cell is considered to be identical and differentially small, of equal material thickness (e.g. for a typical tibial insert, h ~ 10mm). For the models in this thesis, the contact surface area of the cell was chosen as A = 1mm² (based on sensitivity studies), so to a first approximation, using the Young's modulus relationship between modulus, (E), stress (σ) and strain (ε):

$$E = \frac{\sigma}{\varepsilon}$$

With penetration related to strain by the equation:

$$g = \varepsilon \times h$$

And contact pressure taken as a homogenous stress, and so related to normal force by the equation:

$$\sigma = CP = \frac{F_n}{A}$$

So for a basic linear stress/strain relationship, with a given constant value of modulus (typical values for UHMWPE are ~1GPa, depending upon the grade; the NIST standard is 1258 MPa ±22<MPa [167]), the exponent, e, should be taken as unity, whilst the stiffness constant, k, would be:

$$k = \frac{F_n}{g} = \frac{\sigma \times A}{\varepsilon \times h} = E\frac{A}{h}$$

Although this is a very simplistic approach, it is found that in many cases, this linear-elastic model is adequate for a first-order estimate of contact mechanics (especially if there is limited high-stress contact, e.g. edge-loading or cam-loading). There are alternative proposals in existence, e.g. early mathematical

models of knee mechanics used a direct relationship between surface contact pressure and the penetration depth (see the models of Blankevoort, e.g. [218, 219]), and these have subsequently been adopted for use in MBD and rigid-FE based models (e.g. [84, 94, 220]). Note though, that despite differences in terminology (e.g. "elastic foundation", "bed-of-springs", "pressure-overclosure relationship"), and different forms of the equations, these models are all essentially similar; considering the equation presented by Blankevoort:

$$p = \frac{(1-v)E}{(1+v)(1-2v)h}d \quad \text{(from [219])}$$

It is apparent that the pressure term, *p*, can be replaced with force, *F* divided by area, *A*, and the '*d*' term is equivalent to penetration depth '*g*'; so the equation can be re-written as:

$$F = \frac{(1-v)}{(1+v)(1-2v)}E\frac{A}{h}g$$

Then the 'stiffness' term, 'k', is of essentially similar form to the equations used within ADAMS, (although with a correction factor for the Poisson ratio, '*v*'):

$$k = \frac{F}{g} = \frac{(1-v)}{(1+v)(1-2v)}E\frac{A}{h}$$

The FE study by Halloran [94] used rigid-body models of FE based on linear elastic foundation models, and also deformable models of FE. This study reported only small differences between the rigid linear model, and the fully deformable model. Of course, the deformable model can be more accurate, and deformable models become necessary when high loads lead to significant non-linear behaviour or plastic deformation.

To better reflect the relatively incompressible & elastic nature of the polymer, a factor may also be included to account for the Poisson ratio of the material, *v* (typically around ~0.45 for UHMWPE):

$$k = \frac{(1+2v)\sigma \times A}{\varepsilon \times h} = E(1+2v)\frac{A}{h}$$

In reality, UHMWPE has a non-linear stress/strain relationship (modulus varies with strain); a typical experimental relationship is shown in the chart below. In this case, the linear model above cannot be used, and since the exponential curve cannot be fitted exactly, an appropriate regression fit must be chosen.

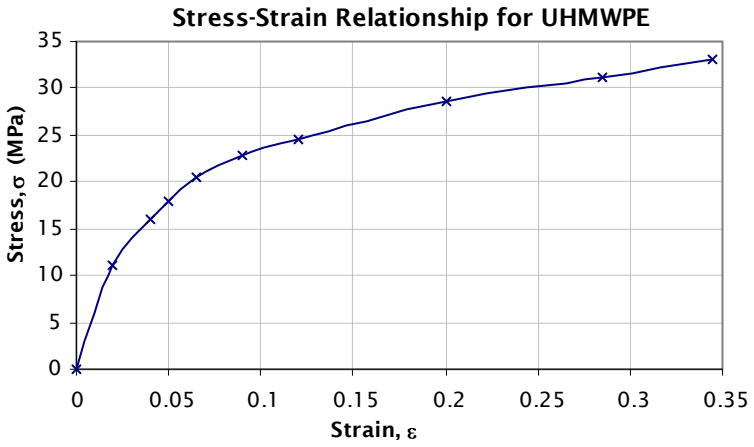

Typical stress-strain relationship for UHMWPE (from [221]).

If the exponential 'IMPACT' function is to reflect this nonlinearity, the exponent will not be '1'; instead, both '*e*' and '*k*' must be fitted to the experimental data. The best result can be achieved by optimising the fit only within the area of interest. High stresses (above 20MPa) should not occur frequently; therefore the fit of the exponential curve may be optimised for the region below this (< ~6% strain). For example, the best fit achieved for the data in the above chart used the values $k \approx 6500$, $e \approx 0.75$ ($R^2 = 0.989$ for $0 \leq \varepsilon \leq 0.06$). Note that these constants are specific to the values of '*h*' and '*A*' selected above (a different arrangement would need new constants). As with other numerical methods, a higher degree of nonlinearity in the contact mathematics will adversely affect solution times. Generally, the solver performance will be better for values of *e* greater than one. This is because the force F_N is set to zero for negative penetration depths (i.e. when the solids are not in contact, there is no contact force). The transition into contact using an exponential-type relationship will always be a continuous function, but the derivatives of this function will *not* be continuous if *e* <= 1, as is illustrated in the following figure.

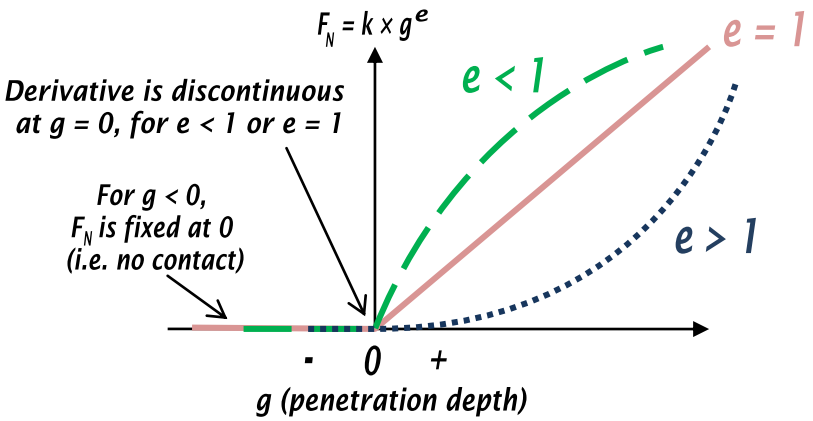

Relationship between F_N and *g*, for different values of *e*. Discontinuous derivatives of this function make numerical solution more challenging; therefore, from the perspective purely of computational-performance, values of e > 1 are to be preferred.

In conclusion, for exploratory and developmental modelling, it may be preferable to use values of e close to, but greater than unity for improved computational performance. For situations where experimental data is not available, the simplified 'linear' elastic approach (i.e. e=1) has been shown to perform acceptably (for more detail on these linear elastic models with rigid-body modelling, the reader is referred to the thesis of Halloran [157]). When better accuracy is needed and good experimental data is available, a custom-fit non-linear model may be selected instead (albeit at the cost of computational-numerical performance).

REFERENCES

1. *World Population Prospects: The 1998 Revision, vol. II, Sex and Age Distribution of the World Population* 1998, The Population Division, Department of Economic and Social Affairs, United Nations Secretariat.
2. *National Joint Registry for England and Wales - 4th Annual Report*. 2007, National Joint Registry (UK).
3. Gray, H., *Anatomy of the Human Body 20th Edition*. 1918: Philadelphia: Lea & Febiger.
4. Richard L. Drake, W.V., Adam W.M. Mitchell., *Gray's anatomy for students*. 2005: Philadelphia; London:Elsevier/Churchill Livingstone.
5. Robinson, J.R., A.M. Bull, and A.A. Amis, *Structural properties of the medial collateral ligament complex of the human knee.* J Biomech, 2005. **38**(5): p. 1067-74.
6. Momersteeg, T.J., L. Blankevoort, R. Huiskes, J.G. Kooloos, J.M. Kauer, and J.C. Hendriks, *The effect of variable relative insertion orientation of human knee bone-ligament-bone complexes on the tensile stiffness.* J Biomech, 1995. **28**(6): p. 745-52.
7. Mommersteeg, T.J., L. Blankevoort, R. Huiskes, J.G. Kooloos, and J.M. Kauer, *Characterization of the mechanical behavior of human knee ligaments: a numerical-experimental approach.* J Biomech, 1996. **29**(2): p. 151-60.
8. Escalante, A., M.J. Lichtenstein, R. Dhanda, J.E. Cornell, and H.P. Hazuda, *Determinants of hip and knee flexion range: results from the San Antonio Longitudinal Study of Aging*. Arthritis Care Res, 1999. **12**(1): p. 8-18.
9. Roach, K.E. and T.P. Miles, *Normal hip and knee active range of motion: the relationship to age.* Phys Ther, 1991. **71**(9): p. 656-65.
10. Hallen, L.G. and O. Lindahl, *The "screw-home" movement in the knee-joint.* Acta Orthop Scand, 1966. **37**(1): p. 97-106.
11. Zuppinger, H., *Die aktive flexion im unbelasteten Kniegelenk: Züricher Habilitationsschrift*. 1904, Bergmann: Wiesbaden. p. 703-63.
12. Pinskerova, V., P. Johal, S. Nakagawa, A. Sosna, A. Williams, W. Gedroyc, and M.A. Freeman, *Does the femur roll-back with flexion?* J Bone Joint Surg Br, 2004. **86**(6): p. 925-31.
13. Banks, S.A., G.D. Markovich, and W.A. Hodge, *In vivo kinematics of cruciate-retaining and -substituting knee arthroplasties.* J Arthroplasty, 1997. **12**(3): p. 297-304.
14. Stiehl, J.B., D.A. Dennis, R.D. Komistek, and H.S. Crane, *In vivo determination of condylar lift-off and screw-home in a mobile-bearing total knee arthroplasty*. J Arthroplasty, 1999. **14**(3): p. 293-9.
15. Dennis, D.A., R.D. Komistek, C.E. Colwell, Jr., C.S. Ranawat, R.D. Scott, T.S. Thornhill, and M.A. Lapp, *In vivo anteroposterior femorotibial translation of total knee arthroplasty: a multicenter analysis.* Clin Orthop Relat Res, 1998(356): p. 47-57.
16. Freeman, M.A.R. and V. Pinskerova, *The movement of the normal tibio-femoral joint*. Journal of Biomechanics, 2005. **38**(2): p. 197-208.
17. Grood, E.S. and W.J. Suntay, *A joint coordinate system for the clinical description of three-dimensional motions: application to the knee.* J Biomech Eng, 1983. **105**(2): p. 136-44.
18. Morlock, M., E. Schneider, A. Bluhm, M. Vollmer, G. Bergmann, V. Muller, and M. Honl, *Duration and frequency of every day activities in total hip patients*. J Biomech, 2001. **34**(7): p. 873-81.
19. Seedhom, B.B. and N.C. Wallbridge, *Walking activities and wear of prostheses.* Ann Rheum Dis, 1985. **44**(12): p. 838-43.
20. Schmalzried, T.P., E.S. Szuszczewicz, M.R. Northfield, K.H. Akizuki, R.E. Frankel, G. Belcher, and H.C. Amstutz, *Quantitative assessment of walking*

activity after total hip or knee replacement. J Bone Joint Surg Am, 1998. **80**(1): p. 54-9.
21. Fregly, B.J., W.G. Sawyer, M.K. Harman, and S.A. Banks, *Computational wear prediction of a total knee replacement from in vivo kinematics.* J Biomech, 2005. **38**(2): p. 305-14.
22. Lafortune, M.A., P.R. Cavanagh, H.J. Sommer, 3rd, and A. Kalenak, *Three-dimensional kinematics of the human knee during walking.* J Biomech, 1992. **25**(4): p. 347-57.
23. *ISO 14243-1:2002 Implants for surgery. Wear of total knee joint prostheses. Loading and displacement parameters for wear-testing machines with load control and corresponding environmental conditions for test*. 2002, The International Organization for Standardization.
24. *ISO 14243-3:2004 Implants for surgery. Wear of total knee joint prostheses. Loading and displacement parameters for wear-testing machines with displacement control and corresponding environmental conditions for test*. 2004, The International Organization for Standardization.
25. Morrison, J.B., *The mechanics of the knee joint in relation to normal walking.* J Biomech, 1970. **3**(1): p. 51-61.
26. Andriacchi, T.P., G.B. Andersson, R.W. Fermier, D. Stern, and J.O. Galante, *A study of lower-limb mechanics during stair-climbing.* J Bone Joint Surg Am, 1980. **62**(5): p. 749-57.
27. Nagura, T., C.O. Dyrby, E.J. Alexander, and T.P. Andriacchi, *Mechanical loads at the knee joint during deep flexion.* J Orthop Res, 2002. **20**(4): p. 881-6.
28. McFadyen, B.J. and D.A. Winter, *An integrated biomechanical analysis of normal stair ascent and descent.* J Biomech, 1988. **21**(9): p. 733-44.
29. Costigan, P.A., K.J. Deluzio, and U.P. Wyss, *Knee and hip kinetics during normal stair climbing.* Gait Posture, 2002. **16**(1): p. 31-7.
30. Bergmann, G., F. Graichen, and A. Rohlmann, *Hip joint loading during walking and running, measured in two patients.* J Biomech, 1993. **26**(8): p. 969-90.
31. Taylor, S.J., J.S. Perry, J.M. Meswania, N. Donaldson, P.S. Walker, and S.R. Cannon, *Telemetry of forces from proximal femoral replacements and relevance to fixation.* J Biomech, 1997. **30**(3): p. 225-34.
32. Taylor, S.J. and P.S. Walker, *Forces and moments telemetered from two distal femoral replacements during various activities.* J Biomech, 2001. **34**(7): p. 839-48.
33. Taylor, S.J., P.S. Walker, J.S. Perry, S.R. Cannon, and R. Woledge, *The forces in the distal femur and the knee during walking and other activities measured by telemetry.* J Arthroplasty, 1998. **13**(4): p. 428-37.
34. Kaufman, K.R., N. Kovacevic, S.E. Irby, and C.W. Colwell, *Instrumented implant for measuring tibiofemoral forces.* J Biomech, 1996. **29**(5): p. 667-71.
35. D'Lima, D.D., S. Patil, N. Steklov, S. Chien, and C.W. Colwell, Jr., *In vivo knee moments and shear after total knee arthroplasty.* J Biomech, 2007. **40**(S1): p. S11-S17.
36. D'Lima, D.D., S. Patil, N. Steklov, J.E. Slamin, and C.W. Colwell, Jr., *The Chitranjan Ranawat Award: in vivo knee forces after total knee arthroplasty.* Clin Orthop Relat Res, 2005. **440**: p. 45-9.
37. D'Lima, D.D., S. Patil, N. Steklov, J.E. Slamin, and C.W. Colwell, Jr., *Tibial forces measured in vivo after total knee arthroplasty.* J Arthroplasty, 2006. **21**(2): p. 255-62.
38. Brown, T.D. and D.T. Shaw, *In vitro contact stress distribution on the femoral condyles.* J Orthop Res, 1984. **2**(2): p. 190-9.
39. Taylor, M. and D.S. Barrett, *Influence of uni-condylar loading on the stresses and kinematics on a total knee joint replacement*, in *International Society of Biomechanics XVIIIth Congress*. 2001: Zurich, Switzerland.

40. Taylor, W.R., M.O. Heller, G. Bergmann, and G.N. Duda, *Tibio-femoral loading during human gait and stair climbing.* J Orthop Res, 2004. **22**(3): p. 625-32.
41. Rand, J.A. and D.M. Ilstrup, *Survivorship analysis of total knee arthroplasty. Cumulative rates of survival of 9200 total knee arthroplasties.* J Bone Joint Surg Am, 1991. **73**(3): p. 397-409.
42. McEwen, H.M.J., J. Fisher, A.A.J. Goldsmith, D.D. Auger, C. Hardaker, and M.H. Stone, *Wear of fixed bearing and rotating platform mobile bearing knees subjected to high levels of internal and external tibial rotation.* Journal of Materials Science: Materials in Medicine, 2001. **12**(10-12): p. 1049-1052.
43. Muratoglu, O.K., H.E. Rubash, C.R. Bragdon, B.R. Burroughs, A. Huang, and W.H. Harris, *Simulated normal gait wear testing of a highly cross-linked polyethylene tibial insert.* J Arthroplasty, 2007. **22**(3): p. 435-44.
44. Barnett, P.I., J. Fisher, D.D. Auger, M.H. Stone, and E. Ingham, *Comparison of wear in a total knee replacement under different kinematic conditions.* J Mater Sci Mater Med, 2001. **12**(10-12): p. 1039-42.
45. *Physical activity and health: A report of the Surgeon General*. 1996, U.S. Department of Health and human services (DHHS).
46. Szoeke, C.E., F.M. Cicuttini, J.R. Guthrie, M.S. Clark, and L. Dennerstein, *Factors affecting the prevalence of osteoarthritis in healthy middle-aged women: data from the longitudinal Melbourne Women's Midlife Health Project.* Bone, 2006. **39**(5): p. 1149-55.
47. Coggon, D., I. Reading, P. Croft, M. McLaren, D. Barrett, and C. Cooper, *Knee osteoarthritis and obesity.* Int J Obes Relat Metab Disord, 2001. **25**(5): p. 622-7.
48. Norman, C.C. and T.A. Kress, *Risk factors for knee osteoarthritis: a review of epidemiological studies*, in *Proceedings of the IASTED International Conference on Biomechanics*, M.H. Hamza, Editor. 2003: Rhodes, Greece. p. 6-11.
49. Robertsson, O., *Knee arthroplasty registers.* J Bone Joint Surg Br, 2007. **89**(1): p. 1-4.
50. *Australian Orthopaedic Association National Joint Replacement Registry. Annual Report.* . 2008, Australian Orthopaedic Association (AOA).
51. *2007 Report - Hip and Knee Replacements in Canada*. 2007, Canadian Joint Replacement Registry (CJRR).
52. *Dansk Knæalloplastik Register (2004-2005)*. 2007, Dansk Selskab for Hofte- og Knæalloplastikkirurgi Dansk Ortopædisk Selskab.
53. *The 2006 Implant Yearbook on Orthopaedic Endoprostheses*. 2006, Finnish Arthroplasty Register.
54. *National Joint Registry - Eight Year Report (Jan 1999 - Dec 2006)*. 2006, New Zealand Orthopaedic Association.
55. *The Norwegian Arthroplasty Register - Report 2008*. 2008, The Norwegian Arthroplasty Register (Helse-Bergen HF, Department of Orthopaedic Surgery, Haukeland University Hospital).
56. *Scottish Arthroplasty Project Annual Report 2008*. 2008, Scottish Arthroplasty Project (NHS Scotland).
57. *Swedish Knee Arthroplasty Register – Annual Report 2008*. 2008, Swedish Knee Arthroplasty Register (Dept. of Orthopedics, Lund University Hospital).
58. Walton, M.J., A.E. Weale, and J.H. Newman, *The progression of arthritis following lateral unicompartmental knee replacement.* Knee, 2006. **13**(5): p. 374-7.
59. Robinson, R.P., *The early innovators of today's resurfacing condylar knees.* J Arthroplasty, 2005. **20**(1 Suppl 1): p. 2-26.
60. Ducheyne, P., A. Kagan, 2nd, and J.A. Lacey, *Failure of total knee arthroplasty due to loosening and deformation of the tibial component.* J Bone Joint Surg Am, 1978. **60**(3): p. 384-91.

61. Fisher, J., H.M. McEwen, P.I. Barnett, C. Bell, M.H. Stone, and E. Ingham, *Influences of sterilising techniques on polyethylene wear.* Knee, 2004. **11**(3): p. 173-6.
62. Kim, Y.H., S.H. Yoon, and J.S. Kim, *The long-term results of simultaneous fixed-bearing and mobile-bearing total knee replacements performed in the same patient.* J Bone Joint Surg Br, 2007. **89**(10): p. 1317-23.
63. Pagnano, M.W. and R.M. Menghini, *Rotating platform knees: an emerging clinical standard: in opposition.* J Arthroplasty, 2006. **21**(4 Suppl 1): p. 37-9.
64. Jacobs, W.C., D.J. Clement, and A.B. Wymenga, *Retention versus sacrifice of the posterior cruciate ligament in total knee replacement for treatment of osteoarthritis and rheumatoid arthritis.* Cochrane Database Syst Rev, 2005(4): p. CD004803.
65. Casey, D., J. Cottrell, E. Dicarlo, R. Windsor, and T. Wright, *PFC Knee Replacement: Osteolytic Failures From Extreme Polyethylene Degradation.* Clin Orthop Relat Res, 2007.
66. Chockalingam, S. and G. Scott, *The outcome of cemented vs. cementless fixation of a femoral component in total knee replacement (TKR) with the identification of radiological signs for the prediction of failure.* Knee, 2000. **7**(4): p. 233-238.
67. Cameron, H.U. and G.A. Hunter, *Failure in total knee arthroplasty: mechanisms, revisions, and results.* Clin Orthop Relat Res, 1982(170): p. 141-6.
68. Waslewski, G.L., B.M. Marson, and J.B. Benjamin, *Early, incapacitating instability of posterior cruciate ligament-retaining total knee arthroplasty.* J Arthroplasty, 1998. **13**(7): p. 763-7.
69. Meding, J.B., E.M. Keating, M.A. Ritter, P.M. Faris, and M.E. Berend, *Long-term followup of posterior-cruciate-retaining TKR in patients with rheumatoid arthritis.* Clin Orthop Relat Res, 2004(428): p. 146-52.
70. Abernethy, P.J., C.M. Robinson, and R.M. Fowler, *Fracture of the metal tibial tray after Kinematic total knee replacement. A common cause of early aseptic failure.* J Bone Joint Surg Br, 1996. **78**(2): p. 220-5.
71. *ISO 14879-1:2000 Implants for surgery. Total knee-joint prostheses. Determination of endurance properties of knee tibial trays*. 2000, The International Organization for Standardization.
72. Kim, J., C.L. Nelson, and P.A. Lotke, *Stiffness after total knee arthroplasty. Prevalence of the complication and outcomes of revision.* J Bone Joint Surg Am, 2004. **86-A**(7): p. 1479-84.
73. Babkin, Y., D. Raveh, M. Lifschitz, M. Itzchaki, Y. Wiener-Well, P. Kopuit, Z. Jerassy, and A.M. Yinnon, *Incidence and risk factors for surgical infection after total knee replacement.* Scand J Infect Dis, 2007: p. 1-6.
74. Murray, D.W. and S.J. Frost, *Pain in the assessment of total knee replacement.* J Bone Joint Surg Br, 1998. **80**(3): p. 426-31.
75. *National Institute of Health Consensus Guidelines: Total knee replacement*. 2004, US Department of Health and Human Services (DHHS).
76. Dawson, J., R. Fitzpatrick, D. Murray, and A. Carr, *Questionnaire on the perceptions of patients about total knee replacement.* J Bone Joint Surg Br, 1998. **80**(1): p. 63-9.
77. Insall, J.N., L.D. Dorr, R.D. Scott, and W.N. Scott, *Rationale of the Knee Society clinical rating system.* Clin Orthop Relat Res, 1989(248): p. 13-4.
78. Strasser, H., *lehrbuch der muskel und gelenkmechanik, III (Textbook of the Muscles and Joint Mechanics)*. 1917: Springer, Berlin.
79. Moeinzadeh, M.H., A.E. Engin, and N. Akkas, *Two-dimensional dynamic modelling of human knee joint.* J Biomech, 1983. **16**(4): p. 253-64.
80. Beynnon, B., J. Yu, D. Huston, B. Fleming, R. Johnson, L. Haugh, and M.H. Pope, *A sagittal plane model of the knee and cruciate ligaments with application of a sensitivity analysis.* J Biomech Eng, 1996. **118**(2): p. 227-39.

81. O'Connor, J.J., T.L. Shercliff, E. Biden, and J.W. Goodfellow, *The geometry of the knee in the sagittal plane.* Proc Inst Mech Eng [H], 1989. **203**(4): p. 223-33.
82. Wismans, J., F. Veldpaus, J. Janssen, A. Huson, and P. Struben, *A three-dimensional mathematical model of the knee-joint.* J Biomech, 1980. **13**(8): p. 677-85.
83. Caruntu, D.I. and M.S. Hefzy, *3-D anatomically based dynamic modeling of the human knee to include tibio-femoral and patello-femoral joints.* J Biomech Eng, 2004. **126**(1): p. 44-53.
84. Bei, Y. and B.J. Fregly, *Multibody dynamic simulation of knee contact mechanics.* Med Eng Phys, 2004. **26**(9): p. 777-89.
85. Maxian, T.A., T.D. Brown, D.R. Pedersen, and J.J. Callaghan, *Adaptive finite element modeling of long-term polyethylene wear in total hip arthroplasty.* J Orthop Res, 1996. **14**(4): p. 668-75.
86. Askew, M.J. and J.L. Lewis, *Analysis of model variables and fixation post length effects on stresses around a prosthesis in the proximal tibia.* J Biomech Eng, 1981. **103**(4): p. 239-45.
87. Lewis, J.L., M.J. Askew, and D.P. Jaycox, *A comparative evaluation of tibial component designs of total knee prostheses.* J Bone Joint Surg Am, 1982. **64**(1): p. 129-35.
88. Sathasivam, S. and P.S. Walker, *Optimization of the bearing surface geometry of total knees.* J Biomech, 1994. **27**(3): p. 255-64.
89. Bartel, D.L., J.J. Rawlinson, A.H. Burstein, C.S. Ranawat, and W.F. Flynn, Jr., *Stresses in polyethylene components of contemporary total knee replacements.* Clin Orthop Relat Res, 1995(317): p. 76-82.
90. Elbert, K., D. Bartel, and T. Wright, *The effect of conformity on stresses in dome-shaped polyethylene patellar components.* Clin Orthop Relat Res, 1995(317): p. 71-5.
91. Godest, A.C., M. Beaugonin, E. Haug, M. Taylor, and P.J. Gregson, *Simulation of a knee joint replacement during a gait cycle using explicit finite element analysis.* J Biomech, 2002. **35**(2): p. 267-75.
92. Knight, L.A., S. Pal, J.C. Coleman, F. Bronson, H. Haider, D.L. Levine, M. Taylor, and P.J. Rullkoetter, *Comparison of long-term numerical and experimental total knee replacement wear during simulated gait loading.* J Biomech, 2007. **40**(7): p. 1550-8.
93. Willing, R.T. and I.Y. Kim, *A Method for Optimizing TKR for Reduced Wear*, in *Transactions of the 54th Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
94. Halloran, J.P., S.K. Easley, A.J. Petrella, and P.J. Rullkoetter, *Comparison of deformable and elastic foundation finite element simulations for predicting knee replacement mechanics.* J Biomech Eng, 2005. **127**(5): p. 813-8.
95. Laz, P.J., S. Pal, J.P. Halloran, A.J. Petrella, and P.J. Rullkoetter, *Probabilistic finite element prediction of knee wear simulator mechanics.* J Biomech, 2006. **39**(12): p. 2303-10.
96. Maher, S.A., J.D. Lipman, L.J. Curley, M. Gilchrist, and T.M. Wright, *Mechanical performance of ceramic acetabular liners under impact conditions.* The Journal of Arthroplasty, 2003. **18**(7): p. 936-941.
97. Walker, P.S., *Friction of Internal Artificial Joints.* Engineering in Medicine, 1972. **1**(4): p. 91-92.
98. Yu, T.-C., C.-H. Huang, C.-H. Hsieh, J.-J. Liau, C.-H. Huang, and C.-K. Cheng, *Fatigue resistance analysis of tibial baseplate in total knee prosthesis--An in vitro biomechanical study.* Clinical Biomechanics, 2006. **21**(2): p. 147-151.
99. Humphreys, P.K., J.F. Orr, and A.S. Bahrani, *Testing of total hip replacements. Endurance tests and stress measurements. Part 1. Endurance tests.* Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine, 1990. **204**(1): p. 29-34.

100. *ISO 7206-4:2002 Implants for surgery. Partial and total hip joint prostheses Part 4: Determination of endurance properties of stemmed femoral components*. 2002, The International Organization for Standardization.
101. Mow, V.C. and R. Huiskes, *Basic orthopaedic biomechanics and mechano-biology*. 3 ed. 2004: Lippincott Williams & Wilkins.
102. Park, J. and R.S. Lakes, *Biomaterials: An Introduction (3rd Edition)*. 2007: Springer.
103. Besong, A.A., J.L. Tipper, E. Ingham, M.H. Stone, B.M. Wroblewski, and J. Fisher, *Quantitative comparison of wear debris from UHMWPE that has and has not been sterilised by gamma irradiation.* J Bone Joint Surg Br, 1998. **80**(2): p. 340-4.
104. Lancaster, J.G., D. Dowson, G.H. Isaac, and J. Fisher, *The wear of ultra-high molecular weight polyethylene sliding on metallic and ceramic counterfaces representative of current femoral surfaces in joint replacement.* Proc Inst Mech Eng [H], 1997. **211**(1): p. 17-24.
105. Turell, M., A. Wang, and A. Bellare, *Quantification of the effect of cross-path motion on the wear rate of ultra-high molecular weight polyethylene.* Wear, 2003. **255**(7-12): p. 1034-1039.
106. Ernsberger, C., D. Whitaker, and J. Chavarria, *UHMWPE Wear Rate As A Function Of Contact Area And Stress*, in *Transactions of the 53rd Annual Meeting, Orthopaedic Research Society* 2007: San Diego, CA.
107. Mazzucco, D. and M. Spector, *Effects of contact area and stress on the volumetric wear of ultrahigh molecular weight polyethylene.* Wear, 2003. **254**(5-6): p. 514-522.
108. Goodfellow, J. and J. O'Connor, *The mechanics of the knee and prosthesis design.* J Bone Joint Surg Br, 1978. **60-B**(3): p. 358-69.
109. Guess, T.M. and L.P. Maletsky, *Computational modeling of a dynamic knee simulator for reproduction of knee loading.* J Biomech Eng, 2005. **127**(7): p. 1216-21.
110. Werner, F., D. Foster, and D.G. Murray, *The influence of design on the transmission of torque across knee prostheses.* J Bone Joint Surg Am, 1978. **60**(3): p. 342-8.
111. Thatcher, J.C., X.M. Zhou, and P.S. Walker, *Inherent laxity in total knee prostheses.* J Arthroplasty, 1987. **2**(3): p. 199-207.
112. Barnett, P.I., H.M.J. McEwen, D.D. Auger, M.H. Stone, E. Ingham, and J. Fisher, *Investigation of wear of knee prostheses in a new displacement/force-controlled simulator.* Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine, 2002. **216**(1): p. 51-61.
113. Walker, P.S. and H.H. Hsieh, *Conformity in condylar replacement knee prosthesis.* J Bone Joint Surg Br, 1977. **59**(2): p. 222-8.
114. Walker, P.S., G.W. Blunn, D.R. Broome, J. Perry, A. Watkins, S. Sathasivam, M.E. Dewar, and J.P. Paul, *A knee simulating machine for performance evaluation of total knee replacements.* J Biomech, 1997. **30**(1): p. 83-9.
115. van Houtem, M., R. Clough, A. Khan, M. Harrison, and G.W. Blunn, *Validation of the soft tissue restraints in a force-controlled knee simulator.* Proc Inst Mech Eng [H], 2006. **220**(3): p. 449-56.
116. Shaw, J.A. and D.G. Murray, *Knee joint simulator.* Clin Orthop Relat Res, 1973(94): p. 15-23.
117. Perry, J., D. Antonelli, and W. Ford, *Analysis of knee-joint forces during flexed-knee stance.* J Bone Joint Surg Am, 1975. **57**(7): p. 961-7.
118. Bourne, R., J. Goodfellow, and J. O'Connor, *A functional analysis of various knee arthroplasties.* Transactions, Orthopaedic Research Society, 1978. **24**: p. 160.
119. Zavatsky, A.B., *A kinematic-freedom analysis of a flexed-knee-stance testing rig.* J Biomech, 1997. **30**(3): p. 277-80.

120. Zachman, N.J., B.M. Hillberry, and D.B. Kettelkamp, *Design of a Load Simulator for the Dynamic Evaluation of Prosthetic Knee Joints.* American Society of Mechanical Engineers (Paper), 1978(78-DET-59): p. 11-11.
121. Maletsky, L.P. and B.M. Hillberry, *Simulating dynamic activities using a five-axis knee simulator.* J Biomech Eng, 2005. **127**(1): p. 123-33.
122. Reeve, A.N., M.A. Strickland, L.P. Maletsky, and M. Taylor. *Verification of a dynamic knee simulator computational model*. in *Proceedings of the ASME 2008 Summer Bioengineering Conference (SBC2008)*. 2008. Marriott Resort, Marco Island, Florida, USA.
123. Guess, T.M. and L.P. Maletsky, *Computational modelling of a total knee prosthetic loaded in a dynamic knee simulator.* Med Eng Phys, 2005. **27**(5): p. 357-67.
124. Hamilton, M.A., M.C. Sucec, B.J. Fregly, S.A. Banks, and W.G. Sawyer, *Quantifying multidirectional sliding motions in total knee replacements.* ASME. Journal of Tribology, 2005. **127**(2): p. 280-6.
125. Kang, L., A.L. Galvin, T.D. Brown, Z. Jin, and J. Fisher, *Quantification of the effect of cross-shear on the wear of conventional and highly cross-linked UHMWPE.* J Biomech, 2008. **41**(2): p. 340-6.
126. Knight, L.A., H. McEwen, J. Fisher, and M. Taylor. *Influence of cross shear on the wear of TKA under various kinematic conditions*. in *52nd Annual Meeting of the Orthopaedic Research Society*. 2006. Chicago, USA: Orthopaedic Research Society.
127. Willing, R.T. and I.Y. Kim, *Validation of a Computational UHMWPE Damage Model for TKR Experiment Simulation*, in *Transactions of the 54th Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
128. Rubinstein, R.Y., *Simulation and the Monte Carlo Method (Wiley Series in Probability and Statistics)*. 1981: Wiley-Interscience
129. Wing Kam Liu , T.B., *Computational Mechanics of Probabilistic and Reliability Analysis*. 1989: Elmepress International.
130. Melchers, R.E., *Structural Reliability Analysis and Prediction, 2nd Edition*. 1999: Wiley
131. *NESSUS Theoretical Manual (Version 7.0)*. 2001, Southwest Research Institute.
132. Keane, A.J. and P.B. Nair, *Computational Approaches for Aerospace Design: The Pursuit of Excellence*. 2005: John Wiley & Sons, Ltd.
133. Madsen, H.O., S. Krenk, and N.C. Lind, *Methods of Structural Safety*. 1986: Prentice-Hall.
134. Perillo-Marcone, A. and M. Taylor, *Effect of Varus/Valgus Malalignment on Bone Strains in the Proximal Tibia After TKR: An Explicit Finite Element Study.* J Biomech Eng, 2007. **129**(1): p. 1-11.
135. Taylor, M. and D.S. Barrett, *Explicit finite element simulation of eccentric loading in total knee replacement.* Clin Orthop Relat Res, 2003(414): p. 162-71.
136. Knight, L.A. and M. Taylor, *The effect of eccentric loading on the wear of total knee arthroplasty*, in *Transactions of the 53rd Annual Meeting, Orthopaedic Research Society* 2007: San Diego, CA.
137. Bei, Y., Dynamic simulation of knee joint contact during human movement, 2003, Thesis submitted for Doctor of Philosophy, Department of Mechanical and Aerospace Engineering, University of Florida
138. Box, G.E.P. and K.B. Wilson, *On the experimental attainment of optimum conditions.* Journal of the Royal Statistical Society Series B, 1951. **13**(1): p. 1-45.
139. Lin, Y.-C., R.T. Haftka, N.V. Queipe, and B.J. Fregly, *A generalized surrogate contact model for dynamic simulations with anatomic joints*, in *Proceedings of the 2006 Summer Bioengineering Conference*. 2006, The American Society of Mechanical Engineers, New York: Amelia Island, Florida, USA.

140. Langley, R.S. *The dynamic analysis of uncertain structures*. in *7th International Conference on recent advances in structural dynamics*. 2000. University of Southampton, UK.
141. Langley, R.S. and C.S. Ribas, *Application of level II reliability theory.* Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering, 1995. **209**(4): p. 291-298.
142. Browne, M., R.S. Langley, and P.J. Gregson, *Reliability theory for load bearing biomedical implants.* Biomaterials, 1999. **20**(14): p. 1285-92.
143. Perez, M.A., J. Grasa, J.M. Garcia-Aznar, J.A. Bea, and M. Doblare, *Probabilistic analysis of the influence of the bonding degree of the stem-cement interface in the performance of cemented hip prostheses.* J Biomech, 2006. **39**(10): p. 1859-72.
144. Mehrez, L., A. New, M.T. Bah, and M. Browne, *Comparative study of probabilistic methods applied to a 3D finite element model of a THR.* J Biomech, 2006. **39**(S1): p. S67-S67.
145. Laz, P.J., S. Pal, A. Fields, A.J. Petrella, and P.J. Rullkoetter, *Effects of knee simulator loading and alignment variability on predicted implant mechanics: a probabilistic study.* J Orthop Res, 2006. **24**(12): p. 2212-21.
146. Eichmiller, F.C., J.A. Tesk, and C.M. Croarkin, *Mechanical Properties of UHMWPE (0.5cm cubes) NIST Reference Material RM 8457. Transactions of the Society for Biomaterials; 27th Annual Meeting*. 2001: St Paul, MN.
147. Eichmiller, F.C., J.A. Tesk, and C.M. Croarkin, *Mechanical Properties of Ultra High Molecular Weight Polyethylene NIST Reference Material RM 8456. Transactions of the Society for Biomaterials; p. 472, 27th Annual Meeting*. 2001: St Paul, MN. p. 472.
148. Jenny, J.Y. and C. Boeri, *Navigated implantation of total knee endoprostheses--a comparative study with conventional instrumentation.* Z Orthop Ihre Grenzgeb, 2001. **139**(2): p. 117-9.
149. Mahaluxmivala, J., M.J. Bankes, P. Nicolai, C.H. Aldam, and P.W. Allen, *The effect of surgeon experience on component positioning in 673 Press Fit Condylar posterior cruciate-sacrificing total knee arthroplasties.* J Arthroplasty, 2001. **16**(5): p. 635-40.
150. Sparmann, M., B. Wolke, H. Czupalla, D. Banzer, and A. Zink, *Positioning of total knee arthroplasty with and without navigation support. A prospective, randomised study.* J Bone Joint Surg Br, 2003. **85**(6): p. 830-5.
151. *US National Health and Nutrition Examination Survey (NHANES)*. 2007, National Center for Health Statistics.
152. Claes, L.E., A. Beyer, W. Krischke, and R. Schmid, *Biomechanical properties of collateral and cruciate ligaments.* 1987: University of Ulm, Ulm, West Germany. p. 22.
153. Mommersteeg, T.J., L. Blankevoort, R. Huiskes, J.G. Kooloos, J.M. Kauer, and J.C. Hendriks, *The effect of variable relative insertion orientation of human knee bone-ligament-bone complexes on the tensile stiffness.* J Biomech, 1995. **28**(6): p. 745-52.
154. Trent, P.S., P.S. Walker, and B. Wolf, *Ligament length patterns, strength and rotational axes of the knee joint. .* Clinical Orthopaedics, 1976. **117**: p. 263-270.
155. Wasmer, G., F.-W. Hagena, T.B. Mittlemeier, M. , and G. Hofman. in *Annual Meeting, European Society of Biomechanics*. 1987. University of Ulm, Ulm, West Germany.
156. Halloran, J.P., A.J. Petrella, and P.J. Rullkoetter, *Explicit finite element modeling of total knee replacement mechanics.* J Biomech, 2005. **38**(2): p. 323-31.
157. Halloran, J.P., Explicit Finite Element Modeling of Joint Replacement Mechanics, 2007, Thesis submitted for Doctor of Philosophy, Department of Mechanical Engineering, University of Denver
158. Haider, H. and P.S. Walker. *Analysis And Recommendations For The Optimum Spring Configurations For Soft Tissue Restraint In Force-Control*

Knee Simulator Testing. in *48th Annual Meeting of the Orthopaedic Research Society*. 2002. Dallas, Texas: Orthopaedic Research Society.
159. DesJardins, J.D., P.S. Walker, H. Haider, and J. Perry, *The use of a force-controlled dynamic knee simulator to quantify the mechanical performance of total knee replacement designs during functional activity.* J Biomech, 2000. **33**(10): p. 1231-42.
160. Fregly, B.J., Y. Bei, and M.E. Sylvester, *Experimental evaluation of an elastic foundation model to predict contact pressures in knee replacements.* J Biomech, 2003. **36**(11): p. 1659-68.
161. *MSC.ADAMS Help Library (Version 2005r2)*. 2005, MSC Software Corporation.
162. Wang, A., *A unified theory of wear for ultra-high molecular weight polyethylene in multi-directional sliding.* Wear, 2001. **248**(1-2): p. 38-47.
163. Archard, J.F., *Contact and Rubbing of Flat Surfaces.* Journal of Applied Physics, 1953. **24**(8): p. 981-988.
164. Maxian, T.A., T.D. Brown, D.R. Pedersen, and J.J. Callaghan, *3-Dimensional sliding/contact computational simulation of total hip wear.* Clin Orthop Relat Res, 1996(333): p. 41-50.
165. Maxian, T.A., T.D. Brown, D.R. Pedersen, and J.J. Callaghan, *A sliding-distance-coupled finite element formulation for polyethylene wear in total hip arthroplasty.* J Biomech, 1996. **29**(5): p. 687-92.
166. Maxian, T.A., T.D. Brown, D.R. Pedersen, H.A. McKellop, B. Lu, and J.J. Callaghan, *Finite element analysis of acetabular wear. Validation, and backing and fixation effects.* Clin Orthop Relat Res, 1997(344): p. 111-7.
167. *Report of Investigation, Reference Material 8456: Ultra High Molecular Weight Polyethylene*. 2003, National Institute of Standards & Technology (NIST, U.S. Dept. of Commerce).
168. Gevaert, M.R., M. LaBerge, J.M. Gordon, and J.D. DesJardins, *The quantification of physiologically relevant cross-shear wear phenomena on orthopaedic bearing materials using the MAX-Shear wear testing system.* Journal of Tribology, 2005. **127**(4): p. 740-749.
169. Fisher, J., H.M.J. McEwen, P.I. Barnett, C.J. Bell, T.D. Stewart, M.H. Stone, and E. Ingham, *(i) Wear of polyethylene in artificial knee joints.* Current Orthopaedics, 2001. **15**(6): p. 399-405.
170. Willing, R.T. and I.Y. Kim, *A Pseudo-Qualitative Method for Measuring Cross-Shearing motions in Total Knee Replacements*, in *Transactions of the 54th Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
171. Kang, L., A.L. Galvin, Z. Jin, and J. Fisher, *Enhanced Computational Prediction of UHMWPE Wear by Incorporating Cross-shear and Contact Pressure into Archard Theory*, in *Transactions of the 54rd Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
172. Jennings, L.M., C.J. Bell, E. Ingham, R.D. Komistek, M.H. Stone, and J. Fisher, *The influence of femoral condylar lift-off on the wear of artificial knee joints.* Proc. IMechE Part H: J. Engineering in Medicine, 2007. **221**: p. 305-314.
173. McEwen, H.M., P.I. Barnett, C.J. Bell, R. Farrar, D.D. Auger, M.H. Stone, and J. Fisher, *The influence of design, materials and kinematics on the in vitro wear of total knee replacements.* J Biomech, 2005. **38**(2): p. 357-65.
174. Haider, H., R.E. Croson, and K.L. Garvin, *Is wear truly lower and is it the main benefit of rotating platform mobile bearing total knees?*, in *Transactions of the 54rd Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
175. Cottrell, J.M., O. Babalola, B.S. Furman, and T.M. Wright, *Stair ascent kinematics affect UHMWPE wear and damage in total knee replacements.* J Biomed Mater Res B Appl Biomater, 2006. **78**(1): p. 15-9.
176. Haider, H., J.N. Weisenburger, R.E. Croson, F. Namavar, and K.L. Garvin, *Concern with adhesion and wear of a Titanium Niobium Nitride coating on*

Total Knee Replacements for metal sensitive patients, in *Transactions of the 54rd Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
177. Williams, P.A., R. Tsukamoto, and I.C. Clarke, *Wear Debris from Sequentially Crosslinked and Crosslinked PE in a Knee Simulator Model*, in *Transactions of the 54rd Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
178. Galvin, A.L., L.M. Jennings, L. Kang, H. McEwen, and J. Fisher, *A Low Conforming, Low Wear Solution in Fixed Bearing Total Knee Prostheses*, in *Transactions of the 54rd Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
179. Pal, S., H. Haider, P. Laz, L. Knight, and P. Rullkoetter, *Probabilistic computational modeling of total knee replacement wear.* Wear, 2007. **264**(7-8): p. 701-707.
180. Essner, A., R. Klein, M. Bushelow, A. Wang, M. Kvitnitsky, and O. Mahoney, *The effect of sagittal conformity on knee wear.* Wear, 2003. **255**(7-12): p. 1085-1092.
181. Knight, L.A., Finite Element Simulation of Surface Wear in Total Knee Joint Replacement, 2009, Thesis submitted for Doctor of Philosophy, Bioengineering Sciences Research Group, University of Southampton
182. Zhao, D., S.A. Banks, D.D. D'Lima, C.W. Colwell, Jr., and B.J. Fregly, *In vivo medial and lateral tibial loads during dynamic and high flexion activities.* J Orthop Res, 2007. **25**(5): p. 593-602.
183. Abdel-Rahman, E.M. and M.S. Hefzy, *Three-dimensional dynamic behaviour of the human knee joint under impact loading.* Med Eng Phys, 1998. **20**(4): p. 276-90.
184. Yu, C.H., P.S. Walker, and M.E. Dewar, *The effect of design variables of condylar total knees on the joint forces in step climbing based on a computer model.* J Biomech, 2001. **34**(8): p. 1011-21.
185. Tumer, S.T. and A.E. Engin, *Three-body segment dynamic model of the human knee.* J Biomech Eng, 1993. **115**(4A): p. 350-6.
186. Shultz, S.J., Y. Shimokochi, A.D. Nguyen, R.J. Schmitz, B.D. Beynnon, and D.H. Perrin, *Measurement of varus-valgus and internal-external rotational knee laxities in vivo-Part II: relationship with anterior-posterior and general joint laxity in males and females.* J Orthop Res, 2007. **25**(8): p. 989-96.
187. Blankevoort, L., R. Huiskes, and A. de Lange, *The envelope of passive knee joint motion.* J Biomech, 1988. **21**(9): p. 705-20.
188. Sutton, L.G., F.W. Werner, T. Hamblin, and J. Clabeaux, *Does Knee Implant Wear Testing Reflect Normal Knee Motion and Loading?*, in *Transactions of the 54rd Annual Meeting, Orthopaedic Research Society*. 2008: San Francisco, CA.
189. Johnson, T.S., M.P. Laurent, J.Q. Yao, and L.N. Gilbertson, *The effect of displacement control input parameters on tibiofemoral prosthetic knee wear.* Wear, 2001. **250**(1-12): p. 222-226.
190. Fleming, B.C., B. Brattbakk, G.D. Peura, G.J. Badger, and B.D. Beynnon, *Measurement of anterior-posterior knee laxity: a comparison of three techniques.* J Orthop Res, 2002. **20**(3): p. 421-6.
191. Maletsky, L.M., R.A. Thibeault, and B.M. Hillberry. *Kinematics of the human knee during walking: A simulator study investigating soft tissue constraints*. in *Proceedings of ASME Bioengineering Division*. 2001. Snow Bird, Utah, USA: ASME.
192. Maletsky, L.P. and B.M. Hillberry, *Computer Modeling of Knee Simulator Tibio-Femoral and Patello-Femoral Loading.* Proceedings of the ASME Dynamic Systems and Control Division, 1997. **DSC 61**: p. 387-392.
193. Maletsky, L.P. and B.M. Hillberry, *Dynamic Control of a Four-Axis, Electrohydraulic Knee Simulator.* Proceedings of the ASME Dynamic Systems and Control Division, 1997. **DSC 61**: p. 447-452.

194. Maletsky, L.P. and B.M. Hillberry, *Loading Evaluation of Knee Joint During Walking Using the Next Generation Knee Simulator*, in *ASME Advances in Bioengineering,*, T.A. Conway, Editor. 2000, ASME: Orlando, Florida (USA). p. 91-92.
195. Rullkoetter, P., S. McGuan, and L. Maletsky. *Development and Verification of a Virtual Knee Simulator for TKR Evaluation*. in *45th Annual Meeting, Orthopaedic Research Society*. 1999. Anaheim, CA (USA): Orthopaedic Research Society.
196. Sheshadri, V.B., P.J. Rullkoetter, and B.M. Hillberry, *In-vitro Measurement of the Six Degree-of-Freedom Kinematics of the Human Knee During Simulated Gait.* Proceedings of the ASME Dynamic Systems and Control Division, 1997. **DSC 61**: p. 375-379.
197. Guess, T.M., Computational Modeling of a Dynamic Knee Simulator, 2003, Thesis submitted for Doctor of Philosophy, Department of Mechanical Engineering, University of Kansas
198. Guess, T.M. and L.P. Maletsky. *Computational modeling of a dynamic knee simulator for reproduction of joint loading*. in *2003 ASME International Mechanical Engineering Congress, Nov 15-21 2003*. 2003: American Society of Mechanical Engineers, New York, NY 10016-5990, United States.
199. Chao, E.Y.S., B.A. MacWilliams, and B. Chan. *Evaluation of a Dynamic Joint Simulator Using a Prosthetic Knee System*. in *Proceedings of ASME Bioengineering Division*. 1994. Chicago, IL, USA: ASME.
200. Lanovaz and Ellis, *Dynamic simulation of a displacement-controlled total knee replacement wear tester.* Proceedings of the I MECH E Part H Journal of Engineering in Medicine, 2008. **222**(5): p. 669-681.
201. Benson, L.C., J.D. DesJardins, M.K. Harman, and M. LaBerge, *Effect of stair descent loading on ultra-high molecular weight polyethylene wear in a force-controlled knee simulator.* Proc Inst Mech Eng [H], 2002. **216**(6): p. 409-18.
202. Ellison, P., D.C. Barton, C. Esler, H.M. McEwen, D.L. Shaw, M.H. Stone, and J. Fisher, *Wear and Creep of Replacement Patellofemoral Joints*, in *53rd Annual Meeting of the Orthopaedic Research Society* 2007: San Diego, CA.
203. Zhao, D., H. Sakoda, W.G. Sawyer, S.A. Banks, and B.J. Fregly. *Can Pin-on-Plate tests be used to predict joint-level wear in Knee Replacements?* in *2006 Summer Bioengineering Conference, June 21-25 2006*. 2006. Amelia Island Plantation, Amelia Island, Florida, USA: The American Society of Mechanical Engineers, New York.
204. Lanovaz, J., T. Edgecombe, C. Elliot, J. Bryant, and R. Ellis, *A Dynamic FE Model to Generate Input Data for a TKR Wear Tester*, in *Transactions of the 52nd Annual Meeting, Orthopaedic Research Society* 2006: Chicago, IL.
205. Dunn, A.C., S.A. Banks, D. Burris, and W.G. Sawyer, *Multi-Directional Sliding Friction With UHMWPE*, in *Transactions of the 53rd Annual Meeting, Orthopaedic Research Society* 2007: San Diego, CA.
206. Essner, A., L. Herrera, S.-S. Yau, A. Wang, J.H. Dumbleton, and M.T. Manley, *Sequentially Crosslinked and Annealed Uhmwpe Knee Wear Debris*, in *51st Annual Meeting of the Orthopaedic Research Society*. 2005, Orthopaedic Research Society: Washington D.C., USA.
207. Arsene, C., M.A. Strickland, P.J. Laz, and M. Taylor, *Comparison of two methods for probabilistic finite element analysis of total knee replacement*, in *8th International Symposium on Computer Methods in Biomechanics and Biomedical Engineering* 2008: Porto, Portugal.
208. Briscoe, A., M.A. Strickland, and M. Taylor, *Medial-lateral Loading And Wear In TKA*, in *16th Congress, European Society of Biomechanics*. 2008: Lucerne, Switzerland.
209. Strickland, M.A., M. Browne, and M. Taylor, *Probabilistic Computer-Aided Analysis of Variables Affecting the Performance of Total Knee Replacement*, in *Biomedical Futures 2006 – Musculoskeletal Biomechanics*.

2006, Royal Academy of Engineering - UK Focus for Biomedical Engineering: Durham, UK.
210. Strickland, M.A., M. Browne, and M. Taylor, *The Effect of Ligament Variability on TKR Performance – a Probabilistic Study*, in *Transactions of the 53rd Annual Meeting, Orthopaedic Research Society* 2007: San Diego, CA.
211. Strickland, M.A., M. Browne, and M. Taylor, *Influence of Wear Algorithm Formulation on Computational-Experimental Corroboration*, in *Transactions of the 54th Annual Meeting, Orthopaedic Research Society* 2008: San Francisco, CA.
212. Strickland, M.A., M. Dressler, T. Render, and M. Taylor, *Holistic Approaches to Pre-clinical TKR Analysis: Computationally-Enriched Experimental Testing*, in *Knee Arthroplasty 2009: From Early Intervention to Revision*. 2009, Institute of Mechanical Engineers: London, UK.
213. Strickland, M.A. and M. Taylor, *in-silico Predictions of TKR Robustness to Wear Variability: A Probabilistic Cross-Design Comparison* in *Transactions of the 55th Annual Meeting, Orthopaedic Research Society* 2009: Las Vegas, NV.
214. Dressler, M., C. Ernsberger, M. Strickland, M. Taylor, T. Render, and M. Heldreth, *Predicting UHMWPE Wear: Evidence for Rapid Decline in Wear Rates Following a Change in Sliding Direction* in *Transactions of the 55th Annual Meeting, Orthopaedic Research Society* 2009: Las Vegas, NV.
215. Petrella, A., V. Patel, P. Laz, and P. Rullkoetter, *A Generalized Model for Cross-Shear Wear in Joint Replacements*, in *Transactions of the 55th Annual Meeting, Orthopaedic Research Society* 2009: Las Vegas, NV.
216. Strickland, M.A. and M. Taylor, *In-silico wear prediction for knee replacements - methodology and corroboration.* Journal of Biomechanics, 2009. **In Press**.
217. Pal, S., Explicit Finite Element Modeling of Joint Mechanics, 2009, Thesis submitted for Doctor of Philosophy, Department of Mechanical Engineering, University of Denver
218. Blankevoort, L. and R. Huiskes, *Validation of a three-dimensional model of the knee.* J Biomech, 1996. **29**(7): p. 955-61.
219. Blankevoort, L., J.H. Kuiper, R. Huiskes, and H.J. Grootenboer, *Articular contact in a three-dimensional model of the knee.* Journal of Biomechanics, 1991. **24**(11): p. 1019-1031.
220. Laz, P.J., S. Pal, J.P. Halloran , A.J. Petrella, and P.J. Rullkoetter. *Probabilistic Finite Element Prediction Of Knee Wear Simulator Mechanics*. in *51st Annual Meeting of the Orthopaedic Research Society*. 2005. Washington, D.C.: Orthopaedic Research Society.
221. DeHeer, D.C. and B.M. Hillberry, *The Effect of Thickness and Nonlinear Material Behavior on Contact Stresses in Polyethylene Tibial Components*, in *38th Annual Meeting of the Orthopaedic Research Society*. 1992, Orthopaedic Research Society. p. 327.