UNIVERSITY OF CAMBRIDGE
Faculty of Mathematics

MATHEMATICAL TRIPOS

2025-26

GUIDE TO COURSES

IN PART IB

This booklet contains informal and non-technical descriptions of courses to be examined in Part IB in the academic year 2025-26, as well as summaries of learning outcomes.

This *Guide to Courses* is intended to supplement the more formal descriptions contained in the booklet *Schedules of Lecture Courses and Form of Examinations*.

These and other Faculty documents for students taking the Mathematical Tripos are available from the undergraduate pages on the Faculty's website at `https://www.maths.cam.ac.uk/undergrad/`

1 Introduction

Each lecture course in the Mathematical Tripos has an official syllabus, or *schedule*, that sets out formally, and in technical terms, the material to be covered. The *schedules* are listed in the booklet *Schedules of Lecture Courses and Form of Examinations* that is available for download at `https://www.maths.cam.ac.uk/undergrad/course/schedules.pdf`. The *Schedules* booklet is the definitive reference for matters of course content and assessment, for students, lecturers and examiners.

The present guide, by contrast, provides an *informal* description of each lecture course in Part IB. These descriptions are intended to be comprehensible without much prior knowledge, and to convey something of the flavour of each course. Summaries of the learning outcomes for each course are also included, along with some suggestions for preparatory reading, if appropriate.

The full learning outcome for Part IB is that you should understand the material described in the formal syllabuses given in the *Schedules* booklet and be able to apply it to the sorts of problems that can be found on Tripos papers from earlier years.

Changes to lecture courses since last year

Two new courses, IB Analysis II and IB Topological Spaces, will run in 2025-26, replacing the courses IB Analysis & Topology and IB Geometry from 2024-25.

2 The Structure of Part IB

The structure of Part IB may be summarised as follows:

- There are four courses of 24 lectures, eight courses of 16 lectures, three courses of 12 lectures, and an additional Computational Projects course (CATAM).

- Five courses are lectured in Michaelmas Term, building on the core material in Part IA, while eight courses are lectured in Lent Term, allowing more specialisation in preparation for Part II.

- Two of the 12-lecture courses are given in Easter Term and may be taken in either the first or second year (*Optimisation* and *Variational Principles*).

- The examination consists of four papers, with Section I ('short') questions and Section II ('long') questions spread as evenly as possible subject to
  - each 24-lecture course having two short questions and four long questions;
  - each 16-lecture course having two short questions and three long questions;
  - each 12-lecture course having two short questions and two long questions;
  - each course having at most one question of each type (long or short) but at least one question of either type on each paper.

  The precise distribution of questions can be found in the *Schedules* booklet.

- Only four short questions and six long questions may be attempted on each paper, with
  - each short (Section I) question marked out of 10 with one beta quality mark;
  - each long (Section II) question marked out of 20 with one quality mark, alpha or beta.

- The Computational Projects course carries 160 marks and no quality marks.

3 Choice of Courses

The Faculty Board has issued the following guidance:

> Part IB of the Mathematical Tripos provides a wide range of courses from which students should, in consultation with their Directors of Studies, make a selection based on their individual interests and preferred workload, bearing in mind that it is better to do a smaller number of courses thoroughly than to do many courses scrappily.

So, you are certainly not expected to take all the courses in Part IB, and the informal course descriptions below are intended to help you start thinking about your choices. It is important to choose courses that you will find rewarding, and to be aware of the consequences of your choices for options in Part II; to this end the *Schedules* booklet contains a table summarising the relationships between courses in Part II and those in Part IB.

In Part IA you were expected to follow four 24-lecture courses each term, i.e. to attend two lectures per day for two terms (total 192 lectures). If you were comfortable with that, then this might be a sensible target for Part IB. Some students prefer to take slightly fewer courses and learn them more thoroughly, while other students may choose to take more. You should check the distribution of questions on the four examination papers (given in the *Schedules* booklet) before making your final choices for revision.

4 CATAM and Preparatory Work

It is important to start work as early as possible on the computational projects (CATAM). You are strongly encouraged to complete the non-examinable Introductory Project 0.1 (see the CATAM manual) over the summer; this will give you valuable practice in programming as well as in producing a coherent write-up. A model answer for this project will be available in Michaelmas Term and comparison of this with your own write-up (before having to submit write-ups for real marks) should be instructive.

You are warned that project work can take much longer than you first expect, and rushing to complete things at the last moment is not a recipe for securing good marks, as well as being a major distraction from your other work, so it is good to get ahead. If you don't know how to program in MATLAB, then you should try to crack this as soon as possible.

Any mathematics that you do over the summer vacation will stand you in good stead for Part IB. It is suggested, first, that if your College expects you to have supervisions in Michaelmas Term on either of the IB Easter Term courses then make sure you are up-to-speed and ready for them, and secondly, that you might wish to do some preparatory reading for one of the Michaelmas Term courses e.g. Analysis & Topology or Quantum Mechanics. The books suggested below are intended to give an idea of the appropriate level and approach for each course. They should all be in your college library, and by browsing there you may find other sources which are just as helpful. More comprehensive reading lists are also given in the *Schedules* booklet.

5 Informal Description of Courses

Linear Algebra   Michaelmas, 24 lectures

The first-year course Vectors and Matrices includes a concrete introduction to vector spaces. Here, vector spaces are investigated from an abstract axiomatic point of view. This has two purposes: firstly to provide an introduction to abstract algebra in an already familiar context and secondly to provide a foundation for the study of infinite-dimensional vector spaces which are required for advanced courses in analysis and physics. One important application is to function spaces and differential and difference operators. A striking result is the Cayley-Hamilton theorem which says (roughly) that any square matrix satisfies the same equation as its eigenvalues (the characteristic equation).

The spaces studied for the first part of the course have nothing corresponding to length or angle. These are introduced by defining an inner product (i.e. a 'dot' product) on the vector space. This is generalised to the notion of a bilinear form ('lengths' do not have to be positive) and even further. There are direct applications to quantum mechanics and statistics.

The last part of the course covers the theory of bilinear and hermitian forms, and inner products on vector spaces. An important example is the quadratic form. The discussion of orthogonality of eigenvectors and properties of eigenvalues of Hermitian matrices has consequences in many areas of mathematics and physics, including quantum mechanics.

There are many suitable books on linear algebra: for example *Finite-dimensional Vector Spaces* by Halmos (Springer, 1974), Birkhoff and MacLane's *Algebra* (Macmillan, 1979) and Strang's *Linear Algebra*

(Academic Press, 1980).

Learning outcomes. By the end of this course, you should:

- understand the concepts of, and be able to prove results in the theory of, real and complex vector spaces;

- understand the concepts of, and be able to prove results in the theory of, linear maps between and endomorphisms of real and complex vector spaces, including the role of eigenvectors and eigenvalues and Jordan canonical form;

- understand, and be able to prove and apply, the Cayley-Hamilton theorem;

- understand, and be able to prove results in the theory of, dual vector spaces;

- understand bilinear forms and their connection with the dual space, and be able to derive their basic properties;

- know the theory of canonical forms for symmetric, alternating and hermitian forms, and be able to find them in simple cases;

- understand the theory of hermitian endomorphisms of a complex inner product space, and know and be able to apply the Gram-Schmidt orthogonalisation process.

Groups, Rings and Modules   Lent, 24 lectures

This course unites a number of useful and important algebraic and geometric ideas by developing three concepts which are fundamental in abstract algebra. Firstly there is the notion of a *group* which you met in Part IA Groups and which is found in so much of mathematics, both pure and applied. The basic concepts of group theory are recalled from the first year and then built upon, resulting in beautiful theorems that reveal much about the structure of finite groups.

Whereas a group has only one operation, a *ring* is a set that is equipped with two operations: that of addition and multiplication, such as the integers. The next third of the course develops this idea in a way that mirrors the approach to groups, as well as considering examples such as fields and the important case of a ring of polynomials in one, and in many, variables.

The last part of the course defines and deals with the notion of a *module*, which can be described as the immediate generalisation of a vector space where the scalars form a ring rather than a field. The advantage of this approach is that it allows proof of general results which can then be used to unify theorems in specific cases, as shown at the end of the course where applications to Jordan Normal Form are given, along with a proof of the classification of finitely generated abelian groups.

For an introduction to groups, J. F. Humphreys, *A course in group theory* (Oxford Science Publications) amongst others is very readable whereas B. Hartley and T. O. Hawkes, *Rings, Modules and Linear Algebra* (Chapman and Hall), although somewhat dry, contains nearly all of the rings part of the course and more than all of the material on modules.

The course also lays the foundations for most of the algebra options in Part II. In particular it is essential for Galois Theory, and highly desirable for areas such as Number Fields and Representation Theory.

Learning outcomes. By the end of this course, you should:

- have a firm understanding of the fundamental concepts of group theory and be comfortable applying these to groups of small order;

- know the definition of a ring, a field and an ideal, and be able to determine whether an ideal is principal, maximal or prime;

- be able to factorise elements in specific rings, including cases where factorisation is non-unique;

- understand the concept of a module and its application to finitely generated abelian groups.

Analysis II

<div align="right">Michaelmas, 24 lectures</div>

In the Analysis I course in Part IA, you encountered for the first time the rigorous mathematical study of the concepts of limit, continuity, differentiability and integrability, applied to functions of a single real variable. This course extends that study in two different ways. First, it introduces the important notion of uniform convergence, which helps to answer the basic question of when we can ensure continuity, differentiability and integrability are preserved when passing to limits of sequences of functions. Then the fundamental ideas of Analysis I are extended from the real line $\mathbb{R}$, first to finite-dimensional Euclidean spaces $\mathbb{R}^n$ and then to still more general 'metric spaces' whose 'points' may be objects such as functions or sets.

The advantages of this more general point of view are demonstrated using Banach's contraction mapping theorem, whose striking applications include a general existence and uniqueness theorem for solutions of differential equations, and the inverse function theorem, a result of fundamental importance.

If you wish to do some vacation reading, W.A. Sutherland's Introduction to Metric and Topological Spaces (OUP, 1975) provides a good introduction to analysis on more general spaces.

Learning outcomes. By the end of this course, you should:

- understand and be able to prove basic results about uniform convergence and continuous functions in $\mathbb{R}^n$;

- understand and be able to prove basic results about differentiability of functions from $\mathbb{R}^n$ to $\mathbb{R}^m$, be able to calculate derivatives and use them to analyse the behaviour of functions;

- understand the notion of uniform continuity of functions and appreciate its significance in the theory of Riemann integration of a function of a single real variable;

- understand the basic theory of metric spaces and continuous functions on metric spaces;

- understand the statement and proof of the contraction mapping theorem and its applications to the solution of differential equations and to the inverse function theorem.

Topological Spaces

<div align="right">Lent, 16 lectures</div>

Topology is the branch of mathematics that is concerned with the properties of a geometric object that are preserved under continuous deformation. In other words, it is the study of shapes defined by the continuous maps that you learned about in Analysis I and II. The first theorems of topology were proved by Euler in the 18th century, but the subject wasn't properly formalised until the beginning of the 20th century. Since then, it has emerged to play a central role in large parts of modern mathematics and physics.

This course serves as a first introduction to topology, building on the metric spaces that you studied in Analysis II. After the definition of a topological space, the key topological ideas of connectedness and compactness are introduced and their applications explained. Fresh views emerge of important results from Analysis I.

The last part of the course moves on to discuss the topology of manifolds, providing a foundation for some of the more advanced topics covered in Part II of the tripos. Plenty of examples are discussed, and some basic properties are proved. The course finishes with a detailed discussion of 2-dimensional manifolds, i.e. surfaces. You will learn how to classify them via Euler characteristic.

Munkres' book Topology (Prentice Hall, 2000) and Lee's Topological manifolds (Springer, 2000) are good sources for this course.

Learning outcomes. By the end of this course, you should:

- appreciate the definition of topological spaces and be able to distinguish between standard topological and non-topological properties;

- understand the topological notion of connectedness and its relation to path-connectedness;

- understand the topological notion of compactness, know its significance in basic analysis and be able to apply it to identify standard quotients of topological spaces;

- understand the notion of a topological manifold, and recognise several examples;

- be familiar with the classification of surfaces, including the concepts of orientability and Euler characteristic.

Variational Principles Easter, 12 lectures

The techniques developed in this course are of fundamental importance throughout physics and applied mathematics, as well as in many areas of pure and applicable mathematics.

The first part of the course considers stationary points of functions on $\mathbb{R}^n$ and extends the treatment in Part IA Differential Equations to deal with *constraints* using the method of *Lagrange multipliers*; e.g. this allows one to determine the stationary points of a function on a surface in $\mathbb{R}^3$.

The second part of the course deals with *functionals* (and functional derivatives) and enables one to find the path that minimises the distance between two points on a given surface (a *geodesic*), the path of a light ray that gives the shortest travel time (satisfying *Fermat's Principle*), or the minimum energy shape of a soap film.

Many fundamental laws of physics (in Newtonian mechanics, relativity, electromagnetism or quantum mechanics) can be expressed as variational principles in a profoundly elegant and useful way that brings underlying symmetries to the fore.

Learning outcomes. By the end of this course, you should:

- understand the concepts of a functional, and of a functional derivative;

- be able to apply constraints to variational problems;

- appreciate the relationship between variational statements, conservation laws and symmetries in physics.

Methods Michaelmas, 24 lectures

This course continues the development of mathematical methods which can be applied to physical systems. The material is fundamental to nearly all areas of applied mathematics and theoretical physics.

The course introduces the important class of ordinary differential equations that are self-adjoint. The equivalent in the complex domain, used in Quantum Mechanics, are Hermitian operators. Self-adjoint equations have nice properties such as having real eigenvalues and orthogonal eigenfunctions, which allow eigenfunction expansions, the prototype being Fourier series. Fourier series generalise, for non-periodic functions, to Fourier transforms which provide a useful way of solving linear differential ordinary and partial differential equations.

Much of the remainder of the course concentrates on second-order partial differential equations: classification into wave, diffusion and Laplace type equations; the fundamental solutions of the three different types; solution by separation of variable which ties in with the earlier work on self-adjoint equations.

The course also introduces the famous Dirac δ, or spike, function and the Green's function which can be regarded as the inverse operator for a differential equation: it is used to express the solution in terms of an integral. It will reappear as a basic tool in quantum field theory.

It is worthwhile to get to grips early with the major new ideas introduced here: Fourier series/transforms and Sturm-Liouville equations. Reasonably friendly accounts can be found in *Mathematical Methods in the Physical Sciences* by Boas (Wiley, 1983), *Mathematical Methods for Physicists* by Arfken (Academic Press, 1985) and *Mathematical Methods for Physicists and Engineers* by Riley, Hobson and Bence (CUP, 98). It is also worthwhile to revise thoroughly the Variational Principles course from the Easter term.

Learning outcomes. By the end of this course, you should:

- be able to apply the theory of Green's functions to ordinary differential equations;

- understand the basic properties of Sturm-Liouville equations;

- be able to apply the method of separation of variables to partial differential equations;

- be able to use standard methods to solve partial differential equations.

- be able to solve wave problems using Fourier analysis and advanced/retarded coordinates.

Complex Methods   Lent, 16 lectures

Complex variable theory was introduced briefly in Analysis I (for example, complex power series). Here, the subject is developed without the full machinery of a pure analysis course. Rigorous justification of the results used is given in the parallel course, Complex Analysis.

The course starts with the definition of analyticity and the Cauchy Riemann equations (which must be satisfied by the real and imaginary parts of a complex function in order for it to be analytic; i.e. in order for it to be expressible as a power series). There follows a brief discussion of conformal mapping with applications to Laplace's equation. Then a heuristic version of Cauchy's theorem leads, via Cauchy's integral formula, to the residue calculus. This is a remarkable technique for evaluating integrals in the complex plane, which can also be used to calculate definite integrals on the real line. It allows the calculation of integrals which one would not have a hope of calculating by other means, as well as remarkably simple and elegant derivations of standard results such as

$$\int_{-\infty}^{\infty} \exp(-x^2/2 + ikx)\, dx = \sqrt{(2\pi)} \exp(-k^2/2) \qquad \text{and} \qquad \int_0^{\infty} (\sin x)/x\, dx = \pi/2\,.$$

An important application is to the theory of Fourier transforms (which were introduced in the Methods course) and Laplace transforms. The transforms are used to represent, for example, a time dependent signal as a sum (in fact, an integral unless the function is periodic) over its frequency components. This is important because one often knows how a system responds to pure frequency signals rather than to an arbitrary input. In many situations, the use of a transform simplifies a physical problem by reducing a partial differential equation to an ordinary differential equation. This is a particularly important technique for numerous branches of physics, including acoustics, optics and quantum mechanics.

For a fairly applied approach, look at chapters 6 and 7 of *Mathematical Methods for Physicists* by Arfken (Academic Press, 1985). This material is also sympathetically dealt with in *Mathematical Methods in the Physical Sciences* by Boas (Wiley, 1983).

Learning outcomes. By the end of this course, you should:

- understand the concept of analyticity;

- be able to use conformal mappings to find solutions of Laplace's equations;

- be able to use the theory of contour integration, including the residue theorem, to evaluate integrals;

- understand the theory of Fourier and Laplace transforms and apply it to the solution of ordinary and partial differential equations.

Complex Analysis   Lent, 16 lectures

This course covers about 2/3 of the material in Complex Methods, from a more rigorous point of view. The main omissions are applications of conformal mappings to solutions of Laplace's equations and the theory of Fourier and Laplace transforms.

The theory of complex variable is exceptionally elegant. It is used in many branches of pure mathematics, including number theory. It also forms one of the guiding models for the modern development of geometry.

A rigorous course not only provides a firm foundation for, and makes clear the underlying structure of, this material but also allows a deeper appreciation of the links with material in other analysis courses — in particular, Metric & Topological Spaces.

An excellent book both for the course and for preliminary reading is Hilary Priestley's *Introduction to Complex Analysis* (OUP, paperback). The books by Stewart and Tall (*Complex Analysis*) and by Jameson (*A First Course in Complex Functions*) are also good.

Learning outcomes. By the end of this course, you should:

- understand the concept of analyticity;

- be able to prove rigorously the main theorems in the course;

- be able to use the theory of contour integration, including the residue theorem, to evaluate integrals.

Quantum Mechanics   Michaelmas, 16 lectures

Quantum mechanics introduces a profound different way of thinking about the physical world, formulated using precise mathematical language. It explains phenomena beyond the reach of classical physics, such as the duality of particles and waves, and the structure and behaviour of atoms, but quantum mechanics is also at work all around us in our daily uses of modern technology.

This course introduces the subject from scratch and deals mainly with the quantum mechanics of a single particle, as described by a complex-valued *wavefunction* obeying the *Schrödinger equation*. For a quantum particle there is no definite trajectory (as determined classically from Newton's Laws) and information about position and momentum must instead be extracted from the wavefunction in terms of probabilities. One consequence of this is the *Heisenberg uncertainty principle*.

The Schrödinger equation is first studied in simple but instructive cases in one dimension, before moving on to three dimensions, culminating in the solution of the Hydrogen atom. The underlying mathematics involves hermitian or self-adjoint (differential) operators whose eigenvalues give the possible outcomes of a physical measurement. Consequently, there are significant overlaps with material in Part IB Methods (and Part IA Vectors and Matrices or equivalently Part IB Linear Algebra), although the treatment in this course is essentially self-contained.

Standard introductory textbooks are *Essential Quantum Physics* by Landshoff, Metherell and Rees (CUP, 2010) and *Quantum Mechanics* by Rae (IOP Publishing, 2002), while *The Quantum Universe* by Hey and Walters (CUP, 1987) contains readable and non-mathematical accounts with lots of pictures, going well beyond the Part IB course.

Learning outcomes. By the end of this course, you should:

- understand the basic theory of quantum mechanics, including: wavefunctions, the Schrödinger equation, observables and operators—measurements, eigenvalues, and expectation values;

- be able to solve, and interpret the solution of, the Schrödinger equation in simple cases, including: 1-dimensional potential wells and steps; the harmonic oscillator; and the hydrogen atom.

Electromagnetism   Lent, 16 lectures

Maxwell's equations of electromagnetism are among the great triumphs of nineteenth century physics. These equations unify the electric and magnetic forces and provide an explanation for many natural phenomena, including the existence of light itself. The equations also hold the seed of the theory of special relativity. This course gives the first opportunity in the Tripos to study a modern physical field theory.

After a brief discussion of electric and magnetic forces, Maxwell's equations are introduced. A key idea is the use of potentials to represent the electric and magnetic fields and it is shown how Maxwell's equations imply the existence of such potential functions. The equations are solved in special cases

of physical interest. First, time-independent situations are covered: for example, point charges, bar magnets, currents in wires. Next, time-varying situations are investigated: for example, induction. It is also shown how Maxwell's equations have wave-like solutions which we identify as light. The course ends with a discussion of special relativity in the context of electromagnetism. When viewed through the lens of relativity, the Maxwell equations become remarkably simple.

The course relies heavily on vector calculus. The latter part of the course also uses the theory of tensors from Part IA Vector Calculus and special relativity from Part IA Dynamics and Relativity. Electromagnetism is important for all of the theoretical physics courses in Part II, and is particularly relevant to General Relativity through its use of 4-vectors and tensors.

Learning outcomes. By the end of this course, you should:

- understand the physical significance of and be able to manipulate Maxwell's equations (including deriving the integral forms);

- solve simple problems in electrostatics including calculation of electrostatic energy, capacity and force;

- derive, and apply to simple situations, the Biot-Savart law;

- use Gauss's law and Ampère's law to calculate electric and magnetic fields in symmetrical situations;

- calculate forces using the Lorentz force;

- derive and apply Faraday's law of induction to simple circuits;

- solve Maxwell's equations to obtain plane waves.

Fluid Dynamics   Lent, 16 lectures

Fluid dynamics investigates the motion of liquids and gases, such as the motion that enables aircraft to fly. Newton's laws of motion apply – acceleration equals force per unit mass – but a subtlety arises because acceleration means the rate of change of velocity of a fluid particle. It does not mean the rate of change of the fluid velocity at a fixed point in space. A special mathematical operator, the material derivative, expresses the required rate of change using vector calculus. The forces entering Newton's laws can be external, such as gravity, or internal, arising from pressure or from viscosity (internal friction). When the viscosity is small enough to be negligible, the motion is often irrotational as well as incompressible: both the curl and divergence of velocity field vanish. In this situation, the fluid velocity can be described by a potential, and standard potential theory applies, including in some cases solutions of Laplace's equation.

The topics studied include jets, bubbles, waves, vortices, flow around aircraft wings, and flow in weather systems. Suitable introductory reading material can be found in Worster's *Understanding Fluid Flow* (CUP) or in Acheson's *Elementary Fluid Dynamics* (Oxford). For background motivation, see also the visionary discussion in the Feynman Lectures on Physics, last two chapters of Volume II (Addison-Wesley).

Learning outcomes. By the end of this course, you should:

- understand the basic principles governing the dynamics of parallel viscous flows and flows in which viscosity in negligible;

- be able to derive and deduce the consequences of the equation of conservation of mass;

- be able solve kinematic problems such as finding particle paths and streamlines;

- be able to apply Bernoulli's theorem and the momentum integral to simple problems including river flows;

- understand the concept of vorticity and the conditions in which it may be assumed to be zero;

- calculate velocity fields and forces on bodies for simple steady and unsteady flows derived from potentials;

- understand the theory of interfacial waves and be able to use it to investigate, for example, standing waves in a container;

- understand fundamental ideas relating to flows in rotating frames of reference, particularly geostrophy.

Numerical Analysis Lent, 16 lectures

An important aspect of the application of mathematics to problems in the real world is the ability to compute answers as accurately as possible subject to the errors inherent in the data presented and the limits on the accuracy of calculation. Numerical analysis is the branch of mathematics studying such computations.

The course commences from approximation theory, focusing on the approximation of functions and data by polynomials, continues with the numerical solution of ordinary differential equations and concludes with the solution of linear algebraic systems. Although computational algorithms form a central part of the course, so do mathematical theories underlying them and investigating their behaviour: computation and approximation at their best should be done with proper mathematical justification.

An Introduction to Numerical Analysis by Suli & Mayers (CUP, 2003) and *Interpolation and Approximation* by Davis (Dover, 1975) are two excellent introductory texts.

Learning outcomes. By the end of this course, you should:

- understand the role of algorithms in numerical analysis;

- understand the role and basic theory (including orthogonal polynomials and the Peano kernel theorem) of polynomial approximation;

- understand multistep and Runge–Kutta methods for ordinary differential equations and the concepts of convergence, order and stability;

- understand the theory of algorithms such as LU and QR factorisation, and be able to apply them, for example to least squares calculations.

Statistics Lent, 16 lectures

Statistics is the study of what can be learnt from data. We regard our data as realisations of random variables, and consider models for the (joint) distribution of these random variables. In this course, we focus entirely on *parametric* models, where the class of distributions considered can be indexed by a finite-dimensional parameter. As a simple example, the family of normal distributions can be indexed by a two-dimensional parameter, representing the mean and variance. Nonparametric models are treated in more advanced courses.

Our aim is to make inference about the unknown parameter by, for example, providing a point estimate, a confidence interval or conducting a hypothesis test. Building on Part IA Probability, this course will present basic techniques of inference, together with their theoretical justification. The final chapter will cover the ubiquitous *linear model*, with its elegant theory of orthogonal projection and application of results from linear algebra.

The most appropriate book for the course is *Statistical inference* by Casella and Berger (Duxbury, 2001).

Learning outcomes. By the end of this course, you should:

- understand the basic concepts involved in point estimation, the construction of confidence intervals and Bayesian inference;

- understand and be able to apply the ideas of hypothesis testing, including the Neyman–Pearson lemma, and generalised likelihood ratio tests, including applications to goodness of fit tests and contingency tables.

- understand and be able to apply the theory of the linear model, including examples of linear regression and one-way analysis of variance.

Markov Chains   Michaelmas, 12 lectures

A Markov process is a random process for which the future (the next step) depends only on the present state; it has no memory of how the present state was reached. A typical example is a random walk (in two dimensions, the drunkard's walk).

The course is concerned with Markov chains in discrete time, including periodicity and recurrence. For example, a random walk on a lattice of integers returns to the initial position with probability one in one or two dimensions, but in three or more dimensions the probability of recurrence in zero. Some Markov chains settle down to an equilibrium state and these are the next topic in the course.

The material in this course will be essential if you plan to take any of the applicable courses in Part II. Further introductory material and notes on the course are available from links on the Study pages on the DPMMS website.

Learning outcomes. By the end of this course, you should:

- understand the notion of a discrete-time Markov chain and be familiar with both the finite state-space case and some simple infinite state-space cases, such as random walks and birth-and-death chains;

- know how to compute for simple examples the n-step transition probabilities, hitting probabilities, expected hitting times and invariant distribution;

- understand the notions of recurrence and transience, and the stronger notion of positive recurrence;

- understand the notion of time-reversibility and the role of the detailed balance equations;

- know under what conditions a Markov chain will converge to equilibrium in long time;

- be able to calculate the long-run proportion of time spent in a given state.

Optimisation   Easter, 12 lectures

A typical problem in optimisation is to find the cheapest way of supplying a set of supermarkets from a set of warehouses: in more general terms, the problem is to find the minimum (or maximum) value of a quantity when the variables are subject to certain constraints. Many real-world problems are of this type and the theory discussed in the course are practically extremely important as well as being interesting applications of ideas introduced earlier in Numbers and Sets and in Vectors and Matrices.

The theory of Lagrange multipliers, linear programming and network analysis is developed. Topics covered include the simplex algorithm, the theory of two-person games and some algorithms particularly well suited to solving the problem of minimising the cost of flow through a network.

Whittle's *Optimisation under Constraints* (Wiley, 1971) gives a good idea of the scope and range of the subject but is a little advanced mathematically; Luenberger's *Introduction to Linear and and Non-linear Programming* (Addison-Wesley, 1973) is at the right level but provides less motivation.

Learning outcomes. By the end of this course, you should:

- understand the nature and importance of convex optimisation;

- be able to apply Lagrangian methods to solve problems involving constraints;

- be able to solve problems in linear programming by methods including the simplex algorithm and duality;

- be able to solve network problems by methods using, for example, the Ford–Fulkerson algorithm and min-cut max-flow theorems.

Computational Projects (CATAM)

This course consists mainly of practical computational projects carried out and written up for submission a week after the beginning of the Lent and Easter terms. For full credit, you do four projects. The first two are prescribed and are submitted soon after the beginning of the Lent term. The remaining two are chosen from a list of projects and are submitted soon after the beginning of the Easter term. There is also a non-examinable project that allows you to practice programming and writing up results, with a model answer provided in Michaelmas Term for comparison with your own answer.

The emphasis in the projects is on understanding the mathematical problems being modelled rather than on the details of computer programming. Some students find the projects somewhat time consuming, especially those who are not used to programming or have not completed the non-examinable project. The CATAM manual will be available over the summer and it would be extremely helpful for you to start as early as possible on the non-examinable project and, if time, the first two examinable projects.

The amount of credit available for the Computational Projects course in Part IB is 160 marks (and no quality marks), which is additional to the marks gained on examination papers. In recent years approximately 99% of Part IB students submitted projects (not necessarily complete).

Learning outcomes. By the end of this course, you should:

- be able to programme using a traditional programming language;

- understand the limitations of computers in relation to solving mathematical problems;

- be able to use a computer to solve problems in both pure and applied mathematics involving, for example, solution of ordinary differential equations and manipulation of matrices.