

Mathematics and Statistics Undergraduate Handbook Supplement to the Handbook

Honour School of Mathematics and Statistics Syllabus and Synopses for Part C 2024–2025 for examination in 2025

Contents

1	Honour School of Mathematics and Statistics	2
1.1	Units	2
1.2	Part C courses in future years	2
1.3	Course list by term	3
2	Statistics units	4
2.1	SC1 Stochastic Models in Mathematical Genetics – 16 MT	4
2.2	SC2 Probability and Statistics for Network Analysis – 16 MT	5
2.3	SC4 Advanced Topics in Statistical Machine Learning – 16 HT	6
2.4	SC5 Advanced Simulation Methods - 16 HT	7
2.5	SC6 Graphical Models – 16MT	8
2.6	SC7 Bayes Methods – 16 HT	9
2.7	SC8 Topics in Computational Biology – 16 HT	10
2.8	SC9 Probability on Graphs and Lattices – 16 MT	12
2.9	SC10 Algorithmic Foundations of Learning – 16 MT	13
3.0	SC11 Climate Statistics – 16 HT	14
3.1	C8.4: Probabilistic Combinatorics - 16 HT	16
4	Mathematics units	17

Every effort is made to ensure that the list of courses offered is accurate at the time of going online. However, students are advised to check the up-to-date version of this document on the Department of Statistics website.

Notice of misprints or errors of any kind, and suggestions for improvements in this booklet should be addressed to the Academic Administrator in the Department of Statistics, academic.administrator@stats.ox.ac.uk.

July 2024

1 Honour School of Mathematics and Statistics

1.1 Units

See the current edition of the Examination Regulations at <https://examregs.admin.ox.ac.uk> for the full regulations governing these examinations. The examination conventions can be found on the Canvas course site.

In Part C,

- each candidate shall offer a minimum of six units and a maximum of eight units from the schedule of units for Part C
- and each candidate shall also offer a **dissertation** on a statistics project (equivalent of 2 units).
- At least 3 of the units taken by Part C students must be assessed by written examination.

At least **two** units should be from the schedule of 'Statistics' units. The USMs for the dissertation and the best six units will count for the final classification.

Units from the schedule of 'Mathematics Department units' for Part C of the Honour School of Mathematics are also available – see Section 3.

This booklet describes the units available in Part C. Information about dissertations/statistics projects will be available on the Department of Statistics Canvas site.

All of the units described in this booklet are "M-level".

Students are asked to register for the options they intend to take by the end of week 10, Trinity Term 2024 using the Mathematical Institute course management portal. <https://courses.maths.ox.ac.uk/course/index.php?categoryid=735>. Students may alter the options they have registered for after this but it is helpful if their registration is as accurate as possible. Students will then be asked to sign up for classes at the start of Michaelmas Term 2024. Students who register for a course or courses for which there is a quota should consider registering for an additional course (by way of a "reserve choice") in case they do not receive a place on the course with the quota.

Every effort will be made when timetabling lectures to ensure that mathematics lectures do not clash. However, because of the large number of options this may sometimes be unavoidable.

1.2 Part C courses in future years

In any year, most courses available in Part C that year will normally also be available in Part C the following year. However, sometimes new options will be added or existing options may cease to run. The list of courses that will be available in Part C in any year will be published by the end of the preceding Trinity Term.

1.3 Course list by term

The 2024-2025 list of Part C courses by term is:

Michaelmas Term

SC1 Stochastic Models in Mathematical Genetics
SC2 Probability and Statistics for Network Analysis
SC6 Graphical Models
SC9 Probability on Graphs and Lattices
SC10 Algorithmic Foundations of Learning

Hilary Term

SC4 Advanced Topics in Statistical Machine Learning
SC5 Advanced Simulation Methods
SC7 Bayes Methods
SC8 Topics in Computational Biology
SC11 Climate Statistics
C8.4 Probabilistic Combinatorics.

2. Statistics Units

2.1 SC1 Stochastic Models in Mathematical Genetics – 16 MT

Level: M-level

Method of Assessment: written examination

Weight: Unit

Recommended Prerequisites

Part A A8 Probability.

SB3.1 Applied Probability would be helpful.

Aims & Objectives

The aim of the lectures is to introduce modern stochastic models in mathematical population genetics and give examples of real world applications of these models. Stochastic and graph theoretic properties of coalescent and genealogical trees are studied in the first eight lectures. Diffusion processes and extensions to model additional key biological phenomena are studied in the second eight lectures.

Synopsis

Evolutionary models in Mathematical Genetics:

The Wright-Fisher model. The Genealogical Markov chain describing the number ancestors back in time of a collection of DNA sequences.

The Coalescent process describing the stochastic behaviour of the ancestral tree of a collection of DNA sequences. Mutations on ancestral lineages in a coalescent tree. Models with a variable population size.

The frequency spectrum and age of a mutation. Ewens' sampling formula for the probability distribution of the allele configuration of DNA sequences in a sample in the infinitely-many-alleles model. Hoppe's urn model for the infinitely-many-alleles model.

The infinitely-many-sites model of mutations on DNA sequences. Gene trees as perfect phylogenies describing the mutation history of a sample of DNA sequences. Graph theoretic constructions and characterizations of gene trees from DNA sequence variation. Gusfield's construction algorithm of a tree from DNA sequences. Examples of gene trees from data.

Modelling biological forces in Population Genetics: Recombination. The effect of recombination on genealogies. Detecting recombination events under the infinitely-many-sites model. Hudson's algorithm. Haplotype bounds on recombination events. Modelling recombination in the Wright-Fisher model. The coalescent process with recombination: the ancestral recombination graph. Properties of the ancestral recombination graph.

Introduction to diffusion theory. Tracking mutations forward in time in the Wright-Fisher model. Modelling the frequency of a neutral mutation in the population via a diffusion process limit. The generator of a diffusion process with two allelic types. The probability of fixation of a mutation. Genic selection. Extension of results from neutral to selection case. Behaviour of selected mutations.

Reading

R. Durrett, *Probability Models for DNA Sequence Evolution*, Springer, 2008
A. Etheridge, Some Mathematical Models from Population Genetics. Ecole d'Eté de Probabilités de Saint-Flour XXXIX-2009, Lecture Notes in Mathematics, 2012
W. J. Ewens, *Mathematical Population Genetics*, 2nd Ed, Springer, 2004
J. R. Norris, *Markov Chains*, Cambridge University Press, 1999
M. Slatkin and M. Veuille, *Modern Developments in Theoretical Population Genetics*, Oxford Biology, 2002
S. Tavaré and O. Zeitouni, *Lectures on Probability Theory and Statistics*, Ecole d'Eté de Probabilités de Saint-Flour XXXI - 2001, Lecture Notes in Mathematics 1837, Springer, 2004

2.2 SC2 Probability and Statistics for Network Analysis – 16 MT

Level: M-level

Method of Assessment: Written examination

Weight: Unit

For this course, 2 lectures and 2 intercollegiate classes are replaced by 2 practical classes. (The total time for this course is the same as for other Part C courses.)

Recommended prerequisites

Part A A8 Probability and A9 Statistics

Aims and Objectives

Many data come in the form of networks, for example friendship data and protein-protein interaction data. As the data usually cannot be modelled using simple independence assumptions, their statistical analysis provides many challenges. The course will give an introduction to the main problems and the main statistical techniques used in this field. The techniques are applicable to a wide range of complex problems. The statistical analysis benefits from insights which stem from probabilistic modelling, and the course will combine both aspects.

Synopsis

Exploratory analysis of networks. The need for network summaries. Degree distribution, clustering coefficient, shortest path length. Motifs.

Probabilistic models: Bernoulli random graphs, geometric random graphs, preferential attachment models, small world networks, inhomogeneous random graphs, exponential random graphs.

Small subgraphs: Stein's method for normal and Poisson approximation. Branching process approximations, threshold behaviour, shortest path between two vertices.

Statistical analysis of networks: Sampling from networks. Parameter estimation for models. Inferring edges in networks. Network comparison. A brief look at community detection.

Reading

R. Durrett, *Random Graph Dynamics*, Cambridge University Press, 2007

E.D Kolaczyk and G. Csádi, *Statistical Analysis of Network Data with R*, Springer, 2014
M. Newman, *Networks*. Oxford University Press

2.3 **SC4 Advanced Topics in Statistical Machine Learning** – 16 HT

Level: M-level

Methods of Assessment: written examination.

Weight: Unit

Recommended prerequisites

The course requires a good level of mathematical maturity. Students are expected to be familiar with core concepts in statistics (regression models, bias-variance tradeoff, Bayesian inference), probability (multivariate distributions, conditioning) and linear algebra (matrix-vector operations, eigenvalues and eigenvectors). Previous exposure to machine learning (empirical risk minimisation, dimensionality reduction, overfitting, regularisation) is highly recommended.

Students would also benefit from being familiar with the material covered in the following courses offered in the Statistics department: SB2.1 (formerly SB2a) Foundations of Statistical Inference and in SB2.2 (formerly SB2b) Statistical Machine Learning.

Aims and Objectives

Machine learning is widely used across the sciences, engineering and society, to construct methods for identifying interesting patterns and predicting accurately from large datasets.

This course introduces several widely used machine learning techniques and describes their underpinning statistical principles and properties. The course studies both unsupervised and supervised learning and several advanced and state-of-the-art topics are covered in detail. The course will also cover computational considerations of machine learning algorithms and how they can scale to large datasets.

Synopsis

Empirical risk minimisation. Loss functions. Generalization. Over- and underfitting. Regularisation.

Support vector machines.

Kernel methods and reproducing kernel Hilbert spaces. Representer theorem. Representation of probabilities in RKHS.

Probabilistic and Bayesian machine learning: Fundamentals of the Bayesian approach.

Gaussian processes. Bayesian optimisation.

Deep learning: Neural networks. Computation graphs. Automatic differentiation. Stochastic gradient descent.

Variational inference. Latent variable models.

Deep generative models. Variational auto-encoders.

Software

Knowledge of Python is not required for this course, but some examples may be done in Python. Students interested in learning Python are referred to the following free University IT online course, which should ideally be taken before the beginning of this course: <https://skills.it.ox.ac.uk/whats-on#/course/LY046>

Reading

C. Bishop, Pattern Recognition and Machine Learning, Springer, 2007

K. Murphy, Machine Learning: a Probabilistic Perspective, MIT Press, 2012

Further Reading

T. Hastie, R. Tibshirani, J. Friedman, Elements of Statistical Learning, Springer, 2009

Scikit-learn: Machine Learning in Python, Pedregosa et al., JMLR 12, pp. 2825-2830, 2011, <http://scikit-learn.org/stable/tutorial/>

2.4 SC5 Advanced Simulation Methods - 16 HT

Level: M-level

Methods of Assessment: This course is assessed by written examination.

Weight: Unit

Recommended Prerequisites

The course requires a good level of mathematical maturity as well as some statistical intuition and background knowledge to motivate the course. Students are expected to be familiar with core concepts from probability (conditional probability, conditional densities, properties of conditional expectations, basic inequalities such as Markov's, Chebyshev's and Cauchy-Schwarz's, modes of convergence), basic limit theorems from probability in particular the strong law of large numbers and the central limit theorem, Markov chains, aperiodicity, irreducibility, stationary distributions, reversibility and convergence. Most of these concepts are covered in courses offered in the Statistics department, in particular prelims probability, A8 probability and SB3.1 (formerly SB3a) Applied Probability.

Familiarity with basic Monte Carlo methods will be helpful, as for example covered in A12 Simulation and Statistical Programming.

Some familiarity with concepts from Bayesian inference such as posterior distributions will be useful in order to understand the motivation behind the material of the course.

Aims and Objectives

The aim of the lectures is to introduce modern simulation methods.

This course concentrates on Markov chain Monte Carlo (MCMC) methods and Sequential Monte Carlo (SMC) methods. Examples of applications of these methods to complex inference problems will be given.

Synopsis

Classical methods: inversion, rejection, composition.

Importance sampling.

MCMC methods: elements of discrete-time general state-space Markov chains theory, Metropolis-Hastings algorithm.

Advanced MCMC methods: Gibbs sampling, slice sampling, tempering/annealing, Hamiltonian (or Hybrid) Monte Carlo, pseudo-marginal MCMC.

Sequential importance sampling.

SMC methods: nonlinear filtering.

Reading

C.P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd edition, Springer-Verlag, 2004

Further reading

J.S. Liu, *Monte Carlo Strategies in Scientific Computing*, Springer-Verlag, 2001

2.5 **SC6 Graphical Models** – 16 MT

Level: M-level

Methods of Assessment: This course is assessed by written examination.

Weight: Unit

Recommended Prerequisites

The basics of Markov chains (in particular, conditional independence) from Part A Probability is assumed. Likelihood theory, contingency tables, and likelihood-ratio tests are also important; this is covered in Part A Statistics. Knowledge of exponential families and linear models (as covered in Part B Foundations of Statistical Inference and Applied Statistics) would be useful, but is not essential.

Aims and Objectives

This course will give an overview of the use of graphical models as a tool for statistical inference. Graphical models relate the structure of a graph to the structure of a multivariate probability distribution, usually via a factorization of the distribution or conditional independence constraints. This has two broad uses: first, conditional independence can provide vast savings in computational effort, both in terms of the representation of large multivariate models and in performing inference with them; this makes graphical models very popular for dealing with big data problems. Second, conditional independence can be used as a tool to discover hidden structure in data, such as that relating to the direction of causality or to unobserved processes. As such, graphical models are widely used as causal models in genetics, medicine, epidemiology, statistical physics, economics, the social sciences and elsewhere.

Students will develop an understanding of the use of conditional independence and graphical structures for dealing with multivariate statistical models. They will appreciate how this is applied to causal modelling, and to computation in large-scale statistical problems.

Synopsis

- Independence, conditional independence, graphoid axioms.
- Exponential families, mean and canonical parameterizations, moment matching; contingency tables, log-linear models.

- Undirected graphs, cliques, paths; factorization and Markov properties, Hammersley-Clifford Theorem (statement only).
- Trees, cycles, chords, decomposability, triangulation, running intersection property. Maximum likelihood in decomposable models, iterative proportional fitting.
- The multivariate Gaussian distribution and Gaussian graphical models.
- Directed acyclic graphs, factorization. Paths, d-separation, moralization. Ancestral sets and sub-models. Decomposable models as intersection of directed and undirected models.
- Running intersection property, Junction trees; message passing, computation of marginal and conditional probabilities, introduction of evidence.
- Causal models, linear structural equations, interventions, the trek rule.
- Average causal effects, adjustment, valid adjustment sets, forbidden projection, and optimal adjustment.

Reading

1. S.L. Lauritzen, Graphical Models, Oxford University Press, 1996.
2. D. Koller and N. Friedman, Probabilistic Graphical Models: Principles and Techniques, MIT Press, 2009.
3. J. Pearl, Causality, third edition, Cambridge, 2013.
4. M.J. Wainwright and M.I. Jordan, Graphical Models, Exponential Families, and Variational Inference, Foundations and Trends in Machine Learning, 2008.
(available for free at https://people.eecs.berkeley.edu/~wainwrig/Papers/WaiJor08_FTML.pdf)
5. A. Agresti. Categorical Data Analysis, 3rd Edition, John Wiley & Sons, 2013.

2.6 SC7 Bayes Methods – 16 HT

Level: M-level

Method of Assessment: Written examination

Weight: Unit

Recommended prerequisites

SB2.1 (formerly SB2a) Foundations of Statistical Inference is desirable, of which 6 lectures on Bayesian inference, decision theory and hypothesis testing with loss functions are assumed knowledge. A12 Simulation and Statistical Programming desirable.

Synopsis

Theory: Decision-theoretic foundations, Savage axioms. Prior elicitation, exchangeability. Bayesian Non-Parametric (BNP) methods, the Dirichlet process and the Chinese Restaurant Process. Asymptotics, and information criteria.

Computational methods: Bayesian inference via MCMC; Estimation of marginal likelihood; Approximate Bayesian Computation and intractable likelihoods; reversible jump MCMC.

Case Studies: extend understanding of prior elicitation, BNP methods and asymptotics through a small number of substantial examples. Examples to further illustrate building

statistical models, model choice, model averaging and model assessment, and the use of Monte Carlo methods for inference.

Reading

C.P. Robert, *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*, 2nd edition, Springer, 2001

Further Reading

A. Gelman et al, *Bayesian Data Analysis*, 3rd edition, Boca Raton Florida: CRC Press, 2014

P Hoff, *A First Course in Bayesian Statistical Methods*, Springer, 2010

DeGroot, Morris H., *Optimal Statistical Decisions*. Wiley Classics Library. 2004.

2.7 SC8 Topics in Computational Biology – 16 HT

Level: M-level

Method of Assessment: Mini-project

Weight: Unit

Recommended Prerequisites:

The lectures attempt to be self-contained but clearly knowledge algorithms, combinatorics and probability theory (A8 Probability and SB3.1 Applied Probability) would be a help. The course requires a good level of mathematical maturity.

Aims & Objectives

Modern molecular biology generates large amounts of data, such as sequences, structures and expression data, that needs different forms of statistical analysis and modelling to be properly interpreted. This course focuses on four topics within this vast area: Molecular Dynamics, Molecule Enumeration, Comparative Biology and Overview of Computational Biology and Computational Neurosciences.

Synopsis:

Modelling the Biosphere - The biosphere is a well-defined physical system that can be described at different levels of coarse-grained. At the most detailed level, it is an extremely complicated system with an energy in- and outflux, a physical system consisting of atmosphere, ocean and crust and in excess of 10 million species. But even at the very coarsest level major insights can be obtained and the greenhouse effect was pretty precisely described 130 years ago based on the simplest of models. Climate Change and Weather Forecasting make models of immense economic importance and are experiencing major development, both in terms of data collection, parameterisation and computational power. Like in epidemics, forecasting and assessment of forecasting is important. Statistical testing and causal analysis of biospheric changes are also receiving major attention. Describing scenarios of Climate Change is of interest to all and the foundation of political and economic decision making.

Advanced Sequence Analysis - Sequences are the classical golden type of data within Bioinformatics/Computational Biology with a 50+ year history. However, there are plenty of open problems that have arisen either due to the explosion of data, incorporation of other information or the rise of Machine Learning. In this sub-topic we go through sequence models including insertion-deletions, substitutions and how to infer parameters, relationships and alignments.

Computational Models of Pattern and Shape - Classically, Multicellular Organisms were identified by their Patterns and Shapes. This has over the last 1-2 centuries been vastly enhanced by molecular/cellular characterisations. This has been a triumph of modern science. The dynamics and evolution of molecular/cellular components and processes has well-developed models associated, but as you proceed to higher levels within an individual, such as tissue and organs and eventually the complete organism, very major challenges remain. Within this sub-topic we will attempt to survey the main Patterns/Shapes observed in Multicellular Organisms, the models that have been used to describe and explain them. Finally, we will discuss attempts to go from Genome to Patterns/Shapes. This subtopic will be one of the main challenges for the biosciences in the coming decades.

Computational Neuroscience - Computational Neuroscience is in massive growth, but has a history going back to the 1940s with publications such as McCulloch and Pitts (1943) paper on neural networks. The present progress is driven by progress on three fronts: (i) experimental data on brains, nerve systems and individual neurons, (ii) increased success in designing artificial neural networks with an increasing variety in architectures with applications in Deep Learning/AI and (iii) the ability to simulate very complex models as models of biological neural networks. Neuroscience and Computational Neuroscience are arguably the most complex areas of present science it could take many decades for the most fundamental questions to be answered. The origin(s) of Neurosystems (and general likelihood/unlikelihood) is still an open question. The definition of consciousness is not only a philosophical question but of practical importance.

The exact plan for lectures is:

W1 The Basics of the Biosphere
W1 Modelling the Dynamics of the Biosphere
W2 Model Properties and Statistics
W2 Climate Change

W3 Models of Sequence Evolution
W3 Alignment
W4 Phylogenetics
W4 Pattern Search in Sequences

W5 Biological & Artificial Neurons
W5 Modelling Small Neural Networks
W6 Large Scale Brain Modelling
W6 Evolution of of the Brain

W7 The Classical Models of Patterns and Shapes
W7 Patterns and Shapes in Animals
W8 Patterns and Shapes in Plants
W8 Functions from Genome to Pattern and Shape

2.8 SC9 Probability on Graphs and Lattices – 16 MT

Level: M-level

Method of Assessment: Written examination Weight: Unit

Recommended Prerequisites

Discrete and continuous time Markov processes on countable state space, as covered for example in Part A A8 Probability and Part B SB3.1 Applied Probability.

Aims and Objectives

The aim is to introduce fundamental probabilistic and combinatorial tools, as well as key models, in the theory of discrete disordered systems. We will examine the large-scale behaviour of systems containing many interacting components, subject to some random noise. Models of this type have a wealth of applications in statistical physics, biology and beyond, and we will see several key examples in the course. Many of the tools we will discuss are also of independent theoretical interest, and have far reaching applications. For example, we will study the amount of time it takes for a random system to reach its stationary distribution (mixing time). This concept is also important in many statistical applications, such as studying the run time of MCMC methods.

Synopsis

- Uniform spanning trees, loop-erased random walks, Wilson's algorithm, the Aldous-Broder algorithm.
- Percolation, phase transitions in \mathbb{Z}^d , specific tools in \mathbb{Z}^2 .
- Ising model, random-cluster model and other models from statistical mechanics (e.g. Potts model, hard-core model).
- Glauber dynamics, mixing times, couplings.

Reading

G. Grimmett, *Probability on graphs: random processes on graphs and lattices*, Cambridge University Press, 2010; 2017 (2nd edition).

B. Bollobás, O. Riordan, *Percolation*, Cambridge University Press, 2006.

T. Liggett, *Continuous time Markov processes: an introduction*, American Mathematical Society, 2010.

D. A. Levin, Y. Peres, E. L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, 2009.

H. Duminil-Copin, *Introduction to Bernoulli percolation*. Lecture notes available online at <https://www.ihes.fr/~duminil/publi/2017percolation.pdf>.

2.9 SC10 **Algorithmic Foundations of Learning** – 16 MT

Level: M-level

Method of Assessment: Written examination

Weight: Unit

Recommended Prerequisites

The course requires a good level of mathematical maturity. Students are expected to have working knowledge of core concepts in probability theory (basic properties of probabilities, such as union bounds, and of conditional expectations, such as the tower property or law of total expectations; basic inequalities, such as Markov's and Jensen's), statistics (confidence intervals, hypothesis testing), and linear algebra (matrix-vector operations, eigenvalues and eigenvectors; basic inequalities, such as Cauchy-Schwarz's and Hölder's). Previous exposure to machine learning (empirical risk minimisation, overfitting, regularisation) is recommended.

Students would benefit from being familiar with the material covered in SB2.1 (formerly SB2a) Foundations of Statistical Inference (in particular, Decision Theory) and in SB2.2 (formerly SB2b) Statistical Machine Learning.

Aims and objectives

The course is meant to provide a rigorous theoretical account of the main ideas underlying machine learning, and to offer a principled framework to understand the algorithmic paradigms being used, along with non-asymptotic methods for the study of random structures in high-dimensional probability, statistics, and optimisation.

Synopsis

- Statistical learning frameworks for prediction, estimation and online learning.
- Probability:
 - Maximal inequalities.
 - Rademacher and Gaussian complexities.
 - Elements of VC theory.
 - Covering and packing numbers.
 - Chaining.
 - Concentration inequalities.
- Statistics:
 - Bayes decision rules.
 - Empirical risk minimisation. Error decomposition: generalisation, optimisation, and approximation.
 - Learning via uniform convergence, margin bounds, and algorithmic stability.
 - Regularisation: explicit (constraints and penalisation) and implicit (algorithmic).
 - Convex loss surrogates.
 - Slow and fast rates.
 - Minimax lower bounds and hypothesis testing.
- Optimisation:
 - Elements of convex theory.
 - Oracle model. Gradient descent. Mirror descent.
 - Stochastic oracle model. Stochastic gradient descent. Stochastic mirror descent. Variance reduction techniques.

- Implicit Bias.
 - Online optimisation.
- Examples
 - Linear predictors, including Boosting.
 - Non-linear predictors, including Support Vector Machines and Neural Networks.
 - High-dimensional estimators for sparse and low-rank problems, including Lasso.
 - Online learning, including multi-armed bandit problems, reinforcement learning and algorithms.

Reading

- Shai Shalev-Shwartz and Shai Ben-David. Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press. 2014
- Sébastien Bubeck. Convex Optimization: Algorithms and Complexity. Foundations and Trends in Machine Learning. 2015
- Ramon van Handel. Probability in High Dimension. Lecture notes available online at <http://www.princeton.edu/~rvan/APC550.pdf>. 2016
- Tor Lattimore and Csaba Szepesvári. Bandit Algorithms. Book available online at <https://tor-lattimore.com/downloads/book/book.pdf>. 2019.
- Martin J. Wainwright. High-Dimensional Statistics. A Non-Asymptotic Viewpoint. Cambridge University Press. 2019.

3.0 SC11 **Climate Statistics** – 16 HT

Level: M-level

Method of Assessment: Written examination

Weight: Unit

Recommended Prerequisites:

Students will need a grounding in probability and statistical theory at least on the level of A8 Probability and A9 Statistics, and will need most of the material of SB1.1 Applied Statistics. There will be heavy use of Linear Algebra, but not assuming proficiency beyond the level of Oxford's first-year courses. It will be helpful for students to have some familiarity with Fourier series, also on the level of the first-year course Fourier series and PDEs), and simulation methods for statistical inference (on the level of SB1.2 Computational Statistics).

Aims and Objectives

This course aims to teach the fundamentals of some statistical concepts and techniques that are relevant for understanding and carrying out research in climate science. It will teach about the varieties of climate data and show how they may be analysed with these techniques, as well as teaching some of the core concepts of climate science, and show students how improvements in climate science have been linked to improvements in statistical methodology.

The topics covered are all core statistical methods that are not at all or hardly covered in our current curriculum. The topics are: time series, multivariate analysis, multivariate

decomposition methods (PCA, CCA, related issues), and extreme values, plus a lecture on the design and interpretation of large simulation models.

Techniques of data analysis in R will be taught, and students will be expected to engage with issues of data analysis. The problem sheets will include computing questions, which may in principle be done in any programming language, though solutions will be provided only in R. They will not be examined on writing code, but on the interpretation of computational outputs.

Synopsis

Introduction: History and background of climate science; the varieties of climate data and climate models.

Time series:

- Exploratory Data analysis and nonparametric smoothing;
- ARIMA models in the time domain;
- Spectral methods;
- Introduction to wavelets. Basic concepts of time-frequency representations, application of Morlet wavelets.

Multivariate analysis:

- MANOVA and MANCOVA;
- Testing with Wilks Λ and Hotelling T^2
- Multivariate regression.

Multivariate decomposition methods:

- Principal Components Analysis;
- Canonical Correlations Analysis;
- Model selection with Mutual Information Criterion and cross validation;
- Predictable Components Analysis.

Extreme values: Basic theory of extreme value distributions, convergence theorems (without detailed proofs), inference for the generalised extreme value distribution.

Computational climate models: Problems of inference from large spatio-temporal simulation models.

Reading

Statistical Methods for Climate Scientists, Timothy M. Delsole and Michael K. Tippett.
Time Series Analysis and its Applications, Robert H. Shumway and David S. Stoffer.

Further Reading:

Climate Change 2021: The Physical Science Basis, The Intergovernmental Panel on Climate Change.

Multivariate Time Series Analysis in Climate and Environmental Research, Zhihua Zhang.

Extreme Value Theory with Applications to Natural Hazards, ed. Nicolas Bousquet and Pietro Bernardera.

Statistical Analysis of Climate Series: Analyzing, plotting, modelling, and predicting with R, Helmut Pruscha.

The Geometry of Multivariate Statistics, Thomas D. Wickens.

3.1 C8.4 Probabilistic Combinatorics - 16 HT

Level: M-level

Method of Assessment: Written examination.

Weight: Unit

Recommended Prerequisites:

B8.5 Graph Theory and A8: Probability. C8.3 Combinatorics is not as essential prerequisite for this course, though it is a natural companion for it.

Overview

Probabilistic combinatorics is a very active field of mathematics, with connections to other areas such as computer science and statistical physics. Probabilistic methods are essential for the study of random discrete structures and for the analysis of algorithms, but they can also provide a powerful and beautiful approach for answering deterministic questions. The aim of this course is to introduce some fundamental probabilistic tools and present a few applications.

Learning Outcomes

The student will have developed an appreciation of probabilistic methods in discrete mathematics.

Synopsis

First-moment method, with applications to Ramsey numbers, and to graphs of high girth and high chromatic number.

Second-moment method, threshold functions for random graphs.

Lovász Local Lemma, with applications to two-colourings of hypergraphs, and to Ramsey numbers.

Chernoff bounds, concentration of measure, Janson's inequality.

Branching processes and the phase transition in random graphs.

Clique and chromatic numbers of random graphs.

Reading

N. Alon and J.H. Spencer, *The Probabilistic Method*, 3rd edition, Wiley, 2008

Further Reading:

B. Bollobás, *Random Graphs*, 2nd edition, Cambridge University Press, 2001

M. Habib, C. McDiarmid, J. Ramirez-Alfonsin, B. Reed, ed., *Probabilistic Methods for Algorithmic Discrete Mathematics*, Springer, 1998

S. Janson, T. Luczak and A. Rucinski, *Random Graphs*, John Wiley and Sons, 2000

M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*, Cambridge University Press, New York (NY), 2005

M. Molloy and B. Reed, *Graph Colouring and the Probabilistic Method*, Springer, 2002

R. Motwani and P. Raghavan, *Randomized Algorithms*, Cambridge University Press, 1995

4 Mathematics units

The Mathematics units that students may take are drawn from Part C of the Honour School of Mathematics. For full details of these units, see the Syllabus and Synopses for Part C of the Honour School of Mathematics, which are available on the web at <https://courses.maths.ox.ac.uk/course/index.php?categoryid=735>

The Mathematics units that are available are as follows:

C1.1	Model Theory	16 MT
C1.2	Godel's Incompleteness Theorems	16 HT
C1.3	Analytic Topology	16 MT
C1.4	Axiomatic Set Theory	16 HT
C2.2	Homological Algebra	16 MT
C2.3	Representation Theory of Semisimple Lie Algebras	16 HT
C2.4	Infinite Groups	16 MT
C2.5	Non-Commutative Rings	16 HT
C2.6	Introduction to Schemes	16 HT
C2.7	Category Theory	16 MT
C3.1	Algebraic Topology	16 MT
C3.2	Geometric Group Theory	16 HT
C3.3	Differentiable Manifolds	16 MT
C3.4	Algebraic Geometry	16 MT
C3.5	Lie Groups	16 MT
C3.6	Modular Forms	16 MT
C3.7	Elliptic Curves	16 HT
C3.8	Analytic Number Theory	16 HT
C3.9	Computational Algebraic Topology	16 HT
C3.10	Additive Combinatorics	16 MT
C3.11	Riemannian Geometry	16 HT
C3.12	Low-Dimensional Topology and Knot Theory	16 HT
C4.1	Further Functional Analysis	16 MT
C4.3	Functional Analytic Methods for PDEs	16 MT
C4.4	Hyperbolic Equations	16 HT
C4.6	Fixed Point Methods for Nonlinear PDEs	16 HT
C4.9	Optimal Transport and Partial Differential Equations	16 MT
C5.1	Solid Mechanics	16
C5.2	Elasticity and Plasticity	16 MT
C5.4	Networks	16 MT
C5.5	Perturbation Methods	16 MT
C5.6	Applied Complex Variables	16 HT
C5.7	Topics in Fluid Mechanics	16 MT
C5.11	Mathematical Geoscience	16 MT
C5.12	Mathematical Physiology	16 MT
C6.1	Numerical Linear Algebra	16 MT
C6.2	Continuous Optimisation	16 HT
C6.5	Theories of Deep Learning	16 MT
C7.1	Theoretical Physics	24MT/16HT
C7.4	Introduction to Quantum Information	16 HT
C7.5	General Relativity I	16 MT
C7.6	General Relativity II	16 HT
C7.7	Random Matrix Theory	16 HT
C8.1	Stochastic Differential Equations	16 MT
C8.2	Stochastic Analysis and PDEs	16 HT
C8.3	Combinatorics	16 MT
C8.4	Probabilistic Combinatorics (see page 13)	16 HT
C8.6	Limit Theorems and Large Deviations in Probability	16 HT
C8.7	Optimal Control	16 HT