

Masters programmes

- Data Science and Analytics
- Machine Learning
- The Internet of Things
- Distributed and Networked Systems
- Computational Finance

with a Year-in-Industry pathway





“The Big Data market will almost double in the coming five years, from nearly \$39 billion in 2015 to more than \$76 billion in 2020.”

The Future of Big Data Analytics – Global Market and Technologies Forecast, 2015-2020.

Big Data and the Internet of Things

Big Data is now part of every sector and function of the global economy – from retailers, to health providers, to financial services and insurance, and to governmental and non-governmental organisations. Planning and strategic decision-making processes rely on large pools of data that need to be captured, aggregated, stored, and analysed. However, people with the right set of skills – data scientists – are in short supply and high demand.

Distributed and networked computation is the paradigm that underpins the software-enabled systems that are proliferating in the modern world, with huge impact in the economy and society, from modern data

centres and cloud computing platforms, to cyberphysical systems, to patient-centred healthcare, to disaster-recovery systems. The demand for specialists is soaring.

The Internet of Things is a new and fast expanding area in Computer Science, technology and engineering: it concerns the new generation of systems of networked devices that are becoming capable of sensing, transmitting and acting on data: smart homes, smart cities, smart industries, and so on. A series of reports, among which several published by Gartner, place the IOT as a new and transformative technology domain that will require millions of developers worldwide by 2020.



Studying Computer Science at Royal Holloway

Royal Holloway is a research-intensive university, ranked 19th in the UK and 129th in the world by the Times Higher Education World University Rankings for 2015-2016, where it is described as 'truly world class'.

Computer Science is an elite department with world-leading researchers. In the latest Research Excellence Framework (2014), we ranked 11th in the UK for the quality of our research output.

We are renowned in particular for our research excellence in Machine Learning, Information Security, and Global Ubiquitous Computing. This means that you will be taught by world-leading academics that are advancing the scientific disciplines and shaping the future of those key areas for society and the economy.

We also work closely with companies such as Facebook, IBM, NASA, GCHQ, Centrica (British Gas, Hive), Cognizant, QinetiQ, Thales UK, Orange Labs (UK), Yahoo! Research, the UK Cards Association, Transport for London and ITSO, among others. Thanks to our multiple connections with companies and organisations, you will also interact with many professionals during your studies, including former Royal Holloway students and post-docs who went on to work for Google, Microsoft, Amazon, and Facebook.

Research in Machine Learning at Royal Holloway started in the 1990's, at which time V. Vapnik and A. Chervonenkis (the inventors of Support Vector Machines) were both professors here. We have developed both fundamental theory and practical algorithms

that have fed into the analytics methods and techniques that are in use today. Current researchers include Alexander Gammerman and Vladimir Vovk – the inventors of conformal predictors theory, a radically new method of estimating the accuracy of each prediction as it is made. – and Chris Watkins, originator of reinforcement learning who developed 'Q-learning', a work that is fundamental to planning and control.

We also carry out world-leading research in all aspects of distributed computing and systems – including design and analysis of algorithms, large-scale and cloud-based systems, fault-tolerance, distributed storage, cloud computing, peer-to-peer, concurrency control, and multi-core computing – and in artificial intelligence, including cognitive and autonomous agents, automated planning, scheduling and domain-independent search control, and applications in surveillance operations, disaster response missions, space operations.

Our Information Security Group (ISG) is a world-leading interdisciplinary research centre, founded in 1990. We are an Academic Centre of Excellence in Cyber Security Research (ACE-CSR) and an EPSRC Centre for Doctoral Training in Cyber Security. Our Smart Card Centre was originally founded in 2002 by Royal Holloway, Vodafone and Giesecke & Devrient, and it has also received funding and practical support from a wide range of other organisations, notably Transport for London, the UK Cards Association, Orange Labs (UK), Visa and ITSO.

Highlights

Royal Holloway provides a unique environment in which to study. We are located close to London, in a campus that is considered to be one of the twelve most beautiful in the world and also the safest in England.

You will have 24/7-access to a dedicated on-site state-of-the-art large-scale data storage and processing facility, and equipment to experiment with interconnected IoT devices.

A weekly seminar series runs in parallel with the academic programme, which includes talks by professionals in a variety of application areas as well as workshops that will train you to find a placement or a job and lead a successful career.

Optional paid internships (up to one year) are available, which will allow you to gain experience and acquire skills that can only be picked up in a real work environment.

Entry requirements

Candidates should have, or expect to gain, at least an upper second-class honours degree (or the UK equivalent) in Computer Science, Economics, Engineering, Mathematics, Physics or other subjects that include a strong element of both mathematics and computing. Industrial experience may compensate for lesser degrees or lack of technical qualification, as will demonstrated programming skills and a strong ability to learn. All applications are treated on an individual basis. Non-native English-speaking applicants should hold IELTS 6.5, or a recognised equivalent qualification.

Fees

Studying for a Masters takes one year (12 months), up to two years if you take an industrial placement. Please check our web site for the fees that apply to your year of entry. If you join a Year-in- Industry programme, you will pay the one-year fee plus a smaller fee in the second year. Please note that you will be asked to make a deposit to accept your offer and receive the documentation that will enable you to obtain a visa.

Scholarships

For international students, a scholarship of £2,000 will be automatically awarded (and deducted from the fees payable) if you hold the equivalent of a UK first-class degree; you don't need to apply. Every year, we make available several other scholarships, some of which depend on your country of origin; please check our web site for the scholarships that apply to you.

Data Science and Analytics

This degree will teach you both the foundational aspects and the practical skills that prepare you for handling and analysing different types of data in different fields, thus responding to the needs of a huge variety of companies and organisations, from retailers, to car manufacturers, health-care providers, or public administration.

Skills acquired

You will acquire the ability to:

- work with methods and techniques such as clustering, regression, support vector machines, boosting, decision trees, and neural networks
- work with structured, unstructured, and time-series data
- work with software packages such as MATLAB and R
- design data processing solutions for data-intensive analytics problems
- design Extract-Transfer-Load (ETL) pipelines
- design Data Warehousing and Decision Support System (DSS) solutions
- work with highly-scalable data-storage solutions, such as MongoDB, Cassandra, HBase, and other NoSQL Data Stores
- work with data-intensive computing technologies, such as Hadoop MapReduce, Spark, Hive, and Pig
- work with Cloud Computing tools, such as Amazon S3, EC2 and Elastic MapReduce

Outline

	Taught modules	Credits
Term 1	Data analysis	20
	Programming for data analysis	10
	Computation with data (*)	10
	Database systems (*)	10
	Electives	(+)
Term 2	Large-scale data storage and processing	20
	Electives	(+)
Exams		
Placement (for Year-in-Industry degree), up to one year		
Individual project / dissertation (12 weeks) – 60 credits		

(*) Compulsory and available only for students who lack background in the corresponding area.

(+) Electives are chosen from the list in page 11 to top up the total term value to 60 credits

Machine Learning

This degree allows you to develop a deeper understanding of the science of systems that can learn from data, what companies such as Facebook, Google, Microsoft and Yahoo require to create, innovate, and define the next generation of search and analysis technologies.

Skills acquired

You will acquire the ability to:

- develop, validate, and use effectively machine learning and statistical models
- work with structured, unstructured, and time-series data
- extract value and insight from data
- work with methods and techniques such as clustering, regression, support vector machines, boosting, decision trees, and neural networks
- develop and use universal prediction algorithms, including universal strategies for dynamic investment
- complement predictions with provably valid measures of accuracy and reliability.
- work with software packages such as MATLAB and R

Outline

	Taught modules	Credits
Term 1	Data analysis	20
	Programming for data analysis	10
	Computation with data (*)	10
	Machine learning (*)	10
	Electives	(+)
Term 2	On-line machine learning	20
	Inference <i>or</i> Applied probability	20
	Electives	(+)
Exams		
Placement (for Year-in-Industry degree), up to one year		
Individual project / dissertation (12 weeks) – 60 credits		

(*) Compulsory and available only for students who lack background in the corresponding area.

(+) Electives are chosen from the list in page 11 to top up the total term value to 60 credits

The Internet of Things

This degree provides you with advanced knowledge and skills in three essential and critical areas: data analytics techniques through which one can make sense from and act on the data that is collected from thousands of sensors; technical and methodological aspects of distributed and networked systems; and cybersecurity aspects such as privacy and safety.

Skills acquired

You will acquire the ability to:

- design and implement middleware services for reliable communication in unreliable networks, and storage solutions for wireless, sensor, and ad hoc-networks
- detect sources of vulnerability in networks of connected devices and deploy the appropriate countermeasures to information security threats
- enforce privacy in “smart” environments
- work with open source and cloud tools for scalable data storage (DynamoDB) and coordination (Zookeeper)
- design end-to-end IoT solutions involving interconnected edge devices and clouds
- work with modern network management technologies (Software-Defined Networking) and standards (OpenFlow)
- work with low-power wireless and mesh networking standards and technologies such as IEEE 802.15.4, ZigBee and XBee
- work with state-of-the-art microcontroller devices and kits, such as Arduino and Tessel, and miniature computing technologies, such as RaspberryPi
- work with methods and techniques such as clustering, regression, support vector machines, boosting, decision trees, and neural networks

Outline

	Taught modules	Credits
Term 1	Data analysis	20
	Interconnected devices	10
	Information security (*)	10
	Advanced distributed systems	20
	Electives	(+)
Term 2	Wireless, sensor and actuator networks	20
	Smart cards, RFIDs and Embedded systems security	20
	Electives	(+)
Exams		
Placement (for Year-in-Industry degree), up to one year		
Individual project / dissertation (12 weeks) – 60 credits		

(*) Compulsory and available only for students who lack background in the corresponding area.

(+) Electives are chosen from the list in page 11 to top up the total term value to 60 credits

Distributed and Networked Systems

This degree is aimed at graduates who already have a background in computer science or engineering and who wish to acquire the depth of knowledge and the skills required to help design, deploy and use the technologies through which systems can operate in networks and in a distributed way. You will be able to choose many electives – including modules in cybersecurity, large-scale data storage and processing, and artificial intelligence – so that you can tailor your degree according to your preference.

Skills acquired

You will acquire the ability to:

- analyse complex distributed systems in terms of their performance, reliability, and correctness
- design and implement middleware services for reliable communication in unreliable networks
- design and implement reliable data communication and storage solutions for wireless, sensor, and ad-hoc networks
- work with open source and cloud tools for scalable data storage (DynamoDB) and coordination (Zookeeper)
- design custom-built application-driven networking topologies using modern tools
- work with modern network management technologies (Software-Defined Networking) and standards (OpenFlow)
- work with relational databases (SQL), NoSQL databases (MongoDb), as well as with Hadoop/Pig scripting
- work with low-power wireless and mesh networking standards and technologies, such as IEEE 802.15.4, ZigBee and XBee
- work with state-of-the-art microcontroller devices and kits, such as Arduino and Tessel, and miniature computing technologies, such as RaspberryPi

Outline

	Taught modules	Credits
Term 1	Interconnected devices	10
	Advanced distributed systems	20
	Electives	(+)
Term 2	Wireless, sensor and actuator networks	20
	Electives	(+)
Exams		
Placement (for Year-in-Industry degree), up to one year		
Individual project / dissertation (12 weeks) – 60 credits		

(*) Compulsory and available only for students who lack background in the corresponding area.

(+) Electives are chosen from the list in page 11 to top up the total term value to 60 credits

Computational Finance

This degree, offered in conjunction with the department of Economics, allows you to specialise in modern quantitative finance and computational methods for financial modelling, which are demanded for jobs in asset structuring, product pricing as well as risk management.

Skills acquired

You will acquire the ability to:

- analyse, critically evaluate, and apply methods of computational finance to practical problems, including pricing of derivatives and risk assessment
- analyse and critically evaluate methods and general principles of computational finance and their applicability to specific problems
- work with methods and techniques such as clustering, regression, support vector machines, boosting, decision trees, and neural networks
- analyse and critically evaluate applicability of machine learning algorithms to problems in finance
- implement methods of computational finance and machine learning using object-oriented programming languages and modern data management systems
- work with software packages such as MATLAB and R
- work with Relational Database Systems and SQL

Outline

	Taught modules	Credits
Term 1	Data analysis	20
	Programming for data analysis	10
	Foundations of finance	20
	Database systems (*)	10
	Electives	(+)
Term 2	Investment and portfolio management	20
	Electives	(+)
Exams		
Placement (for Year-in-Industry degree), up to one year		
Individual project / dissertation (12 weeks) – 60 credits		

(*) Compulsory and available only for students who lack background in the corresponding area.

(+) Electives are chosen from the list in page 11 to top up the total term value to 60 credits

Full list of courses^(*)

Course name	Credits
Advanced Data Communications	10
Advanced Distributed Systems	20
Applied Probability	20
Business Intelligence Systems, Infrastructures and Technologies	20
Computation with Data	10
Computational Optimisation	10
Computer Security	20
Concurrent and Parallel Programming	10
Corporate Finance	20
Cyber Security	20
Data Analysis	20
Data Visualisation and Exploratory Analysis	10
Database Systems	10
Decision Theory and Behaviour	20
Digital Forensics	20
Financial Econometrics	20
Fixed Income Securities and Derivatives	20
Foundations of Finance	20
Inference	20
Information Security	10
Intelligent Agents and Multi-Agent Systems	10

Course name	Credits
Interconnected Devices	10
Introduction to Cryptography	20
Investment and Portfolio Management	20
Large-scale Data Storage and Processing	20
Machine Learning	10
Methods of Bioinformatics	10
Methods of Computational Finance	10
Micro-econometrics	20
Network Security	20
On-line Machine Learning	20
Programming for Data Analysis	10
Security Management	20
Security Technologies	20
Security Testing	20
Semantic Web	10
Smart Cards, RFIDs and Embedded Systems Security	20
Software Security	20
Topics in Applied Statistics	20
Wireless, Sensor and Actuator Networks	20

(*) Subject to availability (some electives may not be taught in some years) and timetabling constraints.

Outline of main courses

Advanced Distributed Systems – The course covers fundamental principles of building modern distributed systems with an emphasis on cloud infrastructure and wireless networking, covering topics such as massive scale, wide distribution, decentralisation, unreliable communication links, component failures and network partitions. It addresses abstract models (such as the synchronous and asynchronous distributed computing models, models for wireless networks), algorithmic techniques (such as distributed coordination, fault-tolerant design of distributed algorithms, synchronization techniques), and practical case studies.

Applied Probability – The course introduces the student to a range of examples of probabilistic methods used to model systems that exhibit random behaviour. Key classes of models covered include Poisson processes, renewal theory, and Markov processes.

Business Intelligence Systems, Infrastructures and Technologies – Business Intelligence refers to the skills, processes, methodologies, technologies, applications, and practices used in order to leverage (gathering, storing, analyzing) an organization's internal and external information assets to support and improve decision-making. This course aims to provide students with a broad understanding of the information assets and the conceptual and technical architectures of information and business intelligence systems in modern organizations, and the necessary background knowledge of, and skills to design, implement and evaluate business intelligence systems and technologies.

Computation with Data – The course covers programming in Java and general principles of object-oriented programming. It is aimed at students with weaker programming skills.

Data Analysis – The course addresses fundamentals of data analysis with a focus on statistical learning. It covers all stages of data analysis from preprocessing to interpreting

the results. It will teach students to develop, validate, and use effectively machine-learning and statistical models and algorithms; including their implementation in R.

Database Systems – This course is aimed at students with a weaker background on the core concepts of data and information management. It covers the entity-relationship model, relational database design, SQL (DDL, DML) and DCL relational data bases.

Foundations of Finance – The course covers the structure of financial markets, the instruments traded and the participants. It provides students with tools to analyse how financial markets function and how problems arise from their operations.

Inference – This course covers the main principles and methods of statistics. It offers students an introduction to parametric estimation and hypothesis testing. It covers key concepts such as maximum likelihood, method of moments, Bayes estimators, sufficiency, unbiasedness, efficiency, asymptotic properties of maximum likelihood estimators.

Interconnected Devices – The course introduces students to IoT design considerations, constraints and interfacing between the physical world and the IoT devices. It then focuses on how the Arduino platform works in terms of the physical board, the libraries and the IDE (integrated development environment). It covers how to program the Arduino via C/C++ code and how to access the pins on the board via the software to control external devices. Finally, it provides hands-on knowledge on how to plug shields into the main Arduino board to perform other functions such as sensing and actuating.

Investment and Portfolio Management – The course covers the underlying theory and empirical evidence in portfolio management in the finance sector. Students will acquire an understanding of how funds are allocated when constructing a portfolio.

Large-scale Data Storage and Processing

– The course covers advanced concepts in distributed systems, networking, databases, and cloud computing. It focuses on principles underlying modern data-intensive computing technologies such as NoSQL Data Stores and Hadoop, and advanced database concepts (OLAP and parallel databases) that are at the heart of ETL, Data Warehousing and DSS platforms. Students acquire hands-on experience in modern data-intensive platforms (such as Hadoop and NoSQL) for solving data analysis problems on real-life datasets.

Machine Learning – The course provides an in-depth coverage of linear and kernel methods of machine learning with a focus on support vector machines and regression. The course enables the students to implement and apply the methods and appraise the results.

Methods of Computational Finance –

The course enables the students to apply mathematical and computational methods for pricing and hedging derivative securities and portfolios of securities.

On-line Machine Learning – The course addresses the on-line framework of machine learning in which the learning system learns and issues predictions or decisions in real time, perhaps in a changing environment. Students will learn how to develop and use universal prediction algorithms, including universal strategies for dynamic investment, and how predictions can be complemented with provably valid measures of their accuracy and reliability.

Programming for Data Analysis – The course covers MATLAB, one of the main systems for numerical programming and data analysis used in industry and academia. At the end of the course students will be able to use MATLAB in solving efficiently important classes of problems arising in data analysis.

Smart Cards, RFIDs and Embedded

Systems Security – This course provides an overview of smart cards/RFIDs/Near Field

Communications (NFC) and properties; introduces applications exploiting smart cards/RFID/NFC including the Internet of Things (IoT); examines benefits, threats and attacks when used as assets for Cyber Security; considers development, manufacture and management of smart cards/RFID/NFC; reviews related standards and security evaluation methodologies for embedded security; and considers/compares related technology e.g. TEE, TPM & Android Host Card Emulation (HCE).

Visualisation and Exploratory Analysis –

This course covers topics in: Construction of informative bivariate plots, including smoothing with loess, visualisation of distributions, and cumulative distributions and QQ plots; Visualisation of multivariate data, including Kernel PCA, canonical correlations analysis (CCA) and Kernel CCA; Dimensional reduction; Exploratory cluster analysis, and metrics for comparison of clusterings; Standard methods for visualisation of relational and graph data, including practical implementations such as Gephi; Principles of selecting modes of presentation for various quantities, and basic principles of colour scale design and glyph choice.

Wireless, Sensor and Actuator Networks –

The course combines formal lectures that will focus on the algorithms and the protocols behind wireless sensor and actuator networks and lab classes that will focus on how to build wireless sensor and actuator networks for a variety of applications. The lectures will cover critical design factors for WSANs, the protocol stack, models and algorithms for WSANs, routing protocols and more advanced open research problems, such as topology control and mobility. The practical classes will cover how to design and build wireless sensor networks and intelligent interactive devices with the ZigBee wireless networking protocol.

Placements

All programmes are offered with an optional paid industrial placement (internship). Because the internship is part of your studies, your visa will be issued to cover the full two-year duration of the programme.

Although the responsibility for finding a placement is ultimately with the student, our Careers Team will help you identify suitable opportunities, make applications and prepare for interviews.

We bring several companies to our campus throughout the year, both for fairs and for delivering advanced topics seminars, which are an excellent opportunity to learn about what they do and discuss possible placements. We also hold special events for companies where we present them our programme, especially the skills that you will have acquired and the benefits that they will have by offering you a placement.

In 2013/14, all students in the Year-in-Industry pathways who were allowed to progress have secured placements at Centrica (four students), UBS (two students), HACT, IMS Health, Shell, and Société Générale.



For the 2014/15 cohort, placements have been secured at Centrica (four students), UBS (two students), BGL Group, the Office for National Statistics (ONS), Gamesys, Rolls Royce, Cartesian, Shell, IMS and BGL. This represents 85% of all students who qualified for a placement. Every employer who placed students in 2013/14 came back to us and asked for one or more students, which is an excellent testimony of the quality of our degrees.

So far, 60% of international students who have secured a placement in the Year-in-Industry have been sponsored to remain working in the UK.

Employers on placement students

Christopher Sherrington, Centrica

“Prasen and Umesh were placed as data scientists with British Gas, attached to the Enterprise Architecture team in order to help us develop data science as discipline within the business. They built a predictive analytics capability into our smart meter installations processes. They built their own tools, planning their own time and presenting their findings to their business customers. They are a great example of the value that can come from partnership between academic institutions and businesses.”

Peter Frampton, UBS

“Asher did some great work in allowing UBS us to extract commercial value and competitive advantage from the poly-structured data held by HR. The first stage deliverable went into production in the UK and EMEA in the fourth month of his rotation with a global rollout currently underway. This was well ahead of our expectations and proves the game-changing potential that this skillset can bring to our organisation and profession.”

Project work

Project work is at the core of our degree programmes, through which students gain hands-on experience in using a wide range of techniques, paradigms, and tools.


The Individual project is a substantial piece of work (worth 60 credits, one third of the degree) that students undertake at the end of their degree for 12 weeks under the supervision of an academic member of staff.

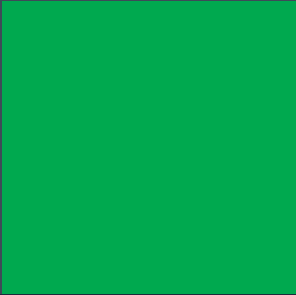
Throughout the course, several modules will require more or less substantial pieces of project work. For example, in CS5324 (Large-Scale Data Storage and Processing), students analysed, in 2013/14, the Enron Corpus, which is a large database of over 600,000 emails generated by 158 employees of the Enron Corporation.


They used MapReduce and Hadoop for cleaning the dataset, extracting a social network graph induced by the individual emails, and analysing its structural properties (such as degree distribution and number of triangles).

They then used Gephi – a social network visualisation software – to glean insight into social relationships between the individual employees within the organisation (such as communities, management hierarchies, etc.). See picture below.

The Department operates a dedicated state-of-the-art computing cluster consisting of 7 CALLEO Infrastructure Servers collectively offer 36 CPU cores, 96GB of RAM, and 64TB of disk storage. The cluster is currently running the full-fledged Hortonworks Hadoop distribution (HDP) as well as the PostgreSQL and MongoDB servers. It also serves as a repository for massive data sets provided by our industrial and academic partners, and is in active use by our advanced degree programme students (MSc/MSci/PhD) and academics as a training and research platform.





 Royal Holloway, University of London
Egham, Surrey, TW20 0EX
T: 01784 434455
royalholloway.ac.uk
msc-admissions@cs.rhul.ac.uk