Scaling up the Ethical Artificial Intelligence MSc Pipeline

Final report to the Office for AI covering frameworks and models of Industry-funded Masters Courses in Machine Learning and related AI skills,
for Grant A1845, May 2019

BCS - The Chartered Institute for IT
First Floor, Block D
North Star House
North Star Avenue
Swindon SN2 1FA

BCS is a registered charity no: 292786

Table of Contents

1 Executive Summary

This report includes:

Recommendations covering	Findings covering
• Postgraduate skills evidenced against standards • Models for MSc programmes that integrate professional development with academic excellence • Diversity and inclusivity • Pipeline from arts and humanities	• Employer demands and needs • University interest, capacity and capability • Student needs • Ethical behaviours • Delivery models

Top priority: Creating a diverse, inclusive pipeline of ethical, competent and talented MSc graduates highly skilled in Machine Learning and Artificial Intelligence

The government's 2017 review[1] of Artificial Intelligence recommended that industry and universities, supported by Government and relevant intermediaries, work together to develop industry-funded AI Masters courses to meet the high demand from industry for skills at this level. Following on from that review the Office for AI commissioned BCS to look at possible frameworks and models for industry funded Masters courses in Machine Learning and related AI skills, which is the subject of this report. The report is based on feedback from an extensive series of consultations with employers, universities, students, government and other relevant stakeholders. There was broad consensus from these consultations that an overarching priority should be the creation of a diverse, inclusive pipeline of ethical, competent and talented people graduating from Masters programmes highly skilled in Machine Learning and Artificial Intelligence. A high-level model illustrating key desirable characteristics of an MSc programme, synthesised from the views expressed in the consultations, is shown in Figure 1 (a logic model derived from this is given in Figure 4).

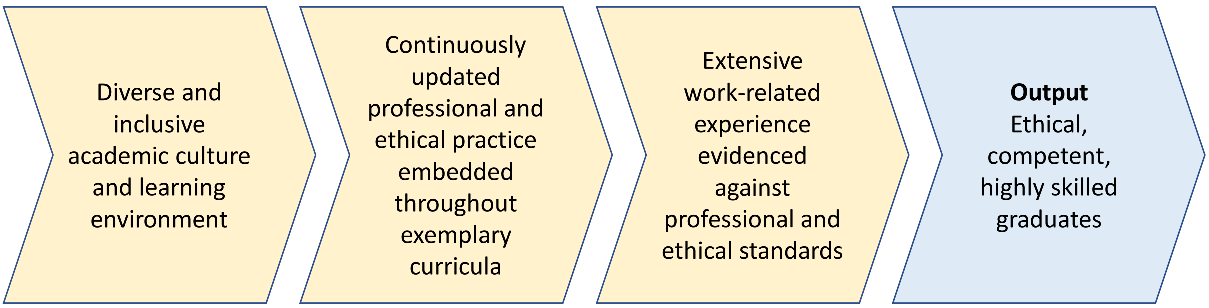

Figure 1: MSc high-level model based on feedback from consultations

The expert opinions and advice in this report are therefore framed in terms of this consensus view and the associated high-level model.

[1] https://www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in-the-uk

Recommended actions following from this report

Government is encouraged to commission a plan of action bringing together universities, employers, professional bodies, the Office for Students, EPSRC, the Office for AI, the Royal Academy of Engineering, the Institute of Coding, the Alan Turing Institute, the Centre for Data Ethics and Innovation and other relevant stakeholders to:

Recommendation 1: Develop and maintain ethical and professional AI standards against which MSc graduates can evidence they have gained the appropriate level of skills required to contribute to the design, development, deployment, management and maintenance of AI products and services.

Recommendation 2: Ensure there are a wide range of MSc courses offered across the UK in Machine Learning and AI based on the model in Figure 1, which provide students with extensive work-related experience that is evidenced and accredited against ethical and professional standards, including through degree apprenticeships at level 7 and MSc programmes with structured work placements.

Recommendation 3: Implement the computing contextualised recommendations from the diversity and inclusivity reports of the Royal Academy of Engineering[21] and the American Astronomical Society[18] to bring about a sustained, systematic, evidence based change in the academic culture and learning environment of postgraduate computing education. Including, amongst others, that university computing departments:

- Develop excellence criteria to be included in inclusive postgraduate admission policies that take account of an applicant's potential to become an outstanding practitioner and consider a range of more broad technical skills beyond those particular to computing.
- Partner with and recruit from undergraduate programmes that produce large numbers of graduates from underrepresented groups.
- Ensure retention support for students is made inclusive by providing students from all backgrounds equal access to an extensive, high quality and properly resourced range of professional development, networking, and mentoring opportunities.
- Are supported to achieve at least Athena Swan Silver award.

Recommendation 4: Ensure there are a broad range of arts and humanities undergraduate degree courses that include quantitative methods as a significant 'with' component offered across the UK (e.g. such as BA in Social Sciences with Quantitative Methods), which are partnered with MSc AI conversion courses for non-STEM graduates.

A timeframe for these actions needs to be carefully worked through with all the relevant stakeholders, but they should be implemented with a sense of urgency for the UK to remain at the forefront of AI.

The rest of the report is divided into sections covering employer demands and needs (Section 2), university interest, capacity and capability (Section 3), student needs (Section 4), ethical behaviours (Section 5) and delivery models (Section 6), which are summarised below.

Employer Demands and Needs (Section 2)

The consultations for this report have consistently shown there is high demand from employers for ethical and competent AI practitioners. Employers have consistently and strongly expressed the view they want a diverse, inclusive, substantial pipeline of talented Machine Learning and AI postgraduates from HEIs[2]. These points are also further discussed in Section 4.1 and Section 6.

Section 2.1 synthesises the overarching professional skills employers have articulated they believe characterise an ideal AI practitioner. Employers highlighted the need for people who are highly skilled at transferring a deep scientific knowledge of AI into business contexts, engineering AI systems that meet business needs, managing the adoption of AI and maximising its value across strategic business units, and ensuring AI systems meet appropriate ethical standards, amongst other desirable skills. Working from these broad characterisations as our starting point, Section 2.1 outlines a high-level set of professional skills split across science skills, data engineering skills, product development skills, business skills and ethical competencies, which have been widely shared with employers for validation. Our findings suggest the level of skills employers identified are more advanced than can realistically be expected of a typical MSc graduate, but they can usefully inform criteria for the types of skills students should develop through work-related experience. This section of the report also includes an evaluation of existing degree level apprenticeships as a vehicle for delivering these types of skills (Section 2.2 and Section 6.1) along with various other options such as industrial placements assessed against industry recognised standards (also discussed in Section 6.2 and Section 6.3).

We found there is a need to create a suitable infrastructure to broker connections between employers, universities and students to help employers determine which university courses best suits their needs, help universities understand how best to provide courses employers and students value, identify how employers can support work experience opportunities as part of MSc courses, and help students find the right employer and university for their needs. The Institute of Coding could be a vehicle for providing such a brokering service.

University interest, capacity and capability (Section 3)

When looking at taught MSc courses that are currently on offer, we found there is extensive capacity throughout the UK to provide high quality Machine Learning and AI MSc courses (see Section 9 for details of the courses considered for this report). This section includes our in-depth analysis of MSc syllabi from across the UK (see Section 3.1 and Section 3.2), from which we conclude there is a growing sense of a shared Machine Learning body of knowledge across universities that is highly relevant to the needs of employers outlined in Section 2. In contrast to our in-depth analysis, feedback from the consultations suggests the rich diversity of MSc courses currently on offer could be confusing for both students and employers. In our view, given such potential for confusion, there is a strong case for succinctly and coherently articulating the offering from universities to students and

[2] Higher Education Institutes

employers through skills frameworks that are tailored to their needs (see also Section 4.2.1, Section 6.4). This is important, because otherwise employers might focus on building relationships with a small number of well-known universities to the possible exclusion of other universities wishing to develop new AI courses, and students may focus on courses that don't fully meet their needs or aptitudes. That could inhibit the capacity of the UK to develop a strong AI practitioner pipeline over the long term.

Our findings suggest there is currently a gap in the capability and capacity of many MSc programmes to develop professional skills through extensive work-related experience, whilst there are also MSc courses which are exemplars of how this can be achieved. Many universities reported that adding lengthy placements, such as for nine months, to an already existing one year MSc course will be challenging. One of those challenges is finding enough employers who are willing and able to provide such placements, especially since the new post-16 T-level qualifications include three months of compulsory work experience that is likely to further reduce the capacity of employers to support placements for MSc programmes. Universities are innovating their teaching practices to develop students' professional skills through a variety of work-related experiences both on and off campus, which could help address the work experience issue and should be widely shared across the HE sector (see Section 4.2.3). For example, some universities have developed innovative, employer informed curricula, which combines significant employer engagement on campus with short periods of focused, structured work experience off campus that builds on the professional development gained through the campus based part of the course.

This section also discusses the relevance of existing structured skills frameworks, such as for example SFIA[4] or DDAT[3], for universities to consider using when supporting student's professional development through work-related experience (see also Section 6.3). For example, BCS, in its role as the UK's professional body for IT, is currently consulting with various universities to see if it is possible to provide a greatly accelerated pathway to Chartered IT Professional status for graduates of AI MSc programmes that include suitable work-related experience.

Student Needs (Section 4)

When we surveyed UK student members of BCS (see Section 4.2.1) they rated certification of both their competency and capability as their most important professional development need. These students have a clear view of their professional development needs and they see demonstrating their capabilities and competencies against industry recognised standards as important for their career development. Based on consultation feedback, our expert opinion is that AI MSc courses should ensure students are able to evidence and articulate how their stated professional development priorities have been achieved on successful completion of the programme.

Throughout all the consultations held to inform this report the lack of diversity and inclusivity in university computing departments has been raised as a priority issue that needs to be addressed (Section 4.1). Strongly held views were consistently expressed that computing departments in HEIs can and should make their subject far more diverse and inclusive, even though that will require significant sustained effort at all levels. In this report we have mainly cited sources that discuss diversity in terms of gender, ethnicity and social background. It is worthwhile noting that for some employers diversity also encompasses nationality, educational background, age profile, and work experience in addition to gender, ethnicity and social background.

There is a growing body of expert opinion, from national organisations such as for example the Royal Academy of Engineering and BCS, that every computing student of whatever gender, ethnicity or social background is held back academically and professionally by a lack of diversity and inclusivity in the culture and learning environments of university computing departments. In our view, for AI MSc programmes to deliver the best student outcomes, fully support the UK Industrial Strategy, and provide value for money, issues around diversity and inclusivity must be dealt with. There is a growing body of research and good practice from a variety of authoritative sources that propose evidence-based solutions, which we believe should be seriously considered for the UK (See Section 4.1.2).

Given the expected impact of AI and Machine Learning to all sectors of the economy and throughout society it is important that there are effective progression pathways for arts and humanities graduates to gain Machine Learning and AI related skills at Masters level, which includes progression to AI MSc conversion courses (as recommended in the AI review[1]). However, emerging research from the Nuffield foundation suggests there are significant learning barriers for students exposed to quantitative methods at undergraduate and postgraduate level if they have not studied them since leaving school[27].

Our conclusion is that ideally non-STEM undergraduates should have opportunities to study quantitative methods that are applicable and relevant within their degree, and which prepare them for progression onto Machine Learning and AI MSc conversion courses that are tailored specifically for them. For that to happen there needs to be a wide range of arts and humanities degrees that include quantitative methods options with their main subject, such as for example a BA in Social Sciences with Quantitative Methods. This idea has already been piloted through the Q-Step project in fifteen universities, which has been running over the last five years. The lessons being learnt from Q-Step should be built on to ensure arts and humanities undergraduate degree programmes with quantitative methods provide appealing, scalable and sustainable progression routes into AI MSc conversion courses. See Section 4.1.3 for further details.

Ethical behaviours (Section 5)

The history of IT is littered with examples of new technologies where security and ethical concerns were seen as ‘add-ons’ that could be deferred until the technology was ready to be deployed, but which in practice turned out to be incredibly hard to add on after the fact. Given the potential impact of AI on so much of the modern world it is essential that AI practitioners have the competencies necessary to ethically design, ethically develop, ethically deploy, ethically manage and ethically maintain AI products and services. At the same time it’s important that AI ethical frameworks do not stifle innovation and do not become a burdensome, over regulated instrument of policy implementation.

Our conclusion, based on the findings from the various consultations, is that it is essential AI MSc graduates should be able to evidence they have the ethical competencies necessary to significantly contribute towards embedding ethical principles throughout the AI product/service lifecycle (see the high level ethical competencies overview in Section 2.1). Ethical practices should be embedded as a core part of AI MSc curriculum, including being a core component of technical lab exercises, assessed projects and integral to core technical subject matter modules. Independent accreditation of AI MSc courses could be one useful mechanism for encouraging embedding 'ethical by design' within curricula, as part of a range of incentives for universities to adopt this approach.

The new Centre for Data Ethics and Innovation in collaboration with BCS will be holding a consultation exercise with stakeholders to gather current professional ethical practice including gathering practice on assurance of ethical standards. This will be used as the basis for future guidance made available to help universities teach 'ethical by design' in MSc programmes.

Delivery Models (Section 6)

Employers stated preference is for students to develop AI related professional skills through extensive work experience, which should be assessed against industry recognised standards (See Figure 4 and Section 6.3). This is somewhat different to how many universities currently provide work related experience. The employer led RITTech standard is one example of how undergraduate IT sandwich degree programmes can provide placements aligned to professional registration. Another example in a different sector where professional standards are embedded within curricula are Chartered Management Degrees, which are level 6 degree apprenticeships. Our conclusion is that models developed for delivering and assessing professional skills in AI MSc courses should consider building on examples such as these.

We found that most large employers would prefer industry funded AI MSc courses be delivered through level 7 degree apprenticeships funded through the apprenticeship levy. Currently there isn't capacity to provide enough MSc degree apprenticeships to achieve the government's stated long-term aspiration of an additional 3,000 places (see Section 6.1). The case study in Section 2.2 illustrates some of the practical challenges universities may need to overcome in establishing MSc degree apprenticeships, but which also suggests those challenges can be overcome. We conclude that for the medium term at least other delivery models will need to be considered while significant MSc degree apprenticeship capacity is developed.

During the consultations for this report many large employers expressed their support for student scholarships as a means of increasing capacity on MSc courses that are not degree-apprenticeships, but which do include extensive work-related experience. Scholarships are seen as a means of providing greater flexibility for students, employers and universities and allows for more agile innovation in curricula than the apprenticeship model.

As pointed out in Section 4.2.3, valuable work-related experience can be provided through a variety of different models. While the traditional nine month placement is regarded by many as the gold-standard, we found it is also possible for students to gain valuable work-related experience through courses where curricula are significantly informed and supported by employer input, and which include short internships or placements, provided these are structured and evidenced through professional skills frameworks.

The Shadbolt review[28] of undergraduate Computer Science degrees stated: '*there is a mandate from employers and the HE sector to strengthen the current accreditation framework so that it is more focused on outcomes and links more closely with employability*.' We finally conclude that accreditation of university AI MSc courses by professional bodies should follow the recommendations of the Shadbolt review and evolve through engagement with employers and universities to provide a mechanism for facilitating innovative work-related experience and professional development in degree programmes.

Consultation Stakeholders

BCS is grateful for the advice, support and feedback from individuals acting in a personal capacity from the following organisations:

- Practitioners working for Amplify, ARM, BAE Systems, Cambridge Consultants, Cisco, DeepMind, Deloitte, HSBC, IBM, Infosys, Lloyds, McKinsey Quantum Black, Microsoft, Nvidia, Ocado
- Officials from DCMS, The Office for AI, DfE, the Office for Students, and BEIS
- The more than 50 university computing departments that responded to consultations through UKCRC and CPHC, as well as the various roundtables and workshops that informed this report
- The Institute of Coding, the Royal Academy of Engineering, Technology Catapults, TechUK, the Data Skills Taskforce, and the Alan Turing Institute
- All the BCS students, professional members, Chartered IT Professionals and Fellows who provided insight and advice during the preparation of this work

Advisory Working Group for this report

This report was written with the advice and guidance of the working group members listed below. Members acted in an individual and not a representative capacity. Members contributed to the project on the basis of their own expertise and good judgement.

- Abhijit Akerkar, Head of AI Business Integration, Lloyds Banking Group
- Professor Jeff Magee, Presidents' Envoy at Imperial College London (previously Dean of the Faculty of Engineering at Imperial College London)
- Jonathan Legh-Smith, Head of Scientific Affairs at BT
- Mike Warriner, Digital CIO at HSBC Retail Bank & Wealth Management (previously Google Engineering Director)
- Dr Nicolaus Henke, Global leader of McKinsey Analytics and Chairman of QuantumBlack
- Rebecca George OBE, Vice Chair and UK Public Sector Leader at Deloitte
- Stephen Pattison CMG, VP Public Affairs at ARM
- Dr Vicky Schneider, academic visitor at University of Cambridge, Honorary Associate Professor University of Melbourne, and Senior Scientific Program Manager at the Amazon Development Centre in Cambridge

The working group was coordinated by Dr Bill Mitchell, Director of Policy at BCS. BCS would like to express our thanks to the group for their invaluable help preparing this report.

2 Employer Demand and Needs

Throughout the preparation for this report employers of all shapes and sizes have strongly agreed there is a need for:

- Masters level graduates capable of ethically applying advanced Machine Learning and other Artificial Intelligence techniques to create significant business value, and have a broad range of professional skills related to science, engineering, product development and business practice (See Section 2.1 for more details)
- creating a diverse and inclusive pipeline of students recruited to Masters level programmes (diversity and inclusivity are covered in Section 4 on Student Needs rather than here)
- a level 7 MSc Apprenticeship as a delivery mechanism for industry funded Masters programmes in Machine Learning and AI related skills (see Section 2.2), which would be in addition to individual scholarships available through bilateral employer university arrangements

Based on employer consultations the key messages relating to development of professional skills are:

- There is general agreement on a desirable set of overarching AI and Machine Learning professional skills that are relevant to the overwhelming majority of employers (which are outlined in Section 2.1).
- Employers tend to have specialist roles, such as for example Data Scientist or Data Engineer, which only require a subset of the overarching desirable high-level skills, and they tend to rely on interdisciplinary teams who can collectively cover the full range of desired skills.
- There are degree level apprenticeships related to Data Science (see Section 2.2) that have significant overlap with the high-level AI Skills Framework, but there are significant gaps that mean they are not an exact fit for the employer needs articulated in the skills framework.
- Unicorns do not exist, i.e. it is extremely unlikely that any one individual will have all of the overarching desired professional skills employers have identified. That also means it is extremely unlikely an MSc course can develop graduates with all these skills. Hence, industry needs are likely to have to be met by a diverse range of MSc programmes tailored to distinct specialist roles.

2.1 First draft high-level AI Skills Framework

This section outlines a draft high-level skills framework for an AI Practitioner, synthesised from the employer consultations for this report. This has been widely circulated across a diverse range of employers engaged in different sectors of the economy. The majority of the feedback from employers is that this is a reasonable overview of desirable AI skills, but will need careful refinement and contextualisation if used to characterise specific skills for particular roles. What this draft framework is not is a set of competencies for the skills that should be achieved by an MSc student during an industrial placement. What it can do is inform the criteria used for determining the appropriate levels of competencies that can be expected from a quality work placement at MSc level.

Science skills for Artificial Intelligence Professionals

An AI Practitioner:

- Understands how AI and particularly Machine Learning algorithms, techniques and methodologies are designed, developed, optimised and applied at scale to achieve business objectives.
- Can select and use appropriate statistical methods for sampling, distribution assessment, bias and error.
- Understands AI and particularly Machine Learning problem structuring methods and can evaluate which method is most appropriate for business needs.
- Applies rigorous scientific methodologies through experimental design, exploratory modelling and hypothesis testing to reach robust conclusions, and can explain how those are reached to internal and external stakeholders.

Data engineering skills for Artificial Intelligence Professionals

An AI Practitioner:

- Has a demonstrable understanding of how to expose data from systems, how to efficiently extract data from potentially heterogeneous source systems, and how to ensure standards of data quality and consistency for processing by AI systems.
- Works with other technologists and analysts to integrate separate data sources in order to map, produce, transform and test new scalable AI products and services that meet user needs.
- Works with other technologists and analysts to understand and make use of different types of data models.
- Understands and can make use of different data engineering tools for repeatable data processing and is able to compare between different data models.
- Understands how to build scalable machine learning pipelines and combine feature engineering with optimisation methods to improve the data product performance.

Product development skills for Artificial Intelligence Professionals

An AI Practitioner:

- Uses a range of professional coding practices to build reliable, reusable, scalable AI products and services to time, quality and budget
- Can work as part of a team to effectively integrate AI technologies into business systems.
- Can take into account non-functional requirements such as system performance and integration requirements as part of an enterprise systems perspective
- Understands the enabling infrastructure required to support AI technologies

- Can demonstrate why AI products and services are valid against user requirements in a manner comprehensible to the relevant internal and external stakeholders.
- Works in accordance with agreed software development standards, including security, accessibility and version control.

Business skills for Artificial Intelligence Professionals

An AI Practitioner:

- Understands the context of the business including its processes, data, priorities and its wider values, objectives and strategy.
- Works collaboratively with domain experts to fully understand the requirements, checking understanding and testing models and solutions throughout the engagement
- Can effectively communicate the value, opportunities and limitations of AI technologies to a range of audiences with varying technical background.
- Uses the most appropriate medium to visualise AI based outputs to tell compelling and actionable stories relevant for business goals.
- Maintains a user focus to design AI solutions that meet user needs, taking account of ethical issues.
- Is familiar with the state of the art of techniques that help in modelling and understanding a business and its operation.
- As part of a team,
    - is able to support the scoping and business priority setting for large or complex changes caused by the adoption of AI, engaging senior stakeholders as required
    - is able to help identify the impact of adopting AI on business value and performance
    - uses the appropriate methods and techniques for the assessment and management of business risks that might result from adopting AI technologies

Ethical concerns for Artificial Intelligence Professionals

An AI practitioner is aware of and takes into account ethical concerns relating to the design, development, deployment, management and maintenance of AI products and services, such as for example

- Unfair or prejudiced bias in data or models
- Potential unconscious bias of AI practitioners and product development teams
- Appropriate level of transparency in design and development of AI models
- The impact of AI on restricting or enhancing user autonomy and wellbeing, whether in the workforce, in the customer base or society at large
- The ability of individuals to have appropriate control over their personal data
- Potential unintended, inappropriate or malicious use of AI products or services

The AI practitioner proactively works with and supports organisational stakeholders to develop appropriate policies, processes and practices to prevent unethical issues arising from the design, development, deployment, or management of AI products and services. They also proactively support their organisation to improve the ethnic and gender diversity and inclusivity of the workforce at all levels.

References

Much of the above draft framework has been adapted or synthesised from these sources:

- Digital, Data and Technology Profession Capability Framework[3]
    - Data scientist: skills they need
    - Data Engineer: skills they need
- Skills Framework for the Information Age[4]
    - Skill: Solution architecture - Level 5
    - Skill: Data management - Level 5
    - Skill: Business analysis - Level 5
    - Skill: Business modelling - Level 5
    - Skill: Requirements definition and management - Level 5
- Google's definition for Professional Data Engineer[5]
- The Life of a Data Engineer[6]
- Summary of the discussion from the July 10th Roundtable hosted at the Alan Turing Institute on behalf of the Office for AI.

2.2 Evaluation of Data Science related apprenticeships

This section looks at degree apprenticeships related to Data Science that are approved for delivery by the Institute for Apprenticeships and Technical Education (IfA). Such apprenticeships have significant overlaps with the high level skills framework in Section 2.1. At the time of writing there are two such apprenticeships:

- Digital And Technology Solution Specialist[7] (Degree Level 7, i.e. MSc), which includes the occupation specialism of Data analytics specialist
- Data Scientist[8] (Degree Level 6, i.e. BSc)

Appendix 7 includes the IfA standards for these. It also includes the detailed evaluation of how much these standards cover the skills in Section 2.1 and vice-versa.

[3] https://www.gov.uk/government/collections/digital-data-and-technology-profession-capability-framework
[4] https://www.sfia-online.org/en/framework/sfia-7/busskills/level-5
[5] https://cloud.google.com/certification/data-engineer
[6] https://www.mastersindatascience.org/careers/data-engineer/
[7] https://www.instituteforapprenticeships.org/apprenticeship-standards/digital-and-technology-solution-specialist-degree/
[8] https://www.instituteforapprenticeships.org/apprenticeship-standards/data-scientist-degree/

Case Study

We include a brief overview of the BSc degree apprenticeship in Digital and Technology Solutions at Queen Mary University of London[9] (QMUL) as a case study.

QMUL currently have apprenticeship employer partners including: ARUP, Bloomberg, Goldman Sachs, BBC, Broadridge, Experian, GSK, Global, Goji Investments, IBM, John Lewis, KPMG, Memiah, Sapphire, and Xantura. The four year degree apprenticeship course includes 24 modules, where six are taken each year, and leads to a degree apprenticeship specialising in software engineering.

In order to comply with IfA regulations, overall 20% of the degree apprenticeship is through taught modules on campus, while the remaining 80% is through work experience with the employer. At QMUL teaching occurs over two terms, each of twelve weeks, and during each term two modules are taught on campus and one at the place of work. To comply with IfA regulations and employer constraints QMUL decided to teach students two full days per week on campus during each teaching term, while the other three days were spent with the employer.

QMUL have been highly innovative in their timetabling and how they have physically provided space in order to deliver the necessarily intensive teaching and lab provision in only 40% of the usual time available. This raised particular challenges, since QMUL decided that apprentices should study alongside undergraduates taking the same modules on standard BSc programmes. This allows QMUL to have a sustainable business model since the viability of apprenticeship modules is not significantly affected by variability in the take up of apprenticeships, since those modules are likely to be viable due to the number of students on standard degree courses who also take them. Plus which, it allows for some limited flexibility for students to transfer between apprenticeships and standard degrees should they choose to. QMUL have also been innovative in their entry requirements, by taking into account previous work experience alongside academic qualifications.

The case study shows there are a range of practical challenges for a university to deal with when introducing a degree apprenticeship. These include timetabling, finding the physical space to deliver adequate teaching and lab facilities within a reduced working week, broadening entry requirements to attract a wider range of students, and sustainable integration of the apprenticeship within existing degree courses. These are apart from the challenges of designing the curriculum, navigating university processes for approval of a new course and attracting employers to partner with them.

Universities and employers will need to work closely together if they want to overcome such challenges and provide worthwhile MSc degree apprenticeships in Machine Learning and AI, but the QMUL case study shows this ought to be possible.

[9] http://eecs.qmul.ac.uk/undergraduate/degree-apprenticeships/

Our evaluation of the Data analytics specialist standard is given in full in Section 7.1. It suggests that it is significantly covered by the draft AI Skills Framework, whilst the Framework also goes beyond the standard in a significant number of places, particularly in the area of ethics. The evaluation of the Data Scientist apprenticeship is given in Section 7.2, which shows it covers a significant amount of the draft AI Skills Framework, but the extent of the knowledge and understanding in the apprenticeship is at level 6, whereas the Skills Framework is meant to be part of a level 7 offering.

It would be worth investigating the establishment of an Apprenticeship Trailblazer group to consider whether an AI occupational specialism can either be created as a standalone Level 7 apprenticeship beyond the level 6 Data Scientist, or an additional AI occupational specialism could be added to an existing standard such as the Digital And Technology Solution Specialist MSc apprenticeship.

3 University Interest, Capacity and Capability

This section looks at the university interest, capacity and capability to teach MSc programmes in areas closely related to Artificial Intelligence and especially Machine Learning that were on offer in the academic year 2018/19. It is based on a representative subset[10] of the existing university MSc course syllabi that are closely related to Machine Learning and Artificial Intelligence across universities in the UK (see Section 9 for the full list).

The key messages on university interest, capacity and capability are:

- There is a significant depth of knowledge in Machine Learning that is commonly taught across many of the university MSc courses we looked at, and which closely matches internationally recognised areas of knowledge in Machine Learning (see Section 3.1 and 3.2).
- There is significant breadth of MSc courses across the UK relevant to Artificial Intelligence and Machine Learning techniques (see Section 9 for illustrative summaries of relevant parts of course syllabi).
- Based on current MSc teaching capacity (see Appendix 8 for details) and past experience of university adaptability it is likely universities will be able to expand to the stated aspiration of an additional 3,000 MSc places over the long term, and will have no difficulty adding a few hundred extra paces in the short term.
- Based on expressions of interest, there are at least fifty universities wishing to be further consulted on future involvement with an industry funded MSc programme.
- Developing extensive work-related experience components of MSc courses to a sufficient capacity to meet employer needs is likely to be challenging (see Section 6.1 on Work Experience Current Capacity in the Delivery Model section for further details). Section 4.2.3 gives details of how universities are developing innovative teaching practices that may help to address this issue.

[10] Which includes details of MSc courses that were provided through consultations with, amongst others, CPHC and UKCRC. It is not exhaustive, and is not definitive.

3.1 Machine Learning canonical problems and approaches to addressing them

There is no current consensus on a core curriculum for MSc covering Machine Learning. For this report we chose to look at how well MSc graduates are likely to be equipped with the knowledge, understanding and skills needed to address the canonical problems summarised in the 2017 Royal Society report on Machine Learning[11], which are:

- Classification - To which category does this data point belong?
- Regression - Given this input from a dataset, what is the likely value of a particular quantity?
- Clustering - Which data points are similar to each other?
- Dimensionality reduction - What are the most significant features of this data and how can these be summarised?
- Semi-supervised learning - How can labelled and unlabelled data be combined?
- Reinforcement learning - What actions will most effectively achieve a desired endpoint?

In other worlds when looking at an MSc curriculum, we used the above as criteria for whether Machine Learning is a major component of an MSc course and to understand the breadth of the course. We used this criteria to look at the MSc courses from the University of Aberdeen, University of Bath, University of Bristol, University of Cambridge, City University of London, Coventry University, De Montfort University, University of Edinburgh, University of Essex, Heriot-Watt University, University of Glasgow, University of Hertfordshire, Imperial College, King's College London, University of Leeds, University of Manchester, University of Portsmouth, Queens University Belfast, Queen Mary, Royal Holloway, University of Southampton, University of St Andrews, University of Surrey, University of Sussex, University of Swansea, UCL, Ulster University, and the University of Warwick.

Some responses to our HEI consultations commented that we did not include in our criteria some of the more advanced Machine Learning approaches that are applicable to these canonical problems, such as for example adversarial machine learning techniques. Some responses questioned whether these canonical problems were better addressed at undergraduate level rather than at MSc.

Section 3.2 shows there are exemplary MSc curriculum that cover these canonical problems, including an extensive range of advanced topics that are at MSc level and not appropriate for undergraduate level. So, while it is true these canonical problems can be covered to some extent in an undergraduate course (and are), they can also provide useful insight into the depth and breadth of MSc courses if other contextual information is taken into account (which we have attempted to do). Generally, the tone of responses was that these canonical problems are useful as far as the remit of this report is concerned, but should not be seen as

[11] https://royalsociety.org/~/media/policy/projects/machine-learning/publications/machine-learning-report.pdf

endorsed by academia as a definitive characterisation of Machine Learning and AI MSc courses.

3.2 An emerging common body of knowledge for Machine Learning taught in MSC courses

There is a significant depth and breadth of knowledge in Machine Learning that is commonly taught across several of the universities we looked at, and which closely matches internationally recognised areas of knowledge in Machine Learning. What's more this would equip graduates with the appropriate knowledge, understanding and skills to address the canonical problems of Machine Learning identified by the Royal Society at MSc standard.

Every university has its own unique style of describing course structures and distinct way of describing module curricula. What's more they do not share exactly the same taxonomy of knowledge in either Artificial Intelligence or Machine Learning. In order to identify what is common to a significant number of universities and the breadth and depth of courses available we mapped some exemplar curricula against two standard texts in Machine Learning (see Appendix 10 for details of the texts, and for details of consultation responses to the suitability of these texts).

We mapped curricula against the subject knowledge described in those texts (see Table 7 for details) and used the terminology of those texts as a standardised framework in an attempt to give as unbiased and transparent opinion as possible, which is also independent of the universities we looked at. There are some standardised classifications of Artificial Intelligence, such as for example the ACM subject classification for computing[12], however we found that the curricula we considered had a better match against the terminology used in the two standard text books we reference.

Based on exemplary curricula descriptions from the universities of Bath, Edinburgh, Exeter and UCL there is clear evidence that the following subject areas are strongly covered in depth and breadth in a variety of MSc courses currently on offer[13]:

- Supervised Learning
- Bayesian Decision Theory
- Bayesian Estimation
- Parametric Methods
- Dimensionality Reduction
- Clustering
- Nonparametric Methods
- Decision Trees
- Neural Networks
- Kernel Machines
- Hidden Markov Models
- Combining Multiple Learners
- Reinforcement Learning

Whilst these are not meant to be a characterisation of the appropriate body of knowledge for an MSc course on Machine Learning, it does at least show there is a significant overlap of

[12] https://www.acm.org/publications/class-2012

[13] These curricular were used because of their depth and breadth of explanation, and should not be regarded as an endorsement of these courses compared to other university courses.

many such courses and that they include topics at an MSc level that address the Royal Society canonical problems.

4 Student Needs

The key messages in this section on student needs are that:

- Improving diversity and inclusivity in university computing departments should be a priority that is recognised as important for the wellbeing, academic success and professionalism of all students. There is a growing body of evidence that suggests these issues can be resolved, provided there is concerted and sustained effort focused on evidence based solutions looking at recruitment, retention and career progression within the academic culture and learning environment of university departments.
- Work-related experience, with strong employer engagement, are extremely effective at ensuring students develop to the fullest extent both academically and professionally.
- University computing departments are developing a wide range of teaching practices that develop students' professional skills, these need to be shared widely and evaluated against industry recognised professional skills frameworks relevant to AI.
- AI MSc qualifications need to clearly demonstrate and articulate graduates have achieved the proper knowledge, understanding and skills for a professional career related to AI. Professional body accreditation should aim to facilitate this, in line with the recommendations of the Shadbolt reivew[28].

These messages are consistent with requirements laid out by the Office for Students in their TEF requirements[14] for universities and those articulated by students in a BCS survey conducted in late autumn 2018 (details of which are included later in the section).

4.1 A diverse and inclusive learning environment and academic culture

Diversity and inclusivity are given special prominence in this report, because the lack of a diverse workforce across IT in general (as shown by the BCS series of reports on diversity[15]) is severely limiting the potential for the UK to fully benefit from AI over the longer term and because of the strength of feeling that has been consistently apparent in consultations. It is also a matter of fundamental social justice that the whole population have the same opportunities to benefit from studying AI at Masters level if they choose. The latest BCS report on social mobility[16] demonstrates that the IT profession provides significant opportunities for people from disadvantaged backgrounds to achieve social mobility compared to their parents, provided they manage to get the right qualifications to enter the profession in the first place.

The following bullet points summarise the main messages in this section on diversity and inclusivity.

[14] https://www.officeforstudents.org.uk/for-students/what-universities-and-colleges-should-do-for-students/
[15] https://www.bcs.org/media/1653/diversity-report-2017.pdf
[16] https://www.bcs.org/media/1652/social-mobility-report.pdf

- The professional and academic development of all computing students is hindered by lack of diversity in university computing departments.
- Current international research suggests that the academic culture and learning environments in computing departments significantly contribute to lack of diversity.
- Case studies such as the Carnegie Mellon Computer Science Department[17] show that it is possible for computing departments to have a diverse and inclusive culture and environment, leading for example to women making up 48% of Carnegie Mellon computer science freshers in 2016.
- Research from the American Astronomical Society[18] proposes concrete initiatives in recruitment, retention and professional development that are likely to improve diversity and inclusivity, which are as applicable to computing as they are to astronomy.
- Improving diversity by creating a pipeline of students recruited from non-STEM degree courses onto AI MSc conversion courses requires innovative solutions, which ensure learning outcomes are attainable by all students. Research[19] shows graduates with no prior exposure to quantitative data science methods don't have sufficient understanding of core concepts to cope with machine learning syllabi of the type found in many AI MSc courses. In particular when encountering quantitative methods for the first time students need extensive support to overcome an '*initial state of [intellectual] confusion and disorientation*' that blocks further progression.
- The Q-Step programme[20] funded by ESRC and the Nuffield Foundation has demonstrated it is possible to embed quantitative methods in social science undergraduate degree programmes, hence developing a new pathway for students without a STEM background into future AI conversion Masters courses. This suggests it should be possible to increase the diversity of postgraduate AI cohorts by recruiting a more diverse range of students from social science with quantitative methods degree programmes, and other humanities degree programmes that include quantitative methods relevant to data science and engineering, such as for example Cognitive Psychology.

For the learning environment and academic culture to be fit for purpose they must provide equal opportunities and support to all students from whatever their social background, gender, or ethnicity. What's more recent work by the Royal Academy of Engineering[21] in 2018 strongly argues that if students are to develop the skills needed to achieve excellence in engineering, then they need to develop those skills within diverse and inclusive teams, fostered within an inclusive, diverse learning environment and culture. In other words,

---

[17] https://www.csd.cs.cmu.edu/news/women-are-almost-half-carnegie-mellons-incoming-computer-science-undergraduates
[18] https://baas.aas.org/community/final-report-of-the-2018-aas-task-force-on-diversity-and-inclusion-in-astronomy-graduate-education/
[19] http://www.nuffieldfoundation.org/sites/default/files/files/Effective%20teaching%20practice%20in%20Q-Step%20Centres_FINAL.pdf
[20] http://www.nuffieldfoundation.org/q-step
[21] https://www.raeng.org.uk/publications/reports/designing-inclusion-into-engineering-education

diverse and inclusive teams achieve more than non-diverse teams. It's not just a matter of social justice that academic culture should be diverse and inclusive, it is necessary for the academic and professional development of all students.

4.1.1 Current state of gender diversity

Currently the undergraduate, postgraduate and research communities in computing are not diverse, with women making up only 18% of undergraduate computer science students (see Figure 2). This seems particularly anomalous when compared with mathematical sciences where 37% of undergraduates are women.

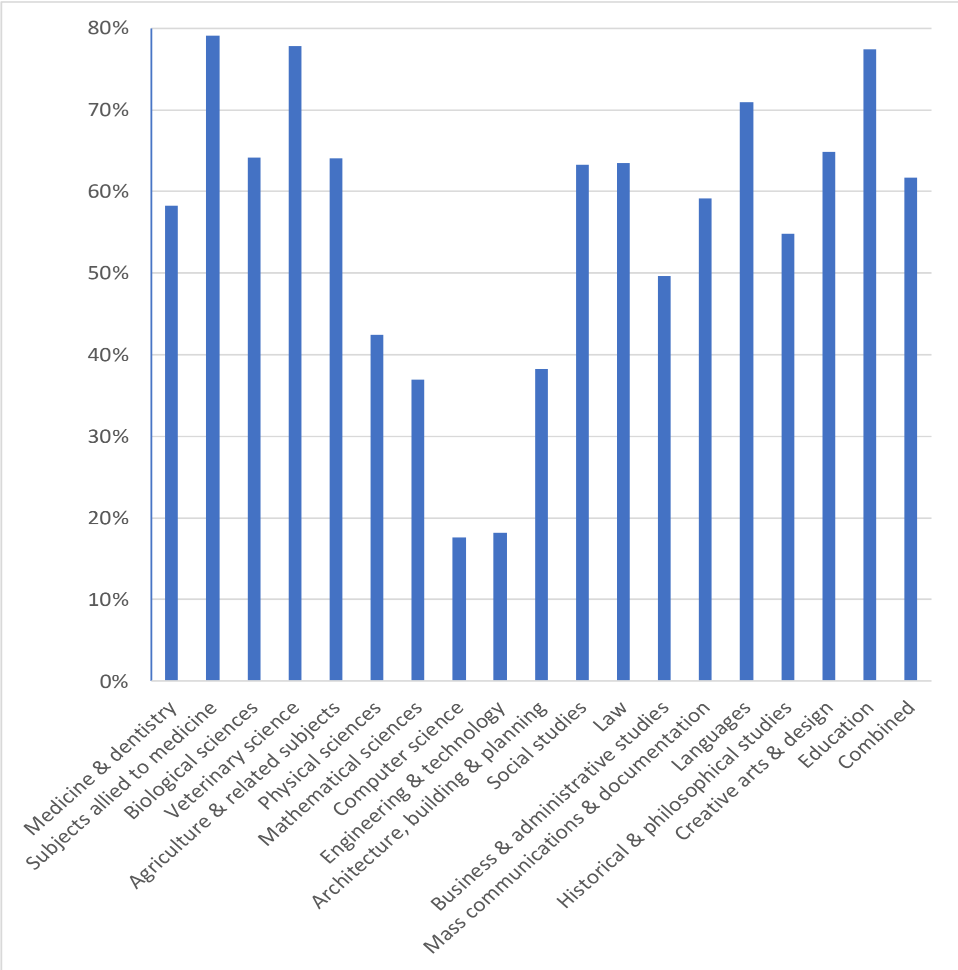

Figure 2: HESA data % of women undergraduates 2017

Research[22] commissioned by EPSRC, BCS, the IET, CPHC, UKCRC and TechWorks, reported that within the ICT academic research community:

"*Discrimination and harassment examples were widely reported in the online survey, while interviews revealed a more complex phenomenon of 'indirect' or 'unquantifiable' experiences of sexism and other prejudice. Women in the research also recounted various instances of their competence being questioned in ways their male colleagues do not experience.*"

[22] https://epsrc.ukri.org/newsevents/pubs/ediinictactionplan/

The report found that female postgraduates were more likely to say that they had felt excluded or inferior due to remarks made by their colleagues (54%, compared with 20% of male postgraduates).

The Athena Swan charter[23] allows universities and university departments to publicly commit to undertake work to address gender equality, and to be awarded Bronze, Silver or Gold as recognition of their systematic efforts to improve equality. Bronze recognises that "*in addition to institution-wide policies, a department is working to promote gender equality and to identify and address challenges particular to the department and discipline*". Silver recognises that "*In addition to the future planning required for Bronze department recognition, Silver department awards recognise that the department has taken action in response to previously identified challenges and can demonstrate the impact of the actions implemented*."

At present 43 computing departments have registered with the Athena Swan Charter, out of roughly 120 such departments in the UK, as shown in the following table:

Aberystwyth University	Bronze Department
Birkbeck College, University of London	Bronze Department
Brunel University London	Bronze Department
Cardiff University	Bronze Department
Durham University	Bronze Department
Durham University	Bronze Department (interim)
Edinburgh Napier University	Bronze Department
Heriot-Watt University	Bronze Department
Imperial College London	Bronze Department
Keele University	Bronze Department
Lancaster University	Bronze Department
Newcastle University	Bronze Department
Open University	Bronze Department
Queen Mary, University of London	Bronze Department
Queen's University Belfast	Silver Department
Royal Holloway, University of London	Bronze Department
Ulster University	Bronze Department
Ulster University	Bronze Department
University College London	Silver Department
University of Aberdeen	Bronze Department
University of Bath	Bronze Department

[23] https://www.ecu.ac.uk/equality-charters/athena-swan/athena-swan-members/

University of Birmingham	Bronze Department
University of Cambridge	Bronze Department
University of Dundee	Bronze Department
University of East Anglia	Bronze Department
University of Exeter	Silver Department
University of Glasgow	Bronze Department
University of Hertfordshire	Bronze Department
University of Kent	Bronze Department
University of Lincoln	Bronze Department
University of Liverpool	Bronze Department
University of Manchester	Bronze Department
University of Nottingham	Bronze Department
University of Oxford	Bronze Department
University of Plymouth	Bronze Department
University of Portsmouth	Bronze Department
University of Salford	Bronze Department
University of Sheffield	Silver Department
University of St Andrews	Bronze Department
University of Surrey	Bronze Department

Table 1: Athena Swan awards to Computing Departments, 2018

Given that universities can apply for Athena Swan awards separately from departments, the fact many computing departments have not applied for an award does not imply they are not supporting the Athena Swan initiative, since they can do so through the initiatives run at institutional level. At the same time, it seems clear the Athena Swan scheme is intended to encourage individual departments to sign up to the charter as well as universities, and in those terms the majority of computing departments have not so far chosen to sign up.

Despite the seemingly bleak picture shown by the data, many computing departments put significant effort into addressing diversity and inclusivity. What the data suggests is that these efforts are not making a significant difference on a national scale. Whilst some of the fundamental gender diversity issues begin in primary school and are endemic throughout society, that does not mean they cannot be addressed and improved in university computing departments as shown by evidence presented in the next section.

4.1.2 Academic culture and learning environment

The fact that the percentage of women gaining Engineering degrees varies massively across different nations shows that societal culture is a major factor in determining whether women go into Engineering. The 2016 report[24] from the Australian Government Chief Scientific Office noted that women account for:

- 40% of engineers in China
- 44% of the engineering graduates in Malaysia
- 58% of engineers in the former USSR

The Royal Academy of Engineering argue[25] that a fundamental change in academic culture and learning environment is needed to ensure underrepresented groups can thrive in university engineering faculties. The same arguments apply equally to computing.

Since 1999 the Computer Science department at Carnegie Mellon University in the US has made fundamental, systemic and sustained changes[26] to their academic culture and learning environment. These included:

- Change of admission policies to include excellence criteria that look for potential leadership qualities and include a range of more broad technical skills beyond those specifically related to computer programming.
- A formalised program of professional, networking, and mentoring opportunities for women.

Equally importantly the Carnegie Mellon initiative did not attempt to change the academic content of the curriculum by adding ‘female friendly’ content. As mentioned at the start of this section, sustained changes such as these resulted in the percentage of women in computing increasing from 10% in 1999 to 48% in 2016.

There are major caveats to bear in mind when considering this example. We should not assume successful strategies that work in a US cultural context can automatically translate to the UK. We should also be careful that the improvement in gender diversity in Carnegie Mellon does not necessarily imply that the total number of women choosing computer science as a degree has increased as a result of those changes; it is conceivable that a greater percentage of those women already intending to study computer science chose Carnegie Mellon because they perceived the department to be more inclusive.

What is clear from the Carnegie Mellon example, the Royal Academy of Engineering research and also from the American Astronomical Society[18] (AAS) is that changes to both the academic culture and learning environments are essential to improve diversity, and there is evidence that such changes if sustained over time can significantly improve matters. In particular AAS have a concrete range of recommendations covering recruitment, retention and professional development that could be adapted for use in computing departments in the UK. Note however, before doing so would require further study to ensure they would be likely to achieve the intended outcomes within a UK context without adversely affecting academic standards.

[24] https://www.chiefscientist.gov.au/wp-content/uploads/OCS-paper-13.pdf
[25] https://www.raeng.org.uk/publications/reports/designing-inclusion-into-engineering-education
[26] http://www.cs.cmu.edu/~lblum/PAPERS/CrossingCultures.pdf

The key points from the existing evidence is that material improvements to diversity and inclusivity can be achieved through changes to recruitment and retention practices:

- Recruitment policies and processes that are inclusive by ensuring entry requirements take account of qualities that measure whether candidates have the potential to grow into great scientists/engineers/practitioners as well as looking at past academic performance.
- Postgraduate programmes should partner with and recruit from undergraduate programmes that produce large numbers of graduates from underrepresented groups.
- Retention support for students is made inclusive by ensuring students from all backgrounds have equal access to a full range of professional development, networking, and mentoring opportunities.

4.1.3 Building diverse pipelines from non-STEM degree subjects into AI MSc conversion courses

The previous sections looked at the potential for improvements to academic culture and learning environments in computing departments. The issue this section considers is attracting a diverse cohort of students to MSc AI conversion courses from non-STEM undergraduate programmes. Developing such a diverse pipeline was proposed by the DCMS and BEIS commissioned independent review[1] into growing the artificial intelligence industry in the UK.

Given the potential impact of machine learning and AI on the social sciences, thanks to the incredible growth of social data now available online, it seems natural that undergraduates in social sciences should be encouraged to study AI at MSc level. There are of course other non-STEM degree programmes that could provide potentially diverse pipelines to AI MSc conversion courses. For example, Cognitive Psychology degrees include modules on statistical research methods that are relevant to AI MSc conversion courses. However, given the potential relevance of machine learning to social science, and the current lack of a pipeline of social science graduates into AI it seems worth focusing discussion on this subject discipline. Plus which, ways of attracting social science students onto AI MSc conversion courses are likely to be equally applicable across a wide range of non-STEM subjects.

In 2014 ESRC and the Nuffield Foundation funded a £19.5m programme, Q-Step[20], across fifteen universities to embed quantitative methods into social sciences undergraduate degree courses (the types of courses available are shown in the table below). The programme is ending in 2019, and as yet it is not clear if there will be a successor. By the end of this pilot phase approximately 1,000 undergraduates a year have started a Q-Step degree pathway. At the same time some Q-Step centres have reported difficulties in recruitment, which means it is important to understand issues that have been identified and their possible resolution where they are relevant to MSc conversion courses in Machine Learning and AI.

Current BA Honours degree courses available 'with' Quantitative Methods
Sociology with Quantitative Methods
Criminology and Sociology with Quantitative Methods
Criminology with Quantitative Methods
Law with Quantitative Research
Politics and International Relations with Quantitative Research
Sociology with Quantitative Research
Social Policy with Quantitative Research
Quantitative Social Sciences
Sociology and Quantitative Methods
Politics, International Studies and Quantitative Methods
Sociology with Quantitative Methods

Although evaluation of the programme has not been completed yet, there are some interesting findings emerging that have been published in a Nuffield report on teaching practice[27], which future AI MSc conversion courses will need to consider. Preliminary findings in the Nuffield report suggest social sciences graduates who have not studied quantitative methods since leaving school may experience significant learning challenges when they first encounter quantitative concepts and methods required for machine learning topics at MSc level, which take a significant time and support for them to overcome.

A key issue the report identified is that students must fully master a series of technical quantitative concepts that build incrementally one upon the other, so that miscomprehending one concept causes a cascade of miscomprehension of new concepts as their correct comprehension is achieved via the technical mastery of previous ones. This often leaves students in an '*initial state of confusion, disorientation and anxiety when learning quantitative methods until the initial superficial understanding of its core concepts comes to be properly integrated and mastered',* which blocks further progression. The Nuffield report suggests overcoming this initial block requires developing '*statistical imagination*' that equips students with the thinking skills needed to imagine '*patterns of probabilistic regularities described in terms of the empirical distribution of variables and their association, and their comparison with 'theoretical' distributions in order to draw inferences'*. Developing this statistical imagination in students takes time, and a clear focus on using the appropriate pedagogy that is designed specifically for non-STEM students.

The kinds of issues raised in the Nuffield report strongly suggests a one year MSc conversion course in Machine Learning and AI may not be attractive to a significant cohort of social science students if they don't have a suitable background in quantitative methods prior to the course, which could undermine the ambition set out in the government's AI review. The whole point of the Q-step programme is to develop a substantial cohort of social science graduates that do have a strong background in quantitative methods. Assuming the programme is at least a partial success, then the obvious conclusion is that the best way of

[27] https://www.nuffieldfoundation.org/sites/default/files/files/Effective%20teaching%20practice%20in%20Q-Step%20Centres_FINAL.pdf

creating a pipeline of social sciences graduates who would benefit from an AI MSc conversion course is to develop a successor to the Q-step programme, and ensure all universities are incentivised to introduce Social Sciences with Quantitative Methods as a degree option, with clear signposting and alignment to further study through AI MSc conversion courses.

4.2 Professional development for postgraduate students

The main messages in this section are:

- When asked about their professional development needs UK students responding to a BCS survey gave scores of more than 4 out of a possible 5 to (see Section 4.2.1).
    - Certification of capability and competence
    - Insight and information about the latest industrial trends
    - Support for career development
    - Ways to externally evidence commitment to professional standards
    - Networking opportunities
- In-depth employer engagement, such as through an extensive placement[28] or through a degree apprenticeship[29], can be highly successful at developing students professional and academic skills and which are aligned with student's own stated needs.
- One of the key determinants for high-quality employer engagement is ensuring they evidence progression based on established professional standards[29, 34] (such as for example the Skills Framework for the Information Age[30], the UK Standard for Professional Engineering Competence[31], or the Digital, Data and Technology Profession Capability Framework[32]).
- Over recent years many university computing departments have developed exemplary practice that supports student professional development[33], and hence improve graduate employability.
- Assurance and validation for AI MSc courses should build on the BCS university academic accreditation requirements[34] that
    - Students should not perceive Legal, Social, Ethical and Professional Issues as peripheral to, or less significant than, technical skills detailed in the syllabus
    - Awareness of professional standards, codes of conduct and relevant legislation must not be separated from the practice of designing and implementing systems
    - Legal, Social, Ethical and Professional Issues should be addressed within core areas of degree programmes rather than in options alone

---

28 https://www.gov.uk/government/publications/computer-science-degree-accreditation-and-graduate-employability-shadbolt-review
29 http://epc.ac.uk/wp-content/uploads/2012/08/Designing-apprenticeships-for-success.pdf
30 https://www.sfia-online.org/en
31 https://www.engc.org.uk/ukspec
32 https://www.gov.uk/government/collections/digital-data-and-technology-profession-capability-framework
33 https://cphc.ac.uk/2017/08/03/gecco-evaluation-report/
34 https://www.bcs.org/media/1209/accreditation-guidelines.pdf

4.2.1 Students stated professional development needs

The following table (Table 2) shows the result of a survey of over four thousand BCS members. The table shows their response to one of the questions in the survey, which asked them to give a score between 0 (of least value to them) and 5 (of greatest value to them) to different options for developing as professionals. To provide more detailed context, the table breaks down the responses to the question as different BCS membership grades[35]. The different options are ranked according to the scores given by UK resident students[36]. Only those items that students scored above 4 out of 5 are listed in the table, as these represent their top priorities.

Priority	Overall	Fellows	MBCS CITP	MBCS	AMBCS	Affiliate	Student UK	Student Non-UK
Certification of my capability and competence	4.03	3.95	4.08	3.98	4.35	3.65	4.22	4.25
Insight and information about what is happening in the industry	3.89	3.65	3.68	3.87	4.03	4.01	4.22	4.21
Support and / or contribution to my career development	3.39	2.81	2.93	3.31	3.98	3.53	4.15	4.37
Ways I can externally evidence I am committed to professional standards	3.88	3.91	3.91	3.87	3.97	3.67	4.11	3.79
Networking opportunities	3.35	3.26	3.1	3.27	3.53	3.51	4.08	4.08

Table 2: BCS survey of members 2018, which item is of most importance to your professional development

What this shows is students have a clear view of their professional development needs and they clearly see demonstrating their capabilities and competencies against recognised standards as important for their career development. AI MSc courses should ensure that students are able to clearly see how their stated priorities will be supported through successful completion of the MSc degree programme.

4.2.2 Employer engagement validated against professional standards

High quality work-related experience, such as through a properly structured and monitored placement, is highly effective at developing students professional and academic skills. The Shadbolt Review[28] found that:

- Of the combined 2011-12 to 2013-14 cohort of UK-domiciled full-time first-degree Computer Sciences graduates, 15% of graduates not on a sandwich course were unemployed six months after graduation, whereas only 6% of graduates on a degree including a sandwich year placement were unemployed six months after graduation.
- Of graduates in employment after six months, 25% of graduates not on a sandwich course were in a non-graduate role, compared with 6% of graduates on sandwich degrees.

[35] AMBCS = associate member BCS, MBCS = full professional member BCS, CITP = Chartered IT Professional MBCS

[36] Roughly three hundred UK resident students took part in the survey

Worthwhile work-related experience for students should support outcomes that are clearly articulated and based on appropriate professional development needs of students. As we have seen in the previous section student's want to evidence their professional development against industry recognised standards.

4.2.3 University teaching practice supporting professional development

Many university computing departments have been developing innovative teaching practice over recent years that supports student professional development, with an emphasis on improving graduate employability. During 2016 CPHC consulted[37] with fifty or so university computing departments to collect exemplary good practice that enhances employability of graduates and development of professional skills. The types of practice reported by CPHC strongly resonate with the earlier findings from the Royal Academy of Engineering in their 2010 report[38] on how universities can enhance engineering graduates' professional skills.

Of the different practices that are being adopted by various universities, examples that are of particular relevance to industry funded AI MSc programmes include:

- Providing employer engagement through work experience is widely seen as one of the most effective ways for students to develop professional skills. Many regard industrial placements running over at least six months as the best way of providing work experience, but there are alternatives that can also be effective, which may be more scalable. For example, employer defined real-world projects[39], where the employer also acts as customer and provides feedback that counts as part of the assessment, can be highly effective as a means of developing students' professional skills.
- A number of universities formally assess industrial placements as a separate module that counts towards the final degree classification. Assessment that took account of regular written feedback from student's in-line manager during the placement helped students consciously think about and articulate the professional skills they were gaining. For the student to best gain meaningful professional development during the placement it's recommended that the student, their course tutor and employer decide placement learning outcomes prior to starting the placement.
- Various universities have found it important to clearly articulate and promote the type and range of professional skills that work placements develop to students from the start of their course. This helps students see the value and relevance of those skills to their personal development as future practitioners. It was also found valuable to explicitly illustrate how those skills are important within a variety of the core academic modules. This can be aided by showing how industry recognised standards of professional practice are relevant to different modules.
- Some universities have taken the approach that professionalism is embedded as an explicit part of all core modules, and forms part of the assessment.
- Students have a better appreciation of the need to develop professional skills when there is clear signposting of career progression opportunities that those skills are

[37] https://cphcuk.files.wordpress.com/2016/01/computinggraduateemployabilitysharingpractice.pdf
[38] https://www.raeng.org.uk/publications/reports/engineering-graduates-for-industry-report
[39] Contributing 60 credits towards the overall assessment

relevant to. For this to be most effective requires signposting to be included across academic modules as well as through separate resources that might be made available online.

- Some universities support students to develop entrepreneurship through on campus initiatives that allow students to create their own start-up companies. Their success at developing business plans, managing product life-cycles, and marketing strategies for the start-up are included within assessment of the final degree. Direct support from employers to help students develop the relevant skills is especially helpful.
- Providing valuable work experience requires academic staff in departments to be sufficiently aware of the professional skills students should be developing through that experience. This can be achieved through academics showcasing to students their own professional development within industry recognised skills frameworks.
- Some universities choose to teach an employer-led curriculum, with significant employer involvement ensuring relevance and currency and including teaching professional skills to interested students through extracurricular activity. The intention being that partnering with employers to inform curricula will offer a more authentic experience of business innovation.
- Incorporating focused, large-scale projects into the curriculum is another mechanism for teaching professional skills to students. Ideally these have an industry client who sets the scope of the project, and the project will require the application of ideas, concepts and techniques from a range of different modules. One innovative approach is to use ‘live projects’. This involves companies giving a live brief to students on campus. Students work on projects in teams, and they receive explicit professional development of skills relevant to their project. The value of the project to students is they are working on something that is genuinely useful to a company and requires them to gain experience of working with a client. Whenever possible, companies provide staff to help with final assessment of the project, as well as give feedback at regular intervals on progress during the project.
- A combination of employer-led curricula, with three month long industrial internships is an approach adopted by some universities. A structured three month internship that builds on professional development during the on-campus part of a course gives students real exposure to working as part of a team that have to deliver against business priorities for a client.
- University practice with regards professional development for students is evolving, which means it is important to systematically gather objective data on student experiences and share lessons learnt with the wider academic community.

The above examples of teaching practice aimed at developing professional development are well worth being considered by universities when developing work experience related opportunities to include in AI MSc programmes.

5 Ethical Behaviours

The main messages for this section are:

- AI MSc graduates should be able to evidence they have the ethical competencies necessary to significantly contribute towards embedding ethical principles throughout the AI product lifecycle, including aspects such as designing, developing, deploying, maintaining and managing AI systems.
- Ethical practices should be embedded as a core part of the curriculum including being a core component of technical lab exercises, assessed projects and integral to core technical subject matter modules.
- Assessment of AI ethical competencies should be a significant component of the overall award of the degree.
- The Centre for Data Ethics and Innovation is currently working with BCS to capture ethical good practice for AI systems engineering, which is intended to be useful for the development of degree curricular.
- Industry funded MSc programmes should have appropriate assurance and validation mechanisms in place to give confidence to employers that graduates are equipped with the appropriate ethical competencies, which could be done for example through accreditation from a professional body.

Developing rigorous, applicable, practical standards for ethical AI that are embedded at the heart of professional AI practice is widely regarded as a priority by employers, academics, government and professional bodies. Table 3 shows the result of a BCS survey of its members asking how important they thought standards of ethical practice are to the IT industry.

Importance of ethical practice	Overall	Fellows	MBCS CITP	MBCS	AMBCS	Affiliate	Student UK	Student Non-UK
Very important	84%	86%	86%	84%	83%	88%	80%	80%
Quite important	13%	9%	11%	13%	12%	10%	12%	19%
Not very important	2%	3%	2%	2%	3%	2%	3%	1%
No opinion	1%	1%	1%	1%	2%	0%	0%	0%
Not at all important	0%	1%	0%	0%	1%	0%	5%	0%

Table 3: BCS Survey asking how important members think standards of ethical practice are to the IT industry

The survey had responses from over four thousand members (details of their level of seniority within the profession are included in the table) and showed that overall 84% said ethical practice is ‘very important’. Interestingly 80% of UK students responding gave a ranking of ‘very important’ compared to 86% for Chartered IT Professionals and Fellows. This may be because as someone reaches a more responsible role, they gain a deeper appreciation of the impact IT systems have on society and the need for ethical practice.

Current State of Ethical Standard	Overall	Fellows	MBCS CITP	MBCS	AMBCS	Affiliate	Student UK	Student Non-UK
Neither high nor low	36%	35%	38%	38%	33%	33%	25%	33%
Quite high	28%	24%	26%	28%	34%	20%	29%	25%
Quite low	20%	28%	22%	20%	16%	20%	17%	15%
Don't know	8%	6%	7%	7%	10%	16%	12%	9%
Very high	4%	1%	3%	3%	5%	8%	12%	14%
Very low	4%	5%	4%	4%	1%	3%	5%	4%
Sum of not high categories	**60%**	**68%**	**64%**	**62%**	**50%**	**56%**	**47%**	**52%**

Table 4: BCS survey asking how members rate the current general standard of ethical practice in the IT industry

The BCS survey also asked members for their perception of the current state of ethical standards in the IT industry, which is shown in Table 4. The bottom line of the table shows the aggregation of the scores for the 'Neither high nor low', 'Quite low' and 'Very Low' categories. Collectively those three categories are termed 'not high'. What this table shows is that 47% of UK students chose 'not high' categories, compared to over 64% of Chartered IT Professionals and Fellows who chose 'not high' categories. So, whilst more senior members of the profession rated ethical practice as more important than students, they also are more likely to perceive that ethical practices are of a 'not high' general standard across the IT industry. That may not be surprising given that people working in the profession have a much more intimate knowledge of day to day practice, whereas students are more removed. Combining the two survey questions it suggests the profession as a whole believe ethics to be of paramount importance, whilst also believing there is room for improvement in day to day practices. This lends weight to the principle that ethics must be integral to AI MSc courses, given the potential for AI to become ubiquitous across all sectors of the economy.

Currently there is much focus from many stakeholders on identifying the right overarching principles of ethical AI. For example, the European Union has just published[40] ethics Guidelines for Trustworthy AI, and the government of Singapore has just published[41] a Model AI Governance Framework for consultation. In the UK the Alan Turing Institute[42] is working with the broader data science community to generate public dialogue about the principles we want to adopt in the UK, and the Centre[43] for Data Ethics and Innovation has been established to advise government on the measures needed to strengthen and improve the way data and AI are used. The Centre is currently working with BCS on plans to capture ethical good practice for AI systems engineering. There is less focus on developing standards that provide solid guidance on how to turn principles into professional practice, which as shown by the BCS survey is something that is very important for the future of the profession. The ideal outcome will be that in future the adoption of ethical standards of

[40] https://ec.europa.eu/futurium/en/ai-alliance-consultation/guidelines#Top
[41] https://www.pdpc.gov.sg/Resources/Model-AI-Gov
[42] https://www.turing.ac.uk/research/data-ethics
[43] https://www.gov.uk/government/consultations/consultation-on-the-centre-for-data-ethics-and-innovation/centre-for-data-ethics-and-innovation-consultation

practice go a long way to ensuring all AI systems are 'ethical by design', analogous to the way we as a nation are working to develop standards that ensure all new online technologies are 'secure by design'.

The fact standards for ethical practice are still being developed means it is especially important that MSc students gain a comprehensive understanding of how ethical concerns could arise in all the stages of a product's lifecycle. They also need the engineering skills necessary to incorporate technological solutions that mitigate against a wide range of ethical concerns that may arise during the design, development, deployment, maintenance and management of AI systems. For example, they may need to develop automated auditing features across a data pipeline to enable compliance with future good practice around data cleansing, add suitable exception handling determined by ethical escalation policies to guarantee fairness of an AI based decision making process within a workflow, or design and implement APIs that capture statistical characteristics of AI models for monitoring potential unethical bias. Those examples are intended to show that engineering choices can profoundly affect the technical capability of a system to monitor, audit, manage and adapt to future ethical concerns. That implies ethics must be taught as a core competency within an AI MSc programme, and assessing ethical competency of graduates is vitally important for developing a pipeline of both talented and ethical AI practitioners.

The BCS accreditation guidelines[44] for university degree courses states that:

> *The relevant legal, social, ethical and professional issues (LSEPIs) should be specifically detailed in the syllabus, mentioned in directions to students on practical assignments and sandwich placements, and not left solely to the discretion of individual lecturers. Whilst LSEPIs should pervade the programme, the central issues of codes of conduct and practice, legislation and ethical standards are important to all information systems engineering practitioners. Therefore, they should be addressed within core areas of the programme rather than in options alone.*

Any industry funded AI MSc programme claiming to equip graduates with ethical competencies should have appropriate assurance and validation mechanisms in place that evidence to employers the programme complies with the above accreditation criterion. They should consider whether professional body accreditation is a useful vehicle for providing that.

[44] https://www.bcs.org/media/1209/accreditation-guidelines.pdf

6 Delivery Models

The key messages for possible delivery models of an industry funded MSc in AI are:

- Extensive work experience evidenced against professional standards is employers' preferred model of ensuring graduates are equipped with appropriate professional competencies. Providing work experience at significant national scale will be challenging over the medium term, both for university departments to create the infrastructure necessary to support significant work experience at scale, and for employers to create sufficiently many worthwhile opportunities.
- Delivery models for industry funded MSc courses will need to be flexible enough to support HEIs adapt to new employer requirements, and to ensure scalability and sustainability of the programme over the long term.
- Given the national aspiration to eventually create three thousand industry funded AI MSc places[45] it will be necessary to create a suitable infrastructure to broker connections between employers, universities and students. This should be designed to help employers appreciate which university courses best suits their needs, help universities understand how best to provide courses employers and students value, and help students find the right employer and university for their needs.
- Most large employers strongly expressed a preference for industry funded AI MSc courses to be delivered through level 7 degree apprenticeships funded through the apprenticeship levy[46]. At the same time many large employers also support providing some MSc places through student scholarships that include other forms of work experience, which provides greater flexibility for students, employers and universities and allows for more agile innovation in curricula.
- SMEs consulted expressed a strong desire to support industry funded AI MSc programmes, but do not tend to have capacity to provide long placements or provide level-7 apprenticeships. They have supported MSc students through scholarships, as well as directly supporting teaching and supervision of students, and are keen to continue to do so.
- Providing professional recognition of MSc graduate competencies by Chartered professional bodies is currently being explored by some universities as a possible means of independently evidencing competencies against professional standards recognised by employers.

6.1 Work experience current capacity

The long term aspiration outlined in the government's independent AI review[45], based on employer anticipated future demand, is the creation of three thousand industry funded places on Machine Learning and AI MSc courses. As explained in Section 2 increasing the supply of AI MSc graduates through level 7 degree apprenticeships has the strong support of employers, but employers are also keen to support scholarships where that enables

[45] https://www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in-the-uk
[46] https://www.gov.uk/government/publications/apprenticeship-levy-how-it-will-work/apprenticeship-levy-how-it-will-work

universities to rapidly develop innovative new courses and provides greater flexibility to meet disparate needs of students.

The Office for Students 2019 insight brief on degree apprenticeships[47] gave an overview of the current state of degree apprenticeships. Table 5 below taken from the insight report shows the number of level 6 and level 7 apprenticeships across all standards for 2017/18.

Framework or standard	Level	Numbers
Chartered manager	6 (degree)	2315
Digital and technology solutions professional	6 (degree)	1310
Chartered surveyor	6 (degree)	815
Registered nurse	6 (degree)	305
Civil engineer	6 (degree)	160
Healthcare science practitioner	6 (degree)	110
Manufacturing engineer	6 (degree)	105
Product design and development engineer	6 (degree)	100
Embedded electronic systems design and development engineer	6 (degree)	95
Aerospace engineer	6 (degree)	85
Nuclear scientist and nuclear engineer	6 (degree)	80
Food industry technical professional	6 (degree)	75
Building services design engineer	6 (degree)	65
Electrical or electronic technical support engineer	6 (degree)	55
Control or technical support engineer	6 (degree)	25
Laboratory scientist	6 (degree)	25
Construction management	6 (degree)	20
Aerospace software development engineer	6 (degree)	10
Broadcast technology higher apprenticeship (BBC)	6 (degree)	10
Non-destructive testing engineer	6 (degree)	5
Chartered legal executive	6	185
Senior insurance professional	6	120
Relationship manager (banking)	6	105
Senior compliance or risk specialist	6	85
Financial services professional	6	55
Licensed conveyancer	6	35
Teacher	6	15
Senior leader	7 (degree)	550
Accountancy or taxation professional	7	3710
Solicitor	7	105
Postgraduate engineer	7	95
Systems engineering	7	30
Academic professional	7	5

Table 5: Level 6 and 7 apprentices by subject in 2017-18

The Digital and technology solutions professional apprenticeship is the second most numerous in the table, note that the Data Scientist apprenticeship does not appear at all (see Section 2.2 for more details of these apprenticeships). These are the two currently

[47] https://www.officeforstudents.org.uk/publications/degree-apprenticeships-a-viable-alternative/

approved apprenticeships that are most relevant to the high level AI skills framework employers have articulated (Section 2.1). These numbers raise questions about how quickly a new AI MSc Apprenticeship might reach three thousand a year if left to expand organically.

Subject	Doctorate research	Masters taught
(1) Medicine & dentistry	2,180	3,935
(2) Subjects allied to medicine	1,645	9,760
(3) Biological sciences	3,420	12,725
(4) Veterinary science	70	75
(5) Agriculture & related subjects	230	1,195
(6) Physical sciences	3,155	5,075
(7) Mathematical sciences	760	2,535
(8) Computer science	1,080	6,760
(9) Engineering & technology	3,465	13,975
(A) Architecture, building & planning	405	5,820
(B) Social studies	2,090	20,225
(C) Law	500	8,200
(D) Business & administrative studies	1,310	54,315
(E) Mass communications & documentation	230	6,445
(F) Languages	1,260	6,275
(G) Historical & philosophical studies	1,470	5,595
(H) Creative arts & design	670	11,990
(I) Education	905	8,565

Table 6: HE qualifiers by subject of study and level of qualification obtained 2017/18

Table 6 shows the numbers of students who graduated from a taught Masters programme in the UK for the period 2017/18 (based on HESA data[48]). Note it excludes masters delivered by research only as those are outside the scope of this report. The table shows there were slightly more than 6.5k taught Masters students graduating in 2018 across all MSc courses covering subject specialisms within Computer Science. The Appendix in Section 8 breaks down the total across universities that offer taught Computer Science MSc courses. The Appendix cross references those departments against ones offering undergraduate M-level courses in Computer Science (e.g. MComp, MEng, MSci) that include a one-year sandwich placement. There are around fifty university departments where they offer one-year sandwich placements and also teach MSc Computer Science courses. This gives some indication of how many departments already have the capacity and infrastructure for offering MSc courses with a placement of a year.

Table 5 together with Table 6 indicate that adding another three thousand places onto AI MSc courses that have an extensive work experience component, either through degree apprenticeships or through work placements, is unlikely to happen over the medium term purely by organic means. At the same time the Office for Students insight report noted that many more universities are now preparing to offer degree apprenticeships, which suggests

[48] https://www.hesa.ac.uk/data-and-analysis/students/table-17

there will be potential capacity in a few years' time that could significantly contribute to the stated aspiration.

6.2 Evidence based work experience

Employers have repeatedly stated in all the consultations that they want MSc graduates to have both academic depth of expertise, and a broad range of applicable professional skills. Figure 3 shows a high level view of a possible AI MSc framework, which was developed at the first roundtable held at the Alan Turing Institute with input and support from industry and academia. Section 2.1 gives more details of the engineering and professional skills alluded to in the illustration. The framework identifies that universities should lead on academic knowledge and understanding, industry should lead on professional skills, and universities working with industry should identify engineering practices for translating AI expertise into industrial scale products and services.

Employers have also been clear they would like students to develop professional skills through work experience, which ideally would be through either an apprenticeship or an extensive placement, but they accept due to local circumstances it may be necessary to provide work-related experience through campus based initiatives. Section 4.2.3 highlights various innovative teaching practices universities have created to help students develop professional skills, which are likely to be relevant to AI MSc programmes. Many of them provide work experience related opportunities on campus that develop a range of professional skills equivalent to those students could gain through work placements. Such skills students develop will enhance their employability, but at the same time they cannot equip a student for every possible role employers may need.

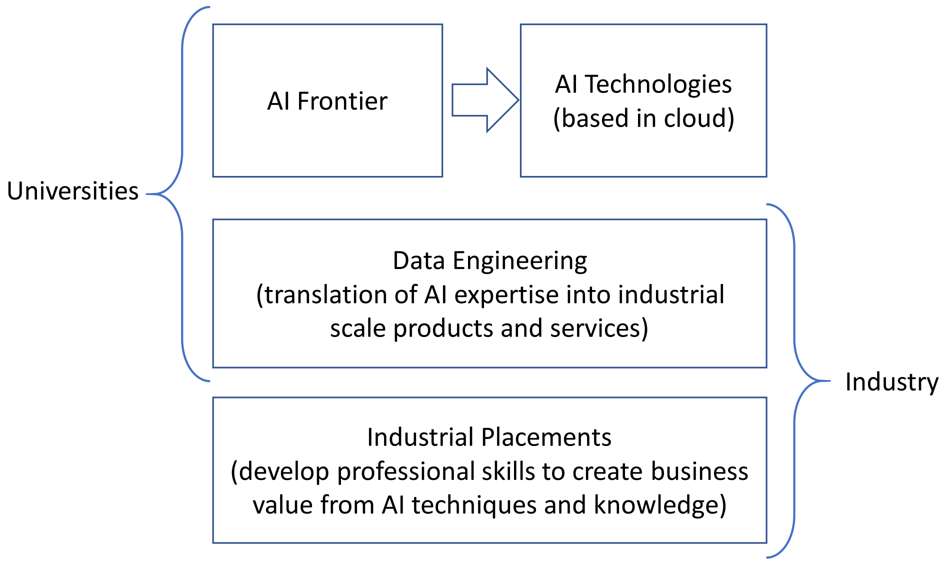

Figure 3: high level AI MSc framework discussed at first Alan Turing Institute roundtable

Employers have historically found it difficult to decipher which skill sets students have developed on a degree, which has often led to employers prioritising recruitment from universities they have an established relationship with. That is a perfectly sensible approach for an individual employer to take in ensuring their own needs are best met, but such an approach is not scalable. It makes it difficult for an employer to rapidly find new graduates with a particular skill set, since it takes time to establish a relationship of trust with a different university department that might be able to meet demand. The approach of bilateral trust-based relationships will also make it difficult for computing departments to rapidly scale or establish new industry funded Masters programmes, since they will have to spend time developing trusted relationships with potential industrial funders before they can expand or establish such programmes.

Employers have been clear throughout the consultations they want work experience related learning to be assessed as part of the overall award of the MSc, and to be evidenced against industry recognised standards. Some universities already do this for some of their undergraduate computing degree courses, based on standards derived from SFIA for example. Universities and employers working together to develop academic and professional standards that are embedded in Masters programmes is a model that allows programmes to be more quickly established and scaled, provided the programmes provide adequate assurance and validation that they do meet the standards and that graduates of those programmes are equipped with the skills defined in those standards. Section 6.3, which follow this one, looks at assurance and validation.

6.3 Assurance and validation through programme accreditation

Accreditation works well when it assures and validates professional and academic good practice within standards that are led by employers and academics. Accreditation does not preclude universities from developing courses that are not accredited, it simply provides an independent mechanism for assuring third parties that publicly available standards have been met. It is important to emphasise that accreditation standards should be led by and promoted by employers collaborating with academics, and that they articulate current good practice developed by the profession itself. That means accreditation can be a vehicle for advancing academic learning and professional development that ultimately benefits society through a profession that can be trusted to work towards ever higher levels of competency and ethical behaviour.

The high level logic model in Figure 4 illustrates how assurance and validation though employer and university led accreditation can support the continuous improvement of both academic and professional standards in education (as well as the IT profession).

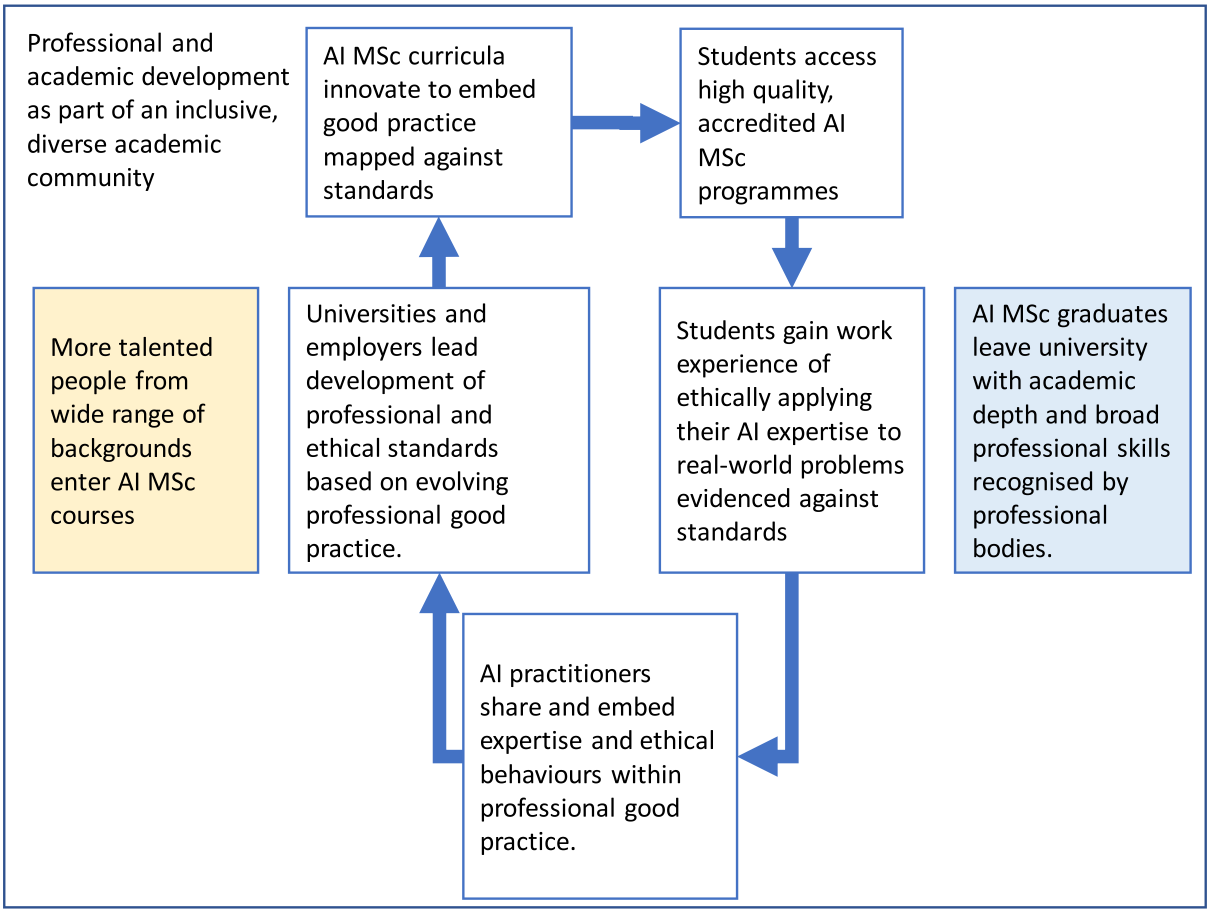

Figure 4: high level logic model of academic and professional development of AI MSc students

One of the purposes of a professional body is to facilitate collaboration between academia and employers to develop professional standards, and then provide a means of accrediting qualifications and certifications that embody those standards in order to establish them at a national level, hence achieving scale. Professional body accreditation should have the overarching aspiration of helping develop and spread good practice that supports a virtuous cycle where employers and academics work together to continuously advance their discipline and profession through the evolution and implementation of standards, as shown in Figure 4. It should provide a mechanism whereby students and employers are assured of the benefit of degrees with respect to their professional needs.

Currently BCS as well as other professional bodies such as IET offer voluntary accreditation to university departments for taught degree courses, both undergraduate and at Masters level. Currently there are over 120 HEIs with BCS accredited degrees[49] from 1999 to the present, which includes nearly all computing departments in UK universities.

As commented on by the Shadbolt review[28], the purpose of BCS academic accreditation is to validate that university departments have the right policies and procedures in place to assure that course curricula as well as teaching and learning practices support students to develop the appropriate knowledge, understanding and skills to start on a professional computing career. It is worthwhile noting that this is not necessarily the same as giving

[49] https://wam.bcs.org/wam/coursesearch.aspx

assurance and validation that every graduate is equipped with the knowledge, understanding and skills needed to embark on a professional computing career. The former is a proxy for the later, but realistically that amounts to validation for the overall cohort rather than being a cast iron guarantee for each and every individual graduate.

Employers providing the majority of funding for an AI MSc course would like to have assurance that every successful graduate has developed appropriate academic knowledge, understanding and professional skills rigorously evidenced against standards they value. For that to be achieved for industry funded AI MSc programmes it will be necessary to build on the current level of professional body academic accreditation to provide individual validation of professional development. This type of accreditation at undergraduate level is already being established through initiatives such as RITTech professional status[50], which is being adopted for sandwich degree placements by some universities. The RITTech standards were defined for apprenticeships, but they are equally valid for use in setting out professional development criteria undergraduates could achieve during a sandwich year placement.

For the student the benefits of measuring their professional development against RITTech is that they gain registration with a Chartered professional body, they are able to readily articulate their professional skills when applying for jobs, and they can see what career progression options are available through those skills because they are mapped to a skills framework that includes career pathways. These are all also benefits for the employer when recruiting from RITTech accredited sandwich degree programmes.

For future industry funded AI MSc programmes it is worth considering whether professional skills graduates develop through work experience can be validated against an industry recognised skills framework in a similar way to RITTech. At present some universities are exploring whether it may be possible to validate level 7 MSc apprenticeship and MSc with a year placement against Chartered IT Professional status (CITP). If that were possible it would provide the same advantages to an MSc programme as the RITTech status does for undergraduate programmes, and potentially allow an MSc student to achieve CITP directly upon graduation.

6.4 Market driven or centrally coordinated

In considering the balance between market forces driving HE provision of AI MSc graduates and the role of coordination between university supply and employer demand it is worth looking at the relevant findings from the Shadbolt review[28]. The Shadbolt review reported, amongst other findings:

- While many employers find that Computer Sciences graduates are well prepared for work, there continues to be a bloc of opinion that suggests that more could be done to improve graduates' skills and work readiness. However, a clear challenge is that employers are often divided on where the problem lies.

[50] https://www.bcs.org/membership/get-registered/professional-registration-for-it-technicians-rittech/registered-it-technician-rittech-standard/

- Employers disagree on what technical skills Computer Sciences students should be taught, although the balance of evidence points to support for HE providers teaching the fundamental principles of Computer Science, and encouraging and enabling students to learn and adapt to new technologies over their careers.
- There is a mandate from employers and the HE sector to strengthen the current accreditation framework so that it is more focused on outcomes and links more closely with employability. It would benefit all stakeholders, including graduates, if employment outcomes, and employability, were to become a more central part of accrediting a degree programme.
- Accreditation should seek to support greater interaction between industry and HE, providing the mechanism to influence the design of degree programmes and an avenue for articulating the changing requirements of industry.

The Wakeham review[51] looking at employability of STEM graduates came to analogous conclusions, including:

- Greater collaboration between business and HE is vital to ensuring appropriately educated and skilled [STEM] graduates. The implied partnership endows each partner with responsibilities that should be explicitly accepted.
- The accreditation of degree programmes under the auspices of professional, statutory and regulatory bodies is an important enabling feature of the current landscape that should be more effectively exploited throughout STEM.

Another consideration when discussing the balance between market driven development and more central coordination is that of scaling the provision of validated work experience at a sufficient pace, which as shown in Section 6.1 is likely to be a challenge if left purely to market forces. One mechanism that would help grow the number of validated work experience opportunities would be a national brokering scheme to help employers, universities and students find the best match. The Institute of Coding[52] (IoC) was in part established to develop specialist skills training in areas of strategic importance, and the IoC includes over thirty university partners at the time of writing as well as professional bodies such as the BCS. Although the remit for the IoC is focused on undergraduate provision, in principle they have the appropriate partners to provide such a brokering function.

From the consultations that have informed this report it's clear the above Shadbolt review findings are equally applicable to industry funded AI MSc programmes as to undergraduate computing degrees. This suggests there is a role for strategically focused coordination between HE and industry through an organisation such as, for example, the IoC, which should build on existing accreditation frameworks from professional bodies, such as for example BCS and the IET, but which also strengthens them to validate graduates professional skills based on industry recognised standards, for example by building on and extending beyond initiatives such as the RITTech case study.

---

51 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/518582/ind-16-6-wakeham-review-stem-graduate-employability.pdf

52 https://www.gov.uk/government/news/prime-minister-announces-20-million-institute-of-coding

Appendices

The rest of the document contains appendices of relevant data, discussions, evaluations and other contextual information that has been used to inform the main body of the report.

7 Appendix: Data Science related degree apprenticeships

This appendix includes the standards related to Data Science apprenticeships that are approved for delivery by Institute for Apprenticeships and Technical Education. These are the Digital And Technology Solution Specialist (Degree Level 7, i.e. MSc), which includes the occupation specialism of Data analytics specialist and Data Scientist (Degree Level 6, i.e. BSc).

It includes an evaluation of how these standards compare to the AI Skills Framework in Section 2.1, and how much each is covered by the other.

7.1 Digital And Technology Solution Specialist Apprenticeship

This is a specialist occupation within the Level 7 approved for delivery apprenticeship of Digital And Technology Solution Specialist. The standard splits into core technical skills, core technical knowledge together with core specialist skills and core specialist knowledge. The core skills and knowledge are, as the name suggests, core to most level 7 digital technology occupations. Details for those can be found on the IFA website, and are not considered here since they are generic.

Below we list the skills and knowledge for this occupation specialism, where each item is annotated with an initial evaluation of how relevant it is to the framework in Section 2.1.

Data Analytics Specialist Apprenticeship Standard	Extent item covered by Draft AI Skills Framework
Skills for Data Analytics Specialist	
Identify and select the business data that needs to be collected and transitioned from a range of data systems; acquire, manage and process complex data sets, including large-scale and real-time data	Item is covered by framework
Undertake analytical investigations of data to understand the nature, utility and quality of data, and developing data quality rule sets and guidelines for database designers	Item is covered by framework
Formulate analysis questions and hypotheses which are answerable given the data available and come to statistically sound conclusions	Item is covered by framework
Conduct high-quality complex investigations, employing a range of analytical software, statistical modelling & machine learning techniques to make data driven decisions solve live commercial problems	Item is covered by framework

Document and describe the data architecture and structures using appropriate data modelling tools, and select appropriate methods to present data and results that support human understanding of complex data sets	Item is covered by framework
Scope and deliver data analysis projects, in response to business priorities, create compelling business opportunities reports on outcomes suitable for a variety of stakeholders including senior clients and management	Item is covered by framework
Technical knowledge for Data Analytics Specialist	
How key algorithms and models are applied in developing analytical solutions and how analytical solutions can deliver benefits to organisations	Item is covered by framework
The information governance requirements that exist in the UK, and the relevant organisational and legislative data protection and data security standards that exist. The legal, social and ethical concerns involved in data management and analysis	Item is covered by framework
The principles of data driven analysis and how to apply these. Including the approach, the selected data, the fitted models and evaluations used to solve data problems;	Item is covered by framework
The properties of different data storage solutions, and the transmission, processing and analytics of data from an enterprise system perspective. Including the platform choices available for designing and implementing solutions for data storage, processing and analytics in different data scenarios	Item is covered by framework
How relevant data hierarchies or taxonomies are identified and properly documented	Item consistent with framework, but not explicitly covered
The concepts, tools and techniques for data visualisation, including how this provides a qualitative understanding of the information on which decisions can be based.	Item is covered by framework

7.2 Data Science Apprenticeship

This is a specialist occupation within the Level 6 approved for delivery apprenticeship of Digital And Technology Solution Specialist. The standard splits into core specialist skills and core specialist knowledge. We list below the specialist skills and knowledge for this occupation specialism, evaluated against the AI Skills Framework outlined in Section 2.1.

Data Science Apprenticeship Standard	Extent item covered by Draft AI Skills Framework
A Data Scientist must understand	
The context of Data Science and the Data Science community in relation to computer science, statistics and software engineering. How differing schools of thought in these disciplines have driven new approaches to data systems	Item goes beyond the framework

How Data Science operates within the context of data governance, data security, and communications. How Data Science can be applied to improve an organisation's processes, operations and outputs. How data and analysis may exhibit biases and prejudice. How ethics and compliance affect Data Science work, and the impact of international regulations (including the General Data Protection Regulation.)	Item is covered by the framework
How data can be used systematically, through an awareness of key platforms for data and analysis in an organisation	Item consistent with framework, but not explicitly covered
How to design, implement and optimise analytical algorithms – as prototypes and at production scale	Item is covered by the framework
The data landscape: how to critically analyse, interpret and evaluate complex information from diverse datasets	Item is covered by the framework
Skills: A Data Scientist is able to	
Identify and clarify problems an organisation faces, and reformulate them into Data Science problems. Devise solutions and make decisions in context by seeking feedback from stakeholders. Apply scientific methods through experiment design, measurement, hypothesis testing and delivery of results. Collaborate with colleagues to gather requirements	Item is covered by the framework
Perform data engineering: create and handle datasets for analysis. Use tools and techniques to source, access, explore, profile, pipeline, combine, transform and store data, and apply governance (quality control, security, privacy) to data	Item is covered by the framework
Identify and use an appropriate range of programming languages and tools for data manipulation, analysis, visualisation, and system integration. Select appropriate data structures and algorithms for the problem. Develop reproducible analysis and robust code, working in accordance with software development standards, including security, accessibility, code quality and version control	Item is covered by the framework
Use analysis and models to inform and improve organisational outcomes, building models and validating results with statistical testing: perform statistical analysis, correlation vs causation, feature selection and engineering, machine learning, optimisation, and simulations, using the appropriate techniques for the problem	Item is covered by the framework
Implement data solutions, using relevant software engineering architectures and design patterns. Evaluate Cloud vs. on-premise deployment. Determine the implicit and explicit value of data. Assess value for money and Return on Investment. Scale a system up/out. Evaluate emerging trends and new approaches. Compare the pros and cons of software applications and techniques	Item goes beyond the framework
Find, present, communicate and disseminate outputs effectively and with high impact through creative storytelling, tailoring the	Item is covered by the framework

message for the audience. Use the best medium for each audience, such as technical writing, reporting and dashboards. Visualise data to tell compelling and actionable narratives. Make recommendations to decision makers to contribute towards the achievement of organisation goals	
Develop and maintain collaborative relationships at strategic and operational levels, using methods of organisational empathy (human, organisation and technical) and build relationships through active listening and trust development	Item goes beyond the framework
Use project delivery techniques and tools appropriate to their Data Science project and organisation. Plan, organise and manage resources to successfully run a small Data Science project, achieve organisational goals and enable effective change.	Item consistent with standard, but not explicitly covered

7.3 Evaluation of AI Skills Framework against Data Apprenticeship Standards

The following table gives an initial evaluation of how much the draft AI Skills Framework is covered by the Data Analytics Specialist Standard and the Data Science Apprenticeship Standard.

Together with the earlier tables in this section, the evaluation shows that the draft AI Skills Framework encapsulates the vast majority of the Data Analytics specialism standard, whilst also going beyond the standard in a significant number of places. For example, the framework is far more extensive on the ethical concerns that need to be mitigated against by an AI practitioner.

The table shows that the Data Science standard encapsulates a great deal of the high level skills in the draft AI Skills Framework. Where these characterisations differ is the level of appropriate knowledge and understanding required for the relevant skills. The AI Skills Framework is predicated on a level of knowledge and understanding provided through a level 7 MSc course, whereas the Data Science apprenticeship is a level 6 BSc qualification.

Draft AI Skills Framework	**Extent item covered by Data Analytics Specialist Standard**	**Extent item covered by Data Science Standard**
Science skills		
Understands how AI algorithms, techniques and methodologies are designed, developed, optimised and applied at scale to achieve business objectives.	Item goes beyond the standard	Item is covered by standard
Can select and use appropriate statistical methods for sampling, distribution assessment, bias and error.	Item is covered by standard	Item is covered by standard

Understands AI problem structuring methods and can evaluate which method is most appropriate for business needs.	Item consistent with standard, but not explicitly covered	Item consistent with standard, but not explicitly covered
Applies rigorous scientific methodologies through experimental design, exploratory modelling and hypothesis testing to reach robust conclusions, and can explain how those are reached to internal and external stakeholders.	Item is covered by standard	Item is covered by standard
Data engineering skills		
Has a demonstrable understanding of how to expose data from systems, how to efficiently extract data from potentially heterogeneous source systems, and how to ensure standards of data quality and consistency for processing by AI systems.	Item consistent with standard, but not explicitly covered	Item consistent with standard, but not explicitly covered
Works with other technologists and analysts to integrate separate data sources in order to map, produce, transform and test new scalable AI products and services that meet user needs.	Item goes beyond the standard	Item consistent with standard, but not explicitly covered
Works with other technologists and analysts to understand and make use of different types of data models.	Item is covered by standard	Item is covered by standard
Understands and can make use of different data engineering tools for repeatable data processing and is able to compare between different data models.	Item is covered by standard	Item is covered by standard
Understands how to build scalable machine learning pipelines and combine feature engineering with optimisation methods to improve the data product performance.	Item goes beyond the standard	Item is covered by standard
Product development		
Uses a range of professional coding practices to build reliable, reusable, scalable AI products and services to time, quality and budget.	Item goes beyond the standard	Item is covered by standard
Can work as part of a team to effectively integrate AI technologies into business systems.	Item is covered by standard	Item is covered by standard
Can demonstrate why AI products and services are valid against user requirements in a manner comprehensible to the relevant internal and external stakeholders.	Item goes beyond the standard	Item is covered by standard
Works in accordance with agreed software development standards, including security, accessibility and version control.	Item is covered by standard	Item is covered by standard
Business skills		

Understands the context of the business including its processes, data, priorities and its wider values, objectives and strategy	Item is covered by standard	Item is covered by standard
Can effectively communicate the value, opportunities and limitations of AI technologies to a range of audiences with varying technical background	Item is covered by standard	Item is covered by standard
Uses the most appropriate medium to visualise AI based outputs to tell compelling and actionable stories relevant for business goals	Item is covered by standard	Item is covered by standard
Maintains a user focus to design AI solutions that meet user needs, taking account of ethical issues	Item is covered by standard	Item is covered by standard
Is familiar with the state of the art of techniques that help in modelling and understanding a business and its operation	Item is covered by standard	Item is covered by standard
As part of a team, is able to support the scoping and business priority setting for large or complex changes caused by the adoption of AI, engaging senior stakeholders as required	Item is covered by standard	Item is covered by standard
As part of a team, is able to help identify the impact on business requirements of adopting AI	Item is covered by standard	Item is covered by standard
As part of a team, uses the appropriate methods and techniques for the assessment and management of business risks that might result from adopting AI technologies.	Item is covered by standard	Item is covered by standard
Ethical concerns		
Unfair or prejudiced bias in data or models	Item goes beyond the standard	Item is covered by standard
Potential unconscious bias of AI practitioners and product development teams	Item goes beyond the standard	Item goes beyond the standard
Appropriate level of transparency in design and development of AI models	Item goes beyond the standard	Item goes beyond the standard
The impact of AI on restricting or enhancing user autonomy and wellbeing, whether in the workforce, in the customer base or society at large	Item goes beyond the standard	Item goes beyond the standard
The ability of individuals to have appropriate control over their personal data	Item goes beyond the standard	Item goes beyond the standard
Potential unintended, inappropriate or malicious use of AI products or services	Item goes beyond the standard	Item goes beyond the standard
Proactively works with and supports organisational stakeholders to develop appropriate policies, processes and practices to prevent unethical issues arising from the design, development or deployment of AI products and services.	Item goes beyond the standard	Item goes beyond the standard

Proactively support their organisation to improve the ethnic and gender diversity and inclusivity of the workforce at all levels.	Item goes beyond the standard	Item goes beyond the standard

This evaluation demonstrates it is worth investigating the feasibility of developing the level 6 Data Science apprenticeship into a level 7 AI practitioner apprenticeship, since at a high level there is very significant overlap between the two.

8 Appendix: Computer Science Taught MSc Graduates 2017/18

The table bellow shows the number of graduates from taught MSc Computer Science courses[53] in the UK for 2017/18. Note this excludes graduates from MSc by research only, and it does not include PhD students. The table also shows how many M-level[54] taught sandwich courses each of those universities offered in 2017/18. The table is intended to give an *approximate* view of the overall national capacity for offering taught MSc courses in Computer Science, including whether departments also offer year long placements. The table shows there were 6735 graduates within the category of taught MSc Computer Science courses in 2017/18. It also shows of universities offering taught MSc courses 50 of them also provide M-level sandwich degree courses in Computer Science.

Note a sandwich course requires a full year long placement, so that column does not count courses where a shorter placement or internship was offered. E.g. UCL in 2017/18 provided three month long summer internships, but not year long placements. The table does not include universities that did not run taught MSc Computer Science courses. E.g. if a university only offered MSc by research for Computer Science it is not included in this particular table.

The data is taken from HESA https://www.hesa.ac.uk/dAta-and-analysis/students/table-19 and from Unistats https://unistats.ac.uk

UKPRN	HE provider	Computer Science Subjects Taught Masters Graduates	M-Level Sandwich Courses (MSci, MComp, MEng)
10000291	Anglia Ruskin University	35	0
10000571	Bath Spa University	5	0
10000824	Bournemouth University	55	0
10000886	The University Of Brighton	30	4

[53] This includes MSc courses specialising in any sub-field of computer science or broadly covering computer science. We aggregate courses in this was because historically universities have often mislabelled subject specialisms with the wrong JISC codes in returns to HESA.

[54] An M-level degree is an integrated degree that combines an undergraduate Bachelors degree with a Masters level one year degree. These are usually four or five year long degrees, which usually depends on whether they involve a placement.

10000961	Brunel University London	45	4
10000975	Buckinghamshire New University	5	0
10001282	The University Of Northumbria At Newcastle	95	6
10001478	City, University Of London	200	4
10001726	Coventry University	140	1
10001883	De Montfort University	70	0
10002718	Goldsmiths College	45	0
10003270	Imperial College Of Science, Technology and Medicine	145	0
10003645	King's College London	135	0
10003678	Kingston University	90	0
10003861	Leeds Beckett University	15	0
10003957	Liverpool John Moores University	30	9
10004048	London Metropolitan University	25	0
10004078	London South Bank University	15	2
10004113	Loughborough University	40	8
10004180	The Manchester Metropolitan University	45	0
10004351	Middlesex University	50	3
10004511	The National Film and Television School	10	0
10004797	Nottingham Trent University	55	5
10004930	Oxford Brookes University	35	0
10005343	Queen's University Belfast	85	3
10005500	The Robert Gordon University	60	0
10005553	Royal Holloway, University of London	165	5
10005790	Sheffield Hallam University	110	1
10006022	Solent University	5	0
10006299	Staffordshire University	35	0
10006566	The University Of West London	25	0
10006840	The University Of Birmingham	155	4
10006841	The University Of Bolton	5	0
10006842	The University Of Liverpool	85	1
10007138	The University Of Northampton	10	0
10007140	Birmingham City University	50	6
10007141	The University Of Central Lancashire	30	8
10007143	University Of Durham	20	0
10007144	The University Of East London	25	0
10007146	The University Of Greenwich	65	1
10007147	University Of Hertfordshire	80	0
10007148	The University Of Huddersfield	20	4
10007149	The University Of Hull	15	1
10007150	The University Of Kent	145	1
10007151	The University Of Lincoln	10	3
10007152	University Of Bedfordshire	65	0
10007154	University Of Nottingham	65	0
10007155	The University Of Portsmouth	65	1
10007156	The University Of Salford	60	0

10007157	The University Of Sheffield	130	14
10007158	The University Of Southampton	160	13
10007159	The University Of Sunderland	15	0
10007160	The University Of Surrey	45	2
10007161	Teesside University	105	3
10007163	The University Of Warwick	80	6
10007164	University Of the West Of England, Bristol	20	10
10007165	The University Of Westminster	45	0
10007166	The University Of Wolverhampton	20	1
10007167	The University Of York	65	0
10007759	Aston University	10	1
10007760	Birkbeck College	65	0
10007762	Glasgow Caledonian University	30	1
10007764	Heriot-Watt University	45	0
10007767	Keele University	15	1
10007768	The University Of Lancaster	45	1
10007772	Edinburgh Napier University	115	0
10007773	The Open University	50	0
10007774	The University Of Oxford	130	0
10007775	Queen Mary University Of London	95	2
10007783	The University Of Aberdeen	25	7
10007784	University College London	500	0
10007785	The University Of Bradford	25	0
10007786	The University Of Bristol	125	0
10007788	The University Of Cambridge	70	0
10007789	The University Of East Anglia	30	0
10007790	The University Of Edinburgh	375	0
10007791	The University Of Essex	40	1
10007792	The University Of Exeter	10	2
10007793	University Of South Wales	40	0
10007794	The University Of Glasgow	210	1
10007795	The University Of Leeds	35	13
10007796	The University Of Leicester	50	3
10007798	The University Of Manchester	240	5
10007799	Newcastle University	155	17
10007800	The University Of the West Of Scotland	35	0
10007801	University Of Plymouth	15	0
10007802	The University Of Reading	20	2
10007803	The University Of St Andrews	85	0
10007804	The University Of Stirling	40	0
10007805	The University Of StrAthclyde	15	0
10007806	The University Of Sussex	50	4
10007807	Ulster University	20	0
10007814	Cardiff University	95	2
10007822	Cranfield University	40	0
10007823	Edge Hill University	25	2

10007833	Glyndŵr University	15	0
10007848	University Of Chester	5	0
10007849	University Of Abertay Dundee	10	0
10007850	The University Of Bath	45	6
10007851	University Of Derby	15	0
10007852	The University Of Dundee	60	0
10007854	Cardiff Metropolitan University	10	0
10007855	Swansea University	35	3
10007856	Aberystwyth University	5	1
10007857	Bangor University	5	0
10007858	University Of Wales Trinity Saint David	5	0

9 Appendix: Examples of syllabi for taught MSc courses related to AI and Machine Learning

The following sections provide details from *some* of the modules from university MSc courses that are particularly relevant to Machine Learning or Artificial Intelligence, together with links to the course websites. The reader should not regard this section as definitive. Details from all the courses that were signposted through the consultations for this report have been included, but by no means do we claim that this list is exhaustive, rather it is included to illustrate the range and depth of syllabi available across universities.

In compiling this section taught MSc courses that are listed as Artificial Intelligence, Machine Learning, Data Science or Data Analytics have been considered for inclusion. Some other courses were also included when they were signposted during consultations.

Universities do not provide MSc course details in a standardised format. In this section we have tried to synthesis as uniform a description of relevant sections of syllabi as is possible, based on information available at the time of writing. This has not been done slavishly, since in some cases universities provided details in a way that was useful to include even though it was different to other universities. We do not attempt to include the entire syllabi of courses, only some of the sections of a syllabi that are likely to be informative within the remit of this report.

Whenever possible details of the modules in courses are taken verbatim from university websites, and we include links to the module descriptions. In some cases where websites do not include relevant details of syllabi, an overview of the programme is included when that was provided through one of the consultations. This section focuses on summarising the syllabi of the modules, and not on the aims, learning outcomes, assessment or teaching methods of those modules. In some cases, learning outcomes are included when they provide the most coherent overview available of module content.

9.1 University of Aberdeen Artificial Intelligence MSc

The University of Aberdeen runs a 12 month taught MSc in Artificial Intelligence[55]. The course covers the following modules.

Foundations of AI (15 Credits) - fundamental techniques of Artificial Intelligence, used in system such as Google Maps, Siri, IBM Watson, as well as industrial automation systems, and which are core to emerging products such as self-driving vehicles. This course will equip the student to understand how such AI technologies operate, their implementation details, and how to use them effectively.

Engineering AI Systems (15 Credits) - provides students with skills to help them engineer AI systems, equipping them with solid programming skills, and using state-of-the-practice languages, tools and technologies.

Machine Learning (15 Credits) - presents the fundamental as well as the most popular Machine Learning theories and algorithms, used in a wide range of applications such as face detection, anomaly detection, and which are core to the design of for instance computer Go player AlphaGo. This course provides the building blocks for understanding and using Machine Learning techniques and methodologies and prepares students to work in data science and general AI systems.

Evaluation of AI Systems (15 Credits) - provides students with knowledge of core evaluation concepts, approaches, tools, techniques and technologies.

Data Mining and Visualisation (15 Credits) - provides students with knowledge of core data mining and visualisation approaches, tools, techniques and technologies.

Natural Language Generation (15 Credits) - provides students with knowledge of core natural language generation concepts, approaches, tools, techniques and technologies.

Knowledge Representation and Reasoning (15 Credits) - An underlying feature of many AI systems concern how knowledge is acquired, represented, and reasoned with. This module provides the theory and practice of knowledge representation and reasoning, also presenting cutting-edge technologies, libraries and tools. At the end of the course students will be able to design, implement and evaluate knowledge-intensive AI systems.

Software Agents and Multi Agent Systems (15 Credits) - Autonomous systems act to achieve goals with no human intervention, and are already found in Tesla's self-driving cars, NASA space probes and systems such as Amazon's Echo. This course provides the student with a solid grounding in the theory and tools which underpin such systems, teaching them both how to develop such systems, and use them effectively as part of a larger product.

[55] https://www.abdn.ac.uk/study/postgraduate-taught/degree-programmes/1034/artificial-intelligence/

9.2 University of Bath MSc in Machine Learning and Autonomous Systems with one year placement

The University of Bath[56] runs a two year MSc in Machine Learning and Autonomous Systems including a placement year. Some of the modules included in the course are as follows.

Statistics for data science (12 credits)

- The laws of probability.
- Discrete and continuous random variables.
- Bayes' Theorem.
- Expectation, variance and correlation.
- Conditional and marginal distributions.
- Common distributions including the normal, binomial and Poisson.
- Statistical estimation including maximum likelihood.
- Hypothesis testing and confidence intervals.

Machine learning 1 (12 credits)

- Numerical optimisation for parameter estimation;
- Algorithmic unsupervised learning (e.g. k-mean clustering and principal component analysis);
- Discriminative approaches to classification and regression;
- Fundamental parametric linear models (e.g. generalised linear models), parametric non-linear models (e.g. decision trees), non-parametric models (e.g. k-nearest neighbours), and ensemble approaches (e.g. boosting).

Intelligent agents (12 credits)

- Agent architectures,
- Agent platforms
- Communication and content languages
- Agent-oriented software engineering,
- Virtual enterprise formation, institutions and norms

Intelligent control and cognitive systems (12 credits)

- Why intelligent control is (computationally) hard, outline
- Review of historic strategies (proof / search based, reactive / dynamic planning, machine learning, hybrids of these).
- Action: mechanisms for sequencing, goal arbitration, problem spaces and contexts. Where do action primitives come from, how does morphology do work for you. Redundancy & degrees of freedom.
- Perception and Learning: sensor fusion, memory, and learning. The beginnings of cognition.

[56] http://www.bath.ac.uk/courses/postgraduate-2018/taught-postgraduate-courses/msc-machine-learning-and-autonomous-systems-including-placement-year

- Introduction to agent-based modelling; the impact of concurrency and society; simulations in policy and science; models, simplicity and explanation.
- Natural intelligence: Evolution and cognitive control, variation in cognitive strategies found in nature, individual variation in nature; perception and action selection in nature.
- Sensing & Action primitives II: Animation and Virtual Reality. Motion capture, segment smoothing. Motion planning and basic AI for games.
- Complex planning systems, achieving multiple goals, agents with emotions and personality. Likeability, believability and engagement.
- Ethics and philosophy of AI, can we build consciousness? What should our users believe about our agents?

Machine learning 2 (12 credits)

- Bayesian approaches to ML, graphical models (e.g. Markov random fields),
- Bayesian non-parametric models (e.g. Gaussian processes)
- Deep learning (e.g. neural networks)
- Time series (e.g. hidden Markov models)
- Sparse models (e.g. compressed sensing)
- Unsupervised learning (e.g. density estimation)

Bayesian Machine Learning (12 credits)

- Bayesian inference
- Probabilistic modelling paradigms
- Programming Bayesian analytic algorithms
- Bayesian modelling
- Bayesian treatments of linear models, neural networks and Gaussian processes

Reinforcement Learning (12 credits)

- Dynamic programming,
- Monte Carlo methods,
- Temporal-difference algorithms,
- Integration of planning and learning,
- Value function approximation,
- Policy gradient methods

9.3 University of Bristol MSc Advanced Computing in Machine Learning, Data Mining and High Performance Computing

The University of Bristol[57] runs a one year taught MSc Advanced Computing in Machine Learning, Data Mining and High Performance Computing. Some of the modules included in the course are as follows.

---

[57] https://www.bris.ac.uk/unit-programme-catalogue/RouteStructureCohort.jsa?byCohort=Y&ayrCode=18%2F19&programmeCode=4COSC010T

Machine Learning (10 credits) - fundamental principles of machine learning and how to build models of data.

Uncertainty Modelling for Intelligent Systems (10 credits) - approaches available in uncertainty modelling, and under what conditions they can be appropriately applied; advanced mathematics underlying different approaches to uncertainty modelling; practical applications of uncertainty modelling.

Applied Data Science (10 credits) - Data ingress and pre-processing; Data storage and data management; Data transformation and integration; Data exploration and visualisation; Data sharing, privacy and anonymisation.

Other **optional modules** that can be taken as part of the course include the following.

Artificial Intelligence with Logic Programming (10 credits) - principles of knowledge representation and automated inference with the Prolog programming language; advanced methods in natural language processing and machine learning which exploit the representation and reasoning power of Prolog.

Neural Information Processing (10 credits) - fundamental concepts related to information processing in the brain and how these theories and methods can be applied in the derivation of principled approaches to learning problems and assimilation; modelling how neurons & the nervous system encode and exchange information; how to apply small and large scale brain networks; how to model learning, including value learning and deep learning and inference.

The course also includes a three months placement as the fourth semester.

9.4 University of Cambridge

The University of Cambridge offer an MPhil in Advanced Computer Science[58] (ACS), and an MPhil in Machine Learning and Machine Intelligence[59].

Advanced Computer Science

ACS is a research preparation degree, specifically designed to provide training for students planning a career in a research laboratory, including companies such as Google DeepMind or Facebook AI Research, or wishing to proceed to a PhD in computer science.

At present, the majority of ACS students go on to research careers in Machine Learning and AI. The university admits students working in other areas of Computer Science, many of which are also relevant in AI research (e.g. type theory, security, compilers, networking). Based on the research dissertation topics in the last two years (2018-19), 65-70% of those students work on AI topics.

The curriculum in the ACS can be tailored to the specific interests of students, allowing specialisation at the Masters level. Each student chooses 5 taught modules, from a

[58] https://www.graduate.study.cam.ac.uk/courses/directory/cscsmpacs
[59] https://www.graduate.study.cam.ac.uk/courses/directory/egegmpmsl

selection in which 13 of the available modules focus on topics in data science, machine learning and AI. Students are also able to take modules on a range of other topics in Advanced Computer Science, but at present the majority of students take all or almost all of their modules in the broad AI/ML.

Machine Learning and Machine Intelligence

The course has an emphasis on machine learning and machine intelligence, including speech and language technology and computer vision. The course aims to teach the state of the art in machine learning and machine intelligence; to give students the skills and expertise necessary to take leading roles in industry; and to equip students with the research skills necessary for doctoral study. The intention of the course is that the end of the programme, students will have acquired:

- a knowledge of the fundamental techniques in machine learning and how to apply these techniques to a range of practical problems;
- a deep understanding of fundamental problems in machine intelligence, including speech and language processing and computer vision, and the technologies that form the current state of the art;
- a comprehensive understanding of techniques, and a thorough knowledge of the literature, applicable to the area of their chosen research topic;
- the ability to critically assess the technical literature in machine learning and machine intelligence and related topics;
- directly marketable skills in computing, machine intelligence, machine learning, and the data sciences;
- collaborative skills through working with other students on the practical exercises and with PhD students and research assistants while carrying out their research project;
- experience in large-scale computing for machine learning and machine intelligence;
- an understanding of how to define and conduct a research project.

9.5 City, University of London

City runs an MSc in Data Science[60]. The course covers the study, integration and application of advanced methods and techniques from:

- Data analysis and machine learning
- Data visualisation and visual analytics
- High-performance, parallel and distributed computing
- Knowledge representation and reasoning
- Neural computation
- Signal processing
- Data management and information retrieval.

Some of the core modules of the course includes:

[60] https://www.city.ac.uk/study/courses/postgraduate/data-science-msc

Principles of Data Science (15 credits)

- Understand the foundations of the data science process, methods and techniques
- Represent and organise knowledge about large heterogeneous data collections
- Use mathematical models and tools for large-scale data analysis and reasoning
- Critically evaluate the choice of data science techniques and tools for particular scenarios

Machine Learning (15 credits)

- Understand the workings of important data science algorithms for learning under uncertainty
- Rationally exploit both statistical and machine learning approaches in applications
- Rigorously assess the validity of inferences and generalisations
- Critically evaluate the choice of algorithms for specific scenarios and requirements

Big Data (15 credits)

- Understand the theory and techniques for data acquisition, cleansing, and aggregation
- Identify and understand the principles and functionalities of big data programming models and tools
- Acquire, process and manage large heterogeneous data collections
- Develop algorithms and systems for information and knowledge extraction from large data collections

Neural Computing (15 credits)

- Understand how to use neural computing and deep learning in an application domain
- Select and apply supervised, unsupervised and hybrid neural networks to different problems and data types
- Critically evaluate a range of neural systems in comparison with a number of learning techniques
- Design and implement neural network models; apply them and evaluate their performance

Visual Analytics (15 credits)

- Learn the principles and rules underlying the design of visual data representations and human-computer interactions
- Understand, adapt and apply representative visual analytics methods and systems for diverse types of data and problems
- Analyse and evaluate the structure and properties of data to select or devise appropriate methods for data exploration
- Combine visualization, interactive techniques, and computational processing to develop practical data analysis for problem solving

Research Methods and Professional Issues (15 credits)

- Understand and apply research methodologies such as inductive and deductive reasoning, explanation and prediction
- Recognise and apply the scientific method and a range of secondary data sources when performing a research task
- Communicate effectively with individuals and groups using a range of media
- Evaluate the legal, ethical and professional dimensions of typical information professions and information industry practices.

9.6 Coventry University

Coventry University offers an MSc in Computer Science[61]. The course explores emerging technologies, such as emergent Artificial Intelligence, machine learning, cloud computing, machine vision and image processing. The course looks at emerging theories, practices, approaches and management of distributed and intelligent computing systems, examining a wide range of case studies to see how applications have been developed and for what purposes.

The course includes a 15 credit module on neural networks, a 15 credit module on computer vision (including extensive coverage of ML tools), and a 15 credit module covering evolutionary and fuzzy systems.

9.7 De Montfort University

De Montfort University runs MSc courses on Intelligent Systems[62], Intelligent Systems and Robotics[63], Data Analytics[64], and Business Intelligence Systems and Data Mining[65].

The MSc Intelligent Systems is of particular relevance to this report. Modules in the course includes:

- Computational Intelligence Research Methods details quantitative and qualitative approaches including laboratory evaluation, surveys, case studies and action research.
- Artificial Intelligence (AI) Programming.
- Mobile Robots discusses the hardware and software architectures used to build mobile robot systems.
- Fuzzy Logic considers the various fuzzy paradigms that have become established as computational tools.
- Artificial Neural Networks appraises neural network computing from an engineering approach and the use of networks for cognitive modelling.

[61] https://www.coventry.ac.uk/course-structure/pg/2018-19/eec/computer-science-msc/
[62] https://www.dmu.ac.uk/study/courses/postgraduate-courses/intelligent-systems/intelligent-systems-msc-degree.aspx
[63] https://www.dmu.ac.uk/study/courses/postgraduate-courses/intelligent-systems-and-robotics/intelligent-systems-and-robotics-msc-degree.aspx
[64] https://www.dmu.ac.uk/study/courses/postgraduate-courses/data-analytics-msc-degree/data-analytics-msc-degrees.aspx
[65] https://www.dmu.ac.uk/study/courses/postgraduate-courses/business-intelligence-systems-and-data-mining/business-intelligence-systems-and-data-mining-msc-degree.aspx

- Computational Intelligence Optimisation (CIO) is a subject that integrates artificial intelligence into algorithms for solving optimisation problems that could not be solved by exact methods. It tackles optimisation problems in engineering, economics, and applied science
- Applied Computational Intelligence considers knowledge-based systems; the historical, philosophical and future implications of AI; then focuses on current research and applications in the area.
- Data Mining, Techniques and Applications examines the tools and techniques needed to mine the large quantities of data generated in today's information age. It provides practical experience as well as consideration of research and application areas

9.8 University of Edinburgh

The University of Edinburgh runs MSc courses[66] in Artificial Intelligence, Data Science, and Cognitive Science. There is a certain amount of overlap in the options for these courses. Below we list some of the more relevant modules available from these courses.

9.8.1 MSc in Artificial Intelligence

The course includes the following modules (which are provided as options within the overall programme structure). Because of the course structure it is possible to almost exclusively take options that are relevant to Machine Learning, which would result in an MSc with as much Machine Learning content as some MSc courses specifically on Machine Learning.

Data Mining and Exploration (10 credits)

- Data preprocessing and cleaning, dealing with missing data
- Data visualization, exploratory data analysis
- Data mining techniques
- Predictive modelling techniques (e.g. SVMs)
- Performance evaluation (e.g. ROC curves)
- Issues relating to large data sets
- Application areas, e.g. text mining, collaborative filtering, retrieval-by-content, web mining, bioinformatics data, astronomy data

Reinforcement Learning (10 credits)

- Reinforcement learning framework
- Bandit problems and action selection
- Dynamic programming methods
- Monte-Carlo methods
- Temporal difference methods
- Q-learning and eligibility traces
- Environment modelling

[66] http://www.drps.ed.ac.uk/18-19/dpt/ptmscaintl1f.htm

- Function approximation for generalisation
- Actor-critic, applications
- Planning in the RL context
- Unsupervised, self-organising networks and RL
- Constructive methods - nets that grow
- Evaluating performance

Automatic Speech Recognition (10 credits)

- Signal analysis for ASR
- Statistical pattern recognition (Bayes decision theory, Learning algorithms, Evaluation methods, Gaussian mixture model, and EM algorithm)
- Hidden Markov Models (HMM)
- Context-dependent models
- Discriminative training
- Language models for LVCSR (large vocabulary continuous speech recognition)
- Decoding
- Robust ASR (Robust features Noise reduction, Microphone arrays)
- Adaptation (Noise adaptation, Speaker adaptation/normalization, Language model adaptation)
- Speaker recognition
- History of speech recognition
- Advanced topics (Using prosody for ASR, Audio-visual ASR, Indexing, Bayesian network)
- Speech recognition applications (including privacy implications)

Neural Information Processing (10 credits)

- Neural coding: reverse correlation, higher order kernels, stimulus reconstruction. Application to the fly visual system.
- Information theory as applied to neural coding: mutual information measures, whitening. Application to retinal and LGN coding.
- Networks based on information-theoretic cost functions: Helmholtz machine, Linsker's info-max principle.
- Independent Component Analysis. Basics, variants of ICA, ICA as model for visual cortex.
- Predictive Coding: Kalman filters. Application to cortical coding
- Bayesian approaches: Stimulus estimation, probabilistic interpretation of populations codes.

Computational Cognitive Neuroscience (10 credits)

- Encoding Information in populations of neurons.
- Decoding Information from populations of neurons.
- Models of Neurons and Networks of Neurons.
- Information transmission and Attention.
- Models of Learning and Plasticity.

- Models of Memory.
- Models of Decision Making.
- Models of Mental disorders.
- The Bayesian Brain.

Natural Language Processing (10 credits)

- Probability and information theory.
- Basic introduction to formal language theoretic tools used in NLP such as finite state transducers, advanced grammar topics.
- Introduction and overview of advanced statistical modelling techniques in NLP such as structured prediction, log-linear models.

Accelerated Natural Language (20 credits)

- Inflectional and derivational morphology
- Finite state methods and Regular expressions
- Word Classes and Parts of speech
- Sequence Models (n-gram and Hidden Markov models, smoothing)
- The Viterbi algorithm, Forward Backward, EM
- Syntactic Concepts (e.g., constituency, subcategorisation, bounded and unbounded dependencies, feature representations)
- Analysis in CFG - Greedy algorithms---Shift-reduce parsing
- Divide-and-conquer algorithms---CKY
- Chart parsing
- Lexicalised grammar formalisms (e.g., TAG, CCG, dependency grammar)
- Logical semantics and compositionality
- Semantic derivations in grammar
- Lexical Semantics (e.g., word senses, semantic roles)
- Discourse and dialogue (e.g., anaphora, speech acts)
- Text classification and sentiment analysis
- Other applications (e.g., machine translation, question answering)

Machine Learning and Pattern Recognition (20 credits)

- Classification and Regression: Linear Regression, logistic regression, Bayes classifiers
- Expanded feature representations: Basis functions, neural networks, kernel methods
- Generalization, regularization and inference: Penalized cost functions, Bayesian prediction, learning theory
- Model selection, pruning and combination: Cross-validation, Bayesian methods, sparsifying regularizers, ensemble methods.
- Representation and metric learning: dimensionality reduction, clustering, feature learning
- Optimization and Inference algorithms: Stochastic gradient descent, simple Monte Carlo ideas, and more specialized methods as required.
- Formulating problems as machine learning, adapting methods to fit problems.
- Numerical and programming issues important for machine learning.

- Ethical issues, such as responsible application of methods and privacy concerns.

Machine Learning Practical (20 credits)

- Feed-forward network architectures
- Optimisation and learning rules
- Regularisation and normalisation
- Neural networks for classification
- Autoencoders
- Convolutional Neural Networks
- Recurrent Neural Networks

This module is run as a two-semester course. Semester 1 develops a deep learning framework based on experiments using the task of classification of handwritten digits using the well-known MNIST dataset. The course uses a Python software framework, and a series of Jupyter notebooks.

Semester 2 is based on small group projects, with a focus on using deep neural networks within the context of a miniproject, using an open source toolkit such as TensorFlow or PyTorch.

Probabilistic Modelling and Reasoning (20 credits)

- Probability: events, discrete variables; joint, conditional probability
- Discrete belief networks, inference
- Continuous distributions, graphical Gaussian models
- Learning: Maximum Likelihood parameter estimation
- Decision theory
- Hidden variable models: mixture models and the EM algorithm; factor analysis; ICA, non-linear factor analysis
- Dynamic hidden variable models: Hidden Markov models; Kalman filters (and extensions)
- Undirected graphical models: Markov Random Fields; Boltzmann machines
- Information theory: entropy, mutual information; source coding, Kullback-Leibler divergence
- Approximate Inference: MCMC, Variational Methods
- Bayesian methods for: Inference on parameters; Model comparison

Robot Learning and Sensorimotor Control (20 credits)

- Machine Learning Tools for Robotics: Regression in High Dimensions; Dimensionality Reduction; Online, incremental learning; Multiple Model Learning
- Optimal Control Approaches: LQR, LQG, Dynamic Programming, Trajectory Optimization
- Adaptive Learning and Control: Predictive Control; Underactuation; Multi-contact modelling and optimization; Constrained Operational Space Control; Hierarchical QP and Stack of task formulation; Trajectory based optimization methods; Re-planning in alternate spaces

- Interaction and Robust Control: Cartesian Impedance Control; Passivity Methods; Lyapunov Stability; LQR-Trees and Sum-of-Squares Programming
- Movement Primitives: Rhythmic vs Point to Point Movements; Dynamical Systems and DMPs; Path Integral Methods; Learning by Demonstration
- Planning and Optimization: Stochastic Optimal Control; Bayesian Inference Planning; RL Apprenticeship Learning and Inverse Optimal Control
- Understanding Human Sensorimotor Control: Force Field and Adaptation; Optimal control theory for Explaining Sensorimotor Behaviour; Cue Integration and Sensorimotor Adaptation; Impedance Control Human(oid) Locomotion and Stability.

9.8.2 MSc in Data Science

The course includes the following modules (which are provided as options within the overall programme structure).

Some of the common modules with the MSc in Artificial Intelligence includes (details of these courses are given above): Data Mining and Exploration, Reinforcement Learning, Machine Learning and Pattern Recognition, Probabilistic Modelling and Reasoning, and Introductory Applied Machine Learning.

Some of the other modules available in this MSc course include: Statistical Learning (10 credits), Modern Optimization Methods for Big Data Problems (10 credits), Large Scale Optimization for Data Science (10 credits), Bayesian Data Analysis (10 credits), Statistical Programming (10 credits), Nonparametric Regression Models (10 credits), Generalised Regression Models (10 credits).

9.8.3 MSc in Cognitive Science

The course includes the following modules (which are provided as options within the overall programme structure).

Some of the common modules with the MSc in Artificial Intelligence includes (details of these courses are given above): Data Mining and Exploration, Reinforcement Learning, Automatic Speech Recognition, Neural Information Processing, Computational Cognitive Neuroscience, Adaptive Learning Environments, Natural Language Processing, Machine Learning and Pattern Recognition, Probabilistic Modelling and Reasoning, and Robot Learning and Sensorimotor Control.

9.9 University of Essex

The University of Essex runs a taught MSc Artificial Intelligence course[67]. Some of the modules taught in this course are outlined below (these are the compulsory modules, other optional modules are also available).

Machine Learning And Data Mining (15 credits)

- Taxonomy of machine learning algorithms
- The inductive bias

[67] https://www.essex.ac.uk/courses/pg00457/1/msc-artificial-intelligence?startdate=18-10&studymode=FullTime

- Data mining
- Decision tree induction
- Naïve Bayes methods
- Bayesian networks
- K-nearest neighbour method
- Support vector machines
- Linear Regression
- Regression trees
- Evaluating learning procedures
- Overfitting and the 'bias-variance trade-off'
- k-means algorithm
- Agglomerative hierarchical methods
- Association rules mining: A priori algorithm
- Reinforcement learning: Q learning
- Multiple learners: Bagging, boosting, forests and stacking

Intelligent Systems And Robotics (15 credits)

- A brief history of robotics, types of robots
- Robot challenges (RoboCup, DARPA Grand Challenge)
- Potential applications of intelligent systems and robotics
- Sensors and Actuators: Sonar, laser scanner, optical encoders; DC motors
- Control: Feedback control; Fuzzy controllers
- Localisation and mapping: Triangulation; Kalman filter
- Behaviour based programming: Robot behaviours; Potential field approach; Behaviour based architecture

Neural Networks And Deep Learning (15 credits)

- Introduction to Artificial Neural Networks: basic concepts and principles of ANNs; biological motivations and brief history of ANNs; neuron models and neural network; architectures; computational power of ANNs in comparison with conventional AI methods; ANN applications
- Basic Learning Rules and Theories: basic issues in neural network learning; derivative-based methods such as error gradient descent learning algorithms; derivative-free methods such as simulated annealing, genetic algorithm, Hebbian learning and competitive learning; the bias-variance dilemma in learning from data
- Feedforward Neural Networks Using Supervised Learning: feedforward neural network architectures and supervised learning; perceptron: architecture, error correction learning, limitations; multilayer perceptron (MLP): architecture, back propagation learning algorithm; radial basis function (RBF) network: architecture, learning algorithm, comparison with MLP
- Self-organising Neural Networks Using Unsupervised Learning: unsupervised learning; adaptive resonance theory (ART) neural network; self-organising map (SOM) neural network

- Recurrent Neural Networks: recurrent neural network architectures; Hopfield neural network - energy function, Hebbian learning, stability analysis
- Deep Neural Networks: concepts and architectures
- ANN Applications and Recent Advances: data collection and preprocessing, classification, regression, prediction, and intelligent control; support vector machine (SVM)- reinforcement learning, neuro-fuzzy networks.

9.10 Heriot-Watt University MSc Artificial Intelligence with Speech and Multimodal Interaction

Heriot-Watt University[68] runs a one year MSc in Artificial Intelligence with Speech and Multimodal Interaction. Some of the modules included in the course are described below.

Data Mining and Machine Learning (15 credits)

- Basic Concepts: datasets, dealing with missing data, classification, supervised vs unsupervised learning.
- Generative Models: naïve Bayes, probabilistic graphical models, cluster analysis (such as k-means clustering, EM algorithm).
- Discriminative Learning: linear regression, decision tree learning, perceptron, advanced models such as multi-layer perceptron and deep learning architectures.

Statistical Modelling and Analysis (15 credits)

- Basic probability concepts: Random variables and their distributions; how distributions relate to sampling scenarios.
- Joint distributions, Sums of random variables, Central limit theorems
- Classical inference: Point estimation, moment estimators and maximum likelihood; Confidence intervals – calculation and interpretation; Hypothesis testing and p-values
- Essentials of Bayesian inference: Priors and posteriors; Credible intervals; Predictive distributions
- Modelling approaches: Regression and ANOVA;
- Multivariate exploratory techniques: Principal Components Analysis + Factor Analysis; Introduction to non-parametric methods
- Practical elements in R or Python

Conversational Agents and Spoken Language Processing (15 credits)

- Introduction to research areas, such as spoken dialogue systems, multi-modal interaction, natural language processing, and human robot interaction.
- Spoken input processing and interpretation.
- Interaction Management.
- Output generation, multimodal fission, speech and gesture synthesis
- System development and evaluation.

[68] http://www.macs.hw.ac.uk/cs/pgcourses/aiws.htm

Artificial Intelligence and Intelligent Agents (15 credits, optional)

- Search algorithms (depth first search, breadth first search, uniform cost search, A* search)
- Constraint satisfaction problems;
- Games (min-max, alpha-beta pruning);
- Logic, resolution, introductory logic programming
- Knowledge representation – logic, rules, frames
- Goal and data-driven reasoning
- Practical rule-based programming
- Overview of main fields of AI (Vision, Learning, Knowledge Engineering)
- Autonomous agents
- Applications of AI
- AI programming

9.11 University of Glasgow MSc in Data Science

The University of Glasgow offers an MSc in Data Science[69] teaching 140 students a year. The course includes modules in Machine Learning, Deep Learning, AI, Recommender Systems, Information Retrieval, Text as Data. Many of the students on the course have funding from the DataLab[70] which ensures they have an internship in Industry as part of their experience.

Some of the modules included in the course are described below.

Machine Learning (learning outcomes)

By the end of the module students will be able to

- Demonstrate knowledge of the major machine learning application areas in, for example Information Retrieval, Human Computer Interaction, Bioinformatics and Computer Vision & Graphics;
- Appreciate some emerging machine learning approaches. For example: non-parametric methods and sampling techniques;
- Explain the principle of learning from data;
- Implement and use machine learning algorithms in Matlab;
- Apply the main machine learning methods: regression, classification, clustering, probability density estimation and dimensionality reduction;
- Explain the strengths and weaknesses of a selection of common algorithms.

Deep Learning (learning outcomes)

By the end of the module students will be able to

- Understand the major technology trends in advanced machine learning;
- Build, train and apply fully connected deep neural networks;

[69] https://www.gla.ac.uk/postgraduate/taught/datascience/?gclid=CjwKCAiAhp_jBRAxEiwAXbniXSgd1OyK58fegBtaUDqRlhz-xuwmp7x1UTHnCNm5lvSL6pfuLrA7zhoCyQAQAvD_BwE&gclsrc=aw.ds

[70] https://www.thedatalab.com/

- Know how to implement efficient, vectorised neural networks in python and understand the underlying backends;
- Apply deep learning methods to new applications;
- Understand the machine learning pipeline, and engineering aspects of training data collation, and the importance of unlabelled data.

Regression Models (learning outcomes)

By the end of the module students will be able to

- formulate Normal Linear Models in vector-matrix notation and apply general results to derive ordinary least squares estimators in particular contexts;
- derive, evaluate and interpret point and interval estimates of model parameters and differences between parameters, including multiple comparisons;
- conduct and interpret hypothesis tests in the context of the Normal Linear Model;
- derive, evaluate and interpret confidence and prediction intervals for the response at particular values of the explanatory variables;
- assess the assumptions of a Normal Linear Model using residual plots;
- estimate and make inferences about the population correlation coefficient;
- calculate and comment on R2
- define and briefly explain the relative advantages and disadvantages of common variable selection procedures - stepwise selection, best subsets and lattices - and implement these rules for model building in particular cases;
- define, calculate and use Akaike's Information Criterion (AIC) for model selection;
- describe the problems created by multicollinearity and heteroscedasticity, formulate strategies for the detection and elimination of these problems and implement these strategies in particular cases;
- describe the circumstances under which variable transformation might be required, describe the Box-Cox and Box-Tidwell procedures for transforming variables, formulate a strategy for implementing the Box-Cox scheme in specific examples;
- state the Gauss-Markov Theorem and the implications and restrictions of this result;
- implement these statistical methods using the R computer package;
- frame statistical conclusions clearly.

9.12 University of Hertfordshire MSc in Artificial Intelligence with Robotics

The University of Hertfordshire[71] runs a two year sandwich MSc in Artificial Intelligence with Robotics. The course covers

- Professional Issues (15 credits)
- Investigative Methods for Computer Science (15 credits)
- Artificial Life with Robotics (30 credits)
- Neural Networks and Machine Learning (30 credits)
- Theory and Practice of Artificial Intelligence (30 credits)
- Course project (60 credits)

---

[71] https://www.herts.ac.uk/courses/artificial-intelligence-with-robotics2

9.13 Imperial College MSc in Computing (Artificial Intelligence)

Imperial College[72] runs a taught MSc in Computing (Artificial Intelligence). Some of the modules included in the course are described below.

Mathematics for Machine Learning

- Bayesian Linear Regression
    - Vector Calculus (e.g., partial derivatives, chain rule, Jacobian)
    - Basic probability distributions (e.g., multivariate Gaussian)
    - Bayes' theorem
    - Conjugate priors
    - Gradient descent
    - Model selection
    - Cross validation
    - Maximum likelihood estimation
    - MAP estimation
    - Bayesian integration
    - Graphical model notation
    - Bayesian linear regression
- Probabilistic PCA
    - Eigenvalues
    - Determinants
    - Basis change
    - Singular value decomposition
    - Gram-Schmidt Orthonormalization
    - Rotations
    - Projections

Argumentation and Multi-agent Systems

- Argumentation
    - Abstract argumentation
    - Assumption-based argumentation
    - Bipolar argumentation
    - Argumentation and preferences
    - Argument mining
- Agent communication
    - Agent lingua franca - general purpose agent communication languages (ACLs)
    - Communication policies and protocols, dialogues
- Interactions amongst agents
    - Argumentation: Conflict resolution, persuasion, etc
    - Games: Strategic interaction
    - Negotiation: The monotonic concession protocol
- Agents/Argumentation and other paradigms

[72] http://www.imperial.ac.uk/computing/prospective-students/courses/pg/specialist-degrees/ai/

  - social welfare (utilitarian, egalitarian social welfare)
  - game theory (concession protocols)
  - decision-making, negotiation

Logic Based Learning

- Formal reasoning: Introduction of deductive, abductive and inductive reasoning, with particular emphasis on abductive reasoning for explanation generation, Bayesian reasoning techniques and ASP computational environment.
- Logic-based learning: Logic-based learning as a form of inductive reasoning. Bottom-up approaches to learning: inverse entailment, bottom clause and bottom generalization. Top-down learning techniques, hypothesis specialization and search heuristics. Meta-level logic-based learning as abductive search. Semantic approaches to logic-based learning, with particular emphasis on monotonic and non-monotonic learning, brave and cautious induction.
- Logic-based learning algorithms and systems: Brief summary of logic programming and answer set programming. Key learning algorithms with particular emphasis on Progol, TopLog, Metagol, TAL, Imparo, ASPAL and ILASP. Soundness and completeness of these learning algorithms. Probabilistic/stochastic inductive learning.
- Applications of logic-based Learning: Theoretical applications of logic-based learning to model elaboration and revision. Integration of logic-based learning with model checking. Refinement of software system models. Learning model revision for adaptive systems. Computational systems biology. Computational Learning. Knowledge from text. Robot planning.

Introductory Machine Learning

- Concept Learning
- Decision Trees
- Artificial Neural Networks
- Evaluating Hypotheses
- Instance Based Learning
- Genetic Algorithms

Probabilistic Inference

- Gaussian processes
- Bayesian optimization
- Sampling techniques
- Variational inference
- Modern topics, e.g., implicit models, normalizing flows, amortized inference, stochastic gradient estimators

The MSc course also includes modules on: Machine Learning for Imaging, Deep Learning, and Probabilistic Model Checking and Analysis. Specific details of the syllabi for these modules was not available at the time of writing.

Imperial College also run and MSc in Machine Learning. The machine learning modules available to study for that course are a subset of the modules in the AI MSc.

9.14 King's College London

KCL offer MSc courses in Artificial Intelligence[73] and Data Science[74].

9.14.1 MSc in Artificial Intelligence

The Artificial Intelligence MSc prepares graduates you work developing intelligent software systems. Students can choose to study a wide-range of topics such as: Agents & Multi-Agent Systems, Pattern Recognition, Neural Networks and Deep Learning, Artificial Intelligence Planning, Nature-Inspired Learning Algorithms, Philosophy & Ethics of Artificial Intelligence, and Data Mining.

Some of the modules in the course include:

- Agents & Multi-Agent Systems (15credits)
- Nature-Inspired Learning Algorithms.
- Pattern Recognition, Neural Networks and Deep Learning (15 credits)
- Philosophy & Ethics of Artificial Intelligence (15 credits)
- Artificial Intelligence Planning (15 credits)
- Data Mining (15 credits)
- Computer Vision (15 credits)
- Artificial Intelligence
- Machine Learning

9.14.2 MSc in Data Science

The Data Science MSc degree will provide graduates with the practical skills needed to assemble collate store, manage and analyse data efficiently. Students will also gain the skills necessary to select the appropriate statistical and computational data modelling and analysis techniques to evaluate data science activities and projects.

Some of the modules in the course include:

- Databases, Data Warehousing & Information Retrieval (15 credits)
- Statistics for Data Analysis (15 credits)
- Data Mining (15 credits)
- Elements of Statistical Learning (15 credits)
- Agents & Multi-Agent Systems (15 credits)
- Nature-Inspired Learning Algorithms (15 credits)
- Pattern Recognition, Neural Networks & Deep Learning (15 credits)
- Network Theory (15 credits)
- Big Data Technologies (15 credits)
- Simulation & Data Visualisation (15 credits)
- Machine Learning (15 credits)

[73] https://www.kcl.ac.uk/study/postgraduate/taught-courses/artificial-intelligence-msc
[74] https://www.kcl.ac.uk/study/postgraduate/taught-courses/data-science-msc.aspx

- Artificial Intelligence (15 credits)

9.15 University of Leeds Advanced Computer Science (Intelligent Systems) MSc

The University of Leeds[75] run a taught Advanced Computer Science (Intelligent Systems) MSc. Some of the modules included in the course are described below.

Bio-Inspired Computing (15 credits)

- Examples of cooperative phenomena in nature.
- Concepts such as emergence, self-organisation and embodiment.
- Genetic algorithms.
- Algorithms for swarm intelligence.
- Biological neural networks.
- Various artificial neural networks and their application (eg, clustering, dimensionality reduction).
- Models in computational and cognitive neuroscience.
- Models of biological computation in computational/cognitive neuroscience and/or bioinformatics.

Knowledge Representation and Reasoning (15 credits)

- Review of logical foundations of knowledge representation including key properties of formal systems (such as soundness, completeness, expressiveness and tractability).
- Principles of Logic Programming.
- Representing and reasoning about time and actions and physical changes (e.g., interval calculus, event calculus).
- Representing space and physical situations (topology, orientation, physical objects).
- Automated inference techniques (e.g., refinements of resolution, relational composition, non-monotonic reasoning).
- Ontology representation languages and tools.
- Semantic web applications.
- Formalisms for representing other aspects of knowledge (e.g., vagueness, uncertainty, belief, desire).

Image Analysis (15 credits)

- Image formation and representations;
- Feature detection;
- Texture analysis;
- Stereo;
- Image restoration;
- Shape representation;
- Object tracking;
- Activity analysis;

[75] http://webprod3.leeds.ac.uk/catalogue/dynmodules.asp?Y=201819&F=P&M=COMP-5400M

- Multi-resolution representations;
- Segmentation;
- Object recognition;
- Convolutional neural networks
- Image compression;
- Applications of image analysis.

Big Data Systems (15 credits, optional)

- The five 'Vs' (Volume, Velocity, Variety, Veracity & Value), technology landscape, and future predictions (data, analysis capacity, business opportunities & employment).
- Application contexts: Structure & properties of data, analysis scenarios and case studies (e.g., nuclear physics (volume), social media (velocity), and medical bioinformatics or consumer retail (variety)).
- Systems architectures, encompassing data acquisition, storage, linkage, computation, security/confidentiality, and end-users.
- Key system components, e.g., Hapoop, MapReduce, parallel databases, SQL vs. NoSQL, algorithm scalability, and exploiting existing infrastructure (e.g., Cloud)

Data Science (15 credits, optional)

- Work context and core skills of a data scientist (problem-solving; statistics; business acumen; communication).
- Data governance: ethics, privacy, regulations, policies, and provenance.
- Analysis lifecycle: problem understanding, data acquisition (data types; record linkage; Open Data), data quality (completeness, correctness, concordance, currency & plausibility), analysis techniques, and communicating the results.
- Practical application using case studies drawn from different application domains such as R and Tableau.
- Scale-up of analysis for Big Data.

Scientific Computation (15 credits, optional)

- Numerical solution of a single nonlinear equation.
- Extension of the algorithms to systems of nonlinear equations and reduction to a series of linear equation systems.
- The concept of nonlinear partial differential equations and example applications.
- The need for reliable, efficient and accurate numerical approximation and how this results in discrete systems of nonlinear equations.
- Efficient direct and iterative solution algorithms for large, sparse, linear equation systems.
- Application to problems from classical fluid mechanics and other nonlinear partial differential equations.

Machine Learning (10 credits, optional)

- Neural networks, decision trees, support vector machines, Bayesian learning, instance-based learning, linear regression, clustering, reinforcement learning, deep learning.
- Methods for evaluating performance.

Intelligent Systems and Robotics (20 credits, optional)

- Robotics: Robotic navigation and localization; Robot arm kinematics; Robot motion planning; Robot software (simulations; robot middleware; control algorithms); Robot ethics: concerns about control, deception, privacy and loss of human contact.
- Perception & Language: Signal Processing, Sensors, Geometry of projection and co-ordinate transformations; Binocular and motion stereo; Surface reconstruction; Recognising and tracking objects; Localisation and Mapping; Communication and Language - introduction to linguistic theory and terminology. Algorithms and techniques for computer-assisted speech/text processing.
- Planning and Control: Representation of knowledge of actions and their effect on the world, and how to construct plans to achieve goal states; path planning.
- Multi agent systems: Examples of cooperative phenomena in nature; concepts such as emergence, self-organisation and embodiment; algorithms for swarm intelligence; multi agent communications.

Data Mining and Text Analytics (10 credits, optional)

- Introduction to data and text theory and terminology.
- Tools and techniques for data-mining and text processing, focusing on applied and corpus-based problems such as data classification by Machine Learning classifiers, collocation and co-occurrence discovery and text analytics.
- Open-source and commercial text mining and text analytics toolkits.
- Web-based text analytics.
- Case studies of current commercial applications in text mining, Beyond English, Arabic data, machine translation, information retrieval, information extraction, text classification.

9.16 University of Manchester MSc Advanced Computer Science: Artificial Intelligence

The University of Manchester run a twelve month taught MSc Advanced Computer Science: Artificial Intelligence. Some of the modules included in the course are described below.

Foundations of Machine Learning (15 credits)

- Classifiers and the Nearest Neighbour Rule
- Linear Models, Support Vector Machines
- Algorithm assessment - overfitting, generalisation, comparing two algorithms
- Decision Trees, Feature Selection, Mutual Information
- Probabilistic Classifiers and Bayes Theorem
- Combining Models - ensemble methods, mixtures of experts, boosting
- Feature Selection - basic methods, plus some tasters of research material

Data Engineering (15 credits)

- An overview of the data life cycle
- Data engineering, modelling and design techniques
- Data storage and warehousing
- Data access and maintenance
- Big Data, Map-Reduce, Hadoop
- Data analytics and visualisation
- Engineering non-traditional data types
- Data standards and data quality

Automated Reasoning and Verification (15 credits)

- Orderings, multi-set
- Propositional reasoning: Language of propositional logic, semantics, truth tables; Satisfiability, validity, equivalence, decidability; Normal forms, CNF, clauses, optimised normalisation; Propositional resolution; DPLL and SAT-solving with backjumping, lemma learning ;Logical modelling
- General first-order reasoning: Language of first-order logic, semantics ;Normal forms, clauses; Herbrand interpretations; Soundness, literal & clause orderings, saturation ; Model construction, Unification for general resolution; Basic general resolution, ordering & selection refinements; Redundancy elimination; Using SPASS (lab)
- Verification: LTL; bounded model checking

As part of their four year MEng course in Artificial Intelligence, the University of Manchester also run undergraduate modules on: Machine Learning and Optimisation, Symbolic AI, and Natural Language Systems.

9.17 University of Portsmouth MSc Data Analytics

The University of Portsmouth run a one year taught MSc in Data Analytics[76]. The course includes the following modules.

Applied Data And Text Analytics (semester 1):

The module is designed to give students the applied knowledge of significant data and text analytics methods. The students will also have hands-on experience using data and text mining toolkits, including how to visualise and interpret the results. The focus in the module will be on the application side and how data mining, as a tool, could be used in different applications. The set exercise assessment is comprised of several tasks (8 to 10), thus the students have approximately 400 to 500 words per task.

Data Management (semester 1):

The students will have an opportunity to learn about emerging trends in modern data management and gain the practical development skills needed in such an expanding field.

[76] https://www.port.ac.uk/study/courses/msc-data-analytics

This module introduces students to the modern data management and database techniques and encourages understanding of the techniques and methods currently used. The module will cover Data governance, Data Architecture, Analysis and Design and Data Security Management. Also, the module will go through the design and development of modern databases designing and development (e.g. NoSQL and cloud relational databases).

Big Data Applications (semester 2):

In the past few years, Data analytics and mining methods need to deal with large amounts of data (Big Data) to generate useful insights. The challenges with Big Data is not just associated with the “volume” of the data but also includes: velocity, variety and veracity. Traditional methods are inadequate to deal with such complex datasets, there are new computational architectures that focus on dealing with Big Data and scaling-out the data processing to meet with the increasing demand.

Business Intelligence (semester 2):

Business intelligence (BI) covers a wide of software applications used to analyse, discover new insights and visualise an organization's raw data. The module will introduce the students to the different types of financial data, reports and decision-making process. The module will teach the students how to effectively use software to visualise business data and use it for decision support. Also, the students will learn the main data warehouse concepts and use. Part of the module will take place in the Bloomberg Suite available at the Portsmouth Business School, this will expose the students to the same data, analytics and software used by financial professionals daily to make a fast-paced investment, trading and financial decisions. This will allow students to monitor and analyse real-time financial market data movements and trade virtually.

9.18 Queens University Belfast MSc Data Analytics

Queens University Belfast run a one year taught MSc in Data Analytics, Some of the modules included in the course are described below.

- Data Analytics Fundamentals
- Databases and Programming Fundamentals
- Data Mining
- Machine Learning
- Frontiers in Data Analytics
- Analytics in Action

9.19 Queen Mary – University of London MSc Artificial Intelligence

Queen Mary University of London run a two year taught MSc in Artificial Intelligence as a two years ‘Thick Sandwich’ degree. Some of the modules included in this course our given below.

Semester 1

- Machine Learning (15 credits)
- Data Mining (15 credits)
- Artificial Intelligence and Games (15 credits)

Further option:

- Introduction to Computer Vision (15 credits)
- Music Perception and Cognition (15 credits)
- Natural Language Processing (15 credits)

Semester 2

Four options from:

- Advanced Robotics Systems (15 credits)
- Music Analysis and Synthesis (15 credits)
- Information Retrieval (15 credits)
- Artificial Intelligence (15 credits)
- Music and Speech Modelling (15 credits)
- Deep Learning and Computer Vision (15 credits)
- Machine Learning for Visual Data Analysis (15 credits)
- Neural Networks and NLP (15 credits)
- Multi-platform Game Development (15 credits)

Semester 3
(compulsory)

- Project (60 credits)

Year 2
(compulsory)

- Industrial Placement Project

9.20 Royal Holloway University of London

Royal Holloway run MSc courses in Machine Learning with a Year in Industry, Data Science and Analytics, and Artificial Intelligence. Some of the modules included in these courses are described below.

9.20.1 MSc courses in Machine Learning with a Year in Industry

The following are the **Core Modules** for the MSc course[77]**:**

Data Analysis

- Algorithm-independent machine learning
- Unsupervised learning and clustering
- Exploratory data analysis

[77] https://www.royalholloway.ac.uk/studying-here/postgraduate/computer-science/machine-learning-with-a-year-in-industry/

- Bayesian methods
- Bayes networks and causality
- Information retrieval and natural language processing

Computation with Data

- Basics of algorithmic thinking and problem solving using programming.
- Object-oriented programming.
- Algorithmic tasks and evaluating programming solutions.

Programming for Data Analysis

- MATLAB (Matrix Laboratory) and WEKA (Waikato Environment for Knowledge Analysis) as tools for machine learning and data mining.
- For MATLAB, you will develop an understanding of how to input and output data using vectors, arrays and matrics; learn techniques in data visualization, including plots in 2 and 3 dimensions, scatter plots, barplots, and histograms; how to implement concepts from linear algebra and statistics, including probability and matrix decompositions.
- For WEKA, how to use the software as a tool for training and testing, predicting generalisation performance, and cross-validation; how to implement decision trees, naïve Bayes classifiers, and clustering methods.

Machine Learning

- Main advantages and limitations of the various approaches to machine learning and examine the features of specific machine-learning algorithms.
- Ideas and algorithms of machine learning applied in other fields, including medicine and industry.

On-line Machine Learning

- On-line framework of machine learning for issuing predictions or decisions in real-time.
- Protocols, methods and applications of on-line learning, covering probabilistic models based on Markov chains and their applications, such as PageRank and Markov Chain Monte-Carlo.
- Time series models, exploring their connections with Kalman filters, and learning models based on the prequential paradigm, including prediction with expert advice, aggregating algorithm, sleeping and switching experts.
- Universal algorithms, their application to portfolio theory, and how prediction within a confidence framework is achieved.

9.20.2 MSc in Artificial Intelligence, starting in 2019/20

The programme[78] will include modules in:

---

[78] https://www.royalholloway.ac.uk/research-and-teaching/departments-and-schools/computer-science/studying-here/postgraduate/

- Artificial Intelligence Principles and Techniques
- Natural Language Processing
- Autonomous Intelligent Systems
- Experimental Design
- Data Analysis
- Programming for Data Analysis
- Machine Learning
- Deep Learning
- Semantic Web
- Visualisation and Exploratory Analysis

9.20.3 MSc in Data Science and Analytics with a Year in Industry

The following are the **Core Modules** for the MSc course[79]**:**

Data Analysis (see above)

Computation with Data (see above)

Programming for Data Analysis (see above)

Database Systems

- Core concepts in data and information management, looking at the role of databases and database management systems in managing organisational data and information.
- Organisational information requirements, model them using conceptual data modelling techniques, convert the conceptual data models into relational data models and verify their structural characteristics using normalisation techniques.
- Designing and implementing a relational database using an industrial database management system, and examine how to manipulate data using SQL.

Large-scale Data Storage and Processing

- Underlying principles of large-scale data storage and processing frameworks.
- Building massive scale analytics solutions, gaining hands-on experience in using large and unstructured data sets for analysis and prediction.
- Techniques and paradigms for querying and processing massive data sets, such as MapReduce, Hadoop, data warehousing, SQL for data analytics, and stream processing.
- Fundamentals of scalable data storage, including NoSQL databases, design, develop, and evaluate an end-to-end analytics solution combining large-scale data storage and processing frameworks.

[79] https://www.royalholloway.ac.uk/studying-here/postgraduate/computer-science/data-science-and-analytics-with-a-year-in-industry/

9.21 University of Southampton MSc course in Artificial Intelligence

The University of Southampton runs a one year taught Artificial Intelligence MSc[80]. Some of the modules included in the course are described below.

Intelligent Agents (15 credits)

- Introduction to agent-based computing: Motivations for agent-based computing - Key concepts and models of reasoning (symbolic, reactive and practical) - Rational decision making and handling uncertainty
- Agent Interactions: Models of coordination (DCOP and the max-sum algorithm) - Models of competitive behaviour (game theory and mechanism design) - Computational markets (auctions)
- Agent design and implementation: Structuring agent models in code - Deploying agents within a simulated environment - Practical reasoning strategies for computational markets.

Foundations of Artificial Intelligence (15 credits)

- Introduction to AI: Flavours of AI - strong and weak, neat and scruffy, symbolic and sub-symbolic, knowledge-based and data-driven. The computational metaphor. What is computation? Church-Turing thesis. The Turing test. Searle's Chinese room argument.
- Search: Finding satisfactory paths: depth-first and breadth-first, iterative deepening, local search and heuristic search. Finding optimal paths: branch and bound, dynamic programming, A*.
- Representing Knowledge: Production rules, monotonic and non-monotonic logics, semantic nets, frames and scripts, description logics.
- Reasoning and Control: Data-driven and goal-driven reasoning.
- Reasoning under Uncertainty: Probabilities, conditional independence, causality, Bayesian networks, belief propagation.
- Machine Learning: Inductive and deductive learning, unsupervised and supervised learning, reinforcement learning, concept learning from examples, Quinlan's ID3, classification and regression trees, Bayesian methods.

Advanced Machine Learning (15 credits, optional)

- Key concepts: Supervised/Unsupervised Learning; Loss functions and generalization; Probability Theory; Elements of Computational Learning Theory
- Kernel Methods for non-linear data: Support Vector Machines; Kernel
- Bayesian methods for using prior knowledge and data: Bayesian inference; Bayesian Belief Networks and Graphical models; Gaussian Processes
- Ensemble Learning: Bagging; Boosting; Random Forest
- Introduction to Deep Learning

Machine Learning Technologies (15 credits, optional)

[80] https://www.ecs.soton.ac.uk/programmes/msc-artificial-intelligence#modules

- Historical Perspective:- Biological motivations: the McCulloch and Pitts neuron, Hebbian learning.
- Conceptual motivations
- Tools in machine learning: Libraries; Implementing and evaluating algorithms
- Supervised Learning: Perceptron Algorithm; Support Vector Classification and Regression; Neural networks/multi-layer perceptrons (MLP)
- Data handling and unsupervised learning: Principal Components Analysis (PCA); K-Means clustering
- Regression and Model-fitting Techniques: Linear regression
- Deep Learning: Deep Neural Networks (CNN, RNN)
- Example applications: Speech, Vision, Natural Language, Bioinformatics.

Deep Learning (15 credits, optional)

- Historical Developments: Deep Belief Networks; CNNs, LeNet and the ImageNet competition;- RNNs
- Learning Algorithms: Initialisation; SGD, Momentum, etc.
- Deep Belief Networks: RBMs
- Auto-encoders: variational; denoising
- CNNs: Architectures; Region Propositions; Semantic Segmentation
- Sequence Modelling: Linear Embeddings; RNNs; LSTMs; GRUs; back-prop through time
- Deep Learning Technologies
- Applications; Computer Vision; Natural Language Processing & Generation; Speech

9.22 University of St Andrews Artificial Intelligence MSc

The University of St Andrews runs an MSc in Artificial Intelligence[81] as an advanced research-led course in the study of artificial intelligence, developing students' skills in logic, constraint programming, language processing, machine learning and neural networks.

Some of the modules included in the course are as follows.

Object-Oriented Modelling

- Design and Programming: introduces and reinforces object-oriented modelling, design and implementation to provide a common basis of skills, allowing students to complete programming assignments within other MSc modules.

Artificial Intelligence Principles:

- foundational knowledge of artificial intelligence (AI) with an overview of AI and its philosophy.

Artificial Intelligence Practice:

[81] https://www.st-andrews.ac.uk/subjects/computer-science/artificial-intelligence-msc/#11728

- practical design and implementation of artificial intelligence (AI), covering techniques in the areas of AI reasoning, planning, doing and learning.

Some of the **optional modules** are:

- Artificial Intelligence in Practice
- Artificial Intelligence Principles
- Database Management Systems
- Data-Intensive Systems
- Knowledge Discovery and Datamining
- Language and Computation
- Software Architecture
- Software Engineering Practice
- Software Engineering Principles

9.23 University of Surrey MSc on Computer Vision, Robotics and Machine Learning

The University of Surrey[82] runs a one year taught MSc course on Computer Vision, Robotics and Machine Learning. Some of the modules included in the course are described below.

AI and AI Programming (15 credits)

- Historical Overview - Definition of artificial intelligence (AI); Application areas; General problem solving versus specific knowledge; Complexity.
- Heuristic Search - Uninformed versus informed search strategies; Formal properties of A*; Minimax game search, alpha-beta pruning.
- Logic and Resolution - Knowledge representation; Propositional and predicate calculus; Inference rules; Clause form; Resolution strategies; Prolog and logic programming.
- Uncertainty Reasoning - Probabilistic reasoning and Bayes theorem; Belief network; Dempster-Shafer theory; Fuzzy logic.
- Basic Prolog - execution model; declarative and procedural meaning; backtracking; arithmetic; list representation; negation as failure and difficulties; simple examples.
- Prolog Programming and Techniques - input/output; meta-logical and extra-logical predicates; set predicates; cuts, program development and style; correctness and completeness; Applications
- Multi-Layer Perceptrons - Convergence theorem; non-separability; LMS algorithms; steepest descent; back-propagation; generalisation; learning factors.
- Radial Basis Function Networks - Multivariable interpolation; regularisation; comparison with MLP; learning strategies.
- Self-Organising Systems - Hebbian learning; competitive learning; SOFM; LVQ.
- Recurrent networks - energy functions; Hopfield net; nonlinear dynamical systems; Liapunov stability; attractors.

[82] https://www.surrey.ac.uk/postgraduate/computer-vision-robotics-and-machine-learning-msc-2018

Image Processing and Deep Learning (15 credits)

- Image Representation and Colour; Geometric Image Transformations; Homogeneous Coordinates.
- Image Warping; Interpolation; Image Quality metrics.
- Image Filtering.; Aliasing; Blurring, Sharping and Edge Detection; Gaussian kernel and its derivatives.; Scale-space pyramids; Thresholding.
- Image Completion and Enhancement; Brightness and Contrast; Poisson Image Editing; Patch based in-painting; Super-resolution.
- Classical object recognition; Interest points, Gradient domain features and Bag of Words.
- Convolutional Neural Networks (CNNs); Siamese Networks; CNN Interpretatability.
- Fully convolutional networks; Up-convolution.; Optical flow; FCNs for Stylization; Segmentation.
- Genereative Adversarial Networks (GANs); Image transformation networks; Super-resolution.

Computer Vision and Pattern Recognition (15 credits)

- Image Processing
- Pattern classification
- Features and Matching
- Shape Description
- Tracking
- Contour models
- Multiview Geometry

Space Robotics and Autonomy (15 credits)

- Manipulator Kinematics
- Manipulator Inverse Kinematics
- Manipulator Differential Kinematics and Space Freeflyer Kinematics.
- Manipulator Dynamics: Mass distribution, Inertia tensor, Parallel axis theorem, Holonomic & non-holonomic systems, Introduction to Lagrange-Euler method and Newton-Euler method.
- Manipulator Motion Control
- Rover Navigation System: introduction to rover navigation problem and major system architectural designs (hieratical, reactive and hybrid). Major navigation functions, localization challenges & strategies, map making and representation, metric path planning, topological path planning, planning algorithms such as A*/D*
- Rover Sensing & Perception: Classification of sensors (e.g. proprioceptive vs. exteroceptive; and passive vs. active), sensor properties, motor sensors, heading sensors, ranging sensors, vision sensors, stereovision, vision processing techniques.

Speech and Audio Processing and Recognition (15 credits)

- Digital speech processing.

- Speech Production - Vocal tract description.
- Speech Perception - The structure of the ear.
- Signal Processing Techniques - Autocorrelation of speech signals.
- Linear Prediction – Z-transform.
- Inverse Filtering of Speech Signal - Separating source from excitation.
- Cepstral Deconvolution- Definition of real cepstrum.
- Audio recording and acoustics – Microphone types and directivity patterns, digital audio acquisition, wave propagation and acoustics, effects of reflections and reverberation.
- Psychoacoustics –Loudness perception, pitch perception, auditory masking, timbre perception, spatial hearing.

9.24 University of Sussex MSc Intelligent & Adaptive Systems

The University of Sussex offers an MSc in Intelligent & Adaptive Systems, which has a significant ML/AI component[83]. The course prepares graduates for research and development in intelligent and adaptive systems, covering theoretical issues and practical techniques for their design and implementation. You course allows students to organise studies around the themes: artificial intelligence and cognitive modelling, robotics and autonomous systems, data science, machine learning and natural language processing, computational biology and consciousness science.

Some of the modules offered in the course are described below.

Adaptive Systems (15 credits)

- cybernetics
- control theory
- self-organisation
- autonomous robotics
- evolutionary and developmental robotics
- dynamical systems approaches to embodied cognition

Machine Learning (15 credits)

- probabilistic and non-probabilistic classification and regression methods
- reinforcement learning approaches including the non-linear variants using kernel methods
- techniques for pre-processing the data (including PCA).

Advanced Natural Language Processing (15 credits)

- word sense disambiguation
- vector space models of semantics
- named entity recognition
- topic modelling
- machine translation

[83] https://www.sussex.ac.uk/study/masters/courses/intelligent-and-adaptive-systems-msc

- hypothesis testing
- data smoothing techniques
- domain adaptation
- generative versus discriminative learning
- semi-supervised learning

Applied Natural Language Processing (15 credits)

- tokenisation
- segmentation
- stemming
- lemmatisation
- part-of-speech tagging
- named entity recognition
- phrasal chunking
- dependency parsing
- document classification
- information retrieval
- information extraction

9.25 University of Swansea MSc course on Data Science

The University of Swansea runs a one and two year taught MSc course on Data Science[84], including machine learning techniques. Some of the modules included in the course are described below.

Big Data and Machine Learning (15 credits)

- Introduction to big data and data mining;
- Data clustering;
- Dimensionality reduction: linear techniques;
- Dimensionality reduction: nonlinear techniques;
- Discriminative analysis;
- Learning theory, including bias and variance theory, innovation process in machine learning;
- Expert systems;
- Unsupervised learning;
- Supervised learning, including parametric and nonparametric methods, neural network, kernels, support vector machine, randomised decision trees;
- Reinforcement and adaptive control;
- Example applications to bioinformatics, health informatics, and web data processing.

Big Data and Data Mining (15 credits)

- Goals of data mining

[84] http://www.swansea.ac.uk/postgraduate/taught/science/mscdatascience/#modules=is-expanded&year-1-level-7-pgt-pi100tf1=is-expanded&year-2-level-7-pgt-pi100tf1=is-expanded&description=is-expanded

- Data and challenges
- Methodology
- Data preparation
- Pattern analysis
- Big data requirements: efficiency and scalability
- Deep learning
- Hadoop

Mathematical Skills for Data Scientists (15 credits)

- Vectors and matrices
- Derivatives and partial derivatives
- Variational calculus (fundamentals)
- Gradient descent
- Least Squares
- Fundamentals of probability
- Standard deviation, variance and covariance
- Bayesian Theorem
- Eigenvalues and eigenvectors, PCA
- Gaussian distribution, T-distribution
- Cross correlation, Chi-square, mahalanobis distance
- Fourier analysis

Visual Analytics (15 credits)

- History and goals of visual analytics.
- Types of data and encodings.
- Data processing and clustering.
- Information visualisation techniques.
- The analytics process and pipeline.

Optional modules in the course (of which three must be chosen, each is worth 15 credits) are:

- Computer Vision and Deep Learning
- Modelling and Verification Techniques
- Data Visualization
- Human Computer Interaction
- High Performance Computing in C/C++
- Graphics Processor Programming
- Operating Systems and Architectures

Computer Vision and Deep Learning is most relevant of those options to AI, and covers:

- Image processing: filtering, object extraction, segmentation, texture analysis.

- Video analysis: camera models and calibration, stereo vision, depth estimation, motion estimation and tracking, local features for tracking.
- Neural networks and Deep Learning: feedforward neural networks, back propagation, convolutional neural network, recurrent neural network, and applications

The second year of the MSc consists of a **research project**. The dissertation covers:

- Discussion of the subject area and its history;
- A literature survey;
- Formulation of scientific questions and the answers to them;
- Theoretical background;
- Description of the approach taken;
- Discussion of issues arising in the undertaking of the project;
- Evaluation of results;
- Progress and achievements of the project;
- Suggestions for further work.

9.26 UCL – University College London

UCL run three MSc courses related to Machine Learning: MSc Computational Statistics and Machine Learning[85], MSc Data Science and Machine Learning[86], and MSc Machine Learning[87].

9.26.1 MSc Machine Learning

Some of the modules included in the course are described below.

Supervised Learning (15 credits)

- Nearest Neighbours
- Linear Regression
- Kernels and Regularisation
- Support Vector Machines
- Gaussian Processes
- Decision Trees
- Ensemble Learning
- Sparsity Methods
- Multi-task Learning
- Proximal Methods
- Semi-supervised Learning
- Neural Networks
- Matrix Factorization
- Online Learning

[85] http://www.cs.ucl.ac.uk/prospective_students/msc_computational_statistics_and_machine_learning/
[86] http://www.cs.ucl.ac.uk/prospective_students/msc_data_science_and_machine_learning/
[87] http://www.cs.ucl.ac.uk/prospective_students/msc_machine_learning/

- Statistical Learning Theory

Advanced Deep Learning and Reinforcement Learning (15 credits)

- The basics of deep learning and reinforcement learning paradigms
- Architectures and optimization methods for deep neural network training
- How to implement deep learning methods within TensorFlow and apply them to data
- The theoretical foundations and algorithms of reinforcement learning
- How to apply reinforcement learning algorithms to environments with complex dynamics

Advanced Topics in Machine Learning (15 credits)

- Definition of a kernel, how it relates to a feature space, The reproducing kernel Hilbert space
- Simple applications: kernel PCA, kernel ridge regression
- Distance between means in RKHS, integral probability metrics, the maximum mean discrepancy (MMD), two-sample tests
- Choice of kernels for distinguishing distributions, characteristic kernels
- Covariance operator in RKHS: proof of existence, definition of norms (including HSIC, the Hilbert-Schmidt independence criterion)
- Application of HSIC to independence testing
- Feature selection, taxonomy discovery.
- Introduction to independent component analysis, kernel ICA
- Large margin classification, support vector machines for classification
- Introduction to supervised learning in the context of statistical learning theory: a taxonomy of learning problems; no free lunch theorem; regularization; model selection; stability and generalization; measures of complexity for hypotheses spaces; sample complexity; generalization bounds

Affective Computing and Human-Robot Interaction (15 credits)

- Emotion theory; What is affect, emotion, mood? Why do we have emotions? Neurological and psychological perspectives; How do humans express and recognise emotions? Emotion expression models, appraisal and causal theories; Affective and social interaction
- Affective computing
- Emotion Recognition: Application of machine learning techniques for adaptive emotion recognition from single modality e.g. facial expressions, biosignals; Adaptive multimodal emotion recognition: signal fusion
- Human-Robot Interaction (HRI): Social robotics: motivation and emotions in robots; Emotion based architecture; Evaluation methods for HRI research; Ethical issues in Affective Computing and HRI research

Applied Machine Learning (15 credits)

- First Order Optimisation methods (gradient descent)

- Second Order Optimisation methods (Newton and Quasi Newton approaches and Conjugate Gradients)
- Methods for solving Large Scale Linear, including Conjugate Gradients
- Automatic Differentiation methods for efficiently computing first and second order gradients
- Classical methods for Regression and Classification including linear and logistic regression
- Methods for Unsupervised Learning including mixture modelling
- Deep Learning Methods for Regression, Classification and Unsupervised Learning
- Recurrent Networks for Time-Series processing
- Matrix and Tensor Factorisation
- Visualisation methods including Autoencoders and tSNE

Approximate Inference and Learning in Probabilistic Models (15 credits)

- Nonlinear, hierarchical (deep), and distributed models.
- Independent component analysis, Boltzmann machines, Dirichlet topic models, manifold discovery.
- Mean-field methods, variational approximations and variational Bayes.
- Expectation propagation
- Loopy belief propagation, the Bethe free energy and extensions.
- Convex methods and convexified bounds.
- Monte-Carlo methods: including rejection and importance sampling, Gibbs, Metropolis-Hastings, anealed importance sampling, Hamiltonian Monte-Carlo, slice sampling, sequential Monte-Carlo (particle filtering).

Introduction to Deep Learning (15 credits)

- Neural networks
- Convolutional networks
- Recurrent networks
- Introduction to dropout, batch normalization, types of hyper-parameter optimization, distributed and constrained computing variants.
- Applications in the area of audio processing and image captioning and vision.

Probabilistic and Unsupervised Learning (15 credits)

- Basics of Bayesian learning and regression.
- Latent variable models, including mixture models and factor models.
- The Expectation-Maximisation (EM) algorithm.
- Time series, including hidden Markov models and state-space models.
- Spectral learning.
- Graphical representations of probabilistic models.
- Belief propagation, junction trees and message passing.
- Model selection, hyperparameter optimisation and Gaussian-process regression.

Statistical Natural Language Processing (15 credits)

- NLP Tasks
    - Language Models
    - Machine Translation
    - Text Classification
    - Sequence Tagging
    - Constituency Parsing
    - Dependency Parsing
    - Information Extraction
    - Machine Comprehension
- NLP and ML methods
    - Structured Prediction
    - Generative Learning
    - Smoothing
    - EM Algorithm
    - Discriminative Learning
    - Deep and Representation Learning

9.26.2 MSc Data Science and Machine Learning

Some of the modules this MSc course includes as options are: Applied Machine Learning (see above), Introduction to Machine Learning (see above), Advanced Deep Learning and Reinforcement Learning (see above), Statistical Natural Language Processing (see above), Supervised Learning (see above).

9.26.3 MSc Computational Statistics and Machine Learning

Some of the modules this MSc course includes as options are: Supervised Learning (see above), Advanced Deep Learning and Reinforcement Learning (see above), Advanced Topics in Machine Learning (see above), Applied Machine Learning (see above), Introduction to Deep Learning (see above), Probabilistic and Unsupervised Learning (see above), Statistical Natural Language Processing (see above).

9.27 Ulster University MSc in Data Science

Ulster University offer an MSc in Data Science[88] (part time). The MSc aims to provide a firm grounding in the core disciplines of data analytics and information processing, partnered with a broad appreciation of aspects of other disciplines where data science can form natural synergistic relationships.

The MSc aims to prepare graduates for a career as a data scientist or business analyst working in any profession where large amounts of data is collected, hence there is a need for skills in data acquisition, information extraction, aggregation and representation, data analysis, knowledge extraction and explanation. These types of skills are intended to support progression into employment in areas such as IT business, security and health sectors, intelligent transport, energy efficiency and the creative industries.

Some of the modules included in the course are described below.

[88] https://www.ulster.ac.uk/courses/201920/data-science-19885

Data Validation and Visualisation

This module introduces the data quality challenges faced by big data. It will present tools and techniques employed to ensure data quality from data collection and computational procedures to facilitate automatic or semi-automatic identification and elimination of errors in large datasets. The module also introduces the topic of understanding and interpreting data through descriptive statistical methods. This will be achieved through a range of techniques such as Statistical metrics, Univariate analysis and Multivariate analysis. Students will develop the knowledge to assess the quality of the data and the skills necessary to perform appropriate data cleaning operations. In addition, students will have an understanding of processing data and interpreting and visualising results.

Machine Learning and Data Modelling

This module covers Machine Learning both conceptually and practically. Students will be introduced to a variety of unsupervised and supervised Machine Learning techniques. Once the core concepts have been introduced they will be given practical experience of their use, application and evaluation through laboratory exercises and a project. The students will develop an in-depth understanding of the potential and scope of applying and evaluating the different forms of Machine Learning. This will allow them to develop a range of applications from simple practical implementations to large scale implementations.

Data Science Foundations

The focus of this module is to present an understanding of key data science concepts, tools and programming techniques. Within the arena of data science, the theory behind the approaches of statistics, modelling and machine learning will be introduced emphasising their importance and application to data analysis. The notion of investigative and research skills will also be introduced through a number of problem solving exercises.

Statistical Modelling and Data Mining

On completing this module, students will be able to compute conditional probabilities and use null hypothesis significance testing to test the significance of results, and understand and compute statistical measures such as the p-value for these tests. Students will be able to apply, evaluate and critically appraise this knowledge in a range of complex real world contexts.

9.28 University of Warwick

The University of Warwick offers an undergraduate degree in Data Science, which is planned to be converted into an MEng degree in the near future. Of particular relevance to this report the university offers an MSc in Data Analytics[89]. Some of the modules included in that course are described below.

Data Mining (15 credits)

- Data pre-processing: handling missing values, basic data transformations.

[89] https://warwick.ac.uk/study/postgraduate/courses-2019/da/

- Rule induction; decision trees; naïve Bayesian probability; neural networks.
- Image processing
- Perceptron and support vector machines.
- Ensemble methods: boosting, bagging & random forests.
- Evaluation: cross validation, ROC.
- Lazy learning: clustering and rule mining; association rule mining.
- Time series.
- Text mining with feature engineering; vector space models.
- Graph mining.

Foundations of Data Analytics (15 credits)

- Introduction to analytics and case studies
- Basic tools - including unix/linux command line tools for data manipulation (sorting, counting, reformatting, aggregating, joining)
- Statistics - the tools from statistics for understanding distributions and probability (means, variance, tail bounds). Hypothesis testing for determining the significance of an observation, and the R system for working with statistical data;
- Databases – including problems found in realistic data: errors, missing values, lack of consistency, and techniques for addressing them. The relational data model, and the SQL language for expressing queries. The NoSQL movement, and the systems evolving around it;
- Regression - predicting new data values via regression models. Simple linear regression over low dimensional data, regression for higher dimensional data via least squares optimization, logistic regression for categoric data;
- Matrices - Matrices to represent relations between data, and necessary linear algerbraic operations on matrices. Approximately representing matrices by decompositions (Singular Value Decomposition and Principal Components Analysis);
- Clustering - Finding clusters in data via different approaches. Choosing distance metrics. Different clustering approaches: hierarchical agglomerative clustering, k-means (Lloyd's algorithm), k-center approximations;
- Classification - Building models to classify new data instances. Decision tree approaches and Naive Bayes classifiers. The Support Vector Machines model and use of Kernels to produce separable data and non-linear classification boundaries. The Weka toolkit;
- Data Structures - Data structures to scale analytics to big data and data streams. The Bloom filter to represent large set values. Sketch data structures for more complex data analysis, and other summary data structures;
- Data Sharing - The ethics and risks of sharing data on individuals. Technologies for anonymising data: k-anonymity, and differential privacy;
- Graphs - Graph representations of data, with applications to social network data. Measurements of centrality and importance. Recommendations in social networks, and inference via relational learning;

Optional modules available for the MSc include:

- High Performance Computing
- Algorithmic Game Theory
- Image and Video Analysis
- Sensor Networks and Mobile Data Communications
- Advanced Computer Security
- Social Informatics
- Natural Language Processing

Of those modules Natural Language Processing is of particular relevance for this report.

Natural Language Processing (15 credits)

- Regular expressions, word tokenization, stemming, sentence segmentation
- N-grams and language models
- Part-of-speech Tagging
- Hidden Markov models and maximum entropy models
- Semantics: lexical semantics, distributional semantics, word sense disambiguation and vector space models
- Spelling correction
- Text classification
- Sentiment analysis
- Information extraction: Named entity recognition, relation extraction
- Information retrieval
- Syntactic parsing
- Semantic parsing
- Question answering and summarisation
- Text processing in social media

10 Appendix: An incomplete body of knowledge for Machine Learning

To test if there was something approaching a common body of knowledge in MSC Machine Learning courses that is appropriate to address the canonical Machine Learning problems we mapped some exemplary curricula descriptions against two standard text books. These are:

- 'Introduction to Machine Learning', Ethem Alpaydin, Adaptive computation and machine learning series, 2014, MIT Press, ISBN-10 0262325748, ISBN-13 9780262028189, 9780262325745. At the time of writing the second edition has been cited 4485 according to Google Scholar[90].
- 'Reinforcement learning: An introduction', Richard Sutton and Andrew Barto, MIT Press, 1998, ISBN-10 0262193981, ISBN-13 9780262193986. New addition available

[90] https://scholar.google.co.uk/scholar?q=Introduction+to+Machine+Learning+Alpaydin

in 2018. At the time of writing this text has been cited 27468 according to Google Scholar[91].

Some responses to our consultations with HEIs pointed out these two texts should not be regarded as defining a comprehensive overview of Machine Learning topics. We would like to thank those who responded with suggested alternative texts. However, there wasn't a significant overlap between the various suggested alternatives, and it seemed unlikely that an analysis based on a new set of text books would significantly alter the conclusion of this report. It would be very interesting to conduct a further survey of MSc courses using the complete list of text books suggested by the consultations, but that is beyond the remit for this report.

The universities we chose as exemplars are Bath, Edinburgh, Essex and UCL. This was based on the clarity and extensiveness of the publicly available curricula descriptions, and should not in any way be interpreted as suggesting other universities either do or do not offer an exemplary MSc. The following table gives our subjective expert opinion of how clearly the main topics in these standard texts can be mapped to the descriptions in the exemplary curriculum.

Machine Learning areas of knowledge	**Extent to which areas visibly signposted in course curriculum**			
	Bath	**Edinburgh**	**Essex**	**UCL**
Supervised Learning	Extensive	Extensive	Extensive	Extensive
Bayesian Decision Theory	Extensive	Extensive	Extensive	Extensive
Parametric Methods	Extensive	In Part	Extensive	Extensive
Multivariate Methods	In Part	In Part	In Part	In Part
Dimensionality Reduction	In Part	Extensive	In Part	Extensive
Clustering	In Part	Extensive	In Part	In Part
Nonparametric Methods	In Part	Extensive	In Part	In Part
Decision Trees	Extensive	In Part	Extensive	Extensive
Linear Discrimination	In Part	In Part	Not Visible	In Part
Neural Networks	Extensive	Extensive	Extensive	Extensive
Local Models	Not Visible	In Part	In Part	Not Visible
Kernel Machines	Not Visible	Extensive	Extensive	Extensive
Graphical Models	Extensive	Not Visible	Not Visible	Not Visible
Hidden Markov Models	Extensive	Extensive	Not Visible	Extensive
Bayesian Estimation	In Part	Extensive	In Part	Extensive
Combining Multiple Learners	Extensive	Extensive	Extensive	Extensive
Reinforcement Learning	Extensive	Extensive	Extensive	Extensive

[91] https://scholar.google.co.uk/scholar?q=Reinforcement+Learning+An+Introduction+Sutton+Barto

Design and Analysis of Machine Learning Experiments	Not Visible	Not Visible	Not Visible	Not Visible

Table 7: Mapping syllabi content against areas of knowledge

The complete list of domain knowledge areas identified in the standard text books are listed in the following tables:

Supervised Learning	
	Learning a Class from Examples
	Vapnik-Chervonenkis Dimension
	Probably Approximately Correct Learning
	Noise
	Learning Multiple Classes
	Regression
	Model Selection and Generalization
	Dimensions of a Supervised Machine Learning Algorithm

Bayesian Decision Theory	
	Classification
	Losses and Risks
	Discriminant Functions
	Association Rules

Parametric Methods	
	Maximum Likelihood Estimation
	Bernoulli Density
	Multinomial Density
	Gaussian Density
	Evaluating an Estimator: Bias and Variance
	The Bayes' Estimator
	Parametric Classification
	Regression
	Tuning Model Complexity: Bias/Variance Dilemma
	Model Selection Procedures

Multivariate Methods	
	Multivariate Data
	Parameter Estimation
	Estimation of Missing Values
	Multivariate Normal Distribution
	Multivariate Classification
	Tuning Complexity
	Discrete Features

	Multivariate Regression

Dimensionality Reduction	
	Subset Selection
	Principal Component Analysis
	Independent Component Analysis
	Feature Embedding
	Factor Analysis
	Singular Value Decomposition and Matrix Factorization
	Multidimensional Scaling
	Linear Discriminant Analysis
	Canonical Correlation Analysis
	Isomap
	Locally Linear Embedding
	Laplacian Eigenmaps

Clustering	
	Mixture Densities
	k-Means Clustering
	Expectation-Maximization Algorithm
	Mixtures of Latent Variable Models
	Supervised Learning after Clustering
	Spectral Clustering
	Hierarchical Clustering
	Choosing the Number of Clusters

Nonparametric Methods	
	Nonparametric Density Estimation
	Histogram Estimator
	Kernel Estimator
	k-Nearest Neighbor Estimator
	Generalization to Multivariate Data
	Nonparametric Classification
	Condensed Nearest Neighbour
	Distance-Based Classification
	Outlier Detection
	Nonparametric Regression: Smoothing Models
	Running Mean Smoother
	Kernel Smoother
	Running Line Smoother
	How to Choose the Smoothing Parameter

Decision Trees	
	Univariate Trees
	Classification Trees
	Regression Trees
	Pruning
	Rule Extraction from Trees
	Learning Rules from Data
	Multivariate Trees

Linear Discrimination	
	Generalizing the Linear Model
	Geometry of the Linear Discriminant
	Two Classes
	Multiple Classes
	Pairwise Separation
	Parametric Discrimination Revisited
	Gradient Descent
	Logistic Discrimination
	Discrimination by Regression
	Learning to Rank

Neural Networks	
	The Perceptron
	Training a Perceptron
	Learning Boolean Functions
	Multilayer Perceptrons
	MLP as a Universal Approximator
	Backpropagation Algorithm
	Nonlinear Regression
	Two-Class Discrimination
	Multiclass Discrimination
	Multiple Hidden Layers
	Training Procedures
	Improving Convergence
	Overtraining
	Structuring the Network
	Hints
	Tuning the Network Size
	Bayesian View of Learning
	Dimensionality Reduction
	Learning Time
	Time Delay Neural Networks

	Recurrent Networks
	Deep Learning

Local Models	
	Competitive Learning
	Online k-Means
	Adaptive Resonance Theory
	Self-Organizing Maps
	Radial Basis Functions
	Incorporating Rule-Based Knowledge
	Normalized Basis Functions
	Competitive Basis Functions
	Learning Vector Quantization
	The Mixture of Experts
	Cooperative Experts
	Competitive Experts
	Hierarchical Mixture of Experts

Kernel Machines	
	Optimal Separating Hyperplane
	The Nonseparable Case: Soft Margin Hyperplane
	v-SVM
	Kernel Trick
	Vectorial Kernels
	Defining Kernels
	Multiple Kernel Learning
	Multiclass Kernel Machines
	Kernel Machines for Regression
	Kernel Machines for Ranking
	One-Class Kernel Machines
	Large Margin Nearest Neighbour Classifier
	Kernel Dimensionality Reduction

Graphical Models	
	Canonical Cases for Conditional Independence
	Generative Models
	d-Separation
	Belief Propagation
	Chains
	Trees
	Polytrees
	Junction Trees
	Undirected Graphs: Markov Random Fields

	Learning the Structure of a Graphical Model
	Influence Diagrams

Hidden Markov Models	
	Discrete Markov Processes
	Hidden Markov Models
	Evaluation Problem
	Finding the State Sequence
	Learning Model Parameters
	Continuous Observations
	The HMM as a Graphical Model
	Model Selection in HMMs

Bayesian Estimation	
	Bayesian Estimation of the Parameters of a Discrete
	Dirichlet Distribution
	Beta Distribution
	Bayesian Estimation of the Parameters of a Gaussian
	Univariate Case: Unknown Mean, Known
	Univariate Case: Unknown Mean, Unknown Variance
	Multivariate Case: Unknown Mean, Unknown Covariance
	Bayesian Estimation of the Parameters of a Function
	Regression
	Regression with Prior on Noise Precision
	The Use of Basis/Kernel Functions
	Bayesian Classification
	Choosing a Prior
	Bayesian Model Comparison
	Bayesian Estimation of a Mixture Model
	Nonparametric Bayesian Modeling
	Gaussian Processes
	Dirichlet Processes and Chinese Restaurants
	Latent Dirichlet Allocation
	Beta Processes and Indian Buffets

Combining Multiple Learners	
	Generating Diverse Learners
	Model Combination Schemes
	Voting
	Error-Correcting Output Codes
	Bagging
	Boosting

	Stacked Generalization
	Fine-Tuning an Ensemble
	Choosing a Subset of the Ensemble
	Constructing Metalearners
	Cascading

Reinforcement Learning	
	Dynamic Programming
	Monte Carlo methods
	Policy gradient methods
	Temporal-Difference Learning
	K-Armed Bandit
	Model-Based Learning
	Value Iteration
	Policy Iteration
	Temporal Difference Learning
	Exploration Strategies
	Deterministic Rewards and Actions
	Nondeterministic Rewards and Actions
	Eligibility Traces
	Generalization
	Partially Observable States

Design and Analysis of Machine Learning Experiments	
	Factors, Response, and Strategy of Experimentation
	Response Surface Design
	Randomization, Replication, and Blocking
	Guidelines for Machine Learning Experiments
	Cross-Validation and Resampling Methods
	Cross-Validation
	Bootstrapping
	Measuring Classifier Performance
	Interval Estimation
	Hypothesis Testing
	Assessing a Classification Algorithm's Performance
	Binomial Test
	Approximate Normal Test
	t Test
	Comparing Two Classification Algorithms
	McNemar's Test
	K-Fold Cross-Validated Paired t Test
	Comparing Multiple Algorithms: Analysis of Variance

	Comparison over Multiple Datasets
	Multivariate Tests