

MASTER'S PROGRAMME
ARTIFICIAL INTELLIGENCE
FACULTY OF SCIENCE
UNIVERSITY OF AMSTERDAM

QANU
Catharijnesingel 56
PO Box 8035
3503 RA Utrecht
The Netherlands

Phone: +31 (0) 30 230 3100
E-mail: support@qanu.nl
Internet: www.qanu.nl

Project number: Q0744

© 2020 QANU

Text and numerical material from this publication may be reproduced in print, by photocopying or by any other means with the permission of QANU if the source is mentioned.



CONTENTS

REPORT ON THE MASTER’S PROGRAMME ARTIFICIAL INTELLIGENCE OF THE UNIVERSITY OF AMSTERDAM..... 5

ADMINISTRATIVE DATA REGARDING THE PROGRAMME..... 5

ADMINISTRATIVE DATA REGARDING THE INSTITUTION..... 5

COMPOSITION OF THE ASSESSMENT PANEL 5

WORKING METHOD OF THE ASSESSMENT PANEL 6

SUMMARY JUDGEMENT..... 9

DESCRIPTION OF THE STANDARDS FROM THE ASSESSMENT FRAMEWORK FOR LIMITED FRAMEWORK ASSESSMENTS..... 11

APPENDICES 19

APPENDIX 1: DOMAIN-SPECIFIC FRAMEWORK OF REFERENCE 21

APPENDIX 2: INTENDED LEARNING OUTCOMES 37

APPENDIX 3: OVERVIEW OF THE CURRICULUM 42

APPENDIX 4: PROGRAMME OF THE SITE VISIT 43

APPENDIX 5: THESES AND DOCUMENTS STUDIED BY THE PANEL 44

This report was finalised on 24 June 2020

REPORT ON THE MASTER'S PROGRAMME ARTIFICIAL INTELLIGENCE OF THE UNIVERSITY OF AMSTERDAM

This report takes the NVAO's Assessment Framework for the Higher Education Accreditation System of the Netherlands for limited programme assessments as a starting point (September 2018).

ADMINISTRATIVE DATA REGARDING THE PROGRAMME

Master's programme Artificial Intelligence

Name of the programme:	Artificial Intelligence
CROHO number:	56981
Level of the programme:	Master
Orientation of the programme:	Academic
Number of credits:	120 EC
Specialisations or tracks:	-
Location(s):	Amsterdam
Mode(s) of study:	Full-time
Language of instruction:	English
Submission deadline NVAO:	1 May 2020

The visit of the assessment panel Artificial Intelligence to the Faculty of Science of the University of Amsterdam took place on 9 and 10 December 2019.

ADMINISTRATIVE DATA REGARDING THE INSTITUTION

Name of the institution:	University of Amsterdam
Status of the institution:	Publicly funded institution
Result institutional quality assurance assessment:	Positive

COMPOSITION OF THE ASSESSMENT PANEL

The NVAO has approved the composition of the panel on 25 September 2019. The panel that assessed the master's programme Artificial Intelligence consisted of:

- Prof. dr. A. (Ann) Nowé . Professor at the Computer Science Department of the Faculty of Science and the Computer Science group of the Engineering Faculty at the Vrije Universiteit Brussel [chair].
- Prof. dr. C. (Cees) Witteveen. Full Professor at the Algorithmics Group of the Faculty of Engineering, Mathematics and Computer Science (EEMCS/EWI), Delft University of Technology.
- Prof. dr. B. (Bart) de Boer. Researcher and professor at the Artificial Intelligence lab of the Vrije Universiteit Brussel.
- Dr. ir. E.D. (Edwin) de Jong. Principal Machine Learning Scientist at ScreenPoint Medical Nijmegen and mentor/coaching A.I. startups at RockStart.
- M. (Maartje) Stokvis MSc. Master graduate Data Science for Decision Making at Maastricht University [student member].

The panel was supported by drs. José van Zwieten, who acted as secretary.

WORKING METHOD OF THE ASSESSMENT PANEL

The site visit to the master's programme Artificial Intelligence at the Faculty of Science of the University of Amsterdam was part of the cluster assessment Artificial Intelligence. Between October 2019 and January 2020 the panel assessed 14 programmes at 6 universities. The following universities participated in this cluster assessment: University of Groningen, Maastricht University / Transnational University Limburg, Radboud University, University of Amsterdam, Vrije Universiteit Amsterdam and Utrecht University.

On behalf of the participating universities, quality assurance agency QANU was responsible for logistical support, panel guidance and the production of the reports. Dr. Marijn Hollestelle was project coordinator for QANU. Marijn Hollestelle, Barbara van Balen, Peter Hildering and José van Zwieten acted as secretaries in the cluster assessment.

During the site visit at the University of Amsterdam, the panel was supported by José van Zwieten, a certified NVAO secretary.

Panel members

The members of the assessment panel were selected based on their expertise, availability and independence. The panel consisted of the following members:

- Prof. dr. A. (Ann) Nowé. Professor at the Computer Science Department of the Faculty of Science and the Computer Science group of the Engineering Faculty at the Vrije Universiteit Brussel [chair];
- Prof. dr. C. (Cees) Witteveen. Full professor at the Algorithmics Group of the Faculty of Engineering, Mathematics and Computer Science, Delft University of Technology;
- Prof. dr. B. (Bart) de Boer. Researcher and professor at the Artificial Intelligence lab of the Vrije Universiteit Brussel;
- Prof. dr. W. (Wiebe) van der Hoek. Interim Executive Pro Vice Chancellor and professor at the Department of Computer Science of the University of Liverpool;
- Prof. dr. F. (Frank) Jäkel. Principal Investigator at the Centre for Cognitive Science of the Technical University Darmstadt;
- Dr. ir. J.F.M. (Hans) Tonino. Associate Professor at the Algorithmics Group of the Faculty of Engineering, Mathematics and Computer Science and Director of Studies Embedded Systems at Delft University of Technology;
- Dr. ir. E.D. (Edwin) de Jong. Principal Machine Learning Scientist at ScreenPoint Medical and coach A.I. startups at RockStart;
- Dr. C.H.M. (Kees) Nieuwenhuis. Technology Manager and member of the bureau of the Chief Technology Officer at Thales Netherlands;
- Dr. A. (Annerieke) Heuvelink-Marck. Senior Scientist Software Concepts at Philips Group Innovation and Research;
- Dr. A.P. (André) Meyer-Vitali. Senior Scientist Data Science group ICT department at the Netherlands Organisation for applied scientific research (TNO);
- M. (Maartje) Stokvis MSc. Master graduate Data Science for Decision Making at Maastricht University [student member];
- F. (Florence) van der Voort BSc. Master student Artificial Intelligence and master student Philosophy: Bioethics and Health at the Vrije Universiteit Amsterdam [student member].

Preparation

On 10 May 2019, the panel chair was briefed by QANU on her role, the assessment framework, the working method, and the planning of site visits and reports. A preparatory panel meeting was organised on 28 August 2019. During this meeting, the panel members received instruction on the use of the assessment frameworks. The panel also discussed their working method and the planning of the site visits and reports.

The project coordinator composed a schedule for the site visit in consultation with the Faculty of Science. Prior to the site visit, the Faculty selected representative partners for the various interviews. See Appendix 4 for the final schedule.

Before the site visit to the University of Amsterdam, QANU received the self-evaluation reports of the programmes and sent these to the panel. A thesis selection was made by the panel's chair and the project coordinator. The selection consisted of 15 theses and their assessment forms for the programmes, based on a provided list of 83 graduates between September 2018 and September 2019. A variety of topics and a diversity of examiners were included in the selection. The project coordinator and panel chair assured that the distribution of grades in the selection matched the distribution of grades of all available theses.

After studying the self-evaluation report, theses and assessment forms, the panel members formulated their preliminary findings. The secretary collected all initial questions and remarks and distributed these amongst all panel members. At the start of the site visit, the panel discussed its initial findings on the self-evaluation report and the theses, as well as the division of tasks during the site visit.

Site visit

The site visit to the University of Amsterdam took place on 9 and 10 December 2019. Before and during the site visit, the panel studied the additional documents provided by the programmes. An overview of these materials can be found in Appendix 5. The panel conducted interviews with representatives of the programmes: students and staff members, the programme's management, alumni and representatives of the Board of Examiners. It also offered students and staff members an opportunity for confidential discussion during a consultation hour. No requests for private consultation were received. The panel used the final part of the site visit to discuss its findings in an internal meeting. Afterwards, the panel chair publicly presented the panel's preliminary findings and general observations.

Report

After the site visit, the secretary wrote a draft report based on the panel's findings and submitted it to the project coordinator for peer assessment. Subsequently, the secretary sent the report to the panel. After processing the panel members' feedback, the project coordinator sent the draft reports to the Faculty in order to have it checked for factual irregularities. The project coordinator discussed the ensuing comments with the panel's chair and changes were implemented accordingly. The report was then finalised and sent to the Faculty and University Board.

Definition of judgements standards

In accordance with the NVAO's Assessment framework for limited programme assessments, the panel used the following definitions for the assessment of the standards:

Generic quality

The quality that, from an international perspective, may reasonably be expected from a higher education Associate Degree, Bachelor's or Master's programme.

Meets the standard

The programme meets the generic quality standard.

Partially meets the standard

The programme meets the generic quality standard to a significant extent, but improvements are required in order to fully meet the standard.

Does not meet the standard

The programme does not meet the generic quality standard.



The panel used the following definitions for the assessment of the programme as a whole:

Positive

The programme meets all the standards.

Conditionally positive

The programme meets standard 1 and partially meets a maximum of two standards, with the imposition of conditions being recommended by the panel.

Negative

In the following situations:

- The programme fails to meet one or more standards;
- The programme partially meets standard 1;
- The programme partially meets one or two standards, without the imposition of conditions being recommended by the panel;
- The programme partially meets three or more standards.

SUMMARY JUDGEMENT

The panel concludes that the master's programme AI from the UvA has a clear technical profile. It has a strong focus on machine learning and its applications in vision, language and information retrieval. This profile is in line with the strong research expertise of the UvA in these aspects of AI. The panel does advise the programme to be explicit on its positioning in the AI field. The programme has formulated clear intended exit qualifications. They are in line with the with the Dublin descriptors. The intended learning outcomes reflect the academic master's level of the programme.

The panel established that the curriculum of the AI master's programme of the University of Amsterdam is designed to realise its final qualifications. The core courses create a clear and interesting curriculum that is in line with the technical profile of the programme. The focus is on the theory and applications of machine learning. The course content is up to standard. In combination with specialised electives, the core courses offer students the opportunity to reach a high level of expertise. Ethical issues related to AI are addressed as well. The panel appreciates the fact that the students can come into contact with the AI industry in one of the elective projects or the Master Project. The programme is challenging but feasible. The panel recommends adding more structure to the graduation phase.

The panel is positive about the quality, English proficiency and teaching skills of the teaching staff of the programme. Student numbers have increased substantially, and the pressure on the teaching staff is high. The reduction in the maximum intake of students seems to be a good and necessary measure to deal with this. The role of senior teaching assistants appears to function well in ensuring small-scale education for a large number of students.

According to the panel, the assessment within the master's programme Artificial Intelligence is up to standard. The assessment policies are in line with the university policy and are well documented. In practice, the responsibilities concerning assessment and grading seem to be less clear. The AI assessment is varied and aligned with the learning objectives. Providing formative feedback appears to be a point of attention. The panel concluded that the assessment of the Master Projects is fitting but still leaves room for some improvement. The Examinations Board has a clear view of its tasks and responsibilities. The AI subcommittee carries out its tasks in a proactive manner.

The panel concluded from the theses it studied and the performance of alumni that the AI students realise the intended learning outcomes. The level of the theses was adequate and in line with the expectations of academic master graduates. Several theses demonstrated a very high level. The panel established that the programme awards a high percentage of cum laude and recommends reconsidering the criteria for this judicium. From the conversation with alumni and the results of an alumni survey, the panel concluded that the students are able to pursue their career in academia or the professional field.



The panel assesses the standards from the *Assessment framework for limited programme assessments* in the following way:

Master's programme Artificial Intelligence

Standard 1: Intended learning outcomes	meets the standard
Standard 2: Teaching-learning environment	meets the standard
Standard 3: Student assessment	meets the standard
Standard 4: Achieved learning outcomes	meets the standard
General conclusion	positive

The chair, prof. dr. Ann Nowé, and the secretary, drs. José van Zwieten, of the panel hereby declare that all panel members have studied this report and that they agree with the judgements laid down in the report. They confirm that the assessment has been conducted in accordance with the demands relating to independence.

Date: 24 June 2020

DESCRIPTION OF THE STANDARDS FROM THE ASSESSMENT FRAMEWORK FOR LIMITED FRAMEWORK ASSESSMENTS

The master's programme Artificial Intelligence (AI) is embedded in the Graduate School of Informatics, part of the Faculty of Science of the University of Amsterdam (UvA). It has a Programme Director, who is responsible for the quality and development of the programme. He works together with the Programme Coordinator, who takes care of organisational matters. The programme has a Programme committee, consisting of three students and three staff members. The Programme committee fulfils its legal role in the quality assurance of the programme, as well as proactively advising the programme management on the quality and development of the programme. To keep the programme aligned with the expectations of the academic and professional field, the Graduate School of Informatics has installed a Professional Advisory Board that meets regularly. It advises on the vision and content of the programmes.

Standard 1: Intended learning outcomes

The intended learning outcomes tie in with the level and orientation of the programme; they are geared to the expectations of the professional field, the discipline, and international requirements.

Findings

The master's programme AI aims to provide students with a research-oriented education of internationally recognised top quality that allows them to become artificial intelligence professionals. According to the self-evaluation report, the AI programme of the UvA distinguishes itself by a strong focus on data-driven approaches and in particular the mathematical foundations of machine learning and its application in language, vision and retrieval. It strives to equip students with the knowledge and skills necessary to address present and future challenges in these areas. This focus is in line with the research profile of the research institutes that are involved in the AI programmes at the UvA.

The general context of the master's programme profile is formed by the KION Frame of Reference which was compiled by the Dutch academic AI programmes. This document formulates the current Dutch consensus on a national frame of reference for Artificial Intelligence in the Netherlands. Graduates of master's programmes Artificial Intelligence are expected to have gained a basic understanding of the eight key areas of Artificial Intelligence, an advanced level of knowledge in some key areas, and specialist knowledge of at least one of the key areas. These are defined as:

- Cognitive Science
- Computational Linguistics
- Context of Artificial Intelligence (History, Philosophy, Ethics)
- Intelligent Autonomous Agents and Multi-Agent Systems
- Interaction (Perception, Human-Computer Interaction, Communication)
- Knowledge Representation and Reasoning
- Machine Learning
- Problem Solving (Search, Decision Making, Optimisation)

The panel established that for the master's programme, the UvA has chosen to emphasise certain aspects of the KION framework (machine learning, mathematics). In its opinion, this positioning is relevant and legitimate in light of the development of the professional field and the research expertise of the UvA. At the same time, some of the key areas of the KION framework are not addressed in the core of the programme. This is the case for Multi-Agent Systems and Knowledge Representation. In that sense, the UvA deviates from the national point of reference for programmes in AI. However, these topics can be chosen as an elective. The panel believes that the programme should be more explicit in this positioning, both to future students and to the KION, to better clarify what can be expected of graduates.



The objectives of the programme are elaborated in eight exit qualifications (see Appendix 1). The exit qualifications address the AI-specific knowledge and skills as well as the general academic and professional skills that the students are expected to acquire. Graduates should be able to apply their knowledge to analyse, design and develop AI systems, while being aware of the social context and consequences of their work. They are expected to be able to contribute to further developments of the theory and application in AI and to express themselves clearly. The panel noticed that the intended knowledge acquisition described in exit qualification 2 could be clearer: it would be preferable to state that graduates should have general knowledge of all of the listed subfields and specialised knowledge in at least one of them. It established that the exit qualifications tie in with the Dublin Descriptors for master's degree programmes and therefore sufficiently indicate what could be expected from students at a master's level. They reflect the content, level and orientation of the master's programme. The panel finds the distinction between the exit qualifications of the bachelor's and the master's programmes to be clear.

Considerations

The panel concludes that the master's programme AI from the UvA has a clear technical profile. It has a strong focus on machine learning and its applications in vision, language and information retrieval. This profile is in line with the strong research expertise of the UvA in these aspects of AI. The panel does advise the programme to be explicit on its positioning in the AI field. The programme has formulated clear intended exit qualifications. They are in line with the with the Dublin descriptors. The intended learning outcomes reflect the academic master's level of the programme.

Conclusion

Master's programme Artificial Intelligence: the panel assesses Standard 1 as 'meets the standard'.

Standard 2: Teaching-learning environment

The curriculum, the teaching-learning environment and the quality of the teaching staff enable the incoming students to achieve the intended learning outcomes.

Findings

The master's programme AI of the UvA has a two-year, 120 EC curriculum. The programme follows the university's academic 8-8-4 schedule. This schedule consists of six periods: each semester has two eight-week periods and a final four-week period. Appendix 3 contains an overview of the curriculum.

All students take 42 EC of compulsory courses and 18 EC constrained choice electives. The compulsory core addresses the focal aspects of AI as stated in the intended learning outcomes: machine learning, computer vision, deep learning, natural language processing, information retrieval and symbolic systems. It also contains the 'Fairness, Accountability, Confidentiality and Transparency in AI' course, which addresses the ethical and social consequences of AI and concepts like 'responsibility' from a technical perspective. It helps the students to make informed decisions when assessing and addressing the potential risks of algorithms. The panel is positive about this substantial attention paid to the ethical and societal matters of AI in the programme. Some of the constrained electives build upon the knowledge from the core courses, such as Machine Learning 2 or Computer Vision 2. There are also courses that offer other perspectives within the AI field, such as Game Theory or Information Theory.

Furthermore, the students take 12 EC of elective courses. The programme listed recommended AI-related electives, some of them from the Vrije Universiteit Amsterdam. The students can also take electives from other programmes, after receiving approval from the Examinations Board. Another option is to use 6 or 12 EC to do an individual project in collaboration with a company or an internship as part of their study programme. The remaining 48 EC of the programme is dedicated to the Master Project. The students are required to prepare a Personal Education Plan (PEP) in which they display the choices they intend to make in their own trajectory. The study advisor assists them in doing so.

The students can update their PEP during the programme. A final version needs to be approved by the Examinations Board before they can apply for their Master's Certificate.

The panel studied the curriculum, as well as the content of a number of courses and projects, and spoke with the students and staff about the content of the curriculum. It concluded that the curriculum is well-designed. The core courses cover an interesting selection of topics, after which the students can specialise in advanced electives. This combination makes it possible to achieve a high level of expertise in line with the technical profile of the programme. The students appear to find what they were looking for in this programme when they applied, and they are satisfied with its level and content. The quality of the study materials is good. The panel appreciates that elective projects and master projects can be done externally. This offers students the possibility to gain experience in the professional applications of AI.

The AI programme uses a combination of teaching formats to realise the integration of theory and practice: lectures, workshops, seminars, lab sessions and project formats. All courses have a combination of lectures and small-scale tutorial sessions in groups of 20-25 students. In these sessions, the students perceive their tutors as accessible and willing to answer questions. They work on a variety of assignments, ranging from programming assignments to poster presentations and case studies. They seem to be satisfied with the quality of the teaching, but they would like to receive more feedback on their assignments.

The master's programme AI has adopted English as the language of instruction. The programme management and teaching staff are convinced that internationalisation of the programme should be pursued, in order to prepare the students properly for developments in the academic and professional field of AI. The rising number of international students contributes in their opinion to a stimulating learning environment. The panel agrees with the policy of the UvA that offering the programme in English, and with an English programme name, ensures the best fit with the international character of the AI domain and the desired balance of the national and international student population.

Thesis Trajectory

48 EC of the programme are dedicated to the Master Project. The students work on an individual project with a research component. They can come with their own proposal or apply for one of the projects that are made available by the AI research groups or external companies. The programme organises a yearly Thesis Fair at which the students can meet companies that have a graduation project available. These external projects have to be approved in advance by staff members and the thesis coordinator. After finding a project, the students draw up a Master Thesis contract. This contract comprises a brief project description, a project start date and intended end date. The names of the supervisor and assessors are also included in the contract. The students have a daily supervisor who is a researcher from the UvA. In the case of an external project, they have a supervisor at the company as well. The majority of the students use this opportunity to do their Master Project externally.

The programme organises coaching sessions for students who work on their Master Project. They are expected to present their project to their peers at the halfway point and to receive feedback in one of these sessions. An interim evaluation by the supervisor and second assessor is also recommended by the programme, but not mandatory. The students finish their project by writing a thesis and defending it in a public session. Apart from the feedback sessions, there is not a standard schedule for the project. The students are expected to manage their planning themselves. The panel recommends organizing more guidance for the Master Projects. Putting the students in charge of their planning can be expected at the master's level. This can, however, affect the workload of supervisors when the students need more time and effort to finalise their project, especially considering the rising student numbers. Furthermore, students coming from abroad usually do not have the time to extend their Master Project. This may create inequality between national and international students. Installing a standard schedule with deadlines for a project plan, interim products and the final thesis may help the students to keep the expected pace with their project.



Student intake and feasibility

The master's programme AI at the UvA is a selective programme. In order to be admitted, applicants need to have a bachelor's degree in AI, Computer Science or an AI major in the Natural and Social Sciences (Bèta Gamma) or Future Planet Studies programmes of the UvA. Their bachelor GPA (Grade Point Average) needs to be at least 6.5, and the allowed study length is the nominal length plus a maximum of three years. An admission committee of five members evaluates all the admission requests. They provide the programme director with a positive or negative advice. Students with a positive advice also get a ranking score on criteria such as GPA, duration of their bachelor's programme, publications, programming skills and motivation. The programme has been confronted with a growing intake from approximately 25 students in 2013-2014 to nearly 200 students in 2018-2019. In order to deal with this, it has introduced a maximum of 120 students from 2019-2020. It is striving to increase the capacity to 150 students per year.

The programme committee is continuously monitoring whether there are specific course units causing study delay. It has several means to monitor the study load of individual courses and of combinations during a period: course evaluations, feedback sessions and curriculum evaluations by students who request their diploma. In some cases, this has led to a recommendation to the course coordinator to adjust the course. In general, the students experience the programme as challenging but doable. They spend 30-40 hours per week on average on their studies.

The self-evaluation report mentioned that the study success of the programme has increased in recent years. The dropout rate is around 10%, and the most recent success rate is 71% graduation within 3 years. The programme aims to realise the intended success rate of the Graduate School of Informatics, which is 75%. This is expected to be realistic as the programme has become more selective in admitting students. The panel discussed the feasibility, study yield and drop-out rates with the management, students and staff during the site visit. There was no mention of specific obstacles regarding the feasibility of the programme. The yield and drop-out rates of this programme are comparable to those of other academic degree programmes at Dutch universities. The panel concluded that the programme management is paying sufficient attention to these topics. Students coming from abroad usually do not have the time to extend their Master Project. In assessing the master's thesis, the time spend on the thesis could very well be taken into consideration. This also relates to the relative high amount of *cum laudes* (see Standard 4).

Teaching staff

The panel looked at the list of teaching staff provided in the self-evaluation report and interviewed a small selection of AI faculty. It concluded that the quality of the teaching staff is good. The courses are delivered by active researchers from four research institutes in the field of computer sciences. They all have a PhD degree. The level of English proficiency of the teaching staff is good, which allows them to appropriately deliver the English-taught curriculum. The panel established that the AI staff possesses the necessary didactic skills. 75% of the lecturers holds a university teaching qualification (UTQ, in Dutch: BKO), and the remaining 25% is completing this qualification trajectory. There are teacher meetings during which the teaching staff is informed of and discusses topics such as assessment and changes in the curriculum.

The pressure on the teaching staff has increased during the last years due to the growth in student numbers in the bachelor's and master's programmes. An increase in the teaching staff is necessary but has only been partially realised. The programme management has arranged a teaching structure that is designed to realise small-scale education, despite the high number of students and scarcity of teachers. It has hired junior lecturers (usually alumni who graduated recently) and teaching assistants who deliver lab group sessions for groups of approximately 20 students. The teaching assistants receive training in didactical skills. Junior lecturers are expected to follow the UTQ trajectory. A new role has been introduced to promote the quality and coherence of the lab groups: Super Teaching Assistant. Junior lecturers can fulfil this role, as can more experienced teaching assistants. This person coordinates the communication between the teaching assistants and the lecturer of the course. He or she makes a lesson plan for the lab group sessions and coordinates the

grading. The students and staff members seem to be content with this system. The panel was impressed by how the programme facilitates small-scale and well-coordinated education in the programme with this approach. The students appreciate this system as well. The teaching assistants are very approachable and can relate well to the students' perspectives and needs in the courses.

Facilities

During the site visit the panel visited the RoboLab and Blue Student Lab, in which the students can work on a voluntary basis and extra curricula with several robots and the associated hardware with students from other disciplines.. During the programme, the students also get support with obtaining access to software they need for their coursework. Individual and collaborative study facilities are available for them at the faculty, but the growth in student numbers makes it harder to gain access to these facilities. In the near future, the programme will move to another location in which education, business partners of the faculty, and research will all have a floor. These facilities are expected to be sufficient for the current student numbers. The panel concluded that the programme offers sufficient facilities to its students.

Considerations

The panel established that the curriculum of the AI master's programme of the University of Amsterdam is designed to realise its final qualifications. The core courses create a clear and interesting curriculum that is in line with the technical profile of the programme. The focus is on the theory and applications of machine learning. The course content is up to standard. In combination with specialised electives, the core courses offer students the opportunity to reach a high level of expertise. Ethical issues related to AI are addressed as well. The panel appreciates the fact that the students can come into contact with the AI industry in one of the elective projects or the Master Project. The programme is challenging but feasible. The panel recommends adding more structure to the graduation phase.

The panel is positive about the quality, English proficiency and teaching skills of the teaching staff of the programme. Student numbers have increased substantially, and the pressure on the teaching staff is high. The reduction in the maximum intake of students seems to be a good and necessary measure to deal with this. The role of senior teaching assistants appears to function well in ensuring small-scale education for a large number of students.

Conclusion

Master's programme Artificial Intelligence: the panel assesses Standard 2 as 'meets the standard'.

Standard 3: Student assessment

The programme has an adequate system of student assessment in place.
--

Findings

The assessment within the AI programme follows the policy of the UvA as documented in the 'UvA Toetskader'. The faculty operationalised this in the policy document 'Handleiding Toetsing FNWI'. The Examinations Board produced 'The Examination Board's Rules and Guidelines' which is updated yearly. According to the first document, the relationship between the Dublin descriptors, final qualifications, learning goals and assessment must be clear. The assessment needs to be reliable, valid, transparent and manageable. These principles have been translated into 22 requirements for assessment at the UvA. Each requirement is assigned to a stakeholder: dean, programme director, examiner and/or the Examinations Board.

The Assessment Plan of the programme describes how these requirements have been translated into assessment practices in the AI programme. This plan is constructed by the programme director. The panel did not find an overview in this assessment plan of when and how the final qualifications are assessed in the various courses of the programme. The programme uses DataNose to structure and



present the relationship between its final qualifications and the various courses. It also has links to the study guides for each course. These study guides do include the modes of assessment.

In daily practice, assessment is a task of the examiners. They need to work on maintaining the good quality of assessment of their courses. The examination of each course is checked by a second examiner or senior teaching assistant. There is a checklist for this peer review that addresses the quality criteria for examinations. Assignments and examinations are graded by use of a pre-established answer key. All assessment materials must be filed in the assessment dossier of the course, which includes the study guide, assignments/exams, answer models and an overview of the grades.

The panel studied the assessment system of the programme, the assessment plan, and some examples of assessment dossiers. It is positive about the assessment policy of the university and the adequate requirements that are set up to stimulate and safeguard the quality of assessment. Assessment dossiers that it studied were not complete in all cases, and during its conversations with the teachers, it gained the impression that in daily practice, the responsibilities are not always clear or strictly followed up. An example is the calibration of grading: in some cases this seems to be delegated to senior teaching assistants although it is supposed to be the responsibility of the course coordinators. Also, assessment forms are not always filled in completely. The panel concluded that the assessment policy is well-designed, but that its implementation by staff members needs more attention.

AI assessment takes the form of written examinations, practical work, written assignments, reports and oral presentations. Most courses are assessed by a combination of formats. Lab groups contain formative assessments, such as programming assignments. The students would like to receive more formative feedback on their assignments and their thesis. The panel is satisfied with the variety of assessment formats. In some courses, the students can compensate practical assignments with written exams. The panel recommends reconsidering this practice.

The panel also discussed the policy regarding fraud. This subject has the attention of both the Examinations Board and of the teaching staff. They mentioned several means to deal with fraud. The programme uses Turnitin for written assignments. The Examinations Board agreed with the panel that additional prevention activities could be useful in order to make sure the students are familiar and compliant with plagiarism regulations.

Thesis trajectory

The Graduation Project is an individual research project that is executed under the responsibility of an experienced researcher who acts as the supervisor. This supervisor is usually the first assessor, with another AI researcher taking the role of second examiner. The appointment of supervisors is done by the graduation coordinator. The Examinations Board has made a list the staff members who are available for graduation supervision. The graduation process includes midterm presentation sessions. When the research is completed, the student writes a master's thesis.

Theses are assessed by the supervisor and second assessor. They perform their assessment independently using their own marking form. These forms indicate the different assessment criteria. The grade is based on their evaluation of the quality of the research, the thesis and the presentation. The process of the graduation project is included as a criterion as well. After their individual assessment of the thesis, the assessors establish the final mark together at the final stage of the defence. From 2019, this grading process has been digitised and is transparent to students.

The panel thinks that the assessment procedure for the Master's Project is fitting. It is positive that there are always two academic examiners involved in the grading, and that they form their opinion separately before establishing the grade together. It observed that there is no fixed weight attached to the subgrades and criteria, which makes it sometimes hard to follow how the final grade was composed from the subgrades. Also, most forms do not contain a written substantiation of the final

grade. The panel recommends including the relative weights of each category on the assessment form for each thesis and providing a qualitative substantiation for the grading. Alternatively, the programme could include a fixed weight of subgrades on the assessment form.

Examinations Board

The quality of AI assessment is safeguarded by an Examinations Board appointed for all programmes within the Faculty of Science. This Board has an external member and is advised by an assessment expert. It has installed an assessment committee that advises the programmes on exam-related matters. There is also a subcommittee that is dedicated to dealing with policy and the handling of fraud and plagiarism. The board has programme-specific subcommittees, with one of them being the subcommittee for the bachelor's programme KI and the master's programme AI. The panel interviewed members of the Examination Board and the AI subcommittee. It concluded that they have a clear view of their tasks and responsibilities and that they work hard to guarantee the quality of assessment faculty-wide.

The AI subcommittee has four members and meets once a month. It deals with students' requests concerning changes in their individual study programme and requests for exemptions. A secretary supports the subcommittee with its work. The subcommittee appoints the programme's examiners and monitors the quality of assessment within the programmes. It has recently introduced meetings with course coordinators. In these meetings, a course coordinator is invited to discuss the assessment and grading of the course. These conversations seem to be appreciated by the staff members, as this gives them an opportunity to discuss difficult considerations in the assessment process and to exchange best practices. Each year, a sample of course coordinators is invited to have such a meeting. The selection is made based on pass rates and student evaluations. The results of these meetings are reported to the Programme Director. The subcommittee intends to monitor each course at least once every five years. For the graduation project, it takes regular samples of theses to monitor their level and the grading process.

Considerations

According to the panel, the assessment within the master's programme Artificial Intelligence is up to standard. The assessment policies are in line with the university policy and are well documented. In practice, the responsibilities concerning assessment and grading seem to be less clear. The AI assessment is varied and aligned with the learning objectives. Providing formative feedback appears to be a point of attention. The panel concluded that the assessment of the Master Projects is fitting but still leaves room for some improvement. The Examinations Board has a clear view of its tasks and responsibilities. The AI subcommittee carries out its tasks in a proactive manner.

Conclusion

Master's programme Artificial Intelligence: the panel assesses Standard 3 as 'meets the standard'.

Standard 4: Achieved learning outcomes

The programme demonstrates that the intended learning outcomes are achieved.

Findings

In order to establish whether the master's programme AI demonstrates that the final qualifications are achieved, the panel studied a selection of fifteen master theses and their assessment forms. It agreed with the assessments of these theses. It concluded that the theses showed that the students are on an academic master's level. They demonstrated that the graduates have the capability to perform a theoretically founded, in-depth research project and to report on it. The panel read several theses of a very high level. In some cases, these projects resulted in publications. The panel concluded that the quality of the studied theses indicated that the graduates have achieved the final qualifications of the master's programme and that the programme makes it possible for them to excel.



The panel was informed that the percentage of *cum laudes* among AI graduates is very high: up to 45% in 2018-2019. It discussed this with the thesis supervisors and programme management. They explained that, in general, the students are ambitious, and the intake procedure enables the programme to select good students. Although it is plausible that the ambitions and the achieved level of the selected students is high, *cum laude* is considered to be a more relative, distinctive judicium according to the panel. It advises the programme to reconsider the *cum laude* criteria. Also, when assessing a thesis, the time spend could be taken into consideration.

The panel interviewed a couple of alumni and employers of AI alumni. They are enthusiastic ambassadors of the programme. Their advanced expertise in machine learning appears to be an appreciated characteristic. The results of an alumni survey in 2019 indicated that after graduation, the graduates easily find a job: 78% of the respondents found a job in less than 1 month after graduating, 90% of them found work in the field of AI. The panel concluded that the programme prepares the students well for a career in the academic or business field of AI.

Considerations

The panel concluded from the theses it studied and the performance of alumni that the AI students realise the intended learning outcomes. The level of the theses was adequate and in line with the expectations of academic master graduates. Several theses demonstrated a very high level. The panel established that the programme awards a high percentage of *cum laude* and recommends reconsidering the criteria for this judicium. From the conversation with alumni and the results of an alumni survey, the panel concluded that the students are able to pursue their career in academia or the professional field.

Conclusion

Master's programme Artificial Intelligence: the panel assesses Standard 4 as 'meets the standard'.

GENERAL CONCLUSION

The panel judged that the master's programme in Artificial Intelligence offered by the University of Amsterdam meets all standards of the NVAO assessment framework for limited programme assessment. It therefore advises positively about the re-accreditation of the programme.

Conclusion

The panel assesses the *master's programme Artificial Intelligence* as 'positive'.

APPENDICES

APPENDIX 1: DOMAIN-SPECIFIC FRAMEWORK OF REFERENCE

Frame of Reference: Bachelor's and Master's Programmes in Artificial Intelligence

For author names and article: https://pure.uva.nl/ws/files/29809664/KION_FoR_2018_Final.pdf¹

October 16th, 2018

This document is an update of the 2013 Frame of Reference as developed by the KION² task force on Curricula for Artificial Intelligence, which was based on:

- Artificial Intelligence Academic Programmes in the Netherlands - A State of the Art report, Quality Assurance Netherlands Universities, 2015³
- Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science, The Joint Task Force on Computing Curricula, Association for Computing Machinery (ACM), & IEEE Computer Society, December 20, 2013, p. 221-229⁴
- The Onderwijs- en Examenregelingen (OER) of the bachelor's and master's programmes in Artificial Intelligence administered by the Dutch Universities.
- Tuning Educational Structures in Europe⁵, European project, 2000-2004.

1 | INTRODUCTION

This document is an update of the 2013 frame of reference for the Dutch University programmes included in the category Artificial Intelligence of the Dutch register of higher education programmes (CROHO)⁶ This frame of reference defines the fields covered by the term Artificial Intelligence as well as the common goals and final qualifications of these programmes.

Artificial Intelligence is a relatively young field. The birth of Artificial Intelligence research is often dated in 1956, when the founding fathers of AI met at the Dartmouth Conference. The history of teaching Artificial Intelligence as a separate discipline is much shorter still, starting in the Netherlands in the early '90's. Consequently, a frame of reference for Artificial Intelligence is still actively developing both in the national and the international context. This document formulates the current Dutch consensus on a national frame of reference for Artificial Intelligence in the Netherlands.

Intelligence is often defined as the ability to reason with knowledge, to plan and to coordinate, to solve problems, to perceive, to learn and to understand language and ideas. Originally these are typical properties and phenomena associated with the human brain, but they can also be investigated without direct reference to the natural system. Both ways of studying intelligence either can or must use computational modelling. The term Artificial Intelligence as used in this document refers to the study of intelligence, whether artificial or natural, by computational means.

1.1 KION: Artificial Intelligence in the Netherlands

The current Dutch Artificial Intelligence programmes were mostly started in the nineties in an interdisciplinary context. Originally they were known under a variety of names such as Cognitive Science (Cognitiewetenschap), Applied Cognitive Science (Technische Cognitiewetenschap), Knowledge Engineering (Kennistechnologie), Cognitive Artificial Intelligence (Cognitieve Kunstmatige Intelligentie) as well as Artificial Intelligence (Kunstmatige Intelligentie).

In 1999, the number of recognised labels in the CROHO was reduced, and the aforementioned study programmes were united under the name Artificial Intelligence⁷. Initially, this was an administrative

¹ The authors like to acknowledge the authors of the 2006 and 2013 Frame of Reference for their work; major parts of this document are still built on their original vision.

² Kunstmatige Intelligentie Opleidingen Nederland

³ www.qanu.nl/en/state-of-the-art-reports (last visited in March 2018)

⁴ www.acm.org/education/curricula-recommendations (last visited in March, 2018)

⁵ www.unideusto.org/tuning/ (last visited in February 2018)

⁶ Centraal Register Opleidingen Hoger Onderwijs

⁷ In Dutch: Kunstmatige Intelligentie



matter that did not influence the content of the curricula. It did mean, however, that from then on cognitive science (as the study of natural intelligence) and artificial intelligence (as a formal approach to intelligence) were shared under the heading of Artificial Intelligence. The above mentioned definition of Artificial Intelligence as the study of natural and/or artificial intelligence by computational means was then agreed upon. The KION (Kunstmatige Intelligentie Opleidingen in Nederland) was formed as a discussion and cooperation platform for the united programmes.

Starting in 2002, all university-level study programmes in the Netherlands were divided into a bachelor's and a master's phase. KION took this as an opportunity to agree upon a common kernel of subjects that would be constituent of every Dutch Artificial Intelligence bachelor's programme, with the aim of advancing an adequate fit of all Dutch bachelor's programmes to all Dutch master's requirements.

Since then, some degree programmes have changed their names for specification and/or marketing purposes. The Human-Machine Communication degree programme in Groningen joined the KION framework soon after the start, in 2004. In 2013, the VU changed the name of its bachelor's in Kunstmatige Intelligentie to Lifestyle Informatics, to better fit their human-oriented approach to AI, which helped to attract a new population of students (including a higher proportion of female students). However, from 2019 on, the bachelor's programme will be taught in English under the name Artificial Intelligence (with a track in Intelligent Systems and a track in Socially Aware Computing). Furthermore, in 2017, Maastricht renamed its bachelor's programme to Data Science & Knowledge Engineering, and changed its master's programme in Operations Research programme to Data Science for Decision Making, to enable more synergy with its master's AI programme. A full list of the degree programmes that are a member of the KION can be found in section 1.2.

During the last decade new developments in Artificial Intelligence (AI) have become increasingly visible to society and the general public. Most appealing successes like IBM's Watson performance and Google's DeepMind victory in AlphaGo, have globally drawn attention. In business, AI's impact on massive data-mining applications in consumer markets may even more revolutionise the use of AI in everyday life.

The successes in the field of AI have not gone unnoticed in the Dutch educational AI programmes. There has been a substantial increase in the intake of virtually all Dutch AI programmes, reflecting the awareness of the growing potential of AI by talented students. In addition, several Dutch programmes are now taught in English, attracting students from all over the world. At some Universities, the substantial growth has led to measures to maintain quality, e.g. by introducing a Binding Study Advice or even by imposing a Numerus Fixus (Radboud University and University of Amsterdam in September 2018; other Universities are likely to follow in 2019).

The prospects of a career in AI, directly or via business-related spin-offs, are very promising; we therefore have to be prepared to face the challenge of keeping quality of our AI programmes during upcoming years, while offering enough capacity to train professionals to fulfil the future needs of society in implementing AI-based solutions.

The following degree programmes are a member of the Kunstmatige Intelligentie Overleg Nederland:

1.2.1 Bachelor's Programmes of the KION

The following Bachelor's programmes are a part of the KION:

- B Artificial Intelligence, Radboud Universiteit Nijmegen (CROHO: 56945)
- B Data Science and Knowledge Engineering, Universiteit Maastricht (CROHO: 50300)
- B Kunstmatige Intelligentie, Rijksuniversiteit Groningen (CROHO: 56981)
- B Kunstmatige Intelligentie, Universiteit van Amsterdam (CROHO: 56981)
- B Kunstmatige Intelligentie, Universiteit Utrecht (CROHO: 56981)
- B Artificial Intelligence, Vrije Universiteit Amsterdam (CROHO: 56983)

1.2.2 Master's Programmes of the KION

- The following Master's degree programmes are a part of the KION:
- M Artificial Intelligence, Radboud Universiteit Nijmegen (CROHO: 66981)
- M Artificial Intelligence, Rijksuniversiteit Groningen (CROHO: 66981)
- M Artificial Intelligence, transnationale Universiteit Limburg (CROHO: 66981)
- M Artificial Intelligence, Universiteit Utrecht (CROHO: 66981)
- M Artificial Intelligence, Universiteit van Amsterdam (CROHO: 66981)
- M Artificial Intelligence, Vrije Universiteit Amsterdam (CROHO: 66981)
- M Data Science for Decision Making, transnationale Universiteit Limburg (CROHO: 60125)
- M Human-machine Communication, Rijksuniversiteit Groningen (CROHO: 60653)

1.3 Aim of this Document

Now that the Dutch Artificial Intelligence programmes are coming up for accreditation in 2019, KION feels that the essence of the 2013 Frame of Reference is still valid, but in definite need of an update. However, this document is not intended purely as a description of the current status quo. Rather, it aims to provide an account of what an Artificial Intelligence programme should provide as a minimum (the communal requirements for every study programme called Artificial Intelligence), and how it can extend this basis to distinguish itself from other Artificial Intelligence programmes.

Agreement among the Dutch Artificial Intelligence programmes upon the contents of this document will advance both the equivalence of these programmes, and the understanding on existing and possible profiles within Artificial Intelligence programmes. Moreover, it is hoped that this document will also be a starting point for defining international standards for Artificial Intelligence programmes.

2 | PROGRAMME CHARACTERISTICS

This section describes definitions regarding the build-up of bachelor's and master's programmes.

2.1 Areas, Courses, Modules and Topics

A bachelor's programme in Artificial intelligence is organised hierarchically into three levels. The highest level of the hierarchy is the area, which represents a particular disciplinary subfield. The areas are broken down into smaller divisions called modules, which represent individual thematic units within an area. A module may be implemented as a complete course, be covered in part of a course, or contain elements from several courses. Each module is further subdivided into a set of topics, which are the lowest level of the hierarchy. The modules that implement the particular programme (or curriculum) are together referred as the 'body of knowledge'.

2.2 Core and Elective Courses

By insisting on a broad consensus in the definition of the core, we hope to keep the core as small as possible, giving institutions the freedom to tailor the elective components of the curriculum in ways that meet their individual needs. The core is thus not a complete programme. Because the core is defined as minimal, it does not, by itself, constitute a complete undergraduate curriculum. Every undergraduate programme must include additional elective courses relating to the body of knowledge. This report does not define what those courses should be, but does enumerate options in terms of modules.

2.3 Assessing the Time Required to Cover a Course

To give readers a sense of the time required to cover a particular course, a metric must be defined that establishes a standard of measurement. No standard measure is recognised throughout the world, but within the European Community agreement has been reached upon a uniform European Credit Transfer System⁸ (ECTS) in which study load is measured in European Credits (ECs). One EC stands for 28 hours of study time and a full year of study is standardised at 60 EC. In this document, we shall use the EC metric as the standard of measurement for study load.

⁸ https://ec.europa.eu/education/resources/european-credit-transfer-accumulation-system_en (last visited on May 4, 2018)



2.4 Coping with Change

An essential requirement of any Artificial Intelligence degree is that it should enable graduates to cope with-and even benefit from-the rapid change that is a continuing feature of the field. But how does one achieve this goal in practice? At one level, the pace of change represents a challenge to academic staff who must continually update courses and equipment. At another level, however, it suggests a shift in pedagogy away from the transmission of specific material, which will quickly become dated, toward modes of instruction that encourage students to acquire knowledge and skills on their own.

Fundamentally, teaching students to cope with change requires instilling an attitude that promotes continued study throughout a career in those students. To this end, an Artificial Intelligence curriculum must strive to meet the following challenges:

- Adopt a teaching methodology that emphasises learning as opposed to teaching, with students continually being challenged to think independently.
- Assign challenging and imaginative exercises that encourage student initiative.
- Present a sound framework with appropriate theory that ensures that the education is sustainable.
- Ensure that equipment and teaching materials remain up to date.
- Make students aware of information resources and appropriate strategies for staying current in the field.
- Encourage cooperative learning and the use of communication technologies to promote group interaction.
- Convince students of the need for continuing professional development to promote lifelong learning.
- Provide students with awareness of potential ethical and legal issues the field of Artificial Intelligence.

3 | SHARED IDENTITY

3.1 Common Role

Apart from the roles academics usually perform in society students of Artificial Intelligence are educated to enrich society with the benefits a formalisation of intelligence and intelligent phenomena can provide. In particular this entails that an alumnus of Artificial Intelligence can contribute to the understanding and exploitation of natural and artificial intelligence. This may lead to new technologies but it may also enrich designs, products, and services with intelligence so that they are more effective, more reliable, more efficient, safer, and often require less natural resources. This role, in combination with the interdisciplinary nature of the field, requires the Artificial Intelligence alumnus to be able to contribute to interdisciplinary teams and, in many cases function as an intermediate who facilitates the interaction of (other) domain specialists.

3.2 Common Requirements

Artificial Intelligence is a broad discipline and many approaches to the study of intelligent phenomena are justified and fruitful. Curricula are therefore often different from their siblings in emphasis, goals, and capabilities of their graduates. Yet they have much in common. Any reputable Artificial Intelligence programme should include each of the following aspects.

1. Essential and foundational underpinnings of the core aspects of intelligence. These must be founded on empirical efforts and based on a formal theory, and they may address professional values and principles. Regardless of their form or focus, the underpinnings must highlight those essential aspects of the discipline that remain unaltered in the face of technological change. The discipline's foundation provides a touchstone that transcends time and circumstances, giving a sense of permanence and stability to its educational mission. Students must have a thorough grounding in that foundation.
2. A foundation in the core concepts of modelling and algorithms for implementing intelligence. The construction and use of models (simplified, abstracted and dynamic representations of some phenomenon in reality) is common to many sciences. In Artificial Intelligence, however, model

building is central: the field of Artificial Intelligence may actually be defined as trying to model aspects of (formal or natural) intelligence and knowledge. Moreover, models within Artificial Intelligence have specific characteristic: they are computational and therefore necessarily mathematical or formal. Artificial Intelligence-graduates must therefore be able to work with (computational) models at different levels of abstraction and understand the recursive nature of models in Artificial Intelligence. This foundation has a number of layers:

- An understanding of, and appreciation for, many of the diverse aspects of intelligence, models of intelligent phenomena, and of algorithms that describe intelligent processes.
- Skills to model intelligent phenomena and appreciate the abilities and limitation of these models, if appropriate in comparison with a natural intelligence counterpart.
- Skills to model and implement intelligent phenomena on a computer, in particular skills to work with algorithms and data-structures in software.
- Skills to design and build systems that are robust, reliable, and appropriate for their intended audience.

3. An understanding of the possibilities and limitations of what intelligent systems can and cannot do. This foundation has a number of levels:

- An understanding of what current state-of-the-art can and cannot accomplish, if appropriate in combination with the accomplishment of the natural system that inspired it;
- An understanding of the limitations of intelligent systems, including the difference between what they are inherently incapable of doing versus what may be accomplished via future science and technology;
- The impact of deploying technological solutions and interventions on individuals, organisations, and society.

4. The identification and acquisition of non-technical skills, including interpersonal communication skills, team skills, and management skills as appropriate to the discipline. To have value, learning experiences must build such skills (not just convey that they are important) and teach skills that are transferable to new situations.

5. Exposure to an appropriate range of applications and case studies that connect theory and skills learned in academia to real-world occurrences to explicate their relevance and utility.

6. Attention to professional, legal and ethical issues such that students acquire, develop and demonstrate attitudes and priorities that honour, protect, and enhance the profession's ethical stature and standing.

7. Demonstration that each student has integrated the various elements of the (under)graduate experience by undertaking, completing, and presenting a capstone project.

3.3 Shared Background for Bachelor's Programmes

Similar to alumni of programmes such as Physics, Computer Science, and Psychology, all Artificial Intelligence bachelors are expected to share a certain amount of support knowledge, domain specific knowledge, specialised domain knowledge, and a set of skills. The content mentioned below ensures a firm common basis that enables AI bachelors of any Dutch university admission to any Dutch master's programme in AI. At the same time, it allows for a wide range of individual and/or institute specific specialisations. The list is an update (extension) of the shared programme agreed upon by the KION platform in 2013.

3.3.1 Core Modules (shared between AI Bachelor's Degree Programmes)

The following topics and skills are part of each of the bachelor's programmes, either as a dedicated course or as a substantial topic within one or more courses. The core modules define the 8 key areas of Artificial Intelligence.

3.3.1.1 Artificial Intelligence (Core) Modules

- Algorithmic Problem Solving (Search, Decision Making, Optimisation)
- Cognitive Science
- Computational Linguistics
- Context of Artificial Intelligence (History, Philosophy, Ethics)
- Intelligent Autonomous Agents and Multi-Agent Systems



- Interaction (Perception, Human-Computer Interaction, Communication)
- Knowledge Representation and Reasoning
- Machine Learning

3.3.1.2 Support Module

- Computer Science
 - Algorithms and Data Structures
 - Programming
- Logic
 - Propositional Logic
 - Predicate Logic

3.3.1.3 Academic Skills

- Mathematics
 - Calculus
 - Discrete Mathematics
 - Linear Algebra
 - Probability Theory
 - Statistics

Apart from curriculum specific skills, the bachelor's programmes support the development of a set of general academic skills. Even though they can be topics in specific modules, they are generally addressed by the appropriate choice of work and assessment methods throughout the curriculum.

- Analytic Skills
- Empirical Methods
- Modelling
- Teamwork
- Written and Oral Communication, Argumentation and Presentation

3.3.2 Elective Modules (within Artificial Intelligence)

The following list of modules is considered as representative of the AI field at this moment. Given that the different AI programs have different priorities in selecting topics, and assigning topics to either the bachelor's or master's, each bachelor's should offer a substantial subset of the following list as part of their bachelor's programme, either as specific course, or as a substantial part of a broader course (i.e. a module).

- Architectures of Cognition and Cognitive Modelling
- Computational and Cognitive Neuroscience
- Computational Intelligence
- Computer Vision
- Data Mining
- Deep Learning
- ELSA (Ethical, Legal and Social Aspects of AI)
- Evolutionary Algorithms (Genetic Algorithms, Evolutionary Computing)
- Language and Speech Technology
- Neural Networks
- Perception (Computational and Natural)
- Reasoning under Uncertainty
- Reinforcement Learning
- Robotics
- Text Mining and Information Retrieval
- Virtual Reality and Gaming
- Web and Artificial Intelligence

4 | BACHELOR'S PROGRAMME ARTIFICIAL INTELLIGENCE

This section is divided into two parts. Section 4.1 describes the roles that a bachelor ought to be able to perform in society. Section 4.2 describes the final qualifications that bachelors in Artificial Intelligence possess in order to fulfil these roles.

4.1 Objectives

The objective of the bachelor's programme is to provide students with a suitable basis for a further career, both in education as well as in employment. The bachelor must be prepared for a number of different roles and opportunities.

4.1.1 Access to Master's Programmes

The bachelor provides the student with the specific knowledge and abilities, exemplified in the form of a bachelor's diploma that allows the bachelor to apply for any master's programme in Artificial Intelligence or other national or international master's programmes, particularly in related disciplines.

4.1.2 Professional Career

The bachelor prepares for a position in which the student can earn his or her own subsistence. In particular it prepares for:

- Supervised work on a national and international academic level;
- Positions in the modern high-tech society, such as functions in knowledge-intensive companies and knowledge intensive parts of the non-profit sector.

4.1.3 Academic Skills

The bachelor provides sufficient training in (scientific) reasoning, conduct, and communication to reach internationally accepted standards of academic skills at that level.

4.1.4 Place in Society

The bachelor's programme provides the bachelor with the knowledge and tools needed to form an informed opinion of the meaning and impact of Artificial Intelligence, and an informed notion of the responsibilities of a specialist in this area.

4.2 Final Qualifications

The objectives of the bachelor can be specified into final qualifications. To comply with international standards these qualifications are presented below in terms of the Dublin descriptors for the bachelor's profile⁹. Together these final qualifications must lead to alumni that exemplify the shared identity defined in section 3.

4.2.1 Knowledge and Understanding

The bachelor demonstrates knowledge and understanding in a field of study that builds upon and supersedes their general secondary education. Knowledge and understanding is typically at a level at which the bachelor, whilst supported by advanced textbooks, is able to include some aspects at the forefront of their field of study.

We distinguish two levels of qualification: a basic understanding, corresponding to knowledge of the essentials and fundamentals of the field in question, such as knowing what the field is, knowing how to apply knowledge in said field, and knowing how to further develop oneself in the field in question, and an advanced understanding, corresponding to in-depth knowledge about a topic in question.

Qualifications:

1. Basic understanding of the (8) key areas in Artificial Intelligence in accordance with the shared identity.
2. Advanced knowledge of at least one of the key areas in Artificial Intelligence, up to a level that without further requirements grants access to a master programme in this area.

⁹ "A Framework for Qualifications of the European Higher Education Area", 2005. (last visited on May 4, 2018)



4.2.2 Applying Knowledge and Understanding

Bachelors can apply their knowledge and understanding in a manner that indicates a professional approach to their work or vocation, and have competences typically demonstrated through devising and sustaining arguments and solving problems and/or designing systems within their field of study. They are able to analyse and model prototypical Artificial Intelligence problems by using known Artificial Intelligence methods and techniques.

Qualifications:

1. The ability to understand, apply, formulate, and validate models from the domains of Artificial Intelligence.
2. The ability to apply knowledge from the key areas of Artificial Intelligence. (as outlined in 3.3.1.1)
3. The ability to apply knowledge from the support modules of Artificial Intelligence (as outlined in 3.3.1.2)
4. Analytical approach to problem solving and design:
 - Ability to comprehend (design) problems and abstract their essentials.
 - Ability to construct and develop logical arguments with clear identification of assumptions and conclusions.
5. The ability to submit an argument in the exact sciences (or humanities) to critical appraisal.
6. Analytical and critical way of thought and ability to apply logical reasoning.
7. Openness to interdisciplinary cooperation and ability to effectively participate therein as an academic professional.
8. The ability to create an effective project plan for solving a prototypical Artificial Intelligent problem in a supervised context.
9. Manage one's own learning and development, including time management and organisational skills.
10. The ability to transpose academic knowledge and expertise into (inter)national social, professional and economic contexts.
11. Readiness to address new problems in new areas, emerging from scientific and professional fields.

4.2.3 Making Judgments

The bachelor has the ability to gather and interpret relevant data (typically within the field of study) and to formulate judgments that include reflection on relevant social, academic or ethical issues.

Qualifications:

1. Ability to critically review results, arguments and problem statements from accepted perspectives in the field of Artificial Intelligence and neighbouring disciplines.
2. Initial competence in search and critical processing of professional literature in Artificial Intelligence.
3. Acquaintance with the standards of academic criticism.
4. Awareness of, and responsible concerning, the ethical, normative and social consequences of developments in science and technology, particularly resulting from Artificial Intelligence.

4.2.4 Communication

The bachelor can communicate information, ideas, problems and solutions to audiences of both domain-specialist and a general audience.

Qualifications:

Academically appropriate communicative skills; the bachelor can:

- Communicate ideas effectively in written form and through the use of Information and Communication Technology,
- Make effective oral presentations, both formally and informally,
- Understand and offer constructive critiques of the presentations of others.

4.2.5 Learning Skills

The bachelor has developed those learning skills that are necessary for a successful further study characterised by a high degree of autonomy (typically in the context of a master or a specialist profession).

Qualifications:

1. Reflection on one's own style of thought and working methods and readiness to take the necessary corrective action.
2. Recognise the need for continued learning throughout a professional career

5 | MASTER'S PROGRAMME ARTIFICIAL INTELLIGENCE

This section is divided into two parts. Section 5.1 describes the roles that a master ought to be able to perform in society. Section 5.2 describes the final qualifications that masters in Artificial Intelligence possess in order to fulfil these roles.

5.1 Objectives

The objective of the master programme is to provide students with a suitable basis for a further career, both in research as well as in the rest of society. The master must be prepared for a number of different roles and careers at key positions in society.

5.1.1 Access to PhD Programmes

The master programme provides the student with the specific knowledge and abilities, exemplified in the form of a master diploma that allows the master access to a PhD programme in a broad range of disciplines, especially in Artificial Intelligence related disciplines.

5.1.2 Professional Career

The master programme prepares for a position in which the student can earn his or her own subsistence. In particular, it prepares for:

- Independent work on an academic level, especially at positions where many of the problems have not been addressed before and where solutions require scientific training
- Key positions in the modern high-tech society, such as higher functions in knowledge-intensive companies and knowledge-intensive parts of the non-profit sector

5.1.3 Academic Skills

The master programme provides sufficient training in independent scientific reasoning, conduct, and communication to reach internationally accepted standards of academic skills at that level. Masters can communicate original ideas in their own language and in English to a public of specialists and non-specialists.

5.1.4 Place in Society

The programme provides the master with the knowledge and tools needed to formulate an informed opinion about the meaning and impact of Artificial Intelligence in society. Masters are able to enrich society with results from contemporary research and oversee the consequences of proposed measures to society and are aware of their responsibility towards society.

5.2 Final Qualifications

The objectives of the master can be specified into final qualifications. To comply with international standards these qualifications are presented below in terms of the Dublin descriptors for the master's profile¹⁰. Together these final qualifications must lead to alumni that exemplify the shared identity defined in section 3.

¹⁰ Framework_for_Qualifications_of_the_European_Higher_Education_Area (last visited on May 4, 2018)



5.2.1 Knowledge and Understanding

The Master demonstrates knowledge and understanding in a field of study that builds upon and supersedes their bachelor's degree. Knowledge, understanding, and abilities are typically at a level at which the master is able to formulate a feasible research plan in one's own specialisation.

We distinguish three levels of qualification: a basic understanding, corresponding to the minimal level of knowledge that is expected of a Bachelor student, an advanced understanding, meaning students must have in-depth knowledge about a topic that they could easily develop to become a specialist, and specialist knowledge, meaning students are highly skilled (and specialised) in the key area in question.

Qualifications:

1. Basic understanding of all (8) key areas of Artificial Intelligence.
2. An advanced understanding in some of the key areas of Artificial Intelligence.
3. Specialist knowledge of at least one of the key areas in Artificial Intelligence, up to a level that the master can appreciate the forefront of research in that field.

5.2.2 Applying Knowledge and Understanding

Masters can apply their knowledge and understanding in a manner that indicates a scientific approach to their work or vocation. They are able to handle complex and ill-defined problems for which it is not a priori known if there is an appropriate solution, how to acquire the necessary information to solve the sub-problems involved, and for which there is no standard or reliable route to the solution.

Qualifications:

1. The ability to formulate a project plan for an open problem in a field related to Artificial Intelligence in general and the own specialisation in particular.
2. The ability to determine the feasibility of a proposal to lead to a solution or design as specified.
3. The ability to contribute autonomously and with minimal supervision to an interdisciplinary project team and to profit from the abilities, the knowledge, and the contributions of other team members.
4. The ability to choose, apply, formulate, and validate models, theories, hypotheses, and ideas from the key areas of Artificial Intelligence.
5. The ability to submit an argument in the exact sciences (or humanities) to critical appraisal and to incorporate its essence in the solution of Artificial Intelligence problems.
6. The ability to translate academic knowledge and expertise into social, professional, economic, and ethical contexts.
7. Awareness of, and responsibility concerning, the ethical, normative and social consequences of developments in science and technology, particularly resulting from original contributions.

5.2.3 Making Judgments

The master is able to formulate an opinion or course of action on the basis of incomplete, limited and in part unreliable information.

Qualifications:

1. Competence in the search and critical processing of all sources of information that help to solve an open and illdefined problem.
2. The ability to demonstrate a professional attitude conform the (international) scientific conduct in Artificial Intelligence.
3. The ability to provide and receive academic criticism conform the standards in one specialism of Artificial Intelligence-research.
4. The ability to formulate an opinion and to make judgments that include social and ethical responsibilities related to the application of one's own contributions.
5. The master is able to judge the quality of his or her work or the work of others from scientific literature.

5.2.4 Communication

The Master can communicate information, ideas, problems and solutions to audiences of specialist in (other) research areas and to a general audience.

Qualifications:

The Master has academically appropriate communicative skills; s/he can:

- Communicate original ideas effectively in written form,
- Make effective oral presentations, both formally and informally, to a wide range of audiences
- Understand and offer constructive critiques of the presentations of others.

5.2.5 Learning Skills

The master has developed those learning skills that are necessary for a successful further career at the highest professional level. The master is able to detect missing knowledge and abilities and to deal with them appropriately.

Qualifications:

1. Being able to reflect upon one's competences and knowledge and, if necessary, being able to take the appropriate corrective action.
2. The ability to follow current (scientific) developments related to the professional environment.
3. Showing an active attitude towards continued learning throughout a professional career.

6 | INTERNATIONAL PERSPECTIVE

As stated in the introduction, this frame of reference is intended not only for the Dutch national context, but also to put the Dutch Artificial Intelligence programmes into an international perspective, and possibly to serve as a starting point for an internationally agreed frame of reference. The latter possibility is of course dependent upon international debate and agreement, and at this moment it is not clear how to bring this about, or whether it will in fact be possible. What we can and will do in this document is provide a comparison between the frame of reference as developed in the previous sections and a number of known related study programmes in other countries. In doing this, we hope to show that the developed frame of reference is up to par from an international perspective as well as the Dutch national one.

Having said this, we must immediately recognise that the Dutch national context appears to be rather special in that we only know of specialised bachelor-level Artificial Intelligence study programmes at one university outside the Netherlands, namely at Edinburgh (United Kingdom), which have a rather different programme structure than the Dutch (and general European) one. In our discussion of the Dutch frame of reference in international perspective, we will therefore add to our comparison with the Edinburgh study programme by a comparison with bachelor's programmes of study programmes in a related field, notably Cognitive Science. Furthermore, we will compare the Dutch bachelor's qualifications with the requirements for enrolment in Artificial Intelligence master programmes in other countries.

A comparison of master programmes is tricky as well. Although, contrary to bachelor's programmes, there are several well-known specialised Artificial Intelligence master programmes outside the Netherlands, study programmes at the master level are much more divergent than at the bachelor's level. A comparison can therefore only be provided in global, subject-independent, terms.

We have drawn up both the bachelor's and master's degree programme comparisons based on the programme descriptions and course lists received from the involved Universities. However, for the purpose of conciseness, we have left out particular details of the programmes that are largely time-dependent and often change from year to year.



6.1 Comparison of Bachelor's Programmes

6.1.1 The Artificial Intelligence Bachelor's Programme in Edinburgh

Edinburgh University (United Kingdom) offers a range of bachelor's degrees related to Artificial Intelligence, one of them in Artificial Intelligence as such, the others in combination with other disciplines (AI & Computer Science, Cognitive Science). An ordinary bachelor's degree consists of 4 years. In order to compare this system with the European standard of a 3-year bachelor's and a 2-year master's programme, we will take the fourth year of the Edinburgh bachelor's programme to be equivalent to the first year of a 2-year master's degree in other European countries, and base our comparison of bachelor's programmes on the first three years.

It should be pointed out that the (first three years of the) AI-related bachelors in Edinburgh show a large variation between them, and an extensive amount of (usually restricted) choices for particular courses within them. In fact, the commonality between the Edinburgh Artificial Intelligence bachelors is smaller than commonality within the Dutch framework. It seems that the wide variation in Edinburgh Artificial Intelligence related bachelor's degrees actually means that the degrees themselves are much more specialised than the Dutch framework proposes, some of them having little or no (cognitive) psychology, others having no mathematics, etcetera.

6.1.2 The Cognitive Science Bachelor's Programme in Osnabrück

The University of Osnabrück (Germany) offers a three-year (180 EC) bachelor's programme in Cognitive Science. The discipline of Cognitive Science is related to Artificial Intelligence, and may in fact be seen as a flavour of Artificial Intelligence, focused somewhat more towards Cognitive Psychology, and somewhat less towards Engineering. The same key knowledge and skills apply in Artificial Intelligence and in Cognitive Science.

Based on studying both programmes, we conclude that the Dutch frame of reference recognises the same AI-specific areas as both Cognitive Science programmes outside the Netherlands. The Dutch frame of reference devotes as much or more attention to any of these areas as any of those Cognitive Science programmes, with the exception of Cognitive Psychology in Linköping. Moreover, the recognition, in the Dutch frame of reference, that each individual study programme has a specific profile in addition to the communal areas appears to hold for both inspected study programmes outside the Netherlands as well.

6.1.3 The Symbolic Systems Bachelor's Programme in Stanford

The University of Stanford offers a programme in Symbolic Systems that has a variant in Artificial Intelligence. The list of core requirements of this programme includes, but is not strictly limited to: single and multivariable calculus, probability theory and statistics, discrete fundamentals, programming, philosophy, cognition and neuroscience, natural language and computation and cognition. Students in Artificial Intelligence should also take courses from the topics of knowledge representation and reasoning, natural language processing, learning and robotics and vision. They offer several courses in these topics. They offer a more in-depth application of numerous of these topics as non-core cognate courses, such as machine learning, motion planning, modal logic, automated reasoning, and more advanced levels of philosophy/linguistics (in reality, they offer a wide variety of supplemental courses).

The course units that are denoted are the core of Symbolic Systems and are supplemented by their Artificial Intelligence variant are very similar to the Dutch framework of reference - all of the topics in the core list of Stanford's programme + variant are reflected in our common core to some degree. There seems to be a slightly bigger focus on Philosophy (3/12th of the Symbolic Systems bachelor's core consists of philosophical foundations). The nature of the setup of compulsory core courses in Stanford does allow students to 'somewhat' skip topics that are important in the eyes of the Dutch framework - for example, the framework in Stanford calls for knowledge on 'Computation and Cognition', which ask of the student to take one course from a list ranging from Theoretical

Neuroscience to Neural Networks and Machine Learning. All in all, the frameworks are similar, but the core of Symbolic Systems appears to be a bit less technical than the Dutch framework.

6.1.4 The Bachelor's Programme Artificial Intelligence at Carnegie Mellon

The Carnegie Mellon University in Pittsburgh has introduced in Fall 2018 the first full Artificial Intelligence bachelor programme of the USA. Their curriculum consists of three cores: a mathematical, computer science and artificial intelligence core. The artificial intelligence core consists of Introductions in AI Representations, Problem Solving, Machine Learning, Natural Language Processing and/or Computer Vision. As electives, one course has to be selected from four clusters: Decision Making & Robotics, Machine Learning, Perception & Language and Human-AI Interaction.

The main difference with the Dutch Framework is the role of Logic; in this framework Propositional Logic and Predicate Logic are explicitly mentioned as support modules, in Pittsburgh logic is part of the Decision Making & Robotics cluster, with courses as Strategic Reasoning for AI and Planning Techniques for Robotics. This is a far more practical approach compared to the theoretical approach of the Dutch Framework.

6.2 Comparison of Master's Programmes

6.2.1 The Artificial Intelligence Master's Programme in Edinburgh

The Artificial Intelligence master programme in Edinburgh spans a full 12-month period and consists of two parts: taught and research. During the taught part (8 months), lectures, tutorials and group practicals are followed. The research part (4 months) consists of a major individual research project on which a dissertation is written. There is also the option of completing only the taught part, in which case, a Diploma will be awarded. MSc courses in Artificial Intelligence in Edinburgh are grouped in four major areas of specialisation:

- Intelligent robotics
- Agents, Knowledge and Data
- Machine Learning
- Natural language processing

Comparing the Edinburgh programmes to the Dutch frame of reference, we can draw the following conclusions:

- The main Artificial Intelligence topics that are in the Dutch framework are also represented in the Edinburgh programmes (as shown in the four different identified areas of specialisation).
- The Edinburgh programmes are 1-year, whereas most Dutch Artificial Intelligence master programmes are 2-year programmes. However, the Edinburgh master programme requires a 4-year honours bachelor's degree.
- The Edinburgh programme knows relatively little study load for practical work. Whereas the minimum length of a Dutch master-thesis ('afstudeerproject') is 30 ECs (half a year), the Edinburgh programme has 4 months for doing practical assignments.
- However, the practical work seems to be more research oriented, whereas in the Dutch programme there is also the option to do a final project in industry.
- The Edinburgh program has an entry requirement on mathematics (During the bachelor degree 60 credits have completed of mathematics.)

6.2.2 The Machine Learning and Machine Intelligence Master's Programme in Cambridge

At the University of Cambridge the master is called Machine Learning and Machine Intelligence. It is a very selective (20 places) two year programme (120 ECTS credits). To apply, the applicants should have a UK First class Honours Degree (equivalent with overall grade of 8/10).

Their programme includes courses such as:

- Deep Learning and Structured Data
- Probabilistic Machine Learning
- Speech Recognition



- Weighted Automata
- Computer Vision

Comparing the Cambridge study programme to the Dutch frame of reference, we can draw the following conclusions:

- the programme has clear programme objectives (9 in total)
- the programme only covers 3 of the 8 core areas (Computational Linguistics, Perception, Machine Learning)
- There is not much choice: the programme consists of eleven core modules
- the programme has only two optional modules: Computer Vision or Natural Language Processing
- As elective the students can pick one elective module from 4th year undergraduate

6.2.3 The Symbolic Systems and Computing Science Master's Programmes in Stanford

Stanford offers two variants that are similar to Artificial Intelligence in the Netherlands: the Symbolic Systems Master and an Artificial Intelligence variant in the Computer Science degree programme. Symbolic Systems is an interdisciplinary programme that combines Computer Science, Psychology, Philosophy, and Linguistics in order to better understand cognition in both humans and machines. Viewing people and computers as symbol processors, the Symbolic Systems programme explores the ways computers and people reason, perceive, and act. Within the Symbolic Systems major, there is a core set of required classes with respect to the aforementioned fields; beyond this core, students choose an area of concentration in order to gain depth.¹¹

The variant that is actually named Artificial Intelligence offers a subset of the disciplines in the Dutch framework, mostly those related to Computer Science. This master's programme focuses on solving problems using computers, and as such mainly contains course units that address topics in Machine Learning, Computer Vision, Natural Language Processing, Data Mining, Robotics and Bioinformatics. There is a framework of courses students can choose from, but this framework is fairly loosely defined (with only very few guidelines), making it hard to distinguish a core beyond the specific attention for the applied computational side of Artificial Intelligence.

Comparing the Stanford study programme to the Dutch frame of reference can only be done by evaluating both the Symbolic Systems Master and the Computer Variant at the same time:

1. The Dutch framework has more formal subjects (such as a stronger focus explicit focus on logic and computer science related topics) than the Symbolic Systems programme. The Dutch framework has more topics related to Cognitive Science and the context of Artificial Intelligence than the Computer Science AI variant.
2. It has already been mentioned that there is much variety between the master programmes
 - Both in the Netherlands and abroad. This is also the case for the programmes at Stanford.
 - The Stanford programmes seem to have a lot of freedom in their choice for electives. In other words, the core of compulsory courses is limited and students have to select many elective courses - thus rejecting the idea behind a broad common core of the programme.

6.2.4 The Cognitive Science Master's Programme in Osnabrück

The university of Osnabrück offers a two-year (120 ec) master of science programme in Cognitive Science. Apart from a 30 ec individual thesis project and a 22 ec team project, students take predominantly elective courses in Cognitive Psychology, Artificial Intelligence, (Computational) Linguistics, Robotics, Neuroscience, Neuroinformatics and Philosophy of Mind and Cognition. The choice of courses largely follows the research group structure. The programme offers a double degree option together with the Cognitive Science programme in Trento.

The Osnabrück programme has a similar focus on cognitive (and neuroscience) aspects of AI as Nijmegen. Nijmegen has a tradition of attracting graduated bachelor students from Osnabrück who are typically directly admitted to the master programme in Nijmegen; in addition, several graduated master students from Osnabrück found a PhD position at the Radboud University. Student mobility

¹¹ <http://symsys.stanford.edu/courses> (last visited on May 4, 2018)

thus suggests that the programmes are relatively comparable in content and quality to the Dutch Framework.

7 | NATIONAL PERSPECTIVE

At its core, the framework of reference serves as a foundation for every Artificial Intelligence degree programme in the Netherlands. This is also what differentiates the degree programmes that are a member of the KION from similar degree programmes: at its very heart, the KION programmes aim to provide a broad foundation that respects the roots of the field, rather than merely offer a specialistic application of Artificial Intelligence. With the growing popularity and diversity of Artificial Intelligence as a field, the need for programmes to adhere to a framework founded in both history and modern-day application is more important than ever.

We have drawn up an analysis that compares AI-related degree programmes in the Netherlands, to see whether they fit the frame of reference. In this way, we hope to highlight the unique position the KION degree programmes hold in the field of Artificial Intelligence: both with respect to delivering broadly educated Artificial Intelligence students, but also with respect to the consistency of the quality and background future employers can expect.

7.1 Bachelor's Programmes

Tilburg University and TU Eindhoven offer a joint BSc degree in Data Science. The focus of this bachelor's degree is on mathematical tools for dealing with big data; there is less emphasis on programming and no broad introduction in Artificial Intelligence. Tilburg University also offers a bachelor in Cognitive Science and Artificial Intelligence; this is part of the School for Humanities. Currently, the programme is under development; the relationship with the KION frame of reference is still open at the time of writing. Finally, TU Delft offers a bachelor Computer Science & Engineering (intelligent data analysis variant). This programme does not offer the broad introduction in Artificial Intelligence that the KION degree programmes do.

7.2 Master's Programmes

Several universities offer a specialisation in Data Science as part of the Computer Science or Information Studies MSc degree. Tilburg University and TU Eindhoven offer a joint MSc degree Data Science and Entrepreneurship. TU Delft offers a MSc degree in Data Science & Technology. These master programmes do not offer the broad overview on AI that the KION degree programmes do, but focus specifically on data science. Tilburg University also offers a one-year MSc degree (in contrast to the two-year KION programmes) in Cognitive Science and Artificial Intelligence.

8 | CONCLUDING REMARKS

Artificial Intelligence is a rapidly developing field. The term Artificial Intelligence does not have the stature of Physics, Psychology, or even Computer Science, due to its relatively recent start as a coherent field of research. Despite this, the recent attention for the successes of the field have ascertained that the field of Artificial Intelligence has made its mark. Internationally, the study of natural and artificial intelligence with computational means is firmly embedded in the fabric of modern Universities.

Modern topics such as gaming, ambient intelligence, ambient awareness, and believable-agent systems are fashionable manifestations of Artificial Intelligence and these and future fashionable spin-offs of Artificial Intelligence will increasingly affect humans. A particularly popular focus of Artificial Intelligence can currently be found in the field of Machine Learning, where possibilities for data analysis provide the world with useful patterns in all sorts of data. These possibilities have helped popularise Artificial Intelligence in the media, but have also raised ethical and legal questions about the field. Future challenges will force products, services, and societies to react faster but remain reliable, to be both flexible and effective, be both efficient and versatile, and to utilise natural resources with maximal benefit. The biggest challenge of all is making the most of this combination of conflicting demands, a challenge that lies very much at the core of in the concept of intelligence.



The Dutch situation is special because of the existence of Artificial Intelligence bachelor's and master's programmes on most of the general universities. This offers the Netherlands a competitive advantage, consistent with its main economic strategy to remain one of the leading "knowledge intensive" economies. This frame of reference explicates how the bachelor's and master's programmes in Artificial Intelligence of Dutch universities contribute to educate alumni that will take a leading role in meeting these future challenges.

APPENDIX 2: INTENDED LEARNING OUTCOMES

Code	Exit Qualifications
1	Has thorough knowledge of the current theories, methods and techniques in the field of Artificial Intelligence
2	Has specialized knowledge of at least one of the following Artificial Intelligence subfields:
2A	Machine Learning
2B	Computer Vision
2C	Deep Learning
2D	Natural Language Processing
2E	Fairness, Accountability, Confidentiality and Transparency in AI
2F	Information Retrieval
2G	Symbolic Systems
3	Has the capability to apply this knowledge to analyse, design and develop AI-systems
4	Can formulate scientific questions and is able to solve problems with the aid of abstraction and modelling
5	Is able to contribute to further developments of the theories, methods and techniques of AI in a scientific context
6	Is able to express him/herself clearly on a technical/mathematical and general level
7	Is aware of the social context and consequences of conducting AI research and work
8	Can obtain an academic position at a university or research centre or scientific/applied position in the industry

Courses	Catalog #	Exit Qualifications															
		1	2	2A	2B	2C	2D	2E	2F	2G	3	4	5	6	7	8	
Computer vision 1	52041COV6Y																
To become familiar with both the theoretical and practical aspects of computer vision and to describe the foundations of image formation, measurement, and analysis																	
To learn physical-, statistical- and deep learning models as the foundation of computer vision and apply them to computer vision problems																	
To learn the fundamental aspects of different computer vision models and to explain their advantages and shortcomings																	
Is able to implement computer vision programs and can evaluate/test these programs to understand the limitations of these vision algorithms																	
To learn how to develop the practical skills necessary to design, model, build, organize, and execute computer vision applications																	
Machine Learning 1	52041MAL6Y																
The student can explain, motivate and distinguish the main areas of machine learning in general and on examples.																	
The student can explain the major statistical learning frameworks/principles together with their advantages and shortcomings.																	
The student knows the major linear and non-linear statistical models together with their advantages and disadvantages, can explain them and the made model assumptions, can reason about and inside them, and can manipulate them.																	
The student can set up learning objectives with the taught models, can train and evaluate them and can assess the quality of fit.																	
The student can implement all the above in Python and apply the learned principles and models to real world problems and data sets.																	
Deep Learning	5204DLFV6Y																
During theory the students will be taught the fundamentals of deep learning and the latest, state-of-the-art models that																	



Courses	Catalog #	Exit Qualifications														
		1	2	2A	2B	2C	2D	2E	2F	2G	3	4	5	6	7	8
empower popular applications, such as Google Photos, Google Translate, Google Text-to-Speech, Google Brain, Facebook Friend Finder, Self-driving cars, etc.																
During the theory part students will learn to re-implement similar models, as well as to build novel ones for new tasks																
During the practicals the students will implement from A to Z the core versions of some of the aforementioned applications. The course will focus on state-of-the-art programming frameworks																
Students will learn what are the most relevant and frequent practical problems, and how they can be addressed in practice																
During the projects, the students will further be able to build real-life, interesting and perhaps even marketable applications. Hopefully and depending on the student's motivation, the applications could lead to demos that can be publicly presented																
Natural Language Processing 1	52401NLP6Y															
Describe computational modelling methods for several levels of language analysis (morphology, syntax, semantics and discourse)																
Discuss the strengths and limitations of these methods																
Construct language processing models for several tasks, such as word representation learning, sentence representation learning, text classification etc.																
Implement supervised (and some unsupervised) estimation procedures for these models																
Evaluate these models experimentally and analyse their performance																
Use the above techniques in NLP applications, such as sentiment analysis																
Fairness, Accountability, Confidentiality and Transparency in AI	5204FACT6Y															
Students can explain, motivate and distinguish the main types of algorithmic harm, in general and in terms of concrete examples in AI is being applied.																
The student can explain the major notions of fairness, accountability, confidentiality, and transparency that have been proposed in the literature, with their strengths and weaknesses																
Be familiar with state-of-the-art algorithmic approaches to fairness, accountability, confidentiality and transparency																
Identify situations in which there is a potential for algorithmic harm and assess the degree to which state-of-the-art algorithmic solutions are effective in addressing the potential harm.																
Information Retrieval 1	52041INR6Y															
Describe, explain, and compare algorithms																
Transform theoretical models to a working system																
Describe and compare different evaluation methods																
Use different evaluation methods to perform experiments																
Analyze results, compare different algorithms, and draw conclusions																
Make decisions and adapt model parameters on the basis of experimentation																
Modify algorithms to suit other information access tasks																
Identify open problems and come up with novel solutions																
Symbolic Systems 1	52042SYS6Y															
Students are expected to master the techniques used in the design of a variety of symbolic systems																
Students develop the ability to apply logic in the context of modelling specific multi-agent systems																
Deep Learning for Natural Language Processing	5204DLNLY															
Gain a good understanding of the deep learning architectures that are used for a spectrum of NLP problems																
Gains practical knowledge in developing one or two models or model variations to address a concrete NLP problem																
Learns how to evaluate and analyze a deep learning model and its performance for specific NLP problems and based on this analysis learns how to formulate scientific questions																
Learns how to present and motivate deep learning modeling choices for specific NLP problems																

Courses	Catalog #	Exit Qualifications															
		1	2	2A	2B	2C	2D	2E	2F	2G	3	4	5	6	7	8	
Information Retrieval 2	52042INR6Y																
At the end of the course, the student is able to understand, experiment with and evaluate advanced retrieval methods and models																	
Machine Learning 2	52042MAL6Y																
Gain an advanced level of understanding of the principles of machine learning																	
Acquire the skills to apply machine learning to complex real world problems																	
Use advanced machine learning techniques to analyze complex data and evaluate the resulting models																	
Reinforcement Learning	5204RELE6Y																
The student is able to describe the main algorithms in Monte Carlo, temporal difference, and model-based and policy-based methods																	
The student is able to describe the main differences between on&off policy learning, TD&Monte Carlo methods, value-based and policy-based methods, tabular and approximate methods; and categorise reinforcement learning methods according to these properties																	
The student is able to implement reinforcement learning algorithm; and able to analyse its performance (or lack thereof) in a given environment. The student is able to apply the learned update rule on given data sets																	
The student is able to compare reinforcement learning algorithms and point out their main advantages and disadvantages. The student is able to choose which reinforcement learning algorithm to select based on characteristics of the given environment																	
The student is able to design experiments to compare reinforcement learning techniques on a given environment. The student is able to critically evaluate reinforcement learning experiments and point out their strong points and weak points																	
Information Theory	5314INTH6Y																
Compute the probability of an event using the most common discrete probability distributions (Bernoulli, binomial, and geometric)																	
Compute inverse probabilities using Bayes' rule																	
Compute the means and variances of commonly used probability distributions																	
Compute the means and variances of sums or products of a random variables with known distributions																	
Bound the probability of an extreme event using inequalities such as the Markov bound, Chebyshev's inequality, or Hoeffding's inequality																	
Compute the entropy of a random variable																	
Compute the mutual information between two random variables																	
Use entropy diagrams to reason about the relative size of the entropies, conditional entropies, and mutual information of two or three random variables																	
Use Jensen's inequality to bound the mean of a random variable defined in terms of convex or concave function of another random variable																	
Construct a d-ary Huffman code for a random variable																	
Use Kraft's inequality to check whether a prefix-free code can be constructed to fit certain codeword lengths																	
Bound the possible rate of lossless compression of output from a given source using Shannon's source coding theorem																	
Define a typical set and reason about its size, probability, and elements																	
Compute the Shannon-Fano-Elias codeword for a sample from a stochastic process																	
Compute the entropy rate of a Markov process																	
Construct a probability model of a communication channel given a verbal description																	
Compute the channel capacity of a channel																	



Courses	Catalog #	Exit Qualifications														
		1	2	2A	2B	2C	2D	2E	2F	2G	3	4	5	6	7	8
Use Shannon's channel-coding theorem to bound the achievable rate of reliable communication over a channel																
Use Bayes' rule to decode corrupted messages sent using an error-correcting code																
Evaluate the rate and reliability of such codes																
Define the jointly typical sets of a source and channel, and use such sets to decode outputs from the channel																
Draw the confusability graph of a given channel, and describe the channel depicted by a given confusability graph																
Compute the independence number and the zero-error capacity of a confusability graph																
Collaborate (under guidance of teachers) in small groups of 3 to 5 (unknown) students on practice problems. In particular, students who are already well-versed in certain topics or skills are encouraged to help their fellow group members in this respect, while weaker students are encouraged to seek help from their peers and the teachers whenever necessary																
Present solutions of mathematical problems orally to peers																
Game Theory	5314GATH6Y															
To become familiar with the basic concepts of noncooperative game theory, including the models of and solution concepts for normal-form games, Bayesian games, extensive games, and imperfect-information games.																
To become familiar with the basic concepts of economic mechanism design as well as their application to the design of simple and combinatorial auctions.																
To become familiar with the basic concepts of cooperative game theory, including the models of and solution concepts for transferable-utility games, hedonic games, and matching markets.																
To acquire the ability to model instances of strategic interaction as well as cooperation between several autonomous agents by using an appropriate game-theoretical formalism.																
To appreciate the usefulness as well as limitations of game theory for the analysis of the strategic behaviour of agents operating in the real world.																
Information Visualization	5204INVI6Y															
Categorize different information visualization techniques, understand their characteristics and knowing when to apply them																
Understand the algorithms underlying information visualization techniques																
Develop Information Visualizations																
Understand visual analytics models																
Gather, select, structure, and visualize data in a structured way																
Use knowledge of perception and cognition to make effective visualizations																
Use design guidelines (e.g. dashboards, aggregation, multiview) for developing effective visualizations																
Design interaction models																
Design of information visualization evaluation schemes																
Communicate data-driven results using storytelling based mechanisms																
Computational Dialogue Modelling	5314CODM6Y															
To become familiar with fundamental concepts in the study of dialogue																
To gain insight into current research directions in computational dialogue modelling																
To acquire the ability to evaluate and contribute to research in computational dialogue modelling																
Computational Social Choice	5314COSC6Y															
To become familiar with fundamental concepts in classical social choice theory																
To gain insight into current research directions in modern computational social choice																

Courses	Catalog #	Exit Qualifications														
		1	2	2A	2B	2C	2D	2E	2F	2G	3	4	5	6	7	8
To acquire the ability to evaluate and contribute to research in computational social choice																
Computer Vision 2	52042COV6Y															
Students will have advanced knowledge of the current theories, methods and techniques in the field of computer vision																
Student will be able to understand and explain state-of-the-art computer vision articles																
Students are able to implement advanced computer vision programs and can evaluate/test these programs to understand the limitations of these vision algorithms																
Students acquire the ability to contribute to academic computer vision research																
Natural Language Processing 2	52042NLP6Y															
Acquire knowledge of advanced NLP techniques, particularly in machine translation																
Acquire awareness of ongoing research and challenges in NLP in general																
Practice reading and presenting scientific articles in NLP																
Develop, program and report basic solutions to example problems within multiple small projects																



APPENDIX 3: OVERVIEW OF THE CURRICULUM

Programme





















Content and study Load of the Components

The MSc AI programme starts with 7 core AI courses (42 EC) in periods 1-6 of the first year. The core curriculum is obligatory for all MSc AI students.

The core courses are: Computer Vision 1, Machine Learning 1, Deep Learning, Natural Language Processing 1, Fairness Accountability Confidentiality and Transparency in AI, Information Retrieval 1, Symbolic Systems 1.

Furthermore, a number of restricted-choice electives are required (at least 18 EC) followed by free-choice electives or additional restricted-choice electives (12 EC).

The programme is completed with the Master Thesis AI (48 EC)

	Sem. 1	Sem. 2	EC
> Computer Vision 1			6
> Machine Learning 1			6
> Deep Learning			6
> Natural Language Processing 1			6
> Fairness, Accountability, Confidentiality and Transparency in AI			6
> Information Retrieval 1			6
> Symbolic Systems 1			6
> <i>Restricted-choice electives: Constrained choice</i>			18
> Master Thesis AI			36
> <i>Free-choice electives: Suggested elective courses</i>			0

APPENDIX 4: PROGRAMME OF THE SITE VISIT

9 December 2019

- 08.30 – 09.00 Arrival committee and welcome
- 09.00 – 10.00 Internal deliberation by the committee
- 10.00 – 11.00 Board and programme management Bachelor KI and Master AI
- 11.00 – 11.15 Internal deliberation by the committee
- 11.15 – 12.00 Students bachelor and alumni KI
- 12.00 – 12.15 Internal deliberation by the committee
- 12.15 – 13.00 Didactical staff bachelor KI (English)
- 13.00 – 14.00 Lunch
- 14.00 – 14.45 Tour Robolab
- 14.45 – 15.00 Internal deliberation by the committee
- 15.00 – 15.45 Exam committee (English)
- 15.45 – 16.00 Internal deliberation by the committee
- 16.00 – 17.00 Students and alumni master AI and advisory board (English)
- 17.00 – 17.15 Internal deliberation by the committee

10 December 2019

- 08.30 – 09.00 Arrival committee
- 09.00 – 09.30 Internal deliberation by the committee
- 09.30 – 10.00 Open consultation
- 10.00 – 10.45 Didactical staff master AI (English)
- 10.45 – 11.00 Internal deliberation by the committee
- 11.00 – 11.45 Interview formeel verantwoordelijken
- 11.45 – 13.45 Opstellen oordelen + lunch
- 13.45 – 14.00 Feedback by the committee
- 14.00 – 14.15 Short break
- 14.15 – 15.00 Development dialogue
- 15.00 – 15.15 Wrap up

APPENDIX 5: THESES AND DOCUMENTS STUDIED BY THE PANEL

Prior to the site visit, the panel studied 15 theses of the master's programme Artificial Intelligence. Information on the selected theses is available from QANU upon request.

During the site visit, the panel studied, among other things, the following documents (partly as hard copies, partly via the institute's electronic learning environment):

- Self-evaluation Report
- Assessment Plan Graduate School of Informatics Masters
- KION frame of reference
- Staff overview
- Rapport Curriculum Commissie
- Rapport Evaluatie Studiesucces
- Teaching and Examination Regulations
- Description of Course components
- Exam protocol
- Rules and guidelines of the Examinations Board
- Annual reports Examinations Board
- Annual reports Programme Committee
- Assessments dossiers and study materials of courses, including Study guides, assessment, results, and evaluation results of the following courses:
 - o Machine Learning 1
 - o Information Retrieval 1
 - o Natural Language Processing 1