

Oplysninger fra CVR

CVR-nummer

29057753

Virksomhedens navn

IT-Universitetet i København

Virksomhedens adresse

Rued Langgaards Vej 7, 2300 København S

0. Ansøgningens type

Der ansøges om

- ☐ Ansøgningsskema 1 bacheloruddannelse 180 ECTS-point
- ☐ Ansøgningsskema 2 kandidatuddannelse 120 ECTS-point
- ☐ Ansøgningsskema 3 kandidatuddannelse 75 ECTS-point
- ☐ Ansøgningsskema 4 erhvervskandidatuddannelse 75-120 ECTS-point
- ☐ Ansøgningsskema 5 erhvervskandidatuddannelse som kombinationsuddannelse (1+2) 120 ECTS-point
- ☒ Ansøgningsskema 6 erhvervskandidatuddannelse som kombinationsuddannelse 75-120 ECTS-point

Ansøgningsrunde

2026-2

1. Grundoplysninger for uddannelsen

Navn på institution

IT-Universitetet i København

Udbudsted

2300 København S

Bynavn

København S

Angiv evt. supplerende bynavn

Uddannelsestype

Erhvervskandidat som kombinationsuddannelse

Fagbetegnelse dansk

Ansvarlig kunstig intelligens

Fagbetegnelse engelsk

Responsible AI

Betegnelse dansk (titel)

Kandidat ([cand.scient.](#)) i ansvarlig kunstig intelligens

Betegnelse engelsk (titel)

Master of Science (MSc) in Responsible AI

Hvilket hovedområde hører uddannelsen under?

Naturvidenskab

Vælg undervisningssprog

- ☐ Dansk
- ☒ Engelsk

Forslag til censorkorps. Angiv hvilket censorkorps, som institutionen ønsker at knytte til uddannelsen

Datalogi

Hvilke adgangskrav gælder til uddannelsen?

Uddannelsen er åben for studerende, der har en bacheloruddannelse med grundlæggende viden inden for datalogi og matematik.

Ansøgere skal kunne dokumentere mindst 75 ECTS-point i kurser inden for datalogi ([f.eks.](#) programmering, softwareudvikling, matematik, algoritmer, databaser, machine learning og relaterede emner).

Uddannelsen er åben for både danske og internationale studerende, der har gennemført engelsk på B-niveau (minimum 3 i snit).

Eksempler på bacheloruddannelser, der er direkte adgangsgivende inkluderer, men er ikke begrænset til:

- Data Science (ITU, AU)
- Softwareudvikling (ITU)
- Datalogi (KU, AU, AAU, SDU)
- Machine Learning og datavidenskab (KU)
- Softwareteknologi (DTU)
- Kunstig intelligens og data (DTU)
- Artificial Intelligence (SDU)

Der vil ikke være nogen retskravsbachelor til Responsible AI.

Beskriv uddannelsens formål og erhvervssigte

Kandidatuddannelsen i Responsible AI har til formål at uddanne kandidater, der kan designe, evaluere og anvende AI-systemer på en teknisk og socialt forsvarlig måde. Uddannelsen giver de studerende et solidt fundament inden for maskinlæring med praktiske kompetencer i datakonstruktion, systemevaluering og interaktionsdesign. Hertil klæder uddannelsen de studerende på til at forstå og navigere i de juridiske, sociale og organisatoriske rammer, AI-systemer indgår i. Uddannelsen henvender sig således til studerende, der vil forbinde teknisk AI-udvikling med ansvarlig og virkelighedsnær anvendelse.

Offentlige og private virksomheder anvender i stigende grad AI-systemer på tværs af sektorer og opgaver. Det skaber en stor og stigende efterspørgsel efter fagpersoner, der kan udvikle systemer, forberede og kvalitetssikre data, evaluere systemers performance og sikre ansvarlig implementering. Dimittenderne fra Responsible AI kvalificerer sig netop til stillinger såsom data scientist og AI-ingeniør samt roller inden for datakvalitet og AI-governance.

Angiv hvilken struktur og konstituerende faglige elementer uddannelsen består af.

Uddannelsen er normeret til 75 ECTS og tilrettelagt som et forløb fordelt på tre semestre. Den består af fem obligatoriske kurser (i alt 45 ECTS) og et valgfrit element på 15 ECTS. Uddannelsen afsluttes med en hovedopgaven på 15 ECTS, der skrives på deltid på erhvervskandidatforhold.

Første semester udgøres af de obligatoriske kurser Foundations of Artificial Intelligence (15 ECTS), Dataset Creation (7,5 ECTS) og Evaluating Systems (7,5 ECTS).

Andet semester indbefatter de obligatoriske kurser Human-AI Interaction (7,5 ECTS) og AI in Society (7,5 ECTS) samt et valgfrit element (15 ECTS).

Som valgfrit element kan de studerende vælge mellem en række af valgfag (se bilag) eller et research project i en virksomhed (med tilknyttet ITU-vejleder samt bedømmelse). Således kan de studerende tone deres uddannelsesmæssige profil med kompetencer indenfor fx IT-sikkerhed, grøn omstilling, virksomhedsforståelse, spildesign mm.

Den afsluttende hovedopgave skrives enten som individuel opgave eller i grupper.

På den sidste del af uddannelsen læser de studerende på deltid samt er i et virksomhedsforløb. I denne periode er de studerende i beskæftigelse i 25 timer om ugen.

Kassogram og kursusbeskrivelser findes i bilaget (Uddannelsens struktur og konstituerende fagelementer).

Angiv antal ECTS-point for hele uddannelsen

75

Angiv normeret studietid i mdr.

Anden normeret studietid

Angiv (anden) normeret studietid i mdr.

17

Her angives beskæftigelseskravet i timer pr. uge (årsnorm)

25

Angiv antal ECTS-point på heltidsforløb (med SU)

60

Heltidskandidatuddannelse som erhvervskandidatuddannelsen giver ret til indskrivning på

Responsible AI (heltidskandidatuddannelse), som søges parallelt med nærværende kombinationsuddannelse, hvorfor den endnu ikke godkendt

Angiv den firecifrede UDD-kode fra Danmarks Statistik, som er tilknyttet den heltidskandidatuddannelse som erhvervskandidatuddannelsen giver ret til indskrivning på

0000

2. Øvrige krav til ansøgningen

Kort beskrivelse af hvordan aftagere og øvrige relevante interessenter er inddraget i processen omkring udvikling af uddannelsen.

Udviklingen af uddannelsen er sket i en løbende proces med inddragelse af både eksterne aftagere og interne interessenter.

Der er i løbet af foråret 2026 indhentet input fra aftagere gennem en spørgeskemaundersøgelse, som modtog 25 besvarelser. Der blev foretaget ni opfølgende, kvalitative interviews med udvalgte respondenter med erfaring indenfor blandt andet finanssektoren, sundhedssektoren, konsulentbranchen, softwareudvikling samt den offentlige sektor. Aftagerne fremhævede behovet for kandidater, der kan kombinere tekniske AI-kompetencer med praktisk forståelse for implementering, compliance og kommunikation. Denne feedback førte bl.a. til, at adgangskravene blev breddet ud til også at omfatte softwareorienterede bacheloruddannelser og til en styrket vægtning af tekniske kompetencer i profilen.

Uddannelsen er desuden blevet drøftet i en række relevante interne fora med deltagelse fra både VIP, TAP, studerende og ledelse. De interne drøftelser bidrog til at skærpe uddannelsens profil og adgangskrav og til at styrke balancen mellem tekniske AI-kompetencer og bredere organisatoriske, samfundsmæssige og regulatoriske perspektiver.

Hertil var uddannelsen en del af koordineringsrunden i Danske Universiteters Uddannelsespolitiske udvalg (RUP).

3. Baggrundsoplysninger til brug for sagsbehandling

Henvisning til evt. tidligere godkendt heltidsuddannelse og dennes takstindplacering

Forslag til takstindplacering

Heltidstakst 3

Begrundelse for takstindplacering

Uddannelsens formål er at give de studerende et stærkt teknisk fundament inden for maskinlæring og kunstig intelligens, så de kan udvikle, evaluere og anvende AI-systemer i praksis. Dette fundament omfatter maskinlæring, datakonstruktion, systemevaluering og kvantitative metoder. Uddannelsens konstituerende kurser er alle teknisk forankrede, og de samfundsmæssige og regulatoriske perspektiver behandles som en integreret del i en teknisk kontekst, hvor de anvendes på konkrete AI-systemer. Uddannelsens kernefaglighed hører dermed under det naturvidenskabelige hovedområde.

IT-Universitetet foreslår på den baggrund, at uddannelsen indplaceres på heltidstakst 3 i lighed med universitetets øvrige data science- og AI-uddannelser.

Beskriv ligheder og forskelle til beslægtede uddannelser, herunder beskæftigelse og eventuel dimensionering

Flere danske kandidatuddannelser deler AI-fundament med Responsible AI, men adskiller sig i deres profil. Uddannelser i data science og datalogi, herunder ITU's egne kandidatuddannelser i Data Science og Computer Science, er brede 120 ECTS-uddannelser, hvor AI udgør ét område blandt flere. Responsible AI er derimod specialiseret med et fokus på datakonstruktion, evaluering af AI-systemer og vurdering af, hvordan modeller udvikles og valideres i praksis.

Andre uddannelser fokuserer på AI i en forretningskontekst, herunder CBS's kommende MSc in Artificial Intelligence for Business og ITU's kommende kandidatuddannelse i Business Analytics og kunstig intelligens. DTU's MSc in Human-Centered Artificial Intelligence er primært ingeniørrettet, mens SDU's MSc in Artificial Intelligence dækker generelle AI-emner.

Responsible AI vil dermed være den eneste 75 ECTS-kandidatuddannelse i Danmark, der er dedikeret til ansvarlig udvikling og anvendelse af AI-systemer. Efterspørgslen efter AI-kompetencer er samtidig høj og stigende på det danske arbejdsmarked, hvorfor beslægtede uddannelser har lav ledighedsprocent. Uddannelsen vil også klæde kandidaterne særligt godt på til stillinger, der arbejder med ansvarlig implementering af AI-systemer, sammenlignet med andre uddannelser.

Vælg forventet opstartsår

2028

Angiv forventet optag på uddannelsen det første år

35

Angiv forventet optag på uddannelsen det andet år

45

Angiv det forventede optag på uddannelsen det tredje år

50

Øvrige bemærkninger til ansøgningen?

Denne uddannelse bliver søgt parallelt med en tilsvarende heltidsuddannelse på 75 ECTS. Optagelsestallene ovenfor angiver den samlede pulje til optag på de to uddannelser. Den specifikke fordeling af pladser mellem de to uddannelser er endnu ikke fastlagt.

☒ ★Det bekræftes hermed at ansøgningen er godkendt af rektor

Navn på kontaktperson

Mads Vincent Hedegaard

Kontaktpersons e-mail

mvhe@itu.dk

Kontaktpersonens telefonnummer

72185431

Vedhæft samlet PDF-fil

Bilag_RAI_75_Kombi_Uddannelsens struktur og konstituerende elementer.pdf 98.1 KB

Bilag: Ansøgning til Responsible AI (kombi)

Uddannelsens struktur og konstituerende fagelementer

The description below is written in English as the programme is taught in English. All course titles and course descriptions are English by default.

Contents:

1. Overview	1
2. Programme Objectives	1
3. Programme Structure	3
4. Course Descriptions	4

1. Overview

The MSc in Responsible AI trains students to design, study, and deploy AI systems in ways that are technically and socially sound. The programme provides a strong foundation in machine learning, combined with practical competencies in dataset construction, systems evaluation, and interaction design. Students learn how data is produced, how claims about AI systems are justified, and how systems operate within legal, social, and environmental constraints.

Graduates understand that AI systems rely on more than algorithms, being shaped by how data is collected and curated, the people who produce and maintain it, the methods used to evaluate performance, and the governance, legal, and social frameworks in which the systems operate. They are therefore able to move between development and critical assessment, making them well suited for roles in industry and the public sector where AI systems must function reliably, transparently, and responsibly in real-world settings.

2. Programme Objectives

The purpose of the MSc programme in *Responsible AI* is to provide students with the scientific qualifications to identify, formulate, analyse, and solve complex problems

relating to the development, evaluation, and deployment of artificial intelligence systems.

The programme prioritises the student's ability to assess, apply, and develop machine learning methods, dataset construction practices, evaluation techniques, and design of interactive systems that underpin reliable and responsible AI systems.

On the background of the student's preceding bachelor's programme, the programme provides the student with the qualifications to define their own academic profile within artificial intelligence and to take independent responsibility for his or her professional development and specialisation.

Within the framework of the programme, the student can acquire the requisite qualifications for specialised positions in industry, the public sector, and organisations developing or governing predictive computing systems.

2.1 Objectives for Learning Output

The graduate will develop research-based knowledge and understanding of, and will be able to reflect on:

- Principles and methods of training machine learning models
- Processes of dataset construction, including collection, annotation, and documentation
- Methods for evaluating AI systems, both quantitative and qualitative
- Principles of human-AI interaction and the design of user-friendly AI systems
- Regulatory, organisational, and societal frameworks shaping the development and deployment of AI
- Theory, scientific methodology and scientific issues within responsible use of artificial intelligence in the above areas at the highest international research level

The graduate will develop the following research-based skills

- Master modern state-of-the-art AI frameworks to design, implement, and evaluate machine learning infrastructures
- Construct, document, and critically assess datasets, including identifying bias and data quality issues
- Design and conduct empirical evaluations of computing systems using appropriate methods
- Design and prototype AI systems that account for human desiderata

- Communicate, visualise, and discuss technical and analytical results to both academic peers and non-specialists

The graduate will develop the following research-based competences

- Design and deploy new solutions to enhance existing complex AI systems, and combine and select relevant analysis methods in complex and unpredictable contexts
- Critically assess the reliability, transparency, and societal implications of computing systems
- Design and carry out state-of-the-art empirical investigations of computing systems
- Independently initiate collaboration and work professionally across technical, organisational, and societal domains in complex and interdisciplinary contexts
- Independently take responsibility for own professional development and specialization

3. Programme Structure

1 st Semester (Fall)	
Foundations of Artificial Intelligence (15 ECTS) <i>Foundational machine learning approaches and domains.</i>	
Dataset Creation (7.5 ECTS) <i>Collection, annotation, documentation, and bias in training data.</i>	Evaluating Systems (7.5 ECTS) <i>Quantitative and qualitative evaluations of predictive systems.</i>
2 nd Semester (Spring)	
Electives (15 ECTS)	
AI in Society (7.5 ECTS) <i>Accountability and regulatory compliance, environmental considerations.</i>	Human-AI Interaction (7.5 ECTS) <i>Design of interactive human-AI systems and experiences.</i>
3 rd Semester (Summer/Fall)	
Final project (15 ECTS) / Job (25 hr./week)	

4. Course Descriptions

Foundations of Artificial Intelligence (15 ECTS)

This course covers contemporary approaches to machine learning, with a focus on deep learning and probabilistic modelling. The course introduces central model classes, learning paradigms, and optimisation techniques that underpin modern ML systems. Students examine supervised, unsupervised, and sequence-based learning, including neural networks, generative models, and transformers, and study how representation, data, and training choices shape model behaviour. The course combines theoretical understanding with hands-on implementation in leading ML frameworks, and applies methods across domains such as natural language processing, computer vision, and health data.

Intended Learning Outcomes:

- Describe contemporary machine learning approaches, including deep learning and probabilistic models.
- Explain key mechanisms in major ML methods, including representation, optimisation, and generalisation.
- Analyse the behaviour and performance of ML models across different tasks and data settings.
- Design and implement ML models in standard frameworks to address complex data science problems.

Dataset Creation (7.5 ECTS)

This course introduces students to the foundations of creating and handling heterogeneous data sources through the steps of privacy-conscious and ethical data collection, annotation, examining ethical issues and sources of bias, making data AI-ready through cleaning and preprocessing, and visualization.

In parallel, the student will also gain experience with discussing, presenting, and visualizing key aspects that can document and inform the data processing steps and effectively describe the continuous production of datasets.

Intended Learning Outcomes:

- Describe key data collection/annotation/visualization methods with regards to their strengths and weaknesses
- Apply appropriate data collection/annotation/visualization methods in order to create novel datasets
- Find suitable connections between dataset properties, analysis methods, and research questions
- Extract insights from the data analysis and present the results with appropriate visualization and reporting

- Create an AI-ready dataset and documentation adhering to appropriate principles (e.g., FAIR)

Evaluating Systems (7.5 ECTS)

This course examines how predictive systems are evaluated and how scientific claims about their performance are established. The course introduces quantitative and qualitative approaches to evaluation, including benchmarking, statistical analysis, and empirical studies of system and user behaviour. Students study how choices in experimental design shape evaluation results, and learn to assess the robustness and reproducibility of reported findings.

Intended Learning Outcomes:

- Describe canonical methods for evaluating predictive systems.
- Explain how evaluation decisions influence results.
- Analyse empirical evidence about system and user performance using quantitative and qualitative analysis methods.
- Plan and carry out evaluations that provide reliable, comparable, and reproducible results.

Human-AI Interaction (7.5 ECTS)

How can interactions between people and AI systems be designed to support effective and meaningful use? This course introduces key concepts and methods in human-AI interaction, including how AI systems present predictions, recommendations, and explanations to users. Students analyse how system behaviour and interface choices shape interaction, and design human-AI interaction approaches while reflecting on their implications for responsible AI use.

Intended Learning Outcomes:

- Describe key concepts and methods in human-AI interaction.
- Explain factors shaping human-AI interaction, including system behaviour, user mental models, and interface design.
- Analyse human-AI interactions using empirical evidence of user behaviour, interpretation, and decision-making.
- Design human-AI interaction approaches and reflect on their assumptions and implications for responsible AI use.

AI in Society (7.5 ECTS)

While algorithmic decision-making is now pervasive, it raises questions of accountability and regulatory compliance, as well as environmental costs of AI models and their overall

impact on society. This course introduces the key principles of AI accountability, transparency and regulatory compliance. It also considers the practical steps that could be taken to reduce environmental impact of systems in production.

Intended Learning Outcomes:

- Describe and discuss core societal challenges due to increased reliance on algorithmic decision-making, non-consensual data collection and aggregation, and generative models
- Identify and apply core legal principles that govern algorithmic decision-making in EU, including provisions on non-discrimination and explanations
- Identify and apply core principles of transparency and accountability, including in algorithmic decision-making and disclosure of generated media
- Quantify and analyse carbon emissions from running machine learning models, and consider ways to reduce them

Electives (15 ECTS)

Students can choose between a programme specific elective, or relevant 7.5 ECTS courses from ITUs existing Master's programmes. Students may furthermore choose to use the elective slot for a 15 ECTS research project within an external organization (supervised by ITU staff and graded).

AI in Organisations (Elective, 7.5 ECTS)

To be effective, a machine learning system needs to address the real needs of the organisation and to be supported by its workflows to correct for any limitations. It also needs to be constantly monitored to keep up with any changes in its context, and it should not expose the organisation to unacceptable security risks and supply chain issues. This course introduces key concepts of MLOps, and security challenges of modern generative AI systems. It also provides an introduction to organisational processes and principles of communication with non-expert stakeholders.

Intended Learning Outcomes:

- Analyse and evaluate how AI systems can support organisational processes, decision-making, and innovation, including their risks and limitations
- Design and implement AI processes to address specific needs of an organisation and engage with contemporary societal and/or organisational challenges
- Design and apply measures necessary for ongoing maintenance of a production model
- Identify possible security vulnerabilities due to the use of generative AI systems

Relevant 7.5 ECTS electives from other ITU MSc programmes include:

- Algorithmic Fairness and Ethics (Data Science)
- Geospatial Data Science (Data Science)
- Computer Systems Performance (Computer Science)
- Security and Privacy (Computer Science)
- Values and Ethics in Design (Digital Design & Interactive Technologies)
- Datafication and Representation (Digital Design & Interactive Technologies)
- AI Processes (Digital Innovation & Management)
- The Digital State (Digital Innovation & Management)
- IT & Green Transition (Digital Innovation & Management)
- Foundations of Game AI (Games)
- User Experience Research (Games)

Final project (15 ECTS)

The programme concludes with a 15 ECTS final project. Student groups are encouraged to select a recent research paper, attempt reproducing its results, and extend the work through additional analysis or experimentation.

A key challenge in machine learning research is low *reproducibility*; that is, claims from research do not hold when investigated by an independent research team. There are many reasons for why scientific studies do not reproduce, such as career incentives, lack of open data practices, and contextual factors. By engaging in a reproducibility format as the final project, the students in *Responsible AI* will learn a highly valuable craft of understanding factors that help make science cumulative, by hands-on investigating how well a recent AI papers claims holds. This way, the students can furthermore work directly with very recent international top-tier research within the time constraints a 15 ECTS project has.