Value and Impact through Synergy, Interaction, and coOperation of Networks of AI Excellence Centres

GRANT AGREEMENT NUMBER: 952070

Deliverable D5.2

Interim Design of European Artificial Intelligence Curriculum

Project title	**Value and Impact through Synergy, Interaction and coOperation of Networks of AI Excellence Centres**
Grant Agreement number	**952070**
Funding scheme	**H2020-ICT-48: Towards a vibrant European network of AI excellence centres, CSA**
Start date of the project and duration	**09/2020 - 08/2024 (48 months)**
Project coordinator name	**University of Leiden**
Deliverable number	**D5.2**
Title of the deliverable	**Interim Design of European Artificial Intelligence Curriculum**
WP contributing to the deliverable	**WP5**
Deliverable type	**Report**
Dissemination level	**Public**
Due submission date	**31 August 2023**
Actual submission date	**23 October 2023**
Partner(s)/Author(s)	**Barry O'Sullivan (UCC), Long Pham (UCC)**
Internal reviewers	**Holger Hoos, Marc Schoenauer, Vit Dockal, Janina Hoppstädter**
Final approval (Executive Board)	**Holger Hoos, Marc Schoenauer, Vit Dockal, Janina Hoppstädter**

Disclaimer

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 952070. This document has been prepared for the European Commission, however, it reflects the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein

History of changes		
When	**Who**	**Comments**
14 April 2023	Long Pham	First Draft
10 August 2023	Long Pham	Second Draft
22 September 2023	Long Pham	Third Draft
15 October 2023	Long Pham	Fourth Draft
20 October 2023	Long Pham	Final

Confidentiality	
Does this report contain **confidential** information?	Yes ☐ No x
Is the report **restricted** to a specific group?	Yes ☐ No x *If yes, please precise the list of authorised recipients*:

Table of Contents

Introduction

The Interim Design of European Artificial Intelligence Curriculum is a deliverable of the Work Package 5 “Development of AI Skills and Education in Europe”. It belongs to the Task 5.2 “Design of a European AI Curriculum & Training Initiatives.”

The Original Plan

The original plan for the task was because AI is deployed more extensively and more intensely, and there will be a compelling need to educate a broader segment of society, starting from an early age, in a broad set of skills including mathematics, coding, and understanding of data and data analytics, as well as in AI.

Complementary cognitive, socio-cultural, and interdisciplinary skills will be required to empower individuals and enable the mastery of sectoral knowledge in combination with new technological solutions.

Thus, education systems must prepare people not only with digital skills but also by strengthening people’s human-centric skills. Educators require specific training and support since it is they who will train the next generations. This task will take an end-to-end view of education for AI in Europe and consider the development of a curriculum from elementary school, through to high school, university, and beyond.

The task was set to identify, integrate, and create new cross-disciplinary AI training initiatives and curricula, using the NoEs as sources of expertise. Activities were planned to support the development of standardized AI curricula suited for integration into either computing/mathematical or non-technical undergraduate and postgraduate programmes and programmes targeted at providing AI skills for non-professionals. We will also catalogue Europe’s current offerings in terms of online courses, taught programmes, and centres for research training in AI. We will run a European AI education and training summit each year, bringing together the very best of Europe’s education initiatives to deliver a training programme with tracks for different levels of participants.

The Changed Plan

In Amendment 2, following the suggestion of the EC, VISION’s Task 5.2 focuses entirely on PhD-level education and ensures that the focus is on delivering the European AI Doctoral Academy (AIDA). Deliverable 5.2 and milestone 7 were moved to M36 to align with the timescale of AIDA and the timescales of the four NoEs.

The change of focus, to ensure that Europe is a world-leading powerhouse in AI, requires a massive input of human intelligence. Europe needs more talents to lead ambitious fundamental world-class research programmes and retain them with attractive research and working conditions. Furthermore, AI systems will augment and complement human capabilities in hybrid work settings. Practitioners in all areas will need to be upskilled with digital competencies to be enhanced in their roles. At the same time, the future workforce will have to be equipped with a new human-centric set of skills that empowers them on a cognitive and a socio-cultural level to face the challenges ahead.

This requires a deep-reaching change in education systems in Europe. Fundamental changes are required in European educational systems and support activities to share experiences and best practices and to benefit from common opportunities. A multi-faceted and multidisciplinary approach is required that mobilises various education and training stakeholders including education institutions

at all levels, social partners, training providers, community colleges, and companies, which will also have to shoulder parts of the AI upskilling and reskilling requirements.

As an activity to support the development of AIDA, an initiative across the set of ICT-48 networks, involving dozens of European universities and research organisations, UCC led the effort to create a standardized AI PhD curriculum suited for integration into either computing/mathematical or non-technical undergraduate and postgraduate programmes. This curriculum, once finalised and agreed upon among key contributors, would serve as a comprehensive framework for advancing AI knowledge and skills among students.

To promote this curriculum, the Task members have been organising events and activities aimed at raising awareness and generating interest in AI education. These events are bringing together experts, educators, and stakeholders from academia and industry to foster collaboration and knowledge exchange. Furthermore, we are actively supporting the development of robust quality assessment mechanisms for AI programs, ensuring that the education provided aligns with the highest standards and meets the evolving needs of the field.

Our commitment extends beyond the initial stages of AIDA's development. We are dedicated to ensuring the long-term sustainability of the academy by providing comprehensive support in all aspects of its growth. This includes offering guidance and resources for faculty members, facilitating collaborations between participating institutions, and leveraging our expertise to address any challenges that may arise. By doing so, we aim to establish AIDA as a prominent institution in the AI education landscape, fostering the growth of skilled AI professionals and advancing the field.

As a key activity of the AIDA's implementation and operation, the below progress report describes what has been done in the Interim Design of the European Artificial Intelligence Curriculum.

European AI PhD Curriculum

The Making of The Curriculum

Under the AIDA's operational structure, an AI Curriculum Committee was formed in August 2022. Since AIDA's ambitious initiative aims at promoting and supporting the development of AI education and research in Europe, the committee included leading European AI educators and researchers. The AI Curriculum Committee members came from ICT-48 networks, which encompass multiple European universities and research organisations. The members, apart from their association with the ICT-48 networks, are prominent AI educators in leading European institutions. Figure 1 showed members of the AI Curriculum Committee, which included representatives from all ICT-48 initiatives: AI4Media, ELISE, HumanE-AI Net, and TAILOR, with management, coordination, and support provided by AIDA and VISION.

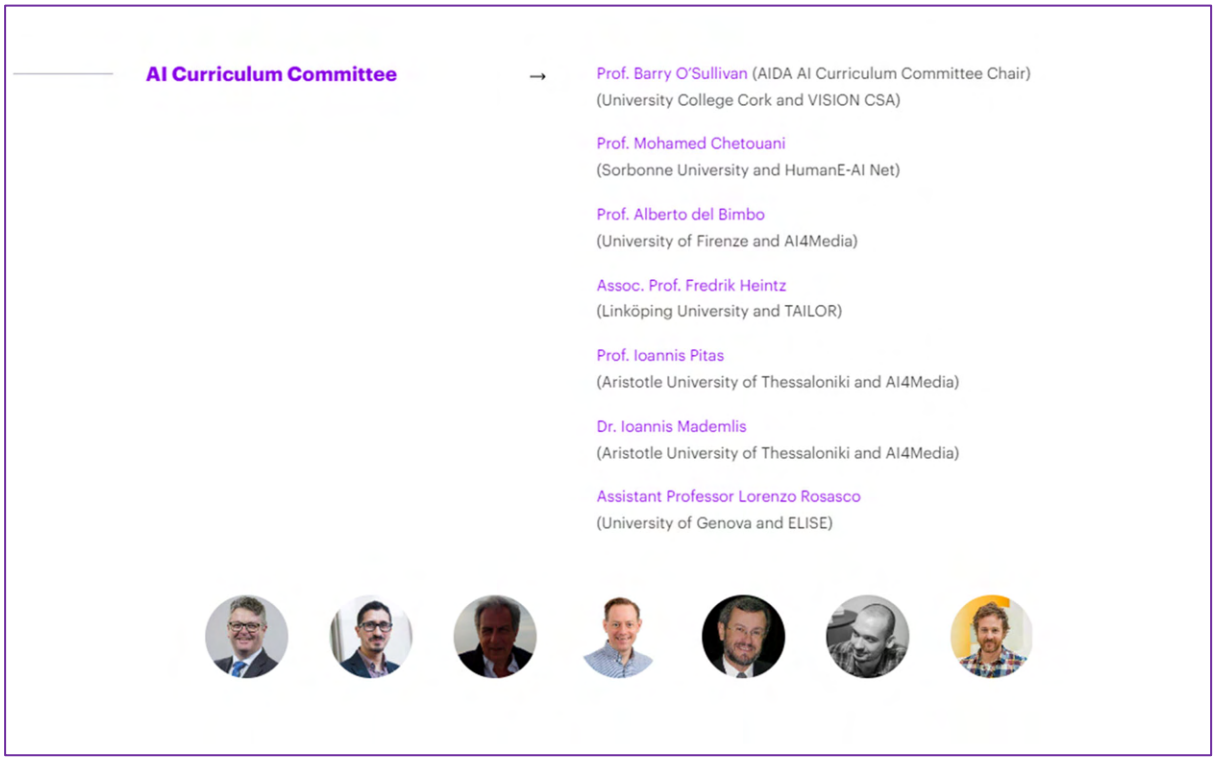

Figure 1: Members of the AI Curriculum Committee were presented on the AIDA Website and are accessible to the public.

Sharing a goal of establishing a standardized AI PhD curriculum, the AI Curriculum Committee discussed proposals for the Curriculum structure, frameworks, formats, and contribution mechanisms from the members of the AI Curriculum Committee and their colleagues.

AIDA AI Curriculum Committee Chair Prof. O'Sullivan, from University College Cork (UCC) and representing VISION CSA, proposed a framework for a curriculum and presented the framework with the committee members on a shared document that is accessible by designated members.

As an original proposal, the framework comprised seven "AI Core Subjects" and four "AI Special Topics." The AI Core Subjects are positioned as applicable to all AI PhD programmes while the AI Special Topics belong to those PhD programmes that fall under the covered areas of the four networks: AI4Media, ELISE, HumanE-AI Net, and TAILOR.

A common framework for subjects was used. Each of the subjects comprises three key components of Level; Learning Outcomes, which include content/knowledge, methodological/skills, transferable/application; and Courses available through AIDA or other online courses linked to the networks that cover elements of the topics' material.

At this interim stage, the AI Core Subjects are:

1. Foundations of Artificial Intelligence
2. AI Paradigms and Representations
3. Deciding and Learning How to Act
4. Machine Learning
5. Computer Vision
6. Natural Language Processing and Analysis

7. AI Ethics and Governance

The four Special Topics are:

1. AI for Media, Society and Democracy, to include:
   a) Music/Sound Analysis and Synthesis
   b) AI and Game Media
   c) Web and Social Media Analysis and Mining
   d) Human-Centred Media Analysis

2. Integrating Approaches for Trustworthy AI to include:
a) Foundations of Trustworthy AI
b) Reasoning and Learning in Social Contexts
c) Automated AI

3. Data-Driven Learning to include:
a) Deep Learning
b) Machine Learning Theory
c) Reinforcement Learning and Sequential Decision-making
d) Distributed and Federated Learning
e) Generative Artificial Intelligence

4 Human-Centric AI, delivered by HumaneAINet, to include:
a) Human-Centred Machine Learning
b) Explainable AI

The interim contributing process took more than eight months starting in Oct 2022, when members have their own access to the shared document that tracked all contributions, comments, feedback, and notes to the Chair of the Committee.

The Committee members meet up in AIDA's monthly meeting organised and coordinated by Prof. Ioannis Pitas of Aristotle University of Thessaloniki and a leader of the AI4Media. The AI Curriculum is a regular item in the monthly meetings for updates, discussions, justifications, and proposals for changes. Also, hundreds of emails were exchanged among the Committee members and their related colleagues for rigorous contribution and feedback.

The Curriculum, shown in its latest form in the Appendix, is truly a living document for the Committee members. Emerging topics, such as Generative AI, were gradually included and updated to reflect the fast-paced development of the global AI landscape.

To date, 30+ contributions have been recorded from the Committee members. The members who contributed also expanded from just one or two from each of the ICT-48 networks to a dozen. While the Curriculum has attracted all networks' participation, the numbers and quality of the contribution varied from each network. Some networks are very active and serious with their contributions, ensuring that their contributions are well-justified and well-resourced, while other networks struggle to get their members' interest or feedback with their own networks' contributions.

Lessons from the Creation of European AI PhD Curriculum

The European AI PhD Curriculum was elevated through a collaborative effort involving four networks of excellence (NoEs). These NoEs brought together leading scientists and educators who evaluated and contributed their respective curricula to shape the development of the Curriculum. The aim was to create a comprehensive framework comprising the core modules and the special topics, all of which were considered essential by the NoEs.

Transparency played a vital role throughout the Curriculum development process. The collaborative nature of the project was facilitated through shared documents that allowed all contributing members to view and access the Curriculum materials. This transparency ensured that all additions, edits, and comments made by the participating scientists and educators were visible to everyone involved. Such openness fostered a sense of inclusivity, allowing for diverse perspectives and expertise to be considered and integrated into the final curriculum.

By leveraging the expertise and knowledge of these four networks of excellence, the European AI Ph.D. curriculum was able to benefit from a wide range of perspectives and experiences. This collaborative and transparent approach ensured that the Curriculum was comprehensive, rigorous, and reflective of the latest advancements and best practices in the field of AI.

The European AI PhD Curriculum received some pioneering contributions from the networks, representing an encouraging level of collaboration in the field. The leading scientists and educators showed their professionalism and expertise with their strong scrutinized and honed their respective curricula to create the first standard AI PhD curriculum for Europe, paving the way for a comprehensive one that would cover every facet of AI research and development.

By involving multiple networks of excellence and encouraging the active participation of experts, the European AI PhD Curriculum benefits from diverse perspectives and ensures that the educational needs of the AI community are met. It provides a strong foundation for doctoral programs in AI across Europe and helps shape the future generation of AI researchers and practitioners.

As described in the making, the Curriculum was discussed and reported in every AIDA management board meeting with progress, directions for development and contributions from academic leaders of the four NoEs; the Curriculum gathered the best minds in the networks to discuss regularly and bring out the most up to date developments in AI. Additional modules and topics were debated and updated regularly according to the very fast-paced AI development in the markets.

The Curriculum, designed toward benchmarking AI Ph.D. Education with a human-centric and ethical AI approach, played a significant role in the discussions held during the AIDA (Artificial Intelligence Doctoral Academy) management board meetings. These meetings served as a platform for academic leaders from the four NoEs to come together and share progress, discuss directions for development, and contribute their expertise to the Curriculum.

The Curriculum contributed to bring together the best minds within the networks to engage in regular discussions and ensure that it captured the most up-to-date developments in the field of AI. Given the rapidly evolving nature of AI technologies and their applications in various markets, it was crucial to regularly update and adapt the Curriculum to stay relevant.

During the management board meetings, the participants debated additional modules and topics to be included in the Curriculum. This allowed for comprehensive coverage of the diverse aspects of AI, ensuring that the Curriculum provided a holistic and thorough understanding of the subject matter.

The involvement of academic leaders from the four NoEs ensured that the discussions were well-informed, and that the Curriculum reflected the latest advancements in AI research, methodologies, and applications. Leveraging the collective expertise of these academic leaders, the Curriculum aimed to equip future AI PhD students with the necessary knowledge and skills to excel in their research and contribute to the field.

The regular discussions and updates to the Curriculum helped it keep pace with the fast-paced development of AI in the markets. By incorporating the most recent trends and breakthroughs, the Curriculum can prepare students for the current and future challenges and opportunities in the field.

Largely, the European AI PhD Curriculum served as a dynamic and collaborative platform that brought together academic leaders to discuss, shape, and update the Curriculum in response to the rapidly evolving landscape of AI. As it might be adopted or used as a reference curriculum for many European universities and educational institutions, the Curriculum can provide a comprehensive education to AI PhD students and ensure that they are equipped with the latest knowledge and skills required to contribute to the field of AI.

Promotion and Dissemination of the European AI PhD Curriculum

Website Presence

European AI PhD Curriculum on AIDA Website

The European AI PhD Curriculum has gathered significant attention from NoEs and others in the European AI ecosystem and its presence and accessibility would be a key to AIDA, the NoEs and European AI educational institutions. At present, there's an "AI PhD Curriculum" webpage on the AIDA website. The webpage will be replaced with new content from the European AI PhD Curriculum after its first launch in December 2023.

The mention of the European AI PhD Curriculum in the making has been promoted through various channels, including the AIDA website, the AI-on-Demand Platform, and other relevant platforms. Its presence was consistently highlighted to ensure awareness before its visibility.

Mock-up designs of the Curriculum showed simpler blocks, only two rather than four at the moment, and more user-friendly for better engagement and navigation to the contents of the presented topics.

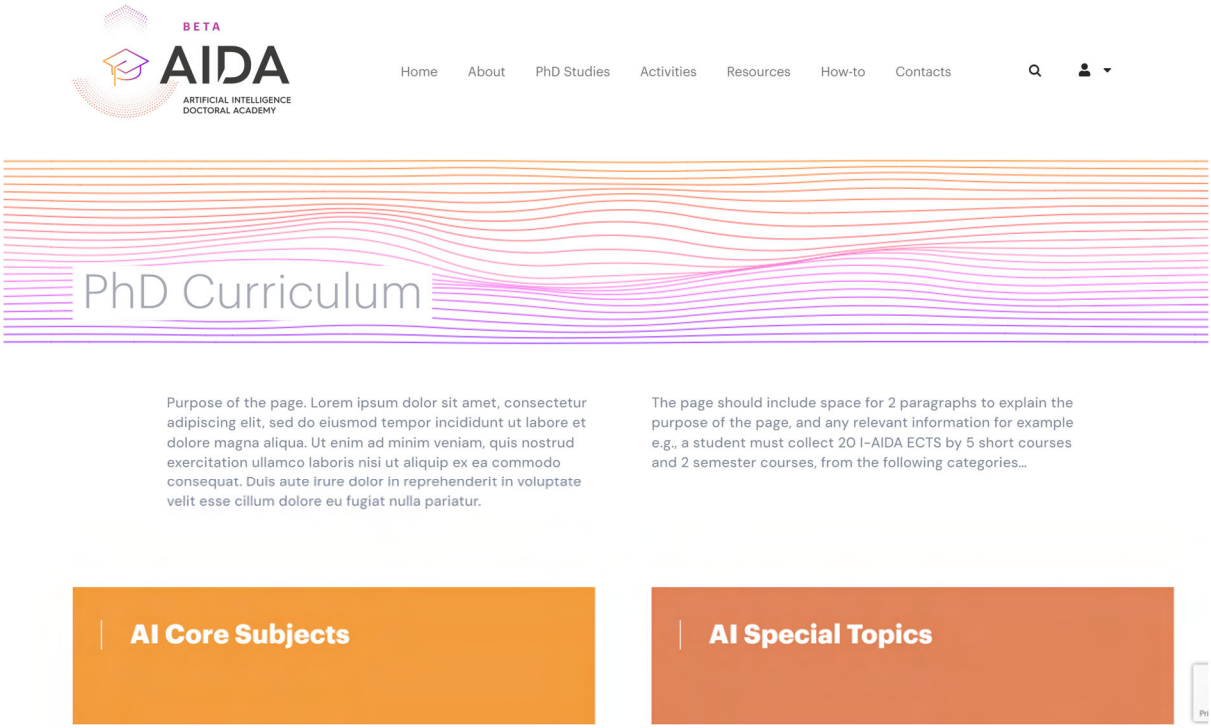

Figure 2: Mock-up webpage design of the Curriculum, where each of the two columns shows the list of topics included.

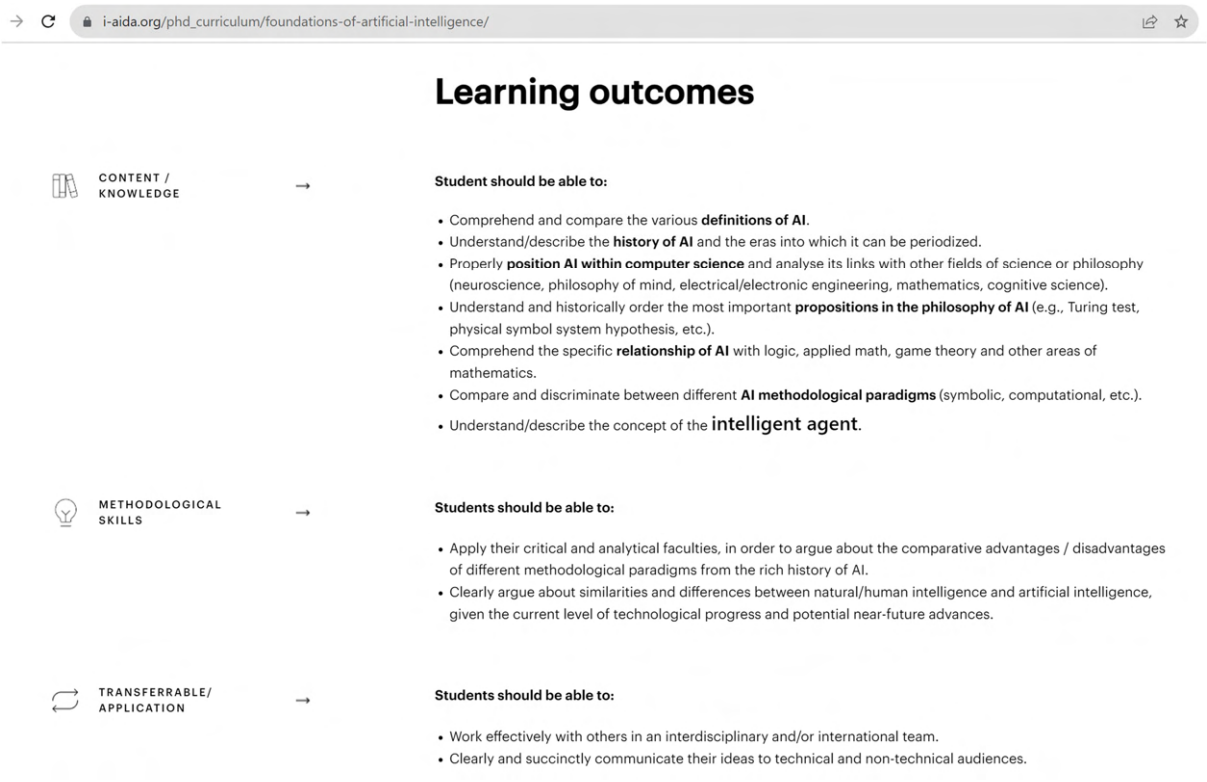

Figure 3: The mock-up presentation on the webpage that shows the content, including the learning outcomes of the Foundations of AI under the AI Core Subjects

European AI PhD Curriculum on Other Websites

The Curriculum is planned to have an introductory presence on NoEs' websites and the AI-on-Demand Platform once it is ready for the public.

The introductory presences will have links that direct the visitors to the Curriculum webpage on the AIDA website. Other websites' associations may arise from those universities or educational institutions that choose the Curriculum as their guidance in their training programmes. The Curriculum's guest presence on the websites of those institutions can be a term or condition for AIDA's sponsorship or support of their programmes.

Other Dissemination Activities

Online Events

In March 2023, the University College Cork (UCC) took the initiative to organise an online event aimed to introduce the PhD initiatives implemented across the Network of Excellence (NoEs). The event also aimed to showcasing the high level of the Curriculum to the public. The news regarding this online event was promoted widely, reaching a wide audience through various EU AI networks and members of the AI-on-Demand Platform. The extensive dissemination of information played a crucial role in generating interest and attracting participants.

The online event proved to be a resounding success, as it received an impressive number of registrations. Nearly 180 individuals expressed their interest in attending and engaging with the event. On the day of the event, almost 90 participants actively participated in the three-hour-long session.

The enthusiastic response demonstrated the immense enthusiasm and eagerness within the AI community to explore and contribute to the European AI PhD Curriculum.

Figure 4: The visual of the AIDA AI Education Online Event in March 2023 which included the introduction of the PhD Curriculum in the making to the audience of nearly a hundred.

Throughout the event, attendees were presented with valuable insights into the various PhD initiatives being pursued across the NoEs. Expert speakers shared their knowledge, research findings, and experiences, offering a comprehensive overview of the Curriculum's high standards. Participants were able to engage in interactive sessions, allowing for fruitful discussions, idea exchanges, and networking opportunities.

The event served as a platform for fostering collaboration, as participants from diverse backgrounds and institutions came together to collectively contribute to the advancement of AI research and development. The connections established during the event paved the way for future collaborations and partnerships, encouraging further progress in the European AI landscape.

Following the event, its impact continued to reverberate throughout the AI community. Attendees shared their experiences and key takeaways on various platforms, sparking conversations and encouraging others to explore the European AI PhD Curriculum. The event's success solidified its position as a pivotal milestone in promoting and establishing the Curriculum as a leading educational framework for AI research within Europe.

A Launching Event for the European AI PhD Curriculum

Moving forward, the European AI PhD Curriculum is expected to witness even greater momentum, as more individuals recognise its value and actively engage with its offerings.

A launching event is planned in the last quarter of 2023 in an online event to maximize the participation of the interested parties. The objectives of the launching event include:

- Bringing the initiative to a larger audience than the NoEs
- Introducing opportunities for institutions to join AIDA
- Making their courses associated with and follow AIDA's standards
- Seeking courses and programmes that may suit AIDA's objectives
- Introducing a benchmark for current and prospective AI students at PhD levels
- Attracting decision-makers to consider using the AIDA curriculum as a benchmark and others.

Meanwhile, ongoing efforts to disseminate information through EU AI networks, the AI-on-Demand Platform, and other channels will continue to play a vital role in raising awareness and attracting talented individuals who are eager to contribute to the cutting-edge field of AI research.

The European AI PhD Curriculum was promoted on AIDA website, AI-on-Demand Platform, and other associated social media channels regularly.

It seeks to establish partnerships with institutions offering AI-related courses and programs, encouraging them to align their curriculum with AIDA's standards. By doing so, AIDA aims to create a cohesive and consistent AI education experience across different institutions, ensuring high-quality standards are maintained and providing students with a more uniform and comprehensive learning experience.

The Interim Design of the EU AI PhD Curriculum has made satisfactory outcomes to date judged by the members of the Curriculum Committee and their NoEs. There is room for improvement and more work to be done until the end of VISION and beyond. The upcoming work includes syncing PhD-level educational resources from relevant European AI platforms and initiatives, expanding and linking AI taxonomy to the Curriculum, and promoting the Curriculum to a diverse audience.

By expanding its reach, aligning courses with AIDA's standards, introducing a benchmark, and attracting decision-makers, the Curriculum aims to provide a comprehensive and standardized educational framework for AI researchers at the doctoral level, fostering excellence and innovation in the field.

Enriching the European AI PhD Curriculum is an ongoing and dynamic process. The Curriculum is designed to adapt to the rapidly evolving field of artificial intelligence, ensuring that students receive comprehensive and up-to-date education. As a living document, the Curriculum is regularly updated to incorporate new courses that align with its objectives and reflect emerging trends in AI research and applications.

The taxonomy associated with the Curriculum is also continuously refined to ensure it accurately captures the diverse and interdisciplinary nature of AI. By regularly updating the taxonomy, the Curriculum remains inclusive of the latest developments and areas of focus within the field. This iterative process allows for the inclusion of emerging topics that are relevant to the Curriculum, ensuring that students are equipped with the most relevant knowledge and skills.

Furthermore, the Curriculum emphasizes the importance of ongoing discussions and expansions regarding emerging topics. AI is a rapidly evolving field, and it is crucial to provide a platform for continuous exploration and dialogue. By fostering a culture of open discussion, the Curriculum can effectively address emerging challenges, trends, and ethical considerations related to AI. This ensures

that students not only have a strong foundation in core AI concepts but also develop critical thinking skills and an awareness of the broader societal implications of AI.

So far, the creation of the European AI PhD Curriculum acknowledges and reflects the dynamic nature of the field and actively encourages the inclusion of new courses and the expansion of the Curriculum's taxonomy to ensure that students receive a comprehensive and relevant education in artificial intelligence. By promoting ongoing discussions and updates, the Curriculum aims to equip future AI researchers and professionals with the knowledge and skills necessary to address the complex challenges and opportunities in the field.

Appendix

AIDA Curriculum

Barry O'Sullivan (VISION, Chair)

Long Pham (VISION)

Ioannis Pitas (AI4Media)

Lorenzo Rosasco (ELISE)

Mohamed Chetouani (HumaneAINet)

Peter Flach and Fredrik (TAILOR)

AI Core Modules

1. Foundations of Artificial Intelligence
2. AI Paradigms and Representations
3. Deciding and Learning How to Act
4. Machine Learning
5. Computer Vision
6. Natural Language Processing and Analysis
7. AI Ethics and Governance

Specialised Topic Pillars

1. AI for Media, Society and Democracy, to include:
    a. Music/Sound Analysis and Synthesis
    b. AI and Game Media
    c. Web and Social Media Analysis/gMining
    d. Human-Centred Media Analysis

2. Integrating Approaches for Trustworthy AI to include:
    a. Foundations of Trustworthy AI
    b. Reasoning and Learning in Social Contexts
    c. Automated AI

3. Data-Driven Learning to include:

   a. Deep Learning
   b. Machine Learning Theory
   c. Reinforcement Learning and Sequential Decision-making
   d. Distributed and Federated Learning
   e. Generative Artificial Intelligence

4. Human-Centric AI, delivered by HumaneAINet, to include:
   a. Human-Centred Machine Learning
   b. Explainable AI

AIDA Core Topic 1	
Foundations of Artificial Intelligence	
Level	This topic presents the foundations, scope, history and methodologies of AI. *Level*: **Foundational, Broad, Theoretical.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Student should be able to: ● Comprehend and compare the various **definitions of AI**. ● Understand/describe the **history of AI** and the eras into which it can be periodized. ● Properly **position AI within computer science** and analyse its links with other fields of science or philosophy (neuroscience, philosophy of mind, electrical/electronic engineering, mathematics, cognitive science). ● Understand and historically order the most important **propositions in the philosophy of AI** (e.g., Turing test, physical symbol system hypothesis, etc.). ● Comprehend the specific **relationship of AI** with logic, applied math, game theory and other areas of mathematics. ● Compare and discriminate between different **AI methodological paradigms** (symbolic, computational, etc.). ● Understand/describe the concept of the **intelligent agent**.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Apply their critical and analytical faculties, in order to argue about the comparative advantages / disadvantages of different methodological paradigms from the rich history of AI. ● Clearly argue about similarities and differences between natural/human intelligence and artificial intelligence, given the current level of technological progress and potential near-future advances.

AIDA Core Topic 1	
Methodological/Skills It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team. ● Clearly and succinctly communicate their ideas to technical and non-technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

AIDA Core Topic 2	
AI Paradigms and Representations	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic covers the challenge of integrating different representations and paradigms for AI in order to enable both learning, reasoning and optimisation. The integrated representations are intended to engender trustworthiness. *Level*: **Intermediate, Broad, Algorithmic, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	● Students should be able to understand the motivations for the need to integrate learning, reasoning and optimisation, and the role of prior knowledge and knowledge representation. ● Students should be able to understand integrated representations for trustworthy AI. ● Students should be able to understand different paradigms that integrate different representations. In particular: ○ Statistical relational AI: the integration of logic and probability/fuzziness for both

AIDA Core Topic 2	
	reasoning and learning. ○ Neurosymbolic AI: integrating logic with neural networks to enable perception and reasoning. ○ Knowledge graphs, ontologies, graph neural networks and embeddings. ○ Constraint satisfaction and optimisation techniques: integrating solvers and learners for better performance and for learning CSP models. ● Students should be able to apply the above methods in perception, spatial reasoning, natural language processing, vision, and other societal/industrial domains
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	● Students should be able to use a wide variety of representations (graphical models, logic, neural networks, knowledge graphs) for both learning and reasoning. ● Students should have insight into the power and limitations of different types of representations. ● Students should be able to combine different representations for a particular AI task. ● Students should be able to use both knowledge and data for a particular AI problem. ● Students should be able to understand and use the above mentioned categories of techniques (StarAI, NeSy, CSP, Knowledge graphs, Ontologies (OWL/RDFS). ● Students should understand the limitations and challenges of the integrated representations and paradigms. ● Students should understand the trustworthiness of these techniques.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences*".	● Students should be able to work effectively with experts in different learning, reasoning and optimisation paradigms. ● Students should be able to collaborate with domain experts to identify suitable integrated learning, reasoning and optimisation techniques for trustworthy AI.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	

AIDA Core Topic 2	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

AIDA Core Topic 3	
Deciding and Learning How to Act	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic covers ways in which AI agents can be empowered with the ability of deliberating autonomously how to act in the world. *Level*: **Intermediate, Broad, Theory, Algorithmic.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	● Understand the different approaches in the fields of Artificial Intelligence and Formal Methods that can be applied in synergy to develop autonomous agents. ● Recognise the mathematical and algorithmic techniques as well as the key challenges to solve sequential decision making problems. ● Integrate data-driven learning methods with model-based reasoning methods for deciding and learning how to act. ● Identify the limitations of current machine learning and reasoning methods to act in the real-world. ● Extended and multi-facet models of the world dynamics and tasks. ● Integrating data-based methods with model-based methods in deciding and learning how to act. ● Learning for reasoners and planners, and reasoners and planners for learning. ● Monitoring and controlling to make actions AI trustworthy in the real world. ● Modelling, reasoning and decision-making under uncertainty, ● Probabilistic/Fuzzy Reasoning ● Multiagent Decision Making ● Probabilistic/Fuzzy Programming

AIDA Core Topic 3	
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	• Program advanced agents using learning and planning techniques for solving sequential decision making tasks that involve other agents. • Analyse autonomy in dynamic, partially observable settings involving a single agent or multiple agents. • Develop methods for optimising control policies in complex sequential decision making problems. • Implement techniques to balance exploration and exploitation in decision making tasks that require learning from the environment while acting on it. • Use linear time logic as a specification language for formulating complex tasks as well as environment properties. • Apply synthesis from LTL and LTLf specifications to solve planning problems in nondeterministic environments.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences*".	• Work effectively with others in an interdisciplinary and/or international team to reach a collective objective by sharing knowledge, learning and building consensus. • Present materials coherently and concisely in written or oral form, with clear use of language to a technical audience. • Students should be able to work effectively with others in an interdisciplinary and/or international team. • Students should be able to design and manage individual projects. • Students should be able to clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Reinforcement Learning, Universitat Pompeu Fabra (Msc. in Intelligent and Interactive Systems)	https://www.upf.edu/web/iis/ReinforcementLearning
Autonomous Systems, Universitat Pompeu Fabra (Msc. in Intelligent and Interactive Systems)	https://www.upf.edu/web/iis/AutonomousSystems
Game-Theoretic Approach to Planning and Synthesis (already offered through AIDA)	https://whitemech.github.io/courses

AIDA Core Topic 3	
Non-Markov Decision Processes and Reinforcement Learning	https://whitemech.github.io/courses

AIDA Core Topic 4	
Machine Learning	
Level	This topic presents the foundations and basic methodologies of machine learning. *Level*: **Foundation, Broad, Theory, Algorithmic, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Student should be able to understand and describe: ● The concept of **learning from data** in contrast to rule-based systems. ● The differences between the various **machine learning paradigms** (supervised learning, unsupervised learning, reinforcement learning, active learning). ● The **statistical foundations** of machine learning (predictor, loss function, statistical risk, Bayes optimal prediction, Bayes risk, bias-variance decomposition, overfitting and underfitting, consistency, regularization, stability, model validation and model selection). ● The basic **linear models** for regression and classification (linear regression and linear classification, logistic regression, Support Vector Machines). ● The theory of **kernel functions** and the implementation of linear models in kernel spaces. ● The notion of **non-parametric learning**, its connection to consistency, the **curse of dimensionality**, and the basic non-parametric algorithms (e.g., k-NN, tree predictors, universal kernels). ● **Probabilistic models** (e.g., Naive Bayes classifiers,

AIDA Core Topic 4	
	graphical models, Gaussian mixture models). ● **Ensemble learning** methods (bagging, boosting, stacking) and their effect on the bias-variance dilemma. ● **Feed-forward neural networks**, expressivity versus network size, hardness of training, backpropagation algorithm. ● Basics of **on-line learning** and **reinforcement learning** (e.g., online gradient descent, multi-armed bandits, Markov Decision Processes with finite and discounted horizon, model-based and model-free RL algorithms). ● Basic **clustering algorithms** (e.g., k-means/k-means++, DBSCAN, hierarchical clustering, spectral clustering, correlation clustering). ● **Model validation and selection** (e.g., cross-validation, nested cross-validation). ● Basic **dimensionality reduction** methods (e.g., LDA, PCA, ISOMAP, LLE).
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Understand the correct methodological approach for implementing learning algorithms and running experiments, including model validation and selection, and reproducibility. ● Understand the basic techniques for hyperparameter tuning and their correct implementation. ● Know the basic evaluation metrics and the basic techniques for preprocessing the datasets, including feature selection, normalization, and transformation (e.g., for handling categorical features).
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences*".	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team. ● Design and manage individual projects. ● Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	

AIDA Core Topic 4	
Machine Learning, QMUL (Semester Course)	https://www.i-aida.org/course/machine-learning/
Non-Markov Decision Processes and Reinforcement Learning, UNIROMA (Semester Course)	https://www.i-aida.org/course/non-markov-decision-processes-and-reinforcement-learning/
Statistical Inference Practice, UCA (Semester Course)	https://www.i-aida.org/course/statistical-inference-practice/

AIDA Core Topic 5	
Computer Vision	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic concerns computer vision for image/video analysis, with its courses ranging from foundational image processing up to specialised deep learning approaches. It covers both theoretical knowledge and its real-world applications (e.g., in autonomous systems perception, medical imaging, etc.). *Level*: **Foundation, Intermediate, Broad, Specialised, Theory, Algorithmic**.
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand/describe **common algorithms** for image segmentation, filtering, registration, search and retrieval, compression/storage, 2D shape description and recognition, as well as algorithms for video motion estimation, description, search, indexing, retrieval, streaming and compression. ● Understand/describe the operation of various **visual sensors**, the image/video **digitization** process and the most common existing image/video **formats**. ● Compare between various **handcrafted and convolutional feature extraction methods** for image/video description / representation. ● Understand/describe **common algorithms** for visual SLAM, face/person/pedestrian detection,

AIDA Core Topic 5	
	object detection/tracking/pose estimation, semantic/instance segmentation, facial expression recognition and activity recognition. ● Understand/describe the process of **image acquisition** by visual sensors and the mathematical modelling of **image formation**. ● Understand/describe **common algorithms** for camera calibration, stereoscopic 3D vision, depth estimation, surface geometry description and object 3D localization, topology description, landmark extraction and registration.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Analyse and develop (in C/C++, MATLAB or Python) the taught computer vision algorithms, by practically applying their gained knowledge in a systematic manner. ● Evaluate the accuracy of computer vision algorithm implementations on suitable datasets, by employing common and appropriate task-specific metrics.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team. ● Design and manage individual projects. ● Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Computer Vision and Machine Learning, AUTH (Web Lecture Series)	https://www.i-aida.org/course/computer-vision-and-machine-learning/
Artificial Intelligence for video streaming platforms, UPB (Short Course)	https://www.i-aida.org/course/artificial-intelligence-for-video-streaming-platforms-2/
Autonomous Systems Perception, AUTH (Semester Course)	https://www.i-aida.org/course/autonomous-systems-perception/

AIDA Core Topic 5	
Computer Vision, AUTH (Semester Course)	https://www.i-aida.org/course/computer-vision/

AIDA Core Topic 6	
Natural Language Processing and Analysis	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic concerns machine learning for Natural Language Processing (NLP). *Level*: **Intermediate, Niche, Algorithmic, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand/describe the **basics** of language modelling, tokenization, normalisation, stemming, lemmatization and Parts-of-Speech (POS) tagging. ● Compare between Bag-of-Words, N-Grams, TF-IDF and learned word embedding-based **representations of text**. ● Understand/describe **common algorithms** for question answering, text classification, text/document summarization, sentiment analysis, sentence similarity estimation, speech recognition and neural machine translation.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Analyse and develop (in C/C++, MATLAB or Python) the taught NLP algorithms, by practically applying their gained knowledge in a systematic manner. ● Evaluate the accuracy of NLP algorithm implementations on suitable datasets, by employing common and appropriate task-specific metrics.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team.

to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	• Design and manage individual projects. • Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
N/A	N/A

AIDA Core Topic 7	
AI Ethics and Governance	
Level - Choose one or more labels from the following: • Foundation, Intermediate, Advanced • Broad, Niche, Specialised • Theory, Algorithmic, Methodological	This topic concerns AI ethics and regulation from a legal and philosophical perspective. The demands, aspects and implications of Trustworthy AI are also covered. *Level*: **Intermediate, Specialised, Theory, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: • Understand/describe current discourse on A**I autonomy, personhood, safety, liability and security**, as well as related **regulations**, their enforcement and oversight. • Comprehend **Dual Use** issues, compare EU with non-EU perspectives, as well as clearly articulate the principles, aspects and implications of **Trustworthy AI**. • Understand/describe current discourse on specific **relevant case studies** (e.g., fake media, fake news, autonomous systems safety, etc.).
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model*	Students should be able to: • Apply their critical and analytical faculties on specific case studies, in order to argue about the need and content of potential AI regulations, as well as to detect possible dangers (to individuals and to the society) from unregulated AI.

selection."	• Clearly argue about how specific AI regulations can/should balance competing concerns (individual liberties, societal good, technical efficiency and economic competitiveness), given the current level of technological progress and potential near-future advances.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: • Work effectively with others in an interdisciplinary and/or international team. • Clearly and succinctly communicate their ideas to technical and non-technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Liability for Damage Caused by Artificial Intelligence, KUL (Short Course)	https://www.i-aida.org/course/liability-for-damage-caused-by-artificial-intelligence/

Pillar 1: AI for Media, Society & Democracy (AI4Media)

Description of the Pillar
AI for Media, Society and Democracy, delivered by AI4Media, includes the following topics: a. Music/Sound Analysis and Synthesis b. AI and Game Media c. Web and Social Media Analysis/Mining d. Human-Centred Media Analysis Modern AI tools/algorithms are of immense importance to the media sector, while they are simultaneously transforming our digitised society and influencing the quality of liberal democracy in multiple ways. This pillar contains educational offerings that are necessary for grasping the basics of such AI algorithms, both at the theoretical and at the applied level. An essential goal of this pillar is to disseminate knowledge in a structured manner about how computer vision, natural language processing and machine/deep learning can be employed for image/video/music/speech/text analysis and understanding. Methods of this type are extremely important to media production/post-production, as well as to applications with exceptional societal value, such as autonomous systems or social media/WWW analysis. Complex network analysis and Semantic Web offerings complement this aspect, while new media such as computer games are also covered thanks to offerings regarding AI algorithms for games. Finally, this pillar concerns the deployment of AI in daily life, exploring issues such as emotion analysis during interactions of AI systems with humans. Thus, the ultimate aim of the "AI for Media, Society & Democracy" pillar is for the student to absorb both the foundations

and more advanced details of relevant AI methods, as well as to grasp their applicability in specific domains relevant to the media industry and the society in general.

Topic 1(a)	
Music/Sound Analysis and Synthesis	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic concerns machine learning for audio/music/sound analysis, synthesis/generation and retrieval, with an emphasis on real-world applications. *Level*: **Foundation**, **Intermediate, Niche, Theory, Algorithmic, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand/describe the basics of **digital audio signal representation and processing**. ● Understand/describe **common algorithms** for music tagging, transcription, indexing, retrieval, synthesis, similarity estimation, tempo estimation and instrument recognition, as well as audio source separation, acoustic scene classification, sound event detection and acoustic anomaly detection.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Analyse and develop (in C/C++, MATLAB or Python) the taught semantic audio analysis algorithms, by practically applying their gained knowledge in a systematic manner. ● Evaluate the accuracy of semantic audio analysis algorithm implementations on suitable datasets, by employing common and appropriate task-specific metrics.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students*	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team. ● Design and manage individual projects. ● Clearly and succinctly communicate their ideas to

should be able to communicate their ideas to technical and non-technical audiences".	technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
N/A	N/A

Topic 1(b)	
AI and Game Media	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic concerns AI for interactive new media such as games. *Level*: **Intermediate, Specialised, Algorithmic, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand/describe the different **creative facets of games** (visuals, audio, narrative, game design, level design and gameplay). ● Students should be able to understand/describe **common algorithms** for procedural content generation (PCG), AI-assisted design (AIAD), as well as game analytics/data mining and affective computing for player modelling.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Analyse and develop (in C/C++, MATLAB or Python) the taught PCG, AIAD and player modelling algorithms, by practically applying their gained knowledge in a systematic manner. ● Evaluate and compare PCG/AIAD/player modelling algorithms, by employing common and appropriate task-specific metrics.
Transferrable/Application It might be commonplace for students to be organised	Students should be able to: ● Work effectively with others in an

into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	interdisciplinary and/or international team. ● Design and manage individual projects. ● Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Player Modelling, UM (Semester Course)	https://www.i-aida.org/course/player-modeling/
Computational Game Creativity, UM (Semester Course)	https://www.i-aida.org/course/computational-game-creativity/

Topic 1(c)	
Web and Social Media Analysis/Mining	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic concerns AI for analysis of WWW and social media data, as well as knowledge mining from large-scale on-line sources. *Level*: **Advanced, Specialised, Theory, Algorithmic.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand/describe **graph-theoretic methods** for **complex network analysis**, basic relevant concepts and algorithms (e.g., for link analysis, centrality measures, etc.) as well as their **applications** in WWW and social media platforms. ● Understand/describe **random graph models**, as well as **algorithms** for community detection, node classification and network information diffusion. ● Understand/describe **Semantic Web** technologies and standards (e.g., RDF, SPARQL, OWL). ● Understand/describe content-based **information retrieval** methods across various modalities.

	• Understand/describe **recommender systems**, the most important relevant concepts and algorithms and their applications in on-line platforms.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to : • Analyse and develop (in C/C++, MATLAB, R or Python) the taught algorithms, by practically applying their gained knowledge in a systematic manner. • Evaluate implementations of the taught algorithms, by employing common and appropriate task-specific metrics. • Use SPARQL.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: • Work effectively with others in an interdisciplinary and/or international team. • Design and manage individual projects. • Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
N/A	N/A

Topic 1(d)	
Human-centred Media Analysis	
Level - Choose one or more labels from the following: • Foundation, Intermediate, Advanced • Broad, Niche, Specialised • Theory, Algorithmic, Methodological	This topic concerns machine learning for affective computing and interaction with humans. *Level*: **Intermediate, Niche, Algorithmic, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome	Students should be able to: • Understand the various **modalities** and ways by which **human emotions** are physically expressed.

might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	• Understand/describe **common unimodal and/or multimodal algorithms** for emotion analysis, facial expression recognition, speech segmentation, speaker recognition, speech emotion recognition, text affect detection, physiological monitoring, body gesture recognition, as well as their **applications** in Human-Computer/Robot Interaction, computer games and education or healthcare settings.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: • Analyse and develop (in C/C++, MATLAB or Python) the taught affective computing algorithms, by practically applying their gained knowledge in a systematic manner. • Evaluate implementations of the taught algorithms, by employing common and appropriate task-specific metrics.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: • Work effectively with others in an interdisciplinary and/or international team. • Design and manage individual projects. • Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
N/A	N/A

Pillar 2: Integrating Approaches for Trustworthy AI (TAILOR)

Description of the Pillar
This pillar of the European AI PhD Curriculum equips students with essential knowledge and skills for developing trustworthy AI systems. Building on the learning outcomes of the AI core pillar, students learn how an integrated approach to learning, optimisation and reasoning can achieve trustworthiness. Foundational topics cover (i) the dimensions of trustworthy AI; and (ii) how to leverage complementary AI paradigms and representations. Advanced topics include (iii) AI agents deciding and learning how to act; (iv) AI agents acting and learning in a society; and (v) ensuring that AI tools and systems are performant, robust and trustworthy.

Topic 2(a)	
Foundations of Trustworthy AI	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic covers the dimensions of Trustworthy AI: (i) Explainability, (ii) Safety, (iii) Fairness, (iv) Accountability and Reproducibility, (v) Privacy, and (vi) Sustainability. *Level*: **Foundation, Broad, Theory, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to understand/describe current discourse on the following questions: ● How can we guarantee user trust in AI systems through explanation? How to formulate explanations as Machine-Human conversation depending on context and user expertise? ● How to bridge the gap from safety engineering, formal methods, verification as well as validation to the way AI systems are built, used, and reinforced? ● How can we build algorithms that respect fairness constraints by design through understanding causal influences among variables for dealing with bias-related issues? ● How to uncover accountability gaps w.r.t. the attribution of AI-related harming of humans? ● Can we guarantee privacy while preserving the desired utility functions? ● Is there any chance to reduce energy consumption for a more sustainable AI and how can AI contribute to solve some of the big sustainability challenges that face humanity today (e.g. climate change)? ● How to deal with properties and tensions of the interaction among multiple dimensions? For instance, accuracy vs. fairness, privacy vs. transparency, convenience vs. dignity, personalization vs. solidarity, efficiency vs. safety and sustainability.
Methodological/Skills	● Students should be able to apply their critical and analytical faculties on specific case studies, in

Topic 2(a)	
These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	order to argue about the need and content of AI trustworthiness issues.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	• Students should be able to work effectively with others in an interdisciplinary and/or international team. • Students should be able to clearly and succinctly communicate their ideas to technical and non-technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

Topic 2(b)	
Reasoning and Learning in Social Contexts	
Level - Choose one or more labels from the following: • Foundation, Intermediate, Advanced • Broad, Niche, Specialised • Theory, Algorithmic, Methodological	This topic covers the foundations, techniques, algorithms and tools for allowing autonomous AI agents to be social and act within societies. It will offer a breadth of understanding in technologies that allow building Social AI systems and a multidisciplinary take on this topic that will impact every aspect of our daily life in the future. *Level*: **Intermediate, Broad, Theory, Algorithmic.**
Learning Outcomes	

Topic 2(b)	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Comprehend that capturing the social aspects of human behaviour is essential in understanding **how people think and how people react to each other**, which is a fundamental step to developing reasoning algorithms that can operate effectively in social contexts. ● Demonstrate a good **understanding of computational models of social reality**. That is, how social contexts determine human behaviour through norms, practices, conventions, rituals and other rules of human social nature. ● Understand current methodologies to model social cognition, collaboration and teamwork. ● Understand /describe theoretical models for cooperation between agents. ● Understand the process of creating systems equipped **with perception and social capabilities** that allows them to adapt to different social contexts and **learn from other agents** in such environments. ● Understand how models of social reality generate emergent behaviour and the impact of such models in agent societies and social networks of multi-agent systems.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Correctly identify different ways to sense the environment and understand how to use off-the-shelf solutions and how to make sense of the captured data. ● Explore the creation of a simple Social AI System, using a perception technology whose data feeds into a reasoning mechanism that outputs social (and intelligent) acts in a context of choice. ● Evaluate social reasoning and learning algorithms in the form of simulations or with a human. ● Analyse the solutions to a problem and critically think about the societal impact.

Topic 2(b)	
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: • Work effectively with others in an interdisciplinary and/or international team. • Design and manage individual projects. • Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

Topic 2(c)	
Automated AI	
Level - Choose one or more labels from the following: • Foundation, Intermediate, Advanced • Broad, Niche, Specialised • Theory, Algorithmic, Methodological	This topic covers meta-level methods to ensure that AI tools and systems are performant, robust and trustworthy. *Level*: **Intermediate, Broad, Theory, Algorithmic.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: • Explain the basic problems solved by AutoAI methods, including (but not limited to) automated algorithm configuration, automated algorithm selection, automated performance prediction, model selection, hyperparameter optimisation and neural architecture search. • Explain, in general and using specific examples, the significance of AutoAI problems and methods for the broader field of AI, including (but not limited to) the areas of machine learning, automated reasoning and

	optimisation. In addition, students should be able to achieve a selection of the following, more specific learning outcomes: ● Demonstrate a working knowledge of Neural Architecture Search, notably how to define search spaces and optimise over these spaces, with both differential and black box methods. ● Identify and define a Hyperparameter Optimization Problem (HPO), specifically in the domain of Algorithm Configuration and Neural Architecture Search. They should also be familiar with hyperparameter importance techniques to interpret different solutions to these problems. ● Be familiar with Gaussian Processes and their modelling capabilities, specifically in the domain of algorithm configuration. ● Assess the strengths and weaknesses of various HPO methods, notably bayesian and evolutionary strategies for doing such. ● Demonstrate knowledge on various speedup techniques to HPO, including leveraging previous information through meta-learning, learning curve prediction and bandit based scheduling techniques. ● Define multiple objectives for an optimization problem and various evolutionary and bayesian techniques for solving such problems. ● Explain Dynamic Algorithm Configuration (DAC) and its difference to Static Algorithm Configuration. They should also be able to demonstrate how to use Reinforcement Learning to solve such optimization problems in DAC. ● Demonstrate knowledge of AutoAI methods for tasks that go beyond supervised learning. This includes knowledge of the underlying theoretical principles and algorithms as well as knowledge of specific tools and systems, including their correct and effective use, strengths and limitations. ● Demonstrate knowledge of AutoAI methods, tools and systems for problems in areas outside of machine learning (i.e., knowledge beyond automated machine learning).

	• Understand the way how AI systems interact with their environment, and what are possible pitfalls of that (e.g., badly calibrated confidence statements, adversarial examples, (un)explainable decisions) • Explain the importance for AI tools and systems to be able to detect situations in which their use becomes problematic (e.g., ineffective or unsafe). • Demonstrate knowledge of techniques and approaches for achieving self-monitoring in at least one major area of AI. • Evaluate AI systems for safety problems in interacting with their environments • Demonstrate awareness of meta-learning, transfer learning, and continual learning techniques that can be leveraged to transfer information from earlier tasks. • Explain how this transfer of knowledge can be used to make AutoML techniques and systems more efficient.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: • Correctly use a range of AutoAI techniques in at least one major area of AI. • Critically assess (in technical and general ways) and explain the limitations of AutoAI methods. • Recognise and explain technical problems that may arise in the use of AutoAI methods.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences*".	Students should be able to: • Work effectively with others in an interdisciplinary and/or international team. • Design and manage individual projects. • Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	

Automated machine learning, Universiteit Leiden	https://studiegids.universiteitleiden.nl/courses/111450/automated-machine-learning
Seminar Trustworthy AI, Universiteit Leiden	(link will appear in March 2023)
Automated Machine Learning, University of Freiburg. (In conjunction with LUHannover, LMU Munich , University of Wyoming)	https://ki-campus.org/courses/automl-luh2021

Pillar 3: Data-Driven Learning (ELISE)

Description of the Pillar
This pillar of the European AI PhD Curriculum equips students with essential knowledge and skills for developing trustworthy AI systems. Building on the learning outcomes of the AI core pillar, students learn ways of reasoning, considering all types of data applicable for almost all sectors of science and industry. Students consider these issues while being aware of data safety and security, and striving to explainable and trustworthy outcomes.

Topic 3(a)	
Deep Learning	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic covers the design, implementation, and optimization of deep learning architectures. *Level*: **Intermediate, Broad, Algorithmic, Methodological.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand and design different deep learning architectures including convolutional and residual networks, transformers and diffusion models. Develop architectures for high dimensional data and also dynamic/time varying data. ● Design and implement different optimization procedures, using differentiable programming (automatic differentiation/ backpropagation), as well as optimization by stochastic gradient methods combined with acceleration and variance reduction techniques (minibatching/normalisation).
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model*	Students should be able to: ● Design and develop deep learning solutions taking advantage of available libraries and compute infrastructure. ● Correctly set the value of hyperparameters for model evaluation and model selection in deep networks

selection."	• Evaluate the accuracy of the derived solutions in a systematic way, using available benchmarks and considering different performance metrics.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: • Work effectively with others in an interdisciplinary and/or international team. • Design and manage individual projects. • Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

Topic 3(b)	
Machine Learning Theory	
Level - Choose one or more labels from the following: • Foundation, Intermediate, Advanced • Broad, Niche, Specialised • Theory, Algorithmic, Methodological	This topic covers the mathematical foundations of machine learning considering different learning paradigms and all facets of machine learning theory, including provable generalisation guarantees as well as optimization and numerical analysis. *Level*: **Foundation, Broad, Theory, Methodological**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should achieve a solid background on the • Theoretical and mathematical aspects needed to analyse the theoretical performance limitations and possible improvements of modern machine learning algorithms under different modelling assumptions on the data and the learning agent. • Computational and algorithmic aspects underlying different machine learning systems. To be able to understand and derive rigorous

	complexity guarantees.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Derive rigorous guarantees on the learning accuracy of different algorithms using a broad range of tools, e.g., from statistics, probability, game theory, optimization, and stochastic processes. ● Derive rigorous guarantees on the computational efficiency of different algorithms using tools from fields including optimization, numerical analysis, and analysis of algorithms.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team. ● Design and manage individual projects. ● Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

Topic 3(c)	
Reinforcement Learning and Sequential Decision-making	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic covers the study and design of machine learning algorithms for online learning, multi-armed bandits and reinforcement learning (RL). *Level*: **Intermediate, Broad, Algorithmic, Methodological.**
Learning Outcomes	

Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand the difference between online and batch learning. ● Describe the main online learning algorithms and understand the analysis of their performance. ● Understand the multi-armed bandit problem, describe the main algorithms, and understand the analysis of their performance. ● Understand the goal of reinforcement learning and the mathematical MDP model. ● Describe the basic evaluation criteria for RL: finite, infinite, and discounted horizon. ● Describe the main algorithms for model-based RL and understand their performance guarantees. ● Describe the main algorithms for model-free RL and understand their performance guarantees. ● Understand value function approximation and deep RL.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Design RL solutions for new problems using a correct MDP abstraction. ● Implement RL algorithms taking advantage of available libraries and simulation environments. ● Evaluate the accuracy of the derived solutions in a systematic way, using available benchmarks and considering different performance metrics.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team. ● Design and manage individual projects. ● Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
	URL
	URL

Topic 3(d)	
Distributed and Federated Learning	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic explores the realm of distributed and federated machine learning. communication efficiency, and scheduling. Furthermore, it also covers the domain of multi-agent systems. Level: **Intermediate, Broad, Algorithmic, Methodological**.
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand and design different deep learning architectures under a distributed environment, including edge computing. ● Understand the challenges of using a computing infrastructure, including communication network, computational resource management, fault tolerance, and privacy concerns derived. ● Understand and design different aspects of multi-agent systems, including distributed problem solving and planning, search algorithms for agents, cooperation and coordination of agents, multi-agent learning, agent negotiation, and agent-oriented software engineering ● Design and implement different optimization procedures, using distributed and federated learning in distributed computing infrastructures.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Derive rigorous guarantees on the learning accuracy of different algorithms using a broad range of tools. ● Derive rigorous guarantees on the computational efficiency of different algorithms. ● Implementing different strategies in a distributed

	computing infrastructure.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences*".	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team. ● Design and manage individual projects. ● Clearly and succinctly communicate their ideas to technical audiences
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
	URL
	URL

Topic 3(e)	
Generative Artificial Intelligence	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic concerns generative artificial intelligence for image/video generation, with its courses ranging from algorithmic tools up to specialised applications. It covers both theoretical knowledge and its real-world applications (e.g., in entertainment). *Level*: **Intermediate, Advanced, Specialised, Theory, Algorithmic**.
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	Students should be able to: ● Understand and design different generative models including variational autoencoders, generative adversarial networks and diffusion models. ● Understand/describe the operation and the prompting procedure of different generative models, as well as their differences/modifications

	for text/image/video generation.
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Design and develop models and algorithms for text/image/video generation using existing programming frameworks. ● Perform model tuning and hyperparameter selection for text/image/video generation. ● Effectively prompt and finetune existing generative models and algorithms. ● Evaluate the quality, diversity and speed of the generative AI models.
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: ● Work effectively with others in an interdisciplinary and/or international team. ● Design and manage individual projects. ● Clearly and succinctly communicate their ideas to technical audiences.
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

Pillar 4: Human-Centred AI (HumaneAINet)

Description of the Pillar
There is broad agreement on a high level definition of "Human-Centred AI" as AI that is "beneficial to humans and society". There is much less agreement on translating that broad vision into specific lower level technological requirements and scientific research questions. This is on one hand due to the fact that the notions of being beneficial to humans and society differ from stakeholder to stakeholder (and between cultures). On the other hand it is difficult (if not impossible) to come up with a 1:1 complete and consistent translation between a notion of "being beneficial" defined on a policy level and concrete technical properties of an AI system. This pillar provides the students with the basic knowledge needed to design, implement, operate and research the next generation of Human-Centred AI systems that are focused on enhancing Human capabilities and optimally cooperating with humans on both the individual and the society levels. The pillar covers fundamental topics of Human-Centred AI related to the interface between AI and Human-Machine Interaction, Computational Social Science, Philosophy, Ethics, Law.

Topic 4(a)	
Human-Centred Machine Learning	
Level - Choose one or more labels from the following: ● Foundation, Intermediate, Advanced ● Broad, Niche, Specialised ● Theory, Algorithmic, Methodological	This topic covers methods to develop Human-Aware ML algorithms and models *Level*: **Intermediate, Broad, Theory, Algorithmic.**
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	● To understand the adaptation of the Machine Learning and Transformers architectures to deal with audio, text and visual data ● To grasp the functioning principles of recent work applying Transformers architecture to multimodal tasks and data. ● To be aware of the human interventions in interactive machine learning ● To understand human teaching strategies and gain knowledge about learning from human feedback, demonstrations and instructions ● …
Methodological/Skills	

These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	• To define • To define computational models of human teaching strategies • <Outcome 2> • <Outcome 3> • …
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences*".	• To discuss applications of human-centered machine learning • <Outcome 2> • <Outcome 3> • …
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

Topic 4(b)	
Explainable AI	
Level - Choose one or more labels from the following: • Foundation, Intermediate, Advanced • Broad, Niche, Specialised • Theory, Algorithmic, Methodological	This topic provides an introduction to Explainable AI with perspectives from different disciplines such as computer science and social science. *Level*: **Intermediate, Broad, Theory, Algorithmic.**
Learning Outcomes	

Content/Knowledge
 	
These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."
	Students should achieve a solid background on the:
	• Exploring the **Motivation and Definition of Explainability**
	  ○ Understanding the importance of explainability in AI systems
	  ○ Defining **related terms** such as

	interpretability and transparency ● Accounts of Explanation in **Social Sciences Literature** ○ Examining the primary perspectives on explanation in social sciences ○ Drawing insights from existing theories and frameworks ● Local and Global Explanation Methods ○ Differentiating between **local and global explanation techniques** ○ Exploring **model-agnostic methods** for generating explanations ● Interpretable Models and Post-hoc Explanations ○ Investigating the concept and utility of **interpretable models** ○ Understanding **post-hoc explanations** and their role in XAI ● Explaining One-shot Decision and Sequential Decision Making Models ○ Contrasting **explainability in one-shot decision scenarios** ○ Analyzing the unique **challenges of sequential decision making models** ● Explainability in Embodied and Non-Embodied AI Systems ○ Investigating the role of explainability in AI systems with **physical embodiment** ○ Analyzing the implications of explainability in non-embodied AI systems ● The Role of Causality and Interactivity in XAI ○ Understanding how **causality** influences explainability in AI systems ○ Exploring the significance of **interactivity** for generating explanations
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	Students should be able to: ● Implementation of Explainable Machine Learning Techniques ○ Understanding and implementing techniques such as **LIME, SHAP, PDP, CNNs with attention**, and **GANs with explainability** ● Familiarity with **XAI Python Libraries** ○ Getting acquainted with popular XAI libraries like **Captum** for practical implementation ● Evaluation of XAI Methods

	○ Applying **quantitative and qualitative metrics** to evaluate XAI methods ○ Incorporating **human-centric evaluations** and conducting user studies ● Application of XAI Methods to **Real-World Datasets** ○ Applying XAI techniques to real-world datasets for practical insights and interpretation
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences"*.	Students should be able to: ● Effectively work with team members from diverse backgrounds and cultures ● Demonstrate proficiency in project planning and management ● Effectively communicate technical ideas to audiences with varying levels of expertise ● Ensure Reproducibility and Collaborative Work ○ Implement best practices for reproducibility in research and project work ○ Utilize tools and platforms for collaborative work and sharing code efficiently (i.e., Github)
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

Topic Description Template for extra copies

Topic (Subject)	
<Name of the topic>	
Level - Choose one or more labels from the following: • Foundation, Intermediate, Advanced • Broad, Niche, Specialised • Theory, Algorithmic, Methodological	
Learning Outcomes	
Content/Knowledge These concern the topic itself. For example, in a Natural Language Processing module, one such learning outcome might be: "*Students should be able to compare TF-IDF representations of text with learned word embeddings*."	• <Outcome 1> • <Outcome 2> • <Outcome 3> • …
Methodological/Skills These concern skills that contribute to designing, building and, above all, evaluating, AI systems. For example, in Machine Learning, a learning outcome might be: "*Students should be able to correctly set the values of hyper-parameters for performance estimation and model selection*."	• <Outcome 1> • <Outcome 2> • <Outcome 3> • …
Transferrable/Application It might be commonplace for students to be organised into teams for group work, e.g.: Students should be able to work effectively with others in a team. Commonplace too will be presentations of their work, e.g.: "*Students should be able to communicate their ideas to technical and non-technical audiences*".	• <Outcome 1> • <Outcome 2> • <Outcome 3> • …
Courses available through AIDA or other online courses linked to the networks that cover elements of this material (if these exist)	
Course Name, Affiliation/Source	URL
Course Name, Affiliation/Source	URL

Appendices

Feel free to include additional materials, links, definitions, glossary, etc.

Notes from TAILOR (Peter Flach)

TAILOR has produced a mapping of PhD-level training offered by TAILOR partners which is accessible here:

https://github.com/TAILOR-UoB/deliverables/blob/gh-pages/deliverables/d95.md

This is now about a year out of date but is intended to be updated on a regular basis.

Towards the end of the document there is some discussion how this could be developed into a curriculum using tracks, themes and skills:

https://github.com/TAILOR-UoB/deliverables/blob/gh-pages/deliverables/d95.md#towards-a-joint-tailor-phd-curriculum

We have also started work on a glossary of terms relevant to PhD curricula. The current draft is reproduced below; note in particular the difference between PhD curriculum (specification) and PhD programme (delivery).

Glossary

Cohort-based doctoral training - PhD programmes are increasingly offered in a cohort-based way, where a group of students work on thematically related PhD topics. Such programmes can be organised within a single University (e.g., Centres for Doctoral Training in the UK) or by a consortium of universities (e.g., the WASP Graduate School in Sweden).

Learning outcome - the measurable skills, abilities, knowledge and values that a **PhD student** will be able to demonstrate as a result of successfully completing a given **training module**.

PhD curriculum - a specification of the structure and content of a **PhD programme**. This typically concentrates on the doctoral training part, which may include **training modules**, self-study, practical assignments, internships and research visits, the acquisition of **transferable skills**, etc. Elements of the doctoral training programme are typically delivered to cohorts of students on the same **PhD programme** (*see* **cohort-based doctoral training**), or to groups of students across different **PhD programmes** but within one University. Completion of parts or all of the doctoral

training programme may be mandatory in order to progress to (or proceed with) the research part of the **PhD programme**, in which case the corresponding modules will be examined through written or oral exams and/or coursework.

PhD degree - a doctoral degree conferred by an accredited institution for higher education (typically a University).

PhD examination - the process by which PhD examiners determine whether a **PhD thesis** is worthy of awarding a **PhD degree**. This can involve one or more of the following:

- written reports from the examiners after reading the thesis;
- a *viva voce* examination behind closed doors;
- a public PhD defence.

The outcome of the PhD examination can be binary (whether the PhD is awarded or not) or may stipulate minor or major corrections of the written thesis.

PhD programme - a combination of research training and supervised research practice, offered by an accredited institution for higher education, with the aim of achieving a **PhD degree** after submitting and successfully defending a **PhD thesis**. **PhD students** need to formally enrol on a PhD programme, typically 3-4 years in duration, for which they may need to pay a tuition fee.

PhD student - a person enrolled on a **PhD programme** with the aim of obtaining a **PhD degree** after several years of research training and practice.

PhD supervisor - an academic with responsibility for guiding a **PhD student** during their PhD studies, and monitoring their progress.

PhD thesis - a written treatise on a chosen research topic, typically 150-200 pages long, submitted at the end of a **PhD programme** to satisfy the requirements of being awarded a **PhD degree**, subject to **PhD examination**.

Training module - part of a **PhD programme** typically consisting of a series of lectures or seminars, and (if mandatory) assessed through written or oral exams and/or practical assignments.

Transferable skill - a skill that is relevant for successful PhD study but more generally applicable, such as writing and presentation skills.