Syllabus of the Master in Artificial Intelligence

Module title	P01 – AI Company Strategy and Project Definition
Academic supervisors	Olivier Bornet Dr. Joël Dumoulin
Learning manager	-
ECTS Credits	10
Type of module	On-the-job learning
Description	The module provides the student with project planning skills. The goal is to define the project(s) that will be developed in module P02. It can be divided in 6 parts: • Familiarize with and understand your company, its corporate culture and strategy. • Determine how the artificial intelligence takes part in your company's strategy • Analyze state-of-the-art development of AI in the field of interest • Design proof(s) of concept to guarantee the success of your project(s) • Set a roadmap to plan your project(s) development • Define the terms of reference of your project(s)
Bibliography	-

Module title	P02 – AI Project(s) Development
Academic supervisors	Olivier Bornet Dr. Joël Dumoulin
Learning manager	-
ECTS Credits	30
Type of module	On-the-job learning
Description	The aim of Module P02-AI Project(s) development is to develop the project(s) the student defines in Module P01 – AI Company strategy and Project(s) definition.
Bibliography	-

Module title	M01 – Linear Algebra and Probabilities
Academic supervisors	Dr. Théophile Gentilhomme, Dr. Ina Kodrasi
Learning manager	Parvaneh Janbakhshi
ECTS Credits	2
Type of module	Fundamental
Description	This module covers Linear Algebra topics from basic matrix/vector operations to singular value decomposition and Probability topics from fundamentals to Markov chains, which are prerequires for most of the AI courses. The module will be directed by examples and intuition rather than formalism. Python language will be used in examples and exercises. Although this module covers most of the basics, it is assumed students have notions and (computer science) bachelor background in linear algebra, probability and programming. Linear Algebra • Review on vector and matrix and geometric interpretations • Review on matrices operations • Solving linear equation: Ax = b • Elimination, permutation, inverse matrix • LU decomposition • Vector space and fundamental subspaces • Rank, solution with null space component • Independence, basis and dimensions • Linear transform and change of basis: Fourier and wavelet • Orthogonality and projection • Least square approximation • Orthogonal basis • Determinant • Eigenvalue decomposition (example of Markov matrices) • Symmetric and positive definite matrices (Covariance matrix) • Singular value decomposition and pseudo inverse • Weighted least square and statistics Probability • Sets, events, probability theory, axioms and derived laws • Conditional probability related rules and theorems (including Bayes!) • Independence • Random variables discrete and continuous • Function of RV • Expectation, variance • Markov chains, steady state, absorption
Bibliography	**Recommended Literature:** Strang, G., (2016). Introduction to linear algebra (Vol. 5). Wellesley, MA: Wellesley-Cambridge Press Bertsekas, D. P., Tsitsiklis, J. N. (2002). Introduction to probability (Vol. 1). Belmont, MA: Athena Scientific.

Module title	M02 – Data structure and algorithms for AI
Academic supervisor	Olivier Bornet
Learning managers	P. Abbet, W. Droz, F. Tarsetti, S. Kayal
ECTS Credits	4
Type of module	Fundamental
Description	The course gives global knowledge in data structure and algorithms. It is organized in 4 parts: 1. Introduction 2. Data structures and algorithms 3. Advanced algorithms 4. Computing tools
Bibliography	-

Module title	M03 – Signal Processing
Academic supervisor	Dr. Michael Liebling
Learning manager	-
ECTS Credits	4
Type of module	Fundamental
Description	1. Signals and Signal Processing • Classification of Signals • Simple Time-Domain Operations, Filtering, sampling 2. Discrete-Time Signals and Systems • Time-Domain Representation • Sampling Rate Alteration 3. Discrete-Time Fourier Transform • The Continuous-Time Fourier Transform • The Discrete-Time Fourier Transform • Discrete-Time Fourier Transform Theorems • Digital Processing of Continuous-Time Signals 4. Discrete-Time Systems • Discrete Time System Examples • Classification of Discrete Time Systems (FIR/IIR) • Frequency Response

	5. Finite-Length Discrete Transforms • Orthogonal Transforms • The Discrete Fourier Transform • Relation Between the Fourier Transform and the DFT and Their Inverses • Circular Convolution • DFT properties and theorems • Computation of the DFT of Real Sequences • Linear Convolution Using the DFT 6. z-Transform • Computation of the Convolution Sum of Finite-Length Sequences • The Transfer Function • Transfer Function Expression, relation to frequency response 7. LTI Discrete-Time Systems in the Transform Domain • characterization of LTI systems, stability • filter design 8. DSP Algorithm Implementation • Computation of the Discrete Fourier Transform • Splines and wavelets • Multirate Filter Banks and Wavelets
Bibliography	Compulsory Literature: Fawwaz T. Ulaby and Andrew E. Yagle, “Signals and Systems,” ISBN: 978-1-60785-486-9 (harcopy)/978-1-60785-487-6 (electronic) Free download of PDF: http://ss2.eecs.umich.edu Additional Literature: Alan V. Oppenheim, Ronald W. Schafer, “Discrete-Time Signal Processing (3rd Edition)” Prentice-Hall Signal Processing Series, ISBN-13: 978-0131988422, ISBN-10: 0131988425 Sanjit Mitra, “Digital Signal Processing,” Mcgraw Hill Higher Education; 4th edition (2010), ISBN10: 0071289461, ISBN-13: 978-0071289467 Martin Vetterli, Jelena Kovacevic, Vivek K. Goyal “Foundations of Signal Processing” 3rd Edition ISBN- 10: 9781107038608, ISBN-13: 978-1107038608

Module title	M04 – Foundations in Statistics for AI
Academic supervisors	Dr. Philip N. Garner, Dr. Ina Kodrasi
Learning managers	Alexandre Bittar, Louise Coppieters
ECTS Credits	4
Type of module	Fundamental
Description	The syllabus is divided into five topics each spanning two weeks. The first two topics are introductory, covering basic statistical modelling. The three subsequent ones are more involved, covering the estimation of parameters (the core of machine learning), inference (the basics of AI), and finally testing. Topic 1: Discrete distributions • Introduction to the course, it's teaching staff and structure • Origins and characteristics of some important discrete distributions • Introductory exercises and python notebook based labs Topic 2: Continuous distributions • Derivation of some important continuous distributions • Introduction to Bayesian concepts including conjugacy • Conclusion of the first two "introductory" topics Topic 3: Estimation • Bayesian estimation • Minimum mean squared error • Maximum likelihood and Maximum a-Posteriori concepts Topic 4: Inference and priors • Bayesian inference and the predictive distribution • Prior elicitation • The non-informative (Jeffreys) prior Topic 5: Testing • Hypothesis testing • Tests based on a normal assumption • Approximating non-normal cases
Bibliography	**Compulsory Literature:** There is no compulsory literature **Additional Literature:**

	There is no single course text; much of the literature will be supplied or is on the web. However, the following may be useful throughout the course. 1. David Williams (2001). Weighing the Odds: A Course in Probability and Statistics. The Edinburgh Building, Cambridge CB2 8RU, UK: Cambridge University Press. ISBN: 978-0-521-00618-7 2. Charles M. Grinstead and J. Laurie Snell (1997). Introduction to Probability. 2 Revised edition (July 1, 1997). American Mathematical Society. ISBN: 978-0821807491. URL: http://math.dartmouth.edu/~prob/prob/prob.pdf 3. Larry Wasserman (2004). All of Statistics: A Concise Course in Statistical Inference. Springer Texts in Statistics. Springer. DOI: 10.1007/978-0-387-21736-9. URL: https://link.springer.com/book/10.1007%2F978-0-387-21736-9 The second of these is freely available; the last is no longer free, but we will provide a pdf copy. In general, if you have done a course in statistics before and still have the course text, it should be useful for this course.

Module title	**M05 – Open Science and Ethics**
Academic supervisors	Dr. S. Marcel, Dr. A. Anjos
Learning managers	Marie-Constance Landelle, François Charlet, Flavio Tarsetti, Dr. Joël Dumoulin
ECTS Credits	2
Type of module	Fundamental
Description	• AI and the Law • AI and Data Protection • AI and Ethics • Reproducibility, What is it? • Data Organization and Evaluation • Version Control with git • Code Sharing with GitLab • Unit Testing and Continuous Integration • Documentation and Reporting • Packaging and Deployment
Bibliography	-

Module title	M06 – Fundamentals in Machine Learning 1
Academic supervisors	Dr. S. Marcel, Dr. A. Anjos, Dr. J. Henderson, Dr. J.M. Odobez
Learning managers	Dr. Pavel Korshunov, Dr. Tiago de Freitas Pereira, Dr. Michael Villamizar, Dr. Olivier Canévet, R. Karimi Mahabadi
ECTS Credits	4
Type of module	Fundamental
Description	• Linear regression • Logistic Regression • Decision Trees • Boosting • Multi-layer Perceptron
Bibliography	Compulsory Literature: C. Bishop, “Pattern Recognition and Machine Learning”, Springer, 2006.

Module title	M07 – Introduction to Image Processing and Computer Vision
Academic supervisors	Dr. Michael Liebling, Dr. Jean-Marc Odobez
Learning managers	Dr. Olivier Canévet, Olivia Mariani
ECTS Credits	4
Type of module	Fundamental
Description	This class covers basic concepts in image and video processing as well as computer vision. Topics include image formation and sampling, image transforms, image enhancement, and image and video compression. Computer vision topics include points of interest, optical flow, and camera calibration. • Introduction to Digital Image processing (imaging types and formats, applications) • Point operations, image histograms • Spatial Filtering and convolutions • Edge detection • 2D Fourier Transforms and representation of images, sampling, and image resizing (low pass filters, pyramids) • Color images and color transformations • Interest points (detection, representation, invariance, matching, RANSAC...) • Calibration • Optical Flow

Bibliography	Compulsory Literature: • Andrew E. Yagle and Fawwaz T. Ulaby “Image Processing for Engineers,” Michigan Publishing, Ann Arbor, Michigan, 2018, ISBN: 978-1-60785-488-3 (hardcopy), 978-1-60785-489-0 (electronic) Additional Literature: • Rafael C. Gonzalez and Richard E. Woods, “Digital Image Processing Third Edition,” Pearson, ISBN 9780131687288, 2008

Module title	**M08 – Fundamentals in Machine Learning 2**
Academic supervisors	Dr. S. Marcel, Dr. A. Anjos, Dr. J. Henderson, Dr. J.M. Odobez
Learning managers	Dr. Pavel Korshunov, Dr. Tiago de Freitas Pereira, Andreas Marfurt, Dr. Olivier Canévet,
ECTS Credits	4
Type of module	Fundamental
Description	• Dimensionality Reduction and Clustering • Kernel Methods and Support Vector Machines • Graphical Models • Exact and Approximate Inference in Bayesian Networks • Probability Distribution Modelling
Bibliography	Compulsory Literature: C. Bishop, “Pattern Recognition and Machine Learning”, Springer, 2006.

Module title	**M09 – Introduction to Speech Processing**
Academic supervisors	Dr. Mathew Magimai Doss
Learning managers	Pavankumar Dubagunta, Julian Fritsch, Enno Hermann, Dr. Ravi Prasad
ECTS Credits	4
Type of module	Fundamental
Description	This course will introduce the students the fundamentals of speech processing and provide them with the key formalisms, models and algorithms to implement speech processing applications such as, speech recognition, speech synthesis, paralinguistic speech processing, multichannel speech processing. Course content Introduction

	why speech processing? speech production, speech perception, basic phonetics Speech signal analysis Sampling, Quantization, Time domain processing, Frequency domain processing, linear prediction, cepstrum, speech coding Practical: Speech signal analysis in Octave and Praat Machine learning for speech processing Static classification, Sequence classification, Regression Practical: Statistical pattern recognition, Hidden Markov models in Octave Automatic speech recognition Dynamic programming, Instance-based speech recognition, Hidden Markov model-based speech recognition, Evaluation measures Practical: Kaldi tutorial Text-to-speech synthesis Concatenative speech synthesis, Statistical parametric speech synthesis, Evaluation measures Practical: Grapheme-to-phoneme conversion, HMM-based speech synthesis Paralinguistics speech processing Emotion, gender, accent, pathological speech assessment, Evaluation measures Practical: OpenSMILE tutorial
Bibliography	Suggested textbooks: 1. B. Gold, N. Morgan and D. Ellis, ``Speech and Audio Signal Processing'', Wiley Publications, 2011. 2. P. Taylor, ``Text-to-Speech Synthesis'', Cambridge University Press, 2009. 3. X. Huang, A. Acero and H-W. Hon, ``Spoken Language Processing: A Guide to Theory, Algorithm and System Development'', Prentice Hall, 2001. 4. B. Schuller and A. Batliner, ``Computational Paralinguistics: Emotion, Affect and Personality in Speech and Language'', Wiley, 2013. Online Literature: SCOOT (https://www.isca-speech.org/iscaweb/index.php/scoot) Software tools: Octave, Praat, Kaldi, HTS, OpenSMILE (needed for practical)

Module title	M10 – Deep Learning
Academic supervisor	Dr. Olivier Canévet
Learning manager	Alexandre Nanchen
ECTS Credits	4
Type of module	Fundamental
Description	The “deep learning” course aims at providing an overview of the deep learning area (theory, methods, applications, and tools) to the students and to use the most important tools via practical sessions. Labs content Bias-variance dilemma through k-NN and polynomial fitting, mini deep-learning framework in numpy, PyTorch basics, MLP on MNIST and CIFAR, LeNet5 on MNIST and CIFAR, optimisation algorithms, finetuning on a pre-trained network, creation and training of model for person detection from top-view depth images, YOLOv3 for object detection, adversarial examples, distribution of activation maps as embeddings.
Bibliography	-

Module title	A01 - Biometrics
Academic supervisor	Dr. Sébastien Marcel
Learning managers	Dr. Vedrana Krivokuca Hahn, Laurent Colbois, Hatef Otroshi Shahreza, Alex Unnervik
ECTS Credits	4
Type of module	Advanced
Description	This module aims to provide an introduction to Biometrics and an understanding of the main biometric modalities (such as face or fingerprint) and associated recognition algorithms. The module addresses also the security and privacy preservation aspects of biometrics. Lecture 1: Introduction to Biometrics Lecture 2: Evaluation of Biometric Systems Lecture 3: Iris Recognition Lecture 4: Face Recognition Lecture 5: Vascular Recognition Lecture 6: Fingerprint Recognition Lecture 7: Multibiometrics Lecture 8: Speaker Recognition Lecture 9: Presentation Attack Detection Lecture 10: Biometric Template Protection
Bibliography	Compulsory Literature:

	Books: 1. Jain, A. K. et al., “Introduction to Biometrics” (2011) 2. Jain, A. K. et al., “Handbook of Biometrics” (2008) Additional Literature: Books: 1. Li, S. Z. and Jain, A. K., “Handbook of Face Recognition” (2011) 2. Maltoni, D. Et al., “Handbook of Fingerprint Recognition“ (2009) 3. Bowyer, K. V. and Burge, M. J. “Handbook of Iris Recognition“ (2016) 4. Marcel, S. et al. “Handbook of Biometric Anti-Spoofing“ (2019) 5. Uhl, A. et al. “Handbook of Vascular Biometrics“ (2019)

Module title	**A02 – Multimodal Computational Sensing of People**
Academic supervisor	Dr. Jean-Marc Odobez
Learning manager	Dr. Michael Villamizar
ECTS Credits	2
Type of module	Advanced
Description	Multimodal processing is at the core of the perception of our world: we see, we hear, we touch, we smell, we taste, and we move. Being able to analyse and combine multimodal streams of data is therefore an inherent ability that should be addressed by artificial intelligence systems, and comprises several core challenges like how to represent multimodal information? how to deal with modality asynchrony? how to fuse complementary or redundant information? how to do co-training of models? The course will cover this topic, with a particular emphasis on the analysis of people and of their behaviors, including interactions, from multimodal sensor streams (with a bias towards vision and audio). We will rely on Bayesian statistical models and deep learning as main modeling tools.
Bibliography	Prerequesite: - machine learning, including Bayesian networks and temporal models (HMMs) - deep learning Documents will be provided along with the course. An example of reference paper is: Multimodal Machine Learning: A Survey and Taxonomy T. Baltrusaitis, C. Ahuja, LP Morency, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol 41(2), pp 423-443, Feb. 2019. Social signal processing: Survey of an emergin domain

	A. Vinciarelli, M. Pantic and H. Bourlard, Image and Vision Computing, Vol 27 (12), pp 1743-1759, Nov. 2009. Automatic nonverbal analysis of social interaction in small groups D. Gatica-Perez Image and Vision Computing, Vol 27 (12), Nov. 2009.

Module title	**A03 – Natural Language Processing**
Academic supervisor	Dr. James Henderson
Learning manager	Andrei Catalin Coman, Melika Behjati
ECTS Credits	4
Type of module	Advanced
Description	The Natural Language Processing module introduces the basic NLP tasks and the main deep learning models currently being applied to them. We will cover learning about words, sentences, and translations between languages, using representation learning, recurrent neural networks, and attention-based models. 1. Introduction to NLP 2. Representation learning for words 3. Practical on word embeddings 4. Modelling text with recurrent neural networks 5. Practical on recurrent neural networks for text 6. Machine translation with sequence-to-sequence models 7. Practical on seq2seq model for machine translation 8. Attention-based models of text 9. Practical on Transformer and pretraining
Bibliography	Compulsory Literature: None. Additional Literature: This course material is based on the Stanford course “Natural Language Processing with Deep Learning” by C. Manning. See http://web.stanford.edu/class/cs224n/ for suggested reading and additional topics.

Module title	A04 – Robotics
Academic supervisor	Dr. Sylvain Calinon
Learning manager	Hakan Girgin
ECTS Credits	4
Type of module	Advanced
Description	The Robotics module aims at discovering the recent trends and challenges of AI in robotics, with a specific focus on robot skill acquisition from demonstration, adaptive control, and human-robot collaboration. Week 1: Introduction - A brief history of robotics and autonomous machines - Learning from demonstration (observational learning, kinesthetic teaching, correspondence problems) - Tools for AI in robotics (simulators, ROS middleware) Week 2: Movement primitives I - Multivariate Gaussian distributions (recap) - Newton's method (recap) - Decomposition as a superposition of basis functions - Bézier curves and Bernstein polynomials Week 3: Movement primitives II - Weighted least squares (recap) - Locally weighted regression - Probabilistic movement primitives - Dynamical movement primitives Week 4: Operational space control - Forward kinematics - Inverse kinematics - Task prioritization and nullspace control Week 5: Human-robot collaboration - Linear dynamical systems - Torque control, dynamic model (inertia, ...), gravity compensation - Impedance control Week 6: Anticipation and planning I - Linear quadratic tracking (LQT) - Model predictive control (MPC) Week 7: Anticipation and planning II - Iterative linear quadratic regulator - LQT applications examples Week 8: Ergodic control for exploration behaviors - Fourier series (recap) - Spatial coverage problems Week 9: Manifolds in robotics

	- Representations of orientation (unit quaternions) - Riemannian geometry - Regression, fusion and control on Riemannian manifolds Week 10: Recap of the topics covered in the course & examples of applications - Products of experts (information fusion, Kalman filter) - Learning by exploration (cross entropy methods) - Task-parameterized models
Bibliography	Compulsory readings: Week 1: Learning from demonstration https://calinon.ch/papers/Calinon-EncyclopediaRobotics2019.pdf (to get familiar with tools for AI in robotics, see also the list of additional recommended material below) Week 2: Movement primitives I Section 2.2 of https://calinon.ch/papers/Calinon_MMchapter2019.pdf Week 3: Movement primitives II Sections 2.1 and 3 of https://calinon.ch/papers/Calinon_MMchapter2019.pdf Week 4: Operational space control (see list of additional recommended material below) Week 5: Human-robot collaboration (see list of additional recommended material below) Week 6-7: Anticipation and planning Section 4.1 of https://calinon.ch/papers/Calinon-Lee-learningControl.pdf Week 8: Ergodic control for exploration behaviors Sections 2.3 and 2.4 of https://calinon.ch/papers/Calinon_MMchapter2019.pdf Week 9: Manifolds in robotics https://calinon.ch/papers/Calinon-RAM2020.pdf Week 10: Examples of applications Section 4.2 of https://calinon.ch/papers/Calinon-Lee-learningControl.pdf Additional recommended material: General robotics textbook: Modern Robotics (available in open source format, including youtube videos and source codes) http://hades.mech.northwestern.edu/index.php/Modern_Robotics http://hades.mech.northwestern.edu/index.php/Modern_Robotics_Videos Robotics simulator targeting AI applications: https://pybullet.org/wordpress/ ROS courses and videos: https://rsl.ethz.ch/education-students/lectures/ros.html Other ROS resources: http://wiki.ros.org/Courses