MS & PhD Program in Artificial Intelligence
(Designed by NCAI)

Background:

The National Center of Artificial Intelligence (NCAI) has top-notch AI scientists/researchers/students selected by Higher Education Commission (HEC) on competitive grounds. Under the umbrella of NCAI and partner universities, the program is proposed and designed by the NCAI. Public/private Universities may use this curriculum as a guide to design their own courses, by acknowledging NCAI and may also provide us with their valuable feedback/suggestion regarding the curriculum. The National Center of Artificial Intelligence (NCAI), the latest technological initiative of the Government of Pakistan has taken the lead by starting MS & PhD Program in Artificial Intelligence to offer specializations in different aspects of AI at the graduate level to help establish and grow the industry in Pakistan.

Program Educational Objectives:

PEO1: Have a strong competence in Artificial Intelligence resulting in successful careers.
PEO2: Pursuing research and innovation and be able to provide modern solutions to technical problems.
PEO3: To apply as well as create Artificial Intelligence based knowledge at par with the developments at both national and international level.

Curriculum:

The proposed curriculum is unified for all NCAI partner universities. For the sake of uniformity and ease of transfer of courses, a national course code has also been defined for each course. This will be treated as a reference for course compatibility between NCAI partner institutions.

Core Courses

National Course Code	Course Name	CHs
AI 801	Artificial Intelligence	(3)
AI 802	Machine Learning	(3)

AI 803	Mathematics for Artificial Intelligence	(3)

MS Elective Courses

National Course Code	Course	CHs
AI 811	Deep Learning	(3)
AI 812	Probabilistic Graphical Models	(3)
AI 813	Multi-agent Systems	(3)
AI 814	Knowledge Representation and Reasoning	(3)
AI 815	Neural Networks	(3)
AI 818	Probabilistic Robotics	(3)
AI 819	Text Analytics	(3)
AI 821	Computer Vision	(3)
AI 822	Advanced Image Processing	(3)
AI 823	Speech Processing	(3)
AI 824	Data Acquisition and Control	(3)
AI 825	Sensors and Sensing	(3)
AI 826	Human Robot Interaction	(3)
AI 827	Simultaneous Localization and Mapping	(3)

AI 828	Complex Adaptive Systems	(3)
AI 829	Social Simulations	(3)
AI 831	Intelligent Systems	(3)
AI 832	Reinforcement Learning	(3)
AI 834	Robotic Grapping and Fixture	(3)
AI 835	Rehabilitation and Assistive Robotics	(3)
AI 836	BioRobotics	(3)
AI 837	Intelligent Transportation Systems	(3)
AI 838	Serious Games	(3)
AI 841	Data Mining	(3)
AI 842	Information Retrieval	(3)
AI 843	Natural Language Processing	(3)
AI 844	Semantic Web	(3)
AI 845	Internet of Things	(3)
AI 846	Computational Creativity	(3)
AI 851	Advanced Signal Processing	(3)
AI 852	Modelling and Simulation	(3)
AI 853	Advanced Programming in Python	(3)

AI 854	Data Analysis and Visualization	(3)
AI 855	Cyber Security	(3)
AI 856	Ethical Machines	(3)

PhD Courses

AI 916	Pattern Recognition	(3)
AI 917	Evolutionary Algorithms	(3)
AI 928	Advanced Computer Vision	(3)
AI 937	Intelligent Systems	(3)
AI 933	Robot Motion Planning	(3)
AI 945	Big Data	(3)
AI 946	Data Mining	(3)
AI 956	Advanced Artificial Intelligence/ Topics in AI	(3)
AI 957	Advanced Signal Processing	(3)
AI 958	Modelling and Simulation	(3)

Detailed Curriculum (Courses)

1. AI 832: Reinforcement Learning

Textbook

- Reinforcement Learning: An Introduction by Sutton and Barto. Also available at http://incompleteideas.net/book/the-book-2nd.html
- Handouts and research articles may also be used by the instructor.

Objective

To realize the dreams and impact of AI requires autonomous systems that learn to make good decisions. Reinforcement learning is one powerful paradigm for doing so, and it is relevant to an enormous range of tasks, including robotics, game playing, consumer modeling and healthcare. This class will provide a solid introduction to the field of reinforcement learning and students will learn about the core challenges and approaches, including generalization and exploration. Through a combination of lectures, and written and coding assignments, students will become well versed in key ideas and techniques for RL. Assignments will include the basics of reinforcement learning as well as deep reinforcement learning — an extremely promising new area that combines deep learning techniques with reinforcement learning. In addition, students will advance their understanding and the field of RL through a final project.

Pre-Requisite

Linear Algebra, Probability, Machine Learning and Deep Learning

Course Outcome

During the course, students will learn to:

- Define the key features of reinforcement learning that distinguishes it from AI and non-interactive machine learning.
- Given an application problem (e.g. from computer vision, robotics, etc), decide if it should be formulated as a RL problem; if yes be able to define it formally (in terms of the state space, action space, dynamics and reward model), state what algorithm (from class) is best suited for addressing it and justification.
- Implement in code common RL algorithms.
- Describe (list and define) multiple criteria for analyzing RL algorithms and evaluate algorithms on these metrics: e.g. regret, sample complexity, computational complexity, empirical performance, convergence, etc.
- Describe the exploration vs exploitation challenge and compare and contrast at least two approaches for addressing this challenge (in terms of

performance, scalability, complexity of implementation, and theoretical guarantees)

Course Outline

Topics	Allocated Periods
• Introduction to Reinforcement Learning • Markov Processes • Policy Search and Iteration • Value Iteration • Policy Evaluation • Model Free Learning: Q-Learning and SARSA • RL with function approximation • Imitation Learning in Large spaces • Exploration/Exploitation • Batch Reinforcement Learning	45

2. AI 812: Probabilistic Graphical Models

Textbook

- Probabilistic Graphical Models: Principles and Techniques by Daphne Koller and Nir Friedman. MIT Press.
- Handouts and research articles may also be used by the instructor.

Objective

Probabilistic graphical models are a powerful framework for representing complex domains using probability distributions, with numerous applications in machine learning, computer vision, natural language processing and computational biology. Graphical models bring together graph theory and probability theory, and provide a flexible framework for modeling large collections of random variables with complex interactions. This course will provide a comprehensive survey of the topic, introducing the key formalisms and main techniques used to construct them, make predictions, and support decision-making under uncertainty.

Pre-Requisite

Basic probability theory, statistics, programming, algorithm design and analysis.

Course Outcome

During the course, students will learn to develop the knowledge and skills necessary to design, implement and apply these models to solve real problems. Be able to describe directed and undirected graphical models, exact and approximate inference methods and estimation of parameters.

Course Outline

Topics	Allocated Periods
• Introduction, Probability Theory, Bayesian Networks • Undirected models • Learning Bayes Nets • Exact Inference; Message Passing • Sampling • MAP Inference; Structured prediction • Parameter Learning • Bayesian Learning; Structure Learning • Exponential families; variational inference	45

3. CSE 860: Artificial Intelligence (3-0)

Textbook: [AIMA] *Artificial Intelligence: A Modern Approach (2nd Edition)*, by Stuart Russel and Peter Norvig, Prentice Hall, 2002
ISBN-10: 0137903952, ISBN-13: 978-0137903955

Reference Books:

• Artificial Intelligence Structures and Strategies for Complex Problem Solving by George F Lugar, Addison Wesley, 4th Ed 2002

• Artificial Intelligence: A New Synthesis by Nils Nilsson, Morgan Kaufman, 1997.

Objective:

The primary objective of this course is to provide an introduction to the basic principles and applications of Artificial Intelligence. Programming assignments are used to help clarify basic concepts. The emphasis of the course is on teaching the fundamentals, and not on providing a mastery of specific commercially available software tools or programming environments. In short, this is course is about the design and implementation of intelligent agents---software or hardware entities that perform useful tasks with some degree of autonomy.

Course Outcome:

Upon successful completion of the course, students will have an understanding of the basic areas of artificial intelligence including problem-solving, knowledge representation, reasoning, decision making, planning, perception and action, and learning. Students will also be able to

design and implement key components of intelligent agents of moderate complexity in C and/or Lisp or Prolog and evaluate their performance. Graduate students are expected to develop familiarity with current research problems, research methods, and the research literature in AI.

Course Outline:

Artificial Intelligence (AI) seeks to understand the mechanisms underlying thought and intelligent behavior, and their embodiment in machines. This course approaches AI by using Intelligent Agents as an integrating perspective on the key topics in intelligent behavior.

Topics	Allocated Periods
Introduction to AI Definitions (Acting Humanly, Cognitive, "laws of Thought", Rational agent approaches Historical perspective Physical symbol system hypothesis **Intelligent agents** Agents and Environment The concept of rationality Performance measures Omniscience, learning and autonomy Nature of environments, Task environments and their Properties **The Structure of agents** Simple Reflex agent Model based agent Goal Based agents Utility based agents Learning agents **Problem solving by searches** Problem solving agents Problem formulating Measuring performance **Search Strategies** **Uninformed Searches** Breadth first Depth first	45

Depth limited Iterative deepening depth-first Bidirectional Comparison of Uninformed problem solving methods **Informed searches** Greedy best-first A* Heuristic Functions (learning, devising) Local search algorithms Hill climbing Simulating annealing Local beam Genetic algorithms **Constraints Satisfaction problems** Basics of CSP Structure of problems Backtracking, Forward chaining Variable and value ordering Intelligent backtracking Local searches for CSPs **Adversarial Searches** Games Optimal decisions Mini-max algorithm Alpha-Beta pruning **Logical agents** Knowledge-based agents Propositional logic (syntax, semantics) Conjunctive/disjunctive Normal form, Horn clauses Reasoning in Propositional logic Resolution Forward and backward chaining Reasoning algorithms First-Order Logic (syntax, semantics) Inference in FOL Conjunctive Normal form Unification and lifting, forward and backward chaining, resolution Comparison of two representational languages	

Machine Learning Introduction ,induction, Types of machine learning Nearest Neighbors Decision Trees Neural Networks Learning Conjunctions Linear and non Linear saparability Evaluating learning algorithms **Natural language Understanding** Levels of language analysis Types of grammar Parsing(Top-down, bottom-up) natural language applications.	

4. EC 803: Computer Vision (3-0)

Textbook: Computer Vision : A Modern Approach, D. Forsyth and J. Ponce, Prentice-Hall, 2001 (FP). ISBN-10: 0130851981, ISBN-13: 978-0130851987

Reference Book:

- Computer Vision: Algorithms and Applications, Richard Szeliski, 2009 (RS). ISBN-10: 1848829345, ISBN-13: 978-1848829343
- Computer Vision, Linda G. Shapiro and George Stockman, Prentice Hall, (2001).
- Computer Vision and Image Processing, Tim Morris, Palgrave Macmillan, (2004).

Objective:

The objective of this course is to understand the basic issues in computer vision and major approaches that address them. Even though Computer Vision is being used for many practical applications today, it is still not a "solved" problem. Hence, definitive solutions are available only rarely; most of the time, we will discuss alternatives and their limitations.

Course Outcome:

After completing the course, the students may expect to have the knowledge needed to read and understand the more advanced topics and current research literature, and the ability to start working in industry or in academic research. However, this course is not designed to be a "cookbook" course that gives just a survey of the methods needed in "practice", nor will it cover "commercial" systems in detail.

Course Outline:

The primary topics to be covered are chapters 13 through 21 in the textbook. Mostly, the material covered in the text book will be followed but in some cases, the material will be supplemented by newer methods, which are available only in the form of research papers or tutorial articles. Following is a list of topics expected to be covered, in anticipated order, and with expected time to be spent on them. This list is intended to be only indicative, the actual topics, the order and the time may vary somewhat depending on various factors including student interests and preparation.

Topics	Allocated Periods
Image Formation - Sampling - Quantization - Geometry - Photometry - Color **Calibration** **Filtering and Convolution** **Feature Extraction** - Corners, edges, regions **Segmentation and grouping** - Parametric fitting, - Hough transform, - Tensor Voting **Stereo** **Structure and motion** **Dense motion and Optical flow** **Range Image Analysis** **Recognition**	45

5. RIME 833: Deep Learning

Textbook

1. Deep Learning. By Ian Goodfellow, Yoshua Bengio, and Aaron Courville, MIT Press, 2016. Also available at www.deeplearningbook.org
Handouts and research articles may also be used by the instructor.

Objective

2. Computer Vision has become ubiquitous in our society, with applications in search, image understanding, apps, mapping, medicine, drones, and self-driving cars. Core to many of these applications are visual recognition tasks such as image classification, localization and detection. Recent developments in

neural network (aka “deep learning”) approaches have greatly advanced the performance of these state-of-the-art visual recognition systems. This course is a deep dive into details of the deep learning architectures with a focus on learning end-to-end models for these tasks, particularly image classification.

Pre-Requisite

3. Linear Algebra, Probability, Machine Learning

Course Outcome

4. During the course, students will learn to implement, train and debug their own neural networks and gain a detailed understanding of cutting-edge research in computer vision and its application for robotics. The final assignment will involve training a multi-million parameter convolutional neural network and applying it on a large dataset.

Course Outline

4. The course covers the techniques and technology to set up the problem of image recognition, the learning algorithms (e.g. backpropagation), practical engineering tricks for training and fine-tuning the networks.

Topics	**Allocated Periods**
• Machine Learning Basics • Deep Feedforward Networks • Regularization for Deep Learning • Optimization for Training Deep Models • Convolutional Neural Networks • Sequence Modeling: Recurrent Neural Networks • Applications of Deep Learning	45

6. RIME 813: Robotic Grasping and Fixturing

Textbook

1. Fundamentals of Robotic Grasping and Fixturing. Caihua Xiong, Han Ding, and Youlun Xiong, World Scientific Publishing Company, 2007.
ISBN-13 978-981-277-183-4, ISBN-10 981-277-183-2
Handouts and research articles may also be used by the instructor.

Objective

2. This course focuses on providing comprehensive information and mathematic models of developing and applying grippers and fixtures in industry, and present long term valuable essential information for the academic researchers who are interested in robotic manipulation as a good reference.

Pre-Requisite

3. Robot Mechanics and Control

Course Outcome

4. This course will furnish the students with a comprehensive insight into robotic grasping and fixturing. It involves study of multifingered robot hand grasp, basic fixture design principle, and evaluating and planning of robotic grasping/fixturing, and focuses on the modeling and applications of Robotic Grasping and Fixturing.

Course Outline:

Topics	Allocated Periods
Robotic Grasp and Workpiece-Fixture Systems • Introduction • Robotic Manipulation and Multifingered Robotic Hands • AMT and Fixtures • Comparison between Grasping and Fixturing **Qualitative Analysis and Quantitative Evaluation of Form-Closure Grasping/Fixturing** • Kinematic Characteristics of Grasping/Fixturing • Discriminances of Form-Closure Grasping/Fixturing • Minimum Number of Contacts with Frictionless • Grasp Evaluation Criteria **Stability Index and Contact Configuration Planning of Force-Closure Grasping/Fixturing** • Description of Contacts with Friction • Conditions of Force Closure Grasp • Grasp Stability Index **Active Grasp Force Planning** • Nonlinear Programming in Grasp • Force Planning Using Neural Networks **Grasp Capability Analysis** • Evaluation of Multifingered Grasp Capability	48

7. RIME 814: Rehabilitative and Assistive Robotics

Textbook

1. Wearable Robots: Biomechatronic Exoskeletons by Jose Pons (Wiley Publishers)

Handouts and research articles may also be used by the instructor.

Objective

2. A wearable robot is a mechatronic system that is designed around the shape and function of the human body, with segments and joints corresponding to those of the person it is externally coupled with. Tele-operation and power amplification were the first applications, but after recent technological advances the range of application fields has widened. These robots have to support the motion of human limbs and are valuable in the process of rehabilitation as strength augmentation. This course will enable the students to get an insight into the design and control challenges of such robots.

Pre-Requisite

3. Linear Algebra, Robot Mechanics and Control

Course Outcome

4. By the end of this course the students will be well-versed in Biomechanics and Human Interaction. Models of Bio-inspiration along with their actuation requirements will also be introduced to the students and design of systems that enable these motions will be focused.

Course Outline

5. The mechanics of the wearable robot and it's biomechanical interaction with the user, including state-of-the-art technologies will be covered. The basis for bioinspiration and biomimetism, general rules for the development of biologically-inspired designs, and how these could serve recursively as biological models to explain biological systems and finally the incorporation of this knowledge in the design of exoskeletons and prosthetic limbs will be taught.

Topics	**Allocated Periods**
• Bioinspiration and Biomimetism • Biological Models • Human Machine Interfacing • Robot Actuation • Sensing and Instrumentation • Compliance in Design • Design of Rehabilitative Robots • Design Considerations in Robot Assistance	45

8. RIME 835: Human-Robot Interaction

Textbook

1. Designing Sociable Robots (Intelligent Robotics and Autonomous Agents series) by Cynthia Breazeal ISBN: 9780262025102. The course will also be research paper based. Depending upon the module, the instructor(s) may also assign a reference book.

Reference Books

2. • Human-Robot Interaction in Social Robotics by Hiroshi Ishiguro and Takayuki Kanda
   • Human-Robot Interactions: Principles, Technologies and Challenges. Editor: Diana Coleman.

Pre-Requisite

3. CSE 860 Artificial Intelligence (or equivalent), Robotics, Strong programming background.

Objective

4. This is a highly research oriented course with the basic objective of bringing the students face to face with the latest cutting edge research in the world. This will play a key role in helping them select suitable topics for their research that would lead to publications in leading international journals and conferences

Course Outcome

5. Students completing this course are expected to have developed a firm knowledge base of critical understanding of the main theories, principles and concepts relating to the domain of Human-Robot Interaction.

Course Outline

6. This course will cover a variety of topics related to social intelligence and socially intelligent robots. Readings/lectures will cover (1) what's known about how this ability arises in human intelligence, and (2) state-of-the-art approaches to building computational systems with this type of social ability.

Topics	Allocated Periods
• **The Vision of Sociable Robots** • **Robot in Society: A Question of Interface** • **Insights from Developmental Psychology** • **Designing Sociable Robots** • **The Physical Robot** • **The Vision System** • **The Auditory System** • **The Motivation System**	45

• **The Behavior System** • **Facial Animation and Expression** • **Expressive Vocalization System** • **12. Social Constraints on Animate Vision**	

9. RIME 836: Probabilistic Robotics

Textbook

1. Probabilistic Robotics. By Sebastian Thrun, Wolfram Burgard, and Dieter Fox, MIT Press (2006), ISBN: 978-0-262-20162-9.

Handouts and research articles may also be used by the instructor.

Objective

2. This course focuses on robot perception and control in the face of uncertainty. Building on the field of mathematical statistics, probabilistic robotics endows robots with a new level of robustness in real-world situations.

Course Outcome

3. This course will furnish the students with a practical experience in robot perception in partially known environments through implementations in pseudo code, detailed mathematical derivations, discussions from a practitioner's perspective, and extensive lists of exercises and class projects

Course Outline:

Topics	Allocated Periods
Introduction • Uncertainty in Robotics • Probabilistic Robotics **Bayes Filter** • Gaussian Filters • Kalman Filter • Extended Kalman Filter (EKF) **Nonparametric Filters** • Histogram Filter • Particle Filter **Localization** • Markov Localization • EKF Localization • Multi-Hypothesis Tracking	45

• Monte Carlo Localization **<u>Occupancy Grid Mapping</u>** **<u>Simultaneous Localization and Mapping</u>** • EKF SLAM • Graph SLAM	

10. RIME 843: Sensors and Sensing

Textbook

1. Handbook of Modern Sensors. By Jacob Fraden, Springer, 2010. ISBN-13: 978-1-4419-6465-6

Reference Books

2. Introduction to Autonomous Mobile Robots. By Roland Siegwart and Illah R. Nourbakhsh, The MIT Press, 2004.ISBN-10: 0-262-19502-X, ISBN-13: 978-0-262-19502-7

Probabilistic Robotics. By Sebastian Thrun, Wolfram Burgard, and Dieter Fox, MIT Press (2006)
ISBN: 978-0-262-20162-9.

Objective

3. This course focuses on various sensors used for Robot Navigation and Control. Its objective is to teach students about sensing and modeling surrounding environment for motion planning and navigation.

Course Outcome

4. This course will furnish the students with an understanding of the design and working of sensors for Robotics while also teaching them about the techniques needed for signal processing of the sensor data.

Course Outline:

Topics	Allocated Periods

• Sensor Characteristics
  • Calibration
  • Accuracy
  • Repeatability
• Data Acquisition from Sensors
• Odometers
• Heading Sensors
• Accelerometer
• Inertial Measurement Unit
• Vision Sensing
• Range Sensing
• Sonar Sensing
• Flow Sensing
• Touch Sensing
• Sensing for Surveillance
• People Sensing
• Multi-Sensor Fusion
	45

11. RIME 837: Simultaneous Localization and Mapping

Textbook

1. Probabilistic Robotics. By Sebastian Thrun, Wolfram Burgard, and Dieter Fox, MIT Press (2006)
ISBN: 978-0-262-20162-9.

Reference Book

2. Introduction to Autonomous Mobile Robots. By Roland Siegwart and Illah R. Nourbakhsh, The MIT Press, 2004. ISBN-10: 0-262-19502-X, ISBN-13: 978-0-262-19502-7.
Handouts and research articles may also be used by the instructor.

Objective

3. This course focuses on Robot Localization and Mapping in unknown or partially known environments.

Course Outcome

4. This course will furnish the students with practical knowledge of SLAM algorithms required for mapping and navigation in unknown of partially known environments.

Course Outline

Topics	Allocated Periods
Introduction to the Simultaneous Localization and Mapping (SLAM) Problem **SLAM with Extended Kalman Filter** • SLAM with Known Correspondences • SLAM with Unknown Correspondences **SLAM with Particle Filter** **Graph-based SLAM** **Biologically-inspired SLAM Solutions** • RatSLAM **Loop-Closure in SLAM** **Hierarchical SLAM** **SLAM Using Vision** • Monocular SLAM • Stereo and Multi-camera SLAM • SLAM using Catadioptric Sensors **Underwater SLAM** **SLAM for UAVs (Unmanned Aerial Vehicles)**	45

12. RIME 817: BioRobotics

Textbook

1. Biorobotics–Methods and Applications, edited by Barbara Webb and Thomas R. Consi.

Objective

3. The objective of this course is to develop expertise of the multidisciplinary field of BioRobotics. The course involves detailed study of Biosignals and the principals of Bio Mechanical Designs.

Course Outcome

4. After studying this course students will be able to use Human Neurological Signals for operations of various kinds of Robots. The course will also enable students to develop Biomatic Robotic systems.

Course Outline

Topics	Allocated Periods
• Introduction to biomechatronics and biorobotics. • Types of biosensors and instruments. • Modeling and design of prosthetic devices. • Principle of electromyography (EMG). • Design of prosthetic devices. • Control of prosthetic devices using EMG – preprocessing and pattern recognition. • Control of prosthetic devices using EMG – control command generation. • Bio-inspired machines. • Robotic rehabilitation. • Brain-controlled robotics. • Functional near-infrared spectroscopy based brain-computer interfaces.	45

13. **RIME 832: Machine Learning (3-0)**

Textbook: Introduction to Machine Learning, by Ethem Alpaydin, MIT Press, 2004.
ISBN-10: 0-262-01211-1, ISBN-13: 978-0-262-01211-9

Reference Book: Machine Learning, by Tom M. Mitchell, McGraw-Hill Science / Engineering / Math; 1 edition (March 1, 1997).
ISBN-10: 0070428077, ISBN-13: 978-0070428072

Objective:

When you have completed this course, you should be able to apply machine learning algorithms to solve both IID and sequential data problems of moderate complexity. You should also be able to read current research papers in machine learning and understand the issues raised by current research in supervised learning.

Pre-Requisite:

CSE 860 Artificial Intelligence (or equivalent)

Course Outcome:

The students graduating from this course are expected to have a thorough knowledge base of Machine Learning algorithms and methodologies as well as their applications in robotics and machine intelligence.

Course Outline:

This course will present an introduction to algorithms for machine learning and data mining. These algorithms lie at the heart of many leading edge computer applications including optical character recognition, speech recognition, text mining, document classification, pattern recognition, computer intrusion detection, and information extraction from web pages. Every machine learning algorithm has both a computational aspect (how to compute the answer) and a statistical aspect (how to ensure that future predictions are accurate). Algorithms covered include linear classifiers (Gaussian maximum likelihood, Naive Bayes, and logistic regression) and non-linear classifiers (neural networks, decision trees, support-vector machines, nearest neighbor methods). The class will also be introduced to techniques for learning from sequential data and advanced ensemble methods such as bagging and boosting.

Topics	Allocated Periods
Overview/Introduction to machine learning Hypothesis spaces Space of Algorithms, Linear Threshold Classifiers Project details Perceptrons Logistic Regression	45

Linear Discriminant Analysis Off-The-Shelf Learning Algorithms Decision Trees Nearest Neighbor Neural networks Bayesian Learning Support Vector Machines Learning Theory Learning Theory finished Bias/Variance Theory & Ensemble Methods Preventing Over-fitting: Penalty and Hold-out methods Hold-Out and Cross-validation Methods Penalty methods: decision trees, neural nets, SVMs Evaluating and Comparing Classifiers Unsupervised Learning	

14. <u>EM 840: Data Acquisition and Control (3-0)</u>

Text Book:

- Transducers and Instrumentation, D.V.S. Mury, Prentice-Hall International, Inc.
- Digital Control System Analysis And Design by C. L. Phillips and H. T. Nagle, Jr., 3rd edition, Prentice-Hall International Inc.

Reference Books:

- Electronics with Digital and Analog Integrated Circuits by Richard J. Higgins Prentice-Hall International, Inc
- TMS320C3x DSP Starter Kit User's Guide by Texas Instruments

Objective:

The objective of this course is to impart theoretical and practical knowledge of advanced data acquisition and control to graduate students.

Course Outcome:

Students after successfully completing of this course will be able to demonstrate:

i. Understanding of all the components required for data acquisition systems.

ii. Knowledge of analyzing the system speed, resolution and accuracy.

iii. Knowledge of designing an efficient data acquisition system for the required applications

iv. Knowledge to use controller design techniques to make the system behavior satisfy specified design objectives

v. Ability to evaluate and test the system performance using digital simulations.

Course Outline:

Topics	Allocated Periods
1. Introduction to Data acquisition 2. Passive and Active Electrical Transducers 3. Signal Conditioning Circuits 4. Digital Interfacing 5. Data Communication and Networks 6. ADC and DAC, Timers And Counters 7. Digital measurements and control programming for real time systems 8. Introduction to Digital Control Systems 9. Digital Controller Design	45

15. RIME 914: Robot Motion Planning (3-0)

Textbook: Planning Algorithms by Steven M. LaValle, Cambridge University Press, 2006.
ISBN-10: 0521862051, ISBN-13: 978-0521862059

Reference Book: Robot Motion Planning, By Jean-Claude Latombe, Kluwer Academic
Publishers. ISBN-10: 079239206X, ISBN-13: 978-0792392064

Objective:

The aim of this course is to teach the students about advanced techniques used for robot motion planning. The course combines the knowledge of robotics with that of artificial intelligence and control theory to give the students a practical overview of the cutting edge methods used in the area of planning algorithms.

Pre-Requisite:

EM 800 Robotics – I (or equivalent)

Course Outcome:

Students completing this course are expected to possess a firm grasp of robot motion planning algorithms.

Course Outline:

The course can broadly be outlined as motion planning, decision theoretic planning and planning under differential constraints.

Topics	Allocated Periods
Discrete Planning Logic Based Planning Methods	45

Configuration Space Sampling Based Motion Planning Combinatorial Motion Planning Time Varying Problems Mixing Discrete and Continuous Spaces Planning for Closed Kinematic Chains Feedback Motion Planning Basic Decision Theory Sequential Decision Theory Sensors and Information Spaces Planning under Sensory Uncertainty Differential Models Sampling Based Planning under Differential Constraints System Theory and Analytical Techniques	

16. <u>RIME 933: Advanced AI/Special Topics in Artificial Intelligence (3-0)</u>

(Any subject on the research interest of the faculty available)

Textbook: The course will mainly be research paper based. Depending upon the module, the instructor(s) may also assign a reference book.

Objective:

This is a highly research oriented course with the basic objective of bringing the students face to face with the latest cutting edge research in the world. This will play a key role in helping them select suitable topics for their research that would lead to publications in leading international journals and conferences.

Pre-Requisite:

CSE 860 Artificial Intelligence (or equivalent)

Course Outcome:

Students graduating from this course are expected to develop a deep insight into cutting edge artificial intelligence technologies with an overview of the underlying theory and methodologies in each subject area. This course will be very helpful for students to take up research theses of good technical value in the cutting edge research areas.

Course Outline:

The course will comprise several modules, each of which would cover one of the cutting edge research topics. Depending upon the interest of the class and the ongoing research projects in the department, some of the following research topics will be taught about. The instructors may add or amend research subjects depending upon the orientation of the class.

Topics	Allocated Periods
Brain simulation Nouvelle AI Embodied AI Situated AI Behavior based AI Embodied Embedded Cognition Embodied cognitive science Computational intelligence (CI) Evolutionary algorithms Swarm intelligence Artificial Immune Systems Immunocomputing Speech Processing and Synthesis	45

17. AI 813: Multi-agent Systems

Textbook

- *Shoham, Y. & Leyton-Brown, K. (2009). Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.* The book is available at http://www.masfoundations.org/
- Handouts and research articles may also be used by the instructor.

Objective

Multiagent systems (MAS) can be defined as loosely coupled networks of problem solvers that interact to solve problems that are beyond the individual capabilities or knowledge of each problem solver. These problem solvers, often called agents, are autonomous and can be heterogeneous in nature.

Research and development in MAS is concerned with the study and construction of a collection of autonomous agents that interact with each other and their environments. The study of such systems goes beyond the study of individual intelligence in its consideration of problem solving with social components.

This course introduces the students to the main topics in the theory and practice of MAS, currently one of the most important and rapidly expanding areas of computer science, having emerged from the study of distributed artificial intelligence (DAI). Multiagent systems have been used as an important means

with which to address the development of large and complex information systems (IS) and decision support systems (DSS). Finally, this course introduces coalitional game theory and its potential applications.

Pre-Requisite

The student should be comfortable with mathematical notation and basic computer algorithms

Course Outcome

During the course, students will learn to:

- Demonstrate/explain the basic concepts of agent-based approach, non-cooperative game theory, multiagent learning, social choice, mechanism design, auctions, and cooperative game theory;
- Develop business and real-world perspectives of multiagent systems;
- Use software tools to develop and test multiagent systems;
- Be aware of future and current trends in MAS research and applications.

Course Outline

Topics	**Allocated Periods**
• Foundations • Intelligent Agents and Multiagent Systems • Multiagent Learning • Social Choice • Mechanism Design • Multiagent Resource Allocation • Coalition Game	45

18. AI 803: Mathematics for Artificial Intelligence

It's a rather unique course offered as core in our comparative cohort. Its need has been felt in the local context where students applying for admission at MS and PhD level have at times relatively weaker mathematical and statistical foundations. Thus, we consider it an important hole filling exercise for their skills. We expect that the usual rantings from the faculty regarding mathematical inadequacy of student cohorts would go down with the introduction of this course at the grad level.

Recommended Books:

1. Linear Algebra and Learning from Data, Gilbert Strang, Wellesley-Cambridge Press, 2019.
2. Math for Machine Learning: Open Doors to Data Science and Artificial Intelligence, Richard Han, CreateSpace Independent Publishing Platform, 2018.
3. Machine Learning: A Probabilistic Perspective, Kevin Murphy, The MIT Press, 2012.

Credit Hours: 3 (3, 0)

Course Objectives:

- To understand the basic mathematics, such as linear algebra and calculus, that is behind many aspects of vision, text processing, and machine learning.
- To understand the basic statistics that can be used in image processing and learning.

Topics / Contents	**Allocated Periods**
The course would cover basic aspects of linear algebra such as scalars, vectors, matrices, and tensors. It would then move on to different matrix operations and discuss eigen values and eigen vectors. It would also discuss matrix decomposition, such as SVD and eigen decomposition. Multivariable calculus would start with the basic concepts of derivates (both first and second order), integration for both discrete and continuous variables, gradient descent and optimization. Probability and statistics would cover topics such as random variables, expected value, variance, different probability distributions, and central limit theorem.	**45**

19. AI 917: Evolutionary Algorithms

Since the publication of Darwin's magnum opus, The Origin of Species in 1859, and the ensuing coining of the term the survival of the fittest, the word evolution has pervaded different disciplines and AI is no exception. In fact it has been observed that certain kind of learning cannot take place within one generation and can only be done over many generations by evolving better and better individuals through the genetic processes. When the computer scientists got cognizant of the concept, a whole new field of evolutionary computation arose, and now many problems in CS are only seen in this context. Among our comparative cohort, only NUS Singapore does not offer it.

Recommended Books:

1. Evolutionary Algorithms, Alain Petrowski, Sana Ben-Hamida, John Wiley & Sons, Inc., 2017.
2. Evolutionary Optimization Algorithms, Dan Simon, John Wiley & Sons, Inc., 2013.
3. Evolutionary Algorithms, Swarm Dynamics and Complex Networks: Methodology Perspectives and Implementation, Ivan Zelinka, Guanrong Chen (Editors), Springer, 2018.

Credit Hours: 3 (3, 0)

Course Objectives:

- To formulate and represent a problem as an evolutionary computation problem.
- To implement and assess different evolutionary algorithms.
- To formulate single and multiple objectives functions.

Topics / Contents	**Allocated Periods**
The topics include evolutionary computation, genetic algorithms, evolution strategies, evolutionary programming, genetic programming, constraint handling, evolutionary neural networks, co-evolution, and swarm intelligence. Other topics include genotypes and phenotypes, schema theorem, convergence velocity perspective, multi-objective functions, self-adaptation in genetic algorithms, and their optimization.	**45**

20. RIME 942: Pattern Recognition

Its another name for machine learning but offers a different perspective. The term has its origins in EE as compared to machine learning which has its origin in CS. It has a more mathematical icing with the underpinning mathematical concepts and terms discussed in detail to bring students up-to-date on this crucial aspect. In our comparative cohort only University of Amsterdam does not offer this course.

Recommended Books:

1. Pattern Recognition and Machine Learning, revised edition, Christopher M. Bishop, Springer, 2016.
2. Pattern Recognition and Classification: An Introduction, Geoff Dougherty, Springer, 2013.

Pattern Classification, 2nd edition, Richard O. Duda, Peter E. Hart, David G. Stork, John Wiley & Sons, Inc., 2001

Credit Hours: 3 (3, 0)

Course Objectives:

- To understand what a pattern is and the basics of what intelligence and learning are.
- To understand feature extraction, their evaluation, and selection.
- To use supervised and unsupervised approaches for identification of patterns in the ata.

Topics / Contents	**Allocated Periods**
The topics include Bayesian decision theory, Bayesian networks, maximum likelihood estimation, dimensionality reduction, feature selection, Bayesian estimation, linear discriminant functions, Support vector machines, expectation maximization algorithm, and non-parametric estimation. Other topics include string matching, bias and variance, bagging and boosting, and unsupervised learning and clustering.	**45**

21. AI 814: Knowledge representation and reasoning

Many modern systems where NLP is used for exploring patterns in the data, or in the cases where multimedia including text is considered, it becomes important to see relationships between different objects and patterns. Thus, ontologies need to be built for reasoning and knowledge representation becomes the key to it. Only Utrecht offers nothing akin to it. The rest all include it in one form or the other.

Recommended Books:

1. Knowledge Representation and Reasoning, 3rd edition, Gerardus Blokdyk, Emereo Pty Limited, 2018.
2. Knowledge Representation and Reasoning with Deep Neural Networks, Arvind Ramanathan Neelakantan, University of Massachusetts Libraries, 2017.
3. Knowledge Representation and Reasoning, Ronald Brachman, Hector Levesque, Elsevier, 2004.

Credit Hours: 3 (3, 0)
Course Objectives:

- To understand what representation is for knowledge.
- To use logic (both propositional and first order) for representation.
- To understand how reasoning can be done in logic, and what entailment is.
- To understand the theory and principles of automated theorem proving.
- To get acquainted to tools that can be used for knowledge representation and reasoning.

Topics / Contents	**Allocated Periods**
Topics include propositional logic, its syntax and semantics. Proof by resolution and automated theorem proving. First-order logic, its syntax and semantics. Forward and backward chaining. Semantic nets. Inheritance in taxonomies. Other topics include other approaches and issues (e.g. predicate logic, fuzzy logic, week and strong slot and filler structures), knowledge acquisition, the frame problem, symbolic reasoning under uncertainty (nonmonotonic reasoning, augmenting a problem Solver), statistical reasoning (e.g. probability and Bays Theorem, Bayesian networks, Dumpster-Shafer theory), building knowledge-based systems.	**45**

22. AI 844: Semantic Web

Only NUS Singapore does not offer any course even remotely related to it within our comparison cohort. The rest all offer it in some context. Semantics is derived from the context, and anywhere where context is being used to identify a word, it falls under semantics. Within the ambit of semantic web it includes information in various forms, such as images and pictures, structured and unstructured information, and videos. Thus, the range of applications are huge.

Recommended Books:

1. Semantic Web Science and Real-World Applications, Lytras, Miltiadis D., Aljohani, Naif, Damiani, Ernesto, IGI Global, 2019.
2. Semantic Web-Based Systems: Quality Assessment Models, Sandeep Kumar, Niyati Baliyan, Springer 2018.
3. Semantic Web: Ontology and Knowledge Base Enabled Tools, Services, and Applications, Amit Sheth, Information Science Reference, 2013.

Credit Hours: 3 (3, 0)
Course Objectives:

- To understand the technical architecture of the semantic web and how it integrates with the traditional World Wide Web.
- To understand the underlying representations of information and do inference on them.
- To be able to use common semantic web tools to design, implement, and verify ontologies.

Topics / Contents	**Allocated Periods**
The topics include RDF and linked data, ontologies, the RDF query language SPARQL, RDF triplestores, understanding datasets, basic OWL modeling, creating an ontology from given requirements, Ontology Design Patterns (ODPs), ODP based modeling, description logics, ontology-alignment and debugging, constraints in RDF, data quality, and data cleaning in RDF.	**45**

23. EM 890: Modeling and Simulation

Mathematical modeling and its simulation is the key to understanding where the system will be successful and where it will fail to perform as per the requirement. Not every scenario can be generated in the lab environment and tested for its efficacy. In most of the cases it is much cheaper to first simulate it and only prototype it once it has been deemed to be workable. That's where modeling and simulation come into the picture. Its also a unique offering within the comparison cohort and we believe it would enhance our success rate dramatically.

Recommended Books:

1. Theory of Modeling and Simulation: Discrete Event & Iterative System Computational Foundations, Bernard P. Zeigler, Alexandre Muzy, Ernesto Kofman, Elsevier Science, 2018.
2. Modeling and Simulation: An Application-Oriented Introduction, Hans-Joachim Bungartz, Stefan Zimmer, Martin Buchholz, Springer, 2014.
3. Modeling and Simulation of Complex Systems: A Framework for Efficient Agent-Based Modeling and Simulation, Robert Siegfried, Springer Vieweg, 2014.

Credit Hours: 3 (3, 0)

Course Objectives:

- To understand what problems may be most suitable to be modeled and simulated as opposed to being physically implemented.
- To understand parameters and constraints that may define the scope of the problem to be modeled and simulated.
- To be able to demonstrate coding and implementation skills in the available tools to simulate some real-world problem, which is not easily expressed using mathematics.

Topics / Contents	**Allocated Periods**
Course contents include topics on system analysis, classification of systems, and system theory basics and its relevance to simulation. For modeling the course includes model classification which refers to conceptualization, abstraction, and simulation of the models. It also includes topics related to simulation systems and languages. It also covers topics such as Petri nets and finite automata and their use in representing models. Both discrete and continuous systems are covered. For continuous systems, numerical methods are included. For testing, validity and verification of models is discussed. Optimization is also made part of it so that the developed systems are optimized as per the set criteria.	**45**

24. AI 822: Advanced Image Processing

Image processing underpins computer vision. No mentionable work can be done in vision without proper grounding in image processing. Even students who have already taken this course at the under-grad level are not familiar with the techniques that fall in the advanced category. It would be useful for both first time and experienced students. It is not offered by any of the universities we have compared our program with but we feel that its inclusion is essential for us.

Recommended Books:

1. Advanced Image and Video Processing Using Matlab, Shengrong Gong, Chunping Liu, Yi Ji, Springer, 2019.
2. Hands-On Image Processing with Python: Expert techniques for advanced image analysis and effective interpretation of image data, Sandipan Dey, Packt Publishing, 2018.
3. Principles of Digital Image Processing: Advanced Methods, Wilhelm Burger, Mark J. Burge, Springer, 2013.

Credit Hours: 3 (3, 0)
Course Objectives:

- To introduce students to advanced theoretical concepts.
- To develop in students the problem solving skills and build in them engineering intuition.
- To inculcate practical skills and know how to be able to build working computer vision based systems.

Topics / Contents	**Allocated Periods**
The topics include color, camera models and calibration, scale in image processing, line and corner detection, and maximally stable extremal regions. It also includes topics related to mathematical morphology such as binary, gray-scale, skeletonization, granulometry, and morphological segmentation. It would also cover image compression, texture, and image registration including rigid, non-rigid, and RANSAC.	**45**

25. AI 843: Natural Language Processing

NLP is one of the main tasks where AI is the key underlying technology. Processing text in main cases is the main problem to be solved specially when data is available in that form, and also when robots have to communicate with humans they may resort to it. Also, its coupling with the speech processing unit would give extra leverage to the robots to do their tasks in a more effective manner. All the universities in the comparison cohort offer this course.

Recommended Books:

1. Natural Language Processing Recipes: Unlocking Text Data with Machine Learning and Deep Learning using Python, Akshay Kulkarni, Adarsha Shivananda, Apress, 2019.
2. Deep Learning in Natural Language Processing, Li Deng, Yang Liu editors, Springer 2018.
3. Python Natural Language Processing, Jalaj Thanaki, Packt Publishing, 2017.

Credit Hours: 3 (3, 0)
Course Objectives:

- To understand vector space models and use of context to derive semantics.

- To understand modern deep learning based algorithms that are used for NLP.
- To be able to process large unstructured and unlabeled text to extract useful information.

Topics / Contents	**Allocated Periods**
The course covers topics such as word vectors, word senses, neural networks and matrix calculus, backpropagation and computational graphs, linguistic structure dependency and parsing, recurrent neural networks and language models, vanishing gradients and fancy RNNs, machine translation, question answering, ConvNets for NLP, subword models, generative models, safety, bias, and fairness. Other topics include words, morphology, and lexicon, noisy channel models and edit distance, classification, parts of speech tags, hidden markov models, Chomsky hierarchy and natural language, context-free recognition, lexical semantics, word sense disambiguation, and interpreting social media.	**45**

26. AI 831/937: Intelligent Systems

Georgia and Amsterdam do not offer it but it may have to do with their peculiar circumstances. It relates to planning, execution, and learning and is essential for us due to the nature of robotics that we are going to do where all the three aspects above would be covered.

Recommended Books:

1. Learning and Execution of Object Manipulation Tasks on Humanoid Robots, Waechter, Mirko, Scientific Publishing, 2018.
2. Machine Learning Methods for Planning, Steven Minton, Elsevier Science, 2014.
3. Using Neural Networks and Dyna Algorithm for Integrated Planning, Reacting and Learning in Systems, National Aeronautics and Space Adm Nasa, Independently Published, 2018.

Credit Hours: 3 (3, 0)
Course Objectives:

- To comprehend the challenges in developing autonomous systems or agents that plan, execute, and learn from the consequences of their actions.

- To understand what planning is, how it is represented, how it is evaluated, and how it is learned.
- To implement a simple complete system which shows autonomous behavior.

Topics / Contents	**Allocated Periods**
The topics include closed-world assumption, linear programming, GPS algorithm, prodigy algorithm, S-Space transformation, constructing graphs of the planning problem, Planning graphs for planning and heuristic search, A* and weighted A* search, multi-heuristic A*, probabilistic path planning, planning under uncertainty, reinforcement learning, experience graphs, learning in planning, deep reinforcement learning, learning cost functions, multi-agent decision making, and short-sighted probabilistic planning.	**45**

27. AI 851 Advanced Signal Processing(3, 0)

Pre-requisite: None

Recommended Books:

1. Advanced Digital Signal Processing and Noise Reduction, Saeed V. Vaseghi, John Wiley & Sons, Limited, 2020.
2. Digital Signal Processing using Matlab, 3rd edition, Robert J. Schilling, Sandra L Harris, Cencage Learning, 2017.
3. Bayesian Signal Processing, 2nd edition, James V. Candy, John Wiley & Sons, Inc., 2016.

Credit Hours: 3 (3, 0)

Course Objectives:

- To be able to use Fast Fourier Transform in a variety of applications such as signal analysis, fast convolution, spectral and temporal interpolation, and filtering.
- To understand how to design digital filters.
- To be able to construct a simple digital communication system.

Topics / Contents	**Allocated Periods**
The topics include digital systems and signals, impulse response, frequency response, discrete-time Fourier transform, z-transform, interpolation and decimation, sample rate changes, polyphase decomposition, random	

signals and vectors, joint random signals and their statistical description, Stationarity and ergodicity, Linear systems with stationary random inputs, Power spectral density, Metric spaces, vector spaces, normed vector spaces, inner projects and inner-product spaces, Hilbert spaces, Projections, Orthogonality, Orthogonalization of vectors, Approximation in Hilbert spaces, The orthogonality principle, Least squares filtering, Linear prediction and AR spectrum estimation, MMSE estimation and MMSE filtering, Weiner filtering, Signal transformation and generalized Fourier series, Scaling functions and wavelets, Wavelet transform and its implementation, Eigenvalues and linear systems/operators, Diagnolization of a matrix, Karhunen-Loeve expansion, Principal component analysis, Pisarenko harmonic decomposition, MUSIC, SVD and applications, Applications of adaptive filtering, LMS algorithm, RLS algorithm.	**45**

28. AI 841 Data Mining (3, 0)

Pre-requisite: None

Recommended Books:

1. Introduction to Data Mining, Pang-Ning Tan, Michael Steinbach, Anuj Karpatne and Vipin Kumar, 2nd Edition, Pearson, 2018.
2. Principles of Data Mining, 3rd edition, Max Bramer, Springer, 2016.
3. Mining of Massive Datasets, 2nd edition, Jure Leskovec, Anand Rajaraman, Jeffrey David Ullman, Cambridge University Press, 2014

Credit Hours: 3 (3, 0)

Course Objectives:

- To determine whether a particular problem is a data mining problem or not.
- To understand the complete lifecycle of a datamining process such as data preparation, modeling, and evaluation.
- To pose a problem as a data mining problem, implement and evaluate it.

Topics / Contents	**Allocated Periods**
This course provides both theoretical and practical coverage of all data mining topics. The topics include: Overview of Data Mining, Data Preprocessing, OLAP and data generalization, Data Cube Computation and Multidimensional Data Analysis, Mining Frequent Patterns,	**45**

Associations, and Correlations, Classification, Cluster Analysis, Outlier Detection, Anomaly Detection; Avoiding False Discoveries.	

29. AI 842 Information Retrieval (3, 0)

Pre-requisite: None

Recommended Books:

1. Learning to Rank for Information Retrieval, Tie-Yan Liu, Springer Berlin Heidelberg, 2014.
2. Multilingual Information Retrieval: From Research To Practice, Carol Peters, Martin Braschler, Paul Clough, Springer, 2012.
3. Modern Information Retrieval. Baeza-Yates Ricardo and Berthier Ribeiro-Neto. 2nd edition, Addison-Wesley, 2011.

Credit Hours: 3 (3, 0)

Course Objectives:

- Introduce the concepts underlying technologies of modern information retrieval systems.
- To study theoretical aspects as well as implementation issues of classical and modern retrieval problems such as search engines.
- Basic and advanced techniques for building text-based information systems

Topics / Contents	**Allocated Periods**
This course provides a variety of basic principles, techniques and modern advances for searching, managing, and mining information. It will cover algorithms, design, and implementation of modern information retrieval systems. The topics include: Search engine architecture, Retrieval models, Retrieval evaluation, Relevance feedback, Link analysis, Search applications, Retrieval system design and implementation, text analysis techniques, retrieval models (e.g., Boolean, vector space, probabilistic, and learning-based methods), search evaluation, retrieval feedback, search log mining, and applications in web information management. IR techniques for the web, including crawling, link-based algorithms, and metadata usage are also discussed	**45**

30. AI 815 Neural Networks (3, 0)

Pre-requisite: None

Recommended Books:

1. Make Your Own Neural Network: A Gentle Journey Through the Mathematics of Neural Networks, and Making Your Own Using the Python Computer Language, Tariq Rashid, CreateSpace Independent Publishing Platform, 2016.
2. Neural network design, 2nd edition, Martin T. Hagan, Howard B. Demuth, Mark Hudson Beale, Martin Hagan, 2014.
3. Neural Networks with R: Smart models using CNN, RNN, deep learning, and artificial intelligence principles, Giuseppe Ciaburro, Balaji Venkateswaran, Packt Publishing, 2017.

Credit Hours: 3 (3, 0)
Course Objectives:

- To introduce the neural networks for regression and classification.
- To give different approaches that can be used to design a neural network based solution.
- To understand how weights are learned and how over-fitting is avoided.
- To demonstrate implementation of neural networks on available tools.
- To demonstrate conception, designing, and solution of real world problems using neural networks.

Topics / Contents	**Allocated Periods**
The topics include the neuron models and basic learning rules, the learning of a single neuron and a single-layered neuron, multilayer neural networks and backpropagation, associative memory, self-organizing feature map, RBF neural networks and support vector machines, and neural network trees.	**45**

31. AI 853 Advanced Programming in Python (3, 0)
Pre-requisite: None
Recommended Books:

1. Advanced Python Programming: The Insider Guide to Advanced Python Programming Systems, Richard Ozer, Python Programming, CreateSpace Independent Publishing Platform, 2017.
2. Introduction to Machine Learning with Python: A Guide for Data Scientists 1st Edition, Andreas C. Müller, Sarah Guido, 2016.
3. Fluent Python: Clear, Concise, and Effective Programming 1st Edition, Luciano Ramalho, 2015

Credit Hours: 3
Course Objectives:

- To be able to add enhancements to classes.

- To implement and run unit tests.
- To create multi-threaded and multi-process applications.

Topics / Contents	Allocated Periods
The topics include Advanced List Comprehensions, Collections Module, Mapping and Filtering, Lambda Functions, Advanced Sorting, Unpacking Sequences in Function, Calls, Modules and Packages, Working with Data, Databases, CSV, Getting Data from the Web, HTML, XML, JSON, Testing and Debugging, Creating Simulations, Testing for Performance, The unittest Module, Classes and Objects, Creating Classes, Attributes, Methods and Properties, Extending Classes, Documenting Classes, Static, Class, Abstract Methods, Decorator.	**45**

32. AI 854 Analysis and Visualization (3, 0)

Pre-requisite: None

Recommended Books:

1. Data Analysis and Visualization Using Python: Analyze Data to Create Visualizations for BI Systems, Dr. Ossama Embarak, Apress, 2018.
2. Python: Data Analytics and Visualization, Phuong Vo.T.H, Martin Czygan, Ashish Kumar, Packt Publishing, 2017.
3. R: Data Analysis and Visualization, Tony Fischetti, Brett Lantz, Jaynal Abedin, Hrishi V. Mittal, Bater Makhabel, Edina Berlinger, Ferenc Illes, Milan Badics, Adam Banai, Gergely Daroczi, Packt Publishing, 2016.

Credit Hours: 3 (3, 0)

Course Objectives:

- To be able to design and critique different visualizations.
- To appreciate the importance of data analytics for visualization.
- To understand the components involved in visualization and also how different data types impact visualization.

Topics / Contents	Allocated Periods
The course covers topics such as data processing and big data, data understanding, data preparation, modeling, evaluation, and analytic techniques. Regarding data the topics included are importance of context, audience's attention, and dissecting model visuals. Predictive modeling includes topics such as linear regression, classification, and clustering. Topics covered related to modeling include	**45**

overfitting, generalization, plain accuracy and its problems, confusion matrix, unbalanced classes, and unequal costs and benefits. Topics related to visualizing model performance include ranking, ROC graphs and curves, and AUC. Other topics include visual encoding, bar chart and pie chart, line chart, highlight tables, scatter plot and trend lines, heatmap, geographic mapping, bullet graph, Gantt chart, and circle view. It also includes table calculation, forecast, log distribution, reference line and table.	

33. AI 823 Speech Processing (3, 0)
Pre-requisite: None
Recommended Books:

1. Audio and Speech Processing with MATLAB, Paul Hill, Taylor & Francis Incorporated, 2018.
2. Language and Speech Processing, Joseph Mariani (Editor), John Wiley & Sons, Inc., 2013.
3. Multilingual Speech Processing, Tanja Schultz, Katrin Kirchhoff, Elsevier, 2006.

Credit Hours: 3
Course Objectives:

- To understand the basic principles of speech perception.
- To understand basic principles of speech recognition, synthesis, and dialogue.
- To be able to generate speech datasets and build applications to recognize speeches and speakers.

Topics / Contents	**Allocated Periods**
The topics include physical and physiological acoustics, perception of speech, phonetics and phonology, signal processing, speech synthesis, prosody and emotions, hidden markov models, language modelling, human-computer communication, and dialogue systems. Other topics include linear-predictive model, cepstral analysis, speech coding, and parameterization of DTW.	**45**

34. AI 945 Big Data Analysis (3, 0)
Pre-requisite: None
Recommended Books:

1. Big Data Analytics: Methods and Applications, Saumyadipta Pyne, B.L.S. Prakasa Rao, S.B. Rao, Springer, 2016.
2. Big Data Science & Analytics: A Hands-On Approach, Arshdeep Bahga, Vijay Madisetti, VPT, 2016.
3. Data Science and Big Data Analytics: Discovering, Analyzing, Visualizing and Presenting Data, EMC Education Services, John Wiley & Sons, 2015.

Credit Hours: 3 (3, 0)
Course Objectives:

- To introduce the concepts and algorithms used for Big Data analytics;
- To introduce the principles of the Hadoop, MapReduce and HBase.
- To enable the students to implement big data technologies

Topics / Contents	**Allocated Periods**
Topics include: overview of Big data, Using Big Data in Businesses such as big data in marketing, analytics, retail, hospitality, consumer good, defense etc., Technologies for Handling Big Data such as Hadoop, functioning of Hadoop, Cloud computing (features, advantages, applications) etc, Understanding Hadoop Ecosystem includes HDFS, MapReduce, YARN, HBase, Hive, Pig, Sqoop, Zookeeper, Flume, Oozie etc., Dig Deep to understand the fundamental of MapReduce and HBase, Understanding Big Data Technology Foundations that is data source layer, ingestion layer, source layer, security layer, visualization layer, visualization approaches etc. Databases and Data Warehouses, Using Hadoop to store data, Learn to Process Data using Map Reduce, Testing and Debugging Map Reduce Applications, Learn Hadoop YARN Architecture, Exploring Hive, Exploring Pig, Exploring Oozie, Learn NoSQL Data Management, Integrating R and Hadoop and Understanding Hive in Detail.	**45**

35. AI 855 Cyber Security (3, 0)
Pre-requisite: None
Recommended Books:

1. Computer and Cyber Security: Principles, Algorithm, Applications, and Perspectives, Brij Gupta, Haoxiang Wang, Taylor & Francis Incorporated, 2018.

2. Cybersecurity for Beginners, Raef Meeuwisse, Cyber Simplicity Limited, 2017.
3. Smart Cities Cybersecurity and Privacy, Danda B. Rawat, Kayhan Zrar Ghafoor, Elsevier Science, 2018.

Credit Hours: 3 (3, 0)
Course Objectives:

- To be able to understand different kinds of cyber security threats and tell how to counter them.
- To understand the mechanism of protecting confidentiality and completeness of data.
- to be able to evaluate the outcome of any counter measure to mitigate threat.

Topics / Contents	**Allocated Periods**
The course starts with cyber security fundamentals discussing in detail cyberspace, cyber security, and what a hacker is. It goes on to discuss different types of malware such as worms, viruses, spyware, and trojans. For cyber security breaches it discusses phishing, identity theft, harassment, and cyberstalking. Among types of cyber attacks it will discuss password attacks, denial of service attacks, passive attacks, and penetration testing. Among prevention tips it would cover two-step verification and legitimacy of websites. It would also discuss different aspects of mobile and social network security. It would finish off with details of different prevention software such as firewalls, virtual private networks, anti-virus and anti-spyware, and routine updates.	**45**

36. AI845 Internet of Things(3, 0)
Pre-requisite: None
Recommended Books:

1. Internet of Things: Principals and Paradigms by Rajkumar Buyya, Amir Vahid Dastjerdi, 1st Edition, Morgan Kaufmann, 2016.

Credit Hours: 3 (3, 0)
Course Objectives:
On completion of the course, the student should be able to:

- Explain in a concise manner how the general Internet as well as Internet of Things work.

- Understand constraints and opportunities of wireless and mobile networks for Internet of Things.
- Use basic measurement tools to determine the real-time performance of packet based networks.
- Analyse trade-offs in interconnected wireless embedded sensor networks.

Topics / Contents	**Allocated Periods**
Introduction of Internet-of-Things, applications in various domain: smart buildings, healthcare, agriculture, urban infrastructure, transportation, assistive tracking for the blind, fundamental design issues for the future Internet, differences between Internet and Internet-of-things, design issues of Internet-of-Things, research challenges, primer on TCP/IP stack, wireless network protocol, medium access control, comparative study of ZigBee, bluetooth, ultro wide band (UWB), IEEE 802.11 a/b/g, Wi-Fi, RFID, capillary networks: data aggregation, 6LowPAN architecture, routing protocol in lossy networks (RPL): performance analysis and evaluation in TinyOS, directed acyclic graph (DAG) construction, parent-child relationship, objective function, minimum rank with hysteresis, constrained application protocol (CoAP).	**45**

37. AI828 Complex Adaptive Systems (3, 0)

Pre-requisite: None

Recommended Books:

Complex and Adaptive Dynamical Systems: A Primer (2011), by Claudius Gros, Springer, 2nd Edition.

Complexity: A Guided Tour (2011), by Melanie Mitchell, Oxford University Press, 1st Edition.

Credit Hours: 3 (3, 0)

Course Objectives:

On completion of the course, the student should be able to:

- Explain a comprehensive, clear and accessible outline to the new area of complex adaptive systems and its application in different areas.
- Understand about the dynamics of cooperation and competition, looking at how and why agents work together to create local patterns of organization.
- Understand the process of self-organization and answer the big questions about how do we get and sustain ordered patterns out of randomness and chaos?

- Explain the process of evolution as a powerful and relentless force that shapes complex adaptive systems on the macro scale.

Topics / Contents	**Allocated Periods**
Complex adaptive systems (CAS) are a broad class of systems consisting of multiple interacting adaptive agents. These systems, which span a wide range of disciplines, have a number of characteristics in common. They are large distributed systems consisting of many self-similar components that interact and adapt. These interactions among the distributed components are self-organizing and produce emergent collective behavior in the system as a whole. CAS tend to be difficult to analyze using traditional analytical models. Agent-based models have been shown to be effective methods for studying CAS. This course will introduce the basic definitions of CAS, discuss example cases of CAS and their features, and implement and analyze computational simulations of CAS. The content include, Overview of CAS, Properties of CAS, Cellular automata, Social systems, Evolution of cooperation, Emergence, Self-organization, Swarm Intelligence, Social networks	**45**

38. AI 819 Text Analytics (3, 0)

Pre-requisite: None

Recommended Books:

Mining Text Data, by Charu C. Aggarwal and ChengXiang Zhai, Springer, 2012.

Foundations of statistical natural language processing, by Manning, Christopher D., and Hinrich Schütze. MIT press, 1999.

Speech & Language Processing, by Dan Jurafsky and James H Martin, Pearson Education India, 2000.

Introduction to Information Retrieval (2007), by Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schuetze, Cambridge University Press, 2007.

Applied Text Analysis with Python (2018), by Benjamin Bengfort, Tony Ojeda and Rebecca Bilbro, O'Reilly Media.

Credit Hours: 3 (3, 0)

Course Objectives:

On completion of the course, the student should be able to:

- have a basic and hands-on understanding of the currently used frameworks and methods for text analytics and natural language understanding, in particular the application of machine learning methods to text analytics.

Topics / Contents	Allocated Periods
Given the dominance of text information over the Internet, mining high-quality information from text becomes increasingly critical. The actionable knowledge extracted from text data facilitates our life in a broad spectrum of areas, including business intelligence, information acquisition, social behavior analysis and decision making. In this course, we will cover important topics in text mining including: basic natural language processing techniques, document representation, text categorization and clustering, document summarization, sentiment analysis, social network and social media analysis, probabilistic topic models and text visualization.	**45**

39. AI846 Computational Creativity (3, 0)

Pre-requisite: None

Recommended Books:

Computational Creativity: The Philosophy and Engineering of Autonomously Creative Systems (2019), by Tony Veale, F. Amílcar Cardoso, Springer, 1st Edition, Kindle Edition

Computational Creativity Research: Towards Creative Machines (2015), by by Tarek R. Besold, Marco Schorlemmer, Alan Smaill, Atlantis Press, 2015 Edition.

Creative Environments: Issues of Creativity Support for the Knowledge Civilization Age, by Andrzej P. Wierzbicki, Yoshiteru Nakamori, Springer, 2007th Edition

Credit Hours: 3 (3, 0)

Course Objectives:

On completion of the course, the student should be able to:

- become familiar with the literature on Computational Creativity.
- become familiar with the state of the art in Computational Creativity.
- acquire experience in designing an Interactive/Autonomous Creative technique or tool.
- construct a program or computer capable of human-level creativity.
- better understand human creativity and to formulate an algorithmic perspective on creative behavior in humans.
- design programs that can enhance human creativity without necessarily being creative themselves.

Topics / Contents	Allocated Periods

Computational creativity is a multidisciplinary field that lies at the intersection of artificial intelligence, cognitive psychology, philosophy, and the arts. The field is concerned with the theoretical and practical issues in the study of creativity. This course is about Computational Creativity with a focus on modeling/discovery & design/invention. Contents of the course include, human creativity, information processing theories of human creativity, interactive tools for augmenting and amplifying human creativity, design of autonomous creative systems for creative tasks.	**45**

40. AI 837 Intelligent Transportation Systems (3, 0)

Pre-requisite: None

Recommended Books:

Sussman, Joseph. Perspectives on Intelligent Transportation Systems (ITS). New York, NY: Springer, 2010.

Mashrur A. Chowdhury, and Adel Sadek, Fundamentals of Intelligent Transportation Systems Planning, Artech House, Inc., 2003.

Pradip Kumar Sarkar, Amit Kumar Jain, Intelligent Transport Systems, PHI learning, 2018.

Credit Hours: 3 (3, 0)

Course Objectives:

On completion of the course, the student should be able to:

- explain transport telemetrics and its increasing significance in transportation planning and management.
- explain scope of transport issues, such as, traffic safety, public transport, advanced vehicle management and control.
- explain the application of information technology and telecommunications to control traffic, inform travelers and drivers, operate public transport, automating payments, handle emergencies and incidents, operate commercial fleets and freight exchange, and automate driving and safety.

Topics / Contents	**Allocated Periods**
Specific topics to be covered in the course include: Identification of transportation problems and costs, Definition and role of Intelligent Transportation Systems, Policy-makers' perspective on ITS, ITS management, Traveler Information Systems, Public transit, bicycles and pedestrians, Eco-friendly and sustainable ITS solutions, ITS	**45**

technologies: Automated highway systems (AHS), Autonomous Vehicles, Intelligent Infrastructures, Evaluation of technologies and large-scale ITS field tests, benefits and costs assessment of ITS, Learning from ITS deployments in various countries, ITS Challenges and Issues: Technical, institutional, funding, and procurement issues, ITS evaluation software, Public and private sector perspectives (institutional and stakeholder issues) on ITS.	

41. AI 829 Social Simulations (3, 0)

Pre-requisite: None

Recommended Books:

Introduction to Computational Social Science: Principles and Applications (2017), by Claudio Cioffi-Revilla, 2nd Ed. Springer

Simulation for the Social Scientist (2005), by Nigel Gilbert, Klaus Troitzsch, 2nd Edition, Open University Press

Monte Carlo Simulation and Resampling Methods for Social Science (2013), by Thomas Carsey, First Edition, SAGE Publications

Interdisciplinary Applications of Agent-Based Social Simulation and Modeling (2014), by Diana Francisca Adamatti, Graçaliz Pereira Dimuro, Helder Coelho, 1st Edition, IGI Global

Credit Hours: 3 (3, 0)

Course Objectives:

On completion of the course, the student should be able to answer:

- When agent-based models are most fruitfully applied,
- How to construct a conceptual agent-based model and then formalize it,
- How to implement the formalized model as custom computer software in Netlogo,
- How to use the Netlogo programs to carry out simulation experiments that will produce scientifically defensible assertions, and
- A broad range of applications of ABM in contemporary sociology research.

Topics / Contents	**Allocated Periods**
In this course we will learn how artificial societies can be specified (agent-based modeling) and then made to exist and evolve in time (computer simulation). The techniques of modeling and simulation provide a systematic way to state and test hypotheses about the microscopic mechanisms (e.g. individual behaviors) that might plausibly be responsible for the emergence of observed macroscopic systemic social patterns. Primary contents covered in the	**45**

course will constitute to answer (i) when agent-based models are most fruitfully applied (ii) how to construct a conceptual agent-based model and then formalize it, (iii) how to implement the formalized model as custom computer software in Netlogo, and (iv) how to use the Netlogo programs to carry out simulation experiments that will produce scientifically defensible assertions. To achieve all this, the course will begin with a few "toy examples", and then a series of research papers in sociology (and the social sciences more broadly) in which agent-based modeling and simulation has played a central role. For each of these papers, the nature of the question would be dissected, the agent-based model, the simulation, the experiments, and the conclusions -- whenever possible, would be attempted to re-create the results using hands-on (computer) laboratory experiments. In the course students will gain an appreciation for the importance of agent-based simulation in sociology research, and begin to become empowered to create their own models and apply simulation techniques towards sociological research questions of interest to us. Agent-based Modelling and Simulation in the Social and Human Sciences (2007), by Denis Phan , Frederic Amblard, The Bardwell Press	

42. AI 838 Serious Games (3, 0)

Pre-requisite: None

Recommended Books:

Serious Games, Foundations, Concepts and Practice (2016), by Dörner, R., Göbel, S., Effelsberg, W., Wiemeyer, J., Springer

Serious Games: Mechanisms and Effects (2009), by Ute Ritterfeld, Michael J. Cody, Peter Vorderer, Routledge Publications

The Art of Serious Game Design: A hands-on workshop for developing educational games: Facilitator guide, by Digital Education Strategies, The Chang School of Continuing Education, Ryerson University

Rules of Play: Game Design Fundamentals, Katie Salen, Eric Zimmerman

What Video Games Have to Teach Us About Learning and Literacy. Second Edition, James Paul Gee Recent research:

Affective Ludology: Scientific Measurement of User Experience in Interactive Entertainment, Ph.D. Thesis, Lennart E. Nacke (2009)

Credit Hours: 3 (3, 0)

Course Objectives:

On completion of the course, the student should be able to:

- Critically appraise legal and ethical aspects of recording and storing personal data gathered in conjunction with gaming activity
- Evaluate user interface design, and to conduct rigorous analysis of human-computer interactions presented by such designs
- Evaluate and judge the effects of video games both in general and for specialized domains
- Have an in-depth understanding of the psychophysiology and psychometrics underlying serious game design
- Understand the range of game mechanics used in games to engage the user
- Understand the effects of games on player behavior
- Identify and critically examine primary literature in the area of serious games

Topics / Contents	**Allocated Periods**
The field of modeling and simulation is large and diverse; modeling and simulation is applied in every discipline to answer questions in research and development and for education and training. Recently there has been a push in the use of specific simulations which have an internal goal. Such simulations are known as games (when applied to education and training, they are known as serious games) and provide an experience that is fun and engaging making them particularly relevant to the current generation of learners. This course is an introduction to simulation with an emphasis on serious games. Primary contents which would be covered in this course will include, Ethical considerations of using games to change behavior, Game interfaces, Universal design, Research methodologies for game interactions, Game telemetry, Game metrics, AI evaluation techniques for "big data" from game telemetry, Evaluation techniques for games research.	**45**

43. AI 856 Ethical Machines (3, 0)
Pre-requisite: None
Recommended Books:
Machine Ethics (2011), by Michael Anderson, Susan Leigh Anderson, 1st Edition, Cambridge University Press
Machine Ethics and Robot Ethics (2016), by Wendell Wallach, Peter Asaro, 1st Edition, CRC Press.

Moral Machines, Teaching Robots Right from Wrong (2010), by Wendell Wallach, Colin Allen, Oxford University Press, 2009
The Technological Singularity (2015), by Murray Shanahan, MIT Press.
Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy (2016), by Cathy O'Neil, Crown Publications
Moral Machines: Teaching Robots Right from Wrong (2010), by Wendell Wallach, Colin Allen, 1st Edition, Oxford University Press.

Credit Hours: 3 (3, 0)

Course Objectives:

On completion of the course, the student should be able to:

- hone their skills in critically exploring the development and implications of AI systems.
- have gained a broader conception of dilemmas in current AI technologies
- have a stronger framework with which to develop their own ethical responsibilities.

Topics / Contents	**Allocated Periods**
Recent advances in computing may place us at the threshold of a unique turning point in human history. Soon we are likely to entrust management of our environment, economy, security, infrastructure, food production, healthcare, and to a large degree even our personal activities, to artificially intelligent computer systems. The autonomous systems raises many complex and troubling questions. How will society respond as versatile robots and machine-learning systems displace an ever-expanding spectrum of blue- and white-collar workers? Will the benefits of this technological revolution be broadly distributed or accrue to a lucky few? How can we ensure that these systems respect our ethical principles when they make decisions at speeds and for rationales that exceed our ability to comprehend? What, if any, legal rights and responsibilities should we grant them? And should we regard them merely as sophisticated tools or as a newly emerging form of life? The goal this course is to equip students with the intellectual tools, ethical foundation, and psychological framework to successfully navigate the coming age of intelligent machines. Primary contents of this course will include, Overview of the history and issues raised by Artificial Intelligence, Philosophical issues, Ethical problems, Legal, social and economic impact, and Future challenges.	**45**