ISAT 344: Intelligent Systems (Sec 1) – 3 credits
Fall 2014

COURSE AND INSTRUCTOR INFORMATION

Meeting Times:	Wednesday Nights from 5-7:30pm in ISAT/CS 343
Open Lab Time:	Monday Nights from 8pm-Late in ISAT/CS 343 and 337
Instructor:	Nicole Radziwill, Ph.D., MBA
Office:	ISAT/CS 325
Phone/SMS:	703.835.6336 (SMS or Email 24/7)
Email:	radziwnm@jmu.edu
Office hours:	Mon 8pm-Late; Mon/Wed 9-Noon by appointment; In Person, Skype or gChat anytime by appt!

IMPORTANT NOTE: First in-class meeting is Wednesday, September 3rd

NATURE OF COURSE CONTENT

COURSE DESCRIPTION

Catalog Course Description: In-depth introduction to current and future intelligent systems, including expert systems, neural networks, hybrid intelligent systems and other intelligent system technologies and their development, uses and limitations. *Prerequisites: ISAT 340 or CS 239.*

Instructor's Course Description: This course is an in-depth introduction to current and evolving intelligent systems, focusing on intelligent agents and featuring expert/knowledge-based systems, artificial neural networks, web science, and other intelligent technologies, as they contribute to building intelligent agents. The course is application oriented and project-team based, and includes an examination of social issues, for example, ethical, legal and policy issues, associated with the development of intelligent systems.

Objectives:

- ***Intelligent systems in context*** Students should be able to relate the historical development of intelligent systems, describe the state of the practice in intelligent systems applications, and explain the philosophical, social, and ethical issues associated with intelligent systems.
- ***Knowledge acquisition and representation*** Students should be able to explain core issues and problems in knowledge acquisition (KA) and knowledge representation (KR), and use several KR techniques to build a knowledge base.
- ***Reasoning*** Students should be able to describe the classical systems of logic used to build inference engines for (reasoning into) intelligent systems, the properties of an ideal inference engine, and ways in which they are extended to cover reasoning under uncertainty.
- ***Expert/knowledge-based Systems*** Students should be able to describe the architecture of a typical expert/knowledge-based system, describe the methodology used to develop expert/knowledge-based systems, and use the methodology to develop an expert system.
- ***Machine Learning*** Students should be able to explain cores issues and problems in machine learning and to describe several machine learning technologies.

- ***Artificial Neural Networks (ANNs)*** Students should be able to describe a typical architecture for an ANN, explain the components and function of an ANN, describe a methodology to develop an ANN, and use the methodology to design and develop an ANN.
- ***Agent-Based Systems*** Students should be able to describe the properties of an ideal intelligent agent, describe the role of existing intelligent systems technologies in building intelligent agents, and describe the state of the practice in agent applications.
- ***Knowledge Engineering*** Students should be able to describe the process of knowledge engineering, including the main knowledge engineering activities such as knowledge acquisition and representation, and discuss the classic difficulties encountered when knowledge engineering an intelligent system.

COURSE STYLE & DELIVERY

This course implements the **10 Principles of the Burning Mind Project** as its core value system. (http://www.burningmindproject.org/the-ten-principles/) As a result, the course is somewhat self-directed, blended (integrating online and in-class components), gift-oriented, and synchronously coordinated with Benton's 340 (Software Engineering) & 348 (Web Development) courses.

- **Blended**: Some of the work can be done online, at your leisure, but there will be scheduled in-class exercises to keep you on track. There will be ample open lab time, with instructor guidance, provided for you to complete lab exercises and extend those lab exercises to create new explorations and new recipes.
- **Gift-Oriented**: In many classes, you may ask "what can I *get* out of taking this class?" However, in this course, we want you to ask the question "what can I *give* to others as a result of my participation in this class?" *Individual gifts* are an important component. As you explore the topics, we request that you identify things you are good at and can contribute to a larger, team project. Put together a one-page poster pop for each individual gift that you think would be useful to contribute to a larger group project. Then, your classmates can find you if they need your skills to build out a particular solution.
- **Synchronously Coordinated**: We also encourage you to partner with M. Benton's ISAT 340/348 project teams to construct applications that leverage multiple skill sets for your individual and group projects.

The course consists of one 2.5 hour combined lecture/lab session each week on Wednesdays.

SCHEDULE AT A GLANCE

Week 1:	Intro to Intelligent Systems - INDEPENDENT WORK
Week 2:	Intelligent Agents
Week 3-5:	Neural Networks
Week 6:	Text Analysis & Spam Filtering (Bayes Classifier)
Week 7:	K-Means Clustering & KNN Classification
Week 8:	Performance Measures: Evaluating & Improving Model Performance
Week 9:	Individual Project Presentations
Week 10:	Search & Community Detection in the Intelligent Web
Week 11:	Decision Trees (C5.0) & Knowledge-Based Systems
Week 12-13:	Work on Group Projects
	Thanksgiving Break
Week 14:	Ethics & Social Context of Intelligent Systems

DETAILED SCHEDULE (subject to minor changes throughout semester)

Week # & Dates	Wednesday
Week 1 8/27/14	*No Class Meetings - Independent Work!* ***Read: Lantz Ch.1*** **Watch introductory videos and take online quiz on Canvas by 11:59pm on Tuesday, 9/2:** • The Fifth Wave - Intelligent Systems: http://www.youtube.com/watch?v=VVQ50ybb7wY • Cognitive Systems - http://www.youtube.com/watch?v=hNZH6xQssmM • Games with a Purpose - http://www.youtube.com/watch?v=OvVAViDtKeA • Peter Norvig - Past, Present, and Future of AI http://www.youtube.com/watch?v=qI623nyCdKE
Week 2 9/3/14	*Agent-Based Systems (Lecture)* • Definitions, Architecture, Ideal Rational Agent, Weak v. Strong Agency, Types of Agents *Agent-Based Systems (Lab Exercises)* • PAGE for Describing Agent Specifications • PEAS for Characterizing Agent Environments • Behavioral Characteristics of Agents
Week 3 9/10/14	***Read: Lantz Ch.7*** *Neural Networks (Lecture/Videos)* • Intro to ANNs, NEUROSCIENCE OF CATEGORIZATION (30 min), How Perceptron Works *Neural Networks (Labs)* • Taxonomy & Training Exercise • Perceptron Exercise
Week 4 9/17/14	*Neural Networks (Lecture)* • The Neuroscientific Basis for Perceptrons and Artificial Neural Networks
Week 5 9/24/14	***Read: Lantz Ch.7*** *Neural Networks for Classification (Lab)* • Multilayer Feedforward NN with AMORE package in R for image recognition *Neural Networks for Regression (Lab)* • Multilayer Feedforward NN with AMORE package in R vs. Logistic Regression, model comparison, and quantitative prediction
Week 6 10/1/14	***Read: Lantz Ch.4*** *Text Analysis (Lab)* • Using the tm, wordcloud, and tau packages in R to extract text metadata for processing in neural networks and other classifiers *Naive Bayes Classifier/Spam Filtering (Lab)* • Implementing intelligent agents that use metadata extracted from text analysis
Week 7 10/8/14	***Read: Lantz Ch.3 & 9*** *K-Means Clustering & KNN Classification (Labs)*

Week 8 10/15/14	***Read: Lantz Ch.10 & 11*** *Performance Measures for Classifiers (Lab)* • Precision, Recall, and F1
Week 9 10/22/14	***Individual Project Presentations***
Week 10 10/29/14	***Read: Marmanis & Babenko Ch.1, 2 or 3 OR $25B Eigenvector*** *depending on group selection* *Intelligent Web (Participative Lecture)* • What is the Intelligent Web? The $25B Eigenvector *Intelligent Web (Lab)* • Compare rankings of documents using PageRank, HITS, Hub & Authority scores in R
Week 11 11/5/14	***Read: Lantz Ch.5*** *Decision Trees vs. Knowledge Based Systems (Lecture)* • KBS Architecture, Business Rule Structure, Product Config & Decision Support • 3 Faces of KBS Specifications, Decision Trees, IF/THEN Business Rules, Forward and Backward Chaining *Decision Trees vs. Knowledge Based Systems (Lab)* • Forward & Backward Chaining • C5.0 Classification Using Decision Trees
Week 12 11/12/14	*Work on Group Projects*
Week 13 11/19/14	*Work on Group Projects*
11/26/13	**Thanksgiving Break**
Week 14 12/3/14	***Read: Big Data is Our Generation's Civil Rights Issue*** *Ethical Analysis of Intelligent Robotic and Weaponry Systems, Smart Materials, Rising Trend of Embedded Intelligent Systems in "Internet of Things"*
Dec 8-12, 2014	***EXAM WEEK: Group Project Presentations***

GOALS OF THE COURSE

COURSE OBJECTIVES

By the end of this course, you will learn some of the basic concepts behind both AI and intelligent systems. You will also learn how to analyze some real-life systems by building and using models of each type. The learning goals for this course are listed below. You will:

1. ***Understand*** and be able to use the basic terminology used by practitioners in the artificial intelligence and intelligent systems areas such as:
   - Artificial intelligence/Intelligent systems
   - Machine learning
   - Agents/Intelligent Agents and Agent-Based Models

- Expert Systems/Knowledge-Based Systems
- Knowledge Representation & Knowledge Engineering
- Neural Networks/Artificial Neural Networks
- Wireless/Wired Sensor Networks
- Swarm Intelligence and Artificial Life
- Speech recognition/Pattern recognition
- Robotics, control systems, unmanned vehicles
- Microbotics, micromanipulation, and smart materials
- How Web search works and why it's intelligent, Search Engine Optimization

2. ***Become familiar with some tools that help you develop and use intelligent systems*** including Corvid, NeuroShell, NetLogo and the statistical software package R
3. ***Become familiar with the journals*** that describe the latest research in artificial intelligence and intelligent systems
4. Describe and ***apply a systematic process/methodology*** for developing expert systems and neural network models

METHODS OF EVALUATION

GRADING

The goal of this course is to *produce artifacts* that demonstrate your understanding of the topics that we cover, and that *provide value* to real clients and/or real people (including your instructors, the students within your learning community, and students who will participate in future learning communities by enrolling in ISAT 344). **Grading is based on *accrued points*** for completing various labs, exams, and projects, to the satisfaction of the instructor and/or instructor-designated proxies. You get as many chances as you like each week to continually improve the quality of your work, but points can only be accrued for successful completion of a particular activity or artifact during the week in which it is scheduled or the week immediately following. There are a total of 50 available points outlined below:

1 Point Each	**2 Points Each** (prereq. in paren.)	**3 Points Each**	**4 Points Each**
1. Intro Quiz (by 9/2*) 2. PAGE Exercise 3. PEAS Exercise 4. Perceptron XLS 5. Attend Dr. H/coloring 6. KBS If/then exercise 7. KBS Decision trees	8. ANN Image Classification (4, 5) 9. ANN Regression (4, 5) 10. Netlogo Exploration (2, 3) 11. Text Fundamentals 12. Naive Bayes Classifier (11) 13. K-Means/KNN 14. HITS/PageRank 15. Performance Measures (8, 11, 13) 16. KBS Chaining (6, 7) 17. KBS Forward Chaining in JESS (6, 7) 18. KBS Backward Chaining in JESS (6, 7) 19. Complete any new lab created by fellow students and approved by instructor 20. Present Group Project at end of semester	**21. Complete and Present Individual Project (*)** 22. Create New Lab 23. Final Exam 85%+ or Oral Exam	**24. Complete Grp Project (*)** - no more than 4 people on a team 25. Teach a Class (requires completing exercises in advance, getting approved by instructor)

(*) = REQUIRED
Grading rubrics for assessment of the team projects will be issued later in the semester.

Grading Scale: 30+ accrued points is an A, 23-29 is a B, 16-22 is a C, 9-15 is a D, 8 or below is an F. Your points are "vested" when you deliver your final project presentation *with your team* during finals week.

REQUIREMENTS & POLICIES

REQUIRED TEXTS and SOFTWARE:

1. Most of our readings and some of our exercises this semester come from an amazing book that you are required to get called **MACHINE LEARNING WITH R** by Brett Lantz, which you order from Amazon: http://www.amazon.com/gp/product/1782162143/ref=as_li_tl?ie=UTF8&camp=1789&creative=390957&creativeASIN=1782162143&linkCode=as2&tag=qualandinnowe-20&linkId=WAVXETO4VE3IL7TV

2. We will read one chapter (provided) from Marmanis & Babenko's "Algorithms of the Intelligent Web"

3. We will also use the **FREE, cross-platform** R Statistical Software which you should download from **http://www.r-project.org**

4. You may also find these pages useful:
http://en.wikipedia.org/wiki/Intelligence
http://en.wikipedia.org/wiki/Philosophy_of_artificial_intelligence
http://en.wikipedia.org/wiki/Inference_engine
http://en.wikipedia.org/wiki/Machine_learning
http://en.wikipedia.org/wiki/List_of_machine_learning_algorithms

ADD/DROP DEADLINES

All of the dates related to adding, dropping, and withdrawing from this course are in the JMU catalog and are posted on the University Registrar's web site. **YOU ARE RESPONSIBLE FOR KNOWING THESE DATES.** Professors are not required to grant grades of "WP" or "WF" after that date. If you have an extraordinary situation you may be granted an "I," but only under extraordinary and unanticipated circumstances that you discuss with me in advance.

COURSE POLICIES AND PROCEDURES

Attendance/Being Late
The way to be successful in this class is to attend regularly and complete the lab exercises that are scheduled for class time each week. Any accommodations (e.g. for sickness, JMU sports) must be made **ahead of time**. As long as there is a justifiable reason that I agree with, I will be as flexible as I can to help you complete the requirements for this course. The most important part is *setting my expectations effectively.* I will check you off on labs during the week that they are scheduled, or the week immediately following, but not after.

Working in Groups
Group work is an important part of the two projects in this course. I strongly encourage collaboration and expect that you cite your sources (even if those sources are fellow classmates).

Homework
Homework will not be turned in for this class. Your completion and understanding of the material will be assessed through your lab exercises and your semester-long team project.

Missed Exams
There are no exams in this course.

Class or Work Missed for Extraordinary Circumstances
Any student who seeks to be excused from class for extraordinary circumstances such as a prolonged illness or the death of a family member has the *personal* responsibility to inform me as soon as knowledge of the absence becomes known, preferably by email. It is the sole responsibility of the student to meet with me *immediately* upon return to campus to work out a schedule of when and how missed work will be made up. If there is a delay by the student of more than 24 hours upon return to campus in contacting me to work out a schedule to make up missed work, and missed work is returned to the class in the interim, the student will not be allowed to make up the missed work and will receive a zero for the missed work.

Special Needs
If you are a student who is registered with the Office of Disabilities, I need to be given written documentation to support your situation in order to provide you with any accommodations (this is required by law). YOU are responsible for pre-arranging with me to provide accommodations (e.g., additional time for an exam). I suggest an email a few days ahead of time so that arrangements can be made. In your email, you need only say that you are contacting me to arrange for your accommodation; you do not need to elaborate in your email. Due to family and professional obligations, I can not necessarily stay later than a scheduled exam time, so it is necessary to work things like this out in advance.

Class Participation and Assigned Reading
Please read the material before coming to class so that you will get the most out of lecture and are able to discuss the day's topics and ask questions. Questioning and discussing in class is expected and will contribute to assessing your final grade for the class.

Honor Code
You are expected to abide by the JMU Honor Code at all times. Examples of academic dishonesty that are violations of the Honor Code include, but are not limited to, the following: turning in work under only your own name that you did not actually do completely yourself (for collaborative work, *always* list the names of your collaborators), plagiarizing other people's words or computer code (and that includes text off the Internet), receiving unauthorized help on an exam, providing unauthorized help on an exam (and that includes talking about an exam before all students have taken it), and misuse of materials that are permitted for an exam. Violations of the JMU Honor Code will be dealt with in accordance with the policy that permits professors, at their own discretion, to assess and penalize students for cheating. All incidents of academic dishonesty will be reported to the Honor Committee, according to the requirements of the university.