

San Pablo Catholic University (UCSP)
Undergraduate Program in
Computer Science
SILABO



CS261. Artificial intelligence (Mandatory)

1. General information

1.1 School	:	Ciencia de la Computación
1.2 Course	:	CS261. Artificial intelligence
1.3 Semester	:	7 ^{mo} Semestre.
1.4 Prerequisites	:	MA203. Statistics and Probability. (4 th Sem)
1.5 Type of course	:	Mandatory
1.6 Learning modality	:	Face to face
1.7 Horas	:	2 HT; 4 HP;
1.8 Credits	:	4
1.9 Plan	:	Plan Curricular 2016

2. Professors

Lecturer

- Juan Carlos Gutiérrez Cáceres <jcgutierrezc@ucsp.edu.pe>
 - PhD in Ciencia de la Computación, Universidad Nacional de San Agustín, Perú, 2013.
 - MSc in Ciencia de la Computación, ICMC-USP, Brasil, 2003.

3. Course foundation

Research in Artificial Intelligence has led to the development of numerous relevant tonic, aimed at the automation of human intelligence, giving a panoramic view of different algorithms that simulate the different aspects of the behavior and the intelligence of the human being.

4. Summary

1. Fundamental Issues 2. Agents 3. Basic Search Strategies 4. Advanced Search 5. Reasoning Under Uncertainty 6. Basic Machine Learning 7. Advanced Machine Learning 8. Natural Language Processing 9. Perception and Computer Vision

5. Generales Goals

- Evaluate the possibilities of simulation of intelligence, for which the techniques of knowledge modeling will be studied.
- Build a notion of intelligence that later supports the tasks of your simulation.

6. Contribution to Outcomes

This discipline contributes to the achievement of the following outcomes:

- 1) Analyze a complex computing problem and to apply principles of computing and other relevant disciplines to identify solutions. (**Usage**)
- 5) Function effectively as a member or leader of a team engaged in activities appropriate to the program's discipline. (**Familiarity**)
- 6) Apply computer science theory and software development fundamentals to produce computing-based solutions. (**Familiarity**)

7. Content

UNIT 1: Fundamental Issues (2)

Competences:

Content

Generales Goals

- Overview of AI problems, examples of successful recent AI applications
- What is intelligent behavior?
 - The Turing test
 - Rational versus non-rational reasoning
- Problem characteristics
 - Fully versus partially observable
 - Single versus multi-agent
 - Deterministic versus stochastic
 - Static versus dynamic
 - Discrete versus continuous
- Nature of agents
 - Autonomous versus semi-autonomous
 - Reflexive, goal-based, and utility-based
 - The importance of perception and environmental interactions
- Philosophical and ethical issues.

- Describe Turing test and the “Chinese Room” thought experiment [Usage]
- Determining the characteristics of a given problem that an intelligent systems must solve [Usage]

Readings: De Castro (2006), Ponce-Gallegos et al. (2014)

UNIT 2: Agents (2)	
Competences:	
Content	Generales Goals
<ul style="list-style-type: none"> • Definitions of agents • Agent architectures (e.g., reactive, layered, cognitive) • Agent theory • Rationality, game theory <ul style="list-style-type: none"> – Decision-theoretic agents – Markov decision processes (MDP) • Software agents, personal assistants, and information access <ul style="list-style-type: none"> – Collaborative agents – Information-gathering agents – Believable agents (synthetic characters, modeling emotions in agents) • Learning agents • Multi-agent systems <ul style="list-style-type: none"> – Collaborating agents – Agent teams – Competitive agents (e.g., auctions, voting) – Swarm systems and biologically inspired models 	<ul style="list-style-type: none"> • List the defining characteristics of an intelligent agent [Usage] • Characterize and contrast the standard agent architectures [Usage] • Describe the applications of agent theory to domains such as software agents, personal assistants, and believable agents [Usage] • Describe the primary paradigms used by learning agents [Usage] • Demonstrate using appropriate examples how multi-agent systems support agent interaction [Usage]
Readings: Nilsson (2001), Russell and Norvig (2003), Ponce-Gallegos et al. (2014)	

UNIT 3: Basic Search Strategies (2)	
Competences:	
Content	Generales Goals
<ul style="list-style-type: none"> • Problem spaces (states, goals and operators), problem solving by search • Factored representation (factoring state into variables) • Uninformed search (breadth-first, depth-first, depth-first with iterative deepening) • Heuristics and informed search (hill-climbing, generic best-first, A*) • Space and time efficiency of search • Two-player games (introduction to minimax search) • Constraint satisfaction (backtracking and local search methods) 	<ul style="list-style-type: none"> • Formulate an efficient problem space for a problem expressed in natural language (eg, English) in terms of initial and goal states, and operators [Usage] • Describe the role of heuristics and describe the trade-offs among completeness, optimality, time complexity, and space complexity [Usage] • Describe the problem of combinatorial explosion of search space and its consequences [Usage] • Compare and contrast basic search issues with game playing issues [Usage]
Readings: Nilsson (2001), Ponce-Gallegos et al. (2014)	

UNIT 4: Advanced Search (18)	
Competences:	
Content	Generales Goals
<ul style="list-style-type: none"> • Stochastic search <ul style="list-style-type: none"> – Simulated annealing – Genetic algorithms – Monte-Carlo tree search • Constructing search trees, dynamic search space, combinatorial explosion of search space • Implementation of A* search, beam search • Minimax search, alpha-beta pruning • Expectimax search (MDP-solving) and chance nodes 	<ul style="list-style-type: none"> • Design and implement a genetic algorithm solution to a problem [Usage] • Design and implement a simulated annealing schedule to avoid local minima in a problem [Usage] • Design and implement A*, beam search to solve a problem [Usage] • Apply minimax search with alpha-beta pruning to prune search space in a two-player game [Usage] • Compare and contrast genetic algorithms with classic search techniques [Usage] • Compare and contrast various heuristic searches vis-a-vis applicability to a given problem [Usage]
Readings: Goldberg (1989), Nilsson (2001), Russell and Norvig (2003), Ponce-Gallegos et al. (2014)	

UNIT 5: Reasoning Under Uncertainty (18)	
Competences:	
Content	Generales Goals
<ul style="list-style-type: none"> • Review of basic probability • Random variables and probability distributions <ul style="list-style-type: none"> – Axioms of probability – Probabilistic inference – Bayes' Rule • Conditional Independence • Knowledge representations <ul style="list-style-type: none"> – Bayesian Networks <ul style="list-style-type: none"> * Exact inference and its complexity * Randomized sampling (Monte Carlo) methods (e.g. Gibbs sampling) – Markov Networks – Relational probability models – Hidden Markov Models 	<ul style="list-style-type: none"> • Apply Bayes' rule to determine the probability of a hypothesis given evidence [Usage] • Explain how conditional independence assertions allow for greater efficiency of probabilistic systems [Usage] • Identify examples of knowledge representations for reasoning under uncertainty [Usage] • State the complexity of exact inference Identify methods for approximate inference [Usage]
Readings: Koller and Friedman (2009), Russell and Norvig (2003)	

UNIT 6: Basic Machine Learning (4)	
Competences:	
Content	Generales Goals
<ul style="list-style-type: none"> • Definition and examples of broad variety of machine learning tasks, including classification • Inductive learning • Simple statistical-based learning, such as Naive Bayesian Classifier, decision trees • The over-fitting problem • Measuring classifier accuracy 	<ul style="list-style-type: none"> • List the differences among the three main styles of learning: supervised, reinforcement, and unsupervised [Usage] • Identify examples of classification tasks, including the available input features and output to be predicted [Usage] • Explain the difference between inductive and deductive learning [Usage] • Describe over-fitting in the context of a problem [Usage] • Apply the simple statistical learning algorithm such as Naive Bayesian Classifier to a classification task and measure the classifier's accuracy [Usage]
Readings: Mitchell (1998), Russell and Norvig (2003), Ponce-Gallegos et al. (2014)	

UNIT 7: Advanced Machine Learning (20)	
Competences:	
Content	Generales Goals
<ul style="list-style-type: none"> • Definition and examples of broad variety of machine learning tasks • General statistical-based learning, parameter estimation (maximum likelihood) • Inductive logic programming (ILP) • Supervised learning <ul style="list-style-type: none"> – Learning decision trees – Learning neural networks – Support vector machines (SVMs) • Unsupervised Learning and clustering <ul style="list-style-type: none"> – EM – K-means – Self-organizing maps • Semi-supervised learning • Learning graphical models • Performance evaluation (such as cross-validation, area under ROC curve) • Application of Machine Learning algorithms to Data Mining (cross-reference IM/Data Mining) 	<ul style="list-style-type: none"> • Explain the differences among the three main styles of learning: supervised, reinforcement, and unsupervised [Usage] • Implement simple algorithms for supervised learning, reinforcement learning, and unsupervised learning [Usage] • Determine which of the three learning styles is appropriate to a particular problem domain [Usage] • Compare and contrast each of the following techniques, providing examples of when each strategy is superior: decision trees, neural networks, and belief networks [Usage] • Evaluate the performance of a simple learning system on a real-world dataset [Usage] • Characterize the state of the art in learning theory, including its achievements and its shortcomings [Usage] • Explain the problem of overfitting, along with techniques for detecting and managing the problem [Usage]
Readings: Russell and Norvig (2003), Koller and Friedman (2009), Murphy (2012)	

UNIT 8: Natural Language Processing (12)**Competences:****Content****Generales Goals**

- Deterministic and stochastic grammars
- Parsing algorithms
 - CFGs and chart parsers (e.g. CYK)
 - Probabilistic CFGs and weighted CYK
- Representing meaning / Semantics
 - Logic-based knowledge representations
 - Semantic roles
 - Temporal representations
 - Beliefs, desires, and intentions
- Corpus-based methods
- N-grams and HMMs
- Smoothing and backoff
- Examples of use: POS tagging and morphology
- Information retrieval
 - Vector space model
 - * TF & IDF
 - Precision and recall
- Information extraction
- Language translation
- Text classification, categorization
 - Bag of words model

- Define and contrast deterministic and stochastic grammars, providing examples to show the adequacy of each [Usage]
- Simulate, apply, or implement classic and stochastic algorithms for parsing natural language [Usage]
- Identify the challenges of representing meaning [Usage]
- List the advantages of using standard corpora Identify examples of current corpora for a variety of NLP tasks [Usage]
- Identify techniques for information retrieval, language translation, and text classification [Usage]

Readings: Nilsson (2001), Russell and Norvig (2003), Ponce-Gallegos et al. (2014)

UNIT 9: Perception and Computer Vision (12)	
Competences:	
Content	Generales Goals
<ul style="list-style-type: none"> • Computer vision <ul style="list-style-type: none"> – Image acquisition, representation, processing and properties – Shape representation, object recognition and segmentation – Motion analysis • Modularity in recognition • Approaches to pattern recognition <ul style="list-style-type: none"> – Classification algorithms and measures of classification quality – Statistical techniques 	<ul style="list-style-type: none"> • Summarize the importance of image and object recognition in AI and indicate several significant applications of this technology [Usage] • List at least three image-segmentation approaches, such as thresholding, edge-based and region-based algorithms, along with their defining characteristics, strengths, and weaknesses [Usage] • Implement 2d object recognition based on contour-and/or region-based shape representations [Usage] • Provide at least two examples of a transformation of a data source from one sensory domain to another, eg, tactile data interpreted as single-band 2d images [Usage] • Implement a feature-extraction algorithm on real data, eg, an edge or corner detector for images or vectors of Fourier coefficients describing a short slice of audio signal [Usage] • Implement a classification algorithm that segments input percepts into output categories and quantitatively evaluates the resulting classification [Usage] • Evaluate the performance of the underlying feature-extraction, relative to at least one alternative possible approach (whether implemented or not) in its contribution to the classification task (8), above [Usage]
Readings: Nilsson (2001), Russell and Norvig (2003), Ponce-Gallegos et al. (2014)	

8. Methodology

1. El profesor del curso presentará clases teóricas de los temas señalados en el programa propiciando la intervención de los alumnos.
2. El profesor del curso presentará demostraciones para fundamentar clases teóricas.
3. El profesor y los alumnos realizarán prácticas
4. Los alumnos deberán asistir a clase habiendo leído lo que el profesor va a presentar. De esta manera se facilitará la comprensión y los estudiantes estarán en mejores condiciones de hacer consultas en clase.

9. Assessment

Continuous Assessment 1 : 20 %

Partial Exam : 30 %

Continuous Assessment 2 : 20 %

Final exam : 30 %

References

- De Castro, L.N. (2006). *Fundamentals of natural computing: basic concepts, algorithms, and applications*. CRC Press.
- Goldberg, David (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison Wesley.
- Koller, Daphne and Nir Friedman (2009). *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press. ISBN: 0262013193.
- Mitchell, M. (1998). *An introduction to genetic algorithms*. The MIT press.
- Murphy, Kevin P. (2012). *Machine Learning: A Probabilistic Perspective*. The MIT Press. ISBN: 0262018020.
- Nilsson, Nils (2001). *Inteligencia Artificial: Una nueva visión*. McGraw-Hill.
- Ponce-Gallegos, Julio et al. (2014). *Inteligencia Artificial*. Iniciativa Latinoamericana de Libros de Texto Abiertos (LATIn).
- Russell, Stuart and Peter Norvig (2003). *Inteligencia Artificial: Un enfoque moderno*. Prentice Hall.