CS607 - Artificial Intelligence
FAQs

Question:	What is Artificial Intelligence about?
Answer:	Artificial intelligence (abbreviated AI, also some times called Synthetic Intelligence) is defined as intelligence exhibited by an artificial entity. Such a system is generally assumed to be a computer. Although AI has a strong science fiction connotation, it forms a vital branch of computer science, dealing with intelligent behavior, learning and adaptation in machines. Research in AI is concerned with producing machines to automate tasks requiring intelligent behavior
Question:	What is the difference between Computational Intelligence and AI?
Answer:	Conventional AI mostly involves methods now classified as machine learning, characterized by formalism and statistical analysis. This is also known as symbolic AI, logical AI, neat AI and Good Old Fashioned Artificial Intelligence (GOFAI). (Also see semantics.) Methods include: Expert systems: apply reasoning capabilities to reach a conclusion. An expert system can process large amounts of known information and provide conclusions based on them. Case based reasoning Bayesian networks Behavior based AI: a modular method of building AI systems by hand. Computational Intelligence involves iterative development or learning (e.g. parameter tuning e.g. in connectionist systems). Learning is based on empirical data and is associated with non-symbolic AI, scruffy AI and soft computing. Methods mainly include: Neural networks: systems with very strong pattern recognition capabilities. Fuzzy systems: techniques for reasoning under uncertainty, has been widely used in modern industrial and consumer product control systems. Evolutionary computation: applies biologically inspired concepts such as populations, mutation and survival of the fittest to generate increasingly better solutions to the problem. These methods most notably divide into evolutionary algorithms (e.g. genetic algorithms) and swarm intelligence (e.g. ant algorithms). With hybrid intelligent systems attempts are made to combine these two groups. Expert inference rules can be generated through neural network or production rules from statistical learning such as in ACT-R. A promising new approach called intelligence amplification tries to achieve artificial intelligence in an evolutionary development process as a side-effect of amplifying human intelligence through technology.
Question:	What is the focus of this course?
Answer:	In this course we will try to envelop some important and basic concepts that will help the students to get an insight into the main topics that Artificial Intelligence deals with.
Question:	Are there any prerequisites to the AI course?
Answer:	In general, the following background is strongly recommended: knowledge of basic computer science principles and skills, understanding of basic propositional logic (predicate logic is a plus) as well as graph theory, ability to understand and analyze fairly complicated algorithms and data structures, familiarity with the basic concepts of probability theory.
Question:	Where is the timetable for this course?
Answer:	Details are provided on the LMS Course Website Page.
Question:	How to contact the lecturer(s) and tutor(s)?
Answer:	You can contact the lecturer(s) and tutor(s) by mailing them at the email address cs607@vu.edu.pk.Unless there is a valid reason, questions sent to the lecturer(s) via e-mail will be redirected to and answered in the appropriate forum on the Discussion Board.

Question:	Is there any dedicated AI news site?
Answer:	The American Association for Artificial Intelligence (AAAI) provides a daily news feed on AI Topics - a most recommended bookmark to anyone interested in the latest AI applications, research, and developments!
Question:	What is the major area's where AI is used?
Answer:	AI systems are now in routine use in economics, medicine, engineering and the military, as well as being built into many common home computer software applications, traditional strategy games like computer chess and other video games. Also AI include control, planning and scheduling, the ability to answer diagnostic and consumer questions, handwriting, speech, and facial recognition. As such, it has become a scientific discipline, focused on providing solutions to real life problems.
Question:	What are the recommended books for this course?
Answer:	There are three books; 1: Artificial Intelligence, 3rd edition by Winston. 2: Artificial Intelligence by Elaine Rich, Kevin Knight. 3: Artificial Intelligence : Structures and Strategies for Complex Problem Solving 5th edition by George Luger.
Question:	What's the difference between strong AI and weak AI?
Answer:	Strong AI makes the bold claim that computers can be made to think on a level (at least) equal to humans. Weak AI simply states that some "thinking-like" features can be added to computers to make them more useful tools... and this has already started to happen (witness expert systems, drive-by-wire cars and speech recognition software). What does 'think' and 'thinking-like' mean? That's a matter of much debate.
Question:	What is the history of AI?
Answer:	For an online timeline of artificial intelligence milestones, see ftp://ftp.cs.ucla.edu/AI/timeline.txt The appendix to Ray Kurzweil's book "Intelligent Machines" (MIT Press,1990, ISBN 0-262-11121-7, $39.95) gives a timeline of the history of AI. Pamela McCorduck, "Machines Who Think", Freeman, San Francisco, CA, 1979. Allen Newell, "Intellectual Issues in the History of Artificial Intelligence", Technical Report CMU-CS-82-142, Carnegie Mellon University Computer Science Department, October 28, 1982
Question:	What has AI accomplished?
Answer:	Quite a bit, actually. In 'Computing machinery and intelligence.', Alan Turing, one of the founders of computer science, made the claim that by the year 2000, computers would be able to pass the Turing test at a reasonably sophisticated level, in particular, that the average interrogator would not be able to identify the computer correctly more than 70 per cent of the time after a five minute conversation. AI hasn't quite lived upto Turing's claims, but quite a bit of progress has been made, including: - Deployed speech dialog systems by firms like IBM, Dragon and Lernout&Hauspie - Applications of expert systems/case-based reasoning: a computerized Lukemia diagnosis system did a better job checking for blood disorders than human experts! - Machine translation for Environment Canada: software developed in the 1970s translated natural language weather forcasts between English and French. Purportedly stil in use. - Deep Blue, the first computer to beat the human chess Grandmaster - Fuzzy controllers in dishwashers, etc. One persistent 'problem' is that as soon as an AI technique trully succeeds, in the minds of many it ceases to be AI, becoming something else entirely. For example, when Deep Blue defeated Kasparov, there were many who said Deep Blue wasn't AI, since after all it was just a brute force parallel minimax search (!)
Question:	What are the branches of AI?
Answer:	There are many, some are 'problems' and some are 'techniques'. Automatic Programming - The task of describing what a program should do and having the AI system 'write' the program. Bayesian Networks - A

	technique of structuring and inferencing with probabilistic information. Natural Language Processing(NLP) - Processing and (perhaps)understanding human ("natural") language Knowledge Engineering/Representation - turning what we know about a particular domain into a form in which a computer can understand it. Planning - given a set of actions, a goal state, and a present state, decide which actions must be taken so that the present state is turned into the goal state Constraint Statisfaction - solving NP-complete problems, using a variety of techniques. Machine Learning - Programs that learn from experience. Visual Pattern Recognition - The ability to reproduce the human sense of sight on a machine. Speech Recogntion - Conversion of speech into text. Search - The finding of a path from a start state to a goal state. Similar to planning, yet different... Neural Networks(NN) - The study of programs that function in a manner similar to how animal brains do. AI problems (speech recognition, NLP, vision, automatic programming,knowledge representation, etc.) can be paired with techniques (NN,search, Bayesian nets, production systems, etc.) to make distinctions such as search-based NLP vs. NN NLP vs. Statistical/Probabilistic NLP. Then you can combine techniques, such as using neural networks to guide search. And you can combine problems, such as posing that knowledge representation and language are equivalent. (Or you can combine AI with problems from other domains.)
Question:	Is AI an inherited branch of different disciplines?
Answer:	It has inherited its ideas, concepts and techniques from many disciplines like philosophy, mathematics, psychology, linguistics, biology etc.
Question:	Who is concerned with Neural Networks?
Answer:	Neural Networks are interesting for quite a lot of very different people: • Computer scientists want to find out about the properties of non-symbolic information processing with neural nets and about learning systems in general. • Statisticians use neural nets as flexible, nonlinear regression and classification models. • Engineers of many kinds exploit the capabilities of neural networks in many areas, such as signal processing and automatic control. • Cognitive scientists view neural networks as a possible apparatus to describe models of thinking and consciousness (High-level brain function). • Neuro-physiologists use neural networks to describe and explore medium-level brain function (e.g. memory, sensory system, motorics). • Physicists use neural networks to model phenomena in statistical mechanics and for a lot of other tasks. • Biologists use Neural Networks to interpret nucleotide sequences. • Philosophers and some other people may also be interested in Neural Networks for various reasons. For world-wide lists of groups doing research on NNs, see the Foundation for Neural Networks's (SNN) page at http://www.mbfys.kun.nl/snn/pointers/groups.html and see Neural Networks Research on the IEEE Neural Network Council's homepage http://www.ieee.org/nnc.
Question:	Isn't there a solid definition of intelligence that doesn't depend on relating it to human intelligence?
Answer:	Not yet. The problem is that we cannot yet characterize in general what kinds of computational procedures we want to call intelligent. We understand some of the mechanisms of intelligence and not others.
Question:	Isn't AI about simulating human intelligence?
Answer:	Sometimes but not always or even usually. On the one hand, we can learn something about how to make machines solve problems by observing other people or just by observing our own methods. On the other hand, most work in AI involves studying the problems the world presents to intelligence rather than studying people or animals. AI researchers are free to use methods that are not observed in people or that involve much more computing than people can do.
Question:	What about IQ? Do computer programs have IQs?
Answer:	No. IQ is based on the rates at which intelligence develops in children. It is the ratio of the age at which a

	child normally makes a certain score to the child's age. The scale is extended to adults in a suitable way. IQ correlates well with various measures of success or failure in life, but making computers that can score high on IQ tests would be weakly correlated with their usefulness. For example, the ability of a child to repeat back a long sequence of digits correlates well with other intellectual abilities, perhaps because it measures how much information the child can compute with at once. However, ``digit span'' is trivial for even extremely limited computers.
Question:	When did AI research start?
Answer:	After World War II, a number of people independently started to work on intelligent machines. The English mathematician Alan Turing may have been the first. He gave a lecture on it in 1947. He also may have been the first to decide that AI was best researched by programming computers rather than by building machines. By the late 1950s, there were many researchers on AI, and most of them were basing their work on programming computers.
Question:	Does AI aim to put the human mind into the computer?
Answer:	Some researchers say they have that objective, but maybe they are using the phrase metaphorically. The human mind has a lot of peculiarities, and I'm not sure anyone is serious about imitating all of them.
Question:	Does AI aim at human-level intelligence?
Answer:	Yes. The ultimate effort is to make computer programs that can solve problems and achieve goals in the world as well as humans. However, many people involved in particular research areas are much less ambitious.
Question:	Where is Fuzzy logic used?
Answer:	Fuzzy logic is used directly in very few applications. The Sony PalmTop apparently uses a fuzzy logic decision tree algorithm to perform handwritten (well, computer lightpen) Kanji character recognition. Most applications of fuzzy logic use it as the underlying logic system for fuzzy expert systems.
Question:	Where are Fuzzy Experts Systems used?
Answer:	To date, fuzzy expert systems are the most common use of fuzzy logic. They are used in several wide-ranging fields, including: Linear and Nonlinear Control Pattern Recognition Financial Systems Operation Research Data Analysi
Question:	Isn't "fuzzy logic" an inherent contradiction? Why would anyone want to fuzzify logic?
Answer:	Fuzzy sets and logic must be viewed as a formal mathematical theory for the representation of uncertainty. Uncertainty is crucial for the management of real systems: if you had to park your car PRECISELY in one place, it would not be possible. Instead, you work within, say, 10 cm tolerances. The presence of uncertainty is the price you pay for handling a complex system. Nevertheless, fuzzy logic is a mathematical formalism, and a membership grade is a precise number. What's crucial to realize is that fuzzy logic is a logic OF fuzziness, not a logic which is ITSELF fuzzy. But that's OK: just as the laws of probability are not random, so the laws of fuzziness are not vague.
Question:	How many assignments will be there for this course?
Answer:	There is an updated course calendar for the details of assignments and quiz at course website link.
Question:	CLIPS Stands for?
Answer:	CLIPS stands for C Language Integrated Production System.
Question:	What is CLIPS and from where I can download it?
Answer:	CLIPS is an expert system tool which provides a complete environment for the construction of rule and

	object based expert systems. Download CLIPS for windows (CLIPSWin.zip) from: http://vulms.vu.edu.pk/Courses/CS607/Downloads/CLIPSWin.zip Also download the complete documentation including the programming guide from: http://www.ghg.net/clips/download/documentation/ The guides that you download will provide comprehensive guidance on programming using CLIP. Here are some of the basics to get you started.
Question:	What is the difference between admissibility and monotonicity?
Answer:	The admissibility is the property of a search method (heuristic) that ensures that we find optimal path to a goal from all possible paths. The monotonicity is the property of a search method (heuristic) that also ensures that every node in the way of finding optimal path also has minimum distance from start state.So we can say that to follow the property of monotonicity our search space is arranged such that whenever we reached a certain node we are certain that this node's distance is minimum from the starting node.
Question:	which software have been used for diagrams?
Answer:	No specific software has been used for searching related diagrams however in later topics Math Type has been used for Mathematical Symbols and there are images from AI related software's like Matlab.
Question:	How to distinguish between the Resolution Strategies • Refraction • Regency • Specificity
Answer:	1. Specificity If a rule has all the conditions of another rule in it then only that rule can be taken for processing and second rule can be ignored, in this way we can say that first rule is more specific (contains more conditions and details) For example Rule I = if there is rain then I will not go to school Rule II = if there is rain and I don't have an umbrella then I will not go to school. Here we can ignore rule I and can take only rule II. 2. Recencity We give priority to most recently used rules over the least recently used rules for deducing conclusions. 3. Refraction If a rule is being used time and again and is not producing any results and causing our program to trap in loops we keep on decreasing its weight and a time comes when we simple ignore this rule and use only other rules.
Question:	Whether clips and/or matlab software will be allowed in the exams or not?
Answer:	No, there will be no need to use clips or matlab in the examination
Question:	What is the difference between Crisp values and Fuzzy values?
Answer:	Crisp values are those values which belong to a fixed set they can't be the part of more than one set for example value four 4 will be the part of set of even numbers it can't be the part of set of odd numbers similarly value 5 will be part of set of odd numbers. But in some situations a value can be taken as the part of more than one sets for example if the exam score of someone is 70 it may be considered as Good by some people and may be considered as Very Good by some other people so this value is member of two sets at a time so we can say that this value is fuzzy value now we will find the membership of this values is both the sets for example its membership in Good Set may by .7 (can be taken as 70%) and in Very Good set it may be .3 (30%) so in case of fuzzy logic we consider fuzzy values there is no absolute value, for example there is no absolute definition of tall man the person with height of 6 feet is also tall and person with height 6.5 feet is also tall but the membership value for the second person height will be greater than the first one in the set of persons who are Tall.
Question:	Can we use another version of clips or other tool as lisp or Proglog?
Answer:	Yes we can use other version of Clips but the mention version of CLIPS in MDB is latest one so it's better to use this version of clips. We are using just Clips in this course so it's better for you to use this one.

CS607 - Artificial Intelligence
Glossary

Action-based Planning :	The goal of action-based planning is to determine how to decompose a high level action into a network of sub actions that perform the requisite task. Therefore the major task within such a planning system is to manage the constraints that apply to the interrelationships (e.g., ordering constraints) between actions. In fact, action-based planning is best viewed as a constraint satisfaction problem. The search for a plan cycles through the following steps: choose a constraint and apply the constraint check; if the constraint is not satisfied, choose a bug from the set of constraint bugs; choose and apply a fix, yielding a new plan and possibly a new set of constraints to check. In contrast, state-based planners generally conduct their search for a plan by reasoning about how the actions within a plan affect the state of the world and how the state of the world affects the applicability of actions.
Adaptive Interface :	A computer interface that automatically and dynamically adapts to the needs and competence of each individual user of the software.
Agent Architecture :	There are two levels of agent architecture, when a number of agents are to work together for a common goal. There is the architecture of the system of agents, that will determine how they work together, and which does not need to be concerned with how individual agents fulfill their sub-missions; and the architecture of each individual agent, which does determine its inner workings. The architecture of one software agent will permit interactions among most of the following components (depending on the agent's goals): preceptors, effectors, communication channels, a state model, a model-based reasoned, a planner/scheduler, a reactive execution monitor, its reflexes (which enable the agent to react immediately to changes in its environment that it can't wait on the planner to deal with), and its goals. The preceptors, effectors, and communication channels will also enable interaction with the agent's outside world.
Agents :	Agents are software programs that are capable of autonomous, flexible, purposeful and reasoning action in pursuit of one or more goals. They are designed to take timely action in response to external stimuli from their environment on behalf of a human. When multiple agents are being used together in a system, individual agents are expected to interact together as appropriate to achieve the goals of the overall system also called autonomous agents, assistants, brokers, bots, droids, intelligent agents, software agents.
AI Effect :	The great practical benefits of AI applications and even the existence of AI in many software products go largely unnoticed by many despite the already widespread use of AI techniques in software. This is the AI effect. Many marketing people don't use the term "artificial intelligence" even when their company's products rely on some AI techniques. Why not? It may be because AI was oversold in the first giddy days of practical rule-based expert systems in the 1980s, with the peak perhaps marked by the Business Week cover of July 9, 1984 announcing, Artificial Intelligence, IT'S HERE. James Hogan in his book, Mind Matters, has his own explanation of the AI Effect: "AI researchers talk about a peculiar phenomenon known as the "AI effect." At the outset of a project, the goal is to entice a performance from machines in some designated area that everyone agrees would require "intelligence" if done by a human. If the project fails, it becomes a target of

	derision to be pointed at by the skeptics as an example of the absurdity of the idea that AI could be possible. If it succeeds, with the process demystified and its inner workings laid bare as lines of prosaic computer code, the subject is dismissed as "not really all that intelligent after all." Perhaps ... the real threat that we resist is the further demystification of ourselves...It seems to happen repeatedly that a line of AI work ... finds itself being diverted in such a direction that ... the measures that were supposed to mark its attainment are demonstrated brilliantly. Then, the resulting new knowledge typically stimulates demands for application of it and a burgeoning industry, market, and additional facet to our way of life comes into being, which within a decade we take for granted; but by then, of course, it isn't AI."
AI Languages and Tools :	AI software has different requirements from other, conventional software. Therefore, specific languages for AI software have been developed. These include LISP, Prolog, and Smalltalk. While these languages often reduce the time to develop an artificial intelligence application, they can lengthen the time to execute the application. Therefore, much AI software is now written in languages such as C++ and Java, which typically increases development time, but shortens execution time. Also, to reduce the cost of AI software, a range of commercial software development tools have also been developed. Stottler Henke has developed its own proprietary tools for some of the specialized applications it is experienced in creating.
Algorithm :	An algorithm is a set of instructions that explain how to solve a problem. It is usually first stated in English and arithmetic, and from this, a programmer can translate it into executable code (that is, code to be run on a computer).
Artificial Intelligence :	Artificial intelligence (AI) is the mimicking of human thought and cognitive processes to solve complex problems automatically. AI uses techniques for writing computer code to represent and manipulate knowledge. Different techniques mimic the different ways that people think and reason (see case-based reasoning and model-based reasoning for example). AI applications can be either stand-alone software, such as decision support software, or embedded within larger software or hardware systems. AI has been around for about 50 years and while early optimism about matching human reasoning capabilities quickly has not been realized yet, there is a significant and growing set of valuable applications. AI hasn't yet mimicked much of the common-sense reasoning of a five-year old child. Nevertheless, it can successfully mimic many expert tasks performed by trained adults, and there is probably more artificial intelligence being used in practice in one form or another than most people realize. Really intelligent applications will only be achievable with artificial intelligence and it is the mark of a successful designer of AI software to deliver functionality that can't be delivered without using AI.
Associative Memories :	Associative memories work by recalling information in response to an information cue. Associative memories can be auto associative or hetero-associative. Auto associative memories recall the same information that is used as a cue, which can be useful to complete a partial pattern. Hetero-associative memories are useful as a memory. Human long-term memory is thought to be associative because of the way in which one thought retrieved from it leads to another. When we want to store a new item of information in our long term memory it typically takes us 8 seconds to store an item that can't be associated with a pre-stored item, but only one or two seconds, if there is an existed information structure with which to associate the new item.
Automated Diagnosis Systems :	Most diagnosis work is done by expert humans such as mechanics, engineers, doctors, firemen, customer service agents, and analysts of various kinds. All of us usually do at least a little diagnosis even if it isn't a major part of our working lives. We use a range of techniques for our diagnoses. Primarily, we compare a current situation with past ones, and reapply, perhaps with small

	modifications, the best past solutions. If this doesn't work, we may run small mental simulations of possible solutions through our minds, based on first principles. We may do more complex simulations using first principles on paper or computers looking for solutions. Some problems are also amenable to quantitative solutions. We may hand off the problem to greater experts than ourselves, who use the same methods. The problem with humans doing diagnosis is that it often takes a long time and a lot of mistakes to learn to become an expert. Many situations just don't reoccur frequently, and we may have to encounter each situation several times to become familiar with it. Automatic diagnosis systems can help avoid these problems, while helping humans to become experts faster. They work best in combination with a few human experts, as there are some diagnosis problems that humans are better at solving, and also because humans are more creative and adaptive than computers in coming up with new solutions to new problems.
Autonomous Agents :	A piece of AI software that automatically performs a task on a human's behalf, or even on the behalf of another piece of AI software, so together they accomplish a useful task for a person somewhere. They are capable of independent action in dynamic, unpredictable environments. "Autonomous agent" is a trendy term that is sometimes reserved for AI software used in conjunction with the Internet (for example, AI software that acts as your assistance in intelligently managing your e-mail). Autonomous agents present the best hope from gaining additional utility from computing facilities. Over the past few years the term "agent" has been used very loosely. Our definition of a software agent is: "an intelligent software application with the authorization and capability to sense its environment and work in a goal directed manner." Generally, the term "agent" implies "intelligence", meaning the level of complexity of the tasks involved approaches that which would previously have required human intervention.
Bayesian Networks: :	A modeling technique that provides a mathematically sound formalism for representing and reasoning about uncertainty, imprecision, or unpredictability in our knowledge. For example, seeing that the front lawn is wet, one might wish to determine whether it rained during the previous night. Inference algorithms can use the structure of the Bayesian network to calculate conditional probabilities based on whatever data has been observed (e.g., the street does not appear wet, so it is 90% likely that the wetness is due to the sprinklers). Bayesian networks offer or enable a set of benefits not provided by any other system for dealing with uncertainty - an easy to understand graphical representation, a strong mathematical foundation, and effective automated tuning mechanisms. These techniques have proved useful in a wide variety of tasks including medical diagnosis, natural language understanding, plan recognition, and intrusion detection. Also called belief networks, Bays networks, or causal probabilistic networks.
Case-based Reasoning: :	Case-based reasoning (CBR) solves a current problem by retrieving the solution to previous similar problems and altering those solutions to meet the current needs. It is based upon previous experiences and patterns of previous experiences. Humans with years of experience in a particular job and activity (e.g., a skilled paramedic arriving on an accident scene can often automatically know the best procedure to deal with a patient) use this technique to solve many of their problems. One advantage of CBR is that inexperienced people can draw on the knowledge of experienced colleagues, including ones who aren't in the organization, to solve their problems. Synonym: Reasoning by analogy.
Classification :	Automated classification tools such as decision trees have been shown to be very effective for distinguishing and characterizing very large volumes of data. They assign items to one of a set of predefined classes of objects based on a set of observed features. For example, one might determine whether a particular mushroom is "poisonous" or "edible" based on its color, size, and gill size.

	Classifiers can be learned automatically from a set of examples through supervised learning
Clustering :	Clustering is an approach to learning that seeks to place objects into meaningful groups automatically based on their similarity. Clustering, unlike classification, does not require the groups to be predefined with the hope that the algorithm will determine useful but hidden groupings of data points. The hope in applying clustering algorithms is that they will discover useful but unknown classes of items. A well-publicized success of a clustering system was NASA's discovery of a new class of stellar spectra. See IQE, GIIF, Web Mediator, Rome Graphics, and data mining for examples of applications that use clustering.
Cognitive Science: :	Artificial intelligence can be defined as the mimicking of human thought to perform useful tasks, such as solving complex problems. This creation of new paradigms, algorithms, and techniques requires continued involvement in the human mind, the inspiration of AI. To that end, AI software designers team with cognitive psychologists and use cognitive science concepts, especially in knowledge elicitation and system design.
Cognitive Task Analysis: :	Cognitive task analysis (CTA) is a systematic process by which the cognitive elements of task performance are identified. This includes both domain knowledge and cognitive processing. Thus, CTA focuses on mental activities that cannot be observed and is in contrast to behavioral task analysis that breaks the task down into observable, procedural steps. CTA is most useful for highly complex tasks with few observable behaviors. Examples of cognitive processing elements include: to decide, judge, notice, assess, recognize, interpret, prioritize, and anticipate. Examples of domain knowledge elements include concepts, principles, and interrelationships; goals and goal structures; rules, strategies and plans; implicit knowledge; and mental models. The results from CTA have various applications such as identifying content to be included within training programs for complex cognitive tasks, research on expert-novice differences in terms of domain knowledge and cognitive processing during task performance, modeling of expert performance to support expert system design, and the design of human-machine interfaces.
Commonsense Reasoning: :	Ordinary people manage to accomplish an extraordinary number of complex tasks just using simple, informal thought processes based on a large amount of common knowledge. They can quickly plan and undertake a shopping expedition to six or seven different shops, as well as pick up the kids from soccer and drop a book back at the library, quite efficiently without logically considering the hundreds of thousands of alternative ways to plan such an outing. Artificial intelligence is far behind humans in using such reasoning except for limited jobs, and tasks that rely heavily on commonsense reasoning are usually poor candidates for AI applications
Computer Vision: :	Making sense of what we see is usually easy for humans, but very hard for computers. Practical vision systems to date are limited to working in tightly controlled environments. Synonym: machine vision.
Constraint Satisfaction: :	Constraints are events, conditions or rules that limit our alternatives for completing a task. Satisfying constraints is particularly important in scheduling complex activities. By first considering applicable constraints, the number of possible schedules to be considered in a search for an acceptable schedule can be reduced enormously, making the search process much more efficient. Constraint satisfaction techniques can be used to solve scheduling problems directly. Constraint satisfaction algorithms include heuristic constraint- based search and annealing.
Control System :	Control systems are an integral part of our daily life. Heater and air conditioners in our homes maintain temperature within our comfort zones. In our cars, things such as power steering, cruise control, and automatic transmission are based on automatic control systems. Air craft, automated

	manufacturing tools, and self guided agricultural equipment all use control systems to stay on track and perform their functions under all kinds of disturbances and fluctuations in the environments they are used.
Data Fusion: :	Information processing that deals with the association, correlation, and combination of data and information from single and multiple sources to achieve a more complete and more accurate assessment of a situation. The process is characterized by continuous refinement of its estimates and assessments, and by evaluation of the need for additional sources, or modification of the process itself, to achieve improved results.
Decision Aids: :	Software that helps humans makes decisions, particularly about complex matters when a high degree of expertise is needed to make a good decision.
Decision Support: :	Decision support is a broad class of applications for artificial intelligence software. There are many situations when humans would prefer machines, particularly computers, to either automatically assist them in making decisions, or actually make and act on a decision. Synonym: intelligent decision support.
Decision Theory: :	Decision theory provides a basis for making choices in the face of uncertainty, based on the assignment of probabilities and payoffs to all possible outcomes of each decision. The space of possible actions and states of the world is represented by a decision tree.
Decision Tree: :	A decision tree is a graphical representation of a hierarchical set of rules that describe how one might evaluate or classify an object of interest based on the answers to a series of questions. Decision trees can be automatically developed from a set of examples and are capable of discovering powerful predictive rules even when very large numbers of variables are involved. These algorithms operate by selecting the test that best discriminates amongst classes/diagnoses and then repeating this test selection process on each of the subsets matching the different test outcomes (e.g., "patients with temperatures greater than 101ºF" and "patients with temperatures less than or equal to 101ºF"). This process continues until all the examples in a particular set have the same class/diagnosis.
Dependency Maintenance: :	Dependency maintenance is the technique of recording why certain beliefs are held, decisions were made, or actions were taken, in order to facilitate revising those decisions, actions, or beliefs in the face of changing circumstances. Several families of truth maintenance systems have been developed to facilitate dependency maintenance in particular kinds of situations (e.g. need to consider many alternate scenarios versus a single scenario, frequency with which assumptions change, etc).
Domain Expert: :	The person who knows how to perform an activity within the domain, and whose knowledge is to be the subject of an expert system. This person's or persons' knowledge and method of work are observed, recorded, and entered into a knowledge base for use by an expert system. The domain expert's knowledge may be supplemented by written knowledge contained in operating manuals, standards, specifications, computer programs, etc., that are used by the experts. Synonym: subject-matter expert (SME).
Domain: :	An overworked word for AI people. "Domain" can mean a variety of things including a subject area, field of knowledge, an industry, a specific job, an area of activity, a sphere of influence, or a range of interest, e.g., chemistry, medical diagnosis, putting out fires, operating a nuclear power plant, planning a wedding, diagnosing faults in a car. Generally, a domain is a system in which a particular set of rules, facts, or assumptions operates. Humans can usually easily figure out what's meant from the context in which "domain" is used; computers could probably not figure out what a

	human means when he or she says "domain."
Emergence: :	Emergence is the phenomenon of complex patterns of behavior arising out of the myriad interactions of simple agents, which may each operate according to a few simple rules. To put it another way, an emergent system is much more than simply the sum of its parts. It can happen without any grand master outside the system telling the individual agents how to behave. Artificial intelligence software running on powerful computers can demonstrate useful emergent behavior as well, such as that demonstrated in automatic scheduling software that creates near-optimal schedules for complex activities subject to many constraints.
Expert System: :	An expert system encapsulates the specialist knowledge gained from a human expert (such as a bond trader or a loan underwriter) and applies that knowledge automatically to make decisions. For example, the knowledge of doctors about how to diagnose a disease can be encapsulated in software. The process of acquiring the knowledge from the experts and their documentation and successfully incorporating it in the software is called knowledge engineering, and requires considerable skill to perform successfully. Applications include customer service and helpdesk support, computer or network troubleshooting, regulatory tracking, autocorrect features in word processors, document generation such as tax forms, and scheduling.
Fuzzy Logic: :	Traditional Western logic systems assume that things are either in one category or another. Yet in everyday life, we know this is often not precisely so. People aren't just short or tall, they can be fairly short or fairly tall, and besides we differ in our opinions of what height actually corresponds to tall, anyway. The ingredients of a cake aren't just not mixed or mixed; they can be moderately well mixed. Fuzzy logic provides a way of taking our commonsense knowledge that most things are a matter of degree into account when a computer is automatically making a decision. For example, one rice cooker uses fuzzy logic to cook rice perfectly even if the cook put in too little water or too much water.
Game Theory: :	Game theory is a branch of mathematics that seeks to model decision making in conflict situations.
Genetic Algorithms: :	Search algorithms used in machine learning which involve iteratively generating new candidate solutions by combining two high scoring earlier (or parent) solutions in a search for a better solution. So named because of its reliance on ideas drawn from biological evolution.
Granularity :	Refers to the basic size of units that can be manipulated. Often refers to the level of detail or abstraction at which a particular problem is analyzed. One characteristic of human intelligence, Jerry R. Hobbs has pointed out, is the ability to conceptualize a world at different levels of granularity (complexity) and to move among them in considering problems and situations. The simpler the problem, the coarser the grain can be and still provide effective solutions to the problem
Heuristic: :	A heuristic is commonly called a rule of thumb. That is, an heuristic is a method for solving a problem that doesn't guarantee a good solution all the time, but usually does. An example of a heuristic would be to search for a lost object by starting in the last place you can remember using it.
Human-Centered Computing: :	Computers and other machines should be designed to effectively serve people's needs and requirements. All too often they're not. Commonly cited examples of this are the difficulty people have in setting up their VCR to record a TV show; and the difficulties people have in setting up a home computer facility, or hooking up to the Internet. Artificial intelligence software can be used to deliver more human-centered computing, improving system usability, extending the powerfulness of human reasoning and enabling greater collaboration amongst humans and machines, and promoting human learning. A goal of human-centered computing is for cooperating humans and machines to compensate for each other's respective weaknesses (e.g., machines to compensate for

	limited human short-term memory and the slowness with which humans can search through many alternative possible solutions to a problems; and for humans to compensate machines for their more limited pattern-recognition capability, language processing, and creativity) in support of human goals. Synonym: mixed initiative planning.
Hybrid Systems: :	Many of artificial intelligence software applications use multiple AI techniques in combination. For example, case-based reasoning may be used in combination with model-based reasoning in an automatic diagnostic system. Case-based reasoning, which tends to be less expensive to develop and faster to run, may draw on an historical databases of past equipment failures, the diagnosis of those, and the repairs effected and the outcomes achieved. So CBR may be used to make most failure diagnoses. Model-based reasoning may be used to diagnose less common but expensive failures, and also to make fine adjustments to the repair procedures retrieved from similar cases in the case base by CBR.
Inference Engine: :	The part of an expert system responsible for drawing new conclusions from the current data and rules. The inference engine is a portion of the reusable part of an expert system (along with the user interface, a knowledge base editor, and an explanation system), that will work with different sets of case-specific data and knowledge
Information Filtering: :	An information filtering system sorts through large volumes of dynamically generated information to present to the user those nuggets of information which are likely to satisfy his or her immediate needs. Information filtering overlaps the older field of information retrieval, which also deals with the selection of information. Many of the features of information retrieval system design (e.g. representation, similarity measures or Boolean selection, document space visualization) are present in information filtering systems as well. Information filtering is roughly information retrieval from a rapidly changing information space.
Intelligent Tutoring Systems: :	Encode and apply the subject matter and teaching expertise of experienced instructors, using artificial intelligence (AI) software technologies and cognitive psychology models, to provide the benefits of one-on-one instruction -- automatically and cost-effectively. These systems provide coaching and hinting, evaluate each student's performance, assess the student's knowledge and skills, provide instructional feedback, and select appropriate next exercises for the student.
KAPPA :	Rule-based object-oriented expert system tool and application developer. KAPPA is written in C, and is available for PCs. See AI Languages and Tools.
Knowledge Elicitation: :	Synonym: knowledge acquisition.
Knowledge Engineering: :	Knowledge engineering is the process of collecting knowledge from human experts in a form suitable for designing and implementing an expert system. The person conducting knowledge engineering is called a knowledge engineer.
Knowledge Representation: :	Knowledge representation is one of the two basic techniques of artificial intelligence, the other is the capability to search for end points from a starting point. The way in which knowledge is represented has a powerful effect on the prospects for a computer or person to draw conclusions or make inferences from that knowledge. Consider the representation of numbers that we wish to add. Which is easier, adding 10 + 50 in Arabic numerals, or adding X plus L in Roman numerals? Consider also the use of algebraic symbols in solving problems for unknown numerical quantities, compared with trying to do the same problems just with words and numbers.
Representations:	The form or structure of databases and knowledge bases for expert and other intelligent systems, so

:	that the information and solutions provided by a system are both accurate and complete. Usually involves a logically-based language capable of both syntactic and semantic representation of time, events, actions, processes, and entities. Knowledge representation languages include Lisp, Prolog, Smalltalk, OPS-5, and KL-ONE. Structures include rules, scripts, frames, endorsements, and semantic networks.
Knowledge-based Systems: :	Usually a synonym for expert system, though some think of expert systems as knowledge-based systems that are designed to work on practical, real-world problems.
LISP :	LISP (short for list processing language), a computer language, was invented by John McCarthy, one of the pioneers of artificial intelligence. The language is ideal for representing knowledge (e.g., If a fire alarm is ringing, then there is a fire) from which inferences are to be drawn.
Machine Learning: :	Machine learning refers to the ability of computers to automatically acquire new knowledge, learning from, for example, past cases or experience, from the computer's own experiences, or from exploration. Machine learning has many uses such as finding rules to direct marketing campaigns based on lessons learned from analysis of data from supermarket loyalty campaigns; or learning to recognize characters from people's handwriting. Machine learning enables computer software to adapt to changing circumstances, enabling it to make better decisions than non-AI software. Synonyms: learning, automatic learning.
Model-based Reasoning: :	Model-based reasoning (MBR) concentrates on reasoning about a system's behavior from an explicit model of the mechanisms underlying that behavior. Model-based techniques can very succinctly represent knowledge more completely and at a greater level of detail than techniques that encode experience, because they employ models that are compact axiomatic systems from which large amounts of information can be deduced.
Natural Language Processing: :	English is an example of a natural language, a computer language isn't. For a computer to process a natural language, it would have to mimic what a human does. That is, the computer would have to recognize the sequence of words spoken by a person or another computer, understand the syntax or grammar of the words (i.e., do a syntactical analysis), and then extract the meaning of the words. A limited amount of meaning can be derived from a sequence of words taken out of context (i.e., by semantic analysis); but much more of the meaning depends on the context in which the words are spoken (e.g., who spoke them, under what circumstances, with what tone, and what else was said, particularly before the words), which would require a pragmatic analysis to extract. To date, natural language processing is poorly developed and computers are not yet able to even approach the ability of humans to extract meaning from natural languages; yet there are already valuable practical applications of the technology.
Neural Networks: :	Neural networks are an approach to machine learning which developed out of attempts to model the processing that occurs within the neurons of the brain. By using simple processing units (neurons), organized in a layered and highly parallel architecture, it is possible to perform arbitrarily complex calculations. Learning is achieved through repeated minor modifications to selected neurons, which results in a very powerful classification system. A problem with neural networks is that it very difficult to understand their internal reasoning process, and therefore to obtain an explanation for any particular conclusion. They are best used, therefore, when the results of a model are more important than understanding how the model works. Neural network software is used to recognize handwriting, and also to control chemical processes to run at desired conditions. Other applications include stock market analysis, fingerprint identification, character recognition, speech recognition, credit analysis, scientific analysis of data, and in neuro physiological research. Neural networks are

	also known as neural nets, connectionism, and parallel associative memory.
Pattern Recognition: :	The use of feature analysis to identify an image of an object. May involve techniques such as statistical pattern recognition, Bayesian analysis, classification, cluster analysis, and analysis of texture and edges. See machine vision.
Plan Recognition :	The goal of plan recognition is to interpret an agent's intentions by ascribing goals and plans to it based on partial observation of its behavior up to the current time. Divining the agent's underlying plan can be useful for many purposes including: interpreting the agent's past behavior, predicting the agent's future behavior, or acting to collaborate with (or thwart) the agent.
Planning and Scheduling: :	Planning is the field of AI that deals with the synthesis of plans, which are partial orders of (possibly conditional) actions to meet specified goals under specified constraints. It is related to scheduling, which is the task of determining when and with what resources to carry out each member of a specific set of actions to satisfy constraints regarding ordering, effectiveness and resource allocation. In 1991, SHAI developed the concept of intelligent entities for planning and scheduling applications. Intelligent entities play the role of managers of various resources, groups of resources, tasks, and projects made up of tasks.
Qualitative Reasoning: :	Inexact reasoning, the opposite of quantitative reasoning
Relevance Feedback: :	Relevance feedback methods are used in information retrieval systems to improve the results produced from a particular query by modifying the query based on the user's reaction to the initial retrieved documents. Specifically, the user's judgments of the relevance or non-relevance of some of the documents retrieved are used to add new terms to the query and to re-weight query terms. For example, if all the documents that the user judges as relevant contain a particular term, then that term may be a good one to add to the original query.
Rule-based System: :	An expert system based on IF-THEN rules for representing knowledge
Signal Filtering :	Signal filtering is a technique for removing the noise or static from a signal so the clear or underlying signal remains. This is a conventional technique commonly used by electrical engineers and others
Simulated Annealing :	Simulated annealing is an optimization method based on an analogy with the physical process of toughening alloys, such as steel, called annealing. Annealing involves heating an alloy and cooling it slowly to increase its toughness. In simulated annealing, an artificial "temperature" is used to control the optimization process, of finding the overall maximum or minimum of a function. As cooling a metal slowly allows the atoms time to move to the optimum positions for toughness, giving time to look for a solution in simulated annealing permits a successful search for the global optimum and avoids being trapped at a local sub optima. It is used, for example, to optimize routing of planes by airlines for most efficient use of the fleet.
Simulation: :	A simulation is a system that is constructed to work, in some ways, analogously to another system of interest. The constructed system is usually made simpler than the original system so that only the aspects of interest are mirrored. Simulations are commonly used to learn more about the behavior of the original system, when the original system is not available for manipulation. It may not be available because of cost or safety reasons, or it may not be built yet and the purpose of learning about it is to design it better. If the purpose of learning is to train novices, then cost, safety, or convenience are likely to be the reasons to work on a simulated system. The simulation may be a

	computer simulation (perhaps a realistic one of a nuclear power station's control room, or a mathematical one such as a spreadsheet for "what-if" analysis of a company's business); or it may be a small-scale physical model (such as a small-scale bridge, or a pilot chemical plant).
Statistical Learning: :	Statistical learning techniques attempt to construct statistical models of an entity based on surface features drawn from a large corpus of examples. These techniques generally operate independent of specific domain knowledge, training instead on a set of features that characterize an input example. In the domain of natural language, for example, statistics of language usage (e.g., word trigram frequencies) are compiled from large collections of input documents and are used to categorize or make predictions about new text. Systems trained through statistical learning have the advantage of not requiring human-engineered domain modeling. This strong dependence on the input corpus has the disadvantage of limiting their applicability to new domains, requiring access to large corpora of examples and a retraining step for each domain of interest. Statistical techniques thus tend to have high precision within a domain at the cost of generality across domains
Supervised Learning: :	Organization and training of a neural network by a combination of repeated presentation of patterns, such as alphanumeric characters, and required knowledge. An example of required knowledge is the ability to recognize the difference between two similar characters such as O and Q. Synonym: learning with a teacher. Contrast with self-organized system; unsupervised learning.
Time Series Analysis: :	A time series is a sequence of observations of a particular variable over time (e.g., the daily closing level of Dow Jones Industrial Average). There are a wide range of statistical and temporal data mining techniques for analyzing such data. Two common uses for this type of analysis are forecasting future events (i.e., time series prediction) and searching a database of previous patterns for sequences that are similar to a particular pattern of interest. This is a conventional statistical technique.
Toy System: :	Small-scale implementation of a concept or model useful for testing a few main features, but unsuitable for complex or real-world problems. For example, a toy rule-based system may contain a few rules to construct an arch out of a number of pre-selected wooden blocks. It is a useful academic approach to unsolved problems. It is not employed in producing practical, real-world solutions.
Truth Maintenance Systems :	Many conventional reasoning systems assume that reasoning is the process of deriving new knowledge from old, i.e., the number of things a person or intelligent software believes increases without retracting any existing knowledge, since known truths never change under this form of logic. This is called monotonic logic. However, this view does not accurately capture the way in which humans think since our actions constantly change what we believe to be true. Humans reason non-monotonically, which means they reason based on incomplete or uncertain information, common sense, default values, changing conditions, and other assertions subject to retraction or revision. Truth maintenance systems seek to emulate the human reasoning process by recording reasons for our beliefs and reasons for retraction or revision of those beliefs, as well as the beliefs themselves. They are particularly useful in keeping track of goals, sub-goals, decisions, tasks, assignments, and design documents on complex projects (such as the design, construction, and testing of a major commercial aircraft) being undertaken by large numbers of people who may work for different organizations in different parts of the world. This is the sort of situation where a decision may be reversed, and all the people who may have to react to that change may not be properly informed. Project management software using TMS can help avoid design problems or wasted effort that can result from this sort of oversight. Also known as Reason Maintenance Systems