DOWNLOAD PDF LIST OF ARTIFICIAL INTELLIGENCE ALGORITHMS

Chapter 1 : Must Read Books on Machine Learning Artificial Intelligence

The below lists various Algorithms and Learning Techniques used in Artificial Intelligence and Machine Learning applications and systems. Learning Algorithms Below we list resources that give details on the different forms of learning algorithms and techniques available.

The company is also planning to expand its international operations to both developed and emerging regions as well as considering the option of making its own AI chips to back this endeavor. Furthermore, the company is currently investing considerably in research and development and boasts that one of the newest algorithms can assess and filter web content with an accuracy of Recently, SenseTime launched a non-profit lab in partnership with its investor and partner Alibaba to nurture artificial intelligence talent mainly in Hong Kong. In fact, the management teams of both companies share the board equally. Nonetheless, Argo AI is open to forming strategic collaborations with other companies. Furthermore, the team at Affirm has created a point-of-sale customer financing technology that utilizes artificial intelligence AI algorithms in performing a credit check on customers, which takes place almost instantaneously. Megvii Megvii, short form for Mega Vision, is one of the notable providers of facial recognition, which makes it a leading competitor of SenseTime. The startup can offer access to credit for individuals with minimal or no credit history. Currently, this service is available in several countries around the globe such as India, Russia, the Czech Republic, Mexico, Poland and Spain. In the commercial field, Kreditech provides its services and technology as a partner to other corporates and financial institutions by offering point-of-sale loans in a similar manner to Affirm. It also white labels its technology to banking institutions. In its recent expansion move into India, Kreditech collaborated with PayU, which is an online payment provider that operates in emerging markets to deliver monthly installments to e-commerce customers, particularly those devoid of a bank card. CloudWalk CloudWalk makes up another Chinese firm that utilizes artificial intelligence, specifically for facial recognition. Aside from using its technology in local Chinese police departs and financial institutions, the startup has been recently rolling out its revolutionary technology in airports all over China. Even so, the company is currently conducting a pilot project for its predictive crime mitigation system, which can make the fiction movie Minority Report a reality in the near future. The team at Zoox ventured into the designing of transportation robots, which can operate devoid of human being supervision, through working on various aspects including proprietary algorithms charging and maintenance station network as well as car design and manufacturing simultaneously. Australian co-founder Tim Kentley-Klay, Zoox has been extra secretive regarding its product, which will only be unveiled to the public in However, leading media outlets can test the Toyotas they utilize to teach the artificial intelligence algorithms. Looking at the above list, it is safe to conclude China, the United States, and the United Kingdom are some of the leading nations in the race for artificial intelligence dominance. As such, more governments and companies can be expected to continue investing in AI startups in an attempt to align themselves with the looming technological revolution. With this growing trend in the emergence of AI-based startups, the future appears to be bright as far as AI adoption is concerned.

Chapter 2 : Machine learning - Wikipedia

Read this introductory list of contemporary machine learning algorithms of importance that every engineer should understand. By James Le, Machine Learning Engineer. It is no doubt that the sub-field of machine learning / artificial intelligence has increasingly gained more popularity in the past couple of years.

View Blog Background â€" How many cats does it take to identify a Cat? In this article, I cover the 12 types of AI problems i. I address the question: Recently, I conducted a strategy workshop for a group of senior executives running a large multi national. In the workshop, one person asked the question: How many cats does it need to identify a Cat? On one level, the answer is very clear: That number is 10 million images.. But the answer is incomplete because the question itself is limiting since there are a lot more details in the implementation â€" for example training on a cluster with 1, machines 16, cores for three days. I wanted to present a more detailed response to the question. Also, many problems can be solved using traditional Machine Learning algorithms â€" as per an excellent post from Brandon Rohrer â€" which algorithm family can answer my question. So, in this post I discuss problems that can be uniquely addressed through AI. This is not an exact taxonomy but I believe it is comprehensive. I have intentionally emphasized Enterprise AI problems because I believe AI will affect many mainstream applications â€" although a lot of media attention goes to the more esoteric applications. What problem does Deep Learning address? What is Deep Learning? Firstly, let us explore what is Deep Learning Deep learning refers to artificial neural networks that are composed of many layers. In contrast, many other machine learning algorithms like SVM are shallow because they do not have a Deep architecture through multiple layers. The Deep architecture allows subsequent computations to build upon previous ones. The presence of multiple layers allows the network to learn more abstract features. Thus, the higher layers of the network can learn more abstract features building on the inputs from the lower layers. A Deep Learning network can be seen as a Feature extraction layer with a Classification layer on top. The power of deep learning is not in its classification skills, but rather in its feature extraction skills. Feature extraction is automatic without human intervention and multi-layered. The network is trained by exposing it to a large number of labelled examples. Errors are detected and the weights of the connections between the neurons adjusted to improve results. The optimisation process is repeated to create a tuned network. Once deployed, unlabelled images can be assessed based on the tuned network. Feature engineering involves finding connections between variables and packaging them into a new single variable is called. Deep Learning performs automated feature engineering. Automated feature engineering is the defining characteristic of Deep Learning especially for unstructured data such as images. This matters because the alternative is engineering features by hand. Deep Learning also suits problems that involve Hierarchy and Abstraction. Abstraction is a conceptual process by which general rules and concepts are derived from the usage and classification of specific examples. For Deep Learning, each layer is involved with detection of one characteristic and subsequent layers build upon previous ones. Hence, Deep Learning is used in situations where the problem domain comprises abstract and hierarchical concepts. Image recognition falls in this category. In contrast, a Spam detection problem that can be modelled neatly as a spreadsheet probably is not a complex problem to warrant Deep Learning A more detailed explanation of this question can be found in THIS Quora thread. Machine Learning Before we explore types of AI applications, we need to also discuss the differences between the three terms AI vs. The holy grail of AI is artificial general intelligence aka like Terminator! AI is evolving rapidly. A range of technologies drive AI currently. Improvements in Deep Learning algorithms drive AI. Deep Learning algorithms can detect patterns without the prior definition of features or characteristics. They can be seen as a hybrid form of supervised learning because you must still train the network with a large number of examples but without the requirement for predefining the characteristics of the examples features. Deep Learning networks have made vast improvements both due to the algorithms themselves but also due to better hardware specifically GPUs Finally, in a broad sense, the term Machine Learning means the application of any algorithm that can be applied against a dataset to find a pattern in the data. This includes algorithms like supervised, unsupervised, segmentation, classification, or

regression. Despite their popularity, there are many reasons why Deep learning algorithms will not make other Machine Learning algor Problems which involve Reasoning based on a complex body of knowledge This includes tasks which are based on learning a body of knowledge like Legal, financial etc. Problems which involve extending a complex body of Knowledge Here, the machine learns a complex body of knowledge like information about existing medication etc. Tasks which involve Planning Many logistics and scheduling tasks can be done by current non AI algorithms. But increasingly, as the optimization becomes complex AI could help. AI meets Re-engineering the corporation! While autonomous vehicles etc get a lot of media attention, AI will be deployed in almost all sectors of the economy. In each case, the same principles apply i. AI will be used to create new insights from automatic feature detection via Deep Learning - which in turn help to optimize, improve or change a business process over and above what can be done with traditional machine learning. I outlined some of these processes in financial services in a previous blog: As per Bernard Marr writing in Forbes: With advances in fields such as image recognition, sentiment analysis and natural language processing, this information is starting to give up its secrets, and mining it will become increasingly big business in In practise, this will mean enhancing the features of ERP and Datawarehousing systems through Cognitive systems. Deep learning has improved computer vision, for example, to the point that autonomous vehicles cars and trucks are viable. But what will be their impact? Autonomous vehicles alone will impact: For example, in Speech recognition, improvements continue to be made and currently, the abilities of the machine equal that of a human. Just six weeks ago, Microsoft engineers reported that their system reached a word error rate of 5. The goal-post continues to be moved rapidly.. Of course, the same ideas can be implemented independently of Watson today. The application of AI techniques to sequential pattern recognition is still an early stage domain and does not yet get the kind of attention as CNNs for example â€” but in my view, this will be a rapidly expanding space. There are already many synergies between AI and Sentiment analysis because many functions of AI apps need sentiment analysis features. The existing reported solutions or available systems are still far from being perfect or fail to meet the satisfaction level of the end users. The main issue may be that there are many conceptual rules that govern sentiment and there are even more clues possibly unlimited that can convey these concepts from realization to verbalization of a human being. Automatic feature learning is the key feature of AI. AI needs many detailed and pragmatic strategies which I have not yet covered here. AI comes with a cost skills, development, and architecture but provides an exponential increase in performance. The winners in AI will take an exponential view addressing very large scale problems i.

Chapter 3 : Twelve types of Artificial Intelligence (AI) problems - Data Science Central

A few common types of artificial intelligence. Artificial intelligence is technology that is designed to learn and self-improve. It is typically used to solve complex problems that are impossible to tackle with traditional code.

As Big Data is the hottest trend in the tech industry at the moment, machine learning is incredibly powerful to make predictions or calculated suggestions based on large amounts of data. So if you want to learn more about machine learning, how do you start? For me, my first introduction is when I took an Artificial Intelligence class when I was studying abroad in Copenhagen. My lecturer is a full-time Applied Math and CS professor at the Technical University of Denmark, in which his research areas are logic and artificial, focusing primarily on the use of logic to model human-like planning, reasoning and problem solving. The textbook that we used is one of the AI classics: At the end of the class, in a team of 3, we implemented simple search-based agents solving transportation tasks in a virtual environment as a programming project. I have learned a tremendous amount of knowledge thanks to that class, and decided to keep learning about this specialized topic. In the last few weeks, I have been multiple tech talks in San Francisco on deep learning, neural networks, data architecture â€" and a Machine Learning conference with a lot of well-known professionals in the field. In this post, I want to share some of the most common machine learning algorithms that I learned from the course. Machine learning algorithms can be divided into 3 broad categories â€" supervised learning, unsupervised learning, and reinforcement learning. A decision tree is a decision support tool that uses a tree-like graph or model of decisions and their possible consequences, including chance-event outcomes, resource costs, and utility. Take a look at the image to get a sense of how it looks like. As a method, it allows you to approach the problem in a structured and systematic way to arrive at a logical conclusion. Naive Bayes Classification Some of real world examples are: To mark an email as spam or not spam Classify a news article about technology, politics, or sports Check a piece of text expressing positive emotions, or negative emotions? Used for face recognition software. Ordinary Least Squares Regression: If you know statistics, you probably have heard of linear regression before. Least squares is a method for performing linear regression. You can think of linear regression as the task of fitting a straight line through a set of points. Ordinary Least Squares Regression Linear refers the kind of model you are using to fit the data, while least squares refers to the kind of error metric you are minimizing over. Logistic regression is a powerful statistical way of modeling a binomial outcome with one or more explanatory variables. It measures the relationship between the categorical dependent variable and one or more independent variables by estimating probabilities using a logistic function, which is the cumulative logistic distribution. Logistic Regression In general, regressions can be used in real-world applications such as: Credit Scoring Measuring the success rates of marketing campaigns Predicting the revenues of a certain product Is there going to be an earthquake on a particular day? SVM is binary classification algorithm. Given a set of points of 2 types in N dimensional place, SVM generates a N â€" 1 dimensional hyperplane to separate those points into 2 groups. Say you have some points of 2 types in a paper which are linearly separable. SVM will find a straight line which separates those points into 2 types and situated as far as possible from all those points. Support Vector Machine In terms of scale, some of the biggest problems that have been solved using SVMs with suitably modified implementations are display advertising, human splice site recognition, image-based gender detection, large-scale image classification

Chapter 4 : Artificial Intelligence Popular Search Algorithms

Learn the fundamentals of Artificial Intelligence (AI), and apply them. Design intelligent agents to solve real-world problems including, search, games, machine learning, logic, and constraint satisfaction problems.

Machine Learning Yearning Nick Bostrom has authored or co-authored over publications, including this book called Superintelligence. But how many of us stop to think about how AI will affect our society? Are we considering the human aspect at all when building AI products and services? If not, we really should. In this thought-provoking book, Nick Bostrom lays down a future scenario where machines reach the superintelligent stage and deliberately or accidentally lead to the extinction of humans. This might sound like a sci-fi movie plot, but the way Mr. Bostrom has laid down his arguments and the thinking behind them will definitely sway you and make you take him seriously. We consider this a must-read for everyone working in the AI space. It is a slightly long read, but well worth it in the end. Ray has described the Singularity is breathtaking and will make you stop in your tracks. Once this happens, machines will be far more intelligent than all of the human species combined. This book by Max Tegmark is an international bestseller and deals with the topic of superintelligence. How can we grow our prosperity through automation, without leaving people lacking income or purpose? How can we ensure that future AI systems do what we want without crashing, malfunctioning or getting hacked? Should we fear an arms race in lethal autonomous weapons? Will AI help life flourish as never before, or will machines eventually outsmart us at all tasks, and even, perhaps, replace us altogether? The Master Algorithm This is one of our favorite books in this list. Can there be just one algorithm that deals with all the aspects of technology? Instead of building AI products for specific functions, can we build one single algorithm for all functions? This thought is quite similar to what Albert Einstein spent the latter years of his life trying to discover. Pedro Domingos is a masterful writer, and he deals with the intricacies of his subject extremely well. Make sure you add this to your reading list! The Amazon links in this article are affiliate links. If you buy a book through this link, we would get paid through Amazon. This is one of the ways for us to cover our costs while we continue to create these awesome articles. Further, the list reflects our recommendation based on content of book and is no way influenced by the commission. End Notes This is just the tip of the iceberg. Books are a wonderful source of knowledge for anyone willing to learn from them. This collection spans various aspects of AI and ML â€” from the mathematics and statistics side to the intangible factors like ethics and impact of society. All of these should be considered together when working on an AI and ML project. Having said that, there is truly no substitute for experience. And as always, if you have any questions or suggestions for us on this article, feel free to share them in the comments section below. We look forward to connecting with you!

Chapter 5 : List of artificial intelligence projects - Wikipedia

Algorithms have the ability to solve a particular problem and its value in artificial intelligence depends on their use. Algorithms are a set of repetitive steps that are guided with simple rules to crunch a complex problem.

Overview[edit] Tom M. Mitchell provided a widely quoted, more formal definition of the algorithms studied in the machine learning field: Machine learning tasks[edit] Machine learning tasks are typically classified into several broad categories: The computer is presented with example inputs and their desired outputs, given by a "teacher", and the goal is to learn a general rule that maps inputs to outputs. As special cases, the input signal can be only partially available, or restricted to special feedback. The computer is given only an incomplete training signal: The computer can only obtain training labels for a limited set of instances based on a budget , and also has to optimize its choice of objects to acquire labels for. When used interactively, these can be presented to the user for labeling. No labels are given to the learning algorithm, leaving it on its own to find structure in its input. Unsupervised learning can be a goal in itself discovering hidden patterns in data or a means towards an end feature learning. Here, it has learned to distinguish black and white circles. Another categorization of machine learning tasks arises when one considers the desired output of a machine-learned system: This is typically tackled in a supervised way. Spam filtering is an example of classification, where the inputs are email or other messages and the classes are "spam" and "not spam". In regression , also a supervised problem, the outputs are continuous rather than discrete. In clustering , a set of inputs is to be divided into groups. Unlike in classification, the groups are not known beforehand, making this typically an unsupervised task. Density estimation finds the distribution of inputs in some space. Dimensionality reduction simplifies inputs by mapping them into a lower-dimensional space. Topic modeling is a related problem, where a program is given a list of human language documents and is tasked to find out which documents cover similar topics. Among other categories of machine learning problems, learning to learn learns its own inductive bias based on previous experience. Developmental learning , elaborated for robot learning , generates its own sequences also called curriculum of learning situations to cumulatively acquire repertoires of novel skills through autonomous self-exploration and social interaction with human teachers and using guidance mechanisms such as active learning, maturation, motor synergies, and imitation. History and relationships to other fields[edit] See also: Timeline of machine learning Arthur Samuel , an American pioneer in the field of computer gaming and artificial intelligence , coined the term "Machine Learning" in while at IBM [11]. As a scientific endeavour, machine learning grew out of the quest for artificial intelligence. Already in the early days of AI as an academic discipline, some researchers were interested in having machines learn from data. They attempted to approach the problem with various symbolic methods, as well as what were then termed " neural networks "; these were mostly perceptrons and other models that were later found to be reinventions of the generalized linear models of statistics. Probabilistic systems were plagued by theoretical and practical problems of data acquisition and representation. Their main success came in the mids with the reinvention of backpropagation. The field changed its goal from achieving artificial intelligence to tackling solvable problems of a practical nature. It shifted focus away from the symbolic approaches it had inherited from AI, and toward methods and models borrowed from statistics and probability theory. Relation to data mining[edit] Machine learning and data mining often employ the same methods and overlap significantly, but while machine learning focuses on prediction, based on known properties learned from the training data, data mining focuses on the discovery of previously unknown properties in the data this is the analysis step of knowledge discovery in databases. Data mining uses many machine learning methods, but with different goals; on the other hand, machine learning also employs data mining methods as "unsupervised learning" or as a preprocessing step to improve learner accuracy. Much of the confusion between these two research communities which do often have separate conferences and separate journals, ECML PKDD being a major exception comes from the basic assumptions they work with: Evaluated with respect to known knowledge, an uninformed unsupervised method will easily be outperformed by other supervised methods, while in a typical KDD task, supervised methods cannot be used due to the unavailability of training data. Relation to

optimization[edit] Machine learning also has intimate ties to optimization: Loss functions express the discrepancy between the predictions of the model being trained and the actual problem instances for example, in classification, one wants to assign a label to instances, and models are trained to correctly predict the pre-assigned labels of a set of examples. The difference between the two fields arises from the goal of generalization: According to Michael I. Jordan , the ideas of machine learning, from methodological principles to theoretical tools, have had a long pre-history in statistics. Some statisticians have adopted methods from machine learning, leading to a combined field that they call statistical learning. Computational learning theory A core objective of a learner is to generalize from its experience. The training examples come from some generally unknown probability distribution considered representative of the space of occurrences and the learner has to build a general model about this space that enables it to produce sufficiently accurate predictions in new cases. The computational analysis of machine learning algorithms and their performance is a branch of theoretical computer science known as computational learning theory. Because training sets are finite and the future is uncertain, learning theory usually does not yield guarantees of the performance of algorithms. Instead, probabilistic bounds on the performance are quite common. The biasâ€”variance decomposition is one way to quantify generalization error. For the best performance in the context of generalization, the complexity of the hypothesis should match the complexity of the function underlying the data. If the hypothesis is less complex than the function, then the model has underfit the data. If the complexity of the model is increased in response, then the training error decreases. But if the hypothesis is too complex, then the model is subject to overfitting and generalization will be poorer. In computational learning theory, a computation is considered feasible if it can be done in polynomial time. There are two kinds of time complexity results. Positive results show that a certain class of functions can be learned in polynomial time. Negative results show that certain classes cannot be learned in polynomial time.

Chapter 6 : Benefits & Risks of Artificial Intelligence - Future of Life Institute

The following is a list of current and past, nonclassified notable artificial intelligence projects.

Russian What is AI? Artificial intelligence today is properly known as narrow AI or weak AI , in that it is designed to perform a narrow task e. Why research AI safety? In the long term, an important question is what will happen if the quest for strong AI succeeds and an AI system becomes better than humans at all cognitive tasks. As pointed out by I. Good in , designing smarter AI systems is itself a cognitive task. How can AI be dangerous? Most researchers agree that a superintelligent AI is unlikely to exhibit human emotions like love or hate, and that there is no reason to expect AI to become intentionally benevolent or malevolent. The AI is programmed to do something devastating: Autonomous weapons are artificial intelligence systems that are programmed to kill. In the hands of the wrong person, these weapons could easily cause mass casualties. The AI is programmed to do something beneficial, but it develops a destructive method for achieving its goal: A key goal of AI safety research is to never place humanity in the position of those ants. Why the recent interest in AI safety Stephen Hawking, Elon Musk, Steve Wozniak, Bill Gates, and many other big names in science and technology have recently expressed concern in the media and via open letters about the risks posed by AI , joined by many leading AI researchers. Why is the subject suddenly in the headlines? The idea that the quest for strong AI would ultimately succeed was long thought of as science fiction, centuries or more away. However, thanks to recent breakthroughs, many AI milestones, which experts viewed as decades away merely five years ago, have now been reached, making many experts take seriously the possibility of superintelligence in our lifetime. While some experts still guess that human-level AI is centuries away, most AI researches at the Puerto Rico Conference guessed that it would happen before Since it may take decades to complete the required safety research, it is prudent to start it now. Because AI has the potential to become more intelligent than any human, we have no surefire way of predicting how it will behave. The best example of what we could face may be our own evolution. Timeline Myths The first myth regards the timeline: A common misconception is that we know the answer with great certainty. In fact, history is full of technological over-hyping. AI has also been repeatedly over-hyped in the past, even by some of the founders of the field. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer. All these surveys have the same conclusion: For example, in such a poll of the AI researchers at the Puerto Rico AI conference , the average median answer was by year , but some researchers guessed hundreds of years or more. Many of the safety problems associated with human-level AI are so hard that they may take decades to solve. When Stuart Russell, author of the standard AI textbook , mentioned this during his Puerto Rico talk , the audience laughed loudly. A related misconception is that supporting AI safety research is hugely controversial. After all, fear sells, and articles using out-of-context quotes to proclaim imminent doom can generate more clicks than nuanced and balanced ones. That scenario combines as many as three separate misconceptions: If you drive down the road, you have a subjective experience of colors, sounds, etc. But does a self-driving car have a subjective experience? Does it feel like anything at all to be a self-driving car? If you get struck by a driverless car, it makes no difference to you whether it subjectively feels conscious. The fear of machines turning evil is another red herring. A superintelligent AI is by definition very good at attaining its goals, whatever they may be, so we need to ensure that its goals are aligned with ours. The beneficial-AI movement wants to avoid placing humanity in the position of those ants. Machines can obviously have goals in the narrow sense of exhibiting goal-oriented behavior: If you feel threatened by a machine whose goals are misaligned with yours, then it is precisely its goals in this narrow sense that troubles you, not whether the machine is conscious and experiences a sense of purpose. To cause us trouble, such misaligned superhuman intelligence needs no robotic body, merely an internet connection â€" this may enable outsmarting financial markets, out-inventing human researchers, out-manipulating human leaders, and developing weapons we cannot even understand. Even if building robots were physically impossible, a super-intelligent and super-wealthy AI could easily pay or manipulate many humans to unwittingly do its bidding. What sort of future do you want? Should we develop lethal autonomous

weapons? What would you like to happen with job automation? Do you prefer new jobs replacing the old ones, or a jobless society where everyone enjoys a life of leisure and machine-produced wealth? Further down the road, would you like us to create superintelligent life and spread it through our cosmos? Will we control intelligent machines or will they control us? Will intelligent machines replace us, coexist with us, or merge with us? What will it mean to be human in the age of artificial intelligence? What would you like it to mean, and how can we make the future be that way? Please join the conversation!

Chapter 7 : A beginner's guide to AI: Algorithms

A complete list of all major algorithms (), in any domain. The goal is to provide a ready to run program for each one, or a description of the algorithm. Programming languages include Java, JavaScript and PHP, C, C++ either in direct form or generated from a Scriptol source.

November 3, 1. Also, we will lesrn all most popular techniques, methods, algorithms and searching techniques. We will use Popular Search Algorithms examples and images for the better understanding. Popular Search Algorithms in Artificial Intelligence 2. Introduction to Popular Search Algorithms in AI In the artificial algorithm, to solve the problem we use the searching technique. The search algorithms help you to search for a particular position in such games. Single Agent Pathfinding Problems There are different types of games. Such as 3X3 eight-tile, 4X4 fifteen-tilepuzzles are single-agent-path-finding challenges. As they are consisting of a matrix of tiles with a blank tile. Thus, to arrange the tiles by sliding a tile either vertically or horizontally into a blank space. Also, with the aim of accomplishing some objective. Search Algorithms Terminology a. Problem Space Basically, it is the environment in which the search takes place. A set of states and set of operators to change those states b. Problem Space Graph We use it to represent problem state. Also, we use nodes to show states and operators are shown by edges. The depth of a problem We can define a length of the shortest path. Space Complexity We can calculate it as the maximum number of nodes that are stored in memory. Time Complexity It is defined as the maximum number of nodes that are created. Admissibility We can say it as a property of an algorithm that is used to find always an optimal solution. Branching Factor We can calculate it as the average number of child nodes in the problem space graph. Depth Length of the shortest path from an initial state to the goal state. Hence, it works very smoothly and fine with a small number of possible states. Requirements for Brute Force Algorithms a. A set of valid operators c. Goal state description 6. And continue through neighboring nodes first. Further, moves towards next level of nodes. Moreover, till the solution is found, generates one tree at a time. As this search can be implemented using FIFO queue data structure. This method provides the shortest path to the solution. Disadvantage It consumes a lot of memory space. As each level of nodes is saved for creating next one. Its complexity depends on the number of nodes. It can check duplicate nodes. Breadth-First Search Algorithm 7. As it stands for Last In First Out. Also, implemented in recursion with LIFO stack data structure. Thus, It used to create the same set of nodes as the Breadth-First method, only in the different order. As the path is been stored in each iteration from root to leaf node. Thus, store nodes are linear with space requirement. With branching factor b and depth as m, the storage space is bm. Disadvantage As the algorithm may not terminate and go on infinitely on one path. Hence, a solution to this issue is to choose a cut-off depth. If the ideal cut-off is d, and if the chosen cut-off is lesser than d, then this algorithm may fail. If the chosen cut-off is more than d, then execution time increases. Its complexity depends on the number of paths. It cannot check duplicate nodes. Depth-First Search Algorithm 8. As till both meets to identify a common state. Moreover, initial state path is concatenated with the goal state inverse path. Each search is done only up to half of the total path. Also, always expands the least cost node. Although, it is identical to Breadth-First search if each transition has the same cost. It explores paths in the increasing order of cost. Uniform Cost search must explore them all. As it performs the DFS starting to level 1, starts and then executes a complete depth-first search to level 2. Moreover, we have to continue searching process till we find the solution. We have to generate nodes till single nodes are created. Also, it saves only stack of nodes. As soon as he finds a solution at depth d, the algorithm ends, The number of nodes created at depth d is bd and at depth d-1 is bd Informed Heuristic Search Strategies Algorithm To increase the efficiency of search algorithm we need to add problem-specific knowledge. We use this to solve large problems with large number of possible states a. Heuristic Evaluation Functions We use this function to calculate the path between two states that a function takes for sliding-tiles games. Also, moves of each tile make from its goal state. Further, adding these number of moves for all tiles. Pure Heuristic Search In order, if heuristic value nodes will expand. Also, creates two lists: First, a closed list of the already expanded nodes; Secondly, an open list created. A node with a minimum heuristic value is expanded, In each iteration. Also, all its child

nodes are created and placed on the closed list. Further, we apply this heuristic function to child nodes. Thus, at last, we have to place it in the open list, as per the heuristic value. Although, we have to save the shortest path while to dispose of the longer ones. Also, avoids expensive expanding path. Although, first expands most promising path. Also, we will use priority queue by increasing f n to implement it. Further, using priority queue we implement it. Disadvantage It can get stuck in loops. It is not optimal. Also, a prospective solution. Further, moves to a neighboring solution. Moreover, returns a valid solution. Hence, the algorithm attempts to better solution by single element of the solution. Although, we take an incremental change as a new solution. As if the change produces a better solution. Moreover, we have to repeat until there are no further improvements. As this function of the problem always, returns a state that is a local maximum.

Chapter 8 : The 10 Algorithms Machine Learning Engineers Need to Know

Artificial Intelligence Popular Search Algorithms - Learning Artificial Intelligence in simple and easy steps using this beginner's tutorial containing basic knowledge of Artificial Intelligence Overview, Intelligence, Research Areas of AI, Agents and Environments, Popular Search Algorithms, Fuzzy Logic Systems, Natural Language Processing, Expert Systems, Robotics, Neural Networks, AI Issues.

Yes Informed Heuristic Search Strategies To solve large problems with large number of possible states, problem-specific knowledge needs to be added to increase the efficiency of search algorithms. Heuristic Evaluation Functions They calculate the cost of optimal path between two states. A heuristic function for sliding-tiles games is computed by counting number of moves that each tile makes from its goal state and adding these number of moves for all tiles. Pure Heuristic Search It expands nodes in the order of their heuristic values. It creates two lists, a closed list for the already expanded nodes and an open list for the created but unexpanded nodes. In each iteration, a node with a minimum heuristic value is expanded, all its child nodes are created and placed in the closed list. Then, the heuristic function is applied to the child nodes and they are placed in the open list according to their heuristic value. The shorter paths are saved and the longer ones are disposed. It avoids expanding paths that are already expensive, but expands most promising paths first. It is implemented using priority queue by increasing f n. Greedy Best First Search It expands the node that is estimated to be closest to goal. It is implemented using priority queue. It is not optimal. Local Search Algorithms They start from a prospective solution and then move to a neighboring solution. They can return a valid solution even if it is interrupted at any time before they end. Hill-Climbing Search It is an iterative algorithm that starts with an arbitrary solution to a problem and attempts to find a better solution by changing a single element of the solution incrementally. If the change produces a better solution, an incremental change is taken as a new solution. This process is repeated until there are no further improvements. Local Beam Search In this algorithm, it holds k number of states at any given time. At the start, these states are generated randomly. The successors of these k states are computed with the help of objective function. If any of these successors is the maximum value of the objective function, then the algorithm stops. The pool is then sorted numerically. The highest k states are selected as new initial states. This process continues until a maximum value is reached. When the metal cools, its new structure is seized, and the metal retains its newly obtained properties. In simulated annealing process, the temperature is kept variable. When the temperature is high, the algorithm is allowed to accept worse solutions with high frequency. End Travelling Salesman Problem In this algorithm, the objective is to find a low-cost tour that starts from a city, visits all cities en-route exactly once and ends at the same starting city. Start Find out all n -1! Possible solutions, where n is the total number of cities. Determine the minimum cost by finding out the cost of each of these n -1! Finally, keep the one with the minimum cost.

Chapter 9 : The World's 10 Largest Artificial Intelligence Startups - Algorithm-X Lab

I have always loved the idea of AI and evolutionary algorithms. Unfortunately, as we all know, the field hasn't developed nearly as fast as expected in the early days.

There are multiple approaches that you might take to create Artificial Intelligence, based on what we hope to achieveâ€¦hackernoon. On the contrary, the artificially intelligent agents that work towards a goal, identify the action or series of actions that lead to the goal. The series of actions that lead to the goal becomes the solution for the given problem. Here, the agent has to consider the impact of the action on the future states. Such agents search through all the possible solutions to find the best possible solution for the given problem. This means we need to program the agent, such that it can clearly classify a state as goal. Since there are many ways to reach that goal, the agent should also be able to evaluate a solution and determine its preference for a solution. The state in which the agent starts or initial condition of the agent. All states that are reachable from initial state by any sequence of actions or all possible states that the agent can take. This is also referred to as State space. All possible actions that the agent can execute. Specifically, it provides the list of actions, that an agent can perform in a particular state. This is also referred to as Action space. This property describes the results of each action taken in a particular state. A way to check, whether a state is the goal. A function that assigns a numeric cost to a path w. A search tree is used to model the sequence of actions. It is constructed with initial state as the root. The actions taken make the branches and the nodes are results of those actions. A node has depth, path cost and associated state in the state space. The search space is divided into 3 regions, namely Explored Frontier Unexplored Search involves moving the nodes from unexplored region to the explored region. Strategical order of these moves performs a better search. The moves are also known as node expansion. Picking the order of node expansion has provided us with different search strategies that are suited for different kind of problems. Different search strategies are evaluated along completeness, time complexity, space complexity and optimality. The time and space complexity are measured in terms of: Uninformed Search This type of search does not use any domain knowledge. This means that it does not use any information that helps it reach the goal, like closeness or location of the goal. The strategies or algorithms, using this form of search, ignore where they are going until they find a goal and report success. The basic uninformed search strategies are: It expands the shallowest node node having lowest depth first. It expands deepest node first. It is DFS with a limit on depth. It expands the node with least cost Cost for expanding the node. Informed Search This type of search uses domain knowledge. It generally uses a heuristic function that estimates how close a state is to the goal. This heuristic need not be perfect. This function is used to estimate the cost from a state to the closest goal. The basic informed search strategies are: Greedy search best first search: Minimize the total estimated solution cost, that includes cost of reaching a state and cost of reaching goal from that state. Search Agents are just one kind of algorithms in Artificial Intelligence. Here, an AI has to choose from a large solution space, given that it has a large action space on a large state space. Selecting the right search strategy for your Artificial Intelligence, can greatly amplify the quality of results. This involves formulating the problem, that your AI is going to solve, in the right way. Some of the real world examples, where they are being used are: Example of applications include tools for driving directions in websites, in-car systems, etc. Find the shortest tour to visit each city exactly once. Although, we are just scratching the surface, this article provides you with an outline of what kind of algorithms drive an AI. Depending on the problem, an Artificial Intelligence can use many other algorithms involving Machine Learning, Bayesian networks, Markov models, etc. So, for being more aware of the world of A. If you liked this article, be sure to show your support by clapping for this article below and if you have any questions, leave a comment and I will do my best to answer. Stuart Russell and Peter Norvig.