DOWNLOAD PDF LEARNING IN HUMANS AND MACHINES

Chapter 1 : What is the Difference Between Machine Learning and Human Learning?

"The combination of human and machine is superior to machine alone or human alone," said Lee. Ultimately, the future requirements of cybersecurity are an interplay of advances in technology, legal and human factors, and mathematically verified trust.

Thoughts on Human Learning vs. Machine Learning Posted by Peter Rudin on January in Essay Picture Credit: Public Domain, Jean Marc Cote, Introduction Learning is the act of acquiring new or reinforcing existing knowledge, behaviors, skills or values. Learning does not happen all at once, but it builds upon and is shaped by previous knowledge. To that end, learning may be viewed as a process, rather than a collection of factual and procedural knowledge. Both human as well as machine learning generate knowledge, one residing in the brain the other residing in the machine. This fact raises the question how we apply what kind of knowledge and how we balance these knowledge resources for optimal results. The following discussion hopefully provides some guidance how we can assess this balance keeping in mind that further progress in machine learning and brain research will impact this discussion. Characteristics of Human Learning Motivation is a important for human learning. It provides a hierarchical model for the cognitive procedures and goals of learning divided into 6 levels where level 1 is the most basic level for teaching knowledge acquisition and level 6 the top with the highest educational requirements to meet the goals of a specific educational program. Mastering a specific level is a prerequisite to move on to the next higher level. The levels are defined as follows: Understanding defined as being able to comprehend facts by comparing and interpreting main ideas within the learned material. Applying defined as the ability to use learned material in a new or unprompted way of abstraction and to solve a newly defined problem. Analyzing defined as the ability to examine a problem area and identify the various components breaking the problem down. Evaluating defined as the ability to make judgments based on criteria or standards or to combine parts to form a new concept or idea. Creating defined as the ability to integrate learning from different areas into a plan for solving a problem and to propose alternative solutions. One of the questions raised is how effective what method of learning is and how long the learned material is retained. Research conducted by Igor Kokcharov, Ph. These applications can also be combined with conventional systems as realized in self driving cars or the management of other robotic devices. To generate knowledge different types of machine learning concepts are used: In supervised learning the algorithms make predictions based on a set of examples. For instance, historical stock prices can be used to provide guesses for future prices. Each example used for training is labeled with the value of interestâ€”in this case the stock price. A supervised learning algorithm looks for patterns in those value labels. Instead, the goal of an unsupervised learning algorithm is to organize the data in some way or to describe its structure. This can mean grouping it into clusters or finding different ways of looking at complex data so that it appears simpler or more organized. From that data we might be able to extract knowledge which we were previously unaware of. In reinforced learning the algorithm gets to choose an action in response to each data point. The learning algorithm also receives a reward signal a short time later, indicating how good the decision was. Based on this, the algorithm modifies its strategy in order to achieve the highest reward. Machine learning is complex and difficult to comprehend partially due to the strong expertise required in mathematics and probability theory. One key driver of advancements in machine learning is brain research and new computational models of brain functions which can be used to define new algorithms. The decision how to proceed is still left up to the person responsible for the cure. In an initiative based on its WATSON cognitive machine learning platform IBM in cooperation with a number of health clinics and hospitals has made significant progress to offer such diagnostic services to doctors including a list of the drugs or treatments currently available to cure the problem. Intuition, consciousness and awareness play an important role in leadership decision making, however simple decisions will be more and more outsourced to machine learning knowledge at level 1 to 4 of the learning hierarchy. Consequently humans have to keep their memorizing brain activity in good health in order to be fit for level 5 and 6 knowledge comprehension. Interactive brain-training as provided by a number of companies is possibly one way to improve our learning

capacity. Sport activities, sleep, healthy diets,reading or meditation provide other means to support brain health. Conclusion Learning has been valued by humans for hundreds of years. The invention of printing in the 15th century by Johannes Gutenberg initiated the democratizing of knowledge and information with an ever increasing momentum. With the internet the distribution of knowledge and information has reached a new dimension. As these courses are typically produced by universities and their teaching staff, the quality of the courses are closely linked to the reputation of the university. Educators and policy makers must get involved in the discussion as how to best apply machine learning in combination with human learning. The consequences are far- reaching as continuous brain research will advance machine learning over the years to come.

Chapter 2 : How to get robots to learn the way humans do

Both human as well as machine learning generate knowledge â€" but there's a big difference between the two. Learning is the act of acquiring new or reinforcing existing knowledge, behaviours, skills or values. Humans have the ability to learn, however with the progress in artificial.

Machines are the creation of humans, and they were created to make their work easier. Humans depend more and more on machines for their day-to-day things. Machines have created a revolution, and no human can think of a life without machines. A machine is only a device consisting of different parts, and is used for performing different functions. They do not have life, as they are mechanical. On the other hand, humans are made of flesh and blood; life is just not mechanical for humans. On the other hand, machines have no feelings and emotions. They just work as per the details fed into their mechanical brain. Humans have the capability to understand situations, and behave accordingly. On the contrary, Machines do not have this capability. While humans behave according to their consciousness, machines perform as they are taught. Humans perform activities as per their own intelligence. On the contrary, machines only have an artificial intelligence. It is a man-made intelligence that the machines have. The brilliance of the intelligence of a machine depends on the intelligence of the humans that created it. Another striking difference that can be seen is that humans can do anything original, and machines cannot. Machines have limitations to their performance because they need humans to guide them. Though machines are very sophisticated, they cannot perform anything original. Machines do not have original thoughts. Another thing that has to be noted is that machines are not superior to humans. Machines do not have life, as they are mechanical. On the other hand, humans are made of flesh and blood; life is not mechanical for humans. Humans have feelings and emotions, and they can express these emotions. Machines have no feelings and emotions. Humans can do anything original, and machines cannot. On the contrary, machines do not have this capability. While humans behave as per their consciousness, machines just perform as they are taught. On the contrary, machines only have an artificial intelligence Search DifferenceBetween. If you like this article or our site. Please spread the word.

Chapter 3 : Difference Between Human and Machine | Difference Between

Humans have the ability to learn, however with the progress in artificial intelligence, machine learning has become a resource which can augment or even replace human learning. Learning does not happen all at once, but it builds upon and is shaped by previous knowledge.

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 4 : Machine Learning: What it is and why it matters | SAS

Another lesson from machine learning and its relationship to privileged information is that machines, and therefore humans, often need to perform a learning task themselves in order to understand its use in a different context, such as seeing someone else perform the same task at a different point in time.

When machines become more intelligent, humans are freed to become more creative. That opens doors to completely new possibilities. If the merchant feeds you a buying suggestion that turns out to be way off the mark, it will alter its process in an attempt to feed you better recommendations next time. Pinterest recently unveiled a machine-learning search tool based on image recognition. Judgment is still solidly handled by humans. Computers free us up to concentrate on more creative tasks. Rather than fearing machines, we should look for ways to harness them so we can better innovate. After all, people used to worry about ATMs taking jobs away from bank tellers. But machines count, collect, and distribute money much better than humans do. And by automating cash withdrawals and check deposits, banks were able to reallocate labor toward more complex tasks, like loan originations and advisory services. The fashion industry is on the verge of a similar evolution. The widespread adoption of this technology would free up fashion specialists to provide more personalized guidance to shoppers and employ brilliant seasonal campaigns. A similar trend is emerging in medicine. Already, machine learning is helping doctors make better diagnoses. Eventually, it could help identify high-risk patients and predict readmissions, thus enabling physicians to better tailor their treatment plans. That could reduce health care costs and improve patient outcomes. If this all seems a little too Pollyannaish, consider your everyday online browsing habits. These sites offer customized content you might never find on your own. For creative types, these discoveries often inspire the development of even better content. A budding producer might discover a fun filming technique she had never thought of thanks to YouTube. A young entrepreneur might learn about a new method of office management on LinkedIn. In just about every industry, people will have to embrace non-human collaborators. Supercomputers can make us superhuman. As machines explore the depths of digital data, humans will be freed to innovate far and wide. Marketers who embrace artificial intelligence have extra human capital to spend on artistic ideas and moving ads. Albert Einstein famously, if apocryphally, used an algorithm to solve the problem of what to wear. His goal was to avoid wasting time and brainpower picking out clothes every morning. Algorithms and machine learning free us up to do other things with our time. But the onus is on us to take advantage of all those new opportunities.

Chapter 5 : Thoughts on Human Learning vs. Machine Learning â€º SINGULARITY

Machine Learning for HumansðŸ¤–ðŸ‘¶ Simple, plain-English explanations accompanied by math, code, and real-world examples. This series is available as a full-length e-book!

Media can only be downloaded from the desktop version of this website. Share Leave a comment Children learn language by observing their environment, listening to the people around them, and connecting the dots between what they see and hear. In computing, learning language is the task of syntactic and semantic parsers. These systems are trained on sentences annotated by humans that describe the structure and meaning behind words. Parsers are becoming increasingly important for web searches, natural-language database querying, and voice-recognition systems such as Alexa and Siri. Soon, they may also be used for home robotics. But gathering the annotation data can be time-consuming and difficult for less common languages. To learn the structure of language, the parser observes captioned videos, with no other information, and associates the words with recorded objects and actions. The approach could expand the types of data and reduce the effort needed for training parsers, according to the researchers. A few directly annotated sentences, for instance, could be combined with many captioned videos, which are easier to come by, to improve performance. In the future, the parser could be used to improve natural interaction between humans and personal robots. The parser could also help researchers better understand how young children learn language. This work is part of bigger piece to understand how this kind of learning happens in the world. Visual learner For their work, the researchers combined a semantic parser with a computer-vision component trained in object, human, and activity recognition in video. Semantic parsers are generally trained on sentences annotated with code that ascribes meaning to each word and the relationships between the words. Some have been trained on still images or computer simulations. The new parser is the first to be trained using video, Ross says. In part, videos are more useful in reducing ambiguity. If the parser is unsure about, say, an action or object in a sentence, it can reference the video to clear things up. The researchers compiled a dataset of about videos depicting people carrying out a number of actions, including picking up an object or putting it down, and walking toward an object. Participants on the crowdsourcing platform Mechanical Turk then provided 1, captions for those videos. They set aside video-caption examples for training and tuning, and used for testing. In training, the researchers gave the parser the objective of determining whether a sentence accurately describes a given video. They fed the parser a video and matching caption. The parser extracts possible meanings of the caption as logical mathematical expressions. The algorithm looks at each video frame to track how objects and people transform over time, to determine if actions are playing out as described. In this way, it determines if the meaning is possibly true of the video. Connecting the dots The expression with the most closely matching representations for objects, humans, and actions becomes the most likely meaning of the caption. The expression, initially, may refer to many different objects and actions in the video, but the set of possible meanings serves as a training signal that helps the parser continuously winnow down possibilities. In short, the parser learns through passive observation: To determine if a caption is true of a video, the parser by necessity must identify the highest probability meaning of the caption. The sentence has to be true of the video. Figure out some intermediate representation that makes it true of the video. Given a new sentence, the parser no longer requires videos, but leverages its grammar and lexicon to determine sentence structure and meaning. One day, I can give you a sentence and ask what it means and, even without a visual, you know the meaning. This is the paper I have been waiting for!

Chapter 6 : Machines Learning about Humans Learning about Machine Learning - Anaconda

The Future of Work: Capital Markets, Digital Assets, and the Disruption of Labor Date: Friday, April 27, MODERATOR: Erik Brynjolffson, PhD '91 DISCUSSANT: Daniel Kahneman, Professor of.

Non-associative learning[edit] Non-associative learning refers to "a relatively permanent change in the strength of response to a single stimulus due to repeated exposure to that stimulus. Changes due to such factors as sensory adaptation , fatigue , or injury do not qualify as non-associative learning. Habituation Habituation is an example of non-associative learning in which the strength or probability of a response diminishes when the stimulus is repeated. The response is typically a reflex or unconditioned response. Thus, habituation must be distinguished from extinction , which is an associative process. In operant extinction, for example, a response declines because it is no longer followed by a reward. An example of habituation can be seen in small song birdsâ€”if a stuffed owl or similar predator is put into the cage, the birds initially react to it as though it were a real predator. Soon the birds react less, showing habituation. If another stuffed owl is introduced or the same one removed and re-introduced , the birds react to it again as though it were a predator, demonstrating that it is only a very specific stimulus that is habituated to namely, one particular unmoving owl in one place. The habituation process is faster for stimuli that occur at a high rather than for stimuli that occur at a low rate as well as for the weak and strong stimuli, respectively. Sensitization Sensitization is an example of non-associative learning in which the progressive amplification of a response follows repeated administrations of a stimulus Bell et al. After a while, this stimulation creates a warm sensation that eventually turns painful. The pain results from the progressively amplified synaptic response of the peripheral nerves warning that the stimulation is harmful. Active learning Experiential learning is more efficient than passive learning like reading or listening. Since understanding information is the key aspect of learning, it is important for learners to recognize what they understand and what they do not. By doing so, they can monitor their own mastery of subjects. Active learning encourages learners to have an internal dialogue in which they verbalize understandings. This and other meta-cognitive strategies can be taught to a child over time. Studies within metacognition have proven the value in active learning, claiming that the learning is usually at a stronger level as a result. Conversely, passive learning and direct instruction are characteristics of teacher-centered learning or traditional education. The research works on the human learning process as a complex adaptive system developed by Peter Belohlavek showed that it is the concept that the individual has that drives the accommodation process to assimilate new knowledge in the long-term memory , defining learning as an intrinsically freedom-oriented and active process. In operant conditioning, a behavior that is reinforced or punished in the presence of a stimulus becomes more or less likely to occur in the presence of that stimulus. Classical conditioning The typical paradigm for classical conditioning involves repeatedly pairing an unconditioned stimulus which unfailingly evokes a reflexive response with another previously neutral stimulus which does not normally evoke the response. Following conditioning, the response occurs both to the unconditioned stimulus and to the other, unrelated stimulus now referred to as the "conditioned stimulus". The response to the conditioned stimulus is termed a conditioned response. The classic example is Ivan Pavlov and his dogs. Meat powder is the unconditioned stimulus US and the salivation is the unconditioned response UR. Pavlov rang a bell before presenting the meat powder. The first time Pavlov rang the bell, the neutral stimulus, the dogs did not salivate, but once he put the meat powder in their mouths they began to salivate. After numerous pairings of bell and food, the dogs learned that the bell signaled that food was about to come, and began to salivate when they heard the bell. Once this occurred, the bell became the conditioned stimulus CS and the salivation to the bell became the conditioned response CR. Classical conditioning has been demonstrated in many species. For example, it is seen in honeybees, in the proboscis extension reflex paradigm. In , Watson published the article "Psychology as the Behaviorist Views," in which he argued that laboratory studies should serve psychology best as a science. Observational learning Observational learning is learning that occurs through observing the behavior of others. It is a form of social learning which takes various forms, based on various processes. In humans, this form of learning seems to not need reinforcement

to occur, but instead, requires a social model such as a parent, sibling, friend, or teacher with surroundings. Imprinting psychology Imprinting is a kind of learning occurring at a particular life stage that is rapid and apparently independent of the consequences of behavior. In filial imprinting, young animals, particularly birds, form an association with another individual or in some cases, an object, that they respond to as they would to a parent. In , the Austrian Zoologist Konrad Lorenz discovered that certain birds follow and form a bond if the object makes sounds. Play activity Play generally describes behavior with no particular end in itself, but that improves performance in similar future situations. This is seen in a wide variety of vertebrates besides humans, but is mostly limited to mammals and birds. Cats are known to play with a ball of string when young, which gives them experience with catching prey. Besides inanimate objects, animals may play with other members of their own species or other animals, such as orcas playing with seals they have caught. Play involves a significant cost to animals, such as increased vulnerability to predators and the risk of injury and possibly infection. It also consumes energy , so there must be significant benefits associated with play for it to have evolved. Play is generally seen in younger animals, suggesting a link with learning. However, it may also have other benefits not associated directly with learning, for example improving physical fitness. Through play, children learn social skills such as sharing and collaboration. Children develop emotional skills such as learning to deal with the emotion of anger, through play activities. As a form of learning, play also facilitates the development of thinking and language skills in children. All types of play generate thinking and problem-solving skills in children. Children learn to think creatively when they learn through play. Play as a form of learning, can occur solitarily, or involve interacting with others. Enculturation Enculturation is the process by which people learn values and behaviors that are appropriate or necessary in their surrounding culture. Multiple examples of enculturation can be found cross-culturally. Collaborative practices in the Mazahua people have shown that participation in everyday interaction and later learning activities contributed to enculturation rooted in nonverbal social experience. The collaborative and helpful behaviors exhibited by Mexican and Mexican-heritage children is a cultural practice known as being "acomedido". Episodic learning is so named because events are recorded into episodic memory , which is one of the three forms of explicit learning and retrieval, along with perceptual memory and semantic memory. He would use semantic memory to answer someone who would ask him information such as where the Grand Canyon is. A study revealed that humans are very accurate in the recognition of episodic memory even without deliberate intention to memorize it. Multimedia learning Multimedia learning is where a person uses both auditory and visual stimuli to learn information Mayer This type of learning relies on dual-coding theory Paivio E-learning and augmented learning[edit] Main article: Electronic learning Electronic learning or e-learning is computer-enhanced learning. A specific and always more diffused e-learning is mobile learning m-learning , which uses different mobile telecommunication equipment, such as cellular phones. Augmented digital content may include text, images, video, audio music and voice. By personalizing instruction, augmented learning has been shown to improve learning performance for a lifetime. Moore [34] purported that three core types of interaction are necessary for quality, effective online learning: In his theory of transactional distance, Moore [35] contented that structure and interaction or dialogue bridge the gap in understanding and communication that is created by geographical distances known as transactional distance. Rote learning Rote learning is memorizing information so that it can be recalled by the learner exactly the way it was read or heard. The major technique used for rote learning is learning by repetition, based on the idea that a learner can recall the material exactly but not its meaning if the information is repeatedly processed. Rote learning is used in diverse areas, from mathematics to music to religion. Although it has been criticized by some educators, rote learning is a necessary precursor to meaningful learning. Deeper Learning Meaningful learning is the concept that learned knowledge e. To this end, meaningful learning contrasts with rote learning in which information is acquired without regard to understanding. Meaningful learning, on the other hand, implies there is a comprehensive knowledge of the context of the facts learned. Informal learning Informal learning occurs through the experience of day-to-day situations for example, one would learn to look ahead while walking because of the danger inherent in not paying attention to where one is going. It is learning from life, during a meal at table with parents, play , exploring, etc. The term formal learning has nothing to do with the formality

of the learning, but rather the way it is directed and organized. In formal learning, the learning or training departments set out the goals and objectives of the learning. Nonformal learning Nonformal learning is organized learning outside the formal learning system. For example, learning by coming together with people with similar interests and exchanging viewpoints, in clubs or in international youth organizations, workshops. Nonformal learning and combined approaches[edit] The educational system may use a combination of formal, informal, and nonformal learning methods. The UN and EU recognize these different forms of learning cf. In some schools, students can get points that count in the formal-learning systems if they get work done in informal-learning circuits. They may be given time to assist international youth workshops and training courses, on the condition they prepare, contribute, share and can prove this offered valuable new insight, helped to acquire new skills, a place to get experience in organizing, teaching , etc. Practicing the moves repeatedly helps build " muscle memory " and speed. Thinking critically about moves helps find shortcuts, which speeds future attempts. Revisiting the cube occasionally helps retain the skill. Tangential learning[edit] Tangential learning is the process by which people self-educate if a topic is exposed to them in a context that they already enjoy. For example, after playing a music-based video game, some people may be motivated to learn how to play a real instrument, or after watching a TV show that references Faust and Lovecraft, some people may be inspired to read the original work. According to experts in natural learning, self-oriented learning training has proven an effective tool for assisting independent learners with the natural phases of learning. The built-in encyclopedias in the Civilization games are presented as an example - by using these modules gamers can dig deeper for knowledge about historical events in the gameplay. The importance of rules that regulate learning modules and game experience is discussed by Moreno, C.

Chapter 7 : Machine learning - Wikipedia

Machines that learn language more like kids do. Computer model could improve human-machine interaction, provide insight into how children learn language.

Print Artificial intelligence AI is at the frontier of a new techno-tsunami that is transforming the way we live and work. Could AI be the solution to solving the big data problem, and bridging the widening workforce gap in the Cyber Security industry? Intelligent machines now have the power to make observations, understand requests, reason, draw data correlations, and derive conclusions. Not only could AI help to effectively detect anomalies and tackle manpower shortage, but it could support rapid incident response operations against zero-day threats. Is AI the answer to patching all the flaws in our security systems? Or is it making IT professionals redundant? Beyond the hype, any future-proof business must consider the applications and implications of this incoming wave. The Power of Machine Learning Traditionally, cyber security has relied on rules-based or signature-based pattern matching. With anti-virus AV for example, researchers at AV companies find malware and generate signatures that can be used to check files on an endpoint to see if they match a signature of known malware. This means that one can only detect malware that is known, and that matches a virus definition or signature. With AI, machine learning can provide an alternative to traditional cybersecurity solutions. Instead of relying on code signatures, machines can analyze the behavior of the programme and use machine learning to find a match, where that behavior is predictive of malicious code. Netflix does a great job at classifying movie genres and giving movie recommendations. Through machine learning, service providers like Netflix, are able to automatically categorize and offer suggestions by aggregating across the entire database of films and users. Ability to Detect and Predict New, Complex Threats Conventional technology is past-centric and depends heavily on known attackers and attacks, leaving room for blind spots when it comes to detecting abnormal events in new-age attacks. The limitations of older defense technologies are now being addressed through machine learning. For example, privileged activity within an internal network can be tracked, and any sudden or significant spike in privileged access activity could denote a possible insider threat. If it is found to be a successful detection, the machine will reinforce the validity of the actions and become more sensitive to detecting similar future patterns. With larger amounts of data and examples, machines can better learn and adapt to spotting anomalies, more quickly and accurately. This is especially useful as cyber attacks are becoming increasingly sophisticated, and hackers are coming up with new and innovative approaches, of which older security technologies would be slow to detect. Ease Burden on Cybersecurity Personnel Machine learning is most effective as a tool when it has access to a large pool of data to learn and analyze from, reducing attack surfaces through predictive analytics. The volume of security alerts that appear daily can be very overwhelming for the security team. Automating threat detection and response helps lighten the load off of cybersecurity professionals who have to contend with prioritizing cybersecurity-related issues and can aid the detection of threats more efficiently than other software-driven methods. As substantial quantities of security data are being generated and transferred over networks every day, it becomes progressively difficult for cybersecurity experts to monitor and identify attack elements rapidly and reliably. This is where AI can come in and expand their monitoring and detecting operations, making sense of the copious data. Machine learning can help cybersecurity personnel respond to scenarios that they have not specifically encountered before, replacing the laborious process of human analysis. AI and machine learning also assist IT security professionals in achieving good cyber hygiene and enforces robust cybersecurity practices. The tables are turned as cybersecurity becomes less about an incessant pursuit of hunting down malicious activity, and more about continuous prevention, prediction, and improvement. It could also become a part of the solution for the widening talent gap in the cybersecurity industry. Limitations of AI and machine learning One of the greatest challenges would be the adoption of AI technology. For a machine learning engine to perform well, it must retrieve the right data, extract the correct features, and cast the appropriate angle on those features. If trained poorly, it will make inaccurate predictions. Such models are only as good as the data that is fed in. Companies who only do end-point detection are missing out as they lack

the data required to leverage on AI. Bad actors could significantly develop their phishing attacks by using AI to circumvent machine learning-based phishing detection systems. In an experiment by Cyxtera, two attackers were able to use AI to improve their phishing attack effectiveness from 0. There is a human attacker behind these threats. Many cybersecurity experts have bold opinions on whether machines should be responsible to manage something as complicated as cybersecurity. Only a human can understand the business context of why an attacker might be after a piece of information and what their motivations are. Machine learning is an effective tool against both known and unknown malware, as it can identify and understand malicious activity when applied properly. However, it should not be the only solution. Ultimately, the future requirements of cybersecurity are an interplay of advances in technology, legal and human factors, and mathematically verified trust. Effective cybersecurity should be about striking a balance between human and machines. Where computers cannot, humans make sense of the data by ensuring machine-suggested actions have business value too. Humans bring the business, legal, and commercial value into decisions, whilst machines have the capacity and speed to analyze and interpret big chunks of data. Both human intelligence and artificial intelligence must work symbiotically for optimal results.

Chapter 8 : Detecting fake face images created by both humans and machines

The book discusses the analysis, comparison and integration of computational approaches to learning and research on human learning. Learning has for some time been an issue of minor importance in the cognitive sciences.

CC0 Public Domain Matthew Hutson, a freelance writer, has published a Feature article in the journal Science outlining progress in getting computers to learn and to think more like human beings. In his article, he suggests that a lot of problems will need to be solved before machines can learn to think the way people do. And at its root, he suggests, it will require figuring out how to get computers to learn both by trial and error and through baked-in features that correspond to instinct. As Hutson notes, deep learning network-based systems have accomplished extraordinary things, such as beating humans at very difficult games or learning to flip hamburgers. But they still lack the ability to apply what they have learned to new and different environments. A chess-playing robot able to beat the best human player in the world would still lose when asked to play a game of checkers with a child. They also lack common sense. As a simple example, Hutson cites asking a robot butler to retrieve a red cup from the cupboardâ€”how should it respond if there are no red cups in the cupboard? Instead, the cabinet is stocked with cups of other colors and plates that are red. A human would most likely choose a cup of another color rather than a red plate that partially matches the request, because she would understand the intended use of the required object. But how do we get robots to do that? To get robots to be more able to handle real-world, random scenarios will likely involve getting them to learn the way humans or other animals learn. This will likely will require going back to the original designâ€”the human brain. Marcus is a developmental cognitive scientist involved in research aimed at studying how humans learn from birth onward. Researchers like Marcus, Hutson suggests, are searching for the means by which humans and other animals are endowed with instinctual behavior. Indeed, Marcus has come up with a list of human instincts that he believes will need to be baked into computers before they ever learn things like causality and assessing cost-benefit situations. Hutson notes that some computer scientists are jumping on such new ideas and he lists companies like Vicarious, in California, and DeepMind in England, which are hard at work trying to implement them. He also cites ongoing research efforts in places like MIT and the University of New South Wales, where teams are trying to learn how the human brain works and how machines can be made to function the same way. Explore further More information: Basic instincts, Science How to get robots to learn the way humans do , May 29 retrieved 12 November from https: Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.

Chapter 9 : Machines that learn language more like kids do | MIT News

Outline 1. Machine Learning and Human Learning 2. Aligning specific results from ML and HL â€¢ Learninggp to predict and achieve rewards â€¢ TD learning â†”Dopamine system in the brain.