Programme Educational Objectives(PEOs)

PEO1 (Core Competency) : Graduates will acquire a strong foundation in mathematical, scientific and engineering fundamentals necessary to formulate, solve and analyze Computer Science and Engineering problems.

PEO2 (Professionalism) : Graduates will practice the profession with ethics, integrity and leadership to relate engineering to global perspective issues and social context.

PEO3 (Higher Studies and Entrepreneurship) : Graduates will be prepared for their careers in the software industry or in higher studies leading to research and for applying the spirit of innovation and entrepreneurship in their career and continuing to develop their professional knowledge on a life long basis.

Programme Outcomes(POs)

PO1: Engineering knowledge: Ability to apply the knowledge of mathematics, physical sciences and computer science and engineering specialization to the solution of complex engineering problems.

PO2: Problem analysis: Ability to identify, formulate and analyze complex real life problems in order to provide meaningful solutions by applying knowledge acquired in computer science and engineering.

PO3: Design/development of solutions: Ability to design cost effective software / hardware solutions to meet desired needs of customers/clients.

PO4: Conduct investigations of complex problems: Use research-based knowledge and research methods including design of experiments, analysis and interpretation of data, and synthesis of the information to provide valid conclusions in the field of computer science and engineering.

PO5: Modern tool usage: Create, select and apply appropriate techniques, resources and modern computer science and engineering tools including prediction and modeling to complex engineering activities with an understanding of the limitations.

PO6: The engineer and society: Apply reasoning informed by the contextual knowledge to assess societal, health, safety, legal and cultural issues and the consequent responsibilities relevant to the professional engineering practice.

PO7: Environment and sustainability: Understand the impact of the professional engineering solutions in societal and environmental contexts, and demonstrate the knowledge of, and need for sustainable development.

PO8: Ethics: Apply ethical principles and commit to professional ethics and responsibilities and norms of the engineering practice.

PO9: Individual and team work: Function effectively as an individual, and as a member or leader in diverse teams, and in multidisciplinary settings.

PO10: Communication: Communicate effectively on complex engineering activities with the engineering community and with society at large, such as, being able to comprehend and write effective reports and design documentation, make effective presentations, and give and receive clear instructions.

PO11: Project management and finance: Demonstrate knowledge and understanding of the engineering and management principles and apply these to one's own work, as a member and leader in a team, to manage projects and in multidisciplinary environments.

PO12: Life-long learning: Recognize the need for, and have the preparation and ability to engage in independent and life-long learning in the broadest context of technological change.

Programme Specific Outcomes (PSOs)

PSO1: Software System Design and Development: The ability to apply software development life cycle principles to design and develop the application software that meet the automation needs of society and industry.

PSO2: Computing and Research ability: The ability to employ modern computer languages, environments and platforms in creating innovative career paths in SMAC (Social, Mobile, Analytics and Cloud) technologies.

K S R INSTITUTE FOR ENGINEERING AND TECHNOLOGY
K.S.R KALVI NAGAR, TIRUCHENGODE - 637 215
DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING
TECH ZEST
Fuel Your Mind
Jan - 2019 - Volume VI / Issue II

Machine Learning / Deep Learning

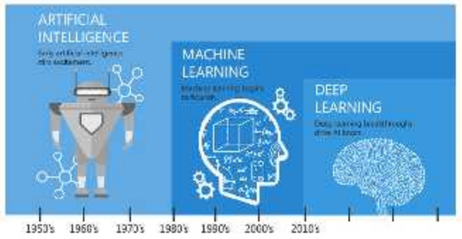

K S R Institute for Engineering and Technology

Vision

To become a globalLy recognized Institution in Engineering Education, Research and Entrepreneurship.

Mission

IM1: Accomplish quality education through improved teaching learning process.
IM2: Enrich technical skills with state of the art laboratories and facilities.
IM3: Enhance research and entrepreneurship activities to meet the industrial and societal needs

Department of Computer Science and Engineering

Vision

To produce globally competitive Computer Science Engineers and Entrepreneurs with moral values.

Mission

DM1 (Quality Education) : Provide quality education to enhance problem solving skills, leadership qualities, team spirit and ethical responsibilities.

DM2 (State of art Laboratory) : Enable the students to adapt to the rapidly changing technologies by providing advanced laboratories and facilities.

DM3 (Research and Development) : Promote research based activities in the emerging areas of techno-environment in order to meet industrial and societal needs.

Machine Learning

The term Machine Learning was coined by Arthur Samuel in 1959, an American pioneer in the field of computer gaming and artificial intelligence and stated that "it gives computers the ability to learn without being explicitly programmed", and in 1997, Tom Mitchell gave a "well-posed" mathematical and r xelational definition that, "A computer program is said to learn from experience E with respect to some task T and some performance measure P, if its performance on T, as measured by P, improves with experience E".

Machine learning is the subset of artificial intelligence (AI) that focuses on building systems that learn—or improve performance—based on the data they consume. Artificial intelligence is a broad term that refers to systems or machines that mimic human intelligence. Machine learning and AI are often discussed together, and the terms are sometimes used interchangeably, but they don't mean the same thing. An important distinction is that although all machine learning is AI, not all AI is machine learning.

Classification of Machine Learning

Machine learning implementations are classified into three major categories, depending on the nature of the learning "signal" or "response" available to a learning system which are as follows:-

Supervised learning - Supervised learning, the algorithm is trained by a dataset that is already labeled and has a predefined output

Unsupervised learning - Unsupervised machine learning involves training based on data that does not have labels or a specific, defined output.

Reinforcement learning - It is about taking suitable action to maximize reward in a particular situation. In reinforcement learning, there is no answer but the reinforcement agent decides what to do to perform the given task. In the absence of training dataset, it is bound to learn from its experience.

Semi-supervised learning - Semi-supervised learning is a class of machine learning tasks and techniques that also make use of unlabeled data for training – typically a small amount of labeled data with a large amount of unlabeled data. Semi-supervised learning falls between unsupervised learning and supervised learning.

S.Devi gayathri , III-CSE

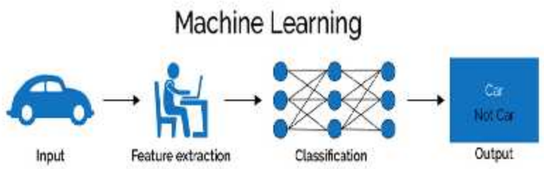

Deep Learning

Deep learning is a subset of machine learning in which the tasks are broken down and distributed onto machine learning algorithms that are organized in consecutive layers. Each layer builds up on the output from the previous layer. Together the layers constitute an artificial neural network that mimics the distributed approach to problem-solving carried out by neurons in a human brain. In other words, it mirrors the functioning of our brains. Deep learning algorithms are similar to how nervous system structured where each neuron connected each other and passing information.

Deep learning models work in layers and a typical model at least have three layers. Each layer accepts the information from previous and passes it on to the next one.

Traditional machine learning models have always been very powerful to handle structured data and have been widely used by businesses for credit scoring, churn prediction, consumer targeting, and so on.

The success of these models highly depends on the performance of the feature engineering phase: the more we work close to the business to extract relevant knowledge from the structured data, the more powerful the model will be.

When it comes to unstructured data (images, text, voice, videos), hand engineered features are time consuming, brittle and not scalable in practice. That is why Neural Networks become more and more popular thanks to their ability to automatically discover the representations needed for feature detection or classification from raw data. This replaces manual feature engineering and allows a machine to both learn the features and use them to perform a specific task.

One of differences between machine learning and deep learning model is on the feature extraction area. Feature extraction is done by human in machine learning whereas deep learning model figure out by itself.

K.Dharani, III-CSE

The $127 billion autonomous vehicle market is being driven by AI.

Student Article

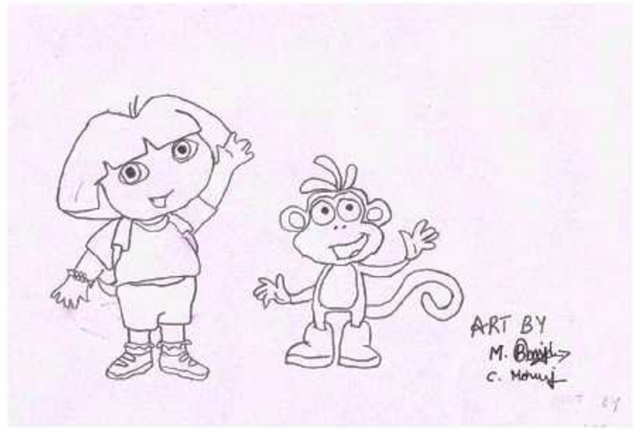

Q1: You have a set of thirty six cards. The cards are six in color (six each) and each color is numbered from 1 to 6. You draw two cards at random. What is probability that they are of a different color and have a different number?

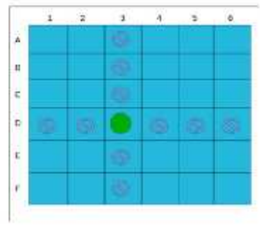

Q2: Which of the following statement(s) correctly represents a real neuron?

A. A neuron has a single input and a single output only
B. A neuron has multiple inputs but a single output only
C. A neuron has a single input but multiple outputs
D. A neuron has multiple inputs and multiple outputs
E. All of the above statements are valid

Q3: "Convolutional Neural Networks can perform various types of transformation (rotations or scaling) in an input". Is the statement correct True or False?

A. True
B. False

Q4: In which neural net architecture, does weight sharing occur?

A. Convolutional neural Network
B. Recurrent Neural Network
C. Fully Connected Neural Network
D. Both A and B

Q5: For an image recognition problem (recognizing a cat in a photo), which architecture of neural network would be better suited to solve the problem?

A. Multi Layer Perceptron
B. Convolutional Neural Network
C. Recurrent Neural network
D. Perceptron

R.Babu, II-CSE

The first AI program, "The Logic Theorist," was developed in 1955 by Newell and Simon.

Machine Learning Applications in Retail

Machine learning in retail is more than just a latest trend, retailers are implementing big data technologies like Hadoop and Spark to build big data solutions and quickly realizing the fact that it's only the start. They need a solution which can analyse the data in real-time and provide valuable insights that can translate into tangible outcomes like repeat purchasing. Machine learning algorithms process this data intelligently and automate the analysis to make this supercilious goal possible for retail giants like Amazon, Target,

According to The Realities of Online Personalisation Report, 42% of retailers are using personalized product recommendations using machine learning technology. It is no secret that customers always look for personalized shopping experiences, and these recommendations increase the conversion rates for the retailers resulting in fantastic revenue. Retailers mine customer actions, transactions, and social date to identify customers who are at a high risk of switching to a competitor. This information is then combined with profitability data so that they can optimize their next best action strategies and personalize end-to-end shopping experience for the customer.

M.Balaji , IV-CSE

Automatically Adding Sounds to Silent Movies

In this task the system must synthesize sounds to match a silent video. The system is trained using 1000 examples of video with sound of a drum stick striking different surfaces and creating different sounds. A deep learning model associates the video frames with a database of pre-rerecorded sounds in order to select a sound to play that best matches what is happening in the scene. The system was then evaluated using a turing-test like setup where humans had to determine which video had the real or the fake (synthesized) sounds. A very cool application of both convolutional neural networks and LSTM recurrent neural networks. A recurrent neural network to predict sound features from videos and then produces a waveform from these features with an example-based synthesis procedure. The system shows that the sounds predicted by our model are realistic enough to fool participants in a "real or fake" psychophysical experiment, and that they convey significant information about material properties and physical interactions.

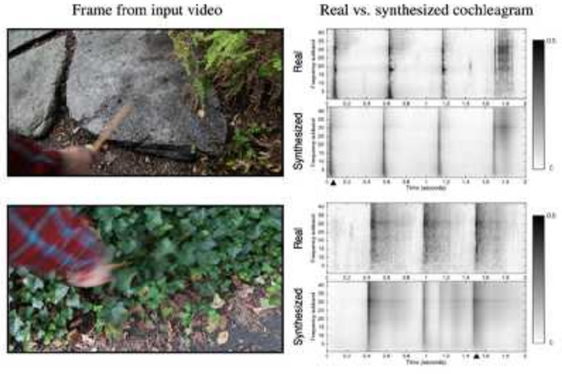

In order to study visually indicated sounds, we collected a dataset containing videos of humans (the authors) probing environments with a drumstick – hitting, scratching, and poking different objects in the scene. We chose to use a drumstick so that we would have a consistent way of generating the sounds. Moreover, since the drumstick does not occlude much of a scene, we can also observe what happens to the object after it is struck. This motion, which we call a reaction. Unlike traditional object- or scene-centric datasets, such as ImageNet or Places where the focus of the image is a full scene, our dataset contains close-up views of a small number of objects.

S.Laxmancibi, III-CSE

Google has a Quantum Artificial Intelligence Lab with an actual quantum computer.

Automatic Machine Translation

This is a task where given words, phrase or sentence in one language, automatically translate it into another language. Automatic machine translation has been around for a long time, but deep learning is achieving top results in two specific areas:

- Automatic Translation of Text.
- Automatic Translation of Images.

Text translation can be performed without any preprocessing of the sequence, allowing the algorithm to learn the dependencies between words and their mapping to a new language. Stacked networks of large LSTM recurrent neural networks are used to perform this translation. As you would expect, convolutional neural networks are used to identify images that have letters and where the letters are in the scene. Once identified, they can be turned into text, translated and the image recreated with the translated text. This is often called instant visual translation.

Google Translate is a free multilingual machine translation service developed by Google, to translate text. It offers a website interface, mobile apps for Android and iOS, and an API that helps developers build browser extensions and software applications. Google Translate supports over 100 languages at various levels and as of May 2017, serves over 500 million people daily. Various functions of google translate are Written Words Translation, Website Translation, Document Translation, Speech Translation, Mobile App Translation, Image Translation and Handwritten Translation. For some languages, Google Translate can pronounce translated text. For some languages, text can be entered via an on-screen keyboard, through handwriting recognition, or speech recognition. Google Translate does not apply grammatical rules, since its algorithms are based on statistical analysis rather than traditional rule-based analysis.

P.N.Koteeshwaran, II-CSE

AI, Internet of Things, Blockchain and cybersecurity are expected to be key trends in 2019

It faces threats from rising protectionism, data flow curbs and fast-changing technological shifts, but the Indian IT industry is keeping its hopes high for the new year with plans afoot for big investments in automation and artificial intelligence.

For the industry body Nasscom, 2018 has been the year of 'Digital at Scale' as IT firms focussed on leveraging new technologies and ensuring sustainability by creating right skills with help from innovation, policies and partnerships.

The year ahead is "punctuated with several transformative opportunities", Nasscom President Debjani Ghosh said.

The industry body has projected exports to grow at 7-9 per cent for 2018-19, almost same as the previous fiscal, but domestic revenue may grow faster at 10-12 per cent and this may make the new year transformative with overseas funds accounting for a lion's share so far.

A recent TeamLease survey said Indian IT sector is geared up to add around 250,000 new jobs in 2019. Automation, AI and robotics are expected to reshape workplace culture and skill requirements with big data analytics, machine learning and AI developers are turning out to be the highest paid areas.

"With bots taking up more mundane tasks, employees are pushed to think 'value' even more - and this has given rise to a set of different skill sets that are required in business as of today," Genpact CHRO Piyush Mehta said.

Mehta said there is also a massive untapped workforce in the non-IT sectors such as healthcare, banking or manufacturing.

According to Nasscom, over 1,200 new advanced technology startups got added to the ecosystem in 2018 with data analytics being the largest contributor and the startup momentum will continue in 2019.

The debate on whether AI will create jobs or take away more roles remains open, but the industry is hoping for bigger deals, stronger growth and better margins in the new year.

R.S.Vasu Mitha, III-CSE

CAPTCHA is short for Completely Automated Public Turing test to tell Computers and Humans Apart

Language Corner

Machine learning is a process to build AI enabled algorithms with which machines are able to learn or produce codes automatically through analyzing the given data. Machine learning is the subset of Artificial Intelligence and again has the intersection with many fields including math and psychology. The following 3 languages are used by the machine learning algorithms such as,

1. R Language
2. Python
3. C language

R Language

R language is combined with lexical scooping, which tends to provide the flexibility in producing statistical models. R is a really powerful language to start with machine learning, as there are many specified GNU packages available. One can surely choose to use R for creating powerful algorithms and plus the R studio has an easy statistical visualisation of your algorithms. Though the language is widely used in academic research and gaining really well recognition in the industry use most recently.

Python

Python language is one of the most flexible languages and can be used for various purposes. Python has gained huge popularity base of this. Python does contain special libraries for machine learning namely scipy and numpy which great for linear algebra and getting to know kernel methods of machine learning. The language is great to use when working with machine learning algorithms and has easy syntax relatively.

C language

The mother of all language is definitely a great programming language to build your predicative algorithms. This language is not a cakewalk and should be only be considered when you have strong fundamentals of computer science and programming languages, however, once you are proficient in C language then there is nothing that can stop you [illegible]

General-Purpose Machine Learning

Darknet - Darknet is an open source neural network framework written in C and CUDA. It is fast, easy to install, and supports CPU and GPU computation.

Recommender - A C library for product recommendations/suggestions using collaborative filtering (CF).

Hybrid Recommender System - A hybrid recommender system based upon scikit-learn algorithms.

neonrvm - neonrvm is an open source machine learning library based on RVM technique. It's written in C programming language and comes with Python programming language bindings.

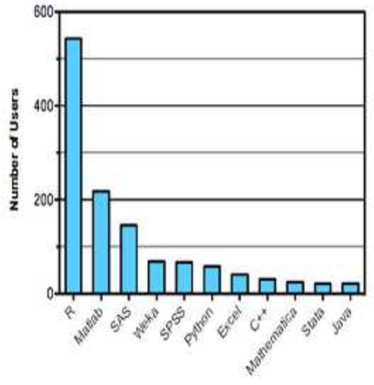

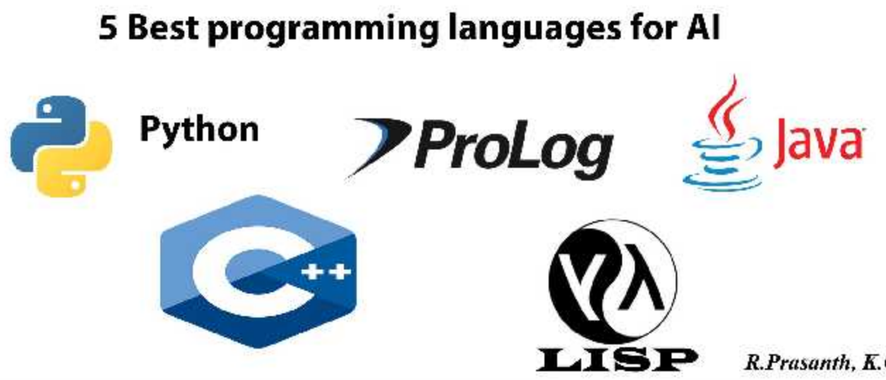

R.Prasanth, K.Gokul, IV-CSE

Machine Learning Applications in Healthcare

Doctors and medical practitioners will soon be able to predict with accuracy on how long patients with fatal diseases will live. Medical systems will learn from data and help patients save money by skipping unnecessary tests. Radiologists will be replaced by machine learning algorithms. Computers and Robots cannot replace doctors or nurses, however the use of life-saving technology (machine learning) can definitely transform healthcare domain. When we talk about efficiency of machine learning, more data produces effective results – and the healthcare industry is residing on a data goldmine.

Personalized treatment has great potential for growth in future, and machine learning could play a vital role in finding what kind of genetic makers and genes respond to a particular treatment or medication. Personalized medication or treatment based on individual health records paired with analytics is a hot research area as it provides better disease assessment. In future, increased usage of sensor integrated devices and mobile apps with sophisticated remote monitoring and health-measurement capabilities, there would be another data deluge that could be used for treatment efficacy. Personalized treatment facilitates health optimization and also reduces overall healthcare costs.

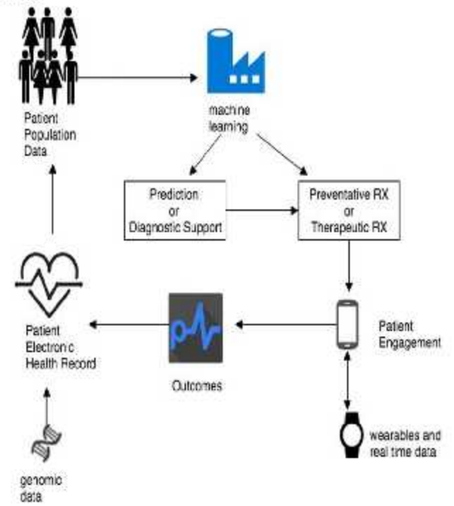

K.Kailash, II-CSE

Search Engine Result Refining

Google and other search engines use machine learning to improve the search results for you. Every time you execute a search, the algorithms at the backend keep a watch at how you respond to the results. If you open the top results and stay on the web page for long, the search engine assumes that the results it displayed were in accordance to the query. Similarly, if you reach the second or third page of the search results but do not open any of the results, the search engine estimates that the results served did not match requirement. This way, the algorithms working at the backend improve the search results. The following diagram represents information retrieval process.

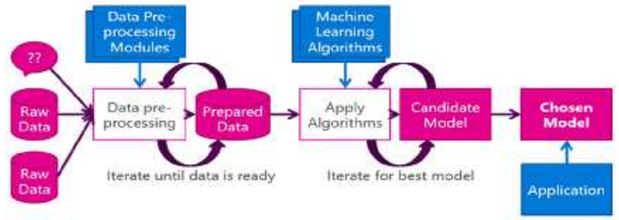

Much of the work in information retrieval can be automated. Processes such as document indexing and query refinement are usually accomplished by computer, while document classification and index term selection are more often performed manually. However, manual development and maintenance of document databases is time-consuming, tedious, and error-prone. Algorithms that "mine" documents for indexing information, and model user interests to help them formulate queries, reduce the workload and can ensure more consistent behavior.

N.Arun Kumar, II-CSE