



Vidya Vikas Education Trust's

Universal College of Engineering

Approved by AICTE, DTE, Maharashtra State Government and Affiliated to Mumbai University
Accredited with B+ Grade by NAAC | Recognised as a Linguistic (Gujarati) Minority Institution

COFFEE & CODE;

An initiative by the Department of Computer Engineering.

CONTENTS

Page 2 -3 : Faculty Activities and Achievements

Page 4 : How Shutterstock finds the perfect photos ?

Page 6 : How Google aims to dominate AI ?

Page 8 : Cover Story

VISION

To be recognized as a department that provides quality technical education and research opportunities that eventually caters to helping and serving the community.

MISSION

- To groom the students to participate in curricular and co-curricular activities by providing efficient resources.
- To motivate the students to solve real-world problems to help the society grow.
- To provide a learning ambience to enhance innovations, team spirit and leadership qualities for students.



“Arise, awake and do not stop until the goal is reached.”

--Swami Vivekananda

In Association with



(Computer Engineering Student Association)

Team : Mr. Ayush Shetty,
Mr. Tejas Gudulekar
Faculty : Mr. Sridhar Iyer,
Mrs. Hezal Lopes,
Ms. Apurva Chaudhari

Faculty Activities and Achievements

One Week Short Term Training Program (STTP) on “Natural Language Processing” In Association with All India Council for Technical Education (AICTE) and ISTE

Mrs. Kanchan Dabre of the Department of Computer Engineering of the Universal College of Engineering had attended a 1 week Short Term Training Program (STTP) on “**Natural Language Processing-Theory and Practices**” at the St. Francis College of Engineering, Borivali.

This One Week AICTE–ISTE approved Short Term Training Programme on “**Natural Language Processing-Theory and Practices**” was held from **2nd December to 6th December 2019**.

The STTP included a number of special sessions organized by recognized experts. The scope of the STTP included basic theory and recent advances in NLP, applications and practical implementation. The proposed STTP on Natural Language Processing-Theory and Practices provided theoretical and practical hands on experience to the participants. The STTP also included discussion on NLP (DLOC-I) Syllabus designed by MU for Final year computer Engineering students.

Our institute believes in promoting quality education by sponsoring teaching faculty to attend such events.



About the Program

Natural language processing (NLP) addresses the problems of automated understanding and generation of natural human language. NLP is an important area of artificial intelligence concerned with the processing and understanding (NLU) of a human language. The goal of NLP and NLU is to process and harness information from a large corpus of text with very little manual intervention. This program focuses on the recent advancements in NLP and its application, which facilitates human machine communication using natural languages

Outcomes:

Following outcomes will be achieved after successful completion of the course:

1. Awareness of NLP from basic to advance.
2. Practical exposure to emerging concepts of NLP like machine translation, semantic analysis, etc.
3. Exposure and practice on current trends and emerging tools of NLP.
4. Generation of new ideas in NLP domain towards solution of real world problems.


Interactive Sessions on
"Education- An Opportunity" & "Menstrual Cycle and Hygiene"
 In Association with
Women Development Cell , UCoE

Mrs Kanchan Dabre had delivered a session for school girls at Nagla Zilla Parishad (ZP) school on "Education an Opportunity" where different examples of inspirational people were demonstrated to the students and also how they used to study and achieve the accolades, how they used to face the tough times, the struggles and overcame the difficulties and made remarkable impressions on the Indian History were also discussed. Also other WDC member Mrs Sonal Borse spread awareness on "Menstrual Hygiene and Yoga ". The Session got over with a Street Play by the UCOE NSS Team followed by sweet distribution and interaction with the parents.



NPTEL SWAYAM CERTIFICATION
"Completion of 12 week online course" on "Deep Learning- Part 1"
 In Association with
NPTEL, SWAYAM and IIT Madras


Mrs. Kanchan Dabre of the Department of Computer Engineering has added another feather to her cap by completing a 12 week course with **71 %** overall score on "Deep Learning- Part 1" in association with the NPTEL as an initiative by the various participating IIT's and IIIT's across India. These short term courses act as both refresher courses and fundamental courses for the overall development of the Teachers as well as the students.



Elite

NPTEL Online Certification

(Funded by the Ministry of HRD, Govt. of India)



This certificate is awarded to
KANCHAN RUFUS DABRE
 for successfully completing the course
Deep Learning - Part 1

with a consolidated score of **71 %**

Online Assignments	19.38/25	Proctored Exam	52.04/75
--------------------	----------	----------------	----------

Total number of candidates certified in this course: **375**


Devendra Jalihal

Prof. Devendra Jalihal
Chairman
 Centre for Continuing Education, IITM


Jul-Oct 2019
(12 week course)

Prof. Andrew Thangaraj

Prof. Andrew Thangaraj
NPTEL Coordinator
 IIT Madras



Indian Institute of Technology Madras



swayam

Roll No: NPTEL19CS85S31820504

To validate and check scores: <https://nptel.ac.in/noc>

Shutterstock Has Trained A Computer To Find You The Perfect Photos

Computer vision can show you images you're actually looking for



Computer vision technology will help you find the best stock images of whatever strikes your fancy. Like the classic "woman eating salad."

I'm picturing a bike next to a fence. It's in a European city somewhere, with narrow cobblestone streets, and the fence is in front of an old-looking brick building. The bike is shiny and blue, with a basket, sort of old fashioned. You can't see the sky, but you can tell it's a somewhat sunny day.

There's no way I could possibly find a picture of a scene like this one on the Internet. Sure, I can type in keywords like "blue bike next to fence in Europe" and it will show me some results that are tangentially related if I'm lucky. My chances are slightly better if I happen to have such an image already at my disposal—that way, I can do a reverse image search and can crawl across sites not limited to English. But oftentimes the results will seem weird, with the wrong feeling or missing key components of the scene in my head.

Computers still can't read our minds. But stock image website Shutterstock has created a whole new way to categorize images. The company's new tool, which launches today on their web site, is one of many innovations in a recent but rapidly growing field called computer vision. And Shutterstock is hoping that it can transform your frustrating process of matching the image in your head to the one on your screen into something that's actually fun.


A PICTURE IS WORTH A THOUSAND WORDS

It's hard to find the right images online because most search engines rely on keywords. If a user is uploading that bike image to Shutterstock's web site, for example, she provides all the keywords. If she's uploading a batch of images that are similar, some of those keywords might not pertain to each individual image.

"All of these keywords together can be strange—that's one of the problems that's inherent when you treat media like a bag of words," says Kevin Lester, the vice president of engineering for search and discovery at Shutterstock, one of the engineers behind the new computer vision tool.

So a lot of image databases fill in those gaps with user behavior. If people searching the words "bike" and "fence" download a particular image more often, that one probably contains those two things in it. It's a simple concept, Lester says, but it's still imperfect.

Computer vision can change all that by eliminating the need for keywords in the first place. Using a series of algorithms, a model can progressively survey each pixel in an image to pick out different features in it—the color, the shapes, the sharpness of the angles. Each calculation is a layer of the deep learning network. At the end of this process, the program generates a single number, a vector. If the model is good, the more similar the number, the more similar the images they quantify. The model trains itself to recognize these features, so the more images plugged into it, the better the model becomes.



COMPUTER VISION

As a field, computer vision has really only been around since 2012, when three researchers from the University of Toronto published a paper that has since been considered a watershed moment for the discipline.

And yet, in just four years, computer vision is crucial for a number of tech companies. Facebook's model can identify faces in pictures with more than 97 percent accuracy; Google's can solve those CAPTCHA puzzles—designed to weed out robots to verify that a user is human—with 99 percent accuracy.

A model for computer vision can be used for a number of different applications, but it's usually trained for a particular task. Shutterstock is using it to detect visually similar images and do a reverse image search.


SEEING LIKE A COMPUTER

One of the main ways people discover images on Shutterstock's web site is in this category called visually similar. It's those images that come up at the bottom when you click on one. Like this:


If the system is relying on keywords, the images it returns are sometimes related, but sometimes not. It's inconsistent and spotty. For Shutterstock's first computer vision model, the engineers used the schematic first outlined in that 2012 paper and trained on the site's 70 million stock images. Even then, it wasn't very good.



"I don't think anyone would consider those extremely similar, other than that the color palette seems to be somewhat consistent," says Lester. The engineers tweaked the model, then gave it weeks to retrain on the data to learn about particular features of images. And it got a little better:



There were a few more iterations, but here are the results turned up in the final version of the tool:



Through internal testing, Shutterstock was able to say that their new visually similar tool was significantly better than their past one that relied on keywords. Now, every time someone clicks on an image on their site (and that happens a lot—the company sells 4.7 images per second), the algorithm searches through the 70 million photos to serve up those it deems most similar. The site also uses the tool on its 4 million film clips, a growing area of business for the company.

Importantly, it does this search in just 200 milliseconds—that's half the time it took from the company's old model. And while a difference of 200 milliseconds might not sound like much, Lester says it makes a world of difference to impatient customers. "When we quicken the speed, we found that people search more, because what we did was reduce the cost of them doing the search, which means they were exploring our site more. And that in turn meant that they were more likely to sign up as customers," he says.

THE LIMIT DOES NOT EXIST

As models like Shutterstock's improve, engineers run into semantic or philosophical issues. At a certain point, Lester says, people differ in how they define images as similar—that's when he'll know that his team can stop improving its model, he says. Then there's the inevitability of offending someone, like when Google's tool labeled a woman as a gorilla last year.

"When the computer makes false assumptions about this image is really this thing and that's a bad, potentially offensive relationship, that's when you start getting into the trouble areas with computer vision," Lester says. To avoid issues like the one Google ran into, Shutterstock's team identified potentially problematic distinctions and retrained its model on those images. If the model is smart enough in those areas, Lester says, it no longer makes those offensive associations.

Eventually sites like Shutterstock can use computer vision to power new types of searching, or new ways to interact with images. Someday you could search for a pair of shoes you see on a celebrity by dragging a box over that part of a photo—you wouldn't need to describe the shoes or even know the name of their wearer.

"When you start changing your discovery experience to be more based on pixels, you can affect your searches in ways that so far the industry hasn't seen," Lester says.

With Shutterstock's new tool launching today, we're one step closer to that reality. And Lester and his colleagues are excited to see how customers will put it to use. "The exciting thing about putting it into the market is to see what people want from it and how they use it," says Lawrence Lazare, the product director for Search and Discovery at Shutterstock. "Sometimes people don't use it how you think they will."

Source : <https://www.popsci.com/shutterstock-is-visualizing-images-in-whole-new-way/>

How Google Aims To Dominate Artificial Intelligence

In November 2007, Google laid the groundwork to dominate the mobile market by releasing Android, an open-source operating system for phones. Eight years later to the month, Android has an 80 percent market share, and Google is using the same trick—this time with artificial intelligence.

Today Google is announcing **TensorFlow**, its open-source platform for machine learning, giving anyone a computer and internet connection (and casual background in deep learning algorithms) access to one of the most powerful machine learning platforms ever created. More than 50 Google products have adopted TensorFlow to harness deep learning (machine learning using deep neural networks) as a tool, from identifying you and your friends in the Photos app to refining its core search engine. Google has become a machine learning company. Now they're taking what makes their services special, and giving it to the world.

INTRODUCING TENSORFLOW, THE ANDROID OF AI

TensorFlow is a library of files that allows researchers and computer scientists to build systems that break down data, like photos or voice recordings, and have the computer make future decisions based on that information. This is the basis of machine learning: computers understanding data, and then using it to make decisions. When scaled to be very complex, machine learning is a stab at making computers smarter. That's the broader, and more ill-defined field of artificial intelligence. TensorFlow is extraordinarily complex, because of its precision and speed in digesting and outputting data, and can unequivocally be placed in the realm of artificial intelligence tools.

Here are the nitty-gritty details: the TensorFlow system uses data flow graphs. In this system, data with multiple dimensions (values) are passed along from mathematical computation to mathematical computation. Those complex bits of data are called tensors. The math-y bits are called nodes, and the way the data changes from node to node tells the overall system relationships in the data. These tensors flow through the graph of nodes, and that's where the name TensorFlow comes from.

Open-sourcing TensorFlow allows researchers and even grad students the opportunity to work with professionally-built software, sure, but the real effect is the potential to inform every machine learning company's research across the board. Now organizations of all sizes—from small startups to huge companies on par with Google—can take the TensorFlow system, adapt it to their own needs, and use it to compete directly against Google itself. More than anything, the release gives the world's largest internet company authority in artificial intelligence. Besides Android, he also likens the platform to Gmail, Google's ubiquitous email application. There are competitors, but Gmail is cleaner and makes more sense in most applications.

While the others, most notably Torch and Theano, do have small groups updating them, it's nothing like the full force of the developers working on Google's machine learning infrastructure. Manning says that while TensorFlow is a huge gift to the community (one capable of reducing time spent optimizing the neural networks by 100 times), they might indirectly benefit from open-sourcing their tools.

"A very small amount of companies have been trying to hire up a very large percentage of the talented people in artificial intelligence in general, and deep learning in particular," Manning says. "Google is not a charity, I'm sure it's also occurred to them that by ceding this, we will have a lot of Ph.D students who will be in universities and already liking Google deep learning tools."

Jeff Dean, one of Google's top engineers and one of the two people who could be listed as an author for TensorFlow (the other is Rajat Monga), is cautious about estimating the adoption in the community. He says that while it's something Google has found immensely useful in their own work, the real test is whether the community will find it as capable. The idea is to provide a tool so the whole community will be able to go from not just ideas, but actual implementations of things more rapidly.

"We're hoping, basically, to accelerate machine learning research and deployment," Dean says. And while this is a big gift the community, the ideal scenario is that the community gives back, and shares what they've made with other researchers (and Google).

"The machine learning community has been really good at polishing ideas, and that's a really good thing, but it's not the same thing as polishing working code associated with research ideas," Dean says.

He also mentions that TensorFlow will help Google interns when they return back to their schools, because they can now access the once-proprietary systems on projects they might not have finished during their time at the company.

The TensorFlow system is a pretty complete package for an individual researcher. The system is a complete, standalone library associated with tools and an Apache 2.0 license, so it can be used in commercial settings. It can be compiled on desktops or laptops, or deployed on mobile (Android first, naturally, and then iOS to come later). It also comes with tutorials and documentation on how to modify and play with the platform.

Manning suggests that the ability to run deep learning algorithms on mobile devices will be an important factor that separates TensorFlow from other open-source systems.

For those who want to use the system as-is, Google is providing a version that researchers can start using right now (as pre-built binaries).

There's also an application programming interface (API), for software developers to train and control their TensorFlow models. And this isn't a knockoff—it's the literal system used in the Google app, and more than 50 other products.

INSIDE GOOGLE'S ARTIFICIAL INTELLIGENCE LAB

Google is opening this platform to the world, which gives us an equal opportunity to peek in and see how the company thinks about developing machine learning systems.

Internally, Google has spent the last three years building a massive platform for artificial intelligence and now they're unleashing it on the world. Although, Google would prefer you call it machine intelligence. They feel that the word artificial intelligence carries too many connotations, and fundamentally, they're trying to create genuine intelligence—just in machines.

It's the model that they've used within the company for years: where any engineer who wants to play with an artificial neural network can fork it off the system and tinker. That's the kind of open structure that allows 100 teams within a company to build powerful machine learning techniques.

"Machine learning is a core, transformative way by which we're re-thinking how we're doing everything," Google CEO Sundar Pichai said on the company's earnings call in October 2015. "We are thoughtfully applying it across all our products, be it search, ads, YouTube, or Play. And we're in early days, but you will see us — in a systematic way — apply machine learning in all these areas."

WELCOME TO GOOGLE, WHERE EVERYTHING IS AI AND AI IS EVERYTHING

It's difficult to lay out a concrete diagram of machine intelligence research at Google, because it's always changing, and saturates nearly every team in the company.

Google's VP of engineering, John Giannandrea, calls this an "embedded model." I met him at one of the many sleek modern moderns at Google's headquarters in sunny Mountain View, California, in the fall of 2015.

I was on a floor technically not open to the public, and when I was left unattended for a moment, an engineer came up to me, noticing I wasn't wearing an employee badge. He asked who I was, and saying I was a writer didn't smooth the situation over. Google prides itself on making its research open to the public, but work in the labs is kept under heavy wraps.

THE VOICE OF THE FUTURE

A rising star in Google's catalog of tools is Voice Search. You've probably run into it before even if you didn't know exactly what it was: it's the little microphone icon in the main Google search bar, which when pressed, let's you speak your search query to Google instead of typing it in. That same little microphone appears in Google's Search app for iPhone and Android, and can be found within the Android search bar itself on many smartphones.

Although superficially thought of as a rival to Siri, Google Voice search has actually become a secondary gateway to Google's vast knowledge base, and to the language recognition team's delight, it's finally getting more popular.

While Google doesn't release the percentage of voice searches in relation to text, it does provide a veritable rabbit hole of statistics: mobile search is now more popular than desktop, mobile voice search has doubled in the last year, about 50 percent of American phone and tablet users know they can ask Google questions, and a third of them actually do it.

"When we started doing speech recognition, users weren't fully confident. They were using it, but you could tell there was hesitation, the technology wasn't as good as it is now," Beaufays says. "Fast forward to nowadays, people are comfortable doing anything possible by voice in their office."

"THERE'S A THIN LINE BETWEEN MAGIC AND MYSTERY, AND WE WANT TO BE ON THE RIGHT SIDE OF IT."

Beyond integration into its core search algorithms, and the expansion into products, Google also has a few moonshots in the works. For that, they rely on Geoff Hinton.

Hinton is one of the foremost thinkers in artificial intelligence—he's often listed in the same sentence as other high-level researchers like Yann LeCun at Facebook, Google's Andrew Ng, and Yoshua Bengio. (In fact, LeCun, Hinton, and Bengio wrote a review in Nature this May on deep learning, which reads like the literal textbook on AI.)

Speaking with Hinton is like talking to someone who lives five years in the future. Our conversation centered around turning documents into thought vectors, so that machines could understand and remember lengthy versions, and reverse engineer the algorithm our brain uses to learn.

"They're getting close to human-level performance. Not in all aspects, but things like object recognition. A few years ago, computer vision people would have told you 'no, you're not going to get to that level in many years.' So that's a reason for being optimistic," Hinton says.

But no matter how well a machine may complement or emulate the human brain, it doesn't mean anything if the average person can't figure out how to use it. That's Google's plan to dominate artificial intelligence—making it simple as possible. While the machinations behind the curtains are complex and dynamic, the end result are ubiquitous tools that work, and the means to improve those tools if you're so inclined.

"There's a thin line between magic and mystery," Google Now's Chennapragada says. "And we want to be on the right side of it."

Source : <https://www.popsci.com/google-ai/>

Excerpts from Swami Vivekananda's Speeches in Chicago

Around 125 years ago, on September 11, Swami Vivekanand delivered an iconic and eloquent speech at the Chicago Convention of Parliament of Religions. Introducing Hinduism to the world in 1893, Swami Vivekanand spoke about intolerance, religion and the need to end all forms of fanaticism. Such was the effect of the remarkable message that he was given a two-minute standing ovation.

Here are extracts from his speech that still continues to play a relevant role in society.

"It fills my heart with joy unspeakable to rise in response to the warm and cordial welcome which you have given us. I thank you in the name of the most ancient order of monks in the world; I thank you in the name of the mother of religions, and I thank you in the name of millions and millions of Hindu people of all classes and sects.

I am proud to belong to a religion which has taught the world both tolerance and universal acceptance. We believe not only in universal toleration, but we accept all religions as true. I am proud to belong to a nation which has sheltered the persecuted and the refugees of all religions and all nations of the earth.

I am proud to tell you that we have gathered in our bosom the purest remnant of the Israelites, who came to Southern India and took refuge with us in the very year in which their holy temple was shattered to pieces by Roman tyranny. I am proud to belong to the religion which has sheltered and is still fostering the remnant of the grand Zoroastrian nation. I will quote to you, brethren, a few lines from a hymn which I remember to have repeated from my earliest boyhood, which is every day repeated by millions of human beings: "As the different streams having their sources in different paths which men take through different tendencies, various though they appear, crooked or straight, all lead to Thee."

The present convention, which is one of the most august assemblies ever held, is in itself a vindication, a declaration to the world of the wonderful doctrine preached in the Gita: "Whosoever comes to Me, through whatsoever form, I reach him; all men are struggling through paths which in the end lead to me."

Sectarianism, bigotry, and its horrible descendant, fanaticism, have long possessed this beautiful earth. They have filled the earth with violence, drenched it often and often with human blood, destroyed civilization and sent whole nations to despair. Had it not been for these horrible demons, human society would be far more advanced than it is now. But their time is come; and I fervently hope that the bell that tolled this morning in honour of this convention may be the death-knell of all fanaticism, of all persecutions with the sword or with the pen, and of all uncharitable feelings between persons wending their way to the same goal."

Here are some other quotes of Swami Vivekananda that are sure to inspire you.

"Arise, awake and do no stop until the goal is reached"

"The world is the great gymnasium where we come to make ourselves strong."

"Condemn none: if you can stretch out a helping hand, do so. If you cannot, fold your hands, bless your brothers, and let them go their own way."

We Celebrate the Birthday of Swami Vivekananda on 12th of January every year as the Vivekananda Jayanti. In 1984, the Government of India declared the same day as the National Youth Day

You can also send your articles to the following email id's:

codecoffeeucoe@gmail.com , sridhar.iyer@universal.edu.in ,
hezal.lopes@universal.edu.in , apurva.chaudhari@universal.edu.in

[Scan Here](#) for our previous editions

