Gayatri Vidya Parishad College of Engineering for Women

DEPARTMENT OF INFORMATION TECHNOLOGY

VOLUME-I

SPIKE-IQ 2k16

FROM THE PRINCIPALS DESK

It’s pleasure to know that GVPCEW is bringing out the magazine of IT department "***SPIKE-IQ***" for the year 2016-2017.

This institution constantly strives in the all-round development of the students through its endless efforts. SPIKE-IQ is one such endeavor providing a wide spectrum of engineering and artistic edifice, swaying from serious thinking to playful inventiveness. The inspiring women students at GVPCEW are brimming with zeal for life empowering themselves with skills and creativity.

I’m happy that there is a dedicated team of staff and students who have brought out ***SPIKE-IQ***. They have presented the stupendous achievements of IT students of GVPCEW in the field of academics, sports and extra- curricular activities.

I extend my heartiest congratulations to the editorial board and all those who have shelved their valuable time to elevate this magazine to unprecedented heights. I wish the readers have a delightful reading. May all our students soar high in uncharted skies and bring glory to the world and their profession with the wings of education.

-Dr.E.V.Prasad

FROM THE EDITORIALS DESK

It gives an immense joy and satisfaction to introduce our very own department magazine- SPIKE-IQ 2K16. Here comes 'SPIKE-IQ 2K16', the magazine of GVPCEW from the IT department. The name of the magazine may see peculiar, but it just means 'the speed at which the technological innovation or advancement is occurring'. So this time, it is the dedication of students, which attempts to bring out the talent concealed within our student community along with teachers. The willingness to share knowledge, concerns and special insights with fellow beings has made this magazine possible. This magazine includes technical articles, biography of a renowned scientist as well as facts regarding computer science, few tricky puzzles with funny corner and exhibits the literary skills and the achievements of students. These contributions have required a generous amount of time and effort. Thank you very much for all the editorial team members who worked for this magazine. It is very glad to take the opportunity of expressing our considerable appreciation to all the contributors of this magazine. Lastly, the contributors and readers of 'SPIKE-IQ 2K16' are always welcome to send us your invaluable feedback and ideas for further improvement of this magazine.

Department Vision

The Department of IT strives to produce competent professionals who are technically sound and ethically strong for the IT industry

Department mission

- Provide quality training that prepares Students to be technically competent for the Industrial and Societal needs.
- Facilitate an environment that promotes continues learning to face the challenges in the IT sector.
- Provide opportunities for learning, leadership and communication skills.

Program Educational Objectives

After successful completion of the program, the graduates will be able to:

- PEO-1: Apply, analyze and solve complex engineering problems using emerging IT technologies with the help of fundamental knowledge in mathematics, science, and engineering.
- PEO-2: Comprehend, Analyse, Design and Create innovative computing products and solutions for real life problems.
- PEO-3: Inculcate the necessary skills to engage in lifelong learning.

Program Specific Outcomes

Engineering Graduates will be able to:

- PSO-1: Develop software applications by analyzing, designing and implementing with cutting edge technology to address the needs of the IT industry.
- PSO-2: Apply the knowledge of Data Science, machine learning, image processing and allied areas to obtain optimized solutions for real time problems

FACULTY ARTICLE

SPEECH RECOGNITION

M.DEEPTHI

ASSISTANT PROFESSOR

INFORMATION TECHNOLOGY

Computer-based processing and identification of human voices is known as speech recognition. It can be used to authenticate users in certain systems, as well as provide instructions to smart devices like the Google Assistant, Siri or Cortana.Essentially, it works by storing a human voice and training an automatic speech recognition system to recognize vocabulary and speech patterns in that voice. In this article, we'll look at a couple of papers aimed at solving this problem with machine and deep learning.

Speech recognition is invading our lives. It's built into our phones, our game consoles and our smart watches. It's even automating our homes. For just $50, you can get an Amazon Echo Dot — a magic box that allows you to order pizza, get a weather report or even buy trash bags — just by speaking out loud:deep learning finally made speech recognition accurate enough to be useful outside of carefully controlled environments.

Andrew Ng has long predicted that as speech recognition goes from 95% accurate to 99% accurate, it will become a primary way that we interact with computers. The idea is that this 4% accuracy gap is the difference between *annoyingly unreliable* and *incredibly useful*. Thanks to Deep Learning, we're finally cresting that peak.

Neural machine translation works by simply feed sound recordings into a neural network and train it to produce text:

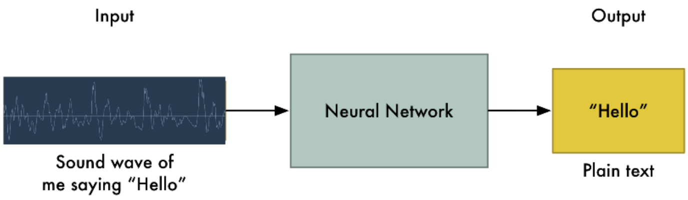

The big problem is that speech varies in speed. One person might say "hello!" very quickly and another person might say "heeeelllllllllllllooooo!" very slowly, producing a much longer sound file with much more data. Both both sound files should be recognized as exactly the same text — "hello!" Automatically aligning audio files of various lengths to a fixed-length piece of text turns out to be pretty hard. We have to use some special tricks and extra precessing in addition to a deep neural network.

Turning Sounds into Bits

The first step in speech recognition is obvious — we need to feed sound waves into a computer.But sound is transmitted as *waves*. How do we turn sound waves into numbers? Let's use this sound clip of me saying "Hello":

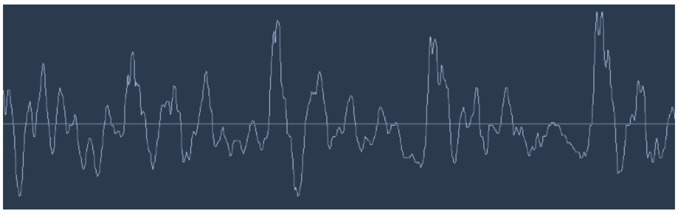

Sound waves are one-dimensional. At every moment in time, they have a single value based on the height of the wave. Let's zoom in on one tiny part of the sound wave and take a look.To turn this sound wave into numbers, we just record of the height of the wave at equally-spaced points

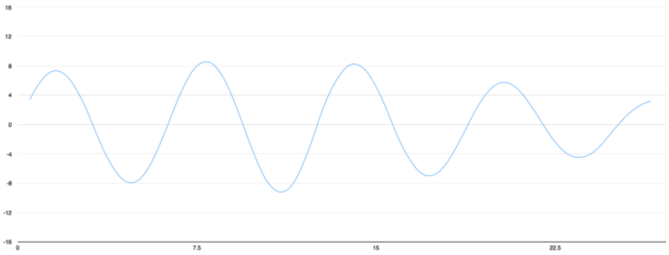

This is called *sampling*.taking a reading thousands of times a second and recording a number representing the height of the sound wave at that point in time. That's basically all an uncompressed .wav audio file is."CD Quality" audio is sampled at 44.1khz (44,100 readings per second). But for speech recognition, a sampling rate of 16khz (16,000 samples per second) is enough to cover the frequency range of human speech. "Hello" sound wave 16,000 times per second. Here's the first 100 samples:

```
[-1274, -1252, -1160, -986, -792, -692, -614, -429, -286, -134, -57, -41, -169, -456, -450, -541, -761, -1067, -1231, -1047, -952, -645, -489, -448
, -397, -212, 193, 114, -17, -110, 128, 261, 198, 390, 461, 772, 948, 1451, 1974, 2624, 3793, 4968, 5939, 6057, 6581, 7302, 7640, 7223, 6119, 5461,
 4820, 4353, 3611, 2740, 2004, 1349, 1178, 1085, 901, 301, -262, -499, -488, -707, -1406, -1997, -2377, -2494, -2605, -2675, -2627, -2500, -2148, -
1648, -970, -364, 13, 260, 494, 788, 1011, 938, 717, 507, 323, 324, 325, 350, 103, -113, 64, 176, 93, -249, -461, -606, -909, -1159, -1307, -1544]
```

Each number represents the amplitude of the sound wave at 1/16000th of a second intervals

A Quick Sidebar on Digital Sampling

But sampling is only creating a rough approximation of the original sound wave because it's only taking occasional readings. There's gaps in between our readings so we must be losing data, right?

According to Nyquist theorem, we know that we can use math to perfectly reconstruct the original sound wave from the spaced-out samples — as long as we sample at least twice as fast as the highest frequency we want to record.

Pre-processing our Sampled Sound Data

We have an array of numbers with each number representing the sound wave's amplitude at 1/16,000th of a second intervals.But trying to recognize speech patterns by processing these samples directly is difficult. Instead, we can make the problem easier by doing some pre-processing on the audio data.start by grouping sampled audio into 20-millisecond-long chunks. Here's first 20 milliseconds of audio (i.e., first 320 samples):

```
[-1274, -1252, -1160, -986, -792, -692, -614, -429, -286, -134, -57, -41, -169, -456, -450, -541, -761, -1067, -1231, -1047, -952, -645, -489, -448
, -397, -212, 193, 114, -17, -110, 128, 261, 198, 390, 461, 772, 948, 1451, 1974, 2624, 3793, 4968, 5939, 6057, 6581, 7302, 7640, 7223, 6119, 5461,
 4820, 4353, 3611, 2740, 2004, 1349, 1178, 1085, 901, 301, -262, -499, -488, -707, -1406, -1997, -2377, -2494, -2605, -2675, -2627, -2500, -2148, -
1648, -970, -364, 13, 260, 494, 788, 1011, 938, 717, 507, 323, 324, 325, 350, 103, -113, 64, 176, 93, -249, -461, -606, -909, -1159, -1307, -1544,
-1815, -1725, -1341, -971, -959, -723, -261, 51, 210, 142, 152, -92, -345, -439, -529, -710, -907, -887, -693, -403, -180, -14, -12, 29, 89, -47, -
398, -896, -1262, -1610, -1862, -2021, -2077, -2105, -2023, -1697, -1360, -1150, -1148, -1091, -1013, -1018, -1126, -1255, -1270, -1266, -1174, -10
03, -707, -468, -300, -116, 92, 224, 72, -150, -336, -541, -820, -1178, -1289, -1345, -1385, -1365, -1223, -1004, -839, -734, -481, -396, -580, -52
7, -531, -376, -458, -581, -254, -277, 50, 331, 531, 641, 416, 697, 810, 812, 759, 739, 888, 1008, 1977, 3145, 4219, 4454, 4521, 5691, 6563, 6909,
6117, 5244, 4951, 4462, 4124, 3435, 2671, 1847, 1370, 1591, 1900, 1586, 713, 341, 462, 673, 60, -938, -1664, -2185, -2527, -2967, -3253, -3636, -38
59, -3723, -3134, -2380, -2032, -1831, -1457, -804, -241, -51, -113, -136, -122, -158, -147, -114, -181, -338, -266, 131, 418, 471, 651, 994, 1295,
 1267, 1197, 1291, 1110, 793, 514, 370, 174, -90, -139, 104, 334, 407, 524, 771, 1106, 1087, 878, 703, 591, 471, 91, -199, -357, -454, -561, -605,
-552, -512, -575, -669, -672, -763, -1022, -1435, -1791, -1999, -2242, -2563, -2853, -2893, -2740, -2625, -2556, -2385, -2138, -1936, -1803, -1649,
 -1495, -1460, -1446, -1345, -1177, -1088, -1072, -1003, -856, -719, -621, -585, -613, -634, -638, -636, -683, -819, -946, -1012, -964, -836, -762,
 -788]
```

Plotting those numbers as a simple line graph gives us a rough approximation of the original sound wave for that 20 millisecond period of time:

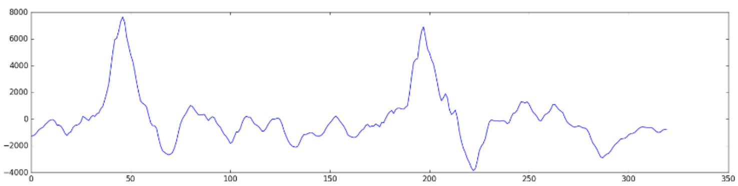

This recording is only *1/50th of a second long*. But even this short recording is a complex mish-mash of different frequencies of sound. There's some low sounds, some mid-range sounds, and even some high-pitched sounds sprinkled in. But taken all together, these different frequencies mix together to make up the complex sound of human speech.To make this data easier for a neural network to process, we are going to break apart this complex sound wave into it's component parts. We'll break out the low-pitched parts, the next-lowest-pitched-parts, and so on. Then by adding up how much energy is in each of those frequency bands (from low to high), we create a *fingerprint* of sorts for this audio snippet.

Imagine you had a recording of someone playing a C Major chord on a piano. That sound is the combination of three musical notes— C, E and G — all mixed together into one complex sound. We want to break apart that complex sound into the individual notes to discover that they were C, E and G. This is the exact same idea.We do this using a mathematic operation called a *Fourier transform*. It breaks apart the complex sound wave into the simple sound waves that make it up. Once we have those individual sound waves, we add up how much energy is contained in each one.The end result is a score of how important each frequency range is, from low pitch (i.e. bass notes) to high pitch. Each number below represents how much energy was in each 50hz band of our 20 millisecond audio clip:

[110.97481594791122, 166.61537247955155, 180.43561044211469, 175.09309469913353, 180.0168691095916, 176.0061997472167, 179.79737781786582, 173.53025213548219, 176.87177119846058, 170.42684732853121
, 159.2602382855659 8, 163.24469810981628, 149.15527353931867, 154.34196586290136, 151.46179061113972, 152.99674239973979, 143.98878156117371, 156.6033737693738, 155.78237530428544, 157.1793094101783
8, 146.28632297509679, 164.37233032929228, 158.1282656446088, 147.23266451005145, 133.26597973863801, 116.5170100028831, 116.85501120577126, 115.40519005123537, 120.85619013711488, 112.4840612316109
1, 111.80244759457571, 92.590676871856431, 105.75863927434719, 95.673146446282971, 90.391748128064208, 79.355818055314899, 86.080143147713926, 84.748200268709567, 83.050569583779065, 86.207180262242
758, 90.252031938154076, 89.361567351948437, 90.917307309643206, 90.746777849123049, 86.726552726337033, 85.709412745066928, 95.938840816664865, 99.092545759170 69, 96.632437741434885, 103.2396123166
6669, 105.80328302591124, 109.53029281234707, 116.46408227060996, 129.20890691592615, 130.43460361780441, 138.15581799444712, 128.25056761852832, 138.14492240466387, 140.0352714810314, 128.151381394
29752, 123.93018478493934, 121.19289035588113, 119.03159255422509, 114.23027889344033, 119.1717342154997, 101.02560719093093, 110.91192243698025, 106.04872005953503, 100.86977927980999, 92.123301579
000341, 94.376766266598295, 97.850709698634489, 113.37126364077845, 110.24526597732718, 113.72249347908021, 120.63960942628063, 122.06482553759932, 117.96716716036715, 120.87682744817975, 125.060973
81947157, 111.57319012901624, 115.54483708595507, 116.99850750130265, 114.40659619324526, 79.869543980883975, 104.83111191845597, 104.66218602004588, 104.91691734582642, 97.143620527536072, 78.43459
781117835, 82.214144782667248, 67.246072805959614, 66.578937262360313, 74.100307226086798, 64.86142301415653, 59.167561212002269, 62.479712687304911, 63.568362396107467, 55.906096471453267, 42.7908
02909362839, 55.693923524361097, 50.776364877715011, 41.196111220671298, 51.062413666348945, 58.493563858289065, 53.081835042922769, 73.060663128159547, 68.21625202122361, 66.7701034934517, 59.76625
124915202, 35.413635503802389, 22.705615809958832, 16.458048045346381, 44.910670465379937, 59.282513769840705, 69.241393677323856, 81.778634874076346, 88.409923803546008, 94.688033733251245, 96.6408
67526244051, 91.806226496828543, 94.570526932206619, 99.250924315589074, 97.899164767741183, 75.176507616277235, 80.947474423758905, 71.859103451990862, 93.863684037461738, 96.757146539348298, 96.52
8614354976241, 99.366456533638413, 102.18717608176904, 102.06596663023235, 101.78493139911082, 103.7883358299547, 99.915220403870748, 107.43478470929935, 104.46449552620618, 105.70789868195298, 101.
10596541338749, 100.75737831526195, 91.742897073196886, 88.307278943069093, 90.936627732905492, 71.134275744339803, 72.504304977841457, 76.233185506299705, 63.281284410272761, 45.380164336858961, 43
.018963766250437, 49.133789791276826, 53.507751009532953, 48.586423555688746, -4.473077611302888 3, 50.833000650183408, 51.003802143009629, 39.577356593427531, 47.096919248906332, 55.442197175664383,
 56.967128095484341, 49.383247263177985]

Each number in the list represents how much *energy* was in that 50hz frequency band.If we repeat this process on every 20 millisecond chunk of audio, we end up with a spectrogram (each column from left-to-right is one 20ms chunk):

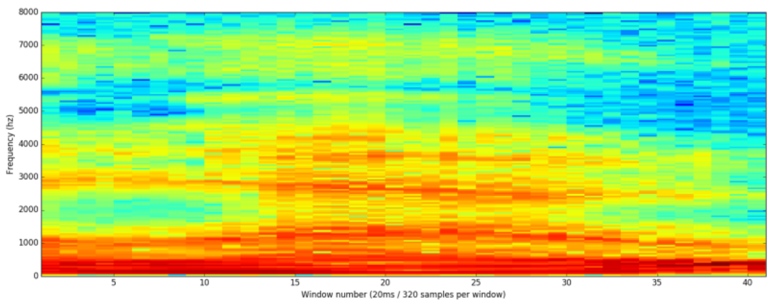

A spectrogram is cool because you can actually *see* musical notes and other pitch patterns in audio data. A neural network can find patterns in this kind of data more easily than raw sound waves. So this is the data representation we'll actually feed into our neural network.

Recognizing Characters from Short Sounds

Now that we have our audio in a format that's easy to process, we will feed it into a deep neural network. The input to the neural network will be 20 millisecond audio chunks. For each little audio slice, it will try to figure out the *letter* that corresponds the sound currently being spoken.

We'll use a recurrent neural network — that is, a neural network that has a memory that influences future predictions. That's because each letter it predicts should affect the likelihood of the next letter it will predict too. For example, if we have said "HEL" so far, it's very likely we will say "LO" next to finish out the word "Hello". It's much less likely that we will say something unpronounceable next like "XYZ". So having that memoryprevious predictions helps the neural network make more accurate predictions going forward.After we run our entire audio clip through the neural network (one chunk at a time), we'll end up with a mapping of each audio chunk to the letters most likely spoken during that chunk. Here's what that mapping looks like for me saying "Hello":

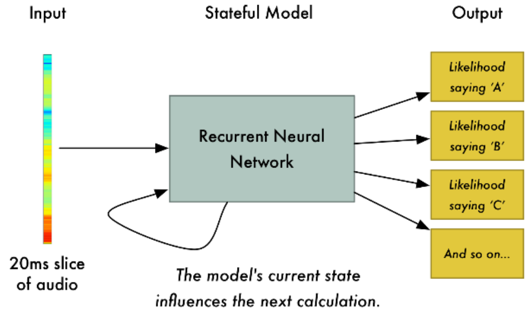

We'll use a recurrent neural network — that is, a neural network that has a memory that influences future predictions. That's because each letter it predicts should affect the likelihood of the next letter it will predict too. For example, if we have said "HEL" so far, it's very likely we will say "LO" next to finish out the word "Hello". It's much less likely that we will say something unpronounceable next like "XYZ". So having that memory of previous predictions helps the neural network make more accurate predictions going forward.

After we run our entire audio clip through the neural network (one chunk at a time), we'll end up with a mapping of each audio chunk to the letters most likely spoken during that chunk. Here's what that mapping looks like for me saying "Hello":

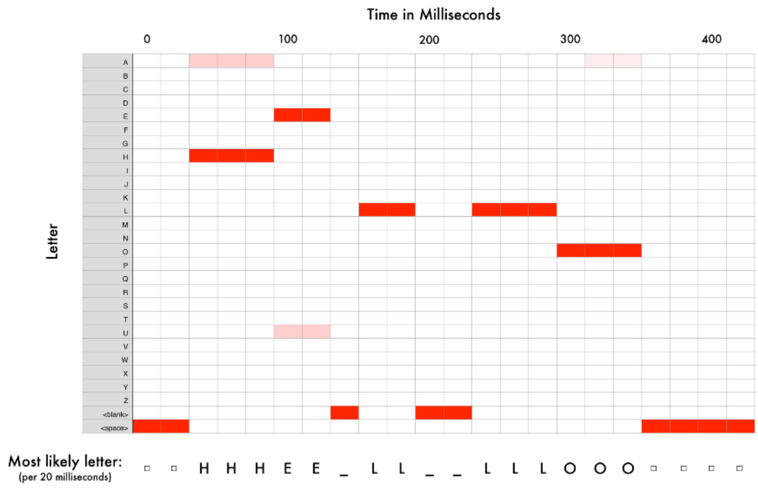

Our neural net is predicting that one likely thing I said was "HHHEE_LL_LLLOOO". But it also thinks that it was possible that I said "HHHUU_LL_LLLOOO" or even "AAAUU_LL_LLLOOO".We have some steps we follow to clean up this output. First, we'll replace any repeated characters a single character:

HHHEE_LL_LLLOOO become

HHHUU_LL_LLLOOO becomes HU_L_LO

AAAUU_LL_LLLOOO becomes AU_L_LO

Then we'll remove any blanks:

- HE_L_LO becomes HELLO
- HU_L_LO becomes HULLO
- AU_L_LO becomes AULLO

That leaves us with three possible transcriptions — “Hello”, “Hullo” and “Aullo”. If you say them out loud, all of these sound similar to “Hello”. Because it’s predicting one character at a time, the neural network will come up with these very *sounded-out* transcriptions. For example if you say “He would not go”, it might give one possible transcription as “He wud net go”.The trick is to combine these pronunciation-based predictions with likelihood scores based on large database of written text (books, news articles, etc). You throw out transcriptions that seem the least likely to be real and keep the transcription that seems the most realistic.

Of our possible transcriptions “Hello”, “Hullo” and “Aullo”, obviously “Hello” will appear more frequently in a database of text (not to mention in our original audio-based training data) and thus is probably correct. So we’ll pick “Hello” as our final transcription instead of the others. Done!Of course it is possible that someone actually said “Hullo” instead of “Hello”. But a speech recognition system like this (trained on American English) will basically never produce “Hullo” as the transcription. It’s just such an unlikely thing for a user to say compared to “Hello” that it will always think you are saying “Hello” no matter how much you emphasize the ‘U’ sound.

Try it out! If your phone is set to American English, try to get your phone’s digital assistant to recognize the world “Hullo.” You can’t! It refuses! It will always understand it as “Hello.”Not recognizing “Hullo” is a reasonable behavior, but sometimes you’ll find annoying cases where your phone just refuses to understand something valid you are saying. That’s why these speech recognition models are always being retrained with more data to fix these edge cases.

Can I Build My Own Speech Recognition System?

One of the coolest things about machine learning is how simple it sometimes seems. You get a bunch of data, feed it into a machine learning algorithm, and then magically you have a world-class AI system running on your gaming laptop’s video card… *Right*?That sort of true in some cases, but not for speech. Recognizing speech is a hard problem. You have to overcome almost limitless challenges: bad quality microphones, background noise, reverb and echo, accent variations, and on and on. All of these issues need to be present in your training data to make sure the neural network can deal with them.Here’s another example: Did you know that when you speak in a loud room you unconsciously raise the pitch of your voice to be able to talk over the noise? Humans have no problem understanding you either way, but neural networks need to be trained to handle this special case. So you need training data with people yelling over noise!

To build a voice recognition system that performs on the level of Siri, Google Now!, or Alexa, you will need a *lot* of training data — far more data than you can likely get without hiring hundreds of people to record it for you. And since users have low tolerance for poor quality voice recognition systems, you can’t skimp on this. No one wants a voice recognition system that works 80% of the time.For a company like Google or Amazon, hundreds of thousands of hours of spoken audio recorded in real-life situations is *gold*. That’s the single biggest thing that separates their world-class speech recognition system from your hobby system. The whole point of putting *Google Now!* and *Siri* on every cell phone for free or selling $50 *Alexa* units that have no subscription fee is to get you to **use them as much as possible**. Every single thing you say into one of these systems is *recorded forever* and used as training data for future versions of speech recognition algorithms. That’s the whole game!

REFERENCES

1. https://medium.com/@ageitgey/machine-learning-is-fun-part-6-how-to-do-speech-recognition-with-deep-learning-28293c162f7a
2. https://heartbeat.fritz.ai/a-2019-guide-for-automatic-speech-recognition-f1e1129a141c
3. http://andrew.gibiansky.com/blog/machine-learning/speech-recognition-neural-networks/
4. https://towardsdatascience.com/recognizing-speech-commands-using-recurrent-neural-networks-with-attention-c2b2ba17c837

https://ai.googleblog.com/2019/03/an-all-neural-on-device-speech.html

VR –The Next Big Thing

P.Kundana

(16JG1A1240)

As impressive and remarkable the technology already is, many people believe it's still in its early days. But it's not we have already become the Techie-dependents. Virtual Reality (VR) means experiencing objects or things through our computers that actually don't really exist. Virtual Reality is already becoming a very in demand concept, and with so many new cameras and devices, it can be easy to feel left behind. That's because VR is still relatively new technology. This Article explains what the best thing VR technology has done recently is!

The Future and What Can We Expect from the Virtual Reality Virtual reality will get more physical the next evolution of Virtual Reality would be where you participate physically in that VR world indeed. And it's not just about sitting down - if you're a quarterback on example, you actually get to throw a football, and thus you can interface with the team. This is a kind of stuff that it's going to happen. Headsets today are doing a great job at catering to your visual senses, and as well a little bit of audio. And that's just 2 of the senses. Since you begin catering to the rest of the senses - temperature-wise, body-wise, and smell, the reality factor of VR becomes more stronger and the virtual piece begins to fade.

Recently Researchers have designed a Virtual Glass. This " virtual glass" wants to trick your taste buds. Researchers at the National University of Singapore have designed a cocktail glass which is capable of tricking your senses into thinking that the water you are drinking is actually, well, anything else. On the rim of the glass are two electrode strips. When these come in contact with your tongue, they send electric pulses, which stimulate your taste buds to mimic different tastes. It's 180 micro amps for a sour taste, 80 micro amps for a bitter taste and 40 micro amps for a bitter taste. A system of sensors and electrodes can digitally transmit the basic color and sourness of a glass of lemonade or any drink to a tumbler of water, making it look and taste like a different drink. This glass also takes care of the smell and sight.

When life hands you digital lemons, make virtual lemonade. And the idea behind this work is to let people share sensory experiences over the internet.

"People are always posting pictures of drinks on social media – what if you could upload the taste as well? That's the ultimate goal," says Nimesha Ranasinghe at the National University of Singapore. Beyond social media sharing, virtual flavorings could help people enjoy digital versions of sugary drinks without consuming calories or damaging their teeth.

The ideas of VR are almost limitless and VR could really be the next Big Thing.

Student Life

The most important factor that affects the student life is the value of time. Being a student we should do everything on time as it never waits for anyone. There are different stages in our life. One of these stages is student days. Student life is considered as the most important period of our life. Our future dreams, desires and hopes depend upon it. Student life is a period of preparations. It is a period of education. At this time, our mind is like clay. Clay is a soft thing and the potter designs various things out of the clay. Like clay, our mind can also be shaped in different ways. Once the pots are made

their shape cannot be changed. Likewise, once our character is formed in one way, it cannot be changed easily. If we make right use and receive good education during our student life, we shall be successful in future. On the other hand, if we aren't serious at this time, we can't achieve our goals. Students, therefore, should be very careful. We must think seriously before every step we take. We must learn new things as much as possible at this period. It cannot be changed easily. If we make right use and receive good education during our student life, we shall be successful in future. On the other hand, if we aren't serious at this time, we can't achieve our goals. Students, therefore, should be very careful. We must think seriously before every step we take. We must learn new things as much as possible at this period. ..

The most important factor that affects the student life is the value of time. Being a student we should do everything on time as it never waits for anyone. Obeying

one's parents and teachers and respecting and loving one's elders are the great

virtues of a student. Another important part of a student's life is his/her social life. A student must have good discipline and he/she must be co-operative with everyone. Social life influences our character in many ways.

Many students enter college expecting good times, friendship and a good sense of direction. They soon find out that colleges come with challenges and struggles because of the great demands and

important for them to manage a proper sticktoit.

Student life is the best part of an individual's life. At this time our main task is to study. We should stop thinking of anything else and concentrate on education. Education must be given the top priority.

Spreading intelligence throughout the cloud

Connected smart machines, such as robots and autonomous vehicles, are fundamental to the evolving Networked Society. Enhanced cloud architecture that can distribute and share machine intelligence will enable smart connected machines to work on an increasingly higher level.

Supported by advancements in artificial intelligence (AI) – particularly in the areas of big data analytics, machine learning and knowledge management – rapid progress has been made in terms of what smart machines can do. Developments in connectivity and cloud technologies are making it possible to distribute and share machine intelligence more easily, at a lower cost, and on a much wider scale than before.

When connected to the cloud, smart machines will be able to use the powerful computational, storage and communication resources of state-of-the-art data centers. Today's intelligent software robotics systems are capable of supporting repetitive administrative tasks with current development pushing toward advisory tasks. Cloudification shifts the capabilities of these systems into a new sphere that includes complex problem-solving and decision-making on a mass-market scale.

Connect, store, compute... and share

Shifting systems into the cloud enables communities of collaborating robots, machines, sensors and humans to process and share information. Each new insight collected within a community can be shared instantly, which increases the effectiveness of collaborative tasks, and improves performance throughout the system, with a common awareness of system state shared by all participants, as well as a shared knowledge base.

A distributed machine intelligence architecture offers lower implementation costs. Sharing a backbone of almost unlimited computational power makes it possible to build lightweight, low-cost robots and smart machines that require a low level of control and a minimum amount of sensors and actuators. Application-specific requirements related to responsiveness and speed will determine whether a local or global cloud is most suitable, and how much intelligence can be distributed.

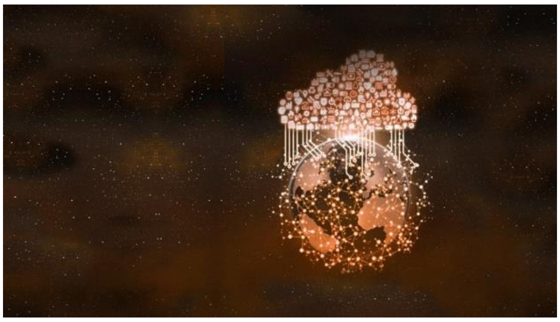

Smart and mobile capabilities virtually everywhere

Intelligent clouds will create new value chains in many industry segments, but some of the forerunners include mining, agriculture, forestry and health care. New opportunities will open up for all organizations and people involved in the supply chain from the manufacturer to the customer. Consider an automated agriculture application. The application remotely controls farm machines to carry out various farming tasks. To harvest mature crops, for example, the system will control the necessary machines to cut, gather and transport them. Each individual machine will take local decisions to ensure secure completion of its set tasks, working in conjunction with all the machines involved in the harvesting. Weather reports gathered from another distributed cloud application are used by the system to carry out harvesting in an optimal way. Contact with the farmer occurs only when participating machines cannot resolve issues themselves.

The harvesting example highlights just one of the many coming applications that will rely on multiple information sources, cloud, and distributed machine intelligence. To ensure scalability and widespread uptake of such applications, the challenge lies in the rapid development and proliferation of universally accessible mobile capabilities. 5G will provide a resilient, high-availability, low-latency network that offers applications with integrated computing and storage resources that are ideally placed to meet latency requirements 5G is well matched to industrial robotics applications because, like other radio technologies, it removes the need for cabling and minimizes infrastructure adaptations, but it also offers identity management, optimum placement of resources, and encryption for security and privacy.

James Gosling

James Arthur Gosling, OC (born May 19, 1955) is a Canadian computer scientist, best known as the founder and lead designer behind the Java programming language.

Education and career

James Gosling received a Bachelor of Science from the University of Calgary and his M.A. and Ph.D. from Carnegie Mellon University, all in computer science. He wrote a version of Emacs called Gosling Emacs (Gosmacs) while working toward his doctorate. He built a multi-processor version of Unix for a 16-way computer system while at Carnegie Mellon University, before joining Sun Microsystems. He also developed several compilers and mail systems there.

Gosling was with Sun Microsystems between 1984 and 2010 (26 years). He is known as the father of the Java programming language. He got the idea for the Java VM while writing a program to port software from a PERQ by translating Perq Q-Code to VAX assembler and emulating the hardware. He left Sun Microsystems on April 2, 2010 after it was acquired by the Oracle Corporation, citing reductions in pay, status, and decision-making ability, along with change of role and ethical challenges. He has since taken a very critical stance towards Oracle in interviews, noting that "during the integration meetings between Sun and Oracle, where we were being grilled about the patent situation between Sun and Google, we could see the Oracle lawyer's eyes sparkle." He clarified his position during the *Oracle v. Google* trial over Android: "While I have differences with Oracle, in this case, they are on the right. Google totally slimed Sun. We were all really disturbed, even Jonathan Schwartz; he just decided to put on a happy face and tried to turn lemons into lemonade, which annoyed a lot of folks on Sun." However, he approved of the court's ruling that APIs should not be copyrightable.

In March 2011, Gosling left Oracle to work at Google. Six months later, he followed his colleague Bill Vass and joined a startup called Liquid Robotics. In late 2016, Liquid Robotics was acquired by Boeing. Following the acquisition, Gosling left Liquid Robotics to work at Amazon Web Services as Distinguished Engineer in May 2017.

He is an advisor at the Scala company Lightbend, Independent Director at Jelastic, and Strategic Advisor for Eucalyptus, and is a board member of DIRTT Environmental Solutions.

He is known for his love of proving "the unknown" and has noted that his favorite irrational number is √2. He has a framed picture of the first 1,000 digits of √2 in his office.

Contributions

Gosling initially became known as the author of Gosling Emacs. He also invented an early Unix windowing system called NeWS, which became a lesser-used alternative to the still used X Window, because Sun did not give it an open source license. He is generally credited with having invented the Java programming language in 1994. He created the original design of Java and implemented the language's original compiler and virtual machine. Gosling traces the origins of the approach to his early graduate student days, when he created a p-code virtual machine for the lab's DEC VAX computer, so that his

UCSD Pascal. In the work leading to Java at Sun, he saw that architecture-neutral execution for widely distributed programs could be achieved by implementing a similar philosophy: always program for the same virtual machine.

For his achievement, the National Academy of Engineering in the United States elected him as a Foreign Associate member. Another contribution of Gosling's was co-writing the "bundle" program, known as "shar", a utility thoroughly detailed in Brian Kernighan and Rob Pike's book *The Unix Programming Environment*.

Honors

- 2002: awarded *The Economist* Innovation Award.
- 2002: awarded *The Flame Award* USENIX Lifetime Achievement Award.
- 2007: made an Officer of the Order of Canada. The Order is Canada's second highest civilian honor. Officers are the second highest grade within the Order.
- 2013: became a fellow of the Association for Computing Machinery.
- 2015: awarded IEEE John von Neumann
- 2019: named a Computer History Museum Fellow for the conception, design, and implementation of the Java programming language.

Books

- Ken Arnold, James Gosling, David Holmes, *The Java Programming Language, Fourth Edition*, Addison-Wesley Professional, 2005, ISBN 0-321-34980-6
- James Gosling, Bill Joy, Guy L. Steele Jr., Gilad Bracha, *The Java Language Specification, Third Edition*, Addison-Wesley Professional, 2005, ISBN 0-321-24678-0
- Ken Arnold, James Gosling, David Holmes, *The Java Programming Language, Third Edition*, Addison-Wesley Professional, 2000, ISBN 0-201-70433-1
- James Gosling, Bill Joy, Guy L. Steele Jr., Gilad Bracha, *The Java Language Specification, Second Edition*, Addison-Wesley, 2000, ISBN 0-201-31008-2
- Gregory Bollella (Editor), Benjamin Brosgol, James Gosling, Peter Dibble, Steve Furr, David Hardin, Mark Turnbull, *The Real-Time Specification for Java*, Addison Wesley Longman, 2000, ISBN 0-201-70323-8
- Ken Arnold, James Gosling, *The Java programming language Second Edition*, Addison-Wesley, 1997, ISBN 0-201-31006-6
- Ken Arnold, James Gosling, *The Java programming language*, Addison-Wesley, 1996, ISBN 0-201-63455-4
- James Gosling, Bill Joy, Guy L. Steele Jr., *The Java Language Specification*,

Addison Wesley Publishing Company, 1996, ISBN 0-201-63451-1
- James Gosling, Frank Yellin, The Java Team, *The Java Application Programming Interface, Volume 2: Window Toolkit and Applets*, Addison-Wesley, 1996, ISBN 0-201-63459-7
- James Gosling, Frank Yellin, The Java Team, *The Java Application Programming Interface, Volume 1: Core Packages*, Addison-Wesley, 1996, ISBN 0-201-63453-8

Java language Environment: A white paper, Sun Microsystems, 1996
- James Gosling, David S. H. Rosenthal, Michelle J. Arden, *The NeWS Book : An Introduction to the Network/Extensible Window System (Sun Technical Reference Library)*, Springer, 1989, ISBN 0-387-96915-2

STUDENTS HUB

We know it takes dedication and courage to learn new skills and explore new career opportunities. We are committed to providing you with all the resources and support you need to achieve your learning goals and advance your career.

That is why, today, we are launching the new **Student Hub** experience to all of our existing and new students. The Student Hub provides you with an effective way to connect with fellow students and receive support and guidance from knowledgeable mentors.

Our goal in launching this new experience is to help you acquire valuable and in-demand skills, and to successfully complete you are challenging Nano degree programs. These are intensive programs, and as you proceed through your curriculum, you will need to set ambitious goals, and achieve important milestones. We are excited to offer new resources that will help you do exactly that.

We have had hundreds of conversations with students about their mentorship and community experience at **SPIKE**. We have analyzed years of data to understand which behaviors produced the most successful outcomes for our students. We have experimented, iterated, tested, and trialed a whole range of solutions. In addition, we have created something special for you. Today marks the culmination of a company-wide effort to offer a powerful new set of resources that will help support your learning goals, and power your success.

SPORTS

Sport's is very important in everyone's life. Some people will have more passion towards sports and participation in sports should always be encouraged. Participation in sport makes us fit, active and healthy. It will develop our social and communication skills. We can explore to new places and people when we go for an competition. It will teach the importance of time in our life because every minute is important in a game. "Healthy mind lives in Healthy body" is so true because for a man to be successful his physical, as well as mental state should be well. Our college "Gayatri Vidya Parishad" helps their students to prove their strength in sports by encouraging them in several activities such as:

EDITORIAL TEAM

- **Dr. Dwiti Krishna Bebarta** - **Head of the Department, Department of IT**
- **B. L. V. Vinay Kumar** - **Assistant Professor, Department of IT**
- **G. S. S. Khyati Priya** - **17JG1A1222**