

COMPUTATIONAL AND DATA SCIENCES

THE OFFICIAL NEWSLETTER OF CDS DEPARTMENT, IISc BANGALORE

cds.iisc.ac.in



The Department of Computational and Data Sciences (CDS) is happy to share the second edition of its monthly newsletter. The broad theme of this newsletter is Deep Learning. We would like to walk you through the life of students, faculty, industry experts and entrepreneurs working in this area. We will also give you a peek into what's happening in and around the department. Have fun reading!

From the Chair's Desk




Prof. Sashikumaar Ganeshan
Chair, CDS, IISc

The Computational and Data Sciences (CDS) department was formed in December 2015 from the academic wing of the Supercomputer Education and Research Centre (SERC). SERC serves as the central computing facility, providing infrastructure to many departments across the institute. With the emergence of Computational Science as one of the most sought after areas, we started MTech in Computational Science in 1999. After branching out from SERC, we started MTech and Research programs in Computational and Data Sciences.

Our department is interdisciplinary and combines the expertise of applied mathematics and computer science streams. The department has two broad research areas - Computational Science, which focuses on numerical schemes and computational models for scientific and industrial applications; and Computer and Data Systems, focusing on high performance scalable systems for scientific computing and big data platforms. We also have a dedicated Machine Learning group working on cutting edge research in the areas of Computer Vision and NLP.

We develop new algorithms for machine learning and pattern recognition, and scale them to large datasets. We also use Neural Network and Machine Learning algorithms to support solutions in the scientific domain. This department is unique in combining scientific computing, machine learning, big data and parallel programming to develop holistic solutions to science and engineering problems.



Faculty Interview

Can you give us a glimpse of your background and research interests?

I completed my Bachelors in Mechanical Engineering from Osmania University, Hyderabad. I was interested in pursuing research since my undergraduate days. I would spend time at the Advanced Systems Laboratory (ASL, DRDO) after my academic activities at college. I had worked on the development of flapping wing micro air vehicles there. I joined Jawaharlal Nehru Centre for Advanced Scientific Research (JNCASR) as a summer student when pursuing my Bachelor's degree, and continued to do Masters in Engineering Mechanics there. My stay at JNCASR got me more interested in computations, numerical simulations and parallel programming. I was very fortunate to get an opportunity to carry out computational fluid dynamics simulations on "EKA", the then 4th fastest supercomputer machine in the world, built by CRL, Pune. After completing Masters at JNCASR, I did another masters programme in Aerospace Engineering from Georgia Tech. After that I completed my PhD from Texas A&M University (TAMU). I worked on high performance computing (HPC) aspects relevant to the simulations of turbulent flows. My PhD dissertation was on asynchronous computational methods for Partial Differential Equations. I completed my PostDoc at Sandia National Laboratories where I focused on massively parallel simulations of gasturbine and scramjet engines. I have also worked on machine learning algorithms for anomaly detection -- to detect events of interest which are typically extreme with respect to the surroundings -- in a parallel/distributed setting. Broadly, my research interests span some areas in mechanical and aerospace Engineering, applied mathematics and computer science.

What are the main research areas your lab would focus on?

I would like to focus on three broad research areas - understanding flow physics that is relevant to problems in aerospace, energy and atmospheric sciences; developing HPC centered numerical



Prof. Konduri Aditya
Assistant Professor
CDS, IISc

methods and HPC software for fluid flow; and using ML based algorithms to analyse the huge data obtained after performing simulations. At the same time I don't want to restrict myself to these areas. I like to work on any problem that interests me if I feel I have the right skill set to work on it.

How was your experience doing Masters at JNCASR? What has motivated you to move back to India and take up a career in academia? Why CDS@IISc?

JNCASR is a nice small institution with strong focus on research. Prof. Roddam Narasimha, with whom I worked at JNCASR, has inspired me a lot. We used to have discussions, not just about research, but on a wide variety of topics. It was a very memorable experience.

My experiences have made me realize that I will have more satisfaction and pleasure working in India. I feel this position at IISc is a very good opportunity and matches my research interests really well.

What courses do you plan to teach at IISc?

I am interested in teaching courses on parallel programming, numerical solutions for partial differential equations and high performance computational fluid dynamics. At present, Mechanical and Aerospace Engineering departments at the Institute offer courses in computational fluid dynamics. But those courses are focussed more on numerical methods for solving the governing equation in different flow regimes. I would like to focus on the use of high performance computing to implement those numerical methods. I will start teaching from January 2020.

What are the qualities that you look for in a student who wants to work in your group?

The most important requirement is to have the passion and interest to learn. Having a background in areas I work on, like experience in programming, numerical methods and fluid mechanics would certainly help. I know it is very difficult to have a combination of all this at an undergraduate level. So I look for students with experience in at least one or two areas that overlap with my interests, and willing to learning other topics as well.

What are your interests outside of work?

I used to play table tennis from my childhood. At TAMU, I played for the university. I want to resume it here too. I like listening to interviews of people from

different walks of life. This helps me get an insight into others' experiences. I like travelling around, especially using public transport, and talk to people around. This makes me feel connected to the outside world.

Your advice to students, something that you learned as a researcher.

Never shy away from working on things that you don't know. Always have the passion to learn the unknown. Keep learning. If you get good opportunities in any area that excites you, don't hesitate to learn and make things work.

Very often opportunities of your interest may not come your way. Do not get discouraged. Work towards creating the necessary opportunities and do not hesitate to ask people for help.

Artificial Intelligence @ CDS

Artificial Intelligence (AI) has shown its resurgence in today's technological world. Various AI techniques are helping solve challenging problems in data science, computer vision, natural language processing (NLP) and many more. CDS department is at the forefront of Deep Learning and AI research in the country; and is one among the renowned research centres in the world. While many of the labs in CDS work on some flavour of machine learning/ deep learning research, three of the labs here, Video Analytics Lab (VAL), Visual Computing Lab (VCL) and Machine and Language Learning Lab (MALL), work exclusively on cutting edge research in the areas of computer vision, NLP and Knowledge Graphs. These labs have had a significant impact in taking forward the research in AI with their highly cited publications in top-tier conferences such as NIPS, EMNLP, ICML, AAI, CVPR, ICCV, ECCV, ACL and many more. Some of the recent publications of these labs are listed here:

- "Zero-Shot Knowledge Distillation in Deep Networks.", Gaurav Kumar Nayak, Konda Reddy Mopuri, Vaisakh Shaj, R. Venkatesh Babu, and Anirban Chakraborty, ICML 2019.
- "Generalizable data-free objective for crafting universal adversarial perturbations.", Konda Reddy Mopuri, Aditya Ganeshan, and R. Venkatesh Babu, PAMI 2019
- "Pictionary-style word-guessing on hand-drawn object sketches: dataset, analysis and deep network models", Ravi Kiran S., Shiv Surya, Trisha M. and R. Venkatesh Babu, PAMI 2019
- "Gray-box Adversarial Training.", Vivek B. S., Konda Reddy Mopuri, and R. Venkatesh Babu, ECCV 2018
- "Adadepth: Unsupervised content congruent adaptation for depth estimation.", Jogendra Nath Kundu, Phani Krishna Uppala, Anuj Pahuja, and R. Venkatesh Babu, CVPR 2018
- "Divide and Grow: Capturing Huge Diversity in Crowd Images with Incrementally Growing CNN", Deepak Babu Sam, Neeraj S., R. Venkatesh Babu, and S. Mukundan, CVPR 2018
- "CAPNet: Continuous Approximation Projection for 3D Point Cloud Reconstruction Using 2D Supervision", Navaneet K.L., Priyanka M., Mayank A., and R. Venkatesh Babu, AAI 2019
- "Zero-shot Word Sense Disambiguation using Sense Definition Embeddings.", Sawan Kumar, Sharmistha Jat, Karan Saxena, and Partha Talukdar, ACL. 2019

- "KVQA: Knowledge-aware Visual Question Answering", Sanket Shah*, Anand Mishra*, Naganand Yadati and Partha Talukdar, AAAI 2019
- "Towards understanding the geometry of knowledge graph embeddings.", Aditya Sharma, and Partha Talukdar, ACL 2018
- "A dual framework for low-rank tensor completion.", Madhav Nimishakavi, Pratik Kumar Jawanpuria, and Bamdev Mishra, NIPS 2018

CDS actively organizes technical talks to inspire budding researchers and facilitate exchange of ideas with researchers around the world. Some of the recent talks related to Deep Learning and AI are listed here:

- "Towards Embodied Visual Intelligence", by Dr. Dinesh Jayaraman, Postdoctoral scholar in EECS at UC Berkeley, December 14, 2018
- "Question Answering in Healthcare: Challenges and Opportunities", by Dr. Chandan K. Reddy, Associate Professor, Virginia Tech, USA, July 04, 2019
- "Weakly-supervised and incremental learning", by Dr. Kartheek Alahari, INRIA, December 26, 2018
- "Towards Interpretable and Responsible AI", by Vaishak Belle, Chancellor's Fellow, University of Edinburgh, UK, July 02, 2019
- "AI for Dynamic Decisioning under Uncertainty: For Real-World Problems in Retail and Financial Trading", by Dr. Ashwin Rao, VP of AI at Target, Adjunct Faculty at Stanford University, April 12, 2019

Research @ VAL



Video Analytics Lab (VAL) focuses on fundamental and applied research in the areas of image and video analysis, computer vision and machine learning. We aspire to develop intelligent solutions to complex visual perception tasks such as object recognition, semantic segmentation, human motion recognition and forecasting, 3D perception from monocular RGB images, crowd counting and HDR image fusion.

While Deep Learning has outperformed the best of classical methods in many fields, yet it is the least understood among all. Improving the interpretability, generalizability, and robustness of deep learning models is a key aspect of research in our lab.

We aspire to understand the flaws in the modern data-driven supervised deep learning approaches. While the use of data-hungry deep learning models has sky-rocketed, we aim towards reducing the data-dependency of deep learning solutions by formalizing the next-generation unsupervised learning techniques.



Prof. R. Venkatesh Babu
Associate Professor
CDS, IISc

The lab also contributes to the society through its strong industry-academia ties and has received several research grants. Active collaboration with prominent organizations such as ISRO, Boeing, Shell, Nvidia and TCS have led to solutions that have a practical use-case in the society. Our students have joined both academia and industry and have extended their continued collaboration with the lab.

We welcome highly motivated and research oriented students for various Computer Vision projects involving the use of Deep Learning. Positions are available for long-term Project Assistants (PA) and short-term Interns. Interested candidates can send an email attaching their CV to val.openings@gmail.com.

ACM Inaugural Event

The IISc ACM Student Chapter is a joint initiative by research students from the Department of Computational and Data Sciences (CDS) and Computer Science and Automation (CSA). This chapter aims to further communication channels with alumni from academia and industry, highlight and recognize research within the institute and collectively bridge the gap between research and industry.

The chapter held its inaugural event on 13 June 2019, and featured Dr. Gargi Banarjee Dasgupta (Director of IBM Research India and the CTO of IBM India and South Asia) as a keynote speaker. Dr. Gargi gave a very interesting talk on Cognitive Technical Support.

Prof. Jayant Haritsa talked about the history of ACM and the activities of ACM India. Mr. Aakash Khochare, the Chair of IISc ACM chapter committee and PhD student at CDS, explained the need of ACM student chapter to strengthen the student body. He further explained that it will provide a platform for information exchange, career counselling and regular talks by alumni and industry experts. The activities outlined will help students present their research work and keep peers informed about the latest developments.

IGCM 2019

Indo-German conference on Computational Mathematics (IGCM) : Dec 2 – 4, 2019

The purpose of this Indo-German conference is to provide an interdisciplinary platform for researchers around the world to present and discuss the most recent innovations, trends, and challenges in the frontier areas of computational mathematics. IGCM aims to foster interactions among High Performance Computing (HPC) research community and to provide a platform to present and discuss challenges in Exascale computing. The conference topics include, but are not limited to, numerical analysis of partial differential equations (PDEs) & stochastic PDEs, efficient and robust numerical schemes for solving complex problems, optimal control and inverse problems, scalable parallel algorithms and hybrid MPI-OpenMP algorithms with GPU computing.

Important Dates:

Last date for abstract submission: 15 Sept 2019


Intimation of accepted papers: 10 Oct 2019

Last date for early registration: 15 Oct 2019

Beginning of conference: 2 Dec 2019

Admission Statistics

The CDS department has seen a surge in number of PhD students admitted over the last few years! Here are the statistics about admissions in M.Tech and PhD programmes.



M.Tech admission statistics

Year	Applications received	Written test shortlist	Interview shortlist	Tentative shortlist	Final shortlist
2017	534	231	82	47	23
2018	489	136	64	32	22
2019	618	362	92	49	-

BCL, 2019

Third Workshop on Brain, Computation and Learning (BCL) : June 24 – 28, 2019

The Third Workshop on Brain, Computation and Learning was jointly organized by several departments in IISc, including CNS, MBU, CSA, CDS, EE and DESE. The event was sponsored by the Pratiksha Trust and IISc. The primary focus of the workshop was to bring Neuroscientists and Computer Scientists together and trigger useful discussions on various interdisciplinary research problems, such as computational approaches to understanding human brain and motivating Deep Learning research through the understanding of cognitive learning methods. The workshop was a great success with over 200 participants from various parts of the country. This was a truly interdisciplinary event with several interesting talks on various aspects of Neuroscience, Machine Learning, Neuromorphic Computation and Neural Circuits by renowned researchers around the globe, including Dr. Mayank Mehta from UCLA and Dr. Pulkit Agarwal from MIT. The poster session on Day-4 was an excellent platform for young researchers from Neuroscience and Machine Learning to get together and exchange ideas on their research.



Dr. Pulkit Agarwal's talk on "Continually Evolving Machines"



Thought provoking Panel Discussions everyday!



Gaurav Nayak from VCL lab explaining his recent ICML paper



Huh, wait a sec... Is that the human brain?!

Placement Zone

My experience during placements

I am a second year M.Tech. student in the department of Computational and Data Sciences. I was interested in Data science for my career and hence chose to attend interviews of only the companies that offer data science profiles. I started my preparation for placements around June last year. My preparation included practicing coding on programming concepts like Dynamic Programming, Graph Algorithms, etc. I also practiced solving general aptitude questions and brushed up on the basics of Machine Learning and Deep Learning. The companies held pre-placement talks a month before the placement season started. These talks helped me understand more about the work done and thus simplified my task of shortlisting the profiles that I was interested in. While most companies had a written test to shortlist candidates for further rounds, some companies like Flipkart shortlisted candidates based on their resume. I attended the written tests conducted by Samsung, Qualcomm, minds.ai, Goldman Sachs, JPMC, IBM Research and Codenation. Written tests primarily tested our coding skills, aptitude skills and ML concepts. I got shortlisted in Flipkart, Samsung Research, Qualcomm, minds.ai and IBM Research for further rounds of interviews. Most of the interviews started with discussion on my dissertation project and moved on to questions on Machine Learning, Probability, discussion on case studies, etc. I got selected in four companies, Flipkart, IBM Research, Qualcomm and Samsung Research and I chose to join Flipkart.

I sincerely thank my advisor, Prof. Venkatesh Babu, CDS department and the institute for all the support that I have received during my masters program and placements.



Abinaya K.
M.Tech student at VAL , CDS

»»—————««
The average number of offers per MTech (course) student is the highest in CDS among all the departments across IISc!! (Only Masters programs with 10 or more students are considered here)
»»—————««

Number of offers	28
Median Salary	21,75,000 INR
Average Salary	23,33,906 INR
Student Count	19
Number of offers per student	1.47

Some of the companies (of interest to CDS) that visited IISc during the placement season 2018-19:

Company name	Job Profile
Goldman Sachs	Analyst
J.P. Morgan Chase & Co.	Analyst
Microsoft	Data Scientist
minds.ai	Neural Network Engineer
Nutanix	System Software Engineer
Nvidia	System Software Engineer
active.ai	AI Researcher
Amazon	Data Scientist
Flipkart	Data Scientist
Samsung R&D	Machine Learning
Qualcomm	Machine Learning
IBM Research	Software Engineer
Codenation	Software Development

Students' Corner



Konda Reddy Mopuri

I have completed my PhD from VAL lab at CDS. I am currently pursuing my postdoctoral research at the School of Informatics, University of Edinburgh, UK. Broadly, I work in areas of Machine Learning (ML) and Computer Vision (CV). More specifically, while at CDS, we worked towards developing explainable and robust AI (Artificial Intelligent) systems. Currently, I am working on the applications of meta learning. I am fortunate to have worked at VAL lab, where I got a chance to explore some of the most interesting and important research problems in ML. It is not easy to keep up with the current pace of applied ML. I believe my experience at VAL has taught me this, which is a key ingredient to effective research.

The institute provides some of the best courses relevant to this field, which motivated me and ensured that I always enjoyed my work. If you are an early PhD student in this field, you must start with thoroughly learning these courses (RTP) before you start solving the research problems. At a personal level, the department has been very special to me. It has provided excellent resources that no other academic lab in the country would have, a great set of motivated people in the corridors and memorable discussions at the green boards with coffee. I truly cherish my memories at IISc. I wish a great learning experience to all the department members.

I am a doctoral researcher at VAL, working with Prof. Venkatesh Babu on Computational Photography, Computer Vision and Machine Learning. My interest is to create new photographic functionalities and experiences that go beyond what is possible with traditional cameras and image processing tools. I am particularly fascinated with alternatives to traditional image processing pipelines, new paradigms for capturing and consuming photos (including stack-based photography), and displays that leverage human perception mechanisms to provide a more immersive visual experience.

I received my M.Tech degree from NIT Rourkela. Prior to my current position, I have also worked at Aricent (now Altran India) and TCS innovation center (intern).



Ram Prabhakar



Abinaya K.

I recently completed my M.Tech. from the department of Computational and Data Sciences. I was associated with VAL lab for my dissertation project. My work was focused on the vulnerability of deep neural networks to small changes in inputs called adversarial perturbations. These adversarial attacks are a threat to the deployment of deep neural networks in the real world. I worked on various ways to craft universal adversarial perturbations, that can successfully fool models for most of the input images. Since access to training data may not always be possible for an attacker, we explored data-free ways of creating universal adversaries. We also worked on universal attacks that do not alter every pixel in the image, but only alter the border pixels. The perturbation looks like a photo frame in this case and fools the model for most of the images.

Meet our alumni

An interview with Prof. Ravi Kiran S., Assistant Professor, IIIT Hyderabad

Can you share some details on your background?

I grew up in a small city called Warangal, currently in the state of Telangana. I obtained a dual-degree (B.Tech + MS) from IIIT-Hyderabad and went on to gather another MS degree from University of Washington, Seattle USA. I worked in US-based R&D labs for 7 years before coming back to India. I joined the PhD program at IISc in 2014 and graduated 4 years later. I spent 6 months at Qualcomm's R&D group in Bengaluru before joining IIIT-Hyderabad as an Assistant Professor in July 2018.



Prof. Ravi Kiran S.
Assistant Professor, IIIT-H

It is not common to see one come back to India after a 7-year long, successful career in the US. What was your motivation behind this?

Working in the US was a lot of fun and very rewarding. But I could never imagine myself settling there. I saw the experience of my aged relatives whose children settled in the US. They would crave for their children to be around and often had no one to look after them when they were unwell. I promised to myself that this would not happen to my parents. Therefore, I decided to come back. I knew I wanted to come back, but I wasn't sure what to do after coming back. My initial plan was to come back and join an R&D lab if possible. During my stay at the US, my wife and I would spend weekends volunteering to organise cultural activities for kids and teenagers. In the process of doing so, I realised that teaching could be a fun experience. At some point, my thoughts morphed into considering teaching as a career choice. Both my parents are teachers, so perhaps that guided my choice subconsciously too!

Can you share your experience at IISc?

I am really grateful for getting a chance to do research at IISc. The first year was extremely challenging, because I was coming back to academia after a long gap and my daughter was a new born. However, thanks to my supportive family, an extremely helpful advisor (Dr. Venkatesh Babu) and of course, my lab-mates at Video Analytics Lab (VAL), I pulled through. I had the opportunity to interact with faculty from other departments as well, and had many insightful discussions with them. Also, the long discussions at Prkruti and chats with my lab-mates and advisor are all memories I cherish a lot.

How do you compare your current role with that of a research position in the industry?

A huge advantage with academia is that you can pursue topics without any corporate agenda attached to them. I find this to be very liberating. The research atmosphere here in IIIT-H is comparable to IISc in many ways. I feel like I am continuing the journey of IISc, just with more responsibilities. Having said that, I am also happy about working in industrial R&D labs before moving back to academia. Industrial labs work on challenging, real-life problems and require you to wear multiple hats, (project lead, developer, manager) and that experience has been extremely valuable.

Your advice to PhD students on how to select their research topic. Is it better to remain focused on a single topic or work on various aspects in a given area?

I went through recent papers from top-tier conferences in my area, shortlisted the ones I found interesting (atleast at the title level) and grouped them into topics. One such group became the eventual topic I worked on. It is hard to say which of the two approaches (single topic vs. multiple topics) works better. Sometimes the exploration involves multiple topics, and a theme or a dominant topic emerges from that. The single topic approach is probably harder because it needs to have sufficient breadth for exploration and to minimize overlap with existing work. But if the topic is novel, then you are in luck and it can be a great opportunity.

A peek into the Industry

Conversation with Dr. Jayavardhan Gubbi, Senior Scientist at TCS Research

What are the broad areas of work at TCS Research? When did TCS start exploring Deep Learning Research? When did you get associated with TCS?

TCS Ltd. was established in 1968. In 1981, India's first software research and development center, Tata Research Development and Design Centre (TRDDC) was setup in Pune. Initially, TRDDC focused on areas such as high performance computing and software systems. In 2007, there was a major gear shift, when TCS decided to focus more on research and started a network of 19 innovation labs around the globe. TCS pursues research in a variety of fields ranging from behavioral, life sciences and social sciences to software systems and services. TCS started research in deep learning around 2011, when people started using deep learning technology extensively.

I joined TCS in 2015 and currently lead the Machine Vision group under Embedded Systems and Robotics Research area. There is a dedicated wing doing core research in Deep Learning and Artificial Intelligence. I work on machine vision, primarily driven by deep learning, in the TCS Innovation lab at IISc. Incidentally, I started my research life as a project assistant at IISc under the guidance of Prof. A. G. Ramakrishnan. I went on to do PhD and PostDoc at the University of Melbourne. My expertise is in the area of pattern recognition, image processing and signal processing.

Can you share some of the cutting edge research areas that your team works on?

In machine vision, our research primarily focuses on four broad areas: Image and Video Processing, 3D processing, remote sensing and spectral vision, and cognitive vision. The cognitive vision division is in the preliminary phase, and aims at the ultimate dream of artificial general intelligence. In the other three areas, we have developed solutions and also deployed them. The image and video processing research is closely related to deep learning.



Dr. Jayavardhan Gubbi
Senior Scientist, TCS Research

Almost all applications need object detection as a primary step. We have also worked on visual change detection. Some of this work was done in collaboration with VAL lab at CDS. This has many industrial applications as it is very important to know how things change over time.

Can you comment on how academic research differs from industry research?

Industry typically requires integrating components into a larger system. We can't deliver a full product with only computer vision models. It has to be integrated with other components like drone systems or NLP models. Yet, computer vision is the key ingredient that is at the heart of most applications. For example, we are currently working on Visual Question and Answer system, which takes in images of tabular data or charts as input and answers questions based on the extracted data. It does more than typical text recognition or image captioning, and uses both NLP and CV. Academic research does not have to meet this requirement. They can work on very specific and independent topics. Also, research projects that we undertake in the industry need to consider the aspect of value creation. Whereas, this is not the case with academic research. Both industry and academic research have their own advantages. Long term research at TCS is usually conducted in collaboration with academic partners through TCS co-innovation network (COIN).

A peek into the Industry

What is the role of the team in IISc? How often does the work done here get integrated into customer offerings?

TCS has really good teams involved in system integration, and this is one of our key strengths. The Indian Passport system, SBI systems and TCS BaNCS are excellent examples to demonstrate our expertise on the systems side. However, core IP driven research was missing earlier. To address this, multiple innovation labs were established. Our lab at IISc is one among them. We work on niche research problems. I work on core research 40% of my time and the rest of my time is dedicated to product driven activities. ChangeNet and VideoNet are results of pure academic research. It was not created based on any of the clients' requirements. However, we also develop solutions based on the requirements of clients. Our work most often gets integrated into customer offerings. The challenges we face during the execution of the client driven projects are often interesting research problems by themselves.

What in your view are the challenges that prevent some of the research ideas from being deployed?

Sometimes, there is a huge gap between how well an algorithm works in theory, and how practical it is to deploy the same on field. We may design an algorithm to make it work for a particular scenario and it may end up being hard to deploy. There are system level constraints that prevent the research from getting deployed. We worked on an end-to-end drone system with 5 different operational modules working really well, for a warehouse. It required the data to be uploaded to the cloud where it runs and generates results. But the warehouse was at a remote location and we could not upload the data. Integrating the system with an edge device made deployment possible and it is integrated to TCS service offerings. Another important thing is the return on investment (ROI). A system may fail 20% of the times, and if a person has to be employed to take care of this 20%, ROI may not be good enough.

What in your view is the future for Deep Learning Research? What are the key focus areas going to be? (from an industry perspective)

Right now, if you see, 80% of the industrial use cases of deep learning are object identification tasks. This is going to stay and improve. The three key focus areas for future research according to me are: meta-learning, which is going to take the center stage of research, reinforcement learning and learning with limited data. The explicability of deep learning techniques is also a major research area. Often, clients are reluctant to deploy deep learning systems in their pipeline as they lack explicability. In areas such as Bio-medical imaging, it is really difficult to get approvals without proper explicability.

What are the opportunities in India for Deep Learning and Computer Vision specializations?

Most of the opportunities available around the globe are available in India also. Deep learning is one of the areas where India is picking up very well in terms of opportunities.

What is it that you look for in a candidate you hire for your team? Can you share some tips to students to make them more industry ready?


The first and foremost thing I look for in a candidate is his attitude. Second thing is flexibility. They should be open to explore different areas as the situation demands.

About getting industry ready, the answer is very simple. Get out of campus during the summer break, and do internships. Engineering at the end of the day is applied and you need to know the real practical problems out there in the field. When it comes to solving practical problems, whatever research we have done may just remain on paper. We need to focus more on the practical problems to be able to deploy research.

Words of Wisdom for the budding entrepreneur in YOU

An interview with the Founder of IntroTUCE, Mr. Srinivas Numburi

What is IntroTUCE all about?

IntroTUCE is a technology and product startup, doing research in IISc Bangalore. We are developing a technique that can potentially revolutionize the way people connect with each other in real-time. This can change the way people celebrate occasions, relish moments, dance, shop and explore places together. Although this is an Augmented Reality (AR) problem, we want to solve it without using any AR gear. This technology will soon be in the hands of billions of smartphone users with the launch of our mobile app, .

Tell us something about yourself

I spent the initial years of my career working with Intel, Microsoft and Amazon in India. I was an engineer with Intel, a program manager in Microsoft and Director at Amazon. Then I moved to London and worked there for about 9 years. During this period, I worked with Barclays Bank and Lloyds Bank, particularly in the internet banking and mobile banking areas. I worked on fraud monitoring, prevention and detection. I have also worked on scaling up mobile apps for a large number of users. I have e-commerce experience from my initial few years of work with Microsoft and Amazon.

How did the idea evolve? How did you decide to take the big plunge?

Personally, I always wanted to do something on my own. Based on my industry experience, a lot of my friends suggested I work on e-commerce or fintech which were gaining popularity. But I was not very keen on working in those areas. I wanted to work on something new, which no one has ever tried so far. I was thinking of many possibilities. Whenever I came up with something new, it would always end up having a narrow scope. I wanted to work on something which has a large impact, on a lot of people. I spent a lot of time thinking of what to work on, almost 3 years. In 2017 I firmed up my thoughts on working in the space of photos and videos as they are pervasive. This is one of the primary needs of any smartphone user, apart from making phone calls.



Mr. Srinivas Numburi
Founder, IntroTUCE

Photo sharing is a well explored area, so I started thinking of what can be changed in photo creation and photo taking. Photo taking is currently limited by the device location and hence I decided to exploit this as an opportunity to develop a new product concept, which is to bring people who are far apart next to each other in real time, using a smartphone. This could potentially revolutionize how we engage with friends and family, how celebrities and fans interact and how product promotions work. There was no looking back after this point, and I was confident this could be the next breakthrough in social networking.

How has your start-up journey been? What were the challenges you faced?

One of the key challenges is to make this app work efficiently on the variety of devices out there in the market while maintaining the quality of output. For example, the integration between CPU and GPU is different in different phones. An I/O operation from CPU to GPU can take ~50ms if it is not well optimized. This could be a huge challenge in real time applications. These are some of the challenges related to the hardware. Bandwidth is another issue. Although our requirement is <1MBps, this itself is quite challenging to guarantee in some places. Lighting conditions around you can impact the quality of the video. These are the challenges where we are spending most of our time. Adapting deep learning models on smartphones is also not a very matured area.

We are trying to cope with technology and hardware limitations to make this possible. We have proved the concept and have a working model, but this needs more polishing before it can go live. There is a lot more optimization to be done.

What is the key change that Introtuce brings in, with respect to the other players like TikTok?

This idea doesn't have a parallel on any social media platform today. Using TikTok one can record videos and post them for people to view and like. We make it possible to bring two or more people in remote locations together, experience the togetherness, shoot videos or click photos together. You can shop and explore places with a friend who is far away. Our technology can also bring about a change in the space of celebrity-fan interaction by providing an opportunity for active engagement with celebrities. We also create a platform for experiential marketing, where the consumers will get involved in brand promotion activities. The existing social media platforms do focus on solving the same problems, but the key difference we bring about is the aspect of real-time active engagement. We solve this AR problem without using AR gear. This makes our solution more accessible to every Smartphone user.

When was this start-up launched? When do you plan to launch this app to public?

We started implementing this idea last August with Anil and Santosh. By December we were able to prove the concept. We got two full stack engineers, Afsar and Madhavi from NIT Raipur, and a UX designer, Akash from VIT to help us with the app development. Our team has 4 people now. In the last couple of months, we could build an app that works on Android. We are working on optimizing it to make it work well with all kinds of phones, while maintaining a good quality of the output. We plan to launch this by September this year. I urge the tech savvy folks to join our early adopters group and share their valuable feedback with us before we launch it to public. Please register at <https://www.introtuce.co/earlyadopters>



My team at Introtuce

Can you share your association with IISc and VAL lab?

When I got started with this idea, I realized that although many of the technologies needed exist, porting them on a smartphone is a challenge. Also, this work requires integration of multiple technologies which is again a challenge. So implementing this idea required not only development, but also research. I thought given this setting, it is best to get associated with a research institution. There was no better place I could think of other than IISc for this. This also helped me come back to India, which I have been wanting to do for a long time. I thought CDS department was the best for this and approached Prof. Phaneendra Yalavarthy, who was the chair at that time. I was very happy to see his prompt and positive response on this. I had several discussions with Prof. Venkatesh Babu, Prof. Anirban Chakraborty, Prof. Govindan Rangarajan, Prof. Yogesh Simmhan and many other professors here. It was very encouraging to get positive feedback from them on the concept, research potential and the application of technology for the benefit of people. They were also eager to collaborate on this. We have been working under the guidance of Prof. Venkatesh Babu from the start and his support for Deep Learning related research is invaluable. Prof. Yogesh Simmhan is helping us with server side system scaling and optimization. It has been very helpful to get their guidance through this journey.

And finally, some food for thought...

Unsupervised Learning: Learning Without Labels

In modern data-driven deep learning paradigms, almost all the state-of-the-art approaches heavily rely on expensive labeled data. This has led to the aphorism that in machine learning, “it’s not who has the best algorithm that wins; it’s who has the most data.” There are many disadvantages of such a strong dependency on labeled information, most important of them being generalizability. Will such a trained model work well on the deployed environment. To account for such a domain-shift (shift in input domain distribution), it is often required to fine-tune the model with additional annotated data from the deployed environment. For example, a road scene segmentation model trained on cities in Germany fails miserably when deployed in San Francisco. This deficiency is often regarded as “dataset-bias”. In such cases, unsupervised learning offers elegant solutions to make use of the unlabelled samples from the deployed environment to extract relevant features, which can improve generalization performance.

Broadly, different learning techniques can be categorized based on two important factors, availability of teacher or expert, and interaction of the agent with the deployed environment. Interaction of the agent can be either active or passive depending on whether the environment reacts to the agent's decision. For example, rule-based game environments versus open-world (exploratory) game environments. Reinforcement learning techniques require access to a reward signal, which is a result of the allowed active interaction with the environment. However, unsupervised learning falls in the category of no access to a teacher, and also no active/ passive interaction with the environment. Hence it is also regarded as self-taught learning similar to the development of human intelligence.

Aiming towards human-level generalization, there is a need to explore adaptable representation learning methods with greater transferability.

	With Teacher	Without Teacher
Active	Reinforcement Learning / Active Learning	Intrinsic Motivation / Exploration
Passive	Supervised Learning	Unsupervised Learning

Transferability can be of two types, domain-transferability and task-transferability. Domain transferability is defined as the robustness of the learned representation across different deployment scenarios mitigating dataset-bias (e.g. synthetic to real environment). Similarly, task-transferability can be interpreted as the robustness of the learned representation across a diverse set of end tasks (e.g. multi-task models for both monocular depth estimation and semantic segmentation). The end goal in unsupervised learning is to obtain representations (or features) which can facilitate the highest level of generalization not only across diverse domains but also across new tasks. As rightly pointed out by Prof. Satinder Singh, “We need to stop thinking about teaching agents to learn tasks and start teaching them to learn **skills**.”

“If intelligence was a cake, unsupervised learning would be the cake, supervised learning would be the icing on the cake, and reinforcement learning would be the cherry on the cake.” - Yann LeCun

In unsupervised learning, the key idea is to recognize the massive degree of disparity between the volume of information between different types of signals (input, labels, and rewards). More specifically, the targets for supervised learning (complete set of labels for ImageNet $\sim \log_2(1000) \cdot 1.28M \rightarrow 12.8$ Mbits) has a lot less information as compared to the input data (ImageNet training set contains $\sim 1.28M$ images each of size $128 \times 128 \rightarrow 500$ Mbits).

Similarly, targets or rewards in reinforcement learning constitute an even lesser fraction of information (a long sequence of actions followed by a single bit of reward, e.g. win or lose only at the end of a Chess game). However, if we stop thinking about signals from the teacher and take into account the signal that is inherently present in the data, then we basically have access to an unlimited supply of the information (vast amount of data available on the internet) without worrying about any annotated information.

RETHINKING GENERALIZATION: Experiments establish that state-of-the-art convolutional networks for image classification trained with stochastic gradient methods easily fit a random labeling of the training data. This phenomenon is qualitatively unaffected by explicit regularization and occurs even if we replace the true images by completely unstructured random noise. Essentially the model (~30M weight) memorizes the labels even in the absence of any common pattern. However, the same network would not be able to memorize the training data if we train a CNN to do density modeling (a method of unsupervised learning). This motivates us to prefer unsupervised learning to achieve a higher degree of generalization.

An overview of different unsupervised learning techniques: In supervised learning, the loss that is minimized to train the deep model is a function of the model's prediction and the corresponding ground-truth label. Whereas, in unsupervised learning we only have access to the input samples X , without any predefined target or task to solve. This is what makes the learning process challenging, as we have to come up with something that indirectly allows us to learn useful representation from the given pool of unlabelled data.

1. Density modeling: Here, the goal is to learn the true distribution from which the data has been drawn. Basically here the model tries to exploit the general patterns or regularities present in the data. Recent generative modeling techniques such as Variational Auto-Encoder (VAE) and Generative Adversarial Networks (GAN) fall under this setting.




Fig-1: Context-encoder, Pathak et al. CVPR'16

2. Self-supervised learning: Different from density modeling where we rely on a generative model to learn reliable features, in self-supervised learning we first define a pretext task and then employ the Neural Network model to solve the proxy task to learn reliable representations. Various pretext tasks have been proposed for self-supervised learning including colorizing grayscale images, image inpainting (Fig-1), image jigsaw puzzle (Fig-2), etc.




Fig-2: Solving jigsaw puzzle, Noroozi et al.

The general pipeline in unsupervised learning is to first learn visual features by training Conv Nets to solve a predefined pre-text task on the unlabeled data. After obtaining a pretrained model from this process, the general trend is to show transfer learning performance on a downstream computer vision task, where only the last few layers are allowed to fine-tune on a set of annotated data (task specific). The performance on these downstream tasks is used to evaluate the quality of the learned features.

References:

Tutorial on Unsupervised Deep Learning - Google DeepMind & Facebook Artificial Intelligence NeurIPS 2018.

Jing, Longlong, and Yingli Tian. "Self-supervised Visual Feature Learning with Deep Neural Networks: A Survey." arXiv preprint arXiv:1902.06162 (2019).

Reference video:

<https://www.youtube.com/watch?v=rjZCjosEFpl>