भारतीय विज्ञान संस्थान

A CSR initiative by kotak Kotak Mahindra Bank

Kotak IISc AI–ML Centre

A Report of Activities
July-September 2024

Kotak IISc AI-ML Centre

CONTENTS

ABOUT THE CENTRE

The Kotak IISc AI–ML Centre (KIAC) was conceived of with the vision of creating a state-of-the-art artificial intelligence–machine learning (AI–ML) Centre at the Indian Institute of Science (IISc). The Centre was established at IISc with a CSR (corporate social responsibility) grant from the Kotak Mahindra Bank Limited (KMBL). A memorandum of understanding (MoU) between both the parties was executed from 1 September 2022 to 31 March 2025. The activities of the Centre commenced from September 2022.

OBJECTIVES

The main objectives of the Centre were chalked out as follows.

1. To develop rigorously-trained human resources who will become future leaders in AI–ML and champion India's surge in this area.
2. To conduct cutting-edge research in this area leading to international visibility of the highest order.
3. To develop innovative, deep tech solutions to meet the current and emerging requirements of not only the country but also the world.

The Centre's activities have been planned and executed in keeping with the above objectives.

ORGANISATION

The governance of the Kotak IISc AI–ML Centre is overseen by the following structure. The Convenor and his team of Professors are responsible for creating technical programmes that are in-line with the MoU between KMBL and IISc. These programmes are presented to the Scientific Advisory Board (SAB) to ensure that the conceived programmes are compatible with the MoU. The SAB reports to the Governing Board (GB). The Project Review Committee (PRC) also reports to the GB and reviews the overall functioning of the Centre.

The members of the three boards/committees have representatives from both IISc and KMBL and meet frequently to review the progress of the Centre and plan future initiatives.

GOVERNING BOARD

Govindan Rangarajan
IISc (Chair)

Navakanta Bhat
IISc

Yadati Narahari
IISc

Kaushal Verma
IISc

Milind Nagnur
KMBL

Himanshu Nivsarkar
KMBL

SCIENTIFIC ADVISORY BOARD

Rajesh Sundaresan
IISc (Chair)

Rajiv Soundararajan
KIAC Convenor

Srikanth Krishnan Iyer
IISc

Bhaskar Kumar
KMBL

PROJECT REVIEW COMMITTEE

S K Satheesh
IISc (Chair)

P S Anil Kumar
IISc

Rajesh Sundaresan
IISc

Himanshu Nivsarkar
KMBL

Prem Sagar Raju Addala
KMBL

KIAC TEAM FACULTY

Rajiv Soundararajan
Convenor, KIAC

Soma Biswas
Professor, IISc

Sriram Ganapathy
Professor, IISc

Chandrasekaran Pandu Rangan
Chair Visiting Professor, KIAC

Viraj Kumar
Visiting Professor, KIAC

KIAC TEAM STAFF

Grace Mathew Abraham
Program Manager

Neetha Ashtakar
Systems Administrator

Geethanjali Monto
Senior Editorial Assistant

Sudha Aithal
Secretary

Manoj Srinivas
Office Assistant

KIAC: JULY–SEPTEMBER 2024

The activities and new initiatives of the Centre in this quarter (July–September 2024) are described in this section. These activities are presented as achievements under the three objectives of KIAC.

On 2 July 2024, delegates from KMBL, Chetan Savla and Prem Sagar Raju Addala, met Govindan Rangarajan, Director of IISc to talk about the way forward for the Centre. They also had a discussion with IISc faculty associated with the Centre.

Meeting with the Director, IISc

Discussion with IISc faculty

Objective 1: Rigorously-trained human resources

In the recently-concluded Asian Engineering Deans' Summit held between 21–22 May 2024 at IISc, David J Srolovitz (Chair, Steering Committee for the 2024 Summit and Dean, Faculty of Engineering, The University of Hong Kong) pointed out, "Since it is already being used in industry, and its use in the industry is growing so rapidly, our responsibility to our students is to make sure they are up to date with the tools of the trade"[1], in his talk on integrating AI into undergraduate engineering education. With the advent of Industry 4.0 and smart manufacturing incorporating technologies such as Internet of Things (IoT), cloud computing, artificial intelligence, and machine learning, the expectations from engineering graduates have changed.[1]

In addition, artificial intelligence is increasingly playing major roles in domains such as healthcare, robotics, education, speech, image, and video analysis. Graduates in these domains who are well versed in AI–ML are in a better position to be selected for high-profile jobs both in academia and industry.

In this context, KIAC has been supporting the academic and research work of many IISc students and candidates from outside the Institute.

Student support

During this quarter, seven students were supported through the Kotak scholarship. This scholarship is awarded to the women students with the top-most JEE Advanced Ranks joining the undergraduate BTech (Mathematics and Computing) programme at IISc. The current awardees are:

1. Sasmita Harini S
2. Shivey Ravi Guttal
3. Lekhya Pillarikuppam
4. Palak Raisinghani
5. Rajul Mahawar
6. M Lasya Priya
7. Durga Naniwadekar

The students supported through the MTech scholarship are:

1. Anmol Gill
2. Himanshu Jain
3. G Pavithra
4. Gaurav Kumar
5. Shashwat
6. Vedantam Srikar
7. Akash Singh

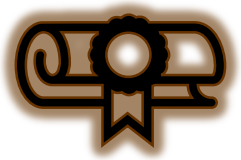

The PhD scholarship awardees are:

1. Shubhankar Gupta
2. Aditya Shirwatkar
3. Priyam Dey
4. Kumar Shubham
5. Rankit Kachroo
6. Tirthajit Baruah

Research support

In this quarter, two persons were supported through the postdoctoral fellowship. They are:

1. Surbhi Shaw
2. Akash Pareek

Thirteen interns and 30 predoctoral fellows were supported in this quarter. These candidates work in myriad research areas ranging from emerging core AI–ML objectives to applications in speech, image, and video analysis; healthcare; robotics; finance; and generative AI. A glimpse into their research projects is provided under Objective 2 in this report.

Impact of the support

The support has helped the candidates (PhDs, interns, and predoctoral fellows) carry out cutting-edge research that has translated into research papers published in top-level journals and presented at A* conferences. This also leads to the fulfilment of the second objective of the Centre, namely, 'To conduct cutting-edge research in this area leading to international visibility of the highest order'. Some testimonials are presented below.

Kumar Shubham
PhD student, Department of Electrical and Communication Engineering, IISc
Recipient of KIAC PhD fellowship and support to present at ICML 2024

"As a KIAC Fellow, I am the recipient of this year's PhD fellowship. I also got a travel grant for the International Conference on Machine Learning in Vienna, Austria. Because of the grant, I was able to connect with a lot of renowned professors working in the same field and get their feedback on my work.

At ICML, I presented my work on cancer and got multiple feedback from my peers and also got to understand the limitations. Additionally, I attended a lot of workshops and sessions, which helped me in understanding the limitations of the current work and also in understanding the most recent technologies that are being used for the same task.

I am really thankful to KIAC for helping me out financially to travel to Vienna and attend the conference."

Shubhankar Gupta
PhD student, AI and Robotics Lab,
Department of Aerospace Engineering, IISc
Recipient of KIAC PhD fellowship

"I think the KIAC fellowship is a great initiative. It is a prestigious thing to have on your resume, first of all, and this fellowship is highly competitive, which really motivates us to work hard and push the boundaries of AI–ML research and submit papers to top journals and conferences. For example, after getting the KIAC fellowship, we have sent two papers to top journals in the world, one is JMLR and the other is TMLR. We are also planning to submit to top conferences in the coming three to four months like AAAI, ICRA, and ICLR.

This initiative not only motivates me, but I think it will motivate all the students to do some competitive research in AI–ML and push the boundaries of AI research."

KIAC alumni have gone forward to take up various positions in academia and the industry.

For example, Ashish Ramayee Asokan (predoctoral fellow) has joined for his PhD at the Carnegie Mellon University, Pennsylvania, and Sai Harsha Mupparaju (predoctoral fellow) is doing his Masters at the NYU Tandon School of Engineering.

Courses offered by KIAC faculty

Distinguished KIAC faculty, Viraj Kumar and Chandrasekaran Pandu Rangan, offer courses on a regular basis for the BTech and MTech students of IISc. In this quarter, Viraj Kumar co-taught a course on 'Algorithms and Programming', which was attended by 140 BTech students. Pandu Rangan taught 'Data Structures and Algorithms' for 66 BTech students; the lab component of the course was handled by Viraj Kumar. Another course was taught by Pandu Rangan for the MTech (AI) students on 'Data Structures and Algorithms', with 68 attendees.

Support to participate in/conduct competitions

The undergraduate students are also encouraged and supported to take part in competitions.

iGEM competition

In the iGEM competition, multidisciplinary student teams from across the world design, build, and test projects using cutting-edge synthetic biology. The student team from IISc, named 'IISc-Software', has 13 BTech students, with a PhD student as their student guide, Debnath Pal (IISc) as the principal investigator (PI) and secondary PIs as Shruthi Viswanath (National Centre for Biological Sciences, Bengaluru) and Manjula Das (Mazumdar Shaw Center for Translational Research, Bengaluru).

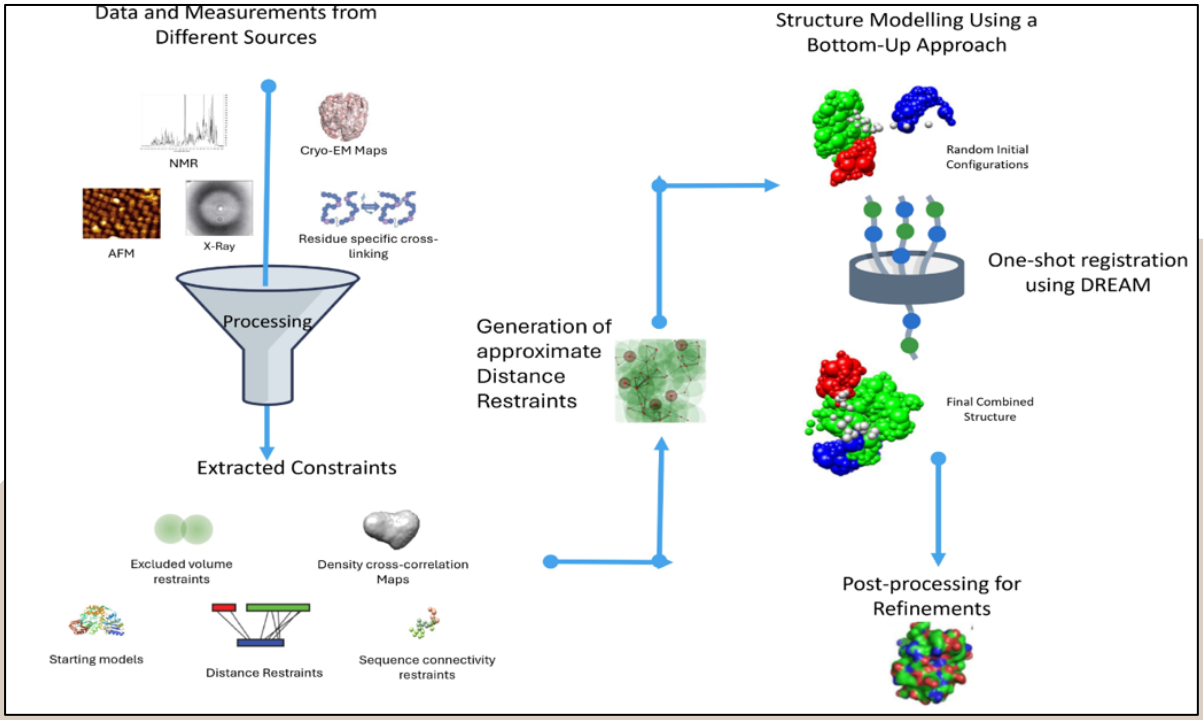

Their project aims to use the integrative modelling platform (IMP) with the DREAM algorithm to build models of macromolecular assemblies using a bottom-up approach, overcoming existing bottlenecks
[details at https://2024.igem.wiki/iisc-software/description].

Ideathon 2024

The Databased Club of IISc, a community of BTech students passionate about diving deep into the world of computer science and exploring the domain of mathematics, conducted 'Ideathon 2024' for the first-year undergraduate students, providing them with a platform to showcase their innovation and problem-solving skills. Over 20 student teams participated in the two-day event, where they had to develop solutions for one of two intriguing problem statements: (i) sustainable campus transportation or (ii) AI-driven early warning system for natural disasters.

The participants were given 24 hours to brainstorm and propose solutions to one of the two problem statements. At the end of the duration, all teams had to upload a PPT file for their presentations to finalise their submission (consisting of the problem identification, proposed solution, implementation and challenges, and impact of the solution). The participants were encouraged to utilise AI-based solutions to tackle the problem statements. All the teams presented their solutions to the judging panel on Day 2, where their work was judged based on criteria such as research, novelty, and feasibility.

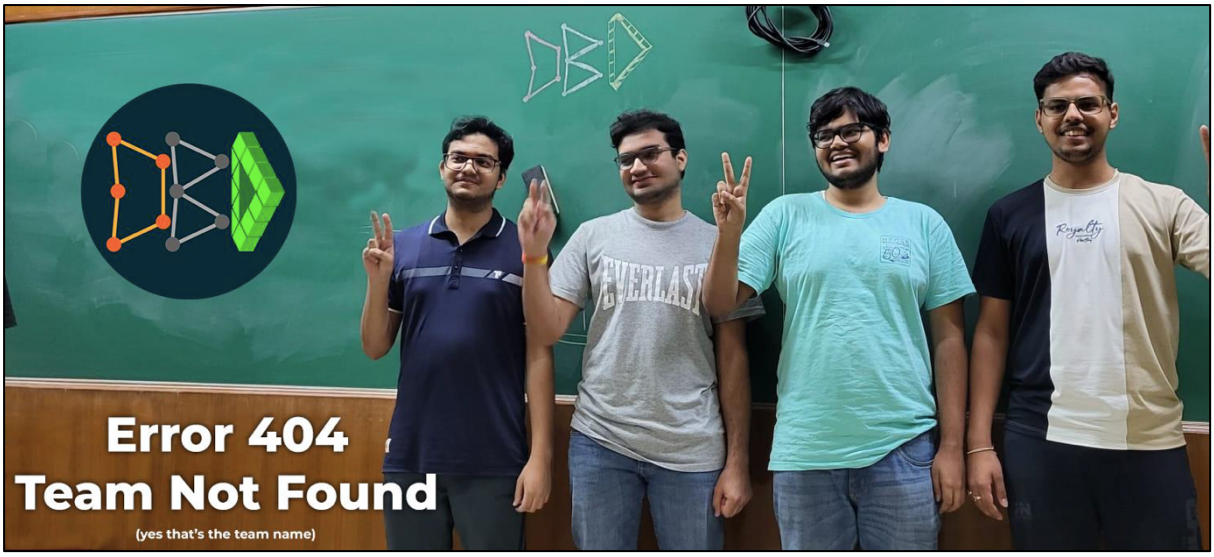

The winning presentation was given by the team ***Error 404 Team Not Found***. They took on the problem statement for 'AI-driven early warning system for natural disasters'. Their solution featured a parametrised approach to utilising a collection of datapoints such as seismic data, ionospheric data, and animal behaviour to predicting natural disasters. They also highlighted many lessons that can be learnt from studies on previous calamities faced by humanity and how existing networking infrastructure can be utilised to deploy better communication systems so that people can be informed of predicted disasters.

The event also featured two talks. The talk by Pandu Rangan Chandrasekaran, Chair Visiting Professor at KIAC, on 'What is Information?', sparked deep reflections on the nature of data, information, and how they intertwine with problem-solving and computational thinking. It introduced freshers to ideas and motivation for information theory and mathematical modelling of information.

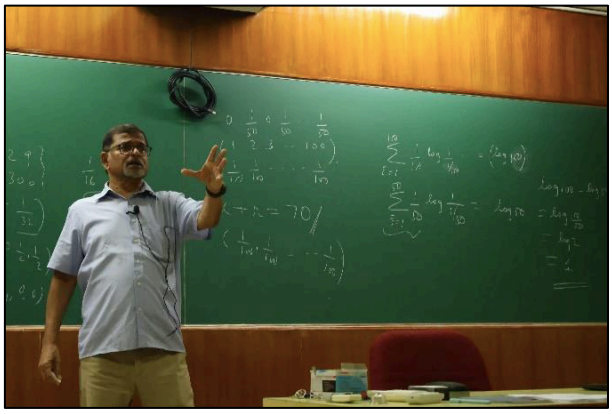

The guest lecture by Sunny Manchanda, Director of the Young Scientist AI Lab at DRDO, was an insightful session on 'Applications of AI'. His talk emphasised the transformative potential of AI, helping participants appreciate the vast possibilities that AI can offer in creating innovative solutions for the real world.

The event concluded with enthusiastic presentations of proposed solutions, displaying the brilliance and creativity of the participants. The efforts of the top teams were recognised during the award ceremony and every participant walked away with valuable experience and knowledge.

Information dissemination and training

In addition to actively supporting research, KIAC facilitates the exposure of students to state-of-the-art research being carried out in academia and industry across the world. Information and training have been imparted to individuals from across the country and the world through AI–ML talks, courses, training sessions, and workshops. Recently, a special seminar series has been initiated from September 2024, named as the 'KIAC Distinguished Seminar Series'. The objective is to invite eminent researchers from all over the country to share their expertise and experiences.

KIAC Distinguished Seminar Series

KIAC launched its Distinguished Seminar Series with a talk by Mitesh M Khapra, Associate Professor in the Department of Computer Science and Engineering and Head of the AI4Bharat Research Lab at IIT Madras.

Title: Towards transforming the Indian language technology ecosystem
Speaker: Mitesh M Khapra
Date: 20 September 2024
Number of attendees: 114

In this talk, Mitesh M Khapra reflected on his team's journey towards transforming the landscape of Indian language technology. He delved into their engineering-heavy approach in addressing the initial scarcity of data for Indian languages, while gradually establishing the necessary human resources to gather high-quality data on a larger scale through *Bhashini*. The objective was to share insights into developing high quality open-source technology for Indian languages. This involves curating extensive data from the internet, constructing multilingual models for transfer learning, and crafting high-quality datasets for fine-tuning and evaluation. He then transitioned into how their experiences can benefit the broader AI community, particularly as India aspires to create large language models (LLMs) for Indic languages.

Bio of the speaker: Mitesh M Khapra is an Associate Professor in the Department of Computer Science and Engineering at the Indian Institute of Technology Madras (IITM). He heads the AI4Bharat Research Lab at IIT Madras which focuses on building datasets, tools, models, and applications for Indian languages. His research work has been published in several top conferences and journals including TACL, ACL, NeurIPS, TALLIP, EMNLP, EACL, and AAAI. He has also served as Area Chair or Senior PC member in top conferences such as ICLR, ACL, and AAAI. Prior to IIT Madras, he was a Researcher at IBM Research India for four and a half years, where he worked on several interesting problems in the areas of statistical machine translation, cross language learning, multimodal learning, argument mining, and deep learning. Prior to IBM, he completed his PhD and MTech from IIT Bombay in January

2012 and July 2008, respectively. His PhD thesis dealt with the important problem of reusing resources for multilingual computation. During his PhD, he was a recipient of the IBM PhD Fellowship (2011) and the Microsoft Rising Star Award (2011). He is also a recipient of the Google Faculty Research Award (2018), the IITM Young Faculty Recognition Award (2019), the Prof B Yegnanarayana Award for Excellence in Research and Teaching (2020), and the Srimathi Marti Annapurna Gurunath Award for Excellence in Teaching (2022).

Kotak IISc AI–ML Talk Series

The Kotak IISc AI–ML talk series facilitates exposure of students to the research work and state-of-the-art in their respective fields of expertise. It includes talks by eminent scientists and researchers from academia and the industry from all over the world, with whom IISc students and faculty can interact and collaborate on various projects. The talks are held at IISc and open to all; non-IIScians can attend upon request.

Two talks were organised during July–September. The first speaker was from the University of California and spoke on 'Scene Understanding for Safe and Autonomous Navigation' while the second speaker was from Harvard University and described 'Protein Sequence Annotation using Language Models'.

Protein Sequence Annotation using Language Models

Speaker: Kumaresh Krishnan, Postdoctoral Fellow, The Eddy Laboratory, Harvard University, USA
Date: 05 August 2024
Number of attendees: 21

Protein function inference relies on annotating protein domains via sequence similarity, often modelled through profile hidden Markov models (profile HMMs), which capture evolutionary diversity within related domains. However, profile HMMs make strong simplifying independence assumptions when modelling residues in a sequence. In this talk, the speaker introduced PSALM (protein sequence annotation with language models), a hierarchical approach that relaxes these assumptions and uses representations of protein sequences learned by protein language models to

enable high-sensitivity, high-specificity residue-level protein sequence annotation. He validated PSALM's performance on a curated set of 'ground truth' annotations determined by a profile HMM-based method and highlighted PSALM as a promising alternative for protein sequence annotation.

Bio of the speaker: Kumaresh Krishnan is a Postdoctoral Fellow in The Eddy Laboratory at Harvard University. He works on machine learning models for annotating, understanding, and analysing protein sequences. He has a PhD from Harvard University where he worked in systems neuroscience, building models of decision making and attentional switching using zebrafish as a model organism. Krishnan's undergraduate and Master's training is in Computer Science and Electrical Engineering from the International Institute of Information Technology (IIIT) Bangalore and he brings this strong computational background to tackle complex real world biological problems.

Scene Understanding for Safe and Autonomous Navigation

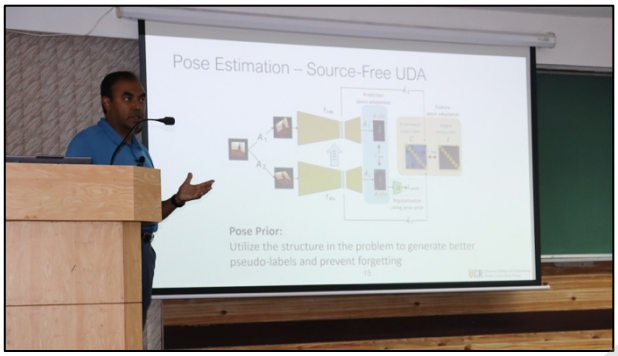

Speaker: Amit K Roy-Chowdhury, Professor, University of California
Date: 18 July 2024
Number of attendees: 38

Autonomous navigation remains one of the most challenging problems in intelligent systems largely because of the close integration of scene understanding and planning that needs to happen. The scene understanding requires analysis of objects and their collections across various scales, from individual people and their actions to wide-area analysis that could span the interactions of these people with many other objects in the scene. An integrated view that is able to span across these ranges of scale is necessary for robust decision making. In this talk, Amit K Roy-Chowdhury considered a variety of scene understanding problems that need to be solved for autonomous navigation to be successful. At the level of individual people, he showed how to estimate the pose of each individual person under challenging real-life conditions such as significant occlusions. At the next higher scale when there are interactions among small groups of individuals and objects, he demonstrated the power of scene graphs to model the semantics of the scene. At a yet higher level, he showed how to track objects across non-overlapping cameras spread over large areas. Robustness to a variety of

operational domains was considered through all of these tasks. In spite of this, it is unlikely that perfect scene understanding will be achieved and any autonomous agent will need to occasionally interact with human experts; he showed how this can be achieved with natural language feedback leveraging upon the power of recently developed vision–language models.

Bio of the speaker: Amit K Roy-Chowdhury received his PhD from the University of Maryland, College Park (UMCP) in 2002 and joined the University of California, Riverside (UCR) in 2004 where he is a Professor and Bourns Family Faculty Fellow of Electrical and Computer Engineering, Cooperating Faculty in Computer Science and Engineering, and Director of the Center for Robotics and Intelligent Systems. He leads the Video Computing Group at UCR, working on foundational principles of computer vision, image processing, and machine learning, with applications in cyber-physical, autonomous and intelligent systems. He has published over 200 papers in peer-reviewed journals and conferences. He has published two monographs: Camera Networks: The Acquisition and Analysis of Videos Over Wide Areas and Person Re-identification with Limited Supervision. He is on the editorial boards of major journals and program committees of the main conferences in his area. He is a Fellow of the IEEE and IAPR, and received the Doctoral Dissertation Advising/Mentoring Award from UCR, and the ECE Distinguished Alumni Award from UMCP.

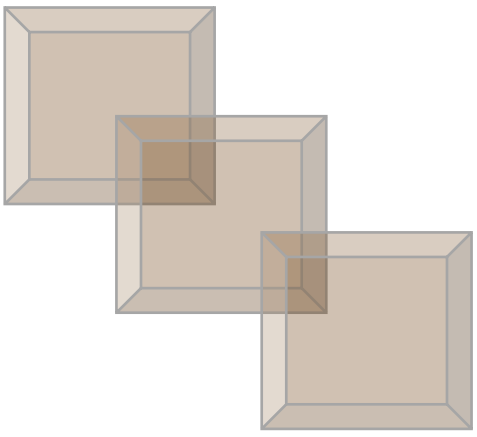

Courses and training sessions by KIAC faculty

Distinguished KIAC faculty offer courses and training sessions for students and faculty.

Pandu Rangan, Chair Visiting Professor at KIAC, taught an online course on 'Introduction to Cryptography' for 42 MTech students at the Indian Institute of Technology Madras. He also conducted an NPTEL (National Programme on Technology Enhanced Learning) course on 'Introduction to Graph Algorithms'; the course was well attended with 1990 participants.

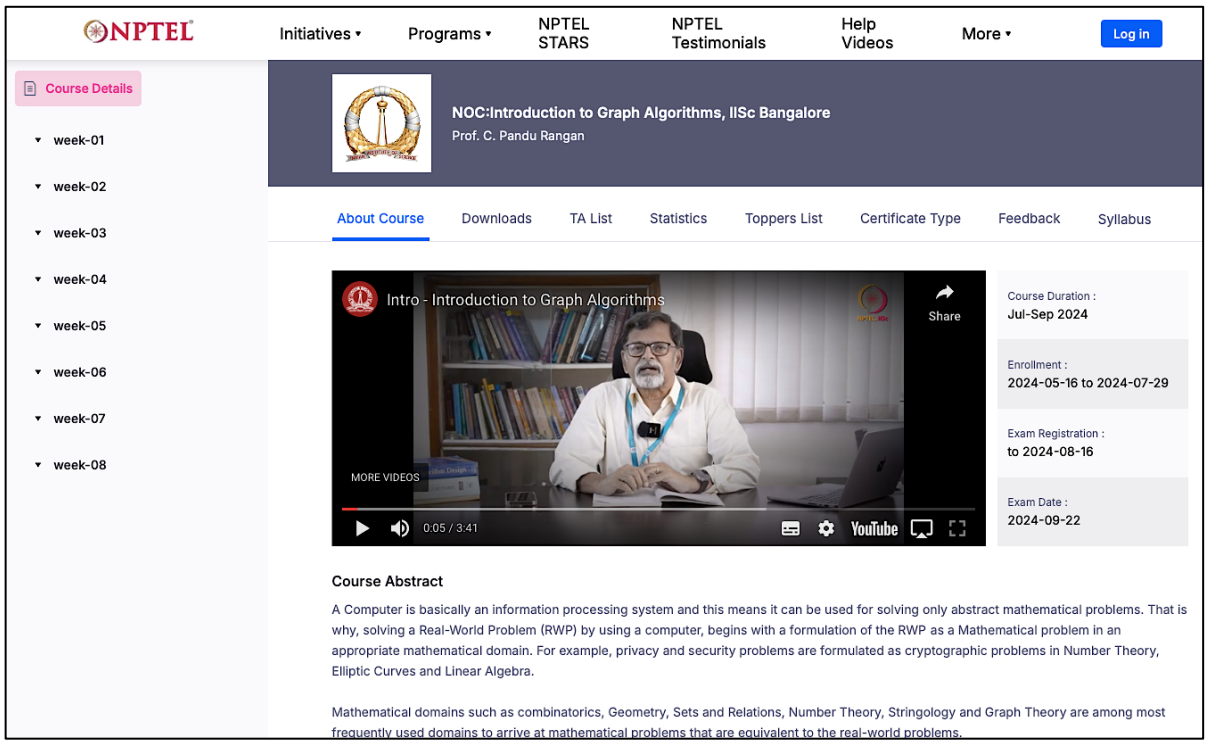

Viraj Kumar, Visiting Professor at KIAC, initiated a six-month faculty development programme (from July 2024 onwards) through an IISc–AICTE (All India Council for Technical Education) partnership. This course targeted faculty in non-computing disciplines. In August 2024, he conducted two one-day faculty training programmes for 50 participants each at (i) M S Ramaiah Institute of Technology, Bengaluru and (ii) Nitte Meenakshi Institute of Technology, Bengaluru.

In addition, Viraj Kumar Presented an ACM India webinar on 'Simple code-writing problems that GenAI cannot solve (at present)' on 26 July 2024 [https://www.youtube.com/watch?v=_yKxTWb7POk] and presented a talk at the IIT Jammu Foundation Program-2024 on 'Learning in the Shadow of “free” GenAI Tools' on 4 August 2024.

Workshops

The Centre supports IISc faculty in conducting workshops in multiple disciplines, thus enabling information dissemination and training to participants from across India and the world. During this quarter, KIAC supported two workshops, as described below.

Cyber-Physical Systems Symposium (CyPhySS 2024)

Date: 25–27 July 2024
Venue: Indian Institute of Science, Bengaluru
Number of participants: 58

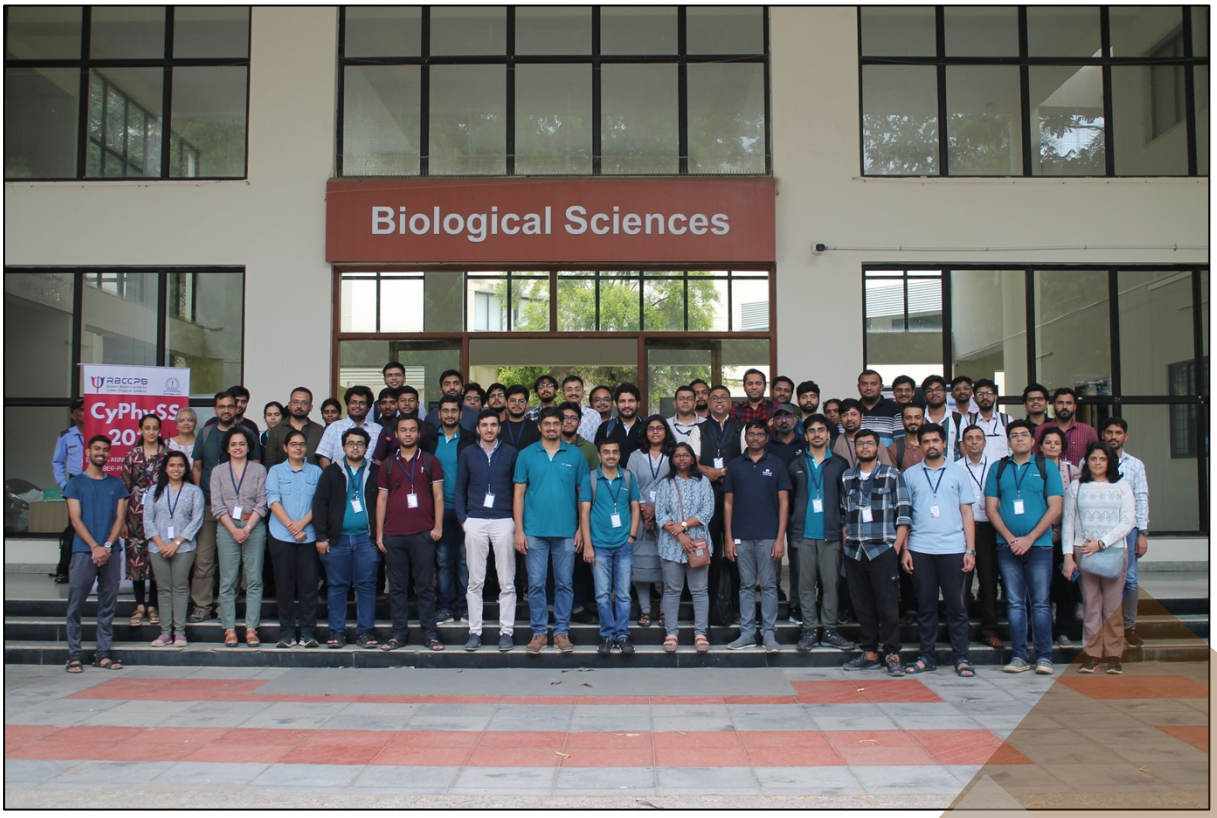

The Cyber-Physical Systems Symposium (CyPhySS), India's largest annual summit focussing on cyber-physical systems (CPS), serves as a pivotal platform for researchers, industry leaders, and academic professionals, both domestic and international, to discuss the latest advancements, research endeavours, and best practices in robotics, control, and artificial intelligence as applied to various sectors including industry, space, agriculture, and socio-technical systems such as energy, mobility, smart grid, and smart cities.

The theme of the eighth Cyber-Physical Systems Symposium (CyPhySS 2024) was 'Robotics, Control and AI' and was hosted by the Robert Bosch Centre for Cyber-Physical Systems at IISc during 25–27 July 2024. The three-day event had many talks, public lectures, tutorial sessions, student presentations, and poster presentations.

The speakers were from around the world and from both academia and industry. The keynote speeches were on 'Security of cyber-physical systems: theory and applications of the dynamic watermarking method' by P R Kumar (Distinguished Professor, Texas A&M University, USA), 'Resilience and distributed decision-making in a renewable-rich power grid' by Anuradha Annaswamy (Professor, Massachusetts Institute of Technology, USA), and 'Bridging the innovation gap: the ongoing story of ARTPARK' by Bharadwaj Amrutur (Professor, IISc and Executive Director, ARTPARK).

There were three tutorial sessions with speeches by Adnane Saoud (Associate Professor, University Mohammed VI Polytechnic, Morocco) on 'Assume-guarantee contracts for control: from continuous to hybrid systems and from safety to temporal logic specifications', by Mayank Baranwal (Senior Scientist, TCS Research) on 'Reinforcement learning: connecting AI innovations to practical solutions', and by Hrishikesh Venkataraman (Associate Professor, IIIT Sri City) on 'Smart transportation and hard real-time cyber physical system: challenges and opportunities'. Some of the invited speakers were Mini C Rai (Professor, University of Lincoln, UK) who spoke on 'In-orbit assembly missions: advancements in robotics, automation and AI' and Arjun Jain (Founder and CEO, Fast Code AI) who spoke on 'Autonomy at scale: where are we headed?'.

The student speakers were from across India and spoke on topics such as 'Three-dimensional nonlinear impact time guidance considering field-of-view constraints', 'Graph-based prediction and planning policy network (GP3Net) for scalable self-

driving in dynamic environments using deep reinforcement learning', 'MRFP: learning generalizable semantic segmentation from Sim-2-Real with multi-resolution feature perturbation', and 'Design and development of a Web-APP based user interface to control and operate an automated mobile robot'.

International Conference on Signal Processing and Communications (SPCOM 2024)

Date: 1–4 July 2024
Venue: Indian Institute of Science, Bengaluru
Number of participants: 115

The 2024 International Conference on Signal Processing and Communications (SPCOM) provides a forum for researchers from academia, research laboratories, and industry to come together to share and learn about current developments in these emerging fields. SPCOM 2024 is the fifteenth in a series of biennial events that have been organised since 1990 at IISc. It was hosted by the Department of Electrical Communication Engineering, IISc during 1–4 July 2024.

This year's conference featured 3 plenaries, 5 tutorials, 20 invited talks, and around 75 contributed papers. The first plenary talk was on 'Transformers for compression and communication' by Krishna Narayanan (Eric D Rubin '06 Professor in Electrical and Computer Engineering, Texas A&M University). Maneesh Agrawala (Forest Baskett

Professor of Computer Science and Director of the Brown Institute for Media Innovation at Stanford University) delivered the second plenary on 'Unpredictable black boxes are terrible interfaces'. The third plenary was by Piya Pal (Associate Professor, ECE, UC San Diego); she spoke on 'Sensing, coding and the quest for super-resolution'. The industry keynote was by Murthy S, Application Engineer (RF and Wireless) at VVDN Technologies, Bengaluru; he spoke on '6G vision and key technologies'.

There were six invited sessions featuring distinguished speakers in the broad areas of information theory, communications, networking, speech and language, applied machine learning and computer vision, and medical imaging. Examples of these talks are 'Artificial general intelligence (AGI)-native wireless systems with common sense: a journey to 6G and beyond', 'Fundamentals of vision-based geolocation', 'Incentivising client participation in federated learning', 'Evaluating LLMs on languages beyond English: challenges and opportunities', and 'Generative models for simulating medical imaging from ultrasound to CT via MRI'.

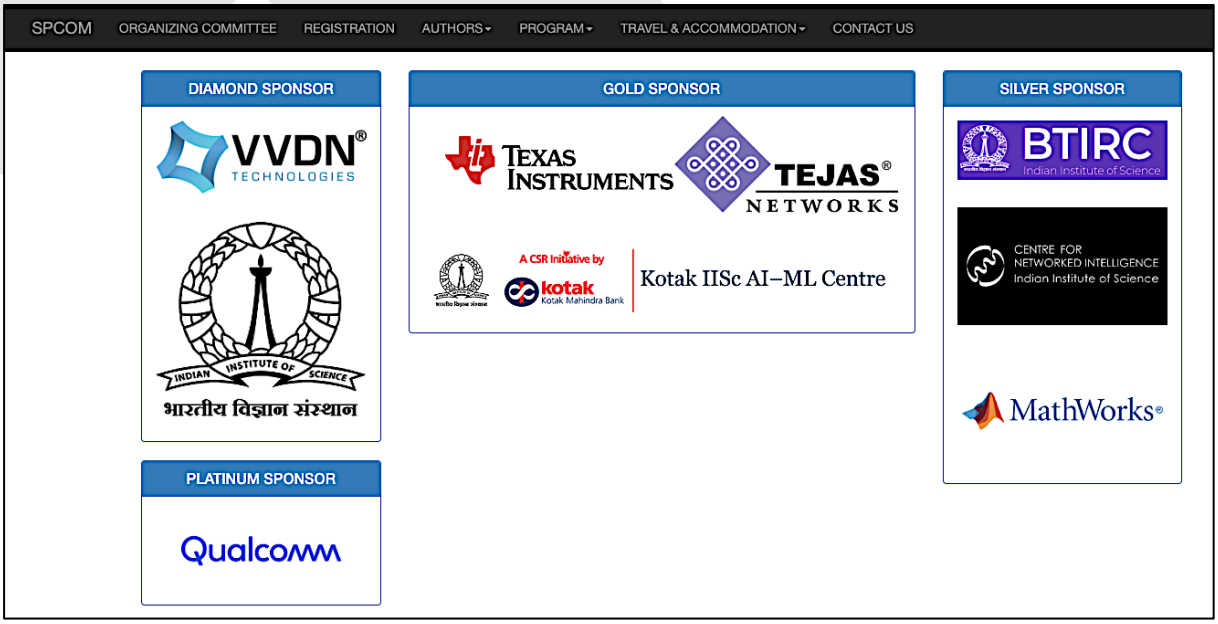

There were five tutorials: (i) Approximate message passing algorithms for high dimensional estimation; (ii) Theory and experiments for peer review and other distributed human evaluations; (iii) The unlimited sensing framework: theory, algorithms, hardware and applications; (iv) IEEE 802.11ad based intelligent joint radar communication transceiver: design, prototype and performance analysis; and (v) Applying artificial intelligence effectively to signal processing applications. These sessions were conducted by experts from India and across the world.

Objective 2: Cutting-edge research of international visibility

According to a study by AI accelerator and ecosystem builder Change Engine, 'The country stands 14th in AI research with a global share of just 1.4% (2018–2023) in terms of paper contribution in the top 10 AI conferences in the world'[2]. The study indicates that India's AI paper growth rate has flattened, while Asian economic tiger countries such as Hong Kong, Singapore, and South Korea have accelerated their growth rate and that India needs to grow AI research by at least 40% a year to get to a 5% global share in five years and create a world-class innovation ecosystem.[2] This can be achieved by extending support for researchers in terms of infrastructure availability and funding (including travel support to attend leading international conferences).

The work described in the following *Times of India* article was presented (with support from KIAC) at the Conference on Computer Vision and Pattern Recognition (CVPR 2024) held at Seattle, USA during 17–21 June 2024 by student Rishubh Parihar. He is one of five people who were supported to attend CVPR 2024.

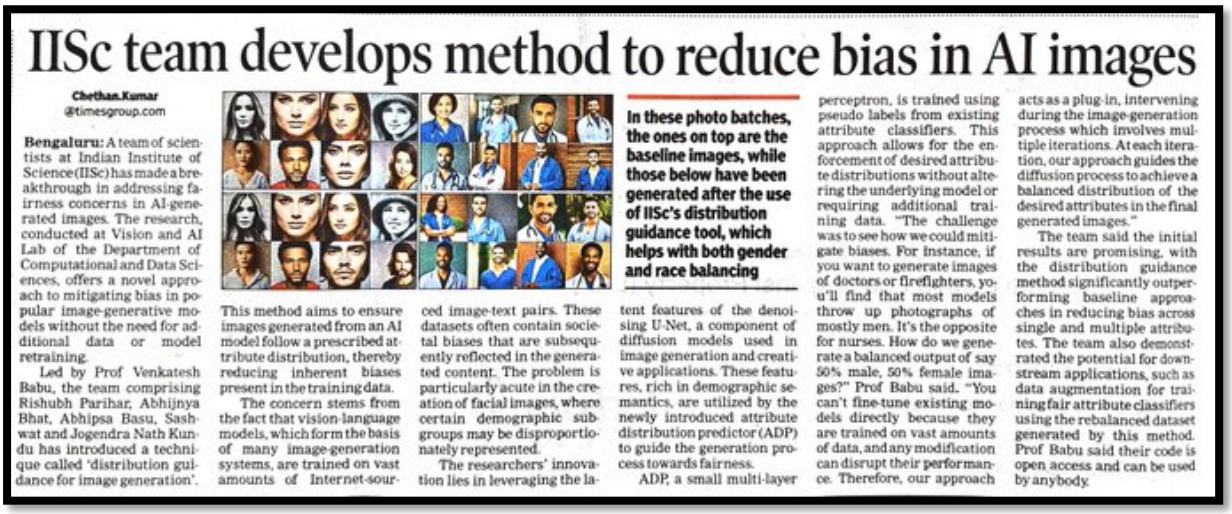

IISc team develops method to reduce bias in AI images

Chethan.Kumar
@timesgroup.com

Bengaluru: A team of scientists at Indian Institute of Science (IISc) has made a breakthrough in addressing fairness concerns in AI-generated images. The research, conducted at Vision and AI Lab of the Department of Computational and Data Sciences, offers a novel approach to mitigating bias in popular image-generative models without the need for additional data or model retraining.

Led by Prof Venkatesh Babu, the team comprising Rishubh Parihar, Abhijnya Bhat, Abhipsa Basu, Saswat and Jogendra Nath Kundu has introduced a technique called 'distribution guidance for image generation'.

In these photo batches, the ones on top are the baseline images, while those below have been generated after the use of IISc's distribution guidance tool, which helps with both gender and race balancing

This method aims to ensure images generated from an AI model follow a prescribed attribute distribution, thereby reducing inherent biases present in the training data.

The concern stems from the fact that vision-language models, which form the basis of many image-generation systems, are trained on vast amounts of Internet-sourced image-text pairs. These datasets often contain societal biases that are subsequently reflected in the generated content. The problem is particularly acute in the creation of facial images, where certain demographic subgroups may be disproportionately represented.

The researchers' innovation lies in leveraging the latent features of the denoising U-Net, a component of diffusion models used in image generation and creative applications. These features, rich in demographic semantics, are utilized by the newly introduced attribute distribution predictor (ADP) to guide the generation process towards fairness.

ADP, a small multi-layer perceptron, is trained using pseudo labels from existing attribute classifiers. This approach allows for the enforcement of desired attribute distributions without altering the underlying model or requiring additional training data. "The challenge was to see how we could mitigate biases. For instance, if you want to generate images of doctors or firefighters, you'll find that most models throw up photographs of mostly men. It's the opposite for nurses. How do we generate a balanced output of say 50% male, 50% female images?" Prof Babu said. "You can't fine-tune existing models directly because they are trained on vast amounts of data, and any modification can disrupt their performance. Therefore, our approach acts as a plug-in, intervening during the image-generation process which involves multiple iterations. At each iteration, our approach guides the diffusion process to achieve a balanced distribution of the desired attributes in the final generated images."

The team said the initial results are promising, with the distribution guidance method significantly outperforming baseline approaches in reducing bias across single and multiple attributes. The team also demonstrated the potential for downstream applications, such as data augmentation for training fair attribute classifiers using the rebalanced dataset generated by this method. Prof Babu said their code is open access and can be used by anybody.

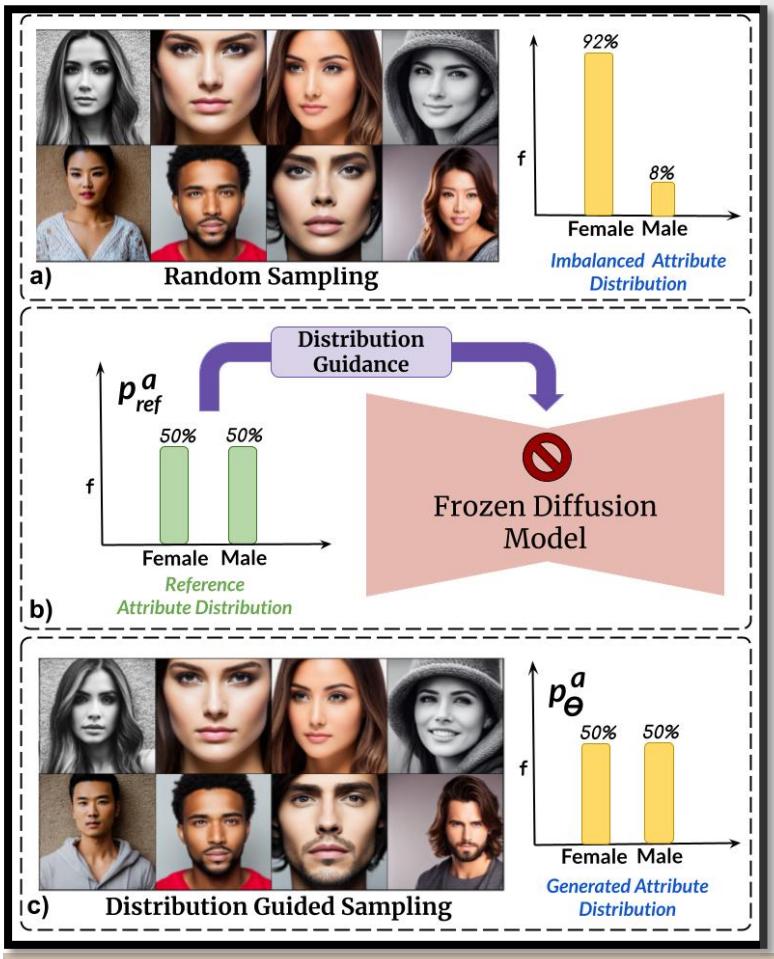

The paper that was presented, titled 'Balancing Act: Distribution-Guided Debiasing in Diffusion Models', put forth a method for debiasing DMs without relying on additional data or model retraining. *This method reduces bias across single/multiple attributes and outperforms the baseline by a significant margin for unconditional and text-conditional diffusion models*. Further, a downstream task of training a fair attribute classifier by rebalancing the training set with generated data was presented.

Presenting accepted papers at A* conferences

In this quarter, KIAC supported eight students/faculty from IISc to present their papers in the following A* conferences:

1. 18th European Conference on Computer Vision (ECCV 2024) at Milan, Italy during 29 September–4 October 2024
2. 62nd annual meeting of the Association for Computational Linguistics (ACL 2024) at Bangkok, Thailand during 11–16 August 2024
3. SIGGRAPH 2024, the premier conference and exhibition on computer graphics and interactive techniques, at Denver during 28 July–1 August 2024
4. 41st International Conference on Machine Learning (ICML 2024) at Vienna, Austria during 21–27 July 2024

ACCEPTED PAPERS PRESENTED AT LEADING INTERNATIONAL CONFERENCES IN ARTIFICIAL INTELLIGENCE		
CONFERENCE	**PAPER TITLE**	**AUTHORS**
ECCV 2024	Text2Place: Affordance-aware Text Guided Human Placement	Rishubh Parihar, Harsh Gupta, Sachidanand and R. Venkatesh Babu
	PreciseControl: Enhancing Text-To-Image Diffusion Models with Fine-Grained Attribute Control	Rishubh Parihar, Sachidanand, Sabariswaran, Tejan and R. Venkatesh Babu
ACL 2024	Evaluating Large Language Models for Health-related Queries with Presuppositions	Navreet Kaur, Monojit Choudhury, Danish Pruthi
SIGGRAPH 2024	Factorized Motion Fields for Fast Sparse Input Dynamic View Synthesis	N. Somraj, K. Choudhary, S. Mupparaju, and R. Soundararajan
ICML 2024	Testing the Feasibility of Linear Programs with Bandit Feedback	Aditya Gangrade, Aditya Gopalan, Venkatesh Saligrama, Clayton Scott
	Unsupervised Parameter-free Simplicial Representation Learning with Scattering Transforms	Hiren Madhu, Sravanthi Gurugubelli, Sundeep Prabhakar Chepuri
	Risk Estimation in a Markov Cost Process: Lower and Upper Bounds	Gugan Thoppe, Prashanth L. A., Sanjay Bhat
	WISER: Weak supervision and supervised representation learning to improve drug response prediction in cancer	Kumar Shubham, Aishwarya Jayagopal, Syed Mohammed Danish, Prathosh AP, Vaibhav Rajan

Attending leading international conferences presents a great opportunity for researchers from India to collaborate with the global community of academicians in their respective fields and leads to collaborations in multiple projects. The papers presented at ICML 2024 were among the 27.5% of accepted long papers in the conference. The public response to the KIAC LinkedIn post is encouraging.

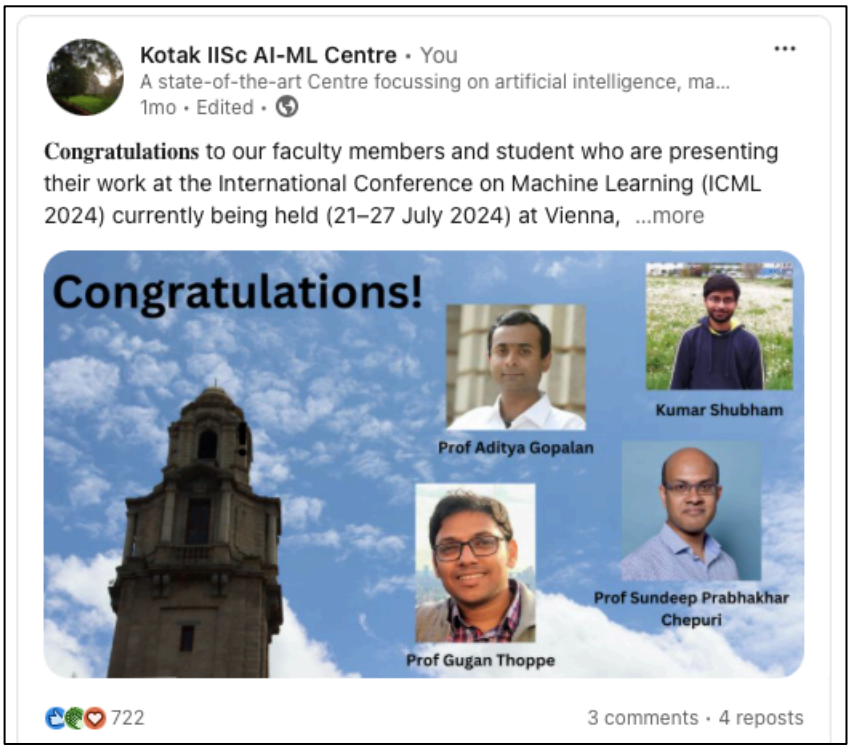

Research projects supported by KIAC

The research supported by KIAC—including the work done by doctoral students, interns, and predoctoral students—encompasses a combination of artificial intelligence and multiple disciplines such as robotics, healthcare, and finance. This is in addition to progress in core domains such as machine learning. Examples of research supported by KIAC, including those presented at leading AI conferences, are highlighted below. The research work falls into six themes: (i) generative AI; (ii) speech, image, and video analysis; (iii) emerging objectives in machine learning; (iv) AI in healthcare; (v) AI in robotics; and (vi) AI in finance.

(i) Generative AI

Generative AI refers to artificial intelligence systems that can generate new, original content from the data that they are trained on. These AI systems access the training data to identify patterns and structures, using machine learning models that mimic human neural networks. The quality of the output depends on the quality of the data and the prompts given by users. The new content can be text, images, sounds, music, animation, videos, and three-dimensional models.

Example of KIAC project: *deepfake detection*

Recently, the generation of highly-realistic fake images/videos has become possible due to significant advances in the field of Gen AI. In addition to its several applications,

Gen AI can also have severe negative effects. The main challenge is to develop a general framework that can detect images/videos generated by different generative models. Since different models have different characteristics, it is difficult for a single model to capture this without *apriori* knowledge of the generative mechanism. Another challenge is that images can be partially edited, which makes it even harder to detect whether the data is real or fake.

One of our predoctoral fellows, Anuska Roy, studies the existing literature in detail to come up with a framework that tries to generalise across different generative models for detecting fake and real data.

(ii) Speech, image, and video analysis

One of the grand challenge problems that the Kotak IISc AI–ML Centre seeks to address concerns 'speech, image, and video analysis'. Voice authentication, speech recognition, and facial recognition can enable a new category of enhanced services and data sharing mechanisms. Users can use smart phones with cameras to authenticate themselves using voice, fingerprints, face recognition or other modalities. Once authenticated, they can then utilise various services or authorise sharing of data.

Example of KIAC project: *A novel multi-resolution feature perturbation (MRFP) technique*

A novel MRFP technique was proposed by one of our PhD students, Aniruddh Sanjoy Sikdar, and his team to randomise domain-specific fine-grained features and perturb the style of coarse features. The experimental results on various urban-scene segmentation datasets clearly indicate that, along with the perturbation of style-information, perturbation of fine-feature components is paramount to learn domain invariant robust feature maps for semantic segmentation models. MRFP is a simple and computationally efficient, transferable module with no additional learnable parameters or objective functions, that helps state-of-the-art deep neural networks to learn robust domain invariant features for simulation-to-real semantic segmentation.

This work was presented as a paper titled 'MRFP: Learning Generalizable Semantic Segmentation from Sim-2-Real with Multi-Resolution Feature Perturbation' at the Conference on Computer Vision and Pattern Recognition (CVPR) 2024 held from 17–21 June 2024 in Seattle, USA.

(iii) Emerging objectives in machine learning

The basic objective in machine learning has been to enable machines to learn from data and to utilise this learning for decision-making in new situations with minimum human intervention. While machine learning applications are on the rise, researchers have raised various concerns. Studies by some of our faculty members and their students address some of these concerns; the related papers were presented at leading conferences with support from KIAC.

Example of KIAC project: *evaluating machine learning models*

The rise in Internet usage has led to the generation of large amounts of data, resulting in the adoption of various supervised and semi-supervised machine learning algorithms that can effectively utilise the massive amount of data to train models. However, before deploying these models in the real world, these must be strictly evaluated on performance measures such as worst-case recall and satisfy constraints such as fairness. Current state-of-the-art empirical techniques offer sub-optimal performance on these practical, non-decomposable performance objectives, while theoretical techniques necessitate training a new model from scratch for each performance objective.

To optimise for the desired objective, one of our professors, Venkatesh Babu Radhakrishnan, and his student team proposed SelMix, a selective mixup-based inexpensive fine-tuning technique for pre-trained models. It was found that SelMix fine-tuning significantly improves the performance of various practical non-decomposable objectives across benchmarks. This study was presented as a paper titled 'Selective Mixup Fine-tuning for Optimizing Non-decomposable Metrics' at the International Conference on Learning Representations (ICLR) 2024 held from 7–11 May 2024 in Vienna, Austria.

(iv) AI in healthcare

The tools of artificial intelligence are being increasingly employed in 'healthcare' – from diagnosis, clinical decision support, electronic health record systems, patient engagement and adherence, and administration. Given the potential impact of AI in the healthcare sector, KIAC supports projects that encompass (a) cancer, (b) diabetes, and (c) air pollution-related health problems.

Example of KIAC project: *drug response prediction*

Recent cancer drug response prediction methods have largely followed the paradigm of unsupervised domain-invariant representation learning followed by a downstream drug response classification step. Introducing supervision in both stages is challenging

due to heterogeneous patient response to drugs and limited drug response data. One of our PhD students, Kumar Shubham, is working in this area and presented a paper titled 'WISER: Weak Supervision and Supervised Representation Learning to Improve Drug Response Prediction in Cancer' at the International Conference on Machine Learning (ICML) 2024 held from 21–27 July 2024 at Vienna, Austria. This paper addressed these challenges through a novel representation learning method in the first phase and weak supervision in the second. The experimental results on real patient data demonstrated the efficacy of the method (WISER) over state-of-the-art alternatives on predicting personalised drug response.

(v) AI in robotics

Robotics is a branch of science and engineering that involves the conception, design, construction, and use of robots, which can assist humans in multiple tasks. Typically, these tasks are repetitive or dangerous for humans to carry out. Artificial intelligence technologies can be integrated into robotic systems to enhance their capabilities and enable them to perform complex tasks. They enable robots to learn from experience, adapt to new situations, and make decisions based on data from sensors using AI techniques such as machine learning, computer vision, and natural language processing.

Example of KIAC project: *human–robot collaboration*

To enhance human–robot collaboration and enable robots to comprehend and respond to human instructions more effectively, one of our interns, Aiman Aatif Bayezeed, is working on a framework that seeks to synergise the power of large language models like GPT with formal languages to elevate the robots' cognitive abilities. This integration empowers robots with advanced linguistic and reasoning capabilities, facilitating seamless communication with humans across various domains.

Through this framework, robots can interpret complex commands, grasp context, and execute tasks with precision, leading to enhanced collaboration in diverse settings such as manufacturing, healthcare, and customer service. Additionally, by leveraging GPT's continuous learning capabilities, the framework can adapt to evolving linguistic nuances and domain-specific knowledge, ensuring sustained performance improvement over time.

(vi) AI in finance

In the finance sector, artificial intelligence technologies can be used to analyse data, automate tasks, and improve decision-making. They can help in personalised customer service, fraud detection, and investment predictions.

Example of KIAC project: *Optimal selection of government securities*

While the banking sector has extensively focused on maintaining sufficient capital through holdings of government securities, there is limited literature addressing the selection and optimisation of these securities under varying market conditions. In particular, the dynamic nature of interest rates, which can undergo regime shifts, presents a significant challenge for treasury management. One of our predoctoral fellows, Aniruddh Srinivasan, is working on a project that aims to address this gap by developing a framework for the optimal selection of government securities based on their duration, with a focus on maximising portfolio performance amidst potential regime changes in interest rates.

The proposed approach involves the use of Markov switching models to capture the probabilistic nature of regime shifts in interest rates. By identifying different regimes, such as high- and low-interest rate environments, the model will enable the construction of portfolios that are resilient to such shifts. The outcome will be an optimised portfolio of government securities that balances the trade-off between risk and return, aligning with the treasury's objective of capital preservation and growth. This research will provide actionable insights for treasuries to enhance their portfolio management strategies in the face of uncertain interest rate environments.

The research projects of the Centre's interns and predoctoral fellows are presented in the tables below.

THEME: GENERATIVE AI

SL. No.	Name	Research projects
1	Sunny Bhati	Latent space exploration for generative AI.
2	Achin Parashar	Synthesising naturalistic 'super-stimuli' using deep generative models.
3	Anuska Roy	Study of existing literature in detail and come up with a framework which tries to generalise across different generative models for detecting fake and real data.
4	Hariharasai Mohan	The aim is to predict the properties of the solvent-2D material interface, evaluating the bulk properties of solvents, experimenting with various solvents, etc. These tasks could be completed quickly using generative AI, thus reducing the need for extensive experimentation.
5	Harish Kumar S	Developing foundation models specifically tailored for probabilistic forecasting and anomaly detection tasks within the energy systems domain.
6	Harshith Reddy Kaila	Developing and demonstrating the performance of an information theory-enhanced machine learning model that generates high-resolution and realistic inflow turbulence while sustaining the temporal correlations characteristic of turbulence. For this, generative adversarial networks (GANs) will be used to generate instantaneous turbulent flow field from a Gaussian random latent vector.
7	Kinjawl Bhattacharyya	Exploring the robust representation spaces in diffusion models and leveraging the same to enhance out-of-distribution (OOD) robustness in critical applications.
8	Prateek Upadhya	Quality assessment of generated images.
9	Sai Teja Tangudu	Quality assessment of generative image models.
10	Tanishq Kekre	Heterogeneous 2.5D chiplet-based architecture for accelerating LLM inference.
11	Sanjana Kiran Tikare	Developing an efficient GAN model and assess its performance on a large-scale time series data from the energy domain.
12	Gyanig Kumar	Investigating and validating a human-in-the-loop approach where human input will be elicited during an image translation task using existing diffusion and GAN (generative adversarial network) models.

THEME: SPEECH, IMAGE, AND VIDEO ANALYSIS

SL. No.	Name	Research projects
1	Navaneeth Sivakumar	Developing actual scale NeRF (neural radiance fields) for three-dimensional reconstruction of objects using few-shot learning approaches.
2	Margamitra Bhattacharya	Streamlining air traffic control (ATC) communication using CPDLC (controller-pilot data link communications) system coupled with automated speech recognition.
3	Rikathi Pal	Using topological methods to better visualise and understand loss landscapes of deep neural networks.
4	Rikhil Gupta	Sparse input novel view synthesis using 3D Gaussian splatting.
5	Rishab Rajesh Sharma	Exploring vision-language models on the lines of CLIP (contrastive language-image pretraining) which was released a few years back.
6	Sushil Jangra	Integration of satellite remote sensing data and PINNs (physics informed neural networks) for deep ocean study.
7	Sachidanand V S	3D inpainting in diffusion using MPI based latent representation.
8	Ram Samarth B B	Developing a parameter-efficient, PDE-based spectral GNN diffusion method.
9	Akansha Sharma	Federated learning: pursuing an innovative approach where the server utilises a pre-trained vision-language model and a limited set of labelled data, while the clients work with only unlabelled data.
10	Ashish Ramayee Asokan	Exploring the utility of large-scale vision -language models such as CLIP from the perspective of black box distillation for improved domain generalisation.

THEME: AI IN ROBOTICS

SL. No.	Name	Research projects
1	Saksham Bhutani	Creating deep learning perception architectures for autonomous systems that are robust to faults and attacks.
2	Anurag Maurya	Imitation learning based spatio-temporal planning for bimanual robot.
3	Chiluveru Akhil	Designing neural control barrier functions, coupled with guarantees that can be achievable through various neural network properties like boundedness and Lipschitz continuity. These efforts aim to craft safety-critical control algorithms tailored for robotic systems with unknown dynamics.
4	Puneeth S	To achieve a holistic, interdisciplinary perspective of robotic bimanual manipulation, the power of machine learning algorithms will be leveraged and human interactive feedback will be incorporated into the learning process.
5	Tashmoy Ghosh	Robotics and AI: investigating hierarchical learning algorithms, exploring the integration and contrast with task and motion planning methods, incorporating imitation learning with human interactive feedback, and applying these approaches to interesting tasks such as (dis)assembly, coffee making, and cleaning.
6	Nitya Ahuja	Automated drone corridor development using AI/ML tools based on the communication infrastructure available in the drone space.
7	Aiman Aatif Bayezeed	Enhancing human-robot collaboration: a hybrid language framework integrating large language models and formal languages.
8	Rankit Kachroo	Developing new and better ways to deal with trajectory data making downstream tasks like clustering more efficient.
9	Rajat Surya	Multi-arm bandit algorithms for multi-robot collaboration.
10	Natish Murugesh	Investigating combining of movement primitives with additional information about the desired impedance profiles (impedance control primitives) for robotic manipulation that leverages learning from indirect observation (e.g., motion capture, Kinect, YouTube videos).

THEME: AI IN HEALTHCARE

SL. No.	Name	Research projects
1	Anuj Attri	Design and implementation of AI-driven solutions tailored to the needs of individuals with diabetes.

THEME: AI IN FINANCE

SL. No.	Name	Research projects
1	Pranoy Varma	Optimal entry and exit problem for a financial asset that follows a mean reverting process.
2	Aniruddh Srinivasan	Developing a framework for the optimal selection of government securities based on their duration, with a focus on maximising portfolio performance amidst potential regime changes in interest rates.

OTHER AI/ML PROJECTS

SL. No.	Name	Research projects
1	Ayman Un Nisa	A machine learning-assisted floor planning technique for 2.5D system (chiplet-based system).
2	Rishabh Sabharwal	Use of power iteration methods to efficiently obtain the top eigenvalues of the Hessian. Additionally, the aim is to adapt discrete-depth optimisers such as Adam and SGD for continuous-depth models represented by neural ODEs.
3	Shrikar Madhu	Class-incremental continual learning in neural networks.
4	Bhakshi Mehul	Computational fluid dynamics in the aerospace industry - generating optimal meshes: utilising neural networks to obtain adjoint solutions directly from the primal solutions, eliminating the need to solve additional adjoint equations and significantly reducing computational costs.
5	Rishabh Sabharwal	Adapting discrete-depth optimisers such as Adam and SGD for continuous-depth models represented by neural ODEs.
6	Jinal Vyas	Exploring the distributed systems, accuracy, performance and reliability aspects of federated learning across edge and cloud resources. It will develop aggregation and client selection techniques to accelerate federated learning.
7	Abhinava Raja	Data-driven driving behaviour analytics in Indian traffic conditions.
8	Venkatesh T	Developing AI/ML techniques to model and predict how the concentrations of important air pollutants change over time in major Indian cities using the Sentinel-5P Satellite data.

Objective 3: Developing innovative, deep tech solutions to meet current and emerging requirements

Call for deep tech solutions

The Centre has planned to release a call for deep tech solutions on 1 October 2024. This call targets IISc faculty and seeks proposals on AI for public good – solutions that can be leveraged by the industry ecosystem or make it to Government missions. The timeline for such solutions has been fixed as March 2025, and the goal for the next two years is to increase the funding for such initiatives. The sponsorship by KIAC for these solutions will be duly acknowledged.

Meanwhile, the Centre's research work and collaborations has led to the development of solutions to global problems and grand challenges. A few cases have been highlighted below.

AI Centre of Excellence in Health

- India has a third of the world's oral cancers. Oral cancer accounts for a third of all cancer deaths in India.
- Over 70% of the cases are diagnosed late due to reduced access to Primary Health Centre (PHC) specialists.
- Two workshops (focussed on AI in oral cancer) conducted in response to a call from the Ministry of Education for consortium proposals on AI for Health from higher education institutions.

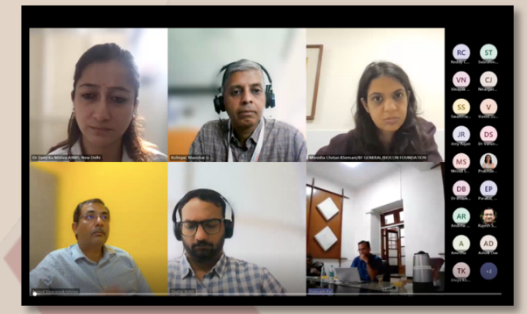

- The IISc-led consortium is one of the four selected for support during the proof-of-concept Phase 1.
- Outcomes of Phase 1:
    - a white light imaging-based, mobile-phone based, AI tool for oral cancer screening for use at scale at the PHCs.
    - An application has been submitted for Class B certification for the Software as a Medical Device (SaMD), to ensure regulatory standards.
    - The open-sourced model and app will serve as digital public goods, fostering grassroot innovation in oral cancer screening.

Farmer Registration and Unified beneficiary InformaTion Systems (FRUITS)

- Work with the e-governance cell of the Government of Karnataka (GoK) to help develop solutions for muzzle print-based authentication of bovine animals (roughly around one crore).
- Technical advice to the AI/ML team at the Centre for e-Governance, GoK.
    - Feature design for robust muzzle print based authentication/identification
    - Design of experimental set up to evaluate solutions

Quality assessment (QA) of AI-generated images

- Multiple e-commerce companies use generative AI to display catalogue content, but the image quality is often poor, leading to poor user experience in online shopping.
- KIAC supported projects on this topic are being adapted to diverse domains such as e-commerce (for use by Flipkart) and virtual try-on (for use by SpreeAI) respectively.

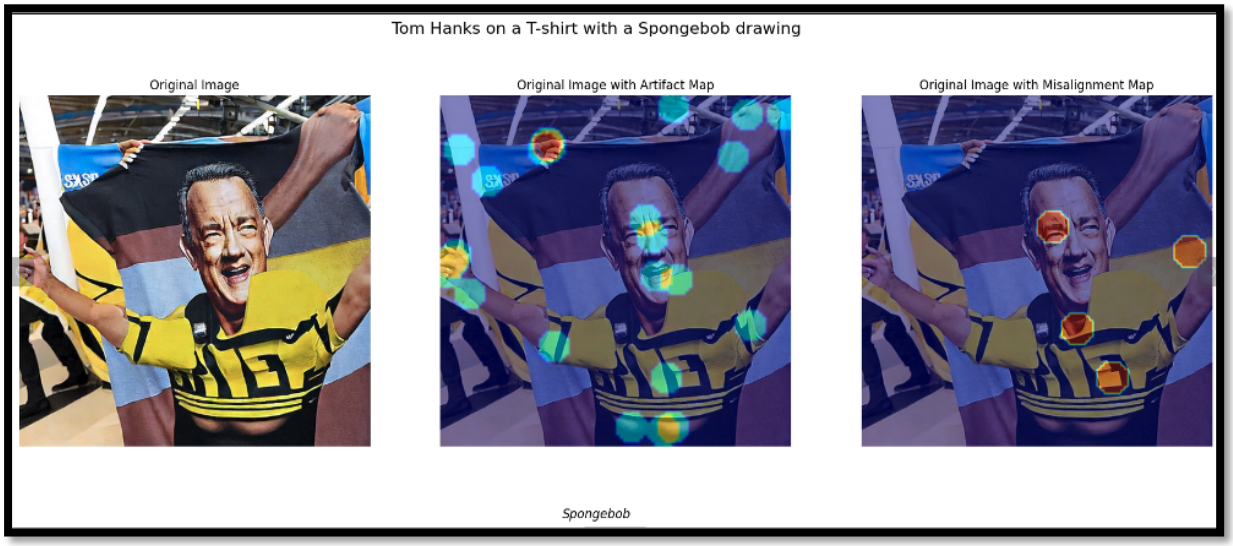

Picture credit: Liang et al., CVPR 2024

Activities and their impact

The beneficiaries and impact of the activities conducted by KIAC in this quarter (July–September 2024) is presented in the table below.

ACTIVITY		BENEFICIARIES	OUTPUT INDICATORS
Student support	Kotak scholarship	7	education, training human resources
	MTech scholarship	7	
	PhD scholarship	7	
Research support	Internship	13	training human resources, developing solutions for current requirements
	Predoctoral fellowship	30	
	Postdoctoral fellowship	2	
Courses and training sessions (by KIAC faculty)	For IISc students	274	education, teaching, skill development
	For IITM MTech students	42	
	NPTEL course	1990	
	For faculty of MSRIT and NMIT	100	
Support to participate in/conduct competitions	iGEM competition	13	skill development, visibility
	Ideathon 2024	80	
KIAC Distinguished Seminar Series	Talk by Mitesh M Khapra	114	education, collaboration, outreach
Kotak IISc AI–ML Talk Series	Talk by Kumaresh Krishnan	21	education, collaboration, outreach
	Talk by Amit K Roy-Chowdhury	38	
Workshops	CyPhySS 2024	58	skill development, education, capacity building
	SPCOM 2024	115	
Support to present accepted papers at A*conferences	ECCV 2024	2	skill development, visibility, enabling publications in prestigious journals and participation in top-level conferences
	ACL 2024	1	
	SIGGRAPH 2024	1	
	ICML 2024	4	
			TOTAL BENEFICIARIES = 2919

References

1. Report on the Asian Engineering Deans' Summit 2024 https://aeds2024.iisc.ac.in/files/Asian_Engineering_Deans_Summit_2024.pdf
2. Sreeradha Basu and Saumya Bhattacharya, India 14th in AI research with just 1.4% share of papers: study, 16 August 2024, https://economictimes.indiatimes.com/tech/artificial-intelligence/india-14th-in-ai-research-with-just-1-4-share-of-papers-study/articleshow/112554830.cms?from=mdr