
Computer Science and Engineering
Indian Institute of Technology Kanpur
06
APR
SUNDAY
INNOVATIONS IN CSE
RESEARCH COLLOQUIUM 2025

Invited Lecture

Innovations in CSE

Research Colloquium 2025

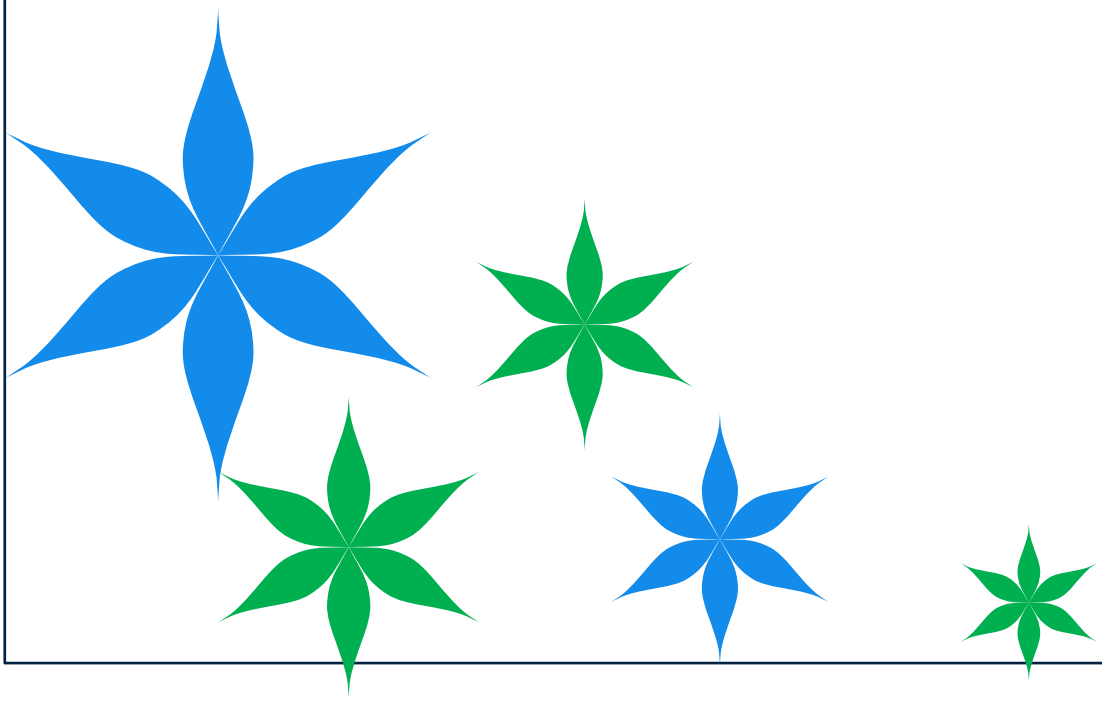

10:00 AM, 06 APRIL 2025,
RM101, Rajeev Motwani Building

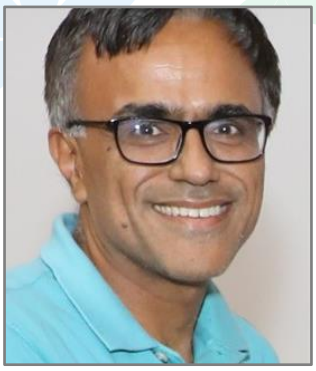

AI/ML-based 6G Wireless Network Design

Prof Rohit Budhiraja

Professor and *Visvesvaraya Fellow*,
Dept of EE, IIT Kanpur

Abstract: 5G networks have recently been deployed all over the world. We have also designed a fully indigenous 5G cellular network in a project funded by the Department of Telecommunications (DoT), Govt of India. The entire hardware and the software for all network layers of this 5G network were developed inhouse. This 5G technology is now deployed at multiple places including at IIT Kanpur and was recently transferred to two companies - Tata-Tejas networks and CDoT. We are now working on enhancing this 5G network towards 6G by incorporating AI/ML in its design from scratch. In this seminar, we will first briefly discuss the architecture of our indigenous 5G network. We will then show how AI/ML is used to address the problems which are specific to the design of 6G wireless networks. We will then then discuss how we are contributing to the overall design of first AI/ML-native 6G network, which will tentatively be deployed in 2030.

Speaker Bio: Rohit Budhiraja is leading the effort to design India's first indigenous 5G network, a technology that was recently transferred to the industry, and 5G+/6G standards. He works closely with the industry, government organizations and R&D labs, is a member of the governing council of TSDSI, India's standard develop organization, and vice chairman of Bharat 6G alliance for 6G use cases and revenue streams. He has 25 years experience building 2G, 3G, 4G and 5G wireless systems. His current research focuses on 5G+/6G networks design using machine learning and optimization-based methods and investigating their performance by prototyping them. He has received the CNR Rao award for architecting and designing the indigenous 5G network, the IIT Kanpur Excellence in Teaching Award, Visvesvaraya Fellowship and Early Career Research Award from the Government of India.

Research Presentations

Cybersecurity and Cyber-physical Systems

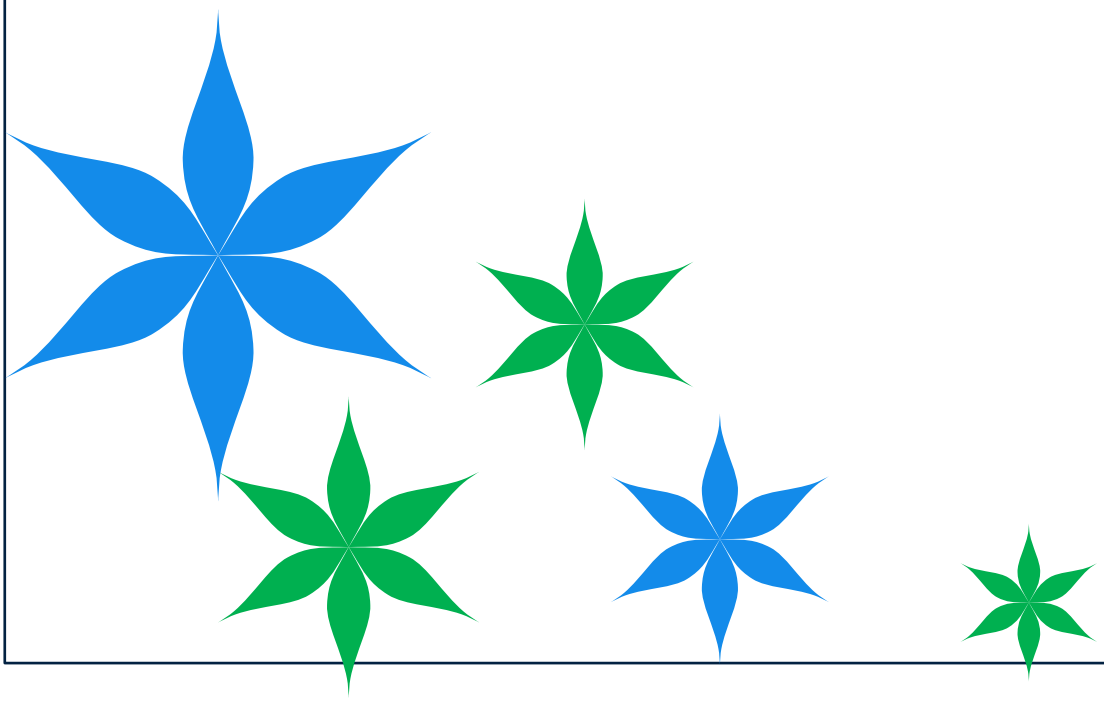

Research Presentations

Dipesh

Title: *Door Knock*: Reverse Engineering the MPSoC Layout through Timing Attack on NoC

TL;DR: The motivation of the Door Knock attack is to find the location of each application on the MPSoC by reverse engineering the MPSoC layout. This attack will relax the adversarial model of the exiting attacks by determining the location of the attacker and the victim on an MPSoC.

Abstract: Multi-Processor Systems-on-Chip (MPSoC) have emerged as highly versatile and efficient platforms suitable for a wide range of applications like multimedia applications and telecommunication architectures. One of the key components in MPSoC is the Network-on-Chip (NoC), which facilitates the interconnection of various processing elements, enabling efficient data communication. Several timing attacks, such as Earthquake attack, P+P Firecracker, and P+P Arrow have been proposed on NoC that exploit the variations in execution times of operations to infer cryptographic keys. In this letter, we propose to leverage the timing attack on NoC to reverse engineer the mapping of each processing element onto the MPSoC architecture. To the best of our knowledge, it is the first work that relies on creating the contention between the requests that are sent to different PEs and reveal the layout by just analysing the reply latency. In the experimental setup, we are able to map PEs for MPSoC consists of Mesh, Torus, Point-to-Point, Ring, and Flattened Butterfly NoC topology with 100% accuracy that can be extremely useful for the attackers in the reconnaissance phase. Further, as the existing mitigation techniques to counter timing attacks are based on the assumption that the contention is created between the packets of secure-insecure domains, they will not be able to mitigate the proposed reverse engineering attack.

Co-authors: Dipesh, Urbi Chatterjee

Email: dipesh@cse.iitk.ac.in

Research Presentations

Ashutosh Deshwal

Title: Through-the-Wall Multi-Person Localization using Translation and Rotation Synthetic Aperture Radar

TL;DR: We propose using synthetic aperture radar with translation and rotation to improve multi-person localization, especially in close proximity and through walls. Our method shows up to 2.19× better localization in experimental scenarios.

Abstract: An emerging application of wireless sensing is locating and tracking humans in their living environments, a primitive that can be leveraged in both daily life applications and emergency situations. However, most proposed methods have limited spatial resolution when multiple humans are in close vicinity. The problem becomes exacerbated when there is no line- of-sight path to the humans. In this paper, we consider multi- person localization of humans in close vicinity of each other. We propose the use of synthetic aperture radar that combines both translation and rotation to increase effective aperture size, leveraging small rhythmic changes in the radar range due to human breathing. We experimentally evaluate the proposed algorithm in both line-of-sight and through-wall cases with three to five humans in the scene. Our experimental results show that: (i) larger synthetic apertures due to radar translation improve multi-person localization, e.g., by 1.42× when the aperture size is increased by a factor of 2×, and (ii) rotation can largely compensate for gains provided by translation, e.g., rotating the radar over 360◦ without changing the aperture size results in 1.22× gains over no rotation. Overall, maximal gains of 2.19× are achieved by rotating and translating over a 2× larger aperture.

Co-authors: Shubham Sinha, Ashutosh Deshwal, Alireza Azizi, Divyanshu Pandey, Nishant Mehrotra, Amitangshu Pal, Ashutosh Sabharwal

Email: ashutosh@cse.iitk.ac.in

Research Presentations

Hrushikesh Chunduri

Title: Trusted Yet Dangerous: The Subversive Role of LOLBins in Cyber Threats

TL;DR: We analyze the use of LOLBins in 5 cyber attacks—Ransomware, Cryptominers, APTs, Info Stealers, and RATs—revealing that 51% rely on LOLBins to evade detection, download payloads, and ensure stealth. Ransomware shows the highest diversity, while APTs exhibit 68% LOLBin usage.

Abstract: The proliferation of advanced detection techniques and the evolution of next-generation firewalls and antivirus engines have led to the increasing sophistication of cyber threats. In this context, malware authors are crafting disguised payloads that mimic benign behavior. To achieve this, the use of Windows signed executables/binaries (Living off the Land Binaries or LOLBins) and libraries has become increasingly relevant as a method to evade antivirus and signature-based detection techniques. These binaries inherently grant attackers a level of trust within the Windows operating system as they are signed by Microsoft. Understanding the specific LOLBins used by different attacks and their variants is crucial for developing effective detection rules and enhanced threat intelligence. Therefore, in this work, we analyze the presence of LOLBins from five distinct cyber attacks through dynamic analysis to determine the ubiquity and role of these Windows signed binaries in these attacks. We observe that the usage of LOLBins is nearly 51% of the payloads across Ransomware, Cryptominers, Advanced Persistent Threats (APTs), Information Stealers, and Remote Access Trojans (RATs)/Trojans. We also identify the distinct roles of the same LOLBins in attack variants in terms of evading defense strategies, downloading payloads, and offering stealth. Notably, ransomware and crypto miner payloads exhibit a higher diversity of utilizing 55 and 30 distinct LOLBins, respectively. Finally, we systematically analyze and compare the usage of LOLBins in Cobalt Strike payloads—a legacy multipurpose tool used by many malware families to evade detection, gather information, and persist within the victim environment. We identify Cobalt Strike to have the highest usage among all categories, at almost 73%.

Co-authors: Hrushikesh Chunduri, P Mohan Anand, P.V Sai Charan and Sandeep Kumar Shukla

Email: hrushikesh22@iitk.ac.in

Research Presentations

NANDA RANI

Title: Tools and Techniques for Advancing Intelligence Extraction and Automated Attribution of APTs

TL;DR: Advanced Persistent Threats (APTs) are stealthy, targeted cyber attacks that require scalable attribution. We introduce TTPXHunter and MalXCap for threat intelligence extraction, and ATTRACT (TTP-based) along with a malware-based framework for accurate APT group attribution.

Abstract: Advanced Persistent Threats (APTs) challenge cybersecurity with their stealthy, persistent, and targeted nature. Manual attribution methods—mapping attack artifacts to frameworks like MITRE ATT&CK—are time-consuming, subjective, and lack scalability as adversarial tactics evolve. Moreover, current tools often fail to integrate diverse data sources such as narrative reports, malware samples, and operational traces, creating gaps in threat intelligence extraction and behavioral attribution. This highlights the need for automated solutions that can extract actionable intelligence and attribute APTs effectively. We present methods and frameworks leveraging natural language processing, machine learning, and behavioral analysis for threat intelligence extraction and APT attribution. For intelligence, we introduce TTPXHunter, a transformer-based model that extracts MITRE-ATT&CK TTPs from unstructured reports, with enhanced support for limited techniques via augmentation and domain-specific pretraining. It supports STIX-formatted output to ensure seamless CTI integration. Complementing this, MalXCap processes over 8,000 malware samples to extract 12 unique malignant capabilities, enriching contextual threat understanding. For attribution, we propose ATTRACT, a TTP-sequence-based framework that models attacker behavior using the Unified Kill Chain and applies a novel similarity metric to match observed TTP patterns. We also introduce a malware-based attribution method that uses TTP and linker timestamp extraction from binaries, supporting attribution across diverse threat groups, including open-world cases. These contributions support scalable, behavior-driven APT attribution by linking diverse threat data sources. This research lays the foundation for next-generation automated threat intelligence systems, offering faster response, improved attribution accuracy, and greater resilience against advanced cyber threats.

Co-authors:

1. TTPXHunter: Nanda Rani, Bikash Saha, Vikas Maurya, and Sandeep Kumar Shukla
2. MalXCap: Bikash Saha, Nanda Rani, and Sandeep Kumar Shukla
3. Chasing the Shadows: Nanda Rani, Bikash Saha, Vikas Maurya, and Sandeep Kumar Shukla
4. Genesis of Cyber Threats: Nanda Rani, Bikash Saha, Ravi Kumar, and Sandeep Kumar Shukla

Email: nandarani@cse.iitk.ac.in

Research Presentations

BIKASH SAHA

Title: Harnessing Large Language Models for Next-Gen Cybersecurity

TL;DR: We present two LLM-powered systems for cybersecurity automation: MaLAware for malware behavior summarization and PARAG for policy-aware intelligence retrieval. Together, they enhance accuracy, efficiency, and interpretability in modern cyber defense.

Abstract: The growing sophistication of cyber threats, complex security policies, and the overwhelming scale of cybersecurity data demand automation in threat intelligence processing and decision support. Traditional approaches relying on manual analysis of security reports, compliance documents, and sandbox outputs are time-consuming and prone to errors. While AI and ML have advanced cybersecurity automation, they often lack contextual reasoning and adaptability. Large Language Models (LLMs) offer new potential, with the ability to process unstructured threat data and assist in scalable security operations. However, their application poses challenges such as hallucination, factual inconsistency, and the need for domain-specific tuning. This research explores the use of LLMs in two distinct cybersecurity applications: MaLAware and PARAG. MaLAware is an LLM-powered system for malware behavior analysis and summarization. It processes sandbox reports to extract behavioral traits and Indicators of Compromise (IoCs), producing interpretable threat summaries. By leveraging LLM reasoning, MaLAware reduces analysis time and expertise requirements while improving triaging accuracy. It supports multiple LLM backends and uses quantization for efficient deployment. PARAG applies Retrieval-Augmented Generation (RAG) to cybersecurity policy interpretation. It transforms regulatory texts, policies, and advisories into structured formats for accurate, context-aware querying. By grounding LLM outputs in policy documents, PARAG enhances factual reliability and assists in automated compliance checks. Evaluated on an organization-agnostic dataset, it improves accessibility to cybersecurity knowledge. Together, these projects demonstrate the practical viability of LLMs in cybersecurity workflows, setting the stage for more intelligent, efficient, and interpretable solutions. Future directions include enhancing adversarial robustness, and integrating explainability for analyst trust.

Co-authors:

1. MaLAware: Bikash Saha, Nanda Rani, and Sandeep Kumar Shukla
2. PARAG: Bikash Saha, Nanda Rani, Joheen Chakraborty, Divyanshu Singh, Soumyo V. Chakraborty and Sandeep Kumar Shukla

Email: bikash@cse.iitk.ac.in

Research Presentations

Supriya Adhikary

Title: RLWE-based IPFE Library and its Application to Privacy-preserving Biometric Authentication

TL;DR: We optimized the Mera et. al's work and propose a fast inner product functional encryption library. As an additional contribution to this work, we design a privacy-preserving biometric authentication scheme using inner product functional encryption primitives.

Abstract: Mera et al. first proposed an inner product functional encryption scheme based on ring learning with errors to improve efficiency. In this work, we optimize the implementation of their work and propose a fast inner product functional encryption library. Specifically, we identify the main performance bottleneck, which is the number theoretic transformation based polynomial multiplication used in the scheme. We also identify the micro and macro level parallel components of the scheme and propose novel techniques to improve the efficiency using open multi-processing and advanced vector extensions 2 vector processor. Compared to the original implementation, our optimization methods translate to 89.72%, 83.06%, 59.30%, and 53.80% improvements in the Setup, Encrypt, KeyGen, and Decrypt operations respectively, in the scheme for standard security level. Designing privacy-preserving applications using functional encryption is ongoing research. Therefore, as an additional contribution to this work, we design a privacy-preserving biometric authentication scheme using inner product functional encryption primitives.

Co-authors: Supriya Adhikary, Angshuman Karmakar

Email: adhikarys@cse.iitk.ac.in

Research Presentations

Shaijal Tripathi

Title: LRHAR: A Lightweight Rule-based Framework for Human Activity Recognition at the Edge

TL;DR: LRHAR is a lightweight, rule-based framework for activity recognition on edge devices. It integrates object detection, tracking, and pose estimation, using logic-driven rules to analyze spatial-temporal relationships for action detection.

Abstract: Activity recognition in videos plays a pivotal role in many surveillance applications, including security, patient monitoring, autonomous driving, etc. Current state-of-the-art approaches often rely on computationally expensive deep learning models, limiting their scalability to low-cost edge devices. In this work, we propose a lightweight framework named LRHAR that leverages logic-driven or rule-based techniques for activity detection. The proposed framework integrates object detection, tracking, and pose estimation to comprehensively understand the surveillance scene. These tasks jointly retrieve minimal yet sufficient information from the cameras. This information is then interpreted by a rule-based system to identify various activities based on the spatial and temporal relationships between the detected poses and objects. Our framework is extremely lightweight, making it suitable for resource-constrained edge devices. Extensive experiments on NTU RGB+D 120 and Wildtrack datasets demonstrate that LRHAR can identify various human actions with a reasonable level of accuracy while still maintaining a decent frame rate above ~8 fps in state-of-the-art edge devices.

Co-authors: Mandar Dhake, Shaijal Tripathi, Shashwati Banerjea, Rakesh Yamjala and Amitangshu Pal

Email: tripjal@cse.iitk.ac.in

Research Presentations

PUJA MONDAL

Title: ZKFault: Fault attack analysis on zero-knowledge based post-quantum digital signature schemes

TL;DR: We analyze the LESS and CROSS signature schemes and devise our attack to recover the secret key using as little as a single fault, which is built on very simple fault assumptions. Also, we proposed various countermeasures to prevent these kinds of attacks.

Abstract: Computationally hard problems based on coding theory, such as the syndrome decoding problem, have been used for constructing secure cryptographic schemes for a long time. Schemes based on these problems are also assumed to be secure against quantum computers. However, these schemes are often considered impractical for real-world deployment due to large key sizes and inefficient computation time. In the recent call for standardization of additional post-quantum digital signatures by the National Institute of Standards and Technology, several code-based candidates have been proposed, including LESS, CROSS, and MEDS. These schemes are designed on the relatively new zero-knowledge framework. Although several works analyze the hardness of these schemes, there is hardly any work that examines the security of these schemes in the presence of physical attacks. In this work, we analyze these signature schemes from the perspective of fault attacks. All these schemes use a similar tree-based construction to compress the signature size. We attack this component of these schemes. Therefore, our attack is applicable to all of these schemes. In this work, we first analyze the LESS signature scheme and devise our attack. Furthermore, we showed how this attack can be extended to the CROSS signature scheme. Our attacks are built on very simple fault assumptions. Our results show that we can recover the entire secret key of LESS and CROSS using as little as a single fault. Finally, we propose various countermeasures to prevent these kinds of attacks and discuss their efficiency and shortcomings.

Co-authors: Puja Mondal, Supriya Adhikary, Suparna Kundu, and Angshuman Karmakar

Email: pujamondal@cse.iitk.ac.in

Research Presentations

Suraj Mandal

Title: Design of a Lightweight Fast Fourier Transformation for FALCON using Hardware-Software Co-Design

TL;DR: In this work, we have designed an efficient hardware-software co-design of FFT for FALCON using Winograd's FFT method. Our proposed framework outperforms the traditional Cooley-Tukey method and is flexible to adopt different instruction sets.

Abstract: Lattice-based post-quantum cryptographic algorithm FALCON needs to execute the time-critical Fast-Fourier Transformation (FFT). Existing works in the literature have explored hardware for FFT of FALCON using Cooley-Tukey. In this work, we have designed an efficient hardware-software co-design of FFT for FALCON using Winograd's FFT method. Winograd's FFT is a widely adopted technique for FFT and reduces the multiplication counts for higher radix FFT than the Cooley-Tukey, with a penalty of some extra addition/subtraction. Our Winograd radix-8 framework for FFT outperforms the traditional Cooley-Tukey method. Moreover, our proposed architecture is flexible in adopting different instruction sets and can also be configured for any type of FFT method with specific instruction sets.

Co-authors: Suraj Mandal, Debapriya Basu Roy

Email: surajmandal@cse.iitk.ac.in

Research Presentations

Vishesh Mishra

Title: Secure and Reliable Approximate Computing for Error-resilient Applications

Co-authors: Vishesh Mishra, Neelofar Hassan, Sparsh Mittal, Rekha Singhal, and Urbi Chatterjee
Email: vishesh.mishra@iiitg.ac.in

Anjali Manoj

Title: A Machine Learning Approach for efficient Side-Channel Analysis

Co-authors: Anjali Manoj, Debapriya Basu Roy, Priyanka Bagade
Email: anjalim@cse.iitk.ac.in

Research Presentations

Theoretical Computer Science

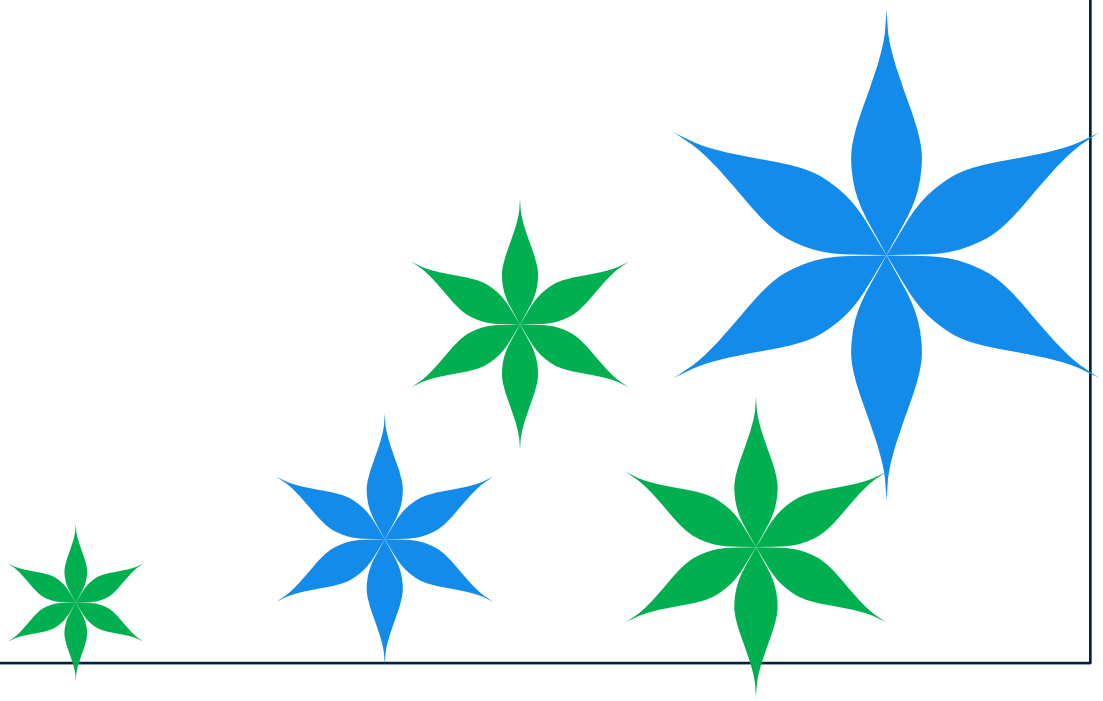

Suronjona Sarma

Title: One-Way Functions and Polynomial Time Dimension
TL;DR: We solved an open problem on the equivalence of randomness notions $\mathrm{K}_{\mathrm{poly}}$ and $\mathrm{cdim}_{\mathrm{P}}$. Polynomial time dimension ($\mathrm{cdim}_{\mathrm{P}}$) uses s-gales for information density, while $\mathrm{K}_{\mathrm{poly}}$ measures compressibility. Assuming one-way functions exist, we proved $\mathrm{cdim}_{\mathrm{P}} > \mathrm{K}_{\mathrm{poly}}$.

Abstract: This paper demonstrates a duality between the non-robustness of polynomial time dimension and the existence of one-way functions. Polynomial-time dimension (denoted $\mathrm{cdim}_{\mathrm{P}}$) quantifies the density of information of infinite sequences using polynomial time betting algorithms called s-gales. An alternate quantification of the notion of polynomial time density of information is using polynomial-time Kolmogorov complexity rate (denoted $\mathrm{K}_{\mathrm{poly}}$). Hitchcock and Vinodchandran (CCC 2004) showed that $\mathrm{cdim}_{\mathrm{P}}$ is always greater than or equal to $\mathrm{K}_{\mathrm{poly}}$. We first show that if one-way functions exist, then there exists a polynomial-time samplable distribution with respect to which $\mathrm{cdim}_{\mathrm{P}}$ and $\mathrm{K}_{\mathrm{poly}}$ are separated by a uniform gap with probability 1. Conversely, we show that if there exists such a polynomial-time samplable distribution, then (infinitely-often) one-way functions exist. Using our main results, we solve a long-standing open problem posed by Hitchcock and Vinodchandran (CCC 2004) and Stull under the assumption that one-way functions exist. We demonstrate that if one-way functions exist, then there are individual sequences X whose poly-time dimension strictly exceeds $\mathrm{K}_{\mathrm{poly}}(X)$, that is $\mathrm{cdim}_{\mathrm{P}}(X) > \mathrm{K}_{\mathrm{poly}}(X)$. Further, we show that the gap between these quantities can be made as large as possible (i.e., close to 1).
Co-authors: Satyadev Nandakumar, Subin Pulari, Akhil S., Suronjona Sarma
Email: suronjona@cse.iitk.ac.in

Tufan Singha Mahapatra

Title: Polynomial Factorization modulo prime powers

Co-authors: Tufan Singha Mahapatra and Nitin Saxena
Email: tufansm@cse.iitk.ac.in

Research Presentations

Computer Systems

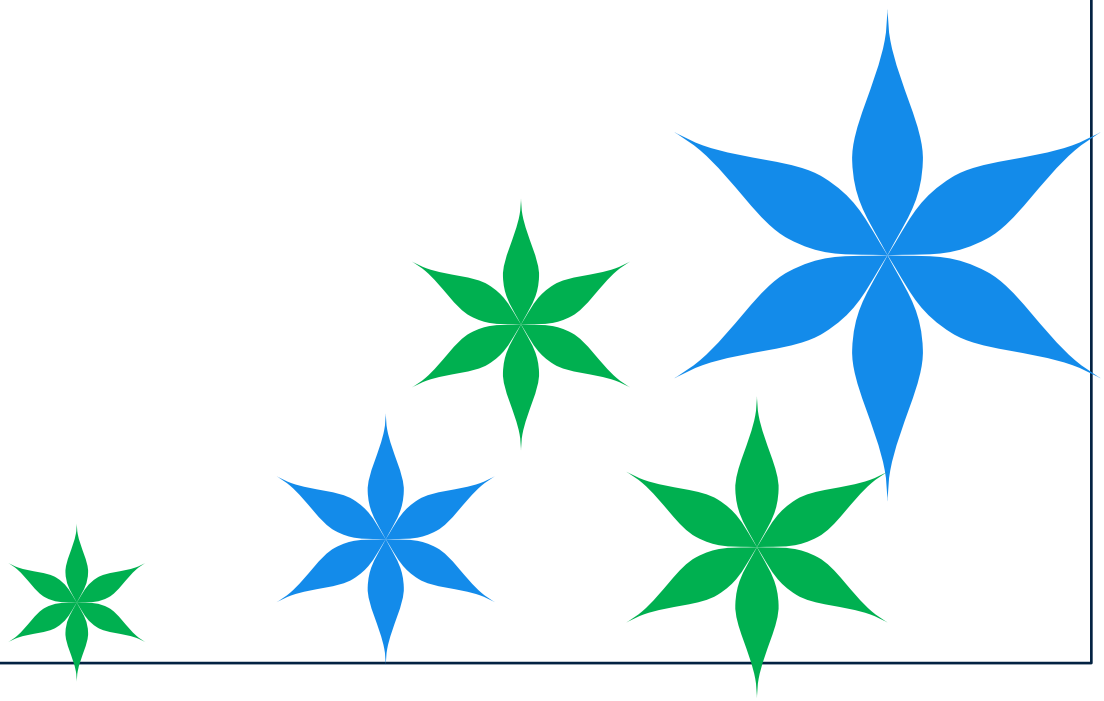

Research Presentations

Suvam Basak

Title: Measuring Orbital Shifts Due to Solar Radiation
TL;DR: Our tool, CosmicDance, tracks orbital shifts in LEO satellites due to solar events. It shows Starlink satellites face short- and long-term decay, even after mild solar events, sometimes crossing neighbouring satellite shells—raising risks of service gaps and orbital instability.

Abstract: Radiation shock waves from solar activities are known to be a menace to spaceborne electronic infrastructure. Recent deployments, like the SpaceX Starlink broadband mega-constellation, open up the possibility of measuring such impact on Low Earth Orbit infrastructure at scale. Our tool, CosmicDance, enables a data-driven understanding of satellite orbital shifts due to solar radiations. CosmicDance could also signal corner cases, like premature orbital decay, that could lead to service holes in such globally spanning connectivity infrastructure. Our measurements with CosmicDance show that Starlink satellites experience both short and long-term orbital decay even after mild and moderate-intensity solar events, often trespassing neighbouring shells of satellites.
Co-authors: Suvam Basak, Amitangshu Pal, Debopam Bhattacherjee
Email: suvambasak@cse.iitk.ac.in

Srinjoy Sarkar

Title: Heterogeneous CPU-GPU Data Structures for Large Workloads
TL;DR: Oversubscribing GPU memory through Unified Virtual Memory leads to poor performance due to far faults and is not suitable for data structure overflowing the GPU's memory. Our work proposes a hashtable design for large-scale applications that overflow GPU memory.

Co-authors: Srinjoy Sarkar, Vipin Patel, Swarnendu Biswas, and Mainak Chaudhuri
Email: srinjoy@cse.iitk.ac.in

Research Presentations

Vipin Patel

Title: Leveraging Cache Coherence to Detect and Repair False Sharing On-the-fly

TL;DR: Performance bugs due to false sharing do not manifest as observable correctness errors, and hence are challenging to detect and repair. Our work proposes an efficient extension of the MESI coherence protocol to eliminate harmful instances of false sharing on-the-fly.

Abstract: Performance bugs due to false sharing do not manifest as observable correctness errors, and hence are challenging to detect and repair. Prior approaches aim to both detect and repair false sharing instances automatically but most of them suffer from one or more of the following drawbacks: (i) high performance overhead due to expensive tracking of shadow memory, (ii) reliance on imprecise hardware events, and (iii) limited applicability and portability. We present extensions to the MESI cache coherence protocol for efficiently identifying and mitigating false sharing instances. The FSDetect protocol tracks the frequency of coherence misses per cache block to identify harmful instances of falsely shared lines while incurring negligible performance overhead. The FSLite protocol extends FSDetect to transparently privatize the falsely shared lines on accesses after detection, thereby eliminating the performance problem arising from false sharing. FSLite maintains coherence by performing precise byte-level updates of privatized blocks at the LLC on termination of privatization. Our simulation results on a variety of multithreaded workloads show that FSDetect can precisely identify all known harmful instances of false sharing. FSLite, on average, improves the performance of applications suffering from false sharing by 1.39X over the unmodified baseline, at the cost of a minimal increase in the chip area. Furthermore, applications running with FSLite stress the network less and show improved energy behavior.

Co-authors: Vipin Patel, Swarnendu Biswas, and Mainak Chaudhuri

Email: vipinpat@cse.iitk.ac.in

Research Presentations

Data Science, AI and Machine Learning

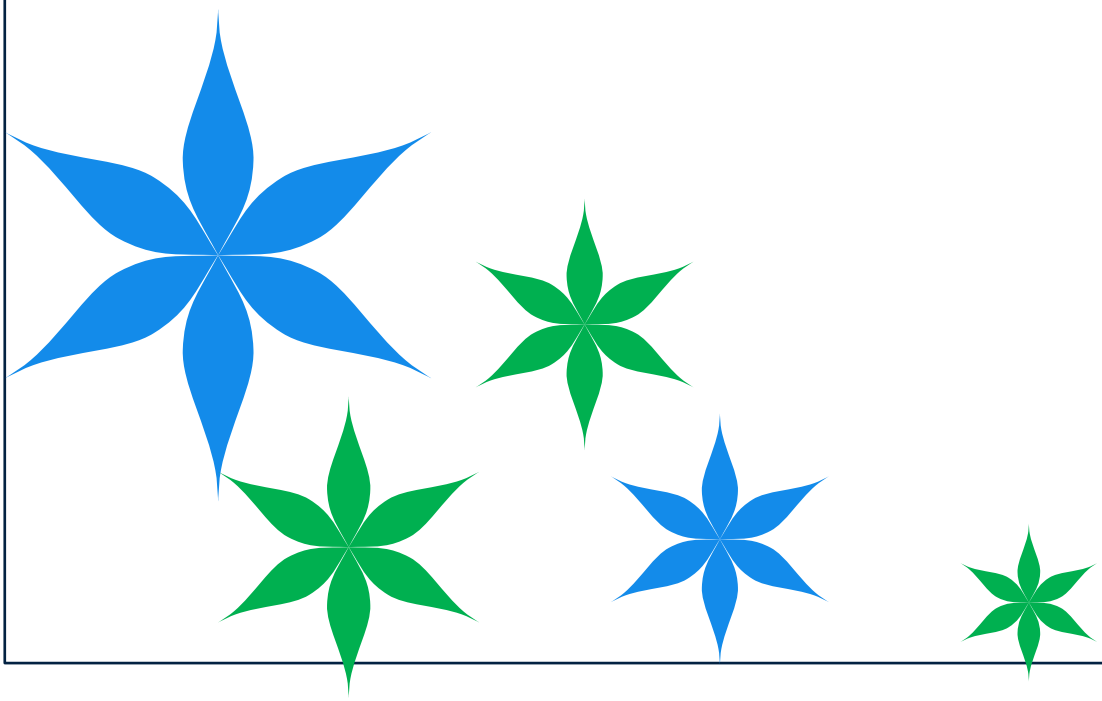

Research Presentations

Utsav Singh

Title: PEAR: Primitive Enabled Adaptive Relabeling for boosting Hierarchical Reinforcement Learning

TL;DR: We effectively leverage expert demonstrations using our adaptive relabeling based approach to deal with non-stationarity in the context of hierarchical reinforcement learning.

Abstract: Hierarchical reinforcement learning (HRL) has the potential to solve complex long horizon tasks using temporal abstraction and increased exploration. However, hierarchical agents are difficult to train due to inherent non-stationarity. We present primitive enabled adaptive relabeling (PEAR), a two-phase approach where we first perform adaptive relabeling on a few expert demonstrations to generate efficient subgoal supervision, and then jointly optimize HRL agents by employing reinforcement learning (RL) and imitation learning (IL). We perform theoretical analysis to bound the sub-optimality of our approach and derive a joint optimization framework using RL and IL. Since PEAR utilizes only a few expert demonstrations and considers minimal limiting assumptions on the task structure, it can be easily integrated with typical off-policy RL algorithms to produce a practical HRL approach. We perform extensive experiments on challenging environments and show that PEAR is able to outperform various hierarchical and non-hierarchical baselines and achieve upto 80% success rates in complex sparse robotic control tasks where other baselines typically fail to show significant progress. We also perform ablations to thoroughly analyze the importance of our various design choices. Finally, we perform real world robotic experiments on complex tasks and demonstrate that PEAR consistently outperforms the baselines.

Co-authors: Utsav Singh, Vinay P Namboodiri

Email: utsavz@cse.iitk.ac.in

Research Presentations

Divyaksh Shukla

Title: Towards Robust Evaluation of Unlearning in LLMs via Data Transformations

TL;DR: This work examines the robustness of existing MUL techniques for their ability to enable leakage-proof forgetting in LLMs, and the effect of data transformation on forgetting, i.e., an unlearned LLMs ability to recall forgotten information when changing the input format.

Abstract: Large Language Models (LLMs) have shown to be a great success in a wide range of applications ranging from regular NLP-based use cases to AI agents. LLMs have been trained on a vast corpus of texts from various sources; despite the best efforts during the data pre-processing stage while training the LLMs, they may pick some undesirable information such as personally identifiable information (PII). Consequently, in recent times research in the area of Machine Unlearning (MUL) has become active, the main idea is to force LLMs to forget (unlearn) certain information (e.g., PII) without suffering from performance loss on regular tasks. In this work, we examine the robustness of the existing MUL techniques for their ability to enable leakage-proof forgetting in LLMs. In particular, we examine the effect of data transformation on forgetting, i.e., is an unlearned LLM able to recall forgotten information if there is a change in the format of the input? Our findings on the TOFU dataset highlight the necessity of using diverse data formats to quantify unlearning in LLMs more reliably.

Co-authors: Abhinav Joshi, Shaswati Saha, Divyaksh Shukla, Sriram Vema, Harsh Jhamtani, Manas Gaur, Ashutosh Modi

Email: divyaksh@cse.iitk.ac.in

Research Presentations

Ayush Pande

Title: ViSt3D: Video Stylization with 3D CNN
TL;DR: This is work is based on Video Stylization which uses 3D CNN architectures in its backbone.

Abstract: Visual stylization has been a very popular research area in recent times. While image stylization has seen a rapid advancement in the recent past, video stylization, while being more challenging, is relatively less explored. The immediate method of stylizing videos by stylizing each frame independently has been tried with some success. To the best of our knowledge, we present the first approach to video stylization using 3D CNN directly, building upon insights from 2D image stylization. Stylizing video is highly challenging, as the video motion, which includes both camera and subject motions, and appearance are inherently entangled in the representations learnt by a 3D CNN. Hence, a naive extension of 2D CNN stylization methods to 3D CNN does not work. To perform stylization with 3D CNN, we propose to explicitly disentangle motion and appearance, stylize the appearance part, and then add back the motion to decode the final stylized video. In addition, we propose a dataset, curated from existing datasets, to train video stylization networks. We also provide an independently collected test set to study generalization of video stylization methods. We provide results on this test dataset comparing the proposed method with 2D stylization methods applied frame by frame. We show successful stylization with 3D CNN for the first time, and obtain better stylization in terms of texture cf. the existing 2D frame by frame methods.
Co-authors: Ayush Pande, Gaurav Sharma
Email: ayushp@cse.iitk.ac.in

Research Presentations

Abhinav Joshi

Title: Towards Quantifying Commonsense Reasoning with Mechanistic Insights

TL;DR: Study to quantify commonsense knowledge acquired in LLMs by performing a rigorous evaluation of real-world activities well understood by humans; and provide ways in which the decision-making about reasoning happening inside these models could be localized and understood.

Abstract: Commonsense reasoning deals with the implicit knowledge that is well understood by humans and typically acquired via interactions with the world. In recent times, commonsense reasoning and understanding of various LLMs have been evaluated using text-based tasks. In this work, we argue that a proxy of this understanding can be maintained as a graphical structure that can further help to perform a rigorous evaluation of commonsense reasoning abilities about various real-world activities. We create an annotation scheme for capturing this implicit knowledge in the form of a graphical structure for 37 daily human activities. We find that the created resource can be used to frame an enormous number of commonsense queries ($\sim 10^{17}$), facilitating rigorous evaluation of commonsense reasoning in LLMs. Moreover, recently, the remarkable performance of LLMs has raised questions about whether these models are truly capable of reasoning in the wild and, in general, how reasoning occurs inside these models. In this resource paper, we bridge this gap by proposing design mechanisms that facilitate research in a similar direction. Our findings suggest that the reasoning components are localized in LLMs that play a prominent role in decision-making when prompted with a commonsense query.

Co-authors: Abhinav Joshi, Areeb Ahmad, Divyaksh Shukla and Ashutosh Modi

Email: ajoshi@cse.iitk.ac.in

Research Presentations

Debkanta Chakraborty

Title: Spatial Multi-omics Analysis Using Deep Generative Modeling

TL;DR: Spatial multiomics integrates high-resolution spatial context with single-cell genomic, transcriptomic, proteomic, and/or epigenomic data, enabling unprecedented insights into cellular heterogeneity, tissue organization, and microenvironmental interactions. Applications span developmental biology, cancer research, and immunology, offering a transformative lens to dissect complex biological systems with both molecular and spatial precision. Here we introduce a deep generative model for spatial multiomics analysis and also we generate simulated datasets using probabilistic modelling.

Abstract: Spatial multiomics is an emerging field that integrates multiple layers of molecular information—such as genomics, transcriptomics, epigenomics, and proteomics—while preserving spatial context within tissues. Unlike traditional single-cell methods that rely on dissociating cells, spatial multiomics retains the native architecture of the tissue, allowing for a deeper understanding of cellular heterogeneity, microenvironmental influences, and spatial gene expression patterns. This approach is particularly valuable in studying complex biological systems, such as tumor microenvironments, neural circuits, and developmental processes.

Recent technological advances, including spatially resolved transcriptomics (e.g., 10x Genomics Visium, MERFISH, Slide-seq) and multiomic integration techniques, have enabled the simultaneous profiling of gene expression, chromatin accessibility, and protein abundance within intact tissues. Computational non-spatial methods such as Seurat, TotalVI etc and spatial methods such as STAGATE, and SpatialGlue have been developed to enhance clustering accuracy, improve spatial domain identification, and optimize latent space representations for better cell-type annotation. These models leverage deep learning and graph-based techniques to capture spatial dependencies, improve cell clustering performance (measured using ARI and NMI), and enhance the robustness of multiomic data integration.

By integrating high-dimensional molecular data with spatial localization, spatial multiomics provides unprecedented insights into cellular communication, tissue organization, and disease mechanisms. This field has broad applications in cancer research, neuroscience, developmental biology, and regenerative medicine. The continued development of more accurate computational frameworks, improved experimental protocols, and scalable data integration strategies will further advance our ability to decode the complex molecular landscape of tissues, leading to new discoveries in precision medicine and targeted therapies.

Here in our project, we integrate graph neural networks (mainly graph attention layers), variational autoencoders and contrastive learning for spatial multiomics analysis. We use spatial graphs to model cell-cell interactions and leverage the generative power of VAE-based latent space regularization. We also generate the simulated spatial multiomics datasets using methods such as Potts Model, Non Spatial Factorization, Probabilistic modeling etc. For real datasets such as human lymph nodes and synthetic in-house datasets, we compare our results with SpatialGlue, Stagate etc for benchmarking our method against the state of the art models.

Co-authors: Jaskaran Singh Walia, Debkanta Chakraborty, Agnish Bhattacharya, Swastik Singhal, Goural Dureja, Vikas Yadav, Hamim Zafar

Email: debkanta@cse.iitk.ac.in

Research Presentations

Musale Krushna Pavan

Title: multiHIVE: Hierarchical Multimodal Deep Generative Model for Single-cell Multiomics Integration
TL;DR: multiHIVE, deep generative model for inferring cellular embeddings by integrating CITE-seq data modalities, with hierarchically stacked latent variables as well as modality-specific latent variables to capture shared and private information from the modalities respectively.

Co-authors: Anirudh Nanduri, Musale Krushna Pavan, Kushagra Pandey, Hamim Zafar
Email: krushna@cse.iitk.ac.in

Priya

Title: PHALCON: phylogeny-aware variant calling from large-scale single-cell panel sequencing datasets
TL;DR: We present PHALCON, a novel tool for scalable mutation detection and tumor phylogeny reconstruction from large-scale single-cell panel sequencing data, while simultaneously modeling errors inherent to such technologies.

Co-authors: Priya, Sunkara B V Chowdary, Aditya Gautam, Hamim Zafar
Email: priya22@iitk.ac.in

Sanjeet Singh

Title: Linguistic Inspired Pose-Stitching for End-to-End Sign Language Translation
TL;DR: This study introduces linguistic template-based augmentation and pose stitching for Sign language translation task. A simple transformer based encoder-decoder surpasses prior models, proving that novel pose-stitched data enhances sign language translation performance.

Co-authors: Abhinav Joshi, Vaibhav Sharma, Sanjeet Singh, Ashutosh Modi
Email: singhsanjeet2601@gmail.com

Research Presentations

Abhinav Anand

Title: Graph-Based Resource Allocation in Wireless Networks Using Policy Gradient RL

Co-authors: Abhinav Anand, Subrahmanya Swamy Peruru, Amitangshu Pal
Email: abhinavanand@cse.iitk.ac.in

Harshit Goel

Title: Un-WEIRDing End-User Programming

Co-authors: Harshit Goel
Email: hgoel@cse.iitk.ac.in

Anubhav Dixit

Title: SAM based object co-segmentation

Co-authors: Anubhav Dixit, Koteswar Rao Jerripothula
Email: anubhavdixit@cse.iitk.ac.in

Visit us at www.cse.iitk.ac.in

Designed by Roop Aparajita Subhra Purushottam on MS PowerPoint
Email: purushot@cse.iitk.ac.in
Web: www.cse.iitk.ac.in/users/purushot
April 2025