My primary research interests focus on **Machine Learning and Optimization**, and their applications in **Energy Systems**. Being a graduate in **Electrical Engineering** from **Indian Institute of Technology (IIT), Dhanbad**, I have always been driven towards working on innovations related to power grids. Besides this, the documentary on AlphaGo and its portrayal of how artificial intelligence (AI) can delve into the uncertain corners beyond what human minds can think of, hooked me to the fascinating **applications of reinforcement learning (RL)**. More recently, my research vision got shaped by the experience at **Microsoft Research (MSR), India** where my work revolves around building **"Technology for Emerging Systems"**. Although RL has several shortcomings in its real-world deployability owing to its assumption of a Markovian world and the curse of being real-world sample inefficient, I believe working on defining safe policies and constrained action value pairs can potentially solve several complex problems in electrical energy sciences. Hence, through graduate studies at Columbia University, I plan to develop reliable AI-enabled solutions robust-enough to drive sustainability goals in next-generation energy systems.

In the recent past, with the evolution of AI, there has been a lot of interest in how to cater to the needs of modern Power Systems. At MSR, I work with Dr. Akshay Nambi, Tanuja Ganu and Dr. S. Kalyanaraman on applying machine learning for optimization and control in distributed energy resources (DERs). Specifically, my work explores various RL techniques, stochastic and traditional optimization algorithms that solve the physics-based constraints associated with power systems.

Renewable energy sources depend heavily on weather and climate changes leading to uncertainty in the power generation values. This makes it difficult for energy operators to make reliable planning and operations decisions. To deal with this problem, we developed a unified framework, **"EnCortex"**, that provides extensible and easy-to-use abstractions to plan, build and execute real-world scenarios efficiently. More recently, my focus has been on expanding the core energy abstractions such as, modeling the local constraints and dynamics involved in the formulation of a Li-ion battery and inheriting the functionalities to similar storage classes like pumped hydro-based systems. Additionally, I then prepare modular scenario-specific gym-compatible environments to pair the required abstractions and tie them with a multitude of extensible optimization algorithms. Also, my current vocation involves making the framework compatible enough to handle a range of use cases in the energy domain. It is now offered to a large global **energy partner of Microsoft** for wide scalability and has been submitted for publication at a **top-tier Systems Conference** [1]. Through the results, I demonstrate that the neural network-based policy learners outperform the traditional algorithms in the presence of forecast errors with sufficient training data.

Energy operators face several challenges while applying classical RL to large-scale real-world environments. Getting a deployed agent to interact with the environment i.e. Learning from Demonstrations (LfD) for solving its exploration-exploitation trade-offs, can prove to be detrimental in critical systems. Moreover, the amount of data collected online while interacting with the environment may also be limited, making it difficult for the RL agent to learn anything useful. Therefore, to deal with these issues:

- I worked on **OfflineRL**, namely on Conservative Q-learning (CQL) algorithm where I proposed first to use Model Predictive Control (MPC) to solve the objective functions of the scenario by maintaining the constraints on the historically available datasets. Then we make our own Markov Decision Process based Datasets which is used as a supervised learning dataset. The performance boost of the OfflineRL approach compared to a simple RL agent trained from scratch showed the superiority of the procedure.
- I experimented on **imitation learning** or more specifically, **behaviour cloning** to have the environment interactions performed as a pre-processing step. I used a trained RL expert to extend its learning trajectories and demonstrations to a new agent working on a similar environment. This benefits the energy operators to train the optimizers without having the need for large datasets.

Therefore, I hope to leverage the experience gained while working on this project to further research and engineer large-scale projects, improve the current modalities of optimization and RL algorithms and its applicability to different decision making problems.

Having interests in solving challenges faced in the wide adoption of e-Mobility, I worked as a research intern at MSR on modeling energy consumption (EC) of batteries in electric vehicles (EV). EV owners are generally concerned about untimely battery drainage causing **"range-anxiety"**. To address this, I identified the challenges faced in real-world EC modeling and proposed the nature of data required to understand the phenomenon, and developed a two-stage approach (by extensive feature engineering using domain knowledge) to predict the EC of an EV before the start of the trip. The explainability and interpretability of the proposed models, alongside the boost in performance over non-explainable neural network models, led to a **successful collaboration** with a large e-bus company in India. Also, the work recently got published in **ACM COMPASS'22** [2].

Through the research experience gained across domains specific to applied RL and distributed control in energy systems, I would like to improve on the **theoretical aspects** of optimization to pursue advanced research. In this direction, I have always striven for academic excellence - during my Bachelor's; I graduated as the **Silver Medalist** of my cohort. I believe the experiences gained while working on these research projects have prepared me to work on open problems alongside being persistent. Working with researchers has helped me to interpret better and address the reviews and comments in paper submissions. I believe that valuable soft skills like time management and collaboration are vital for survival in challenging academic settings.

Columbia University's rich collaborative environment and interdisciplinary research ideology perfectly align with my research interests and motivations. Having an interest in contributing to the machine learning community with my experiences across energy systems, I have envisaged certain possible research directions in which I would like to work on during my graduate studies:

1. Advancing the understanding of **multi-agent reinforcement learning** systems :
   - Agents working **in cooperation**: where agents are the various resources of a producer solving cooperatively to optimize for a common objective, abiding by the global and the local constraints.
   - In a **competitive** setting: simulate an energy market scenario where different participants bid volumes and prices, thus agents acting against each other to have their bids accepted.
2. The environment can never be deterministic since the decisions are made on forecasted values before logging the actual values of a generation. Hence, this calls for work on the following two problems:
   - Build accurate forecasting algorithms reducing forecasting errors to produce optimal results. Generate uncertainty values, thus providing a **relaxation-bucket** for accepted sub-optimality in the results.
   - Design various Model-free algorithms for **robust control** that work on not a sufficiently accurate environment but bring out the best possible results subject to the uncertainty in the forecasts.
3. I want to delve deeper into understanding the risk-averse, constrained, worst-case criteria to be included in the exploration process of any learning agent. Hence, working on the **safety** issues of individual agent-environment interactions while considering the uncertainties involved across parameters is another possible research direction I want to explore.
4. I see potential applications of convex, non-convex optimization, stochasticity, and game theory applications in **sequential decision-making** of dynamic volume allocation and market-bidding opportunities.
5. Energy systems require managing physical and cyber variables to monitor, communicate and control the evolving complexities. This calls for solving a new challenge: detecting security threats in the Cyber-Physical systems (CPS). The availability of fewer training instances under zero-day attacks motivates me to explore **reliable one-shot learning Deep RL** techniques as another exciting research direction.

These are some specific research directions that I like, but in general I am open to exploring a broader spectrum of problems in developing new modalities of optimization algorithms and control for any autonomous systems. Over the years, I have been influenced by the exciting research done by **Dr. Shipra Agrawal** and **Dr. Vineet Goyal** in optimization under uncertainty, reinforcement learning and data-driven decision-making. I have a strong desire to work with them, while collaborating with **Dr. Bolun Xu** and **Prof. Daniel Bienstock**, to advance the understanding of uncertainties present and tackle the respective loopholes existing in current smart grids for safe and reliable decision-making. Along similar lines, I also aspire to work with **Dr. Christian Kroer** in his Coffee and Convexity Lab, to build AI-enabled solutions addressing large-scale economic complexities present in modern-day electricity markets, revenue management and develop methods for robust ML applications.

I believe the Ph.D. program at Columbia University will provide me the valuable exposure to grow as a researcher and push me towards my future goal of leading my own "Technology for Emerging Systems" based research laboratory. In summary, I believe I bring with me, research experience, industry-tested programming and engineering skills, soft skills, and, most importantly, an unquenchable thirst for knowledge and excellence. Therefore, a Ph.D. from Columbia University is the next significant achievement I eagerly anticipate in my life.

[1] EnCortex: A General, Extensible and Scalable Framework for Decision Management in New-age Energy Systems, Vaibhav Balloli*, **Millend Roy***, Anupam Sobti, Tanuja Ganu, Akshay Nambi. (Under Review).

[2] Reliable Energy Consumption Modeling for an Electric Vehicle Fleet, **Millend Roy**, Akshay Nambi, Anupam Sobti, Tanuja Ganu, Shivkumar Kalyanaraman, Shankar Akella, Jaya Subha Devi, and S A Sundaresan. In: *ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS). COMPASS'22. Seattle, WA, USA*