RESEARCH STATEMENT
Amutheezan Sivagnanam (Ph.D., Informatics, The Pennsylvania State University)

My research lies at the intersection of deep reinforcement learning, operations research, and optimization, with a strong emphasis on developing efficient algorithms to solve complex, real-world problems. My work mixes deep reinforcement learning with classic optimization and lots of real-world data so the systems make smart decisions in seconds and keep costs down. My research has been published in peer-reviewed A* conferences and deployed in real-world applications within the transportation domain.

My interest in machine learning began when I worked with HL7 messaging protocols used in the healthcare industry. I developed an algorithm to track resource shortages by analyzing admission and registration messages and proposed a way to analyze doctors' diagnoses for early warning of emerging diseases. Later, I explored how natural language processing and data modeling could enhance this work. Accordingly, during my undergrad thesis, I focused on classifying the sentiment of tweets using label-efficient semi-supervised learning. I introduced a novel co-training based solution approach with efficient feature vector combination that outperformed existing state-of-the-art approaches in terms of F1-scores.

During my Ph.D., I focused on formulating and solving complex combinatorial optimization problems using heuristic and meta-heuristic as well as deep reinforcement learning (DRL). These included dynamic vehicle routing with advanced requests and real-time emergency responder repositioning, systems characterized by stochasticity, partial observability, and hard temporal constraints. Prior solution approaches were either too slow or lack in solution quality with respected well known evaluation metrics in their respective settings. In contrast, I developed DRL-based agents that learn near-optimal policies via environment interaction and leverage neural approximations for rapid inference during deployment.

Accordingly, my research addressed three core challenges in sequential decision-making under uncertainty: (i) **Proactive resource reallocation**: introduced a Deep Deterministic Policy Gradient (DDPG) based solution approach to proactively reposition ambulances in anticipation of future demand. Further, I utilize a Transformer-based actor network that handles variable dimensional complex and discrete state-and-action space, to train the decision agent to make real-time decisions. Eventually, I replace the computationally expensive online tree search methods with trained DRL decision agent (ii) **Online Booking with Flexibility**: formulated a model combining two major components (i.e., online booking and offline vehicle routing), where I use a trained DRL agent to determine the the tight pickup time from user specified broad pickup windows at the time of booking and use off-the-shelf offline VRP solver to generate high-quality route plans. This system balances efficient utilization of limited resources while maintaining real-time responsiveness. (iii) **Dynamic Routing with Continuous Reoptimization**: devise a mathematical model to support advance booking in dynamic routing, where I use simple insertion heuristics to determine the acceptance/rejection of request in real-time and utilize anytime VRP solver to optimize the routes between arrival or requests. The result was significantly improved service rate in high-demand environments.

Currently, my research focuses on applying probabilistic modeling to integrate multi-modal transportation systems, aiming to enhance accessibility, efficiency, and equity in urban mobility. In this work, I focus on predicting user choices between end-to-end on-demand transit over cost-effective multi-modal transit, supported by high-frequency bus-transit services. In order to learn the optimal policy, I utilize Bayesian optimization techniques to find the optimal decision making strategy. This work is motivated by the need to design transportation networks that are not only cost-effective but also responsive to the evolving needs of diverse communities especially underserved populations.

Building on my expertise in real-time control systems, my future research will focus on extending deep reinforcement learning algorithms to more complex, multi-agent, and hierarchical decision making settings. Additionally, I plan to explore offline reinforcement learning and safe exploration techniques, especially in domains where exploration is costly or risky, such as emergency response and urban planning. This work will contribute to the development of robust, generalizable controllers that can operate effectively under uncertainty and in the presence of incomplete information.