Research Statement

Balajee Kannan
bkannan@ri.cmu.edu

My research goal is to design intelligent systems capable of completing complex tasks in dynamic, unstructured environments in as user friendly manner as possible. To this end, I seek to understand and actively pursue research in the areas of fault-tolerance, machine learning, multi-robot/agent systems and human-robot interaction.

As intelligent systems become pervasive, it is absolutely essential for every researcher in this field to understand and enable effective human-robot teams engaged in complex tasks. There is a fundamental difference in the way humans interact with robots that can only be understood to thorough analysis and experimentation. Additionally, it is important that the developed systems are capable of autonomously adapting to the variations in the operating environment while maintaining the overall objective of assisting humans to accomplish tasks. Such a system involves complex coordination and cooperation between humans and robots, with the robots displaying ability to learn from experience, adapt to changing environment and seamlessly integrate information to-and-from humans . This belief has been at the core of my research focus through my Ph.D. and beyond.

Current Research

Coordinating Human-Robot Teams in Unstructured Environments

Robotic solutions are addressing problems in increasingly complex domains in which a single agent acting alone cannot achieve the required goals. Disaster response, agriculture, mining, and construction all involve several spatially distributed tasks that must be executed by teams of autonomous agents such as robots and/or humans. My principal research objective is to significantly advance the state-of-the-art in autonomously coordinated team performance in complex unstructured environments. We have shown that for such scenarios, a combination of market-based and operations research-based coordination and planning techniques can be effectively used to coordinate human-robot teams operating in unstructured environments. Additionally, robustness of operation can be improved through team monitoring, interventions by humans, and interventions by specialized agents for error recovery. Humans will also be able to take low-level control of a robot (or robots) to correct problems and then have the team continue autonomously. The coordination mechanism we build will also have enhanced flexibility in several dimensions including customizability for different operations, fluidity of team members, and options for planning mechanisms thus allowing for highly dynamic operating environments. In addition, I plan to create a framework of metrics for evaluating team performance. This work will include enhancing the infrastructure for evaluating performance and making strategic changes before, during, and after execution using appropriate metrics, and incorporating metrics into costing, recruiting, and modeling and exposing faults and errors that cannot be detect by individual robots.

Dissertation Research

The objective of my research, in part, is to outline a framework for developing a turn-key solution for fault diagnosis in complex teams of heterogeneous mobile robots. The approach, called LeaF – Learning based Fault-diagnostic framework for multi-robot teams – consists of two essential parts: a partial causal model for representing the various faults in the system and an adaptive case-based learning algorithm for diagnosing previously unknown-faults based on experience and to automatically update their causal models. In addition to developing a generic fault-diagnostic system, a further contribution of my research is the development of metrics to measure the overall performance and the fault-tolerance exhibited by the systems. The developed metrics are designed to be application-independent and can be applied to any fault-diagnostic architecture.

Approach: LeaF – Learning-based adaptive Fault-tolerance

For a robot team, the faults that occur during the course of operation provide an opportunity for the system to learn about its environment, thereby adapting to it. Based on this key principal, I designed a fault diagnostic architecture called LeaF. When an *unknown* fault is encountered, the system uses a case-based learning approach to adapt and categorize it and subsequently add it to the *a priori* causal model for future use. The use of learning allows the system to identify and categorize unexpected faults, enable team members to learn from problems encountered by others, and make intelligent decisions regarding the environment. By using a hybrid approach of learning and model building, each individual member of the team will be able to make the best possible decision regarding successful task execution.

LeaF is designed to be a modular architecture that combines typical robot control processes with modules for adaptive fault diagnosis and recovery. The identification of similarities between faults is achieved using a method called Lazy Induction of Descriptions (LID), originally developed by Armengol and Plaza. LID builds a symbolic similarity description, *similitude*, for the fault, finding the best match to one or more nodes in the causal model. LID identifies relevance by only selecting the nodes with similar characteristics to that of the encountered fault. In order to identify the relevance of the elements in the similitude set, I developed a diagnosis algorithm that determines which of the existing faults have similar characteristics or symptoms to the encountered fault. In addition, I use a probabilistic method, based on the frequency of occurrence of faults, to further streamline the diagnosis process and to prune the unnecessary nodes from the causal model. This approach provides a better reasoning for fault diagnosis and reduces overall system complexity and the time spent in fault diagnosis.

Evaluation Metrics: The ability to measure the extent of fault-tolerance exhibited by the system would provide the designer with a useful analysis tool for better understanding the system as a whole. I have developed metrics for calculating the influence of fault-tolerance towards system performance and outlined potential methods for analyzing these obtained measures towards evaluating the true capability of a multi-robot system.

Plans for Future Research

Incorporate prediction capabilities in LeaF: A problem that I am interested in exploring is the idea of incorporating fault avoidance in addition to fault-tolerance. If a system is capable of predicting a fault, then the said fault can be potentially predicted and subsequently avoided. This would especially be useful for autonomous intelligent systems operating in highly unstructured and dangerous environments where the cost of fault-tolerance maybe prohibitive. The idea then would be to extend the adaptability component of LeaF to attempt to predict faults based on deviations in measured sensor data. This work would involve several areas of traditional computer science including robotics, distributed artificial intelligence, machine learning, networking, etc., fused with statistical methods like probability analysis, hidden Markov modeling etc.

Explore potential Human-Robot Interaction issues: I am keen on exploring the role of humans interacting with robots in a peer-to-peer or supervisory capacity for diagnosing and predicting faults. Researchers have long identified the advantage and need for involving humans as team members in autonomous robot applications. I am interested in expanding LeaF to facilitate human intervention in predicting and preventing potential hazardous faults that can not be easily identified by an autonomous systems. This brings up a host of unresolved issues, like: How and when to best utilize human input? humans as part of robot team, or in a more supervisory role, Whom to "trust" more: humans or robots, How to autonomously regulate "trust", etc.