The Homeland Security Center for Dynamic Data Analysis at Rutgers University

Center for Dynamic Data Analysis

Optimization & Learning for Decision-making under Uncertainty

Mission

There is a wide variety of practical decision-making settings in which a decision maker has the ability to collect a finite amount of information before he or she must render a decision. Examples of important applications in homeland security include sampling the air for nuclear material, determining when to examine people and containers for smuggled material, managing mobile sensors, determining an action plan during a potential disease outbreak, and selecting transportation or evacuation routes during an emergency. In each instance, the decision maker has the ability to gather additional information (potentially costing time and money) in order to make a more informed decision and avoid the potential costs of an incorrect decision. What (and how much) additional information should he or she collect before, for example, sounding an alarm?

This project targets the general theme of optimal and adaptive learning in a dynamic environment. Our work focuses on developing methods that enable us to solve problems such as i) determining what data to collect to learn efficiently about the state of the world and ii) detecting change, and the cause of change, as quickly as possible. We have developed an information collection method called the knowledge gradient policy, which is an easily implemented rule that tells us how to sample information on competing alternatives to learn which is best. Applications arise in sampling the air for nuclear material, determining when to examine people and containers for smuggled material, trying out new scanning technologies, and managing sensors. We have also developed nearly optimal sequential change detection and identification algorithms for nonstationary data streams influenced by hidden factors, when false-alarm and misdiagnosis can be kept strictly small.

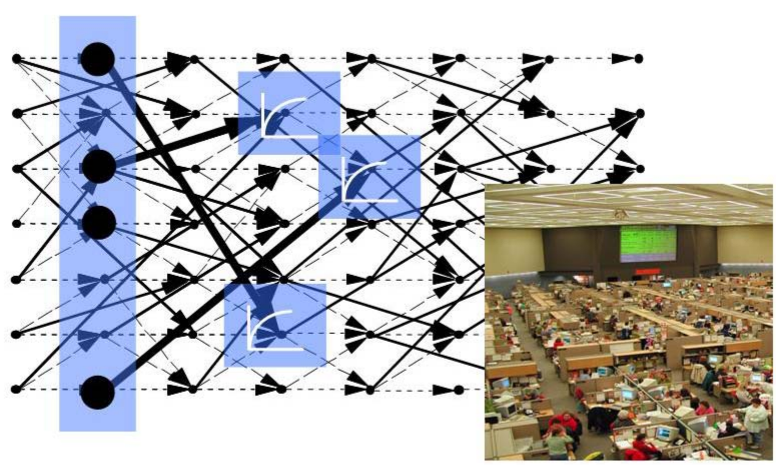

Many command and control applications require making sequential decisions about resource allocation in the presence of uncertainties that are evolving through time. We often model and solve such problems as large-scale dynamic resource allocation problems and apply techniques for iteratively learning needed information.

Benefit: *Our project is developing logic for optimally collecting information and for detecting fundamental changes in streams of data. These methods are applicable to a broad range of detection problems, such as determining when a nuclear detector has failed or when a new pattern of activity emerges.*

Collaborator(s):

- *Warren B. Powell*
- *Savas Dayanik*
- *Paul Kantor*

Funded by: *US Department of Homeland Security*

Early Development — Lab Prototype — Commercial Product

For more information, contact:
Warren Powell, powell@princeton.edu
(dydan.rutgers.edu)