

Guest Editorial: Special Issue on Constrained Decision-Making in Robotics

Marco Pavone · Stefano Carpin

the date of receipt and acceptance should be inserted later

As the complexity of tasks envisioned for robotic systems increases, their decision making module needs to be capable of concurrently evaluating and trading off multiple, possibly contradicting and stochastic objectives. For example, a rescue robot might be required to plan trajectories so as to maximize the probability of success to reach a critical location and, at the same time, minimize the duration of the traversal. Other more general examples include trading off information gathering versus energy expenditures, computation time versus optimality of a decision, or learning versus safety. A natural framework to address this class of problems is *constrained* decision-making, whereby a decision maker seeks to optimize a given cost function (often stochastic) while keeping other costs (usually involving risk assessments) below given bounds. In the last decade, the operations research community has made significant strides on the topics of constrained decision-making (notoriously more challenging than the unconstrained counterpart) and risk assessment in dynamic scenarios. The result is a comprehensive theory and a set of algorithmic tools for this class of decision-making problems. Yet, despite their relevance, these results have seen limited application within the robotics domain.

Accordingly, in July 2014 we organized the workshop “*Constrained decision-making in robotics: models, algorithms, and applications*” during the Robotics Science and Systems (RSS) conference in Berkeley, CA.

Marco Pavone
Department of Aeronautics and Astronautics
Stanford University
E-mail: pavone@stanford.edu

Stefano Carpin
School of Engineering
University of California, Merced
E-mail: scarpin@ucmerced.edu

The workshop had three objectives: (1) to convene together researchers working in the areas of decision-making, risk theory, and robotics, (2) to inform robotic researchers about the state of the art in constrained decision-making and modern risk theory, and (3) to formulate a research agenda on the topics of risk modeling and constrained decision-making for robotic applications.

After the workshop, participants were invited to submit an extended version of their work. The call for contributions was also extended to the broader research community and widely circulated. Eventually, six papers were selected for publication and are included in this special issue. Collectively, these papers covers several robotic application domains for constrained decision-making, ranging from sensors networks to autonomous vehicles and space robotics.

The first three papers deal with a *deterministic* problem setup. Specifically, the paper “Cognitive Robots Learning Failure Contexts through Real-World Experimentation” by Sertac Kapinar and Sanem Sariel studies how to endow robots with the ability to learn about their limitations of performance when it comes to task execution, and to capitalize on such knowledge to formulate risk-averse decisions subject to performance constraints. The authors tackle the problem by using an Inductive Logic Programming (ILP) framework, whereby hypotheses are formulated that relate execution contexts to action outcomes. Hypotheses are continuously updated based on online observations and are used to model risks, constraints, and limitations on task execution. On top of this knowledge base, a planner formulates decision sequences fulfilling the constraints derived from the learned hypotheses. The theoretical framework is experimentally demonstrated and validated both on a mobile robot and on a 7-DOF robotic arm perform-

ing object manipulation tasks.

The paper “A Mechanism for Real-Time Decision Making and System Maintenance for Resource Constrained Robotic Systems through ReFrESH” by Yanzhe Cui, Richard M. Voyles, Josh T. Lane, Akshay Krishnamoorthy, and Mohammad H. Mahoor addresses the problem of reliability and integrity of complex robotic systems operating in critical domains, like urban search and rescue. The proposed ReFrESH system periodically monitors system performance and degradation, searching for faulty modules. Once a failure is identified, new task configurations are formulated so that they are achievable given the system feasibility constraints computed at run time. The decision-making process follows a greedy approach in a constrained search space whose structure is determined by the identified faulty modules. The method is validated on a two-robot system performing a visual servoing task.

The paper “Computationally-Aware Control of Autonomous Vehicles: A Hybrid Model Predictive Control Approach” by Kun Zhang, Jonathan Sprinkle, and Ricardo G. Sanfelice considers the design of nonlinear model predictive control algorithms under computation constraint, with a focus on controlling a ground vehicle. The strategy is to consider two models of the vehicle for model predictive control, one coarse (allowing fast computation) and the other more accurate. The authors introduce a metric called uncontrollable divergence and, using this metric, propose a mechanism to select the model to use in the predictive controller. The map of uncontrollable divergence plotted over the state space gives the criterion to judge where the coarse model can be tolerated when a high update rate is preferred (e.g., at high speed and small steering angles), and where the high-fidelity model is required to avoid obstacles or make tighter curves (e.g., at large steering angles). Based on this metric, a mechanism for switching between multiple predictive controllers is developed in order to lower the controller’s return time while maintaining predictive accuracy. This work is relevant to many kinds of problems in cyber-physical systems where the runtime of a computer-in-the-loop affects the performance of the system.

The last three papers deal with a *stochastic* problem setup. Specifically, the paper “Constrained Decision-making for Low-count Radiation Detection by Mobile Sensors” by Jianxin Sun and Bert Tanner considers the problem of detecting weak mobile radioactive sources using multiple robots. The authors focus on the detection problem. In other words, assuming that the emis-

sion characteristics and trajectories of the sources are known, the problem is to compute an optimal control policy to place the sensors so that the missed detection rate is minimized while the number of false positives is constrained below a given bound. The proposed solution consists of two layers. The high-level layer determines a collective strategy for the sensors. The low-level layer implements the strategy on each of the individual mobile sensor and accounts for the constraints imposed by the environment (e.g., obstacle avoidance) and the sensors themselves (e.g., motion constraints.) The proposed methodology is demonstrated both in simulation and on a three-dimensional testbed involving both an aerial and a ground platform carrying a radioactive sensor.

The paper “Decentralized Probabilistic Density Control of Autonomous Swarms with Safety Constraints” by Nazl Demir, Utku Eren, and Behçet Açıkmeşe considers the problem of controlling a swarm of robots in a decentralized fashion while ensuring safety constraints (chiefly, collision and conflict avoidance). The idea is to let the robots move probabilistically (according to statistically independent decisions) and specify the time evolution of their probabilistic density distribution by using a Markov chain. The safety constraints, e.g., density upper bound constraints for collision avoidance, are formulated as equivalent linear inequality conditions on the Markov chain matrix by using the duality theory of convex optimization. The authors then propose i) an offline method based on linear matrix inequalities, and ii) an online method based on quadratic programming for the synthesis of a Markov matrix that guides the swarm to a desired steady-state final distribution, while fulfilling the aforementioned safety constraints. Specifically, the offline method synthesizes constant Markov matrices, while the online method utilizes real-time estimates of the swarm density distribution to continuously update the Markov matrix in order to maximize the convergence rate.

The paper “Chance-Constrained Dynamic Programming with Application to Risk-Aware Robotic Space Exploration” by Masahiro Ono, Marco Pavone, Yoshiaki Kuwata, and J. (Bob) Balaram presents a solution algorithm for dynamic programming problems with joint probabilistic (chance) constraints. In such problems one desires to compute a control policy that minimizes an expected cost while guaranteeing that the probability of constraint violation is within a user-specified risk bound. The approach is to (conservatively) reformulate a joint chance constraint as a constraint on the expectation of a summation of indicator random vari-

ables, which can be incorporated into the cost function by considering a dual formulation of the optimization problem. As a result, the primal variables can be optimized by standard dynamic programming, while the dual variable is optimized by a root-finding algorithm that converges exponentially. The algorithm is demonstrated on three optimal control problems, namely a path planning problem, a Mars entry, descent, and landing problem, and a Lunar landing problem.

Papers included in this special issue are just a small, yet representative sample of the findings produced by a rapidly growing research community. With robotics experiencing an explosive growth in the breadth of applications, more and more complex domains are being considered, and this necessarily leads to situations where multiple objectives and constraints should be accounted for. In our opinion, this represents one of the most fertile areas for both practitioners and theoreticians.