Automatica 50 (2014) 2987–2997

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Consistent approximation of a nonlinear optimal control problem with uncertain parameters ☆

Chris Phelps [a], Qi Gong [a,1], Johannes O. Royset [b], Claire Walton [a], Isaac Kaminer [c]

[a] Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA, United States

[b] Department of Operations Research, Naval Postgraduate School, Monterey, CA, United States

[c] Department of Mechanical and Astronautical Engineering, Naval Postgraduate School, Monterey, CA, United States

ARTICLE INFO

Article history:
Received 30 April 2013
Received in revised form
6 March 2014
Accepted 11 July 2014
Available online 28 October 2014

Keywords:
Optimal control
Computational methods
Optimization
Nonlinear system
Search theory

ABSTRACT

This paper focuses on a non-standard constrained nonlinear optimal control problem in which the objective functional involves an integration over a space of stochastic parameters as well as an integration over the time domain. The research is inspired by the problem of optimizing the trajectories of multiple searchers attempting to detect non-evading moving targets. In this paper, we propose a framework based on the approximation of the integral in the parameter space for the considered uncertain optimal control problem. The framework is proved to produce a zeroth-order consistent approximation in the sense that accumulation points of a sequence of optimal solutions to the approximate problem are optimal solutions of the original problem. In addition, we demonstrate the convergence of the corresponding adjoint variables. The accumulation points of a sequence of optimal state-adjoint pairs for the approximate problem satisfy a necessary condition of Pontryagin Minimum Principle type, which facilitates assessment of the optimality of numerical solutions.

1. Introduction

In the last decades, a variety of computational algorithms have been developed for solving constrained nonlinear optimal control problems, including Euler (Polak, 1997, chap. 4), Runge–Kutta (Kameswaran & Biegler, 2008; Schwartz & Polak, 1996), and Pseudospectral (Gong, Kang, & Ross, 2006; Kang, 2010; Ross & Karpenko, 2012). These computational optimal control methods have achieved great success in many areas of control applications (Bedrossian, Bhatt, Kang, & Ross, 2009; Bedrossian, Karpenko, & Bhatt, 2012; Chung, Polak, Royset, & Sastry, 2011; Li, Ruths, Yu, & Arthanari, 2011). In a standard nonlinear optimal control problem, the objective functional is of the Bolza type, which consists of an end cost as well as an integral over the time domain. In this paper we are interested in a class of non-standard optimal control problems in which the objective functional involves an expectation of a Bolza-type cost functional over a space of stochastic parameters. This class of problems is defined in the following.

Problem B. Determine the function pair $\{x, u\}$ with $x \in W_{1,\infty}([0,1];\mathbb{R}^{n_x})$, $u \in L_\infty([0,1];\mathbb{R}^{n_u})$ that minimizes the cost functional

$$J = \int_\Omega \left[F\left(x(1),\omega\right) + G\left(\int_0^1 r(x(t),u(t),t,\omega)dt \right) \right] p(\omega)d\omega$$

subject to the dynamics

$$\dot{x}(t) = f(x(t),u(t)), \tag{1}$$

initial condition $x(0) = x_0$, and the control constraint $g(u(t)) \le 0$ for all $t \in [0,1]$.

In Problem B, $W_{1,\infty}([0,1];\mathbb{R}^{n_x})$ is the space of all essentially bounded functions with essentially bounded distributional derivatives, which map the interval $[0, 1]$ into the space $\mathbb{R}^{n_x}$, and $L_\infty([0,1];\mathbb{R}^{n_u})$ is the set of all essentially bounded functions. The function p is a continuous probability density function for the stochastic parameter $\omega \in \Omega \subset \mathbb{R}^{n_\omega}$ and we allow r to be vector valued: that is, $r : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^1 \times \mathbb{R}^{n_\omega} \mapsto \mathbb{R}^K$, $G : \mathbb{R}^K \mapsto \mathbb{R}$.

☆ This work is supported by US Office of Naval Research under Grant N0001412WX21229. The material in this paper was partially presented at the 51st IEEE Conference on Decision and Control (CDC), December 10–13, 2012, Maui, Hawaii, USA. This paper was recommended for publication in revised form by Associate Editor Michael V. Basin under the direction of Editor Ian R. Petersen.

E-mail addresses: cdphelps@soe.ucsc.edu (C. Phelps), qigong@soe.ucsc.edu (Q. Gong), joroyset@nps.edu (J.O. Royset), cwalton@soe.ucsc.edu (C. Walton), kaminer@nps.edu (I. Kaminer).

[1] Tel.: +1 831 459 3753; fax: +1 831 459 4482.

http://dx.doi.org/10.1016/j.automatica.2014.10.025
0005-1098/

Problem B can be viewed as a generalization of the standard nonlinear optimal control problem where the cost function does not involve the stochastic parameter ω. Such a problem formulation allows a broad range of existing control problems to be extended to incorporate parameter uncertainty. For instance, in a number of optimal control applications such as asset protection (Ding, Rahmani, & Egerstedt, 2009) and target tracking (Quintero, Papi, Klein, & Chisci, 2010), the objective functional depends on other agents whose behavior may involve parameter uncertainty. Another application which can be addressed using this formulation is optimal path planning in uncertain environments, such as aircraft routing in a threat environment (Zabarankin, Uryasev, & Pardalos, 2002) or navigating an unmanned surface vehicle in a riverine environment (Gadre, Du, & Stillwell, 2012). Problem B is also closely related to *ensemble control* problems studied in, e.g. Ruths and Li (2010, 2012), where the uncertainty appears in both the cost function and the state dynamics.

Our main motivation to study such non-standard optimal control problems is from the topic of *optimal search for uncertain targets*. Work on search theory can, in general, be divided into two categories depending on how the target is modeled. Mangel (1989) provides a review of the components of the problem and various models used. In the first category, the motion of the target is given by a Markov process. Hellman (1970) and Mangel (1981, 1982) address the problem of computing the posterior distribution of the target's position. Necessary and sufficient conditions for a search plan to be optimal are developed in Hellman (1972), Ohsumi (1991) and Saretsalo (1973). The second category considers targets whose dynamics are conditionally deterministic, which means that the motion of the target depends on a stochastic parameter, and if the value of this parameter is known, the location of the target will be known for each time instance. Such conditionally deterministic targets are considered in Chung et al. (2011), Foraker (2011), Foraker, Royset, and Kaminer (submitted for publication), Lukka (1977), Phelps, Gong, Royset, and Kaminer (2012), Pursiheimo (1976), Royset and Sato (2010) and Sato and Royset (2010), where optimal search plans are given by the solutions to some optimal control problems with objective functionals involving an integral over a space of stochastic parameters, as well as the typical integral over the time-domain. Such optimal search models with conditionally deterministic targets belong to the non-standard optimal control problem considered in this paper, i.e., Problem B.

To briefly demonstrate how the search for conditionally deterministic targets can be modeled as Problem B, consider the problem of a searcher looking for a moving target in order to maximize the probability of detecting the target over some time horizon $[0, T]$ (without loss of generality we assume the time horizon is $[0, 1]$ as other time horizons can be handled by rescaling the time parameter). Let the searcher trajectory, $x(t)$, be determined by the dynamical system (1) with initial condition x_0. We assume that the target's motion is conditionally deterministic. In other words, there exists a random vector $\omega \in \Omega \subset \mathbb{R}^{n_\omega}$, such that the trajectory of the target conditioned on ω is given by $y(\cdot, \omega)$. It is assumed that the probability density of ω over Ω is known to the searchers and is given by $p : \Omega \mapsto \mathbb{R}^+$. The final component of the search model is a function describing the effectiveness of the searcher. Let $\tilde{r} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} \mapsto \mathbb{R}$ be the instantaneous rate of detection such that the probability of detection in a sufficiently small interval $[t, t+\Delta t]$, conditioned on ω, is given by $\tilde{r}(x(t), y(t, \omega))\Delta t$. The rate function $\tilde{r}$ is chosen to model the qualities of sensors such as acoustic and sonar sensors. Denote $P(t)$ to be the probability of non-detection at time instance t conditioned on ω. Then

$$P(t + \Delta t) = P(t)(1 - \tilde{r}(x(t), y(t, \omega))\Delta t).$$

As $\Delta t \to 0$ we get

$$P(t) = \exp\Big(-\int_0^t \tilde{r}(x(\tau), y(\tau, \omega))dt\Big).$$

Thus the probability that the target is not detected in the time interval $[0, 1]$ is given by the integral

$$J = \int_\Omega \exp\Big(-\int_0^1 \tilde{r}(x(t), y(t, \omega))dt\Big)p(\omega)d\omega.$$

The problem of finding the trajectory for the searcher which minimizes the probability of not detecting the target can now be framed as a special case of Problem B, with cost functional given by $J[\cdot]$ defined above. Detailed derivation of optimal search models including the construction of detection rate function $\tilde{r}$ can be found in Chung et al. (2011), Foraker (2011) and Foraker et al. (submitted for publication), as well as in Section 5 where an example of an optimal search problem is solved.

Given the difficulty in solving standard nonlinear optimal control problems, it is not surprising that the inclusion of the expectation of the cost functional over the parameter space, combined with the nonlinear dynamics and control constraints, makes Problem B particularly challenging. In the literature, some aspects of Problem B are considered, usually in simplified settings. Early studies into the search problem consider simplified searcher dynamics or conditionally deterministic targets subject to additional special restrictions (Lukka, 1977; Pursiheimo, 1976; Stone, 1977). For example, a necessary condition for optimality is developed in Pursiheimo (1976) for a type of optimal search problem with discrete parameter space. In Lukka (1977), a necessary condition for optimality in the continuous-space setting is derived for a single integrator linear dynamics and a box control constraint. More recent works consider general constrained nonlinear dynamics. In Chung et al. (2011) a numerical algorithm is provided to calculate an optimal solution for a special case of search for a target moving at a constant velocity in a channel. Foraker (2011) and Foraker et al. (submitted for publication) use a composite-Simpson integration scheme to discretize a two-dimensional parameter space and develop a computational method for solving a reduced version of Problem B. Foraker (2011) and Foraker et al. (submitted for publication) also analyze the performance of the computational method using Polak's consistent approximation theory (Polak, 1997, Section 3.3). Ruths and Li (2012) consider an optimal ensemble control problem, which is more general than Problem B in the sense that the uncertain parameter appears in both the cost function and the state dynamics. Consistency and convergence results are developed in Ruths and Li (2012) for a particular computational method based on a LGL-pseudospectral approximation in both the parameter and time domains.

In this paper we propose a computational framework for the solution of the uncertain optimal control Problem B. Based on the numerical approximation of the integral over the stochastic parameters in the objective functional, the considered uncertain optimal control problem can be approximated by a sequence of standard nonlinear optimal control problems, which can in turn be solved using existing computational methods such as Runge–Kutta (Kameswaran & Biegler, 2008; Schwartz & Polak, 1996) and pseudospectral (Gong et al., 2006) approaches. To ensure meaningful results in this computational framework, it is essential to guarantee that the discretization schemes provide valid approximations to the original non-standard optimal control Problem B. Indeed, even for standard optimal control problems, there are counterexamples showing that an inappropriately designed discretization may not be convergent (Cullum, 1972). In this paper, we show that the proposed computational framework approximates the optimal solution to the non-standard optimal control problem under mild assumptions. In particular, we show in Section 3 that the approximation based on the discretization process satisfies a zeroth-order consistency property. That is, accumulation points of a sequence of optimal solutions to the approximate problem are optimal solutions to the original uncertain optimal control problem. We contrast this condition to consistency and convergence results on