ARTICLE IN PRESS

Automatica ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Automatica

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

Technical communique

Finite-horizon inverse optimal control for discrete-time nonlinear systems☆

Timothy L. Molloy[a], Jason J. Ford[a], Tristan Perez[a,b]

[a] School of Electrical Engineering and Computer Science, Queensland University of Technology (QUT), Brisbane, QLD 4000, Australia
[b] Institute of Future Environments (IFE), Queensland University of Technology (QUT), Brisbane, QLD 4000, Australia

ARTICLE INFO

Article history:
Received 29 July 2016
Received in revised form 3 July 2017
Accepted 10 September 2017
Available online xxxx

Keywords:
Optimal control
Discrete-time systems
Nonlinear systems

ABSTRACT

In this note, we consider the problem of computing the parameters (or weights) of an optimal control objective function given optimal closed-loop state and control trajectories. We establish a method of inverse optimal control that exploits the discrete-time minimum principle. Under a testable matrix rank condition, our proposed method is guaranteed to recover the unknown objective-function parameters of finite-horizon discrete-time nonlinear optimal control problems.

1. Introduction

The problem of inverse optimal control arises in several areas of study in science and engineering including robotics (Aghasadeghi & Bretl, 2014; Aghasadeghi, Long, & Bretl, 2012; Hatz, Schlöder, & Bock, 2012; Maillot, Serres, Gauthier, & Ajami, 2013; Mombaur, Truong, & Laumond, 2010), and biomedical engineering (Priess, Conway, Choi, Popovich, & Radcliffe, 2015). For example, inverse optimal control has been applied successfully in the analysis of biological systems such as human locomotion (Aghasadeghi & Bretl, 2014; Hatz et al., 2012; Mombaur et al., 2010), human-posture control (Pauwels, Henrion, & Lasserre, 2014), and human-controlled aircraft motion (Maillot et al., 2013). Despite its broad potential applications, the theory of inverse optimal control in discrete-time finite-horizon settings has received limited attention. In this note, we develop a novel method of inverse optimal control for discrete-time finite-horizon systems.

Inverse optimal control is the problem of determining the unknown objective function (or alternatively the objective-function parameters) of an optimal control problem from optimal state and control trajectories (Hatz et al., 2012; Mombaur et al., 2010; Priess et al., 2015). Most theoretical treatments of inverse optimal control reported in the literature have focused on the infinite-horizon linear quadratic regulator (LQR) problem—see Boyd, Ghaoui, Feron, and Balakrishnan (1994, Section 10.6), Priess et al. (2015) and references therein. In this inverse LQR problem, existing approaches either assume knowledge of the optimal feedback gain matrix (Boyd et al., 1994, Section 10.6), or first recover it from the optimal state and control trajectories (Priess et al., 2015). With knowledge of the optimal feedback gain matrix, the objective-function parameters (i.e., weighting matrices) of the quadratic objective function are found by solving linear matrix inequalities (Boyd et al., 1994; Priess et al., 2015).

In order to solve finite-horizon inverse optimal control problems in general nonlinear continuous-time systems, Mombaur et al. (2010) proposed a bilevel (or nested) optimisation approach. This bilevel approach involves repeatedly solving optimal control problems with candidate objective-function parameters as part of a numeric optimisation. Due to its reliance on numeric optimisation, there are no theoretical results characterising the performance of this bilevel approach, and its computational expense is significant (cf. Johnson, Aghasadeghi, and Bretl, 2013 and Mombaur et al., 2010).

Recently, methods of inverse optimal control for nonlinear systems that avoid repeatedly solving the optimal control problem have been proposed on the basis of the Karush–Kuhn–Tucker conditions in discrete-time (cf. Keshavarz, Wang, & Boyd, 2011; Molloy, Tsai, Ford, & Perez, 2016; Panchea & Ramdani, 2015; Puydupin-Jamin, Johnson, & Bretl, 2012), and Pontryagin's minimum principle and the Hamilton–Jacobi–Bellman equation in continuous-time (cf. Hatz et al., 2012; Johnson et al., 2013; Pauwels et al., 2014). For continuous-time differentially flat

☆ This work was supported by an Australian Research Council Discovery Grant (DP14010089). The material in this paper was partially presented at the 55th IEEE Conference on Decision and Control, December 12–14, 2016, Las Vegas, NV, USA. This paper was recommended for publication in revised form by Associate Editor A. Pedro Aguiar under the direction of Editor André L. Tits.

E-mail addresses: t.molloy@qut.edu.au (T.L. Molloy), j2.ford@qut.edu.au (J.J. Ford), tristan.perez@qut.edu.au (T. Perez).

https://doi.org/10.1016/j.automatica.2017.09.023
0005-1098/

systems, Aghasadeghi and Bretl (2014) established theoretical conditions that guarantee the recovery of the unknown objective-function parameters. However, there appear to be few, if any, existing methods or theoretical results that guarantee the exact recovery of the objective-function parameters in finite-horizon discrete-time nonlinear inverse optimal control problems.

The main contribution of this note is the proposal of a method of finite-horizon discrete-time inverse optimal control. Unlike previous approaches (e.g., Molloy et al., 2016), our method exploits the discrete-time minimum principle and we are able to establish conditions that guarantee the recovery of the objective-function parameters. Furthermore, our method generalises beyond the classical discrete-time LQR setting, and does not require reconstruction of the optimal control law or the solution of candidate optimal control problems.

2. Problem formulation

Consider the deterministic discrete-time system

$$x_{k+1} = f(x_k, u_k), \quad x_0 \in \mathbb{R}^n \tag{1}$$

for $0 \leq k \leq K-1$ where $f(\cdot,\cdot) : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}^n$ is a given (possibly nonlinear) time-invariant function, $x_k \in \mathbb{R}^n$ are state vectors, and $u_k \in U_k$ are control inputs constrained to values in the sets $U_k \subset \mathbb{R}^m$. Let us define the objective function

$$V_K\left(x_{0K}, u_{0K-1}, \theta\right) \triangleq F(x_K, \theta) + \sum_{k=0}^{K-1} L(x_k, u_k, \theta),$$

with parameters $\theta \triangleq [\theta_1, \theta_2, \ldots, \theta_N]' \in \Theta \subset \mathbb{R}^N$. Here, we use $x_{0K} \triangleq \{x_0, x_1, \ldots, x_K\}$ and $u_{0K-1} \triangleq \{u_0, u_1, \ldots, u_{K-1}\}$ to denote the state and control trajectories respectively. We assume that the stage $L(\cdot,\cdot,\cdot)$ and terminal $F(\cdot,\cdot)$ objective functions are of the form

$$L(x_k, u_k, \theta) \triangleq \sum_{i=1}^{N_L} \theta_i L_i(x_k, u_k), \tag{2}$$

and

$$F(x_K, \theta) \triangleq \sum_{j=N_L+1}^{N} \theta_j F_j(x_K) \tag{3}$$

where $L_i(\cdot,\cdot)$ and $F_j(\cdot)$ are differentiable real-valued basis functions for $1 \leq i \leq N_L$ and $N_L < j \leq N$.

In the discrete-time finite-horizon optimal control problem, we are given the basis functions $L_i(\cdot,\cdot)$ for $1 \leq i \leq N_L$ and $F_j(\cdot)$ for $N_L < j \leq N$, the parameters θ, and the initial state x_0, and we solve the optimisation problem (Bertsekas, 1995):

$$\begin{aligned} \inf_{u_{0K-1}} \quad & V_K\left(x_{0K}, u_{0K-1}, \theta\right) \\ \text{s.t.} \quad & x_{k+1} = f(x_k, u_k), \quad k = 0, 1, \ldots, K-1 \\ & u_k \in U_k, \quad k = 0, 1, \ldots, K-1. \end{aligned} \tag{4}$$

In the inverse optimal control problem, we are given optimal closed-loop state $x^*_{0K} \triangleq \left\{x^*_0, x^*_1, \ldots, x^*_K\right\}$ and control $u^*_{0K-1} \triangleq \left\{u^*_0, u^*_1, \ldots, u^*_{K-1}\right\}$ trajectories that solve the optimal control problem (4) with unknown parameters $\theta = \theta^* \in \Theta$. Our aim is then to recover the unknown parameters θ^* from these optimal closed-loop trajectories under the following assumption.

Assumption 1. The open-loop system dynamics $f(\cdot,\cdot)$, the constraint sets U_k, and the basis functions $L_i(\cdot,\cdot)$ and $F_j(\cdot,\cdot)$ for $1 \leq i \leq N_L$ and $N_L < j \leq N$ are known and specify the optimal control problem (4) with $\theta = \theta^*$.

In this note, we also consider the inverse optimal control problem where we are given multiple pairs of optimal closed-loop state and control trajectories (with potentially different horizons and initial states) that are optimal under (4) with $\theta = \theta^*$. We shall let $x^{*,\ell}_{0K_\ell} \triangleq \left\{x^{*,\ell}_0, x^{*,\ell}_1, \ldots, x^{*,\ell}_{K_\ell}\right\}$ and $u^{*,\ell}_{0K_\ell} \triangleq \left\{u^{*,\ell}_0, u^{*,\ell}_1, \ldots, u^{*,\ell}_{K_\ell-1}\right\}$ denote the ℓth pair of optimal state and control trajectories with horizon K_ℓ for $\ell = 1, 2, \ldots, M$. Finally, we note that the parameters θ^* will only be recoverable up to an unknown scaling factor $0 < r < \infty$ since if the trajectories x^*_{0K} and u^*_{0K-1} are optimal under (4) with $\theta = \theta^*$, then they are also optimal under (4) with $\theta = r\theta^*$ for all $0 < r < \infty$.

3. Inverse optimal control

In this section, we propose a method of inverse optimal control that exploits the discrete-time minimum principle of Bertsekas (1995, Proposition 3.3.2).

3.1. Discrete-time minimum principle

In order to present the discrete-time minimum principle, let us define the Hamiltonian associated with (4) for any $\theta \in \Theta$ as the function

$$H(x_k, u_k, \lambda_k, \theta) \triangleq L(x_k, u_k, \theta) + \lambda_k' f(x_k, u_k) \tag{5}$$

for $0 \leq k \leq K-1$ where $\lambda_k \in \mathbb{R}^n$ are costate (or adjoint) variables. Let us also define the column vectors of first-order partial derivatives of $H(x_k, u_k, \lambda_k, \theta)$ with respect to x_k and u_k (and evaluated at x_k and u_k) as $\nabla_x H(x_k, u_k, \lambda_k, \theta) \in \mathbb{R}^n$ and $\nabla_u H(x_k, u_k, \lambda_k, \theta) \in \mathbb{R}^m$, respectively. We shall similarly use $\nabla_x f(x_k, u_k) \in \mathbb{R}^{n\times n}$ and $\nabla_u f(x_k, u_k) \in \mathbb{R}^{m\times n}$ to denote matrices containing the first-order partial derivatives of $f(x_k, u_k)$ with respect to x_k and u_k, respectively.

The discrete-time minimum principle is established under the following standard assumption on the convexity of the constraint sets U_k (see Bertsekas, 1995, Section 3.3.3 for a discussion of this assumption).

Assumption 2. The constraint sets $U_k \subset \mathbb{R}^m$ are convex for all $0 \leq k \leq K-1$.

Under Assumption 2, Proposition 3.3.2 of Bertsekas (1995) states that if Assumption 1 holds so that x^*_{0K} and u^*_{0K-1} are solutions to the optimal control problem (4) for a given $\theta^* \in \Theta$, then the trajectories x^*_{0K} and u^*_{0K-1} satisfy (1), and there exist costate variables $\lambda_k \in \mathbb{R}^n$ for $0 \leq k \leq K-1$ that satisfy the backwards recursion

$$\lambda_{k-1} = \nabla_x H\left(x^*_k, u^*_k, \lambda_k, \theta^*\right) \tag{6}$$

for $1 \leq k \leq K-1$ with the boundary condition

$$\lambda_{K-1} = \nabla_x F\left(x^*_K, \theta^*\right). \tag{7}$$

Furthermore, Proposition 3.3.2 of Bertsekas (1995) also establishes that the controls $u^*_k \in U_k$ satisfy the optimality condition

$$\nabla_u H\left(x^*_k, u^*_k, \lambda_k, \theta^*\right)' \left(u_k - u^*_k\right) \geq 0 \tag{8}$$

for all $u_k \in U_k$ and all $0 \leq k \leq K-1$. We shall exploit the optimality condition (8) to propose a method of inverse optimal control.

3.2. Proposed method

We require the following assumption to propose our method.

Assumption 3. The input u^*_k belongs to the interior (not the boundary) of the constraint set $U_k \subset \mathbb{R}^m$ for all $0 \leq k \leq K-1$.

Please cite this article in press as: Molloy, T.L., et al., Finite-horizon inverse optimal control for discrete-time nonlinear systems. Automatica (2017), https://doi.org/10.1016/j.automatica.2017.09.023.