



Towards global optimal control via Koopman lifts[☆]

Mario Eduardo Villanueva^{a,*}, Colin N. Jones^b, Boris Houska^a

^a School of Information Science and Technology, ShanghaiTech University, China

^b Automatic Control Laboratory, EPFL, Switzerland



ARTICLE INFO

Article history:

Received 2 March 2020

Received in revised form 14 October 2020

Accepted 21 February 2021

Available online 16 July 2021

Keywords:

Pontryagin's maximum principle

Koopman operators

Symplectic geometry

Optimal feedback control

ABSTRACT

This paper introduces a framework for solving time-autonomous nonlinear infinite horizon optimal control problems, under the assumption that all minimizers satisfy Pontryagin's necessary optimality conditions. In detail, we use methods from the field of symplectic geometry to analyze the eigenvalues of a Koopman operator that lifts Pontryagin's differential equation into a suitably defined infinite dimensional symplectic space. This has the advantage that methods from the field of spectral analysis can be used to characterize globally optimal control laws. A numerical method for constructing optimal feedback laws for nonlinear systems proceeds by computing the eigenvalues and eigenvectors of a matrix that is obtained by projecting the Pontryagin–Koopman operator onto a finite dimensional space. We illustrate the effectiveness of this approach by computing accurate approximations of the optimal nonlinear feedback law for a Van der Pol control system, which cannot be stabilized by a linear control law.

© 2021 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

During the last decades algorithms and software for computing local solutions of nonlinear optimal control problems have reached a high level of maturity (Biegler, 2007; Diehl et al., 2002; Zavala & Biegler, 2009), enabling its deployment in industrial applications (Qin & Badgwell, 2003). By now, such local solutions can be computed within the milli- and microsecond range. Moreover, auto-generated implementations of real-time local optimal control solvers run on embedded hardware systems (Houska, Ferreau, & Diehl, 2011; Mattingley & Boyd, 2009). In contrast to these developments in local optimal control, algorithms for locating global minimizers of non-convex optimal control problems, can hardly be claimed to be ready for widespread industrial application. There are at least two reasons for this. Firstly, generic global optimization methods can often only be applied to problems of modest dimensions. And secondly, their run-time usually exceeds the run-time of local solvers by orders of magnitude. Nevertheless, in the last years there have been a number of

promising developments in the field of global optimal control, which are reviewed next.

Dynamic Programming (DP) (Bellman, 1957), which proceeds by approximately solving the Hamilton–Jacobi–Bellman (HJB) equation (Frankowska, 1993), is a historically important method able to find globally optimal feedback laws. In Luss (1990) and Grüne and Semmler (2002) tailored discretization grids for DP implementations have been developed that can successfully solve nonlinear optimal control problems, as long as the dimension of the state of the system is small. For higher dimensional state-spaces these methods are, however, not applicable as the complexity of storing and processing the value functions during the DP recursion grows exponentially with the number of states. Other deterministic methods for global optimization problems involving nonlinear differential equations are based on Branch-and-Bound (BB) (Chachuat, Singer, & Barton, 2006; Esposito & Floudas, 2000; Lin & Stadtherr, 2007; Singer & Barton, 2006) and its variants, including α -BB methods (Diedam & Sager, 2017; Papamichail & Adjiman, 2002). These BB algorithms have in common that they can effectively solve problems with a small number of decision variables. Unfortunately, implementations of branch-and-bound search easily run out of memory in higher dimensional spaces due to exponentially growing search trees. In particular, as the discretization of control inputs typically leads to a large number of optimization variables, BB-methods are usually not suited for solving optimal control problems.

An alternative to Branch-and-Bound are the so-called Branch-and-Lift (BL) methods (Houska & Chachuat, 2014). In contrast to BB, these methods never discretize the control inputs directly.

[☆] MEV and BH acknowledge financial support via ShanghaiTech University. Grant F-0203-14-012. CNJ would like to acknowledge support received from the Swiss National Science Foundation under the RISK project (Risk Aware Data-Driven Demand Response), grant number 200021 175627. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Kok Lay Teo under the direction of Editor Ian R. Petersen.

* Corresponding author.

E-mail addresses: meduardov@shanghaitech.edu.cn (M.E. Villanueva), colin.jones@epfl.ch (C.N. Jones), borish@shanghaitech.edu.cn (B. Houska).

Instead, these methods branch over orthogonal projections of the control function in lower dimensional subspaces of increasing dimension, until an ϵ -suboptimal global control input is found. A rigorous analysis of the mathematical properties of such BL methods can be found in Houska and Chachuat (2019). However, in the context of developing practical implementations of BL one faces numerical challenges. In particular, these methods require the computation of accurate enclosures of moment-constrained reachable sets of nonlinear differential equations. At the current state of research, the lack of generic methods for computing such moment-constrained reachable set enclosures limits the applicability of BL. Although, currently, some problems have been solved successfully with BL by using enclosure propagation methods that are tailored to particular applications (Feng, Villanueva, Chachuat, & Houska, 2017; Houska & Chachuat, 2014).

In order to mitigate some of the limitations of existing global optimal control methods, this paper proposes a completely new framework for constructing methods for solving nonlinear infinite horizon optimal control problems. Here, the main idea is to analyze the spectrum of a Koopman operator that is associated with Pontryagin's differential equation—under the assumption that Pontryagin's optimality condition (Pontryagin, Boltyanskii, Gamkrelidze, & Mishchenko, 1962) is satisfied at the optimal solution. To understand the contributions of this paper, it is important to first be aware of the fact that Koopman operators have originally been introduced almost 90 years ago—in fact, by Koopman himself (Koopman, 1931). This theory has later been extended for more general nonlinear differential equations in Mezić (2005), where the concept of Koopman mode analysis has been introduced. A more detailed analysis of these Koopman mode-based methods from the field of ergodic theory has, however, only appeared much later (Arbabi & Mezić, 2017). Notice that a mathematical analysis of finite dimensional approximations of the Koopman operator can be found in Govindarajan, Mohr, Chandrasekaran, and Mezić (2019) and a variety of related applications of approximate Koopman mode analysis methods can be found in Budišić, Mohr, and Mezić (2012). Moreover, in Korda and Mezić (2018) Koopman mode estimation heuristics for data driven control have been introduced.

These approaches for lifting nonlinear differential equations into an infinite dimensional space using Koopman operators are closely related to similar reformulations based on Perron–Frobenius operators—which are, in fact, dual to the Koopman operator (Ding, 1998). This is relevant in the context of this paper as Perron–Frobenius operators can be used to reformulate nonlinear optimal control problems as linear programs (Vinter, 1993). For example, in Lasserre, Henrion, Prieur, and Trelat (2008) nonlinear optimal control problems are lifted into an infinite dimensional space using occupation measures. The corresponding method for global optimal control could be called a “first-lift-then-optimize” approach. First, the optimal control problem is “lifted” by reformulating it as an infinite dimensional linear programming problem. The resulting optimization problem is then solved in a second step using a generic LMI hierarchy (Lasserre et al., 2008). In contrast to this, the current paper proposes a “first-optimize-then-lift” approach, which formulates first the Pontryagin's necessary conditions for optimality and then lifts the corresponding condition into an infinite dimensional space. The related contributions are summarized below.

Contribution. The key idea of this paper is to introduce a Koopman operator that is associated with the Pontryagin differential equation of a nonlinear infinite-horizon optimal control problem—in the following called a *Pontryagin–Koopman operator*—which can be used to construct globally optimal feedback laws. Section 2 introduces infinite horizon optimal control problems and briefly reviews Pontryagin's necessary condition for

optimality. This review section is followed by a presentation of the main contributions of this article, which can be outlined as follows.

- (1) Section 3 uses ideas from symplectic geometry (Arnold & Novikov, 2001) to characterize the structural properties of Pontryagin–Koopman operators. Section 4 leverages on these structural properties in order to perform a symplectic spectral analysis that finally leads to a complete characterization of globally optimal feedback control laws (see Section 5, Theorem 3). Although such an exploitation of the global structure of Pontryagin's differential equation is only possible under relatively strong technical assumptions that will be introduced throughout this paper, it should be noted that this is the first time that the symplecticity of Pontryagin–Koopman operators is analyzed and used to derive practical characterizations of optimal control laws. In this sense, Theorem 3 can be regarded as the main theoretical contribution of this paper.
- (2) In order to avoid misunderstandings, it is mentioned clearly that this paper does not claim to solve all the numerical issues regarding the discretization of Pontryagin–Koopman operators that would have to be solved to develop a generic software for global optimal control. However, Section 6 illustrates the practical applicability of the proposed theoretical framework by designing an optimal regulator for a controlled Van der Pol oscillator. Although this implementation is only based on a naive Galerkin projection of the Pontryagin–Koopman operator onto a finite dimensional Legendre basis, the proposed framework is successfully applied to construct accurate approximations of nonlinear globally optimal feedback laws, for a system that cannot be stabilized by a linear control law.

The above contributions are relevant for the future of optimal control algorithm developments, as they pave the way towards the development of practical procedures for approximating globally optimal control laws of a very general class of nonlinear systems. Therefore, Section 7 does not only conclude the paper, but also outlines the potential and relevance of the proposed framework for future research in control.

Notation. The distance of a point $x \in \mathbb{R}^{n_x}$ to a trajectory $\phi : \mathbb{R} \rightarrow \mathbb{R}^n$ is denoted by

$$\text{dist}(x, \phi) = \inf_{t \in \mathbb{R}} \|x - \phi(t)\|.$$

We denote with \mathbb{L}_p^n the set of (potentially complex-valued) Lebesgue measurable functions $\varphi : \mathbb{R}^n \rightarrow \mathbb{C}$ whose p th power of the absolute value is integrable on \mathbb{R}^n . Moreover, we use the notation $\mathbb{W}_{k,p}^n$ to denote the Sobolev space of k -times weakly differentiable functions on \mathbb{R}^n , whose weak derivatives up to order k are all in \mathbb{L}_p^n . The symbol M^\top denotes the Hermitian transpose of the matrix $M \in \mathbb{C}^{n \times m}$. The symbol ∇ denotes the gradient operator. The associated second order derivative operator is denoted by $\nabla^2 = \nabla \nabla^\top$. Last but not least, the support of a function $\varphi : \mathbb{R}^n \rightarrow \mathbb{C}$ is denoted by

$$\text{supp}(\varphi) = \{x \in \mathbb{R}^n \mid \varphi(x) \neq 0\}.$$

2. Infinite horizon optimal control

This section introduces infinite horizon optimal control problems and briefly summarizes existing methods for analyzing the local stability properties of optimal periodic limit orbits (see Section 2.2). Moreover, Sections 2.3 and 2.4 review Pontryagin's necessary optimal condition thereby introducing the notation that is used throughout the paper.

2.1. Problem formulation

This paper considers infinite horizon optimal control problems of the form

$$V(x_0) = \min_{x,u} \int_0^\infty [l(x(t), u(t)) - l^*] dt$$

$$\text{s.t.} \begin{cases} \forall t \in [0, \infty), \\ \dot{x}(t) = f(x(t), u(t)) \\ x(0) = x_0. \end{cases} \quad (1)$$

Here, $x : \mathbb{R} \rightarrow \mathbb{R}^{n_x}$ denotes the state trajectory and $u : \mathbb{R} \rightarrow \mathbb{R}^{n_u}$ denotes the control input. The initial value $x_0 \in \mathbb{R}^{n_x}$ is assumed to be given. Throughout this paper the following assumptions are imposed.

Assumption 1. The function $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ is twice continuously differentiable and globally Lipschitz continuous in x .

Assumption 2. The function $l : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}$ is twice continuously differentiable.

The constant $l^* \in \mathbb{R}$ in (1) denotes the optimal average cost,

$$l^* = \lim_{T \rightarrow \infty} \min_{x,u} \frac{1}{T} \int_0^T l(x(t), u(t)) dt$$

$$\text{s.t.} \begin{cases} \forall t \in [0, T], \\ \dot{x}(t) = f(x(t), u(t)) \\ x(0) = x_0, \end{cases}$$

assuming that this limit exists. Notice that l^* does not need to be known explicitly in order to solve (1), as adding constant offsets to the objective does not affect the solution of an optimization problem. In this paper, the constant l^* is merely introduced for mathematical reasons, such that $V(x_0)$ remains finite and well-defined on infinite horizons under mild regularity assumptions that shall be introduced further below.

Remark 1 (Limitations). As much as the problem formulation (1) considers a rather general class of nonlinear systems and potentially non-convex stage costs, it should be mentioned that an explicit analysis of state- and control constraints is beyond the scope of this paper. Moreover, we assume that (1) is a time-autonomous optimal control problem, which means that f and l may depend on $x(t)$ and $u(t)$, but they are not allowed to explicitly depend on the time t . As we will see below, this assumption on the time-autonomy of (1) is crucial for the developments in this paper.

2.2. Limit behavior and periodic orbits

In practical instances, one is often interested in whether optimal solutions for the state trajectory of (1) converge to an optimal steady-state or, in more generality, to an optimal periodic limit orbit. These optimal steady states or more general periodic orbits are defined as follows.

Definition 1. The function $(x_p, u_p) \in \mathbb{W}_{1,1}^{n_x} \times \mathbb{L}_1^{n_u}$ is called an optimal periodic orbit, if there exists a period $T > 0$ such that for all $t \in [0, \infty)$

1. $x_p(t + T) = x_p(t)$ and $u_p(t + T) = u_p(t)$,
2. $\dot{x}_p(t) = f(x_p(t), u_p(t))$ for all $t \in [0, T]$, and
3. $\frac{1}{T} \int_0^T l(x_p(t), u_p(t)) dt = l^*$.

For the special case that x_p and u_p are constant functions, this pair is called an optimal steady-state.

In order to prepare the following analysis, it is helpful to introduce the shorthands

$$A(t) = f_x(x_p(t), u_p(t)), \quad B(t) = f_u(x_p(t), u_p(t)),$$

$$Q(t) = l_{xx}(x_p(t), u_p(t)), \quad R(t) = l_{uu}(x_p(t), u_p(t)),$$

$$S(t) = l_{xu}(x_p(t), u_p(t))$$

assuming that (x_p, u_p) is an optimal periodic orbit. Here, f_x and f_u denote the partial derivatives of f with respect to x and u and an analogous notation is then also used for the mixed second order derivatives of l . The above matrix-valued functions can be used to construct sufficient conditions under which an optimal periodic orbit is locally stabilizable. Here, one relies on the theory of periodic Riccati differential equations (Bittanti, Colaneri, & De Nicolao, 1991) of the form

$$-\dot{P}(t) = P(t)A(t) + A(t)^T P(t) + Q(t) - (P(t)B(t) + S(t))R(t)^{-1}(P(t)B(t) + S(t))^T$$

$$P(0) = P(T) > 0. \quad (2)$$

It is well-known (Bittanti et al., 1991) that if Assumption 1 and 2 are satisfied and if R has full rank, then the existence of a periodic and positive definite function P satisfying (2) is sufficient to ensure that the solution x^* of (1) converges to the optimal periodic orbit x_p as long as $\text{dist}(x_0, x_p)$ is sufficiently small—that is, if x_0 is in a small neighborhood of the optimal orbit x_p . However, this statement is of a rather local nature. This means that, if we wish to understand the global behavior of system (1), an analysis of the periodic Riccati equation is not sufficient. The following sections focus on analyzing global solutions of (1), under the assumption that $\text{dist}(x_0, x_p)$ is not necessarily small.

2.3. Pontryagin's differential equations

Pontryagin's maximum principle (Pontryagin et al., 1962) can be used to derive necessary conditions for the minimizers of (1). The first order variational optimality condition is summarized as follows. Let $H : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}$ be the Hamiltonian function of (1),

$$H(x, u, \lambda) = \lambda^T f(x, u) + l(x, u). \quad (3)$$

If Assumption 1 and 2 hold, H is, by construction, twice continuously differentiable in all variables. Next, let

$$u^*(x, \lambda) = \underset{u}{\text{argmin}} H(x, u, \lambda) \quad (4)$$

denote the associated parametric minimizer of H . At this point, we introduce the following regularity assumption.

Assumption 3. The parametric minimizer u^* in (4) satisfies the second order sufficient condition, $H_{uu} > 0$.

Notice that for the practically relevant special case that f is affine in u and l strongly convex in u , Assumption 3 always holds whenever Assumption 1 and 2 are satisfied.

Next, if Assumption 1, 2, and 3 are satisfied, it is well-known (Pontryagin et al., 1962) that any minimizer (x, u) of (1) necessarily satisfies Pontryagin's differential equation

$$\dot{x}(t) = f(x(t), u^*(x(t), \lambda(t))) \quad (5)$$

$$\dot{\lambda}(t) = -\nabla_x H(x(t), u^*(x(t), \lambda(t)), \lambda(t)) \quad (6)$$

for a co-state function $\lambda \in \mathbb{W}_1^{n_x}$. In the following, we summarize this differential equation in the form

$$\dot{y}(t) = F(y(t)), \quad (7)$$

where $y = [x^T, \lambda^T]^T$ denotes the stacked state and F a stacked version of the right-hand side functions of (5) and (6).

2.4. Necessary boundary and limit conditions

Besides Pontryagin’s differential equation (7), optimal solutions of (1) satisfy necessary boundary conditions. First, of course, the initial condition $x(0) = x_0$ must hold. Moreover, the co-state λ satisfies a necessary limit condition, which can be summarized as follows.

Proposition 1. *Let Assumptions 1, 2, and 3 hold. Moreover, let (x^*, u^*) be a primal optimal solution of (1) converging to an optimal periodic orbit $(x_p, u_p) \in \mathbb{W}_{1,1}^{n_x} \times \mathbb{X}_{-1}^{n_x}$ at which the periodic Riccati differential equation (2) admits a positive definite solution. Then, there exists a periodic function $\lambda_p \in \mathbb{W}_{1,1}^{n_x}$ such that the associated co-state λ^* necessarily satisfies*

$$\lim_{t \rightarrow \infty} \|\lambda^*(t) - \lambda_p(t)\| = 0.$$

Notice that Proposition 1 is—at least in very similar versions—well-known in the literature (Liberzon, 2012; Pontryagin et al., 1962). However, as this result is important for understanding the developments in this paper, the following proof briefly recalls the main argument supporting why this proposition holds.

Proof. Since (2) admits a positive definite solution, the value function V in (1) is well-defined and differentiable in a neighborhood of the optimal orbit x^* . Thus, (1) can be replaced by an equivalent finite-horizon optimal control problem as long as V is used as a terminal cost. Now, Pontryagin’s principle for optimal control problems with Mayer terms (Pontryagin et al., 1962) yields the boundary condition

$$\lambda^*(t) = \nabla V(x^*(t))$$

for any horizon length $t > 0$. Clearly, since x^* converges to the optimal periodic orbit x_p , λ^* must converge to the associated dual limit orbit $\lambda_p(t) = \nabla V(x_p(t))$. \square

In the next sections we will see that Proposition 1 implies that the states of Pontryagin’s differential equation for stabilizable systems must evolve along certain stable manifolds, which can be characterized using a Koopman mode analysis (Mezić, 2005).

3. Symplectic Koopman operators

This section analyzes the symplectic structure of the flow associated with Pontryagin’s differential equation. These symplectic structures are needed to understand the properties of Pontryagin Koopman operators. In fact, Section 3.2 introduces a symplectic test space—that is, a space of observables—in which the Pontryagin–Koopman operator inherits certain symplecticity properties of its underlying incompressible flow field. Moreover, Section 3.3 leverages on ideas from the field of symplectic geometry (Arnold & Novikov, 2001) in order to work out the symplectic dual of the Pontryagin–Koopman operator. Notice that these developments are the basis for those in Section 4, which use a symplectic spectral analysis in order to characterize globally optimal control laws.

3.1. Symplectic flows

Let $\Gamma_t : \mathbb{R}^{2n_x} \rightarrow \mathbb{R}^{2n_x}$ denote the flow of Pontryagin’s differential equations such that

$$\frac{d}{dt} \Gamma_t(z) = F(\Gamma_t(z)) \quad \text{and} \quad \Gamma_0(z) = z$$

for all $z \in \mathbb{R}^{2n_x}$. If Assumptions 1, 2, and 3 hold, then Γ_t is a well-defined, continuously differentiable function. In the mathematical literature (Arnold & Novikov, 2001) it is well-known that Γ_t is

a so-called symplectic flow. In order to reveal this symplectic structure it is helpful to introduce the block matrix

$$\Omega = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix} \in \mathbb{R}^{2n_x \times 2n_x}.$$

Now, the structural properties of Γ_t can be summarized as follows.

Lemma 1. *Let Assumptions 1, 2, and 3 be satisfied. Then, the function $\frac{\partial}{\partial z} \Gamma_t$ satisfies the equation*

$$\forall t \in \mathbb{R}, \quad \left[\frac{\partial}{\partial z} \Gamma_t \right]^\top \Omega \left[\frac{\partial}{\partial z} \Gamma_t \right] = \Omega, \tag{8}$$

that is, Γ_t is a symplectic function.

Proof. Let us first recall that the right-hand side F of Pontryagin’s differential equation is given by (5) and (6), which can also be summarized as

$$F = \begin{bmatrix} f \\ -H_x^\top \end{bmatrix} \quad \text{with} \quad y = \begin{bmatrix} x \\ \lambda \end{bmatrix}. \tag{9}$$

This implies that the derivative of F with respect to y can be written in the form

$$F_y = \begin{pmatrix} f_x - f_u H_{uu}^{-1} H_{ux} & -f_u H_{uu}^{-1} f_u^\top \\ -H_{xx} + H_{xu} H_{uu}^{-1} H_{ux} & -f_x^\top + H_{xu} H_{uu}^{-1} f_u^\top \end{pmatrix}. \tag{10}$$

Here, we have used Assumptions 1 and 2 to ensure that the second derivatives of H with respect to x and u exist. Moreover, we have used Assumption 3, which implies that H_{uu} is invertible and, as a consequence, that the implicit function theorem holds. In particular, we have

$$u_x^* = -H_{uu}^{-1} H_{ux} \quad \text{and} \quad u_\lambda^* = -H_{uu}^{-1} H_{u\lambda} = -H_{uu}^{-1} f_u^\top.$$

In this form, it becomes clear that the matrix ΩF_y is symmetric and we arrive at the intermediate result

$$\Omega F_y = F_y^\top \Omega^\top = -F_y^\top \Omega. \tag{11}$$

In order to proceed, we write the first order variational differential equation for Γ_t in the form

$$\frac{d}{dt} \left[\frac{\partial}{\partial z} \Gamma_t \right] = F_y \left[\frac{\partial}{\partial z} \Gamma_t \right] \quad \text{with} \quad \left[\frac{\partial}{\partial z} \Gamma_0 \right] = I. \tag{12}$$

Now, the main idea of this proof is to show that the function

$$\Delta(t) = \left[\frac{\partial}{\partial z} \Gamma_t \right]^\top \Omega \left[\frac{\partial}{\partial z} \Gamma_t \right] - \Omega \tag{13}$$

vanishes, that is $\Delta(t) = 0$. Here, we first have $\Delta(0) = 0$ by construction, since $\left[\frac{\partial}{\partial z} \Gamma_0 \right] = I$. Moreover, the derivative of Δ with respect to time is given by

$$\begin{aligned} \dot{\Delta}(t) &= \frac{d}{dt} \left[\frac{\partial}{\partial z} \Gamma_t \right]^\top \Omega \left[\frac{\partial}{\partial z} \Gamma_t \right] + \left[\frac{\partial}{\partial z} \Gamma_t \right]^\top \Omega \frac{d}{dt} \left[\frac{\partial}{\partial z} \Gamma_t \right] \\ &\stackrel{(12)}{=} \left[\frac{\partial}{\partial z} \Gamma_t \right]^\top \left[F_y^\top \Omega + \Omega F_y \right] \left[\frac{\partial}{\partial z} \Gamma_t \right] \stackrel{(11)}{=} 0 \end{aligned}$$

Thus, in summary, we must have $\Delta(t) = 0$ for all $t \in \mathbb{R}$, which yields the statement of this lemma. \square

The following corollary summarizes two immediate consequences of Lemma 1 which are both equivalent to stating that Γ_t is an incompressible flow.

Corollary 1. Let Assumptions 1, 2, and 3 be satisfied. Then, Γ_t is an incompressible flow; that is, we have

$$\forall t \in \mathbb{R}, \quad \det \left(\left[\frac{\partial}{\partial z} \Gamma_t \right] \right) = 1. \tag{14}$$

Moreover, the divergence of the associated vector field F vanishes, $\text{div}(F) = \nabla^\top F = 0$.

Proof. By taking the determinant on both sides of (8) and using that $\det(\Omega) = 1$, we find that

$$\det \left(\left[\frac{\partial}{\partial z} \Gamma_t \right] \right)^2 = 1 \iff \det \left(\left[\frac{\partial}{\partial z} \Gamma_t \right] \right) = \pm 1.$$

Since $\frac{\partial}{\partial z} \Gamma_0 = I$ and since $\frac{\partial}{\partial z} \Gamma_t$ is continuous, this is only possible if (14) holds. Next, by taking the logarithm on both sides of (14) and differentiating with respect to time, we find¹

$$\begin{aligned} 0 &= \frac{d}{dt} \log \left(\det \left(\left[\frac{\partial}{\partial z} \Gamma_t \right] \right) \right) \\ &= \text{Tr} \left(\left[\frac{\partial}{\partial z} \Gamma_t \right]^{-1} \frac{d}{dt} \left[\frac{\partial}{\partial z} \Gamma_t \right] \right) \stackrel{(12)}{=} \text{Tr}(F_y) = \text{div}(F). \end{aligned}$$

This completes the proof of the corollary. \square

3.2. Pontryagin-Koopman operators

The developments from the previous section can be used to analyze the structural properties of the Pontryagin-Koopman operator, which are defined to be the Koopman operator that is associated with the flow Γ_t of Pontryagin's differential equation, or, more formally:

Definition 2 (Pontryagin-Koopman Operator). We use the notation $U_t : \mathbb{W}_{1,2}^{2n_x} \rightarrow \mathbb{W}_{1,2}^{2n_x}$ to denote the Pontryagin-Koopman operator of Γ_t , which is defined such that

$$\forall \Phi \in \mathbb{W}_{1,2}^{2n_x}, \forall t \in \mathbb{R}, \quad U_t \Phi = \Phi \circ \Gamma_t$$

with \circ denoting the composition operator.

It is well known (Koopman, 1931; Mezić, 2005) that U_t is a linear operator satisfying $U_{t_1+t_2} = U_{t_1} U_{t_2}$ for all $t_1, t_2 \in \mathbb{R}$. Moreover, since the relation

$$U_0 \Phi = \Phi \circ \Gamma_0 = \Phi$$

holds for all $\Phi \in \mathbb{W}_{1,2}^{2n_x}$, the operator $U_0 = \text{id}$ corresponds to the identity operator. Next, if we substitute $t = t_1 = -t_2$ in the above equation, we find the relation

$$\text{id} = U_0 = U_t U_{-t}.$$

Consequently, the operator U_t is not only linear but also invertible and its inverse is given by $U_t^{-1} = U_{-t}$.

In order to introduce a notion of symplecticity in the space of observables of the Pontryagin-Koopman operator, we introduce the bilinear form

$$\omega : \mathbb{W}_{1,2}^{2n_x} \times \mathbb{W}_{1,2}^{2n_x} \rightarrow \mathbb{R},$$

which is defined as

$$\omega(\varphi, \Phi) = \int_{\mathbb{R}^{2n_x}} \nabla \varphi(z)^\top \Omega \nabla \Phi(z) dz$$

for all $\varphi, \Phi \in \mathbb{W}_{1,2}^{2n_x}$. Because we have $\Omega^\top = -\Omega$, the bilinear form ω is skew-symmetric, i.e.

$$\omega(\varphi, \Phi) = -\omega(\Phi, \varphi),$$

and $(\mathbb{W}_{1,2}^{2n_x}, \omega)$ is a symplectic space (Arnold & Novikov, 2001).

¹ Alternatively, the equation $\text{Tr}(F_y) = 0$ can also directly be found by substituting the explicit expression (10) from the proof of Lemma 1.

Theorem 1. Let Assumptions 1, 2, and 3 be satisfied. Now, U_t is for all $t \in \mathbb{R}$ a symplectic operator in the space $(\mathbb{W}_{1,2}^{2n_x}, \omega)$. This means that we have

$$\omega(U_t \varphi, U_t \Phi) = \omega(\varphi, \Phi)$$

for all $\varphi, \Phi \in \mathbb{W}_{1,2}^{2n_x}$.

Proof. Assumptions 1, 2, and 3 imply that the function Γ_t and its inverse $\Gamma_t^{-1} = \Gamma_{-t}$ are both continuously differentiable. This implies in particular that the equivalence

$$\Phi \in \mathbb{W}_{1,2}^{2n_x} \iff [\Phi \circ \Gamma_t] \in \mathbb{W}_{1,2}^{2n_x}$$

holds. Next, the definition of the operator U_t and the chain rule for differentiation imply that

$$\nabla[U_t \Phi] = \nabla[\Phi \circ \Gamma_t] = \left[\frac{\partial}{\partial z} \Gamma_t \right]^\top (\nabla \Phi) \circ \Gamma_t \tag{15}$$

holds for all $\Phi \in \mathbb{W}_{1,2}^{2n_x}$. Furthermore, because Assumptions 1, 2, and 3 hold, it follows from Lemma 1 that

$$\left[\frac{\partial}{\partial z} \Gamma_t \right]^\top \Omega \left[\frac{\partial}{\partial z} \Gamma_t \right] = \Omega.$$

Let us multiply this equation with $\left[\frac{\partial}{\partial z} \Gamma_t \right] \Omega$ from the left and substitute $\Omega^2 = -I$. This yields

$$\left[\frac{\partial}{\partial z} \Gamma_t \right] \Omega \left[\frac{\partial}{\partial z} \Gamma_t \right]^\top \Omega \left[\frac{\partial}{\partial z} \Gamma_t \right] = - \left[\frac{\partial}{\partial z} \Gamma_t \right].$$

Next, we multiply the latter equation with $\left[\frac{\partial}{\partial z} \Gamma_t \right]^{-1} \Omega$ from the right and use the relation $\Omega^2 = -I$ once more, in order to arrive at the equation

$$\left[\frac{\partial}{\partial z} \Gamma_t \right] \Omega \left[\frac{\partial}{\partial z} \Gamma_t \right]^\top = \Omega. \tag{16}$$

Thus, by substituting the previous equations, we get

$$\begin{aligned} \omega(U_t \varphi, U_t \Phi) &= \int_{\mathbb{R}^{2n_x}} \nabla[U_t \varphi](z)^\top \Omega \nabla[U_t \Phi](z) dz \\ &\stackrel{(15),(16)}{=} \int_{\mathbb{R}^{2n_x}} \nabla \varphi(\Gamma_t(z))^\top \Omega \nabla \Phi(\Gamma_t(z)) dz. \end{aligned} \tag{17}$$

In order to further transform this integral, we need to introduce the change of variables $z' = \Gamma_t(z)$. Because Corollary 1 ensures that

$$\left| \det \left(\frac{\partial}{\partial z} \Gamma_t(z) \right) \right| = 1,$$

this change of variables is volume preserving. This implies that

$$\omega(U_t \varphi, U_t \Phi) = \int_{\mathbb{R}^{2n_x}} \nabla \varphi(z')^\top \Omega \nabla \Phi(z') dz' = \omega(\varphi, \Phi)$$

for all $\varphi, \Phi \in \mathbb{W}_{1,2}^{2n_x}$. Thus, U_t is a symplectic operator. \square

3.3. Symplectic duality

In order to prepare the analysis of the properties of the spectrum of symplectic Koopman operators, it is helpful to introduce a notion of duality. However, instead of defining duality in a Hilbert space, we propose to introduce the following notion of symplectic duality.

Definition 3. Let $A : \mathbb{W}_{1,2}^{2n_x} \rightarrow \mathbb{W}_{1,2}^{2n_x}$ be a linear operator. If there exists a linear operator $A^* : \mathbb{W}_{1,2}^{2n_x} \rightarrow \mathbb{W}_{1,2}^{2n_x}$ with

$$\omega(\varphi, A\Phi) = \omega(A^* \varphi, \Phi)$$

for all $\varphi, \Phi \in \mathbb{W}_{1,2}^{2n_x}$, then we call A^* the symplectic adjoint operator of A .

In the following, we consider a linear differential operator $L : \mathbb{W}_{1,2}^{2n_x} \rightarrow \mathbb{W}_{1,2}^{2n_x}$ given by

$$L\Phi = F^\top \nabla \Phi \tag{18}$$

for all $\Phi \in \mathbb{W}_{1,2}^{2n_x}$. This operator is found by differentiating the definition of the Pontryagin–Koopman operator U_t with respect to time, such that $L = \dot{U}_t$.

Lemma 2. *Let Assumptions 1, 2, and 3 hold. The operator L admits a symplectic adjoint L^* in the symplectic space $(\mathbb{W}_{2,2}^{2n_x}, \omega)$. This adjoint is given by*

$$L^* = -L. \tag{19}$$

Proof. For any function $\Phi \in \mathbb{W}_{2,2}^{2n_x}$, we have

$$\nabla(L\Phi) = \frac{dF}{dy} \nabla \Phi + \nabla^2 \Phi F, \tag{20}$$

which follows by differentiating (18) on both sides. Next, we know from Theorem 1 that U_t is a symplectic operator. Thus, we can differentiate the equation

$$\omega(U_t \varphi, U_t \Phi) = \omega(\varphi, \Phi)$$

on both sides with respect to t , which yields

$$\frac{d}{dt} \omega(U_t \varphi, U_t \Phi) = 0.$$

Now, if $\varphi, \Phi \in \mathbb{W}_{2,2}^{2n_x}$, we can expand the derivative on the right as

$$\begin{aligned} 0 &= \frac{d}{dt} \omega(U_t \varphi, U_t \Phi) \\ &\stackrel{(17)}{=} \frac{d}{dt} \int_{\mathbb{R}^{2n_x}} \nabla \varphi(\Gamma_t(z))^\top \Omega \nabla \Phi(\Gamma_t(z)) dz \\ &= \int_{\mathbb{R}^{2n_x}} F(\Gamma_t(z))^\top \nabla^2 \varphi(\Gamma_t(z)) \Omega \nabla \Phi(\Gamma_t(z)) dz \\ &\quad + \int_{\mathbb{R}^{2n_x}} \nabla \varphi(\Gamma_t(z))^\top \Omega \nabla^2 \Phi(\Gamma_t(z)) F(\Gamma_t(z)) dz. \end{aligned}$$

Using the change of variables $z \leftarrow \Gamma_t(z)$ —recalling the condition $\det\left(\frac{\partial}{\partial z} \Gamma_t\right) = 1$ —and resorting terms leads to

$$\begin{aligned} &\int_{\mathbb{R}^{2n_x}} F(z)^\top \nabla^2 \varphi(z) \Omega \nabla \Phi(z) dz \\ &= - \int_{\mathbb{R}^{2n_x}} \nabla \varphi(z)^\top \Omega \nabla^2 \Phi(z) F(z) dz. \end{aligned} \tag{21}$$

Moreover, we recall that the equation

$$\Omega \frac{dF}{dy} = \left(\frac{dF}{dy}\right)^\top \Omega^\top = - \left(\frac{dF}{dy}\right)^\top \Omega \tag{22}$$

also holds (see Eq. (11)). Thus, in summary, we have

$$\begin{aligned} &\omega(\varphi, L\Phi) \\ &= \int_{\mathbb{R}^{2n_x}} [\nabla \varphi^\top \Omega \nabla [L\Phi]](z) dz \\ &\stackrel{(20)}{=} \int_{\mathbb{R}^{2n_x}} \left[\nabla \varphi^\top \Omega \left(\frac{dF}{dy} \nabla \Phi + \nabla^2 \Phi F \right) \right](z) dz \\ &\stackrel{(21), (22)}{=} - \int_{\mathbb{R}^{2n_x}} \left[\left(\frac{dF}{dy} \nabla \varphi + \nabla^2 \varphi F \right)^\top \Omega \nabla \Phi \right](z) dz \\ &= \omega(-L\varphi, \Phi). \end{aligned}$$

Because the latter equation holds for all $\varphi, \Phi \in \mathbb{W}_{2,2}^{2n_x}$, we find that $L^* = -L$ is the symplectic adjoint of L , as claimed by the statement of this lemma. \square

Remark 2. In many articles on Koopman operators, a so-called transport differential equation of the form

$$\frac{\partial}{\partial t} \phi = -\text{div}(F\phi)$$

is considered. If one introduces suitable boundary conditions, this advection PDE can be interpreted as the derivative of a Perron–Frobenius operator that is dual to the Koopman operator U_t with respect to the standard \mathbb{L}_2 -scalar product (Ding, 1998). Notice that in our case, the symplectic adjoint operator

$$L^* : \phi \rightarrow -F^\top \nabla \phi = -\text{div}(F\phi) + \underbrace{\text{div}(F)}_{=0} \phi$$

happens to coincide with the right-hand operator of the above advection PDE. Thus, physically, one could interpret the operator L^* as an advection operator of the incompressible flow field F recalling that $\text{div}(F) = 0$.

4. Spectral analysis

In this section, we show that the symplectic structures of the Koopman operator, as analyzed in the previous section, have important consequences on its set of eigenvalues. Moreover, Section 5 uses these spectral properties of the Koopman operator to characterize global optimal control laws that are associated with the infinite horizon optimal control problem (1).

4.1. Eigenfunctions and eigenvalues

As already mentioned in the introduction, the spectrum of general Koopman operators has been analyzed by many authors; for example in Arbabi and Mezić (2017), Budišić et al. (2012), Mezić (2005). In the context of this paper, we call a weakly differentiable function Ψ an eigenfunction of the differential Pontryagin–Koopman operator L , if the Lebesgue measure of the set $\text{supp}(\Psi)$ is either unbounded, or, if it exists, is not equal to 0, and

$$L\Psi = \kappa \Psi \iff \nabla \Psi^\top F = \kappa \Psi$$

for a potentially complex eigenvalue $\kappa \in \mathbb{C}$. The fact that the symplectic adjoint of the operator L is given by $L^* = -L$ has important consequences on its spectrum, which can be summarized as follows.

Theorem 2. *Let Assumptions 1, 2, and 3 hold and let $\Psi_1, \Psi_2 \in \mathbb{W}_{2,2}^{2n_x}$ be two eigenfunctions of L with eigenvalues $\kappa_1, \kappa_2 \in \mathbb{C}$. If $\omega(\Psi_1, \Psi_2) \neq 0$, then we must have $\kappa_1 = -\kappa_2$.*

Proof. Because the functions $\Psi_1, \Psi_2 \in \mathbb{W}_{2,2}^{2n_x}$ are eigenfunctions of L , they satisfy

$$L\Psi_1 = \kappa_1 \Psi_1 \quad \text{and} \quad L\Psi_2 = \kappa_2 \Psi_2. \tag{23}$$

Thus, since ω is a bilinear form, we have

$$\begin{aligned} \kappa_2 \omega(\Psi_1, \Psi_2) &\stackrel{(23)}{=} \omega(\Psi_1, L\Psi_2) \\ &\stackrel{(19)}{=} -\omega(L\Psi_1, \Psi_2) \\ &\stackrel{(23)}{=} -\kappa_1 \omega(\Psi_1, \Psi_2) \end{aligned}$$

and, after resorting terms,

$$(\kappa_1 + \kappa_2) \omega(\Psi_1, \Psi_2) = 0.$$

Thus, if $\omega(\Psi_1, \Psi_2) \neq 0$, we must have $\kappa_1 + \kappa_2 = 0$. This is equivalent to the statement of the theorem. \square

Remark 3. The statement of [Theorem 2](#) is formally not directly applicable for eigenfunctions Ψ_1, Ψ_2 of L that are locally weakly twice differentiable but whose derivatives are not square integrable, such that Ψ_1, Ψ_2 are not elements of the Sobolev space $\mathbb{W}_{2,2}^{2n_x}$. However, in practice, one is usually interested in constructing eigenfunctions of L on a compact domain $C \subset \mathbb{R}^{2n_x}$, in the following called the region of interest, such that

$$\forall x \in C, \quad [L\Psi](x) = \kappa\Psi(x). \quad (24)$$

This region of interest C can, for example, model a-priori bounds on the optimal primal and dual trajectories of (1). Consequently, because we are simply not interested in how the flow Γ_t and the associated eigenfunctions Ψ are defined outside of the set C , these functions can simply be redefined arbitrarily for $x \notin C$, for example, such that the desired eigenfunctions of L satisfy $\Psi \in \mathbb{W}_{2,2}^{2n_x}$ by construction. Thus, for the purpose of this paper, it is not restrictive at all to assume that the derivatives of the eigenfunctions of L are square integrable. Notice that this technique is also illustrated by our tutorial case study in [Section 6](#), where we explain how to choose C and how to discretize (24) on C .

Let $\sigma(L) \subseteq \mathbb{C}$ denote the spectrum of L ; that is, the set of eigenvalues of the linear operator L . In the following, we use the notation $\Psi_\kappa \in \mathbb{W}_{2,2}^{2n_x}$ to denote an eigenfunction that is associated with an eigenvalue $\kappa \in \sigma(L)$. Moreover, for a function $q : \mathbb{R}^{2n_x} \rightarrow \mathbb{R}^{2n_x}$, we use the shorthand notation

$$\omega(q, q) = \begin{pmatrix} \omega(q_1, q_1) & \omega(q_1, q_2) & \dots & \omega(q_1, q_{2n_x}) \\ \omega(q_2, q_1) & \omega(q_2, q_2) & \dots & \omega(q_2, q_{2n_x}) \\ \vdots & \vdots & \ddots & \vdots \\ \omega(q_{2n_x}, q_1) & \omega(q_{2n_x}, q_2) & \dots & \omega(q_{2n_x}, q_{2n_x}) \end{pmatrix}$$

to denote the matrix that is obtained by evaluating the skew symmetric bilinear form ω for all possible combinations of the components of q . We call q a skew-orthogonal function in the symplectic space $(\mathbb{W}_{1,2}^{2n_x}, \omega)$, if it satisfies $\omega(q, q) = \Omega$. Notice that the construction of such skew orthogonal functions is straightforward by using the standard skew-symmetric variant of the Gram-Schmidt algorithm ([Arnold & Novikov, 2001](#)).

Definition 4. The operator L is said to admit a spectral decomposition with respect to a given function q , if there exists a generalized function $a : \sigma(L) \rightarrow \mathbb{C}^{2n_x}$ such that

$$q = \int_{\sigma(L)} \Psi_\kappa a(\kappa) d\kappa. \quad (25)$$

Notice that the above definition uses the ‘‘control engineering notation’’ for generalized functions, which means that we use a as if it was a standard function, although this notation suppresses the distributional nature of a . Thus, in mathematical terms, this notation has to be translated as ‘‘ a represents a distribution of order 0; that is, a linear operator on $\sigma(L)$ that is Lipschitz continuous with respect to the L_∞ -norm’’.² The following statement is a consequence of [Theorem 2](#).

Corollary 2. Let [Assumptions 1, 2, and 3](#) hold. Let the operator L admit a spectral decomposition with respect to a skew-orthogonal function q . Then, there exist at least $2n_x$ eigenvalues

$$\kappa_1^+, \kappa_2^+, \dots, \kappa_{n_x}^+, \kappa_1^-, \kappa_2^-, \dots, \kappa_{n_x}^- \in \sigma(L),$$

such that $\kappa_i^+ = -\kappa_i^-$, where κ_i^+ has a non-negative real part for all $i \in \{1, 2, \dots, n_x\}$.

² Notice that, for example, a could be a Dirac distribution, which is not a function in the traditional sense but, by construction, a Lipschitz continuous linear operator.

Proof. By substituting (25) in the equation $\omega(q, q) = \Omega$, we find that the equation

$$\Omega = \int_{\sigma(L)} \int_{\sigma(L)} \underbrace{\omega(\Psi_{\kappa'}, \Psi_\kappa) a(\kappa') a(\kappa)^\top}_{\text{rank 1}} d\kappa d\kappa', \quad (26)$$

holds. Let us have a closer look at the terms in (26). Clearly, the matrix Ω on the left has full rank. But, on the other side, we have an integral over the rank 1 matrices $a(\kappa') a(\kappa)^\top$. This integral term can only have full rank, if there are at least $2n_x$ pairs of eigenvalues $(\kappa', \kappa) \in \sigma(L) \times \sigma(L)$ for which

$$\omega(\Psi_{\kappa'}, \Psi_\kappa) \neq 0.$$

But now [Theorem 2](#) implies that all these pairs must be such that $\kappa' = -\kappa$ and, after sorting all eigenvalues with respect to their real-part, we find that there must be at least n_x eigenvalues with non-negative real part and n_x associated mirrored eigenvalues with non-positive real-part, as claimed by the statement of this corollary. \square

Remark 4. Notice that [Corollary 2](#) makes a statement about the spectrum of L under the rather strong assumption that this operator admits a spectral decomposition with respect to at least one skew orthogonal function. Thus, one should further ask the question under which assumptions the existence of such a spectral decomposition can be guaranteed for at least one such skew orthogonal function. In full generality, this question is difficult to answer, but sufficient conditions for the existence of (much more general and, in our context, sufficient) spectral decompositions can be found in [Arbabi and Mezić \(2017\)](#), [Budišić et al. \(2012\)](#), [Mauroy and Mezić \(2016\)](#), [Mezić \(2005\)](#), which use ideas from the field of ergodic theory ([Wiener & Wintner, 1941](#)) as well as Yoshida’s theorem ([Yosida, 1978](#)). For example, if the monodromy matrix, $\frac{\partial}{\partial z} \Gamma_T(x_p(0))$, of a periodic optimal orbit with period length $T > 0$ is diagonalizable, one can ensure that a spectral decomposition is possible. This sufficient condition follows by applying [Proposition 3](#) in [Mauroy and Mezić \(2016\)](#) to the Pontryagin differential equation. A more complete review of such results from the field of functional analysis would, however, go beyond the scope of the present paper.

5. Optimal feedback control laws

The theoretical results from the previous sections can be used to derive optimality conditions, which, in turn, can be used to develop practical numerical algorithms for solving (1). Let us introduce the set

$$\sigma^+(L) = \{\kappa \in \sigma(L) \mid \text{Re}(\kappa) > 0\}$$

of unstable eigenvalues and its associated invariant manifold

$$\mathcal{M}^+ = \left\{ [x^\top, \lambda^\top]^\top \in \mathbb{R}^{2n_x} \mid \begin{array}{l} \forall \kappa \in \sigma^+(L), \\ \Psi_\kappa(x, \lambda) = 0 \end{array} \right\}.$$

Since we assume that the eigenfunctions Ψ_κ are weakly differentiable, we may assume without loss of generality that the functions Ψ_κ are also continuous—otherwise there exists a continuous function $\tilde{\Psi}_\kappa(x) = \Psi_\kappa(x)$ for almost every $x \in \mathbb{R}^{2n_x}$ and we can use $\tilde{\Psi}_\kappa$ instead of Ψ_κ recalling that we work with Sobolev spaces, in which such arguments are indeed possible.

In the following, we say that μ^* is a regular optimal control law of (1), if the closed-loop trajectories,

$$\dot{x}^*(t) = f(x^*(t), \mu^*(x^*(t))) \quad \text{with} \quad x^*(0) = x_0,$$

are minimizers of (1) at which Pontryagin’s necessary optimality conditions are satisfied such that x^* converges to an optimal periodic orbit at which the conditions of [Proposition 1](#) are satisfied.

Theorem 3. Let Assumptions 1, 2, and 3 hold. We further assume (without loss of generality) that the eigenfunctions $\Psi_\kappa \in W_{2,2}^{2n_x}$ are continuous. Let $\mu^* : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_u}$ be a regular optimal control law of (1). Then, there exists a function $\Lambda : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_x}$ such that

$$[x^\top, \Lambda(x)^\top]^\top \in \mathcal{M}^+ \quad \text{and} \quad \mu^*(x) = u^*(x, \Lambda(x)).$$

Proof. Since L is a time-autonomous infinitesimal generator of U_t , the eigenfunctions of L also satisfy (Mezić, 2005)

$$U_t \Psi_\kappa = e^{\kappa t} \Psi_\kappa \quad (27)$$

for all $\kappa \in \sigma(L)$. The remainder of the proof is divided into two parts. In the first part, we show that (27) implies that the optimal periodic limit orbit is a subset of the manifold \mathcal{M}^+ . In the second part of the proof, we use this property of the limit orbit to construct a globally optimal feedback law.

Part I: Let (x_p, λ_p) denote the optimal periodic limit orbit of (1) with period time T . If we would have $(x_p(t), \lambda_p(t)) \notin \mathcal{M}^+$ for a time $t \in [0, T]$, we must have

$$\Psi_\kappa(x_p(t), \lambda_p(t)) \neq 0$$

for at least one $\kappa \in \sigma^+(L)$. Since the optimal periodic orbit satisfies Pontryagin's differential equation, this implies in particular that

$$\begin{aligned} |\Psi_\kappa(x_p(t+T), \lambda_p(t+T))| &\stackrel{(27)}{=} |e^{\kappa T} \Psi_\kappa(x_p(t), \lambda_p(t))| \\ &> |\Psi_\kappa(x_p(t), \lambda_p(t))| \end{aligned} \quad (28)$$

as κ has a strictly positive real part. But this is a contradiction, since (x_p, λ_p) is periodic. Thus, in summary, we must have

$$\forall t \in [0, T], \quad (x_p(t), \lambda_p(t)) \in \mathcal{M}^+. \quad (29)$$

Part II: Let (x, λ) denote an optimal solution of (1) with $x(0) = x_0$. Now, if we would have $(x(0), \lambda(0)) \notin \mathcal{M}^+$, then there would have to exist at least one $\kappa \in \sigma^+(L)$ for which

$$\Psi_\kappa(x(0), \lambda(0)) \neq 0.$$

But if this would be the case, then we would also have

$$\lim_{t \rightarrow \infty} |\Psi_\kappa(x(t), \lambda(t))| = \infty, \quad (30)$$

as Ψ_κ is strictly unstable. But this limit statement is in conflict with (29), since Ψ_κ is continuous and we assume that $(x(t), \lambda(t))$ converges to the optimal periodic limit orbit. Thus, in summary, there exists for every x_0 a $\lambda_0 = \lambda(0)$ with $(x_0, \lambda_0) \in \mathcal{M}^+$ and the corresponding map from x_0 to λ_0 can be denoted by Λ . The associated optimal control input is given by

$$\mu^*(x_0) = u^*(x_0, \Lambda(x_0)).$$

This is already sufficient to establish the statement of this theorem, as the optimal feedback law must be time-autonomous. \square

Notice that Theorem 3 can be used to systematically search for globally optimal solutions of (1). Here, one of the key observations is that if, in addition to the assumptions of Theorem 3, the assumptions of Corollary 2 are also satisfied and if none of the non-negative eigenvalues $\kappa_1^+, \kappa_2^+, \dots, \kappa_{n_x}^+$ happens to be on the imaginary axis, then the parametric equation system

$$\forall \kappa \in \sigma^+(L), \quad \Psi_\kappa(x_0, \lambda_0) = 0. \quad (31)$$

consists of at least n_x independent equations while the number of variables, $\dim(\lambda) = n_x$, is equal to n_x , too. Thus, this equation system can, in many practical instances, be expected to admit a finite number of parametric solutions $\lambda_0 = \Lambda(x_0)$ only. This means that, if all the above regularity assumptions are satisfied, Theorem 3 singles out a finite number of candidate control

Input:

- The right-hand side function $F = [f^\top, -H_x]^\top$ of Pontryagin's differential equation and the parametric minimizer u^* of the Hamiltonian function, as defined in (4).

Algorithm:

- (1) Choose a sufficiently large $N \geq 2n_x$ and orthogonal basis functions $\varphi_1, \dots, \varphi_N \in \mathbb{W}_{1,2}^{2n_x}$.

- (2) Compute the coefficients of the matrix $M \in \mathbb{R}^{N \times N}$,

$$M_{i,j} = \langle L\varphi_i, \varphi_j \rangle = \langle F^\top \nabla \varphi_i, \varphi_j \rangle,$$

for all $i, j \in \{1, \dots, N\}$.

- (3) Compute the n_x eigenvalues $\kappa_1, \dots, \kappa_{n_x} \in \mathbb{C}$ of M with largest real parts and their associated left eigenvectors $v_1, \dots, v_{n_x} \in \mathbb{C}^N$, such that the equation $M^\top v_i = \kappa_i v_i$ holds for all $i \in \{1, \dots, n_x\}$.

- (4) Compute all parametric solutions $\lambda_0 = \tilde{\Lambda}(x_0)$ of the equation system

$$\forall i \in \{1, \dots, n_x\}, \quad \sum_{j=1}^N v_{i,j} \varphi_j(x_0, \lambda_0) = 0$$

and compute the associated candidate control laws

$$\tilde{\mu}(x_0) = u^*(x_0, \tilde{\Lambda}(x_0)).$$

- (5) If there is more than one candidate control law $\tilde{\mu}$, simulate the associated closed-loop systems to pick the one that leads to minimum cost.

Output:

- The (approximately) optimal control law $\tilde{\mu} \approx \mu^*$.

Fig. 1. Optimal control via Koopman lifts.

laws $\mu^*(x) = u^*(x, \Lambda(x))$ —at least one of which must be globally optimal.

Notice that a practical implementation of the above procedure requires one to discretize the linear differential operator L . As already mentioned in the introduction, a complete overview and analysis of possible numerical procedures, which could lead to such a practical implementation is beyond the scope of this paper. Nevertheless, Fig. 1 outlines a prototype algorithm based on a Galerkin projection (Ayuso & Marini, 2009) of L .

Remark 5. As already mentioned in Remark 1, the proposed Pontryagin–Koopman operator based framework is only applicable to time-autonomous optimal control problems, because (27) does not hold if F depends on t . This is a general limitation of Koopman-operator mode decomposition based methods (Mezić, 2005).

Remark 6. Notice that many existing local and global optimization methods focus on solving finite horizon optimal control problems (Biegler, 2007; Chachuat et al., 2006; Diehl et al., 2002;

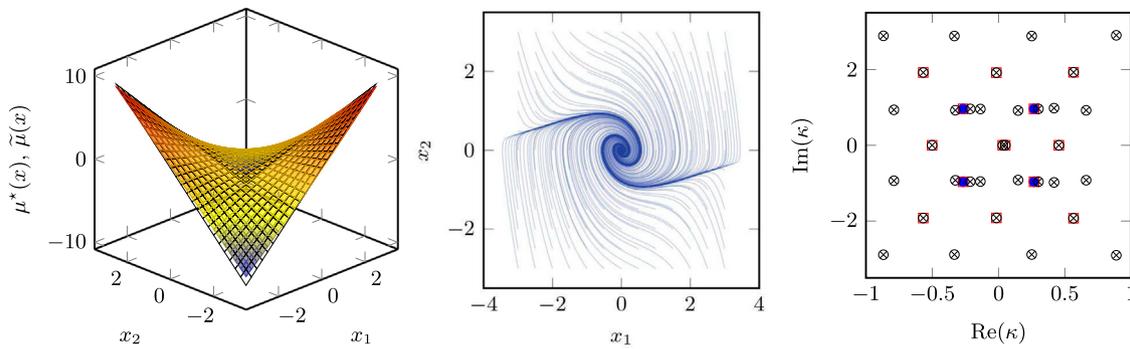


Fig. 2. Left: Comparison between the optimal feedback law $\mu^*(x) = -x_1x_2$, and the approximate feedback law $\tilde{\mu}(x)$ in (36). Middle: trajectories of the closed-loop system under the approximate feedback law. Right: Eigenvalues of the projection $M \in \mathbb{R}^{N \times N}$ for $N = 4$ (blue circles), $N = 15$ (red squares), and $N = 35$ (black circles with a cross).

Papamichail & Adjiman, 2002). This is in contrast to the proposed Koopman–Pontryagin operator based algorithm focusing on time-autonomous infinite horizon optimal control problems. The main reason for this is that the control law that generates the optimal trajectories of a finite horizon optimal control problem is in general time-varying. This means that a time-autonomous characterization of such optimal feedback laws – as established in Theorem 3 for infinite horizon optimal control problems – is in this form not possible for general optimal control problems with finite time horizons.

6. Numerical example

This section illustrates how Theorem 3 can be used to construct accurate approximations of globally optimal control laws of (1) for a Van der Pol oscillator system with

$$f(x, u) = \begin{pmatrix} x_2 \\ x_1 - \frac{1}{2}(1 - x_1^2)x_2 + x_1u \end{pmatrix} \quad (32)$$

$$\text{and } l(x, u) = \frac{1}{2}(x_2^2 + u^2). \quad (33)$$

The associated system, $\dot{x}(t) = f(x(t), u(t))$, has a linearly uncontrollable equilibrium at $x = (0, 0)^T$. Notice that for this particular example an explicit solution of the Hamilton–Jacobi–Bellman equation is known (Rodrigues, Henrion, & Cantwell, 2016): it turns out that the optimal value function V and the globally optimal feedback law μ^* are given, respectively, by

$$V(x) = \frac{1}{2}(x_1^2 + x_2^2) \quad \text{and} \quad \mu^*(x) = -x_1x_2. \quad (34)$$

In the following, this explicit solution is, however, only used to assess the accuracy of the proposed method; that is, the numerical procedure below neither knows the above expression for V nor for μ^* .

Remark 7. The function f in (32) is not globally Lipschitz continuous and, consequently, Assumption 1 is violated. However, as further elaborated below, we are interested, at least in this example, in approximating the control law within a compact region of interest. Consequently, we can simply replace the function f by a globally Lipschitz continuous function that coincides with f within this region of interest—without effecting the following numerical results.

6.1. Galerkin discretization

In order to illustrate the practical applicability of the developments in Section 5, we introduce a simple Galerkin discretization of the operator L . Let $\varphi_1, \dots, \varphi_N \in \mathbb{W}_{1,2}^{2n_x}$ be orthogonal functions

with respect to the standard \mathbb{L}_2 -scalar product $\langle \cdot, \cdot \rangle$ on $\mathbb{W}_{1,2}^{2n_x}$. Next, if we compute the coefficients

$$M_{i,j} = \langle L\varphi_i, \varphi_j \rangle$$

for all $i, j \in \{1, \dots, N\}$, the matrix M^T can be interpreted as a Galerkin approximation of the operator L over the subspace spanned by $\varphi_1, \dots, \varphi_N$. In our implementation, we set φ_i to the i th multivariate Legendre polynomial on the 4-dimensional compact interval box $C = [-\frac{1}{2}, \frac{1}{2}]^4$ and we set $\varphi_i(x, \lambda) = 0$ outside of this domain; that is, for $(x, \lambda) \notin C$. Consequently, our discretization can only be expected to be accurate inside our region of interest C , but other choices for C and for the basis functions would be possible, too.

Remark 8. Standard Galerkin methods are, in general, numerically unstable when applied to advection operators. Consequently, although this method happens to yield reasonable approximations for our particular example, such naive discretization schemes cannot be recommended in general (Roos, Stynes, & Tobiska, 2008). More advanced discretization schemes for linear advection operators can be found in the modern PDE literature; see, for example (Ayuso & Marini, 2009; Brenner & Scott, 2005; Roos et al., 2008), where one can also find an in-depth discussion of the numerical accuracy and convergence properties of such methods. A more complete discussion of such numerical discretization methods for Pontryagin–Koopman operators, is, however, beyond the scope of this paper.

6.2. Approximations of the optimal feedback law

The above Galerkin approximation of the operator L can be used to construct approximate eigenfunctions. In detail, if $a \in \mathbb{R}^N$ is a left eigenvector of M with eigenvalue $\kappa \in \mathbb{C}$, then

$$\Psi = \sum_{i=1}^N a_i \varphi_i \implies L\Psi \approx \kappa\Psi$$

is an approximation of an eigenfunction of L . The right plot in Fig. 2 shows the spectrum of the matrix M for different choices of N (blue circles: $N = 4$, red squares: $N = 15$, and black circles with a cross: $N = 35$). In order to understand the structure of this spectrum it is helpful to recall Corollary 2, which predicts that there exists at least 2 eigenvalues κ_1, κ_2 of L such that $-\kappa_1$ and $-\kappa_2$ are also eigenvalues of L . Notice that such symmetric eigenvalue pairs are indeed present in the spectrum of M , although M is only a Galerkin approximation of L .

In order to further illustrate how the above spectral analysis of M can be used to construct approximations of the globally optimal control law μ^* , we can compute an approximation of the

manifold \mathcal{M}^+ by using the approximate eigenfunctions instead of the exact ones (see Theorem 3). For example, for $N = 4$, the matrix M has the eigenvalues

$$\kappa_{1,2}^+ = \frac{1}{48}(\beta_m \pm \beta_p \sqrt{-1})$$

$$\text{with } \beta_p = \sqrt{+983 + 96\sqrt{143}}$$

$$\text{and } \beta_m = \sqrt{-983 + 96\sqrt{143}} \quad (35)$$

and the associated Galerkin approximation of the globally optimal control law is given by

$$\tilde{\mu}(x) = (12 - \sqrt{143})x_1^2 + \frac{(11 - \beta_m)}{2}x_1x_2. \quad (36)$$

The left plot in Fig. 2 shows $\tilde{\mu}$ and compares it to the optimal feedback law μ^* . In fact, the squared integral error over C is approximately 6×10^{-5} . Quite remarkably, this approximate feedback law can even be used to control the system for initial values outside of C . The plot in the middle of Fig. 2 shows the corresponding trajectories of the closed-loop system that are obtained by using the approximately optimal feedback law $\tilde{\mu}$.

7. Conclusions

This paper has presented an analysis of infinite horizon nonlinear optimal control problems, whose minimizers satisfy Pontryagin's necessary conditions of optimality. The proposed formalism is based on Pontryagin–Koopman operators, which have been shown to possess a symplectic structure, as revealed by Theorem 1. Moreover, Theorem 2 and Corollary 2 have established conditions under which the spectrum of the differential Pontryagin–Koopman operator contains at least $2n_x$ mirrored eigenvalues. This spectral structure is used in Theorem 3 to characterize optimal control laws.

The theoretical findings of this paper have been applied to construct accurate approximations of a globally optimal control law for a Van der Pol oscillator, which illustrates the potential of the proposed Pontryagin–Koopman operator based framework for the design of global optimal control algorithms. Here, it needs to be highlighted that, in contrast to dynamic programming methods, which rely on the discretization of nonlinear Hamilton–Jacobi–Bellman PDEs, the proposed framework for global optimal control relies on the computation of eigenfunctions of a linear differential operator. This opens the door to an application of linear algebra methods and tailored discretization schemes for linear PDEs, which have never been considered for the computation of optimal control laws. Therefore, the development of tailored, structure-exploiting projection and linear algebra methods for symplectic Pontryagin–Koopman operators and their application to global optimal control can be regarded as a promising direction for future research.

References

- Arbabi, H., & Mezić, I. (2017). Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the Koopman operator. *SIAM Journal on Applied Dynamical Systems*, 16(4), 2096–2126.
- Arnold, V. I., & Novikov, S. P. (2001). *Dynamical systems IV, Symplectic geometry*. Springer.
- Ayuso, B., & Marini, L. D. (2009). Discontinuous Galerkin methods for advection-diffusion-reaction problems. *SIAM Journal on Numerical Analysis*, 47, 1391–1420.
- Bellman, R. E. (1957). *Dynamic programming*. Princeton: Princeton University Press.
- Biegler, L. T. (2007). An overview of simultaneous strategies for dynamic optimization. *Chemical Engineering and Processing*, 46, 1043–1053.
- Bittanti, S., Colaneri, P., & De Nicolao, G. (1991). The periodic Riccati equation. In *The Riccati equation*. Springer.

- Brenner, S., & Scott, R. L. (2005). *The mathematical theory of finite element methods*. Springer.
- Budišić, M., Mohr, R., & Mezić, I. (2012). Applied koopmanism. *Chaos*, 22(4), Article 047510.
- Chachuat, B., Singer, A. B., & Barton, P. I. (2006). Global methods for dynamic optimization and mixed-integer dynamic optimization. *Industrial and Engineering Chemistry Research*, 45(25), 8373–8392.
- Diedam, H., & Sager, S. (2017). Global optimal control with the direct multiple shooting method. *Optimal Control Applications & Methods*, <http://dx.doi.org/10.1002/oca.2324>.
- Diehl, M., Bock, H. G., Schlöder, J. P., Findeisen, R., Nagy, Z., & Allgöwer, F. (2002). Real-time optimization and nonlinear model predictive control of processes governed by differential-algebraic equations. *Journal of Process Control*, 12(4), 577–585.
- Ding, J. (1998). The point spectrum of Frobenius–Perron and Koopman operators. *Proceedings of the American Mathematical Society*, 126(5).
- Esposito, W. R., & Floudas, C. A. (2000). Deterministic global optimization in nonlinear optimal control problems. *Journal of Global Optimization*, 17, 96–126.
- Feng, X., Villanueva, M. E., Chachuat, B., & Houska, B. (2017). Branch-and-lift algorithm for obstacle avoidance control. In *Proceedings of the 56th IEEE conference on decision and control* (pp. 745–750).
- Frankowska, H. (1993). Lower semicontinuous solutions of Hamilton–Jacobi–Bellman equations. *SIAM Journal on Control and Optimization*, 31(1), 257–272.
- Govindarajan, N., Mohr, R., Chandrasekaran, S., & Mezić, I. (2019). On the approximation of koopman spectra for measure preserving transformations. *SIAM Journal on Applied Dynamical Systems*, 18(3), 1454–1497.
- Grüne, L., & Semmler, W. (2002). Using dynamic programming with adaptive grid scheme to solve nonlinear dynamic models in economics. In *Computing in economics and finance 2002 no. 99*.
- Houska, B., & Chachuat, B. (2014). Branch-and-lift algorithm for deterministic global optimization in nonlinear optimal control. *Journal of Optimization Theory and Applications*, 162(1), 208–248.
- Houska, B., & Chachuat, B. (2019). Global optimization in Hilbert space. *Mathematical Programming, Series A*, 173, 221–249.
- Houska, B., Ferreau, H. J., & Diehl, M. (2011). An auto-generated real-time iteration algorithm for nonlinear MPC in the microsecond range. *Automatica*, 47(10), 2279–2285.
- Koopman, B. O. (1931). Hamiltonian systems and transformations in Hilbert space. *Proceedings of the National Academy of Sciences of the United States of America*, 17(5), 315–318.
- Korda, M., & Mezić, I. (2018). Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, 93, 149–160.
- Lasserre, J. B., Henrion, D., Prieur, C., & Trelat, E. (2008). Nonlinear optimal control via occupation measures and LMI relaxations. *SIAM Journal on Control and Optimization*, 47, 1649–1666.
- Liberzon, D. (2012). *Calculus of variations and optimal control theory: A concise introduction*. Princeton University Press.
- Lin, Y., & Stadtherr, M. A. (2007). Deterministic global optimization of nonlinear dynamic systems. *AIChE Journal*, 53(4), 866–875.
- Luss, R. (1990). Optimal control by dynamic programming using systematic reduction in grid size. *International Journal of Control*, 51(5), 995–1013.
- Mattingley, J., & Boyd, S. (2009). *Automatic code generation for real-time convex optimization*. In *Convex optimization in signal processing and communications*. Cambridge University Press.
- Mauroy, A., & Mezić, I. (2016). Global stability analysis using the eigenfunctions of the Koopman operator. *IEEE Transactions on Automatic Control*, 61(11), 3356–3369.
- Mezić, I. (2005). Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1–3), 309–325.
- Papamichail, I., & Adjiman, C. S. (2002). A rigorous global optimization algorithm for problems with ordinary differential equations. *Journal of Global Optimization*, 24, 1–33.
- Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., & Mishchenko, E. F. (1962). *The mathematical theory of optimal processes*. New York: John Wiley & Sons.
- Qin, S. J., & Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7), 733–764.
- Rodrigues, L., Henrion, D., & Cantwell, B. J. (2016). Symmetries and analytical solutions of the Hamilton–Jacobi–Bellman equation for a class of optimal control problems. *Optimal Control Applications & Methods*, 37(4), 749–764.
- Roos, H.-G., Stynes, M., & Tobiska, L. (2008). *Robust numerical methods for singularly perturbed differential equations: Convection-diffusion-reaction and flow problems (vol. 24)*. Springer Science & Business Media.
- Singer, A. B., & Barton, P. I. (2006). Global optimization with nonlinear ordinary differential equations. *Journal of Global Optimization*, 34(2), 159–190.

- Vinter, R. (1993). Convex duality and nonlinear optimal control. *SIAM Journal on Control and Optimization*, 31(2), 518–538.
- Wiener, N., & Wintner, A. (1941). Harmonic analysis and ergodic theory. *American Journal of Mathematics*, 63(2), 415–426.
- Yosida, K. (1978). *Functional analysis*. Berlin: Springer.
- Zavala, V. M., & Biegler, L. T. (2009). The advanced step NMPC controller: Optimality, stability and robustness. *Automatica*, 45, 86–93.



Mario Eduardo Villanueva is a postdoctoral researcher at the School of Information Science and Technology at ShanghaiTech University. He received a master and Ph.D. in chemical engineering from Imperial College London in 2011 and 2016, respectively. He was a postdoctoral researcher at Texas A&M University in 2016. Mario Villanueva is the recipient of the 2016 Dudley Newitt Price and the 2018 as well as 2019 SIST Excellent Postdoc Award. His research interests include set based computing, robust and stochastic control, global optimization and system theoretic tools for data

analysis.



Colin N. Jones received the bachelor's and master's degrees in electrical engineering and mathematics from the University of British Columbia, Vancouver, Canada and the Ph.D. degree in 2005 from the University of Cambridge, U.K, for his work on polyhedral computational methods for constrained control. Since 2017, he has been an Associate Professor with the Automatic Control Laboratory, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland. He was a Senior Researcher with the Automatic Control Lab, ETH Zurich, Switzerland, until 2010. He has authored or coauthored

more than 200 publications and was awarded an ERC starting grant to study the optimal control of building networks. His current research interests include high-speed predictive control and optimization, and the control of green energy generation, distribution and management.



Boris Houska is an assistant professor at the School of Information Science and Technology at ShanghaiTech University. He received a diploma in mathematics from the University of Heidelberg in 2007, and a Ph.D. in Electrical Engineering from KU Leuven in 2011. From 2012 to 2013 he was a postdoctoral researcher at the Centre for Process Systems Engineering at Imperial College London. Boris Houska's research interests include numerical optimization and optimal control, robust and global optimization, as well as fast model predictive control algorithms.