



HAL
open science

Symmetry reduction for dynamic programming

John Maidens, Axel Barrau, Silvère Bonnabel, Murat Arcak

► **To cite this version:**

John Maidens, Axel Barrau, Silvère Bonnabel, Murat Arcak. Symmetry reduction for dynamic programming. *Automatica*, 2018, 10.1016/j.automatica.2018.08.024 . hal-01982793

HAL Id: hal-01982793

<https://minesparis-psl.hal.science/hal-01982793>

Submitted on 16 Jan 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Symmetry reduction for dynamic programming [★]

John Maidens ^a, Axel Barrau ^b, Silvère Bonnabel ^c, and Murat Arcaç ^a

^a*Department of Electrical Engineering & Computer Sciences, University of California, Berkeley, 253 Cory Hall, Berkeley, CA, 94720 USA*

^b*SAFRAN TECH, Groupe Safran, Rue des Jeunes Bois - Châteaufort, 78772 Magny Les Hameaux, France*

^c*MINES ParisTech, PSL Research University, Centre for Robotics, 60 bd Saint-Michel, 75006 Paris, France*

Abstract

We present a method of exploiting symmetries of discrete-time optimal control problems to reduce the dimensionality of dynamic programming iterations. The results are derived for systems with continuous state variables, and can be applied to systems with continuous or discrete symmetry groups. We prove that symmetries of the state update equation and stage costs induce corresponding symmetries of the optimal cost function and the optimal policies. We then provide a general framework for computing the optimal cost function based on gridding a space of lower dimension than the original state space. This method does not rely on explicitly transforming the state update equations and can therefore be applied in circumstances where this is difficult. We illustrate these results on two six-dimensional optimal control problems that are computationally difficult to solve by dynamic programming without symmetry reduction.

Key words: Optimal control; Dynamic programming; Invariant systems; Model reduction; Path planning; Medical applications

1 Introduction

The dynamic programming algorithm for computing optimal control policies has, since its development, been known to suffer from the “curse of dimensionality” (Bellman, 1957). Its applicability in practice is typically limited to systems with four or five continuous state variables because the number of points required to grid a space of n continuous state variables increases exponentially with the state dimension n . This complexity has led to a collection of algorithms for approximate dynamic programming, which scale to systems with larger state dimension but lack the guarantees of global optimality of the solution associated with the original dynamic programming algorithm (Bellman and Dreyfus, 1959; Bertsekas, 2012; Powell, 2007, 2016).

In practice, many real-world systems exhibit symmetries that can be exploited to reduce the complexity of

system models. Symmetry reduction has found applications in fields ranging from differential equations (Clarkson and Mansfield, 1994; Bluman and Kumei, 2013) to model checking (Emerson and Sistla, 1996; Kwiatkowska et al., 2006). In control engineering, symmetries have been exploited to improve control of mechanical systems (Bloch et al., 1996; Bullo and Murray, 1999), develop more reliable state estimators (Barrau and Bonnabel, 2014), study the controllability of multiagent systems (Rahmani et al., 2009) and to reduce the complexity of stability and performance certification for interconnected systems (Arcak et al., 2016; Rufino Ferreira et al., 2017). Symmetry reduction has also been applied to the computation of optimal control policies for continuous-time systems in (Grizzle and Marcus, 1984; Ohsawa, 2013) and Markov decision processes (MDPs) in (Zinkevich and Balch, 2001; Narayanamurthy and Ravindran, 2007).

In this paper, we present a theory of symmetry reduction for the optimal control of discrete-time, stochastic nonlinear systems with continuous state variables. This reduction allows dynamic programming to be performed in a lower-dimensional state space. Since the computational complexity of a dynamic programming iteration increases exponentially with state dimension, this reduction significantly decreases computational burden. Fur-

[★] A preliminary version of this work was presented at the 2017 American Control Conference. Corresponding author J. Maidens.

Email addresses: maidens@eecs.berkeley.edu (John Maidens), axel.barrau@mines-paristech.fr (Axel Barrau), silvere.bonnabel@mines-paristech.fr (Silvère Bonnabel), arcac@eecs.berkeley.edu (Murat Arcaç).

ther, our proposed method does not rely on an explicit transformation of the state update equations, making the method applicable in situations where a such a transformation is difficult or impossible to find analytically.

We present two theorems that summarize our method of symmetry reduction. Theorem 4 describes how symmetries of the system dynamics imply symmetries of the optimal cost and optimal policy functions. Theorem 6 then describes a method of computing the cost function based on reduced coordinate system that depends on fewer state variables.

This paper builds on the work we presented in the conference paper (Maidens et al., 2017). The most substantial improvement is the additional theoretical results presented in Sections 2.3 and 3.2. The conference version presented an *ad hoc* symmetry reduction for a magnetic resonance imaging (MRI) application, but did not provide a general methodology for computing the coordinate reduction. This paper addresses this shortcoming by presenting a general method based on the moving frame formalism, which leads to the general symmetry reduction formula presented in Theorem 6. Additionally, the MRI example has been reworked to match this new formalism, and the numerical implementation and graphs of the numerical solution have been improved. We have also included two new extensions of this formalism to the case of equivariant costs in Section 3.3 and to the synchronization problem of stochastic dynamic systems on matrix groups in Section 5, along with examples to illustrate the algorithm in these contexts.

This paper is organized as follows: in Section 2 we introduce notation and provide background information both on dynamic programming for optimal control, and on the mathematical theory of symmetries. In Section 3, we derive our main theoretical results, that is, we prove that control system symmetries induce symmetries of the optimal cost function and optimal control policy, and then leverage the result to present a general method of performing dynamic programming in reduced coordinates. In Section 4 we apply symmetry reduction to compute the solution of an optimal control problem arising in dynamic MRI acquisition. In Section 5 we apply the algorithm to the problem of coordination on Lie groups. Simulations are performed on a cooperative control problem for two Dubins vehicles. Code to reproduce the computational results in this paper is available at <https://github.com/maidens/Automatica-2017>.

2 Dynamic Programming and Symmetries

In this section, we first recall the main features of dynamic programming for optimal control of stochastic discrete time systems. Then we introduce our problem and provide the reader with a primer on the classical theory

of symmetries. We also introduce the notion of invariant control systems with invariant costs.

2.1 Dynamic programming for optimal control of stochastic systems

We consider a discrete-time dynamical system

$$x_{k+1} = f_k(x_k, u_k, w_k), \quad k = 0, 1, \dots, N-1 \quad (1)$$

where $x_k \in \mathcal{X} \subseteq \mathbb{R}^n$ is the system state, $u_k \in \mathcal{U} \subseteq \mathbb{R}^m$ is the control variable to be chosen at time k , $w_k \in \mathcal{W} \subseteq \mathbb{R}^\ell$ are independent continuous random variables each with density p_k , and $N \in \mathbb{Z}_+$ is a finite control horizon. Associated with this system is an additive cost function

$$g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, u_k, w_k)$$

that we wish to minimize through our choice of u_k . We define a *control system* to be a tuple $\mathcal{S} = (\mathcal{X}, \mathcal{U}, \mathcal{W}, p, f, g, N)$ where $p = \prod_{k=0}^{N-1} p_k$ is the joint density of the random variables w_k .

We consider a class of control policies $\pi = \{\mu_0, \dots, \mu_{N-1}\}$ where $\mu_k : \mathcal{X} \rightarrow \mathcal{U}$ maps observed states to admissible control inputs. Given an initial state x_0 and a control policy π , we define the expected cost under this policy as

$$J_\pi(x_0) = \mathbb{E} \left[g_N(x_N) + \sum_{k=0}^{N-1} g_k(x_k, \mu_k(x_k), w_k) \right].$$

An optimal policy π^* is defined as one that minimizes the expected cost:

$$J_{\pi^*}(x_0) = \min_{\pi \in \Pi} J_\pi(x_0)$$

where Π denotes the set of all admissible control policies. The optimal cost function, denoted $J^*(x_0)$, is defined to be the expected cost corresponding to an optimal policy.

An optimal policy π^* and the optimal cost function J^* can be computed using the dynamic programming algorithm. We quote the following result due to Bellman from (Bertsekas, 2005):

Proposition 1 (Dynamic Programming) *For every initial state x_0 , the optimal cost $J^*(x_0)$ of the basic problem is equal to $J_0(x_0)$, given by the last step of the following algorithm, which proceeds backward in time from period $N-1$ to period 0:*

$$\begin{aligned} J_N(x_N) &= g_N(x_N) \\ J_k(x_k) &= \min_{u_k \in \mathcal{U}} \mathbb{E} \left[g_k(x_k, u_k, w_k) + J_{k+1}(f_k(x_k, u_k, w_k)) \right] \quad (2) \\ & \quad k = 0, 1, \dots, N-1, \end{aligned}$$

where the expectation is taken with respect to the probability distribution of w_k . Furthermore, if $u_k^* = \mu_k^*(x_k)$ minimizes the right hand side of (2) for each x_k and k , then the policy $\pi^* = \{\mu_0^*, \dots, \mu_{N-1}^*\}$ is optimal.

2.2 Invariant system with invariant costs

We first recall the definition of a transformation group for a control system, as in (Martin et al., 2004; Jakubczyk, 1998; Respondek and Tall, 2002). See Olver (1999) for the more general theory.

Definition 2 (Transformation group) A transformation group on $\mathcal{X} \times \mathcal{U} \times \mathcal{W}$ is set of tuples $h_\alpha = (\phi_\alpha, \chi_\alpha, \psi_\alpha)$ parametrized by elements α of a Lie group \mathcal{G} having dimension r , such that the functions $\phi_\alpha : \mathcal{X} \rightarrow \mathcal{X}$, $\chi_\alpha : \mathcal{U} \rightarrow \mathcal{U}$ and $\psi_\alpha : \mathcal{W} \rightarrow \mathcal{W}$ are all C^1 diffeomorphisms and satisfy:

- $\phi_e(x) = x$, $\chi_e(u) = u$, $\psi_e(w) = w$ when e is the identity of the group \mathcal{G} and
- $\phi_{a*b}(x) = \phi_a \circ \phi_b(x)$, $\chi_{a*b}(u) = \chi_a \circ \chi_b(u)$, $\psi_{a*b}(x) = \psi_a \circ \psi_b(x)$ for all $a, b \in \mathcal{G}$ where $*$ denotes the group operation and \circ denotes function composition.

To simplify notation we will sometimes suppress the subscripts α . In the present paper, we will consider the following class of systems and cost functions.

Definition 3 (Invariant control system with invariant costs) A control system \mathcal{S} is \mathcal{G} -invariant with \mathcal{G} -invariant costs if for all $\alpha \in \mathcal{G}$, $x_k \in \mathcal{X}$, $u_k \in \mathcal{U}$ and $w_k \in \mathcal{W}$ we have:

$$\begin{aligned} \phi^{-1} \circ f_k(\phi(x_k), \chi(u_k), \psi(w_k)) &= f_k(x_k, u_k, w_k), \\ &k = 0, 1, \dots, N-1 \\ g_k(\phi(x_k), \chi(u_k), \psi(w_k)) &= g_k(x_k, u_k, w_k), \\ &k = 0, 1, \dots, N-1, \\ g_N(\phi(x_N)) &= g_N(x_N), \text{ and} \\ p_k(\psi(w_k)) |\det D\psi(w_k)| &= p_k(w_k) \\ &k = 0, 1, \dots, N-1. \end{aligned}$$

The rationale is simple: For any fixed $\alpha \in \mathcal{G}$, consider the change of variables $X_k = \phi_\alpha(x_k)$, $U_k = \chi_\alpha(u_k)$, $W_k = \psi_\alpha(w_k)$. Then, we have

$$X_{k+1} = f_k(X_k, U_k, W_k), \quad k = 0, 1, \dots, N-1,$$

and for $k = 0, 1, \dots, N-1$ we have also $g_k(X_k, U_k, W_k) = g_k(x_k, u_k, w_k)$. As a result, if u_1, \dots, u_{N-1} is a series of controls that minimize $J(x_0)$, then one can expect U_1, \dots, U_{N-1} to minimize $J(X_0)$, under some assumptions on the noise. As a result, the optimal control problem needs only be solved once for all initial conditions belonging to the set $\{\phi_\alpha(x_0) | \alpha \in \mathcal{G}\}$, reducing the

initial n dimensional problem to a $n-r$ dimensional problem. The present paper derives a proper theory for such symmetry reduction in dynamic programming, and provides various examples of engineering interest.

2.3 Cartan's moving frame method

To find a reduced coordinate system in which to perform dynamic programming, we will use the moving frame method of Cartan (Cartan, 1937). In general, this method only results in a local coordinate transformation as it relies on the implicit function theorem. However, for many practical problems the transformation computed using this method extends globally. To simplify the exposure we will present a ‘‘global’’ version of the method.

We briefly introduce the moving frame method following the presentation in (Bonnabel et al., 2008). Consider an r -dimensional transformation group (with $r \leq n$) acting on \mathcal{X} via the diffeomorphisms $(\phi_\alpha)_{\alpha \in \mathcal{G}}$. Assume we can split ϕ_α as $(\phi_\alpha^a, \phi_\alpha^b)$ with r and $n-r$ components respectively so that ϕ_α^a is an invertible map. Then, for some c in the range of ϕ^a , we define a coordinate cross section to the orbits $\mathcal{C} = \{x : \phi_\alpha^a(x) = c\}$. This cross section is an $n-r$ -dimensional submanifold of \mathcal{X} . Assume moreover that for any point $x \in \mathcal{X}$, there is a unique group element $\alpha \in \mathcal{G}$ such that $\phi_\alpha(x) \in \mathcal{C}$. Such α will be denoted $\gamma(x)$, and the map $\gamma : \mathcal{X} \rightarrow \mathcal{G}$ will be called moving frame.

A moving frame can be computed by solving the normalization equations:

$$\phi_{\gamma(x)}^a(x) = c.$$

Define the following map $\rho : \mathcal{X} \rightarrow \mathbb{R}_{n-r}$ as

$$\rho(x) = \phi_{\gamma(x)}^b(x).$$

Note that, for all $\alpha \in \mathcal{G}$ we have $\rho(\phi_\alpha(x)) = \rho(x)$, that is, the components of ρ are *invariant* to the group action on the state space. Further, due to our assumptions, the restriction of ρ to \mathcal{C} is injective. We denote this restricted function $\bar{\rho}$, and it will serve as a reduced coordinate system to solve the invariant optimal control problem.

3 Main Results

In order to combat the ‘‘curse of dimensionality’’ associated with performing dynamic programming in high-dimensional systems, we describe a method to reduce the system’s dimension by exploiting symmetries in the dynamics and stage costs.

3.1 Symmetries imply equivalence classes of optimal policies

Theorem 4 (Symmetries of the optimal cost and policy) *Let \mathcal{G} be a group and let \mathcal{S} be a \mathcal{G} -invariant control system with \mathcal{G} -invariant costs. Then the optimal cost functions $J_k(x_0)$ satisfy the symmetry relations*

$$J_k = J_k \circ \phi_\alpha$$

for any $k = 0, \dots, N$ and any $\alpha \in \mathcal{G}$. Furthermore, if $\pi^* = \{\mu_0^*, \dots, \mu_{N-1}^*\}$ is an optimal policy then so is $\tilde{\pi}^* := \{\chi_\alpha \circ \mu_0^* \circ \phi_\alpha^{-1}, \dots, \chi_\alpha \circ \mu_{N-1}^* \circ \phi_\alpha^{-1}\}$ for any $\alpha \in \mathcal{G}$.

PROOF. First, note that

$$J_N(x_N) = g_N(x_N) = g_N(\phi(x_N)) = J_N(\phi(x_N)).$$

Now, suppose that for some $k \in \{0, \dots, N-1\}$ we have $J_{k+1}(x_{k+1}) = J_{k+1}(\phi(x_{k+1}))$ for all $x_{k+1} \in \mathcal{X}$. Then for any $x_k \in X$, and $u_k \in \mathcal{U}$ we have

$$\begin{aligned} & \mathbb{E} \left[g_k(x_k, u_k, w_k) + J_{k+1}(f_k(x_k, u_k, w_k)) \right] \\ &= \int_{\mathcal{W}} \left[g_k(x_k, u_k, w_k) + J_{k+1}(f_k(x_k, u_k, w_k)) \right] p_k(w_k) dw_k \\ &= \int_{\mathcal{W}} \left[g_k(\phi(x_k), \chi(u_k), \psi(w_k)) + J_{k+1}(\phi^{-1} \circ f_k(\phi(x_k), \chi(u_k), \psi(w_k))) \right] p_k(w_k) dw_k \\ &= \int_{\mathcal{W}} \left[g_k(\phi(x_k), \chi(u_k), \psi(w_k)) + J_{k+1}(f_k(\phi(x_k), \chi(u_k), \psi(w_k))) \right] p_k(w_k) dw_k \\ &= \int_{\mathcal{W}} \left[g_k(\phi(x_k), \chi(u_k), \psi(w_k)) + J_{k+1}(f_k(\phi(x_k), \chi(u_k), \psi(w_k))) \right] p_k(\psi(w_k)) |\det D\psi(x_k)| dw_k \\ &= \int_{\psi(\mathcal{W})} \left[g_k(\phi(x_k), \chi(u_k), \tilde{w}_k) + J_{k+1}(f_k(\phi(x_k), \chi(u_k), \tilde{w}_k)) \right] p_k(\tilde{w}_k) d\tilde{w}_k \\ &= \int_{\mathcal{W}} \left[g_k(\phi(x_k), \chi(u_k), w_k) + J_{k+1}(f_k(\phi(x_k), \chi(u_k), w_k)) \right] p_k(w_k) dw_k \\ &= \mathbb{E} \left[g_k(\phi(x_k), \chi(u_k), w_k) + J_{k+1}(f_k(\phi(x_k), \chi(u_k), w_k)) \right] \end{aligned}$$

Therefore,

$$\begin{aligned} J_k(x_k) &= \min_{u_k \in \mathcal{U}} \mathbb{E} \left[g_k(x_k, u_k, w_k) + J_{k+1}(f_k(x_k, u_k, w_k)) \right] \\ &= \min_{u_k \in \mathcal{U}} \mathbb{E} \left[g_k(\phi(x_k), \chi(u_k), w_k) + J_{k+1}(f_k(\phi(x_k), \chi(u_k), w_k)) \right] \\ &= \min_{\tilde{u}_k \in \chi(\mathcal{U})} \mathbb{E} \left[g_k(\phi(x_k), \tilde{u}_k, w_k) + J_{k+1}(f_k(\phi(x_k), \tilde{u}_k, w_k)) \right] \\ &= J_k(\phi(x_k)). \end{aligned}$$

Thus $J^* = J^* \circ \phi$. Now, if $\pi^* = \{\mu_0^*, \dots, \mu_{N-1}^*\}$ is an optimal policy and we denote $\tilde{x}_k = \phi(x_k)$ then for any

$k \in \{0, \dots, N-1\}$ we have

$$\begin{aligned} J_k(\tilde{x}_k) &= J_k(x_k) \\ &= \mathbb{E} \left[g(x_k, \mu_k^*(x_k), w_k) + J_{k+1}(f_k(x_k, \mu_k^*(x_k), w_k)) \right] \\ &= \mathbb{E} \left[g_k(\phi(x_k), \chi(\mu_k^*(x_k)), w_k) + J_{k+1}(f_k(\phi(x_k), \chi(\mu_k^*(x_k)), w_k)) \right] \\ &= \mathbb{E} \left[g_k(\phi(x_k), \chi \circ \mu_k^* \circ \phi^{-1}(\phi(x_k)), w_k) + J_{k+1}(f_k(\phi(x_k), \chi \circ \mu_k^* \circ \phi^{-1}(\phi(x_k)), w_k)) \right] \\ &= \mathbb{E} \left[g_k(\tilde{x}_k, \chi \circ \mu_k^* \circ \phi^{-1}(\tilde{x}_k), w_k) + J_{k+1}(f_k(\tilde{x}_k, \chi \circ \mu_k^* \circ \phi^{-1}(\tilde{x}_k), w_k)) \right] \end{aligned}$$

Thus $\tilde{\pi}^* := \{\chi \circ \mu_0^* \circ \phi^{-1}, \dots, \chi \circ \mu_{N-1}^* \circ \phi^{-1}\}$ is an optimal policy. \square

3.2 Dynamic programming can be performed using reduced coordinates

Theorem 4 readily implies the problem can be reduced, as all states along an orbit of \mathcal{G} are equivalent in terms of cost, and that there are equivalence classes of optimal policies. So it suffices to only consider the cost corresponding to a single representative of each equivalence class, and to find a single representative of the optimal policy within each class. This can now easily be done using the injective map $\bar{\rho} : \mathcal{C} \rightarrow \mathbb{R}^{n-r}$.

For $\bar{x} \in \bar{\rho}(\mathcal{C}) \subset \mathbb{R}^{n-r}$, let $z \in \mathcal{C}$ be such that $\bar{x} = \bar{\rho}(z)$, and define

$$\bar{J}_k(\bar{x}) = J_k(z).$$

The following result shows that the functions J_k on the n -dimensional space $\mathcal{X} \subseteq \mathbb{R}^n$ are completely determined by the values of \bar{J}_k on the subset $\bar{\rho}(\mathcal{C})$ of \mathbb{R}^{n-r} .

Corollary 5 *For any $x \in \mathcal{X}$ and $k = 0, \dots, N$, the cost function J_k for the full problem can be computed in terms of the lower-dimensional cost function \bar{J}_k as*

$$J_k(x) = J_k(\phi_{\gamma(x)}(x)) = \bar{J}_k(\bar{x}),$$

where $\bar{x} := \bar{\rho}(\phi_{\gamma(x)}(x))$ is well defined as $\phi_{\gamma(x)}(x) \in \mathcal{C}$.

It is thus sufficient to have evaluated \bar{J} at all points of $\bar{\rho}(\mathcal{C}) \subset \mathbb{R}^{n-r}$ to be able to instantly evaluate J at any point of \mathcal{X} .

Theorem 6 (Dynamic programming in reduced coordinates) *The reduced coordinates are in one to one correspondance with the cross-section \mathcal{C} . For any \bar{x} , let $z \in \mathcal{C}$ satisfy $\bar{\rho}(z) = \bar{x}$. Then in the reduced coordinates, the sequence \bar{J}_k can be computed recursively via*

$$\bar{J}_k(\bar{x}) = \min_{u_k \in \mathcal{U}} \mathbb{E} \left[g_k(z, u_k, w_k) + \bar{J}_{k+1}(\bar{\rho}(f_k(z, u_k, w_k))) \right].$$

PROOF. We have

$$\begin{aligned}\bar{J}_k(\bar{x}) &= J_k(z) = \min_{u_k \in \mathcal{U}} \mathbb{E} \left[g_k(z, u_k, w_k) + J_{k+1}(f_k(z, u_k, w_k)) \right] \\ &= \min_{u_k \in \mathcal{U}} \mathbb{E} \left[g_k(z, u_k, w_k) + \bar{J}_{k+1}(\rho \circ f_k(z, u_k, w_k)) \right]. \quad \square\end{aligned}$$

3.3 Case of equivariant costs

So far, we have considered the costs to be invariant. However, equivariance can obviously also be considered. Indeed, the cost g_k is said to be equivariant if there exists a family of diffeomorphisms φ_α such that $g_k(\phi_\alpha(x_k), \chi_\alpha(u_k), \psi_\alpha(w_k)) = \varphi_\alpha \circ g_k(x_k, u_k, w_k)$. As we want the cost function J to be equivariant too, we will need $\varphi_\alpha(\cdot)$ to be linear. Thus we will simply assume that φ_α is of the form $\varphi_\alpha(J) = l(\alpha)J$, that is, it is a scaling of the cost, where $l: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$. For simplicity's sake, we consider here the problem to be noise free. Along the lines of the preceding sections it is easily proved that

$$J_k(\phi(x)) = \varphi \circ J_k(x)$$

as already noticed in (Alvarez and Stokey, 1998) for the case of homogeneous costs. Symmetry reduction can then be applied. We now give two tutorial examples.

Example 1 Consider the linear system

$$x_{k+1} = Ax_k + Bu_k$$

with quadratic costs $g_k = x_k^T Q x_k + u_k^T R u_k$. The system is invariant to scalings, $\phi_\alpha(x) = \alpha x$, $\chi_\alpha(u) = \alpha u$, and the cost is equivariant letting $\varphi_\alpha(J) = \alpha J$, where $\alpha \in \mathcal{G} = \mathbb{R}_{>0}$. The unit sphere is a cross section to the orbits, and the normalization equations yield $\gamma(x) = 1/\|x\|$. Applying the results above, we see that the controls that minimize $J(x_0)$ are $\|x_0\|u_1^*, \dots, \|x_0\|u_{N-1}^*$, where u_1^*, \dots, u_{N-1}^* are those minimizing $J(\frac{x_0}{\|x_0\|})$. This agrees with the well known fact that the optimal controller for the problem above is the linear quadratic controller, and is indeed of the linear form $u_k = -F_k x_k$.

Example 2 Consider the following system and costs

$$x_{k+1} = Ax_k + Bu_k, \quad g_k = h(x_k) + \|u_k\|_1$$

where $\|u_k\|_1$ denotes the L^1 norm of u_k and h is a map satisfying $h(ax) = ah(x)$ for $a > 0$. Such costs may arise when one tries to force some controls to zero to create sparsity, a method known as L^1 regularization. This problem is challenging, particularly for nonconvex h . But according to the theory above, it is sufficient to solve it numerically for initial conditions lying on the unit sphere of the state space.

4 Application to MRI Fingerprinting

Magnetic resonance imaging (MRI) has traditionally focused on acquisition sequences that are static, in the sense that sequences typically wait for magnetization to return to equilibrium between acquisitions. Recently, researchers have demonstrated promising results based on dynamic acquisition sequences, in which spins are continuously excited by a sequence of random input pulses, without allowing the system to return to equilibrium between pulses. Model parameters corresponding to T_1 and T_2 relaxation, off-resonance and spin density are then estimated from the sequence of acquired data. This technique, termed magnetic resonance fingerprinting (MRF), has been shown to increase the sensitivity, specificity and speed of magnetic resonance studies (Ma et al., 2013; Davies et al., 2014).

This technique could be further improved by replacing randomized input pulse sequences with sequences that have been optimized for informativeness about model parameters. To this end, we present a model of MR spin dynamics that describes the measured data as a function of T_1 and T_2 relaxation rates and the sequence of radio-frequency (RF) input pulses, used to excite the spins.

The following model was introduced in the conference paper (Maidens et al., 2017). In this paper, an optimal control was computed via dynamic programming on a very sparse six-dimensional grid. Now using our symmetry reduction technique, we exploit symmetry reduction to provide a much more accurate optimal input sequence computed on a finer five-dimensional grid.

We model the spin dynamics via the equations

$$\mathbf{x}_{k+1} = U_k \begin{bmatrix} \theta_2 & 0 & 0 \\ 0 & \theta_2 & 0 \\ 0 & 0 & \theta_1 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} 0 \\ 0 \\ 1 - \theta_1 \end{bmatrix} \quad (3)$$

where the states $x_{1,k}$ and $x_{2,k}$ describe the transverse magnetization (orthogonal to the applied magnetic field) and $x_{3,k}$ describes the longitudinal magnetization (parallel to the applied magnetic field). To simplify the presentation, off-resonance is neglected in this model. Control inputs $U_k \in \text{SO}(3)$ describe flip angles corresponding to RF excitation pulses that rotate the state about the origin. Between acquisitions, transverse magnetization decays according to the parameter $\theta_2 = e^{-\Delta t/T_2}$ and the longitudinal magnetization recovers to equilibrium (normalized such that the equilibrium is $x_0 = [0 \ 0 \ 1]^T$) according to the parameter $\theta_1 = e^{-\Delta t/T_1}$ where Δt is the sampling interval.

We assume that data are acquired immediately following the RF pulse, allowing us to make a noisy measurement of the transverse magnetization. We also assume that the

$$\begin{aligned}
f_k(x_k, U_k, w_k) &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ U_k & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & U_k \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \theta_2 & 0 & 0 & 0 & 0 & 0 \\ 0 & \theta_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & \theta_1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \theta_2 & 0 & 0 \\ 0 & 0 & 0 & 0 & \theta_2 & 0 \\ 0 & 0 & 1 & 0 & 0 & \theta_1 \end{bmatrix} x_k + \begin{bmatrix} 0 \\ 0 \\ 1 - \theta_1 \\ 0 \\ 0 \\ -1 \end{bmatrix} \\
g_k(x_k, U_k, w_k) &= -x_k^T \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{\gamma} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{\gamma} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} x_k
\end{aligned} \tag{4}$$

measured data are described by a multivariate Gaussian random variable

$$\mathbf{y}_k = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \mathbf{x}_k + v_k$$

where v_k is a zero-mean Gaussian noise with covariance $\begin{bmatrix} \gamma & 0 \\ 0 & \gamma \end{bmatrix}$. This model results from a time discretization of the Bloch equations (Bloch, 1946; Nishimura, 2010) under a time scale separation assumption that specifies that the RF excitation pulses act on a much faster time scale than the relaxation time constants T_1 and T_2 . A simplified two-state version of this model was considered in (Maidens et al., 2016b), where the transverse magnetization was modelled using a single state describing the magnitude of $[x_{1,k}, x_{2,k}]^T$.

We see from the model (3) that magnetization in the transverse direction decays while magnetization in the longitudinal direction grows. However only the transverse component of the magnetization can be measured. Thus there is a trade-off between making measurements (which leads to loss of magnetization) and magnetization recovery. This is the trade-off that we hope to manage through the optimal design of an input sequence U_k .

We wish to quantify the informativeness of an acquisition sequence based on the information about the T_1 relaxation parameter θ_1 that is contained in the resulting data set. More formally, we wish to choose $U_k \in \text{SO}(3)$ to maximize the Fisher information about θ_1 contained in the joint distribution of $Y = (\mathbf{y}_0, \dots, \mathbf{y}_N)$. The Fisher information \mathcal{I} can be expressed as a quadratic function

of the sensitivities of \mathbf{x}_k with respect to θ_1 :

$$\mathcal{I} = \sum_{k=0}^N \frac{\partial}{\partial \theta_1} \mathbf{x}_k^T \begin{bmatrix} 1/\gamma & 0 & 0 \\ 0 & 1/\gamma & 0 \\ 0 & 0 & 0 \end{bmatrix} \frac{\partial}{\partial \theta_1} \mathbf{x}_k$$

where the sensitivities $\frac{\partial}{\partial \theta_1} \mathbf{x}_k$ satisfy the following sensitivity equations:

$$\frac{\partial}{\partial \theta_1} \mathbf{x}_{k+1} = U_k \begin{bmatrix} \theta_2 & 0 & 0 \\ 0 & \theta_2 & 0 \\ 0 & 0 & \theta_1 \end{bmatrix} \frac{\partial}{\partial \theta_1} \mathbf{x}_k + U_k \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}.$$

It should be noted that for system (3), the objective function \mathcal{I} has many local optima as a function of the input sequence U_k . Thus, in contrast with (Maidens et al., 2016a) which consider optimal experiment design for hyperpolarized MRI problems, for this model, local search methods provide little insight into what acquisition sequences are good. In contrast with the MRI model presented in (Maidens and Arcak, 2016), where global optimal experiment design heuristics are developed for linear dynamical systems, in this model the decision variables U_k multiply the state vector \mathbf{x}_k , making the output \mathbf{y}_k a nonlinear function of the sequence $U = (U_0, \dots, U_{k-1})$. Thus we must use dynamic programming to find a solution.

4.1 Model

To present this problem in the formalism we have introduced, we define an augmented state vector

$$x_k = \begin{bmatrix} \mathbf{x}_k \\ \frac{\partial}{\partial \theta_1} \mathbf{x}_k \end{bmatrix} \in \mathbb{R}^6.$$

We can write the dynamics of the augmented state as a control system with f and g defined in Equation (4). This system has a one-dimensional group of symmetries defined by

$$\phi_\alpha(x_k) = \begin{bmatrix} \cos(\alpha) & -\sin(\alpha) & 0 & 0 & 0 & 0 \\ \sin(\alpha) & \cos(\alpha) & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \cos(\alpha) & -\sin(\alpha) & 0 \\ 0 & 0 & 0 & \sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} x_k$$

$$\chi_\alpha(U_k) = \begin{bmatrix} \cos(\alpha) & -\sin(\alpha) & 0 \\ \sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 1 \end{bmatrix} U_k \begin{bmatrix} \cos(\alpha) & \sin(\alpha) & 0 \\ -\sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$\psi_\alpha(w_k) = w_k$$

for any $\alpha \in \mathbb{R}/2\pi\mathbb{Z}$.

4.2 Dynamic programming in reduced coordinates

To perform dynamic programming in a reduced coordinate system, we begin by defining the cross-section $\mathcal{C} = \{x : x_1 = 0, x_2 > 0\}$, and computing the moving frame $\gamma(x)$. To do so, we solve

$$0 = \phi_{\gamma(x)}^a(x) = x_1 \cos \gamma(x) - x_2 \sin \gamma(x).$$

Isolating γ yields

$$\gamma(x) = \text{atan2}(x_1, x_2).$$

Next, we compute the invariants $\rho(x)$ using

$$\begin{aligned} \rho(x) &= \phi_{\gamma(x)}^b \\ &= \begin{bmatrix} \sin(\text{atan2}(x_1, x_2)) & \cos(\text{atan2}(x_1, x_2)) & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \cos(\text{atan2}(x_1, x_2)) & -\sin(\text{atan2}(x_1, x_2)) & 0 \\ 0 & 0 & 0 & \sin(\text{atan2}(x_1, x_2)) & \cos(\text{atan2}(x_1, x_2)) & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} x \\ &= \begin{bmatrix} \sqrt{x_1^2 + x_2^2} \\ x_3 \\ \frac{1}{\sqrt{x_1^2 + x_2^2}}(x_2 x_4 - x_1 x_5) \\ \frac{1}{\sqrt{x_1^2 + x_2^2}}(x_1 x_4 + x_2 x_5) \\ x_6 \end{bmatrix} \end{aligned}$$

Further, ρ restricted to the cross-section \mathcal{C} is injective with inverse $\bar{\rho}^{-1} : \mathbb{R}_+ \times \mathbb{R}^4 \rightarrow \mathcal{C}$ given by $\bar{\rho}^{-1}(\bar{x}) = \begin{bmatrix} 0 & \bar{x}_1 & \bar{x}_2 & \bar{x}_3 & \bar{x}_4 & \bar{x}_5 \end{bmatrix}^T$. The theory above tells us we can thus solve the optimal stochastic control problem in a 5 dimensional state space, reducing the original 6 dimensional problem of 1 dimension.

4.3 Results

To implement this algorithm, we discretize the reduced five-dimensional state space and two-dimensional input space via grids of size $6 \times 10 \times 15 \times 15 \times 15$ and 16×8 respectively. The code was written in the Julia language and parallelized to allow evaluation of J_k in parallel across grid points (Maidens et al., 2016b). The implementation is publicly available at <https://github.com/maidens/Automatica-2017>.

Optimal input and state trajectories for the model corresponding to the initial condition at the equilibrium $x_0 = [0 \ 0 \ 1 \ 0 \ 0 \ 0]^T$ are plotted in Figures 1 and 2.






Fig. 1. Optimal input sequence for the MR fingerprinting model. The angles α , β and δ represent rotations about the z , y and x axes respectively, resulting in an control input $U_k = R_z(\alpha_k)R_y(\beta_k)R_x(\gamma_k)$.

In contrast with the results from (Maidens et al., 2016b) where we considered a simplified version of the model, for this full model we no longer find that the optimal flip angle sequence converges to a cyclic pattern, rather it appears irregular. However, state sequence of longitudinal magnetizations and transverse magnetization magnitudes appears to converge to a constant sequence. This is likely because in this work we assumed Gaussian noise in the inputs in contrast with the Rician noise assumed in the previous work, therefore it is no longer necessary to conserve magnetization across multiple time steps before generating a reliable measurement.



(a) Magnetizations



(b) Sensitivities

Fig. 2. Optimal state sequence for the MR fingerprinting model. Here we have plotted the longitudinal and transverse components of both the magnetization (states x_1 , x_2 , and x_3) and the sensitivities (states x_4 , x_5 , and x_6) where the transverse component is computed as the Euclidean norm of the vectors (x_1, x_2) and (x_4, x_5) respectively.

5 Optimal Formation Control on Lie Groups

We now apply the theory presented in Section 2.2 where the state space $\mathcal{X} \subseteq \mathbb{R}^n$ is the Cartesian product of matrix Lie groups. Note that, straightforward modifications arise along the way as the state space and noise space are not vector spaces as in the theory above. The methodology is then applied to the synchronization of two non-holonomic cars in the presence of uncertainties.

5.1 General problem

We model the system as a collection of K agents, where the state of each agent evolves on a r -dimensional matrix Lie group \mathcal{G} . We assume that the evolution of the state of agent j proceeds according to the equation

$$X_{k+1}^j = X_k^j M(u_k^j) W_k^j \quad (5)$$

where X_k , $M(u_k)$, W_k are all square matrices belonging to \mathcal{G} , u_k is a control that lives in some finite dimensional vector space, and W_k is the noise. The control objective is to reach a desired configuration, that is, a desired value for the relative configurations of the agents $(X^1)^{-1}X^2, \dots, (X^{K-1})^{-1}X^K$, see e.g., Sarlette et al. (2010) for more information.

Systems of this form are naturally invariant to left multiplication of all X^j by some matrix $A \in \mathcal{G}$:

$$\phi_A(X) = \begin{bmatrix} AX^1 \\ \vdots \\ AX^K \end{bmatrix}$$

where $X = (X^1, \dots, X^K) \in \mathcal{G}^K$. Letting $\chi(u^1, \dots, u^K) \equiv (u^1, \dots, u^K)$, $\psi(W^1, \dots, W^K) \equiv (W^1, \dots, W^K)$, and the costs be of the form $\tilde{g}((X^1)^{-1}X^2, \dots, (X^{K-1})^{-1}X^K) + h(u^1, \dots, u^K)$, we get an invariant system with invariant costs.

One can define a cross section to the orbits by letting the first agent coordinates be equal to the identity matrix, that is, $\mathcal{C} = \{X \in \mathcal{G}^K : X^1 = I\}$. The normalization equation are given by $I = \phi_{\gamma(X)}^a(X) = \gamma(X)X^1$, hence the moving frame is given by $\gamma(X) = (X^1)^{-1}$. The invariants are computed as

$$\rho(X) = \phi_{\gamma(X)}^b = \begin{bmatrix} (X^1)^{-1}X^2 \\ \vdots \\ (X^1)^{-1}X^K \end{bmatrix}$$

The optimal stochastic control problem can then be solved in the reduced coordinate system defined by ρ , reducing the state space from dimension Kr to $(K-1)r$.

5.2 Application: cooperative formation control for two stochastic Dubins vehicles

We consider two identical Dubins vehicles each with dynamics

$$\begin{aligned} z_{k+1} &= z_k + v_k \cos \theta_k \\ y_{k+1} &= y_k + v_k \sin \theta_k \\ \theta_{k+1} &= \theta_k + \frac{1}{L} v_k \tan s_k + w_k \end{aligned}$$

where y_k and z_k denote the two-dimensional position of the vehicle, θ_k denotes the heading of the vehicle, v_k is a velocity input, s_k is a steering angle input, and w_k is independent, identically-distributed zero-mean Gaussian noise with variance σ^2 , and L is a parameter that determines the vehicle's steering radius.

These dynamics can be embedded in the three-dimensional special Euclidean matrix Lie group $\mathcal{G} = SE(2)$, by defining the state

$$X_k = \begin{bmatrix} \cos \theta_k & -\sin \theta_k & z_k \\ \sin \theta_k & \cos \theta_k & y_k \\ 0 & 0 & 1 \end{bmatrix},$$

input matrix

$$M(v_k, s_k) = \begin{bmatrix} \cos(\frac{1}{L}v_k \tan s_k) & -\sin(\frac{1}{L}v_k \tan s_k) & v_k \\ \sin(\frac{1}{L}v_k \tan s_k) & \cos(\frac{1}{L}v_k \tan s_k) & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

and noise matrix

$$W_k = \begin{bmatrix} \cos w_k & -\sin w_k & 0 \\ \sin w_k & \cos w_k & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

with state update equation of the form (5).

We wish to compute a control policy for a two-vehicle system, with states X^1 and X^2 , where the controls can only take a finite number of values, and with terminal cost

$$J(X_0^1, X_0^2) = \mathbb{E} \left[g_N((X_N^1)^{-1} X_N^2) \right]$$

where $g_N(X) = \arccos(X_{11})^2 + |\sqrt{X_{13}^2 + X_{23}^2} - 1|$, that is, we want the vehicles to have the same heading, and follow each other at unit distance. Thanks to the theory developed above, the stochastic control problem is reduced from problem with a six dimensional state space to a problem with a three dimensional state space only.

For numerical simulations, the cost functions \bar{J}_k were computed on a fixed grid of dimension $51 \times 51 \times 65$ using turning radius parameter $L = 1$, input sets $v_k \in \{-0.1, 0, -0.1\}$ and $s_k \in \{-1, 0, -1\}$. Globally optimal input and state trajectory sequences corresponding to the initial condition $x_0 = [0.1 \ 0 \ \frac{1}{2}\pi \ -0.1 \ 0 \ \frac{3}{2}\pi]^T$ are shown in Figures 3 and 4. These are compared against a deterministic version of the model with $w_k = 0$ in Figures 5 and 6.

6 Conclusion

We have presented a method of reducing the complexity of dynamic programming for systems in which the state dynamics, stage costs and transition probabilities are invariant under a group of symmetries. This allows us to




Fig. 3. Optimal input sequence for cooperative stochastic Dubins vehicle model with $\sigma = 0.3$.




Fig. 4. Optimal state sequence for the cooperative stochastic Dubins vehicle model with $\sigma = 0.3$.

compute globally optimal control policies for systems of moderate state dimension. We have applied this technique to compute globally optimal trajectories to a six-dimensional original MRI model with a one-dimensional group of symmetries and for a six-dimensional stochastic Dubins vehicle model with a three-dimensional group of symmetries by reducing the dimension of the state space to five and three dimensions respectively. Since computation time for dynamic programming depends exponentially on the state space dimension, this technique enables the computation of optimal control policies for systems in which it was previously infeasible.

Acknowledgements

Research supported in part by the National Science Foundation under grant ECCS-1405413.




Fig. 5. Optimal input sequence for cooperative deterministic Dubins vehicle model.




Fig. 6. Optimal state sequence for the cooperative deterministic Dubins vehicle model.

References

- F. Alvarez and N. L. Stokey. Dynamic programming with homogeneous functions. *Journal of economic theory*, 82(1):167–189, 1998.
- M. Arcak, C. Meissen, and A. Packard. Symmetry reduction. In *Networks of Dissipative Systems: Compositional Certification of Stability, Performance, and Safety*, pages 55–62. Springer, 2016.
- A. Barrau and S. Bonnabel. The invariant extended Kalman filter as a stable observer. *Arxiv preprint*, page arXiv:1410.1465, 2014.
- R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- R. Bellman and S. Dreyfus. Functional approximations and dynamic programming. *Mathematical Tables and Other Aids to Computation*, 13(68):247–251, 1959.
- D. P. Bertsekas. *Dynamic Programming and Optimal Control, Volume I*. Athena Scientific, 3rd edition, 2005.
- D. P. Bertsekas. *Dynamic Programming and Optimal Control, Volume II: Approximate Dynamic Programming*. Athena Scientific, 4th edition, 2012.
- A. M. Bloch, P. S. Krishnaprasad, J. E. Marsden, and R. M. Murray. Nonholonomic mechanical systems with symmetry. *Arch. Rational Mech. Anal.*, 136:21–99, 1996.
- F. Bloch. Nuclear induction. *Phys. Rev.*, 70:460–474, 1946.
- G. Bluman and S. Kumei. *Symmetries and Differential Equations*. Applied Mathematical Sciences. Springer, 2013.
- S. Bonnabel, P. Martin, and P. Rouchon. Symmetry-preserving observers. *IEEE Transactions on Automatic Control*, 53(11):2514–2526, 2008.
- F. Bullo and R. M. Murray. Tracking for fully actuated mechanical systems: a geometric framework. *Automatica*, 35(1):17–34, 1999.
- É. Cartan. *La théorie des groupes finis et continus et la géométrie différentielle: traitées par la méthode du repère mobile*. Cahiers scientifiques. Gauthier-Villars, 1937.
- P. A. Clarkson and E. L. Mansfield. Symmetry reductions and exact solutions of a class of nonlinear heat equations. *Physica D: Nonlinear Phenomena*, 70(3):250 – 288, 1994.
- M. Davies, G. Puy, P. Vandergheynst, and Y. Wiaux. A compressed sensing framework for magnetic resonance fingerprinting. *SIAM Journal on Imaging Sciences*, 7(4):2623–2656, 2014.
- E. A. Emerson and A. P. Sistla. Symmetry and model checking. *Formal Methods in System Design*, 9(1):105–131, 1996.
- J. Grizzle and S. Marcus. Optimal control of systems possessing symmetries. *IEEE Transactions on Automatic Control*, 29(11):1037–1040, 1984.
- B. Jakubczyk. Symmetries of nonlinear control systems and their symbols. In *Canadian Math. Conf. Proceed*, volume 25, pages 183–198, 1998.
- M. Kwiatkowska, G. Norman, and D. Parker. Symmetry reduction for probabilistic model checking. In T. Ball and R. B. Jones, editors, *Computer Aided Verification: 18th International Conference, CAV 2006, Seattle, WA, USA*, pages 234–248. Springer, 2006.
- D. Ma, V. Gulani, N. Seiberlich, K. Liu, J. L. Sunshine, J. L. Duerk, and M. A. Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–192, 03 2013.
- J. Maidens and M. Arcak. Semidefinite relaxations in optimal experiment design with application to substrate injection for hyperpolarized MRI. In *Proceedings of the American Control Conference (ACC)*, pages 2023 – 2028, 2016.
- J. Maidens, J. W. Gordon, M. Arcak, and P. E. Z. Larson. Optimizing flip angles for metabolic rate estima-

- tion in hyperpolarized carbon-13 MRI. *IEEE Transactions on Medical Imaging*, 35:2403 – 2412, 2016a.
- J. Maidens, A. Packard, and M. Arcaç. Parallel dynamic programming for optimal experiment design in nonlinear systems. In *Proceedings of the Conference on Decision and Control (CDC)*, pages 2894 – 2899, 2016b.
- J. Maidens, A. Barrau, S. Bonnabel, and M. Arcaç. Symmetry reduction for dynamic programming and application to MRI. In *Proceedings of the American Control Conference (ACC)*, 2017. To appear.
- P. Martin, P. Rouchon, and J. Rudolph. Invariant tracking. *ESAIM: Control, Optimisation and Calculus of Variations*, 10(1):1–13, 2004.
- S. M. Narayanamurthy and B. Ravindran. Efficiently exploiting symmetries in real time dynamic programming. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, pages 2556–2561. Morgan Kaufmann, 2007.
- D. G. Nishimura. *Principles of Magnetic Resonance Imaging*. Lulu, 2010.
- T. Ohsawa. Symmetry reduction of optimal control systems and principal connections. *SIAM Journal on Control and Optimization*, 51(1):96–120, 2013.
- P. J. Olver. *Classical invariant theory*, volume 44. Cambridge University Press, 1999.
- W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley-Interscience, 2007.
- W. B. Powell. Perspectives of approximate dynamic programming. *Annals of Operations Research*, 241(1): 319–356, 2016.
- A. Rahmani, M. Ji, M. Mesbahi, and M. Egerstedt. Controllability of multi-agent systems from a graph-theoretic perspective. *SIAM Journal on Control and Optimization*, 48(1):162–186, 2009.
- W. Respondek and I. A. Tall. Nonlinearizable single-input control systems do not admit stationary symmetries. *Systems & Control Letters*, 46(1):1–16, 2002.
- A. Rufino Ferreira, C. Meissen, M. Arcaç, and A. Packard. Symmetry reduction for performance certification of interconnected systems. *IEEE Transactions on Control of Networked Systems*, 2017.
- A. Sarlette, S. Bonnabel, and R. Sepulchre. Coordinated motion design on lie groups. *IEEE Transactions on Automatic Control*, 55(5):1047–1058, 2010.
- M. Zinkevich and T. Balch. Symmetry in Markov decision processes and its implications for single agent and multi agent learning. In *In Proceedings of the 18th International Conference on Machine Learning*, pages 632–640. Morgan Kaufmann, 2001.