This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2026.3727165

Bounding the Error of Value Functions in Sobolev Norm Yields Bounds on Suboptimality of Controller Performance

Morgan Jones, Matthew M. Peet

***Abstract*—Optimal feedback controllers for nonlinear systems can be derived by solving the Hamilton-Jacobi-Bellman (HJB) equation. However, because the HJB equation is a nonlinear partial differential equation, numerical methods typically provide only approximate solutions. While numerical error bounds on approximate HJB solutions are often available, these bounds do not necessarily translate into guarantees on the suboptimality of the resulting controllers. In this paper, we establish that the suboptimality of the resulting controller is bounded by the L^∞ norm of the HJB residual, which is, in turn, bounded by numerical error in the value function as measured in the Sobolev $W^{1,\infty}$ norm. This implies that convergence of value functions in $W^{1,\infty}$ result in controllers that yield a cost that is arbitrarily close to the true minimum. In contrast, we demonstrate that such guarantees do not hold when the value function error is measured in weaker norms, such as the Sobolev $W^{1,p}$ norm for finite p. These results apply to systems governed by Lipschitz continuous dynamics over a finite time horizon with compact input space.**

I. Introduction

Consider finite horizon Optimal Control Problems (OCPs) in continuous time. The objective of an OCP is to determine an input that guides the system dynamics to minimize a given cost function over a specified time horizon. Arguably one of the most popular approaches for solving OCPs is Dynamic Programming (DP), which requires solution of the Hamilton-Jacobi-Bellman (HJB) Partial Differential Equation (PDE). The solution of the HJB PDE, known as the Value Function (VF), is then used to design an optimal controller.

Unfortunately, solving the HJB PDE is notoriously difficult. In general, the HJB PDE does not typically possess a classical smooth solution, and it wasn't until the 1980s that a suitable alternative, known as the viscosity solution, was introduced [1]. To approximate this viscosity solution/value function, several classes of numerical methods have been developed. Such methods include the (mixed) finite elements approach [2] for which one may obtain an approximate VF with error bound in terms of the first order mixed L^2 norm (assuming the Cordes condition holds). Other examples of this class of methods includes discretization approaches such as [3] or [4], which yield an approximate VF with an L^∞ error bound – a bound which convergse as the level of discretization increases. Semi-Lagrangian (SL) methods [5] approximate the solution by discretizing the OCP directly, rather than using finite difference approximations of the HJB PDE derivatives. SL methods synthesize feedback controls via a discrete-time lookahead procedure, distinct from the standard instantaneous feedback policy used in continuous control considered in this paper, see Eq. (4). While bounds on the local one-step optimality gap can be derived for SL schemes under certain assumptions (see Chapter 8 in [5]), these cannot be used in the context of continuous-time controllers/policies defined over the entire trajectory (as in Theorem 1 from this paper). Alternative notable numerical HJB PDE solution schemes include the method of characteristics [6], max-plus methods [7], grid based methods [8], [9], Neural Network methods [10], [11], policy iteration [12], [13] and value iteration [14].

An important emerging field for approximately solving the HJB PDE is that of Physics Informed Neural Networks (PINNs) [15]–[17]. In this class of methods, a neural network is used to parameterize solutions of HJB PDE and the residual, as evaluated at collocation (training data) points, is minimized. However, as shown in [18], minimizing the Euclidean norm of the HJB residual may result in PINN methods failing to converge to the VF. We show, minimizing the HJB residual in the L^∞ norm does yield controllers with arbitrarily good performance, however, we also show that even if a method converges to the VF in some sense, such as in the $W^{1,p}$ norm for finite $p < \infty$, this may result in a controller with poor performance no matter how tight the VF approximation is, unless $p = \infty$. This is illustrated in Example 2 from Section V.

While the use of approximate VFs to construct controllers has been studied, such controllers typically: apply only to OCPs with specific structure (e.g. affine in input [19]); do not have associated performance bounds; and/or have conservative assumptions like the differentiability of the VF [10], [20]–[23]. The work of [24] does provide a performance bound but requires the assumption that there exists a transformation that linearizes the HJB PDE. Alternatively [25] proposes a bilinear SOS optimization framework without any convergence guarantees. The SOS and moment based approaches [26]–[28] yield controllers that converge to optimal performance under specific assumptions, however, this convergence is not related to VF approximation. By contrast, the results of this paper provide a general performance bound independent of the approach used to approximate the VF.

For discrete time OCPs over infinite time horizons it has been shown in Proposition 6.1 from [29] (and also presented in [30]) that if a function J, is ε-close to the true VF under the uniform norm, then the sub-optimality of the controller synthesized using J is bounded by $2\alpha\varepsilon/(1-\alpha)$, where $\alpha > 0$ is the discount factor of the OCP. In this paper we seek to derive an analogous result for the optimality gap of controllers synthesized using approximate VFs for continuous time OCPs over finite time horizons. The main result of the paper, given in Theorem 1, shows that controller performance is bounded by $C\varepsilon$ where $\varepsilon > 0$ is the VF approximation error in the $W^{1,\infty}$ norm and $C > 0$ is a known constant depending on the problem data.

M. Jones is with the School of Electrical and Electronic Engineering, The University of Sheffield, e-mail: morgan.jones@sheffield.ac.uk

M. Peet is with the School for the Engineering of Matter, Transport and Energy, Arizona State University, e-mail: mpeet@asu.edu

Authorized licensed use limited to: Arizona State University. Downloaded on October 27,2026 at 09:29:22 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2026.3727165

Notation: For $x \in \mathbb{R}^n$ we denote the Euclidean norm by $\|x\|_2 = \sqrt{\sum_{i=1}^n x_i^2}$. For $r > 0$ and $x \in \mathbb{R}^n$ we denote the open ball $B_r(x) := \{y \in \mathbb{R}^n : \|x-y\|_2 < r\}$ and closed ball $\bar{B}_r(x) := \{y \in \mathbb{R}^n : \|x-y\|_2 \leq r\}$. For $p \in [1,\infty)$ we denote the set of p-integrable functions by $L^p(\Omega) := \{f : \Omega \to \mathbb{R} \text{ measurable } : \int_\Omega |f|^p < \infty\}$, in the case $p = \infty$ we denote $L^\infty(\Omega) := \{f : \Omega \to \mathbb{R} \text{ measurable } : \operatorname{ess\,sup}_{x\in\Omega} |f(x)| < \infty\}$. We use the notation a.e to mean almost everywhere, meaning the set in which a property does not hold has measure zero. For $\alpha \in \mathbb{N}^n$ we denote the partial derivative $D^\alpha f(x) := \Pi_{i=1}^n \frac{\partial^{\alpha_i} f}{\partial x_i^{\alpha_i}}(x)$. Let $C(\Omega,\Theta)$, $AC(\Omega,\Theta)$ and $Lip(\Omega,\Theta)$ be the set of continuous, absolutely continuous and uniformly Lipschitz continuous functions respectively with domain $\Omega \subseteq \mathbb{R}^n$ and image $\Theta \subseteq \mathbb{R}^m$. We denote the set of i continuously differentiable functions by $C^i(\Omega,\Theta) := \{f \in C(\Omega,\Theta) : D^\alpha f \in C(\Omega,\Theta) \text{ for all } \alpha \in \mathbb{N}^n,\ \|\alpha\|_1 \leq i\}$. For $k \in \mathbb{N}$ and $1 \leq p \leq \infty$ we denote the Sobolev space of functions with weak derivatives by $W^{k,p}(\Omega) := \{u \in L^p(\Omega) : D^\alpha u \in L^p(\Omega) \text{ for all } \alpha \in \mathbb{N}^n,\ \|\alpha\|_\infty \leq k\}$. For $u \in W^{k,p}(\Omega)$ we denote the Sobolev norm $\|u\|_{W^{k,p}(\Omega)} := \begin{cases} \left(\sum_{|\alpha|\leq k} \int_\Omega |D^\alpha u(x)|^p dx\right)^{\frac{1}{p}} & \text{if } 1 \leq p < \infty \\ \sum_{|\alpha|\leq k} \operatorname{ess\,sup}_{x\in\Omega}\{|D^\alpha u(x)|\} & \text{if } p = \infty. \end{cases}$ In the case $k = 0$ we have $W^{0,p}(\Omega) = L^p(\Omega)$ and thus we use the notation $\|\cdot\|_{L^p(\Omega)} := \|\cdot\|_{W^{0,p}(\Omega)}$. For $V \in C^1(\mathbb{R}\times\mathbb{R}^n,\mathbb{R})$ we denote $\nabla_x V := (\frac{\partial V}{\partial x_2}, ..., \frac{\partial V}{\partial x_{n+1}})^\top$. For a function $u \in C(\Omega,\mathbb{R})$ and $x \in \Omega \subset \mathbb{R}^N$ we define the super/sub differential sets

$$D^+u(x) := \left\{p \in \mathbb{R}^N : \limsup_{y\to x, y\in\Omega} \frac{u(y)-u(x)-p^\top(y-x)}{|y-x|} \leq 0\right\}.$$

$$D^-u(x) := \left\{p \in \mathbb{R}^N : \liminf_{y\to x, y\in\Omega} \frac{u(y)-u(x)-p^\top(y-x)}{|y-x|} \geq 0\right\}.$$

We denote the union of super/sub differential sets as $D^\pm u(x) = D^+u(x) \cup D^-u(x)$. For $u \in \text{Lip}(\Omega,\mathbb{R})$, we define the *Clarke generalized gradient* at $x \in \Omega$ as

$$\partial_C u(x) := \left\{p \in \mathbb{R}^n : p^\top v \leq \limsup_{y\to x, h\downarrow 0} \frac{u(y+hv)-u(y)}{h} \text{ for } v \in \mathbb{R}^n\right\}.$$

For functions J over $\mathbb{R}\times\mathbb{R}^n$, elements $(p,q) \in D^\pm J(t,x)$ or $(p,q) \in \partial_C J(t,x)$ explicitly separate the time differential $p \in \mathbb{R}$ from the spatial differential $q \in \mathbb{R}^n$.

II. Optimal Control Problems

Consider the following family of Optimal Control Problems (OCPs), each initialized by $(t_0,x_0) \in [0,T]\times\mathbb{R}^n$,

$$\begin{aligned} &V^*(t_0,x_0) := \inf_u \int_{t_0}^T c(t,x(t),u(t))dt + g(x(T)) \text{ subject to:} \\ &x(t_0) = x_0, \quad \dot{x}(t) = f(t,x(t),u(t)) \text{ for } t \in [t_0,T], \\ &u : [t_0,T] \to \mathbb{R}^m \text{ is a measurable function,} \\ &u(t) \in U \text{ for all } t \in [t_0,T], \end{aligned} \tag{1}$$

where $c : \mathbb{R}\times\mathbb{R}^n\times\mathbb{R}^m \to \mathbb{R}$, $g : \mathbb{R}^n \to \mathbb{R}$, $f : \mathbb{R}\times\mathbb{R}^n\times\mathbb{R}^m \to \mathbb{R}^n$, $U \subset \mathbb{R}^m$, and $T > 0$. Each OCP (1) is associated with a tuple $\{c,g,f,U,T\}$. A measurable function, $u : [t_0,T] \to \mathbb{R}^m$, is an *admissible input* for OCP (1) if $u(t) \in U$. We refer to V^*, defined in Eq. (1), as the Value Function (VF) of the OCP. Throughout this paper we consider the following standard class of OCPs for which it can be shown that the VF is Lipschitz continuous and solution maps exist.

Definition 1. *We say OCP* (1) *corresponding to tuple* $\{c,g,f,U,T\}$ *is class L if the following hold:*

1) $U \subset \mathbb{R}^m$ *is a compact set.*
2) f*,* c *and* g *are uniformly Lipschitz continuous in the state variable and uniformly continuous in the time and input variable. Specifically, there exist* $\alpha_f,\alpha_c,\alpha_g > 0$ *such that* $\|f(t,x,u)-f(t,y,u)\|_2 \leq \alpha_f\|x-y\|_2$*,* $\|c(t,x,u)-c(t,y,u)\|_2 \leq \alpha_c\|x-y\|_2$ *and* $\|g(x)-g(y)\|_2 \leq \alpha_g\|x-y\|_2$ *for all* $x,y \in \mathbb{R}^n$*,* $t \in [0,T]$ *and* $u \in U$*.*
3) f*,* c *and* g *are bounded by* $\mathcal{O}(\|x\|_2)$*. Specifically, there exists* $\beta_f,\beta_c,\beta_g > 0$ *such that* $\|f(t,x,u)\|_2 \leq \beta_f(1+\|x\|_2)$*,* $\|c(t,x,u)\|_2 \leq \beta_c(1+\|x\|_2)$ *and* $\|g(x)\|_2 \leq \beta_g(1+\|x\|_2)$ *for all* $x \in \mathbb{R}^n$*,* $t \in [0,T]$ *and* $u \in U$*.*

Solution Map. Under the first two items of Definition 1, standard global existence and uniqueness theorems (see, e.g., [31, Thm. 3.2]) guarantee that for any initial condition (t_0,x_0) and any admissible input, the initial value problem $\dot{x}(t) = f(t,x(t),u(t)),\quad x(t_0) = x_0$ admits a unique absolutely continuous solution defined on the interval $[t_0,T]$. We denote this unique trajectory by $x_{u,x_0}(\cdot)$, and refer to the mapping $(u,x_0) \mapsto x_{u,x_0}$ as the solution map of the system.

Definition 2 (Admissible & Optimal Controller)**.** *Given an OCP* (1)*, we say a measurable function* $\pi(t,x)$ *is an* ***admissible controller*** *for the OCP if, for any* $(t_0,x_0) \in [0,T]\times\mathbb{R}^n$*,*

$$\dot{x}(t) = f(t,x(t),\pi(t,x(t))), \quad x(t_0) = x_0,$$

has a unique solution and $\pi(t,x(t)) \in U$ *for all* $t \in [t_0,T]$*. For* $t_0 = 0$*, we denote the corresponding admissible input as* $u_{\pi,x_0}(t) = \pi(t,x(t))$*. Furthermore, an admissible controller,* π*, is* ***optimal*** *if* u_{π,x_0} *solves OCP* (1) *with* $t_0 = 0$*.*

The Dynamic Programming (DP) method allows one to obtain an optimal controller by solving HJB PDE:

$$\begin{aligned} &\frac{\partial}{\partial t}V(t,x) + \inf_{u\in U}\{H(t,x,u,\nabla_x V(t,x))\} = 0 \\ &\qquad\qquad \text{for all } (t,x) \in [0,T]\times\mathbb{R}^n, \\ &V(T,x) = g(x) \quad \text{for all } x \in \mathbb{R}^n, \end{aligned} \tag{2}$$

where we define the Hamiltonian as

$$H(t,x,u,q) := c(t,x,u) + q^\top f(t,x,u). \tag{3}$$

Given a differentiable solution to the HJB PDE (2), V^*, an optimal controller may be constructed by minimizing the Hamiltonian [6]:

$$\pi^*(t,x) \in \arg\inf_{u\in U} H(t,x,u,\nabla_x V^*(t,x)). \tag{4}$$

Note, the HJB PDE may not possess a classical differentiable solution. However, it can be shown under mild assumptions that the HJB PDE will always admit a "viscosity solution" – See, e.g. [1]. Furthermore, for class L (Defn. 1) OCPs, it can be shown that the VF (as defined in Eq. (1)) is the unique viscosity solution of the HJB PDE (2) and is Lipschitz continuous.

Authorized licensed use limited to: Arizona State University. Downloaded on October 27,2026 at 09:29:22 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2026.3727165

Lemma 1 (Page 555 [32])**.** *Suppose tuple* $\{c, g, f, U, T\}$ *defines an OCP of class L (Defn. 1). Then* $V^* : [0,T]\times\mathbb{R}^n \to \mathbb{R}$*, as defined in Eq.* (1)*, is the unique viscosity solution of the HJB PDE* (2) *and* $V^* \in Lip([0,T]\times\mathbb{R}^n, \mathbb{R})$.

Eq. (4) can be used to determine the value of the optimal controller, $\pi^*(t,x)$, at points where $\nabla_x V^*(t,x)$ exists. Furthermore, Lemma 1 implies that for class L (Defn. 1) OCPs, V^* is Lipschitz continuous and hence differentiable almost everywhere. Unfortunately, however, segments of the optimal closed-loop trajectory may coincide with the measure zero regions of state space in which V^* is not differentiable . In such cases, we consider the approach used in [33] to define sub/super differentials ($D^{\pm}$) of V^* at these points.

Definition 3. *Suppose OCP* (1) *is class L (Defn. 1). We say a function* $J \in Lip([0,T]\times\mathbb{R}^n, \mathbb{R})$ *is an* ***admissible candidate value function*** *of OCP* (1) *if*

1) $D^{\pm}J(t,x) = \partial_C J(t,x)$ *for all* $(t,x) \in [0,T]\times\mathbb{R}^n$.
2) *There exists an admissible controller* $\pi_J : [0,T]\times\mathbb{R}^n \to U$ *and measurable functions* $(p_J(t,x), q_J(t,x)) \in D^{\pm}J(t,x)$ *such that*

$$(p_J(t,x), q_J(t,x), \pi_J(t,x)) \in \arg\inf_{(p,q,u)\in D^{\pm}J(t,x)\times U} \{p + H(t,x,u,q)\} \quad \text{for all } (t,x) \in [0,T]\times\mathbb{R}^n, \tag{5}$$

where H *is given in Eq.* (3)*.*

The use of sub/super differentials in defining the controller π_J, as in Definition 3, will allow us to extend the applicability of the main result given in Theorem 1 to OCPs where the VF is only Lipschitz continuous and not continuously differentiable. However, this definition requires that the union of sub- and superdifferentials of the candidate VF coincide with the Clarke generalised gradient. This condition is automatically satisfied when J is continuously differentiable, since $D^{\pm}J(t,x) = \partial_C J(t,x) = \{\nabla J(t,x)\}$ (Lemma 1.8 in [34], Theorem 10.8 in [35]). Examples of non-differentiable admissible candidate value functions include ridges or valleys – e.g. $J(x) = |x|$ or $J(x) = -|x|$, for which $D^{\pm}J(t,x) = \partial_C J(t,x) = [-1,1]$.

More generally, semiconcave and semiconvex functions are admissible (Prop. 4.7 in [34]), a class that includes VFs arising from a wide range of optimal control problems (see Chapter 4.2 in [34]). Examples of non-admissible candidate value functions (where $D^{\pm}J \neq \partial_C J$) include functions with saddle points (e.g., $J(t,x) = |x_1 - t| - |x_2 - t|$, where $D^{\pm}J(t,[t,t]^\top) = \emptyset$) and functions with singular oscillations such as $J(x) = x^{3/2}\sin(1/\sqrt{x})$ – for which $D^{\pm}J(0) = \{0\}$ but $\partial_C J(0) = [-1,1]$ (see Exercise 13.10 in [35]).

III. A Performance Bound On Controllers Constructed Using Approximate Value Functions

Having defined classes of admissible candidate VFs, we may now proceed to the main result of the paper. Of course, given an OCP, if a viscosity solution, V^* to the HJB PDE (2) is known, an admissible controller can be verified as optimal using [33, Thm. 3.5]. However, finding a HJB PDE viscosity solution is challenging and hence we consider the case where a function, J, is obtained which is close to V^* with respect to some norm, X – i.e. $\|J - V^*\|_X \leq \epsilon$, but does not necessarily satisfy the HJB PDE. This naturally raises the question: how well does a controller constructed from an approximate VF perform and does the choice of approximation norm matter?

To formalize this analysis, then, we first define the optimality gap (loss) of an initial condition, x_0 and admissible input, u, associated with an OCP (1) with $t_0 = 0$:

$$L(u, x_0) := \int_0^T c(s, x_{u,x_0}(s), u(s))ds + g(x_{u,x_0}(T)) - V^*(0, x_0), \tag{6}$$

where V^* is the VF of the OCP. Clearly, $L(u,x_0) \geq 0$ for any admissible u.

Given any admissible candidate value function, J, we may define an associated admissible controller, π_J, as in Eq. (5) (or equivalently Eq. (4) if J is differentiable). Our goal is then to bound the optimality gap $L(u_{\pi_J,x_0}, x_0)$ in terms of $\|V^* - J\|_{W^{1,\infty}([0,T]\times\mathbb{R}^n)}$, where V^* is the true VF of the OCP. This result can be stated as follows (with proof deferred).

Theorem 1. *Suppose the OCP* (1) *corresponding to tuple* $\{c, g, f, U, T\}$ *is of class L (Defn. 1), and denote the associated value function by* V^**. Let* J *be an admissible candidate value function (Definition 3), with associated admissible input* u_{π_J,x_0}*. Then for any* $x_0 \in \mathbb{R}^n$*, we have*

$$L(u_{\pi_J,x_0}, x_0) \leq C\|J - V^*\|_{W^{1,\infty}([0,T]\times B_R(0))}, \tag{7}$$

$$\text{where } C := 2\max\{1, T, T\beta_f(1+R)\}, \tag{8}$$

L (the loss) is defined in Eq. (6)*,* $R := (1 + \|x_0\|_2)e^{\beta_f T} - 1$*, and* β_f *is as in Defn. 1.*

To establish Theorem 1, we first show that for any candidate value function there is an associated OCP for which the resulting admissible controller is optimal.

Lemma 2. *Consider OCP* (1) *with tuple* $\{c, g, f, U, T\}$*. For any admissible candidate value function* J *(Definition 3), suppose* π_J *satisfies Eq.* (5)*. Then* π_J *is optimal for the OCP defined by the tuple* $\{\tilde{c}, \tilde{g}, f, U, T\}$*, where*

$$\tilde{c}(t,x,u) := c(t,x,u) - E_J(t,x), \quad \tilde{g}(x) := J(T,x), \tag{9}$$

$$E_J(t,x) := \inf_{(p,q,u)\in D^{\pm}J(t,x)\times U} \{p + c(t,x,u) + q^\top f(t,x,u)\}.$$

Proof. For any admissible input $u : [0,T] \to U$, let us denote the associated cost as $\tilde{W}_u(t_0,x_0) := \int_{t_0}^T \tilde{c}(t, x_{u,x_0}(t), u(t))dt + \tilde{g}(x_{u,x_0}(T))$. We show π_J is optimal by showing that

1) $J(0,x_0) \leq \tilde{W}_u(0,x_0)$ for any admissible input, u.
2) $J(0,x_0) = \tilde{W}_{u_{\pi_J,x_0}}(0,x_0)$.

To begin, consider (π_J, p_J, q_J) satisfying Eq. (5). Then $(p_J(t,x), q_J(t,x)) \in D^{\pm}J(t,x)$ for all $(t,x) \in [0,T]\times\mathbb{R}^n$ and we have

$$\begin{aligned} &p_J(t,x) + \inf_{u\in U}\{\tilde{c}(t,x,u) + q_J(t,x)^\top f(t,x,u)\} \\ &= \inf_{u\in U}\{p_J(t,x) + c(t,x,u) + q_J(t,x)^\top f(t,x,u)\} - E_J(t,x) \\ &= \inf_{(p,q,u)\in D^{\pm}J(t,x)\times U}\{p + c(t,x,u) + q^\top f(t,x,u)\} - E_J(t,x) \\ &= E_J(t,x) - E_J(t,x) = 0. \end{aligned} \tag{10}$$

Authorized licensed use limited to: Arizona State University. Downloaded on October 27,2026 at 09:29:22 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2026.3727165

Now, for any admissible input, u, x_{u,x_0} is absolutely continuous which implies by Lemma 5 (found in the appendix) that $\frac{d}{dt}J(t,x_{u,x_0}(t)) = p + q^\top f(t,x_{u,x_0}(t),u(t))$ a.e for $t \in [0,T]$ and for any $(p,q) \in D^{\pm}J(t,x_{u,x_0}(t))$. Since $(p_J(t,x_{u,x_0}(t)), q_J(t,x_{u,x_0}(t))) \in D^{\pm}J(t,x_{u,x_0}(t))$, we have

$$\begin{aligned}
&\frac{d}{dt}J(t,x_{u,x_0}(t)) \qquad (11)\\
&= p_J(t,x_{u,x_0}(t)) + q_J(t,x_{u,x_0}(t))^\top f(t,x_{u,x_0}(t),u(t))\\
&= p_J(t,x_{u,x_0}(t)) + \tilde{c}(t,x_{u,x_0}(t),u(t))\\
&+ q_J(t,x_{u,x_0}(t))^\top f(t,x_{u,x_0}(t),u(t)) - \tilde{c}(t,x_{u,x_0}(t),u(t))\\
&\geq p_J(t,x_{u,x_0}(t)) - \tilde{c}(t,x_{u,x_0}(t),u(t))\\
&+ \inf_{v\in U}\{\tilde{c}(t,x_{u,x_0}(t),u) + q_J(t,x_{u,x_0}(t))^\top f(t,x_{u,x_0}(t),v)\}\\
&= -\tilde{c}(t,x_{u,x_0}(t),u(t)) \text{ a.e } t\in[0,T].
\end{aligned}$$

Where the last equality follows from Eq. (10). Integrating Eq. (11) over $[0,T]$ and using $J(T,x) = \tilde{g}(x)$ we deduce $J(0,x_0) \leq \tilde{W}_u(0,x_0)$ for any admissible input.

Now consider the admissible input, u_{π_J,x_0}, obtained from the candidate value function J and associated controller, π_J. To reduce notational complexity we denote the trajectory associated with this input, $x_{u_{\pi_J,x_0},x_0}$, as x_J. Then, by Eq. (10) and since it follows from Eq. (5) that the infimum in Eq. (10) is attained by u_{π_J,x_0} we have

$$\begin{aligned}
&p_J(t,x_J(t)) + \tilde{c}(t,x_J(t),u_{\pi_J,x_0}(t))\\
&\quad + q_J(t,x_J(t))^\top f(t,x_J(t),u_{\pi_J,x_0}(t))\\
&= p_J(t,x_J(t)) +\\
&\quad \inf_{v\in U}\{\tilde{c}(t,x_J(t),v) + q_J(t,x_J(t))^\top f(t,x_J(t),v)\} = 0,
\end{aligned}$$

and hence as per Eq. (11), $\frac{d}{dt}J(t,x_J(t)) = -\tilde{c}(t,x_J(t),u_{\pi_J,x_0}(t))$. We conclude by integrating that $J(0,x_0) = \tilde{W}_{u_{\pi_J,x_0}}(0,x_0)$ – establishing that the associated control π_J is optimal. □

Before deriving our controller performance bound (given in Theorem 1) we next prove E_J, as defined in Eq. (9), is lower-semi continuous for any admissible candidate value function, J – allowing us to apply Lemma 4 (found in the appendix).

Lemma 3. *Consider OCP* (1) *with tuple* $\{c,g,f,U,T\}$. *Let J be an admissible candidate value function (Definition 3), then* $E_J(t,x) := \inf_{(p,q,u)\in D^{\pm}J(t,x)\times U}\{p + c(t,x,u) + q^\top f(t,x,u)\}$ *is lower semi-continuous.*

Proof. To show E_J is lower semi-continuous, we show that for any bounded sequence $(t_n,x_n) \to (t,x)$, we have $E_J(t,x) \leq \liminf_{n\to\infty} E_J(t_n,x_n)$.

First, by definition of $\liminf$ there exists a subsequence $\{t_{n_k},x_{n_k}\}_k$ such that $\lim_{k\to\infty} E_J(t_{n_k},x_{n_k}) = \liminf_{n\to\infty} E_J(t_n,x_n)$. Now, for each k in the subsequence, let $\{p_k,q_k,u_k\}$ be the sequence of associated minimizers so that:

$$E_J(t_{n_k},x_{n_k}) = p_k + H(t_{n_k},x_{n_k},u_k,q_k),$$

where H is given in Eq. (3). Let L_J be a uniform Lipschitz constant for $J \in Lip([0,T]\times\mathbb{R}^n,\mathbb{R})$. It follows that $\partial_C J(t,x) \subset \bar{B}_{L_J}(0)$ for $(t,x)\in[0,T]\times\mathbb{R}^n$ (See, e.g. [35] (Prop 10.5)). Since $(p_k,q_k,u_k) \in D^{\pm}J(t_{n_k},x_{n_k}) \times U = \partial_C J(t_{n_k},x_{n_k}) \times U \subset \bar{B}_{L_J}(0)\times U$ for all k, and both $\bar{B}_{L_J}(0)$ and U are compact, the sequence $\{p_k,q_k,u_k\}$ is uniformly bounded. We now, using the Bolzano–Weierstrass theorem, extract a convergent subsequence $\{p_{k_j},q_{k_j},u_{k_j}\}$ of $\{p_k,q_k,u_k\}$ where $\lim_{j\to\infty}(p_{k_j},q_{k_j},u_{k_j}) \to (p^*,q^*,u^*)$. Because U is closed we get $u^* \in U$ and Prop. 10.10 in [35] implies $(p^*,q^*) \in \partial_C J(t,x)$.

By the continuity of the Hamiltonian H:

$$\lim_{j\to\infty}(p_{k_j} + H(t_{n_{k_j}},x_{n_{k_j}},u_{k_j},q_{k_j})) = p^* + H(t,x,u^*,q^*).$$

Hence,

$$\begin{aligned}
\liminf_{n\to\infty} E_J(t_n,x_n) &= \lim_{k\to\infty} E_J(t_{n_k},x_{n_k}) = \lim_{j\to\infty} E_J(t_{n_{k_j}},x_{n_{k_j}})\\
&= \lim_{j\to\infty}(p_{k_j} + H(t_{n_{k_j}},x_{n_{k_j}},u_{k_j},q_{k_j}))\\
&= p^* + H(t,x,u^*,q^*).
\end{aligned}$$

Since $(p^*,q^*,u^*) \in \partial_C J(t,x)\times U = D^{\pm}J(t,x)\times U$, for any (t,x), by definition of E_J,

$$E_J(t,x) \leq p^* + H(t,x,u^*,q^*) = \liminf_{n\to\infty} E_J(t_n,x_n).$$

□

We now bound the optimality gap of using a controller/policy synthesized from an approximated VF based on the HJB PDE residual error.

Proposition 1. *Suppose the OCP* (1) *corresponding to tuple* $\{c,g,f,U,T\}$ *is of class L (Defn. 1). Then for any any admissible candidate value function J (Definition 3) with associated admissible controller,* π_J *Eq.* (5)*, we have that*

$$\begin{aligned}
&L(u_{\pi_J,x_0},x_0) \leq 2T\delta_1 + 2\delta_2, \qquad (12)\\
&\delta_1 = \operatorname*{ess\,sup}_{x\in B_R(0),t\in[0,T]}\left|\frac{\partial}{\partial t}J(t,x) + \inf_{u\in U}\{H(t,x,u,\nabla_x J(t,x))\}\right|,\\
&\delta_2 = \operatorname*{ess\,sup}_{x\in B_R(0)}|J(T,x) - g(x)|,
\end{aligned}$$

where L (the loss) is defined in Eq. (6)*,* $R := (1 + \|x_0\|_2)e^{\beta_f T} - 1$ *and* β_f *is as in Defn. 1.*

Note: J is Lipschitz so the derivatives, $\frac{\partial}{\partial t}J(t,x)$ and $\nabla_x J$, in δ_1 are defined a.e and hence the $\operatorname{ess\,sup}$ is well-defined.

Proof. By Lemma 2 we have that u_{π_J,x_0} solves the OCP associated with tuple $\{\tilde{c},\tilde{g},f,U,T\}$, where $\tilde{c}$ and $\tilde{g}$ are defined in terms of E_J as in Eq. (9). Denoting its associated trajectory simply as x_J (rather than $x_{u_{\pi_J,x_0},x_0}$), we have for any admissible input $u:[0,T]\to U$ that

$$\begin{aligned}
&\int_0^T \tilde{c}(s,x_J(s),u_{\pi_J,x_0}(s))ds + \tilde{g}(x_J(T)) \qquad (13)\\
&\leq \int_0^T \tilde{c}(s,x_{u,x_0}(s),u(s))ds + \tilde{g}(x_{u,x_0}(T)).
\end{aligned}$$

By substituting $\tilde{c}(t,x,u) = c(t,x,u) - E_J(t,x)$ and $\tilde{g}(x) = J(T,x)$ into Inequality (13) and re-arranging terms, we have

Authorized licensed use limited to: Arizona State University. Downloaded on October 27,2026 at 09:29:22 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2026.3727165

$$\begin{aligned}
&\int_0^T c(s, x_J(s), u_{\pi_J,x_0}(s))ds + g(x_J(T)) \qquad (14)\\
&\quad - \int_0^T c(s, x_{u,x_0}(s), u(s))ds - g(x_{u,x_0}(T))\\
&\le \int_0^T E_J(s, x_J(s)) - E_J(s, x_{u,x_0}(s))ds\\
&\quad + g(x_J(T)) - J(T, x_J(T))\\
&\quad + J(T, x_{u,x_0}(T)) - g(x_{u,x_0}(T))\\
&\le T \operatorname*{ess\,sup}_{s\in[0,T]}\{|E_J(s, x_J(s)) - E_J(s, x_{u,x_0}(s))|\}\\
&\quad + 2 \sup_{y\in B_R(0)}\{|g(y) - J(T, y)|\}\\
&\le 2T \operatorname*{ess\,sup}_{(s,y)\in[0,T]\times B_R(0)}\{|E_J(s, y)|\} + 2 \operatorname*{ess\,sup}_{y\in B_R(0)}\{|g(y) - J(T, y)|\}.
\end{aligned}$$

Here the first inequality is simply Eq. (13). To establish the second inequality, we bound the integral using the monotonicity of the Lebesgue integral (Hölder's inequality). This operation is well-defined because the mapping $t \mapsto E_J(t, x(t))$ is measurable, which follows automatically from the fact that E_J is lower semi-continuous (Lemma 3) and the trajectories are absolutely continuous. Furthermore, the essential supremum is finite because E_J is bounded, a direct consequence of the uniform boundedness of the Clarke generalized derivatives due to the Lipschitz continuity of J (Prop 10.5 [35]). The third inequality relaxes the bound by taking the supremum over a fixed spatial domain rather than the specific trajectories. This is done through the use of Lemma 4 (found in the appendix) and because $x_{u,x_0}(t)$ is contained within the closure of $B_R(0)$ for all $t \in [0, T]$ and any admissible input $u : [0, T] \to U$ (since under Defn. 1 $\|x_{u,x_0}(t)\|_2 \le (1 + \|x_0\|_2)e^{\beta_f T} - 1$ for all $t \in [0, T]$ by the Gronwall Bellman inequality, see [36]). Furthermore, since V^* and J are continuous, their essential supremum over $B_R(0)$ coincides with the standard supremum.

Since J is Lipschitz continuous, it is differentiable almost everywhere by Rademacher's theorem, and hence,

$$E_J(t, x) = \frac{\partial}{\partial t}J(t, x) + \inf_{u\in U} H(t, x, u, \nabla_x J(t, x)) \qquad (15)$$
$$\text{for } (t, x) \in [0, T] \times B_R(0) \text{ a.e,}$$
$$\text{implying} \quad \delta_1 = \operatorname*{ess\,sup}_{(s,y)\in[0,T]\times B_R(0)}\{|E_J(s, y)|\}. \qquad (16)$$

Finally, since Inequality (14) holds for all admissible u, taking the supremum over u and recognizing that the infimum with respect to admissible inputs of the cost yields the optimal value function $V^*(0, x_0)$, we obtain Eq. (12) as required. □

We next obtain a bound on the optimality gap by connecting the HJB residual to the Sobolev approximation error between the candidate and true VF. This constitutes the paper's main result, previously stated in Theorem 1.

Proof of Theorem 1. Throughout this proof we denote the set of points where a Lipschitz function, $W \in Lip([0, T]\times\mathbb{R}^n, \mathbb{R})$, is differentiable by,

$$S_W := \{(t, x) \in [0, T] \times B_R(0) : W \text{ is differentiable at } (t, x)\}.$$

From Prop. 1, Eq. (16) and $V^*(T, x) = g(x)$ it follows that

$$\begin{aligned}
L(u_{\pi_J,x_0}, x_0) &\le 2T \operatorname*{ess\,sup}_{(s,y)\in[0,T]\times B_R(0)}\{|E_J(s, y)|\} \qquad (17)\\
&\quad + 2 \operatorname*{ess\,sup}_{y\in B_R(0)}\{|V^*(T, y) - J(T, y)|\}.
\end{aligned}$$

We now split the remainder of the proof into three parts. In Part 1, we derive an upper bound for $E_J(s, y)$. In Part 2, we find a lower bound for $E_J(s, y)$. In Part 3 we combine these to bound $\operatorname{ess\,sup}_{(s,y)\in[0,T]\times B_R(0)}\{|E_J(s, y)|\}$ in Eq. (17) and hence obtain Eq. (7) to complete the proof.

Part 1 of Proof: For each $(s, y) \in S_{V^*} \cap S_J$ consider any

$$k^*_{s,y} \in \arg\inf_{u\in U}\{H(s, y, u, \nabla_x V^*(s, y))\}.$$

Such a $k^*_{s,y}$ exists for each fixed $(s, y) \in S_{V^*} \cap S_J$ by the extreme value theorem[1]. It follows from Eq. (15) that

$$\begin{aligned}
E_J(s, y) &= \frac{\partial}{\partial s}J(s, y) + \inf_{u\in U} H(s, y, u, \nabla_x J(s, y)) \qquad (18)\\
&\le \frac{\partial}{\partial s}J(s, y) + H(s, y, k^*_{s,y}, \nabla_x J(s, y)).
\end{aligned}$$

Moreover, since V^* is the viscosity solution to the HJB PDE by Lemma 1, by definition we have that

$$\frac{\partial}{\partial s}V^*(s, y) + H(s, y, k^*_{s,y}, \nabla_x V^*(s, y)) = 0. \qquad (19)$$

Combining Eqs. (18) and (19), applying Hölder's inequality ($a^\top b \le \|a\|_\infty\|b\|_1 \le \|a\|_2\|b\|_1$), using the bound $\|f(t, x, u)\|_2 \le \beta_f(1+\|x\|_2)$ (Defn. 1), and the fact $x_{u,x_0}(t) \in B_R(0)$ (by Gronwall Bellman), it follows that

$$\begin{aligned}
E_J(s, y) &\le \frac{\partial}{\partial s}J(s, y) - \frac{\partial}{\partial s}V^*(s, y)\\
&\qquad + (\nabla_x J(s, y) - \nabla_x V^*(s, y))^\top f(s, y, k^*_{s,y})\\
&\le \left|\frac{\partial}{\partial s}J(s, y) - \frac{\partial}{\partial s}V^*(s, y)\right|\\
&+ \sup_{(t,x,u)\in[0,T]\times B_R(0)\times U}\|f(t, x, u)\|_2 \sum_{i=1}^n \left|\frac{\partial}{\partial y_i}(J(s, y) - V^*(s, y))\right|\\
&\le \max\left\{1, \beta_f(1 + R)\right\}\Bigg(\left|\frac{\partial}{\partial s}J(s, y) - \frac{\partial}{\partial s}V^*(s, y)\right|\\
&\qquad + \sum_{i=1}^n \left|\frac{\partial}{\partial y_i}(J(s, y) - V^*(s, y))\right|\Bigg). \qquad (20)
\end{aligned}$$

Part 2 of Proof: If $(s, y) \in S_J \cap S_{V^*}$ and π_J satisfies Eq. (5), then by Eq. (15)

$$E_J(s, y) = \frac{\partial}{\partial s}J(s, y) + H(s, y, \pi_J(s, y), \nabla_x J(s, y)). \qquad (21)$$

Moreover, since V^* is the viscosity solution to the HJB PDE (2), by Lemma 1, we have that

$$\begin{aligned}
&\frac{\partial}{\partial s}V^*(s, y) + H(s, y, \pi_J(s, y), \nabla_x V^*(s, y)) \qquad (22)\\
&\ge \frac{\partial}{\partial s}V^*(s, y) + \inf_{u\in U} H(s, y, u, \nabla_x V^*(s, y)) = 0.
\end{aligned}$$

Combining Eqs (21) and (22) it follows by a similar argument to Eq. (20) that,

$$\begin{aligned}
&E_J(s, y)\\
&\ge -\max\left\{1, \beta_f(1 + R)\right\}\Bigg(\left|\frac{\partial}{\partial s}J(s, y) - \frac{\partial}{\partial s}V^*(s, y)\right|\\
&\qquad + \sum_{i=1}^n \left|\frac{\partial}{\partial y_i}(J(s, y) - V^*(s, y))\right|\Bigg). \qquad (23)
\end{aligned}$$

[1]This follows because $U \subset \mathbb{R}^m$ is compact, c and f are continuous, and $\nabla_x V^*$ is independent of u and bounded by Rademacher's Theorem.

Authorized licensed use limited to: Arizona State University. Downloaded on October 27,2026 at 09:29:22 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2026.3727165

Part 3 of Proof: From Eqs (20) and (23) it follows for any $(s,y) \in S_J \cap S_{V^*}$ we have

$$|E_J(s,y)| \leq \max\left\{1, \beta_f(1+R)\right\}\left(\left|\frac{\partial}{\partial s}(J(s,y)-V^*(s,y))\right| + \sum_{i=1}^{n}\left|\frac{\partial}{\partial y_i}(J(s,y)-V^*(s,y))\right|\right).$$

and hence since $([0,T]\times B_R(0))/(S_J \cap S_{V^*})$ has zero measure it follows that,

$$\begin{aligned}\operatorname*{ess\,sup}_{(s,y)\in[0,T]\times B_R(0)}\{|E_J(s,y)|\} &\leq \max\left\{1, \beta_f(1+R)\right\} \\ &\times \operatorname*{ess\,sup}_{(s,y)\in[0,T]\times B_R(0)}\left(\left|\frac{\partial}{\partial s}(J(s,y)-V^*(s,y))\right| + \sum_{i=1}^{n}\left|\frac{\partial}{\partial y_i}(J(s,y)-V^*(s,y))\right|\right). \quad (24)\end{aligned}$$

Combining Inequalities (17) and (24) we obtain (7), completing the proof.

$$\begin{aligned}L(u_{\pi_J,x_0},x_0) &\leq 2T\max\left\{1,\beta_f(1+R)\right\} \\ &\times \operatorname*{ess\,sup}_{(s,y)\in[0,T]\times B_R(0)}\left(\left|\frac{\partial}{\partial s}(J(s,y)-V^*(s,y))\right| + \sum_{i=1}^{n}\left|\frac{\partial}{\partial y_i}(J(s,y)-V^*(s,y))\right|\right) \\ &+ 2\sup_{y\in B_R(0)}\{|V^*(T,y)-J(T,y)|\} \\ &\leq C\|J-V^*\|_{W^{1,\infty}([0,T]\times B_R(0))},\end{aligned}$$

where $C := 2\max\{1, T, T\beta_f(1+R)\}$. □

IV. Interpretation of Performance Bounds

Prop. 1 shows that direct minimization of the HJB residual is ideal for controller synthesis. In practice, however, directly minimizing this residual is often computationally intractable due to the residual itself containing an implicit inner optimization problem ($\inf_{u\in U}$). Consequently, many numerical methods must construct candidate VFs via other means. Theorem 1 bridges this gap: if a numerical scheme approximates the true VF, securing a small error in the $W^{1,\infty}$ norm provides a guarantee that the resulting controller will perform well. Later, we will see in Example 2 that a small error in the weaker $W^{1,p}$ norm for $p < \infty$ alone (without HJB residual minimization) is not strong enough to produce good controller performance.

Importantly, while a small $W^{1,\infty}$ error is a sufficient condition for good controller performance, Example 1, presented next, demonstrates that it is not strictly necessary. It is possible to construct a candidate VF with an arbitrarily large $W^{1,\infty}$ error that still yields a arbitrarily near optimal controller. This occurs when the gradient approximation error is orthogonal to the system's vector field, effectively hiding the error from the Hamiltonian. In such cases, the HJB residual remains zero despite the divergent gradient. However, generic numerical approximation schemes do not generally restrict gradient errors exclusively to directions orthogonal to the system's vector field. Therefore, Theorem 1 provides a universal performance guarantee, bounding controller performance by the Sobolev VF error.

Example 1. *Consider an OCP with tuple $\{c, g, f, U, T\}$ defined by $c \equiv 0$, $g(x) = \|x\|_2^2$, $U = [-1,1]$, $T = 1$ and $f(x,u) = [1,-1]^\top u\|x\|_2^2$. Let V^* be the true VF and for any $\epsilon > 0$, consider the admissible candidate value function $J_\varepsilon(t,x) := V^*(t,x) + \sqrt{\varepsilon}\sin((x_1+x_2)/\varepsilon)$. First, we observe that $\nabla_x J_\varepsilon(t,x) = \nabla_x V^*(t,x) + \varepsilon^{-1/2}\cos((x_1+x_2)/\varepsilon)[1,1]^\top$, so $\lim_{\epsilon\to 0}\|\nabla_x J_\varepsilon - \nabla_x V^*(t,x)\| = \infty$ – implying the error in the $W_{1,\infty}$ norm becomes arbitrarily large.*

Despite this divergence, the HJB residual meanwhile becomes arbitrarily small. To see this, we note that $\frac{\partial}{\partial t}J_\varepsilon(t,x) = \frac{\partial}{\partial t}V^*(t,x)$ and $\nabla_x J_\varepsilon^\top f = \nabla_x V^{*\top} f + \varepsilon^{-1/2}\cos((x_1+x_2)/\varepsilon)[1,1][1,-1]^\top u\|x\|_2^2 = \nabla_x V^{*\top} f$. We conclude that

$$\begin{aligned}&\left|\frac{\partial}{\partial t}J_\epsilon(t,x) + \inf_{u\in U}\{H(t,x,u,\nabla_x J_\epsilon(t,x))\}\right| \\ &= \left|\frac{\partial}{\partial t}V^*(t,x) + \inf_{u\in U}\{H(t,x,u,\nabla_x V^*(t,x))\}\right| = 0.\end{aligned}$$

Furthermore, $\lim_{\epsilon\to 0}\|J_\varepsilon(t,\cdot) - V^*(t,\cdot)\|_{L^\infty([0,T],\mathbb{R}^n)} = 0$, implies $\lim_{\varepsilon\to 0} J_\varepsilon(T,x) = V^*(T,x) = g(x)$. Applying Prop. 1, we conclude that $\lim_{\epsilon\to 0} L(u_{\pi_{J_\epsilon}},x_0,x_0) = 0$ – implying that the performance of the sequence of resulting controllers converges to optimality.

Finally, we note that the construction of the bounds in Prop. 1 and Thm. 1 use the Gronwall Bellman inequality to bound growth of the solution – resulting in an exponential dependency in time-horizon, T. Such exponentials may be avoided if a bound on the forward reachable set is known a priori. Furthermore, because the suboptimality of any controller cannot exceed the gap between the worst-case (maximum cost) VF, we denote as $\mathcal{V}^*$, and the optimal VF, V^*, we can refine Eq. (7):

$$L(u_{\pi_J,x_0},x_0) \leq \min\left\{C\|J-V^*\|_{W^{1,\infty}([0,T]\times B_R(0))}, \mathcal{V}^*(0,x_0) - V^*(0,x_0)\right\}. \quad (25)$$

V. Numerical Examples

Example 2. *Let us consider the OCP associated with the tuple $\{c, g, f, U, T\}$, where $c(t,x,u) \equiv 0$, $g(x) = x$, $f(t,x,u) = xu$, $U = [-1,1]$, and $T = 1$. It was shown in [6] that the VF associated with this tuple can be found analytically:*

$$V^*(t,x) = \begin{cases}\exp(t-1)x & \text{if } x > 0,\\ \exp(1-t)x & \text{if } x < 0,\\ 0 & \text{if } x = 0.\end{cases} \quad (26)$$

Let us now consider the following candidate VFs:

$$\begin{aligned}\hat{V}_1(t,x) &= V^*(t,x) - 100\varepsilon^\delta\eta((x-0.75)/\varepsilon), \quad (27)\\ \hat{V}_2(t,x) &= V^*(t,x) + \varepsilon^2\sin(10^6 x),\end{aligned}$$

where $\eta(x) := \begin{cases}\exp\left(\frac{1}{x^2-1}\right) & \text{if } x\in(-1,1)\\ 0 & \text{otherwise.}\end{cases}$ is the standard bump function from mollification theory and $\delta \in (0,1)$. It is well known that $\eta \in C^\infty(\mathbb{R},\mathbb{R})$ [32] with compact support; hence $C_{0,p} := \|\eta\|_{L^p(\mathbb{R})} < \infty$ and $C_{1,p} := \|\eta'\|_{L^p(\mathbb{R})} < \infty$ for all $p \geq 1$, including $p = \infty$. Therefore[2]:

[2]See full derivation details on Arxiv [37]

Authorized licensed use limited to: Arizona State University. Downloaded on October 27,2026 at 09:29:22 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2026.3727165

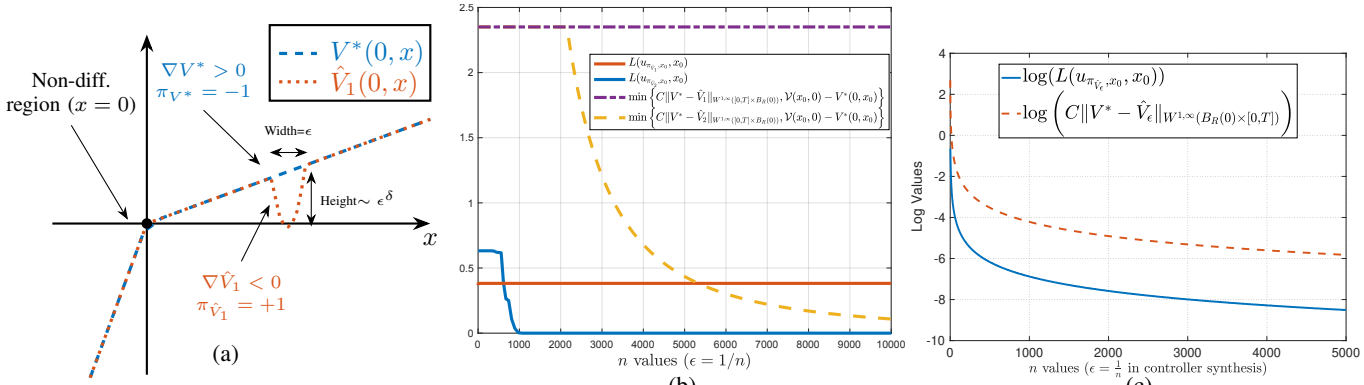


Fig. 1: (a) Graph showing plots of V^* and $\hat{V}_1$ from Eqs (26) and (27) respectively. (b) Graph showing numerically calculated performance gaps, $L(u_{\pi_{\hat{V}_1}},x_0,x_0)$ and $L(u_{\pi_{\hat{V}_2}},x_0,x_0)$, associated with the approximate VFs given in Eq. (27) for decreasing $\varepsilon > 0$, along with the theoretical bounds for these performance gaps given in Eqs (29) and (30). (c) Numerically calculated logarithms of performance gap, $L(u_{\pi_{\hat{V}_\varepsilon}},x_0,x_0)$, associated with the approximate VFs given in Eq. (33) for decreasing $\varepsilon > 0$, along with the logarithm of the theoretical performance gap bound given in Eq. (35).

$$\|V^*-\hat{V}_1\|_{W^{1,p}([0,T]\times\mathbb{R})}\leq 100(\varepsilon^{\delta+1/p}C_{0,p}+\varepsilon^{\delta-1+1/p}C_{1,p}),$$
$$\|V^*-\hat{V}_2\|_{W^{1,\infty}([0,T]\times\mathbb{R})} = (1+10^6)\varepsilon^2. \tag{28}$$

Applying the performance bound in Eq. (25) *for* $x_0 = 1$ *and* $\beta_f = 1$*, we deduce:*

$$L(u_{\pi_{\hat{V}_1}},x_0,x_0)\leq \min\left\{400e\left(C_{0,\infty}\varepsilon^\delta+\frac{C_{1,\infty}}{\varepsilon^{1-\delta}}\right),(e-e^{-1})\right\}, \tag{29}$$

$$L(u_{\pi_{\hat{V}_2}},x_0,x_0) \leq \min\left\{4e(1+10^6)\varepsilon^2,(e-e^{-1})\right\}. \tag{30}$$

Observe via Eq. (28) *that for any finite* $p < \infty$*, we can select* $\delta \in (1-1/p,1)$ *such that* $\hat{V}_1 \to V^*$ *in the* $W^{1,p}$ *norm. Despite this, the synthesized controller fails to yield trajectories that converge to the optimal cost. This discrepancy arises because* $\hat{V}_1$ *does not converge to* V^* *in the stronger* $W^{1,\infty}$ *norm. As illustrated in Fig. 1a, for every* $\varepsilon > 0$ *and* $\delta \in (0,1)$*, the bump function perturbation creates a local minimum in* $\hat{V}_1$ *that acts as a "trapping region." Consequently, the induced controller stabilizes the trajectory around the center of the bump* ($x = 0.75$) *rather than optimally driving the state to minimize the terminal cost* $g(x) = x$*.*

This phenomenon is confirmed numerically in Fig. 1b for $\delta = 0.9$*, where the cost does not converge to zero as* $\varepsilon \to 0$*. The right-hand side of the bound in Eq.* (29) *remains constant at* $e - e^{-1} \approx 2.35$ *as* $\varepsilon \to 0$*. While this provides a valid upper bound, it indicates that the controller is not guaranteed to achieve optimal performance.*

Conversely, Eq. (30) *demonstrates that the controller synthesized by* $\hat{V}_2$ *will converge to the optimal cost, since* $L(u_{\pi_{\hat{V}_2}},x_0,x_0) \to 0$ *as* $\varepsilon \to 0$*. In Fig. 1b, we observe that* $u_{\pi_{\hat{V}_2}},x_0$ *numerically converges to optimal performance at a faster rate than the theoretical error bound derived in Eq.* (30)*.*

Example 3. *Let us consider the OCP associated with the following tuple* $\{c,g,f,U,T\}$*, where* $c(t,x,u) = (1+2(T-t))x^2 + 2(T-t)|x|$*,* $g(x) = 0$*,* $f(t,x,u) = -x+u$*,* $U = [-1,1]$*, and* $T = 1$*. Since,* $\inf_{u\in[-1,1]}\{\nabla_x V(t,x)u\} = -|\nabla_x V(t,x)|$ *for any differentiable* V *it follows that the associated HJB PDE is*

$$\frac{\partial}{\partial t}V(t,x)+(1+2(T-t))x^2+2(T-t)|x| - \nabla_x V(t,x)x - |\nabla_x V(t,x)| = 0. \tag{31}$$

We can then deduce $V(t,x) = x^2(T-t)$ *is the VF for this OCP by substitution into Eq.* (31)*. Hence the optimal controller is given by*

$$\pi^*(t,x) = -\operatorname{sign}(2x(T-t)) = \begin{cases} -1 \text{ if } x > 0, \\ 0 \text{ if } x = 0, \\ 1 \text{ if } x < 0. \end{cases} \tag{32}$$

Let us next consider a sequence of approximate VFs

$$\hat{V}_\varepsilon(t,x) = x^2(T-t)+\varepsilon x. \tag{33}$$

The candidate VF, $\hat{V}_\varepsilon$*, yields the following controller*

$$\hat{\pi}_{\hat{V}_\varepsilon}(t,x) = -\operatorname{sign}(2x(T-t)+\varepsilon) = \begin{cases} -1 \text{ if } x > \frac{-\varepsilon}{2(T-t)}, \\ 0 \text{ if } x = \frac{-\varepsilon}{2(T-t)}, \\ 1 \text{ if } x < \frac{-\varepsilon}{2(T-t)}. \end{cases} \tag{34}$$

For $x_0 = 0$ *we have that* $R := (1+\|x_0\|_2)e^{\beta_f T} - 1 = e-1$*. Hence,* $\|V^* - \hat{V}_\varepsilon\|_{W^{1,\infty}([0,T]\times B_R(0))} \leq (R+1)\varepsilon \to 0$ *but* $\lim_{\varepsilon\to 0}\|\pi^* - \hat{\pi}_{\hat{V}_\varepsilon}\|_{L^\infty([0,T]\times B_R(0),\mathbb{R})} \neq 0$ *since* $\pi^*(t,0) = 0$ *but* $\hat{\pi}_{\hat{V}_\varepsilon}(t,0) = -1$ *for all* $\varepsilon > 0$*. Despite the fact that the controller* $\hat{\pi}_{\hat{V}_\varepsilon}$ *does not converge to the optimal controller uniformly, Theorem 1 still shows that the associated performance/cost of applying the controller converges to zero. This is shown numerically in Fig. 1c. The figure also shows the theoretical performance bound found by substituting* $R = e-1$*,* $\beta_f = 1$ *into Eq.* (7) *to get*

$$L(u_{\pi_{\hat{V}_\varepsilon}},x_0,x_0) \leq 2e^2\varepsilon. \tag{35}$$

VI. Conclusion

We have shown that when approximate value functions are used to synthesize controllers, the suboptimality of the resulting controller is bounded by the L^∞ norm of the HJB residual, which is, in turn, bounded by the Sobolev $W^{1,\infty}$ norm of the error between the candidate value function and true value function. Consequently, numerical methods for solving the HJB equation which converge in the $W^{1,\infty}$ norm will yield controllers which minimize the associated cost arbitrarily close to the true optimum. In contrast, we also showed that convergence in $W^{1,p}$ for finite p does not necessarily guarantee optimal performance. These findings suggest that future work should prioritize minimizing the $W^{1,\infty}$ distance to the true value function or L^∞ HJB residual error over other metrics and objectives.

Authorized licensed use limited to: Arizona State University. Downloaded on October 27,2026 at 09:29:22 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Automatic Control. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TAC.2026.3727165

REFERENCES

[1] M. G. Crandall and P.-L. Lions, "Viscosity solutions of hamilton-jacobi equations," *Transactions of the American mathematical society*, vol. 277, no. 1, pp. 1–42, 1983.

[2] D. Gallistl, T. Sprekeler, and E. Süli, "Mixed finite element approximation of periodic Hamilton–Jacobi–Bellman problems with application to numerical homogenization," *arXiv preprint arXiv:2010.01647*, 2020.

[3] D. Kalise and K. Kunisch, "Polynomial approximation of high-dimensional Hamilton–Jacobi–Bellman equations and applications to feedback control of semilinear parabolic PDEs," *SIAM Journal on Scientific Computing*, vol. 40, no. 2, pp. A629–A652, 2018.

[4] Y. Achdou, F. Camilli, and I. Capuzzo Dolcetta, "Homogenization of Hamilton–Jacobi equations: numerical methods," *Mathematical models and methods in applied sciences*, vol. 18, no. 07, pp. 1115–1143, 2008.

[5] M. Falcone and R. Ferretti, *Semi-Lagrangian approximation schemes for linear and Hamilton—Jacobi equations*. SIAM, 2013.

[6] D. Liberzon, *Calculus of variations and optimal control theory: a concise introduction*. Princeton University Press, 2011.

[7] W. M. McEneaney, "A curse-of-dimensionality-free numerical method for solution of certain HJB PDEs," *SIAM journal on Control and Optimization*, vol. 46, no. 4, pp. 1239–1276, 2007.

[8] W. Kang and L. C. Wilcox, "Mitigating the curse of dimensionality: sparse grid characteristics method for optimal feedback control and HJB equations," *Computational Optimization and Applications*, vol. 68, no. 2, pp. 289–315, 2017.

[9] K. Kunisch, S. Volkwein, and L. Xie, "HJB-POD-based feedback design for the optimal control of evolution problems," *SIAM Journal on Applied Dynamical Systems*, vol. 3, no. 4, pp. 701–722, 2004.

[10] M. Abu-Khalaf and F. L. Lewis, "Nearly optimal control laws for nonlinear systems with saturating actuators using a neural network HJB approach," *Automatica*, vol. 41, no. 5, pp. 779–791, 2005.

[11] Y. Tassa and T. Erez, "Least squares solutions of the hjb equation with neural network value-function approximators," *IEEE transactions on neural networks*, vol. 18, no. 4, pp. 1031–1041, 2007.

[12] D. Vrabie and F. Lewis, "Neural network approach to continuous-time direct adaptive optimal control for partially unknown nonlinear systems," *Neural Networks*, vol. 22, no. 3, pp. 237–246, 2009.

[13] Y. Meng, R. Zhou, A. Mukherjee, M. Fitzsimmons, C. Song, and J. Liu, "Physics-informed neural network policy iteration: Algorithms, convergence, and verification," *arXiv preprint arXiv:2402.10119*, 2024.

[14] G. Xiao and H. Zhang, "Convergence analysis of value iteration adaptive dynamic programming for continuous-time nonlinear systems," *IEEE Transactions on Cybernetics*, vol. 54, no. 3, pp. 1639–1649, 2023.

[15] Y. Liu, L. Cai, Y. Chen, and B. Wang, "Physics-informed neural networks based on adaptive weighted loss functions for hamilton-jacobi equations," *J. Mathematical biosciences and engineering: MBE*, vol. 19, no. 12, 2022.

[16] A. Shilova, T. Delliaux, P. Preux, and B. Raffin, "Revisiting continuous-time reinforcement learning. a study of hjb solvers based on pinns and fems," in *Sixteenth European Workshop on Reinforcement Learning*, 2023.

[17] A. Mukherjee and J. Liu, "Bridging physics-informed neural networks with reinforcement learning: Hamilton-jacobi-bellman proximal policy optimization (hjbppo)," *arXiv preprint arXiv:2302.00237*, 2023.

[18] C. Wang, S. Li, D. He, and L. Wang, "Is L^2 physics informed loss always suitable for training physics informed neural network?," *Advances in Neural Information Processing Systems*, vol. 35, pp. 8278–8290, 2022.

[19] A. M. Ribeiro, A. R. Fioravanti, A. Moutinho, and E. C. de Paiva, "Control design based on sum of squares programming for non-affine in input systems," in *2020 IEEE 6th International Conference on Control Science and Systems Engineering (ICCSSE)*, pp. 130–135, IEEE, 2020.

[20] Y. Jiang and Z.-P. Jiang, "Global adaptive dynamic programming for continuous-time nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 60, no. 11, pp. 2917–2929, 2015.

[21] S. Baldi, P. A. Ioannou, and E. B. Kosmatopoulos, "A scalable iterative convex design for nonlinear systems," in *2012 American Control Conference (ACC)*, pp. 979–984, IEEE, 2012.

[22] S. Baldi, G. Valmorbida, A. Papachristodoulou, and E. B. Kosmatopoulos, "Piecewise polynomial policy iterations for synthesis of optimal control laws in input-saturated systems," in *2015 American Control Conference (ACC)*, pp. 2850–2855, IEEE, 2015.

[23] Y. Zhu, D. Zhao, X. Yang, and Q. Zhang, "Policy iteration for H_∞ optimal control of polynomial nonlinear systems via sum of squares programming," *IEEE transactions on cybernetics*, vol. 48, no. 2, pp. 500–509, 2017.

[24] Y. P. Leong, M. B. Horowitz, and J. W. Burdick, "Optimal controller synthesis for nonlinear dynamical systems," *arXiv preprint arXiv: 1410.0405*, 2014.

[25] T. Cunis, J.-P. Condomines, and L. Burlion, "Sum-of-squares flight control synthesis for deep-stall recovery," *Journal of Guidance, Control, and Dynamics*, pp. 1–14, 2020.

[26] M. Korda, D. Henrion, and C. N. Jones, "Controller design and value function approximation for nonlinear dynamical systems," *Automatica*, vol. 67, pp. 54–66, 2016.

[27] P. Zhao, S. Mohan, and R. Vasudevan, "Control synthesis for nonlinear optimal control via convex relaxations," in *2017 American Control Conference (ACC)*, pp. 2654–2661, IEEE, 2017.

[28] Y. Chen, M. Ahmadi, and A. D. Ames, "Optimal safe controller synthesis: A density function approach," in *2020 American Control Conference (ACC)*, pp. 5407–5412, IEEE, 2020.

[29] D. P. Bertsekas and J. N. Tsitsiklis, "Neuro-dynamic programming: an overview," in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 1, pp. 560–564, IEEE, 1995.

[30] S. P. Singh and R. C. Yee, "An upper bound on the loss from approximate optimal-value functions," *Machine Learning*, vol. 16, pp. 227–233, 1994.

[31] H. K. Khalil, *Nonlinear Systems*. Upper Saddle River, NJ: Prentice Hall, 3rd ed., 2002.

[32] L. C. Evans, *Partial differential equations*, vol. 19. American Mathematical Society, 2010.

[33] X. Y. Zhou, "Verification theorems within the framework of viscosity solutions," *Journal of mathematical analysis and applications*, vol. 177, no. 1, pp. 208–225, 1993.

[34] M. Bardi, I. C. Dolcetta, *et al.*, *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations*, vol. 12. Springer, 1997.

[35] F. Clarke, *Functional analysis, calculus of variations and optimal control*, vol. 264. Springer, 2013.

[36] J.-B. Caillau, R. Ferretti, E. Trélat, and H. Zidani, "An algorithmic guide for finite-dimensional optimal control problems," in *Handbook of Numerical Analysis*, vol. 24, pp. 559–626, Elsevier, 2023.

[37] M. Jones and M. Peet, "Bounding the error of value functions in sobolev norm yields bounds on suboptimality of controller performance," *arXiv preprint arXiv:2502.15421*, 2026.

VII. APPENDIX

Lemma 4. *Let $E : [0,T] \times \mathbb{R}^n \to \mathbb{R}$ be a lower semicontinuous function. For any measurable function $x : [0,T] \to B_R(0) \subset \mathbb{R}^n$, for some $R > 0$, the following inequality holds:*

$$\operatorname*{ess\,sup}_{t\in[0,T]} E(t, x(t)) \leq \operatorname*{ess\,sup}_{(t,y)\in[0,T]\times B_R(0)} E(t, y). \tag{36}$$

Proof. Since E is lower semicontinuous, its strict superlevel sets are open, meaning any exceedance of the essential supremum along $(t, x(t))$ would contradictorily require an exceedance on a non-empty open euclidean set which by definition has strictly positive Lebesgue measure (see Arxiv version for more details [37]). □

Lemma 5. *Consider functions $J \in Lip([0,T] \times \mathbb{R}^n, \mathbb{R})$ and $x \in AC([0,T], \mathbb{R}^n)$ satisfying $\dot{x}(t) = f(t, x(t))$ almost everywhere on $[0,T]$. Then, for almost every $t \in [0,T]$ and any $(p,q) \in D^{\pm}J(t, x(t))$, we have*

$$\frac{d}{dt}J(t, x(t)) = p + q^\top f(t, x(t)). \tag{37}$$

Proof. The composite function $t \mapsto J(t, x(t))$ is absolutely continuous, guaranteeing its classical derivative exists almost everywhere. At points of differentiability, the time-superdifferential reduces to the singleton $\{\frac{d}{dt}J(t, x(t))\}$ [34, Lemma 1.8], which, combined with the chain rule inclusion for superdifferentials [34, Proposition 2.7], directly implies the desired equality (see Arxiv version for more details [37]). □

ACKNOWLEDGEMENTS

This research was supported by the NSF Grant 2606245.

Authorized licensed use limited to: Arizona State University. Downloaded on October 27,2026 at 09:29:22 UTC from IEEE Xplore. Restrictions apply.
© 2026 IEEE. All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.