Analysis of extremum value theorems for function spaces in optimal control under numerical uncertainty

Pavel Osinenko[a], Stefan Streif[a]

[a]*Automatic Control and System Dynamics Lab; Technische Universität Chemnitz, 09107 Chemnitz, Germany*

Abstract

The extremum value theorem for function spaces plays the central role in optimal control. It is known that computation of optimal control actions and policies is often prone to numerical errors which may be related to computability issues. The current work addresses a version of the extremum value theorem for function spaces under explicit consideration of numerical uncertainties. It is shown that certain function spaces are bounded in a suitable sense, i. e., they admit finite approximations up to an arbitrary precision. The proof of this fact is constructive in the sense that it explicitly builds the approximating functions. Consequently, existence of approximate extremal functions is shown. Applicability of the theorem is investigated for finite–horizon optimal control, dynamic programming and adaptive dynamic programming. Some possible computability issues of the extremum value theorem in optimal control are shown on counterexamples.

Keywords: Extremum, Approximate, Optimal Control, Dynamic Programming

1. Introduction

Optimal control represents an important part of control theory. Typically, one seeks for an optimal function over a state space (also called *control policy*) so as to minimize a given cost functional. It is, however, not in general possible to compute optimizing control policies exactly due to limitations of numerical procedures which may have certain effects on the system behavior. The current work shows how, under some mild and practicable assumptions, approximate optimal control policies can still be explicitly computed. The proofs are done constructively, i. e., they entail certain ways of computing the objects in question. Constructive results are not unusual in control engineering and are often desired: for instance, Banaschewski & Mulvey [3] gave a constructive proof of the Stone-Weierstrass theorem, which is used in a number of applications to effectively find approximations to specific functions in a suitable basis. The famous Sontag's formula [21] is the core of the Sontag's constructive proof of the Artstein's theorem on nonlinear stabilization. Sepulchre, Jankovic & Kokotovic [20] developed this methodology further into a vast variety of constructive methods of finding specific stabilizing controllers.

Going back to the problem of optimal control and related effects which may occur due to numerical uncertainty, consider the following simple example of a discrete-time system whose dynamical behavior is switched by a binary decision variable u:

$$x_{k+1} = \begin{cases} (\frac{1}{2} + b)x_k, & u_k = 1, \\ (\frac{1}{2} + c)x_k, & u_k = -1, \end{cases} \quad x_0 = 1, u_k \in \{1, -1\}, \quad (1)$$

where b, c are real numbers which may, e. g., represent some physical quantities. Let an infinite-horizon cost function be defined as:

$$\min_{\{u_k\}_k} \quad J = \sum_{k=0}^{\infty} x_k^2.$$

Suppose, for the sake of the example, that one of the numbers b, c is zero and the other one is positive. Then, by the virtue of the system dynamics (1), the

optimal control policy is $u^* = \{1, 1, 1, \dots\}$ and the corresponding optimal state sequence is $x^* = \{1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots\}$. In this case, the optimal cost is $J^* = 2$. Thus, if one could find the optimal control policy, i. e., the optimal control action at each time step, then either $J^* = \frac{2}{1-2c} = 2$ or $J^* = \frac{2}{1-2b} = 2$ by the geometric sum and hence either $b = 0$ or alternatively $c = 0$. However, in practice, there may occur a numerical uncertainty between the exact values of b, c and their representations in a computational device, usually as rational numbers. One of the possible simple ways to consider these representations of b, c is in the form of Cauchy sequences $\{b(n)\}_n, \{c(n)\}_n$ which are regular in the following sense:

$$\forall n, m \in \mathbb{N}$$
$$|b(n) - b(m)| \leq \tfrac{1}{n} + \tfrac{1}{m},$$
$$|c(n) - c(m)| \leq \tfrac{1}{n} + \tfrac{1}{m}.$$

In practice, the system (1) may contain some particular approximations $b(n'), c(n'), n' \in \mathbb{N}$ where n' is the precision of the computational device. In the current work, all the proofs are done by working directly with the representations of real numbers which helps address numerical uncertainty. The said rational approximations may well come from, e. g., a measurement, which always has a finite precision, or from some computational algorithm, such as model identification. Therefore, to computationally check whether $b = 0$ or alternatively $c = 0$, approximations $b(n), c(n)$ for all $n \in \mathbb{N}$ must be compared. Such an unbounded search is, however, not technically possible. Therefore, different optimal control policies might result depending on precision – in this case, a particular number n'.

The same issue may appear when minimizing, e. g., the following particular cost function:

$$J(u) = \min\{u^2 + b, (u - 1)^2 + c\}.$$

By the virtue of the numbers b and c as described above, it follows that $\min J = 0$. However, if an optimal control action u^* could be computed exactly, such that $J(u^*) = \min J$, then either $u^* = 0$ or $u^* = 1$. It would be equiva-

lent to deciding whether b or c is exactly zero which is not always technically possible. This has been typically demonstrated in simple counter-examples [7], whose more detailed description may be found in Appendix. Particular examples of peculiar phenomena related to numerical uncertainty and floating-point arithmetic may also be found in [18].

As shown in the example above, optimality in general may fail to be achieved depending on the representation of system parameters. To address these issues, the present work seeks to show existence of optimal control in an approximate format by explicitly considering numerical uncertainty. The proofs are done constructively and in the setting of [8] since it offers convenient tools for keeping track of the number representations. The details are given in the next section.

It should be noted that, classically, the extremum value theorem states the following:

Theorem 1. *If a function f is continuous on a compact interval $[a, b]$, then there exist $x, y \in [a, b]$ such that $f(x) = \sup f$ and $f(y) = \inf f$.*

Some constructive approaches to Theorem 1 were addressed, e.g., in [5, 12] with additional assumptions on the function f. These assumptions, are, however, not always easy to verify practically, especially when one wants to apply the theorem to function spaces. Instead of strengthening the conditions of the theorem, an approximate format is considered in the present work which is sufficient for practical applications of optimal control. To achieve this, it is shown, that certain function spaces admit finite approximations. The proof is based on constructing finite approximations explicitly.

Another consequence of the new results shows also what at best can be achieved in general when addressing optimal control. The major implication is a theoretical limit at which any numerical algorithm may perform. Whereas exact optimal control policies are not achievable in general, the new result demonstrates principal possibility of computing approximate optimal control policies up to prescribed accuracy provided that the optimization problem satisfies certain conditions which are, as will be shown in the case study of Section

4, practicable. The next section discusses the important preliminaries needed to prove the main theorem of Section 3.

2. Preliminaries

In this section, the definitions and some basic technical results necessary for derivation of the new approximate extremum value theorem for function spaces are recalled. For a comprehensive description, refer, for example, to Bishop & Bridges [8], Bridges & Richman [10], Bridges & Vita [11], Ye [27], Schwichtenberg [19]. A **real** number x in the current work well be characterized by its rational approximations in the following *regular* Cauchy sequence format:

$$\forall n, m \in \mathbb{N}, |x(n) - x(m)| \leq \frac{1}{n} + \frac{1}{m},$$

where $x(n)$ is some operation that produces the $n-$th rational approximation to x. The inequalities on real numbers are defined as follows:

$$(x \leq y) \triangleq \forall n \in \mathbb{N}, x(n) \leq y(n) + \frac{2}{n},$$
$$(x < y) \triangleq \exists n \in \mathbb{N}, x(n) < y(n) - \frac{2}{n}.$$

In the second definition, the number n is also called a *witness*. Such objects are said to certify the respective formulas. They can be used by computational devices. Further, the maximum of two real numbers is defined as follows: $\max\{x, y\}(n) \triangleq \max\{x(n), y(n)\}$. The basic properties of it can be proven, but in general it cannot be decided whether $\max\{x, y\} = x$ or $\max\{x, y\} = y$. However, the following simple technical result can be easily proven:

Proposition 2. *For any two real numbers x, y satisfying $x \leq y$, it follows that* $\max\{x, y\} = y$.

Proof. It suffices to show that

$$\forall n, |\max\{x(n), y(n)\} - y(n)| \leq \frac{2}{n}.$$

It follows that $\forall n \in \mathbb{N}, \max\{x(n), y(n)\} \leq y(n) + \frac{2}{n}$ from the condition of the proposition which implies $\forall n \in \mathbb{N}.x(n) \leq y(n) + \frac{2}{n}$. On the other hand, $\forall n \in \mathbb{N}, \max\{x(n), y(n)\} \geq y(n) \geq y(n) - \frac{2}{n}$, and the result follows. $\square$

Remark 3. Properties of the minimum are derived similarly.

A metric space (X, ρ) is a set X together with an operation $\rho : X \times X \to \mathbb{R}$ that satisfies the usual axioms of a metric. A metric space (X, ρ) is **totally bounded** if for all natural k, there exists a finite set of unequal points $\{x_1, \ldots, x_n\} \subset X$ such that for any $x \in X$, there exists an $x_i \in \{x_1, \ldots, , x_n\}$ with $\rho(x, x_i) \leq \frac{1}{k}$. Such a finite set is also called a $\frac{1}{k}$–**approximation** to X. A subset A of a metric space (X, ρ) is **located** if it is non-empty and for any x in X, the metric $\rho(x, A) \triangleq \inf\{\rho(x, y) : y \in A\}$ can be effectively computed. A totally bounded subset of a metric space is also located (see Proposition 2.2.9 in [11]). A metric between two subsets A and B is defined as $\rho(A, B) \triangleq \inf\{\rho(x, y) : x \in A, y \in B\}$.

A (uniformly continuous) **function** from a totally bounded metric space (X, ρ) to a metric space (Y, σ) is a pair consisting of an operation $x \mapsto f(x), x \in X$ and an operation $\omega : \mathbb{Q} \to \mathbb{Q}$ called **modulus of** (uniform) **continuity** such that:

$$\forall \varepsilon \in \mathbb{Q}, \forall x, y \in X, \rho(x, y) \leq \omega(\varepsilon) \implies \sigma(f(x), f(y)) \leq \varepsilon.$$

A function is **Lipschitz continuous** if $\forall x, y \in X, \sigma(f(x), f(y)) \leq L \cdot \rho(x, y)$ for some rational $L > 0$. The set $\mathcal{F}$ of (all) uniformly continuous functions from a totally bounded metric space (X, ρ) to a metric space (Y, σ) together with the metric $\tau(f, g) \triangleq \sup_{x \in X} \sigma(f(x), f(y))$ for any $f, g \in \mathcal{F}$ is called the **function space** from X to Y. A function space $\mathcal{F}$ is **equicontinuous** if there exists a common modulus of continuity for all f in $\mathcal{F}$. Further, $\mathcal{F}$ is the space of uniformly Lipschitz and uniformly bounded functions whenever there exists a common Lipschitz constant L and respectively a common bound $K \in \mathbb{Q}, K > 0$ such that $\forall f \in \mathcal{F}, \|f\| \leq K$. A uniformly continuous **functional** F on a totally bounded function space $\mathcal{F}$ is an operation $f \mapsto F[f] \in \mathbb{R}$ with a modulus of continuity α such that $|F[f] - F[g]| \leq \frac{1}{k}$ whenever $\tau(f, g) \leq \alpha\left(\frac{1}{k}\right)$.

The symbol x^i denotes the i−th coordinate of the point x in $\mathbb{R}^n$. The two common norms on $\mathbb{R}^n$ are the d_2–norm: $\|x\|_2 = \left(\sum_{i=1}^n \left(x^i\right)^2\right)^{1/2}$, and the d_∞−norm (or maximum norm): $\|x\|_\infty = \max_i |x^i|$. The subscripts "2" and "∞"

may be omitted whenever the type of the norm is clear from the context. The corresponding metric between any two points x, y is defined as $\|x - y\|$. For the metrics d_2 and d_∞, the following holds: $\| \bullet \|_\infty \leq \| \bullet \|_2 \leq \sqrt{n} \| \bullet \|_\infty$. A real space $\mathbb{R}^n$ with the metric d_∞ will be also denoted as $(\mathbb{R}^n, d_\infty)$. A (rational) closed **ball** $\bar{\mathcal{B}}(b, K)$ in $\mathbb{R}^n$ with a radius $K \in \mathbb{Q}, K > 0$ centered at $b \in \mathbb{Q}^n$ is the set $\{x : x \in \mathbb{R}^n \wedge \|x - b\| \leq K\}$. For example, with the d_∞–metric, $\bar{\mathcal{B}}(b, K)$ is effectively a hypercube with 2^n vertices with rational coordinates. On the reals $\mathbb{R}$, it is a compact interval. Clearly, a closed ball $\bar{\mathcal{B}}(b, K)$ is located. A **regular partition** with a step $\delta = \frac{K}{k}, k \in \mathbb{N}$ on a closed ball $\bar{\mathcal{B}}(0, K)$ in $\mathbb{R}^n$ is a finite set of points $\{b_i\}_{i=1}^N \subset \bar{\mathcal{B}}(0, K), N \in \mathbb{N}$ with the coordinates satisfying all the combinations of the form $b_i^j := \pm n_{ij}\delta, n_{ij} \in \{0, \ldots, k\}$. Notice that the number N of partition points depends on the dimension n and the step δ. For instance, a regular partition on $\bar{\mathcal{B}}\left(\frac{1}{2}, \frac{1}{2}\right)$ in $\mathbb{R}$ – which is the unit interval $[0, 1]$ – with a step $\delta = \frac{1}{4}$ is the finite set of points $\left\{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\right\}$. Clearly, regular partitions on nontrivial closed balls exist and they may be considered as witnesses for total boundedness. That is, for any approximation to a closed ball, there exists a regular partition with the same property that for any point in the set there exists a point in the partition that is close to the given point up to the given precision. The following simple technical result can be easily proven:

Proposition 4. *Let $P = (p_1, \ldots, p_N)$ be a regular partition with a step $\delta = \frac{K}{k}, k \in \mathbb{N}$ on a closed ball $\bar{\mathcal{B}}(0, K) \subset (\mathbb{R}^n, d_\infty)$. Then, for any $x \in \bar{\mathcal{B}}(0, K)$, there exists a closed ball $\bar{\mathcal{B}}(p_i, \delta), p_i \in P$ such that $x \in \bar{\mathcal{B}}(p_i, \delta)$.*

Proof. Since x is a tuple of n real numbers $(x^1, \ldots, x^n)$, an m−th rational approximation to x is a tuple of rational numbers $x(m) := (x^1(m), \ldots, x^n(m))$. Indeed, for any $m' \in \mathbb{N}$, $\|x(m) - x(m')\| = \max_i |x^i(m) - x^i(m')| \leq \frac{1}{m} + \frac{1}{m'}$ since $|x^i(m) - x^i(m')| \leq \frac{1}{m} + \frac{1}{m'}$ for all $i \in \{1, \ldots, n\}$. Let $m := \lceil \frac{4}{\delta} \rceil$, then $\|x(m) - x\| \leq \frac{\delta}{2}$. Compute all distances $\|x(m) - p_i\|, p_i \in B$. If $\|x(m) - p_i\| < \frac{\delta}{2}$, then $x \in \bar{\mathcal{B}}(p_i, \delta)$. If there are more than one such balls, pick the one with the smallest index. If $\|x(m) - p_j\| = \frac{\delta}{2}$ for some indices $j = j_1, \ldots, j_L$, then pick the smallest such j and conclude that $x \in \bar{\mathcal{B}}(p_j, \delta)$. $\square$

Remark 5. In the current constructive setting, it cannot be deduced whether a point in a real metric space $(\mathbb{R}^n, d_\infty)$ belongs to a subset A or to a subset B if $A \cap B$ has a dimension less than n. However, comparison of a real number with a non-trivial interval is decidable, i.e., whether a real number in the interval $[a, b] \subset \mathbb{R}$ belongs to a non-trivial interval $I_1 \subseteq [a, b]$ or to $I_2 \subseteq [a, b]$ if and only if $I_1 \cap I_2$ is a non-trivial interval. In the proposition above, this fact is generalized to overlapping hypercubes.

Further, an important result, called *constructive Arzela–Ascoli's lemma*, which is due to Bishop & Bridges [8, p. 100], is recalled:

Proposition 6. *Let $\mathcal{F}$ be an equicontinuous function space from a totally bounded metric space (X, ρ) to a metric space (Y, σ). Suppose that for any finite $\frac{1}{k}$–approximation $\{x_1, \ldots, x_N\}$ to X, the set $A := \{(f(x_1), \ldots, f(x_N)) : f \in \mathcal{F}\} \subset Y^N$ is totally bounded. Then, $\mathcal{F}$ is totally bounded.*

Proof. Let ω be the continuity modulus for the function space $\mathcal{F}$. Let $k \in \mathbb{N}$, and let $\{x_1, \ldots, , x_N\}$ be an $\omega\left(\frac{1}{3k}\right)$–approximation to X. By assumption, the set A is totally bounded. Let $\{f_1, \ldots, , f_M\}$ be a set of functions in $\mathcal{F}$ such that the points $a_i := (f_i(x_1), \ldots, f_i(x_N))$, $i = 1, \ldots, M$ form a finite $\frac{1}{4k}$–approximation to X. Then, for an arbitrary $f \in \mathcal{F}$, it follows that there exists an f_i such that $\sum_{j=1}^{N} \sigma(f_i(x_j), f(x_j)) \leq \frac{1}{4k}$. For an arbitrary $x \in X$, there exists an x_j such that $\rho(x, x_j) \leq \omega\left(\frac{1}{3k}\right)$. Then,

$$\sigma(f_i(x), f(x)) \leq$$
$$\sigma(f_i(x), f_i(x_j)) + \sigma(f_i(x_j), f(x_j)) + \sigma(f_i(x_j), f(x)) \leq$$
$$\frac{1}{3k} + \frac{1}{4k} + \frac{1}{3k}.$$

It follows that $\tau(f_i, f) \leq \frac{1}{k}$. Therefore, $\{f_1, \ldots, , f_M\}$ is a finite $\frac{1}{k}$–approximation to $\mathcal{F}$ whence $\mathcal{F}$ is totally bounded. $\qquad \square$

3. Main results

Based on the preliminaries of the previous section, the new result on approximate optimal control policies can be derived. This is made in two steps. First,

it is shown that certain function spaces, which represent the sets of admissible control policies, admit finite approximations provided that they satisfy certain assumptions which are, however, applicable in practice.

3.1. *Finite approximations to function spaces*

The central theorem of this section is stated as follows:

Proposition 7. *Let $\mathcal{F}$ be the space of uniformly Lipschitz and uniformly bounded real–valued functions on a totally bounded metric space (X, ρ). Then $\mathcal{F}$ is totally bounded.*

Proof. Let X_0 be a finite subset of X consisting of unequal points $\{x_1, \ldots, , x_N\}$, $N \in \mathbb{N}$. Suppose that L and K are the uniform Lipschitz constant and uniform bound for $\mathcal{F}$ respectively. First, show that the subset

$$Y := \{(f(x_1), \ldots, f(x_N)) : f \in \mathcal{F}\}$$

of $\mathbb{R}^N$ with the product metric is totally bounded. To this end, let $P = (p_1, \ldots, , p_M), M \in \mathbb{N}$ be a regular partition of $[-K, K]$ with a step $\delta := \frac{1}{k}, k \in \mathbb{N}$. Let f be any function from the function space $\mathcal{F}$ and fix some arbitrary $n \in \mathbb{N}$. Construct a piece–wise linear function $\varphi : X \longrightarrow \mathbb{R}$ such that $\forall x, y \in X.|\varphi(x) - \varphi(y)| \leq L\rho(x, y)$ and $\forall x_i \in X_0.|f(x_i) - \varphi(x_i)| \leq \frac{1}{nN}$. By the product metric, the latter condition would imply that $\left\| (f(x_1), \ldots, f(x_N)) - (\varphi(x_1), \ldots, \varphi(x_N)) \right\| \leq \frac{1}{n}$. First, the image of φ on X_0 is constructed inductively. By Proposition 4, for any $x \in X$ and $f \in \mathcal{F}$, there exists a $p_i \in P$ such that $|f(x) - p_i| \leq \delta$. Suppose that $f(x_1), f(x_2)$ are within some closed balls $\bar{\mathcal{B}}(p_{j_1}, \delta), \bar{\mathcal{B}}(p_{j_2}, \delta), j_1, j_2 \in \{1, \ldots, M\}$ respectively. Let $\varphi(x_1) := p_{j_1}$. Observe that since $|f(x_1) - f(x_2)| \leq L\rho(x_1, x_2)$, it follows that $|p_{j_1} - p_{j_2}| \leq L\rho(x_1, x_2) + 2\delta$. Notice that p_{j_1} and p_{j_2} are rational numbers. It can, therefore, be assumed that either $p_{j_1} - p_{j_2} > 2\delta$, or $p_{j_1} - p_{j_2} < -2\delta$, or $|p_{j_1} - p_{j_2}| \leq 2\delta$. The first two cases are analogous whence one may assume that $p_{j_1} - p_{j_2} > 2\delta$. Let $\varphi(x_2) := p_{j_2} - 2\delta$. This setting ensures the Lipschitz condition. Indeed,

$$\left| \varphi(x_1) - \varphi(x_2) \right| = p_{j_2} - 2\delta - p_{j_1} \leq \left| p_{j_1} - p_{j_2} \right| - 2\delta \leq L\rho(x_1, x_2).$$

On the other hand, since $f(x_2) \in \bar{\mathcal{B}}(p_{j_2}, \delta)$ and by the setting of $\varphi(x_2)$ it follows that $f(x_2) \in \bar{\mathcal{B}}(\varphi(x_2), \delta + 2\delta)$. If $\left| p_{j_2} - p_{j_1} \right| \leq 2\delta$, then setting $\varphi(x_2) := p_{j_2}$ ensures the same conditions. Suppose now that, at the step i, $f(x_i) \in \bar{\mathcal{B}}(\varphi(x_2), \delta + 2(i-1)\delta)$. Assume that $f(x_{i+1}) \in \bar{\mathcal{B}}(p_{j_{i+1}}, \delta)$, $j_{i+1} \in \{1, \ldots, M\}$. Following exactly the same procedure, one may pick the next value of φ so that the approximation radius grows by 2δ whereas the Lipschitz condition is satisfied. After the step N, it holds that $f(x_N) \in \bar{\mathcal{B}}(\varphi(x_N), \delta + 2(N-1)\delta)$. Therefore, setting k equal to $2nN^2$ ensures $\left\| (f(x_1), \ldots, f(x_N)) - (\varphi(x_1), \ldots, \varphi(x_N)) \right\| \leq \frac{1}{n}$. Now, extend φ to the whole space X. To this end, let

$$\psi(x) := \frac{1}{2} \left[\max_i \left(\varphi(x_i) - L\rho(x, x_i) \right) + \min_i \left(\varphi(x_i) + L\rho(x, x_i) \right) \right], \forall x \in X.$$

Proposition 2 implies that $\psi(x_j) = \varphi(x_j)$. It follows from the fact that $\forall j, \varphi(x_i) - L\rho(x_j, x_i) \leq \varphi(x_j)$. To see that $\psi(x)$ is an L–Lipschitz function, observe that $\rho(x, x_i)$ is a 1–Lipschitz function of x whence $\varphi(x_i) - L\rho(x, x_i)$ is an L–Lipschitz function of x for each i. Therefore, $\max_i \left(\varphi(x_i) - L\rho(x, x_i) \right)$ and $\min_i \left(\varphi(x_i) + L\rho(x, x_i) \right)$ are uniformly continuous functions with the same Lipschitz constant L. The factor $\frac{1}{2}$ ensures that $\psi(x)$ is L–Lipschitz. Let $\varphi(x) := \max\{\min\{\psi(x), K\}, -K\}$. Due to the properties of the minimum and maximum [8, p. 23], it follows that φ is an L–Lipschitz continuous function satisfying $\varphi(x) \leq K$ and $\varphi(x) \geq -K$. Thus, φ belongs to the function space $\mathcal{F}$ and approximates f at the points X_0 arbitrarily close. Since φ is uniquely defined by its values at X_0, and the values $\{\varphi(x_i) : x_i \in X_0\}$ take place in a finite set P, whereas the distances between the function values at each two points of X_0 have fixed bounds, there are finitely many such functions. Further, since f was arbitrary, it follows that Y is totally bounded. By the constructive Arzela–Ascoli's proposition 6, the function space $\mathcal{F}$ is totally bounded. $\square$

Corollary 8. *Let $\mathcal{F}$ be the space of uniformly Lipschitz and uniformly bounded functions from a totally bounded metric space (X, ρ) to $\mathbb{R}^m$, $m \in \mathbb{N}$ with the d_∞–metric. Then $\mathcal{F}$ is totally bounded.*

Proof. The proof amounts to the same procedure, as in the proof of the theorem,

done for each coordinate separately since

$$\forall i = 1, \ldots, m, \left| x^i - y^i \right| \le \varepsilon \iff \|x - y\|_\infty = \max_i \left| x^i - y^i \right| \le \varepsilon$$

for any x and y in $(\mathbb{R}^m, d_\infty)$ and $\varepsilon > 0$. $\qquad\square$

Remark 9. The same result applies if there is a uniform bound and uniform Lipschitz constant for each dimension separately: $\exists (K_1, \ldots, K_m), \forall f \in \mathcal{F}, \forall x \in X, |f^i(x)| \le K_i, i = 1, \ldots, m$ and $\exists (L_1, \ldots, L_m), \forall f \in \mathcal{F}, \forall x, y \in X, \big|f^i(x) - f^i(y)\big| \le L\rho(x, y), i = 1, \ldots, m$. The proof is by rescaling of the hypercuboid with the side lengths $(2K_1, \ldots, 2K_m)$ centered at the origin.

Corollary 10. *Let $\bar{\mathcal{C}}^1$ be the space of uniformly bounded functions from a compact set $X \subset \mathbb{R}^n, n \in \mathbb{N}$ to $\mathbb{R}^m, m \in \mathbb{N}$ with the d_∞–metric, and suppose that the derivatives of the functions in $\bar{\mathcal{C}}^1$ are uniformly bounded. Then $\bar{\mathcal{C}}^1$ is totally bounded.*

Proof. Let $\mathcal{F}$ denote the space of functions as described in the theorem. Fix an arbitrary function g from $\bar{\mathcal{C}}^1$. Clearly, g is a function in $\mathcal{F}$ whence $\bar{\mathcal{C}}^1 \subset \mathcal{F}$. The converse is not true. It suffices to show that $\bar{\mathcal{C}}^1$ is dense in $\mathcal{F}$. Fix an arbitrary function $f \in \mathcal{F}$ and a number $k \in \mathbb{N}$. Following the construction as in the proof, one can derive a piece–wise linear function $\varphi : \mathbb{R}^n \to (\mathbb{R}^m, d_\infty)$ that approximates f on X up to the precision $\frac{1}{2k}$. Further, one can construct an analytic function φ_{2k} that approximates φ up to the precision $\frac{1}{2k}$ and has the same Lipschitz constant (see details in Appendix). Therefore, φ_k approximates f up to the precision $\frac{1}{k}$. Therefore, the set $\bar{\mathcal{C}}^1$, as a dense subset of a totally bounded set $\mathcal{F}$, is itself totally bounded [10, p. 28]. $\qquad\square$

Corollary 11. *Let $\bar{\mathcal{C}}^N, N \in \mathbb{N}$ be the space of uniformly bounded functions from a compact set $X \subset \mathbb{R}^n, n \in \mathbb{N}$ to $\mathbb{R}^m, m \in \mathbb{N}$ with the d_∞–metric, and suppose that the derivatives of the functions in $\bar{\mathcal{C}}^N$ up to order N are uniformly bounded. Then $\bar{\mathcal{C}}^N$ is totally bounded.*

3.2. Approximate extrema

In this section, based on the construction of approximating functions of the previous section, the new constructive version of the approximate extremum value theorem for function spaces is stated. The implication of it is that, under certain assumptions, approximate control policies may be computed up to a prescribed accuracy.

Theorem 12. *Let $\mathcal{F}$ be the space of uniformly Lipschitz and uniformly bounded functions from a totally bounded metric space (X, ρ) to $\mathbb{R}$, and let J be a uniformly continuous functional from $\mathcal{F}$ to $\mathbb{R}$. Then, for any $k \in \mathbb{N}$, there exists an $f \in \mathcal{F}$ such that $J[f] - \frac{1}{k} \le \inf J$.*

Proof. Since $\mathcal{F}$ is totally bounded by Theorem 7 and J is uniformly continuous, $\inf J$ exists. Let α be the continuity modulus of J and $\mathcal{F}_0 = \{f_1, \ldots, f_N\}$ be an $\alpha\left(\frac{1}{8k}\right)$−approximation to $\mathcal{F}$. Consider all finitely many $\{J[f_i](8k)\}, i = 1, \ldots, N$. Let $J[f_j](8k)$ be the smallest one, and such that j is the smallest index if there are more than one such indices. Observe that $\left|J[f_j] - J[f_j](8k)\right| \le \frac{1}{4k}$ and $\forall f \in \mathcal{F}, \left\|f_j - f\right\| \le \alpha\left(\frac{1}{8k}\right) \implies \left|J(f) - J[f_j]\right| \le \frac{1}{4k}$ whence $J[f_j](8k) - \frac{1}{2k} \le J[f](8k)$. Therefore, $J[f_j](8k) - \frac{1}{2k} \le J[f]$ and consequently $J[f_j] - \frac{1}{k} \le J[f]$. The same holds trivially if $\left\|f_j - f\right\| > \alpha\left(\frac{1}{8k}\right)$. Since f is arbitrary, $J[f_j] - \frac{1}{k}$ is a lower bound of J and so, in particular, $\inf J \ge J[f_j] - \frac{1}{k}$. $\square$

Corollary 13. *Let $\mathcal{F}$ be the space of uniformly Lipschitz and uniformly bounded functions from a totally bounded metric space (X, ρ) to $\mathbb{R}$, and let J be a uniformly continuous functional from $\mathcal{F}$ to $\mathbb{R}^m, m \in \mathbb{N}$ with the d_∞−metric. Then, for any $k \in \mathbb{N}$, there exists an $f \in \mathcal{F}$ such that $J[f] - \frac{1}{k} \le \inf J$.*

Remark 14. To compute an extremal function which yields a $\frac{1}{k}$-infimum of J, where k describes the specified precision, construct all possible piece-wise linear functions over the regular partition of step $\frac{1}{2kN^2}$ where N is as in Theorem 12 by preserving the common Lipschitz constant. Smoothen the constructed functions, if necessary, as per Corollary 10. Then, choose a one f_j which satisfies $J[f] - \frac{1}{k} \le \inf J$.

Remark 15. Theorem 12 describes the worst-case scenario a numerical algorithm can perform in general. Various numerical approaches exist and they may be numerically fast, but the best one can expect in general is the result as in the statement of the theorem. The implication for optimal control is that optimality may fail to be achieved in general. Instead, approximate optimal control policies can be effectively computed, provided that the system and the cost function satisfy the assumptions in the statement of the theorem. These assumptions are, however, practicable as justified by physical nature of the control problems and demonstrated in the next section on finite-horizon optimal control, dynamic programming and adaptive dynamic programming.

Remark 16. The statement for the supremum is equivalent.

4. Case study: optimal control

In this section, the derived version of a constructive extremum value theorem in application to finite–horizon optimal control, dynamic programming (DP), and adaptive dynamic programming (ADP), is discussed.

4.1. Finite-horizon optimal control

Classical theorems of existence of extremal solutions to functional optimization problems essentially rely on Bolzano-Weierstrass's theorem that every bounded sequence has a convergent subsequence. One first shows that the function space in question is compact, and applies the sequential compactness argument. There is, unfortunately, no constructive way to find a convergent subsequence. Therefore, approximate solutions are investigated in this section. Recall the problem of minimization of the following cost functional:

$$J[u] := \varphi\left(x\left(t_1\right)\right) + \int_{t_0}^{t_1} \mathcal{L}(x(t), u(x(t)), t)dt \tag{2}$$

subject to $\dot{x}(t) := f(x(t), u(t), t), x\left(t_0\right) = x_0$. Here, $\mathcal{L}$ is the running cost, or Lagrangian, which is usually a positive-definite function of x, u, t. Assume that the state space $X \subset \left(\mathbb{R}^n, d_\infty\right), n \in \mathbb{N}$ is compact. With the d_∞–metric, two

states x and y are close whenever their respective components $x^i, i = 1, \ldots, n$ and $y^i, i = 1, \ldots, n$ are close. Therefore, a state trajectory $x(t)$ is uniformly continuous whenever each state component $x^i(t)$ is uniformly continuous. It can be assumed that $u(x) \in (\mathbb{R}^m, d_\infty), m \in \mathbb{N}$ for any $x \in X$. Let $\mathcal{U}$ denote the set of admissible control policies, i. e., those which yield state trajectories within X. In (2), the starting time $t_0 \in \mathbb{Q}$ and the final time $t_1 \in \mathbb{Q}$ are assumed fixed. Suppose that $f : X \times U \times \mathbb{R} \to X$ satisfies the Lipschitz condition for x and u on $X \times U$ in the following sense:

$$\|f(x, u, t) - f(y, v, t)\| \le L_f \max\{\|x - y\|, \|u - v\|\} \tag{3}$$

for some rational $L_f > 0$. Then, the constructive theorem of existence and uniqueness of solutions of the initial value problem $\dot{x}(t) = f(x(t), u(t), t), x(t_0) = x_0, t \in [t_0, t_1]$ applies [19, 12.4] provided that $u(t)$ is continuous. Further, φ is assumed to be uniformly continuous on $X \times X$. The Lagrangian should be also uniformly continuous on $X \times U$:

$$\exists \omega_\mathcal{L} : \mathbb{Q} \to \mathbb{Q}, \forall k \in \mathbb{N}, \forall x, y \in X, \forall u, v \in U, \forall t \in [t_0, t_1],$$
$$\max\{\|x - y\|, \|u - v\|\} \le \omega_\mathcal{L}\left(\frac{1}{k}\right) \implies$$
$$|\mathcal{L}(x, u, t) - \mathcal{L}(y, v, t)| \le \frac{1}{k}.$$

Consider two control policies $u(x), v(x), x \in X$ in $\mathcal{U}$ and the respective state trajectories:

$$x_u(\tau) = \int_{t_0}^{\tau} f(x(t), u(x(t)), t)dt,$$

$$x_v(\tau) = \int_{t_0}^{\tau} f(x(t), v(x(t)), t)dt$$

for an arbitrary $\tau \in [t_0, t_1]$. It follows that

$$\left\| \int_{t_0}^{\tau} \left(f(x(t), u(x(t)), t) - f(x(t), v(x(t)), t) \right) dt \right\| \leq$$

$$\sup_{t_0 \leq \tau \leq t_1} \| f(x(t), u(x(t)), t) - f(x(t), v(x(t)), t) \| \cdot (\tau - t_0) \leq$$

$$L_f \cdot \sup_{x \in X} \| u(x) - v(x) \| \cdot (\tau - t_0) \leq$$

$$L_f \cdot \| u - v \| \cdot (t_1 - t_0).$$

Therefore, if $\sup_{x \in X} \| u(x) - v(x) \| \leq \frac{1}{k}, k \in \mathbb{N}$, which is to say that $\| u - v \| \leq \frac{1}{k}$, then $\left\| x_u - x_v \right\| \leq L_f \frac{t_1 - t_0}{k}$. Consequently, for $k \geq L_f (t_1 - t_0)$, and $\| u - v \| \leq \omega_{\mathcal{L}} \left(\frac{1}{k} \right)$ it follows that: $\left| \mathcal{L} \left(x_u(t), u(t), t \right) - \mathcal{L} \left(x_v(t), v(t), t \right) \right| \leq \frac{1}{k}$. If $k < L_f (t_1 - t_0)$, then $\left| \mathcal{L} \left(x_u(t), u(t), t \right) - \mathcal{L} \left(x_v(t), v(t), t \right) \right| \leq L_f \frac{t_1 - t_0}{k}$ whence the continuity modulus is easily derived. Further, it holds that:

$$\left| \int_{t_0}^{t_1} \left(\mathcal{L}(x(t), u(x(t)), t) - \mathcal{L}(x(t), v(x(t)), t) \right) dt \right| \leq \frac{t_1 - t_0}{k},$$

In case if $\mathcal{U}$ is a located subset of the space of uniformly bounded and uniformly Lipschitz functions on X with a uniform bound and uniform Lipschitz constant, respectively (which can physically be dictated by the fact that any control policy has limits on magnitude and rate of change with respect to the state), by Theorem 7 and Lemma 4.3. in [27], $\mathcal{U}$ is totally bounded. Since J is a uniformly continuous functional from $\mathcal{U}$ to $\mathbb{R}$, by Theorem 12, for any $k \in \mathbb{N}$, there exists a control policy $u^*(x), x \in X$ such that $J[u^*] - \frac{1}{k} \leq J[u]$ for any other control policy $u(x), x \in X$.

This implies that control policies, which yield approximate optima of the cost functional, can be effectively computed provided that the controllers have bounds on the magnitude and rate of change of the controls which is satisfied in practice due to physical nature. The next section considers infinite-horizon optimal control in the framework of dynamic programming.

4.2. Dynamic programming

In dynamic programming, optimization problems in the following form are considered:

$$\sup_{u \in U} \{r(x, u) + \gamma V(f(x, u))\}, \forall x \in X, \tag{4}$$

In (4), u is taken from a totally bounded set $U \subset \mathbb{R}^m, m \in \mathbb{N}$. In the case when each component $u^i, i = 1, \ldots, m$ is an independent function, the d_∞–metric may be assumed on $\mathbb{R}^m$. The function $r : X \times U \longrightarrow \mathbb{R}$ is a positive–definite *utility function (or running cost)* that describes the instantaneous cost whereas V is the *value function* that describes the cumulative cost. The functions r and V are assumed to be bounded on their domains. The parameter $0 < \gamma < 1$ is called *discounting factor*. Finally, the *dynamic programming operator* is introduced:

$$T[V](x) := \sup_{u \in U} \{r(x, u) + \gamma V(f(x, u))\}, \forall x \in X. \tag{5}$$

The operator T acts on the space of continuous and bounded functions. The *Hamilton–Jacobi–Bellman* equation is defined by the fixed–point of T. The natural question is whether T yields continuous functions, whether the extrema of $r(x, u) + \gamma V(f(x, u))$ over U exist, and whether they are continuous in x in a certain sense. The answer is given by the Berge's Theorem of the Maximum [4, p. 115]. It shows what kind of continuity of the extrema is preserved if the optimization problem is continuous in a certain sense. First, recall the definition of hemi–continuity that generalizes the notion of continuity to *multi–functions*, i. e., functions from a set to the power set of another set. A multi–function is called *compact–valued* if its values are compact sets.

Definition 17. (Upper hemi–continuity) Let $\Gamma : X \to U$ be a multi–function such that $\Gamma(x)$ is closed for all x in X. Then, Γ is called upper hemi–continuous at $x \in X$ if for any sequence $\{x_n\}_n$ in X, u in U and sequence $\{u_n\}_n$ such that $u_n \in \Gamma(x_n)$, it follows that

$$\left(\lim_{n \to \infty} x_n = x \wedge \lim_{n \to \infty} u_n = u \right) \implies u \in \Gamma(x).$$

Definition 18. (Lower hemi–continuity) A multifunction $\Gamma : X \to U$ is called lower hemi–continuous at $x \in X$ if for any sequence $\{x_n\}_n$ in X such that $\lim_{n \to \infty} x_n = x$, any $u \in \Gamma(x)$, there exists a subsequence $\{x_{nk}\}_k \subset \{x_n\}_n$ such that there exist $u_k \in \Gamma(x_{nk})$ with $\lim_{k \to \infty} u_k = u$.

Theorem 19. *(Theorem of the Maximum) Let f be a jointly continuous function from a product of two metric spaces $(X, \rho) \times (U, \sigma)$ to $\mathbb{R}$, and $\Gamma : X \to U$ be a compact–valued upper and lower hemi–continuous multi–function. Then, the function $h(x) := \sup\limits_{u \in \Gamma(x)} f(x, u)$ is continuous, and $u^*(x) := \arg \sup\limits_{u \in \Gamma(x)} f(x, u)$ is a non–empty, compact–valued, upper hemi–continuous multi–function.*

In the setting of (4), the multi–function Γ is taken as a constant multi–function $\Gamma(x) \equiv U$ with the assumption that any control action is available at any state. The proof of Theorem 19 essentially uses the classical extremum value theorem that is not valid constructively. A constructive analysis of the Maximum Theorem has been done by Tanaka [23]. To prove a constructive version of Theorem 19, Tanaka introduces the notion of a function with *sequentially locally at most one maximum*. Such a condition is, however, hard to verify in practice. To summarize, approximate extrema of (4) cannot be shown uniformly continuous in general, let alone found exactly. However, a weaker result holds:

Proposition 20. *Let f be a uniformly continuous function from a product of two totally bounded metric spaces $(X, \rho) \times (U, \sigma)$ to $\mathbb{R}$. Then, the function $h(x) := \sup\limits_{u \in U} f(x, u)$ is uniformly continuous.*

Proof. Fix any $x, y \in X$ and $k \in \mathbb{N}$. By Theorem 12, there exists $u^* \in U$ such that $f(x, u^*) + \frac{1}{k} \geq \sup\limits_{u \in U} f(x, u)$. Let ω be the continuity modulus of f. Fix any $y \in X$ such that $\rho(x, y) \leq \omega\left(\frac{1}{k}\right)$. It follows that $\forall u \in U, |f(x, u) - f(y, u)| \leq \frac{1}{k}$. In particular, $|f(x, u^*) - f(y, u^*)| \leq \frac{1}{k}$ whence $f(y, u^*) + \frac{1}{k} \geq f(x, u^*)$. Therefore, $f(y, u^*) + \frac{2}{k} \geq \sup\limits_{u \in U} f(x, u)$. From the continuity condition, it follows that $f(y, u) \leq f(x, u) + \frac{1}{k}$ whence $f(y, u^*) + \frac{3}{k} \geq f(y, u)$ for all $u \in U$. It means that $f(y, u^*) + \frac{3}{k} \geq \sup\limits_{u \in U} f(y, u)$. Finally, $\left| \sup\limits_{u \in U} f(x, u) - \sup\limits_{u \in U} f(y, u) \right| \leq$

$|f(x, u^*) - f(y, u^*)| + \frac{2}{k} = \frac{3}{k}$. It follows that h is a uniformly continuous function with a modulus $k \mapsto \omega\left(\frac{3}{k}\right)$. $\square$

It can be proven that T is indeed an operator that sends uniformly continuous functions to uniformly continuous functions provided that r, V, f are uniformly continuous functions. In turn, uniformly continuous functions on totally bounded sets are bounded. Denote the set of uniformly continuous functions from X to X by $\mathcal{V}$. Assuming the supremum norm on $\mathcal{V}$, it is easy to show that T is a contraction mapping. The proof is already constructive. First, observe that $r(x, u) + \gamma V(f(x, u)) \le r(x, u) + \gamma W(f(x, u))$ whenever $V(x) \le W(x)$ for any x in X. Therefore,

$$\sup_{u \in U} \{r(x, u) + \gamma V(f(x, u)))\} \le \sup_{u \in U} \{r(x, u) + \gamma W(f(x, u)))\}.$$

Consequently, $T[V] \le T[W]$. It constitutes the monotonicity of T which is the first Blackwell's sufficient condition for a contraction mapping [9]. The second condition requires that T be discounting:

$$T[V + a](x) = \sup_{u \in U} \{r(x, u) + \gamma V(f(x, u)) + \gamma a\} = T[V](x) + \gamma a$$

for any $a \ge 0$. It follows that T is a contraction mapping with a modulus γ. By the Banach fixed point theorem [17], T has a unique fixed point. The proof of the theorem is essentially constructive and provided by the algorithm:

$$V_n := T[V_{n-1}], n \in \mathbb{N}$$

starting from an arbitrary uniformly continuous function V_0 from which it immediately follows that $\{V_n\}_n$ is a Cauchy sequence with the modulus defined from:

$$\left\|T^n[V_0] - V^*\right\| \le \frac{\gamma^n}{1 - \gamma}\left\|T[V_0] - V_0\right\|,$$

where $T^n[V] = \underbrace{T[T[\ldots T[V]]]}_{n \text{ times}}$ and V^* is the fixed point. Thus, V_n converges to V^* uniformly and V^* is in turn a uniformly continuous function. The only difference to the classical theorem is that $\mathcal{V}$ must be non-empty: $\exists V_0 \in \mathcal{V}$. Using

this fact, one can perform *value iteration* starting from any $V_0 \in \mathcal{V}$ and stopping at some V_n such that a convergence criterion $||V_n - V^*|| \leq \varepsilon$ is satisfied. Notice that the number of steps n can be directly determined from the desired accuracy ε. Having found a suitable approximate V_n, an optimal control policy is of interest. In the constructive setting, a control policy must be a uniformly continuous function. In real–world applications, each control action has physical limits for magnitude and rate of change. Thus, one may assume that the space of admissible control policies $\mathcal{U}$ is a located subset of the space of uniformly bounded and uniformly Lipschitz functions on X with a common bound and Lipschitz constant, respectively. By Theorem 7, the latter space is totally bounded, and since $\mathcal{U}$ is located, it is also totally bounded [27, Lemma 4.3.]. To cope with the problem of continuity of extrema in the state variable, it is suggested to consider the following relaxed optimization problem for V_n:

$$\sup_{u \in \mathcal{U}} \inf_{x \in X} \left\{ r\left(x, u(x)\right) + \gamma V_n(f(x, u(x))) \right\}. \tag{6}$$

It follows that $J : \mathcal{U} \to \mathbb{R}$ defined by

$$J[u] := \inf_{x \in X} \left\{ r\left(x, u(x)\right) + \gamma V_n(f(x, u(x))) \right\}$$

is a uniformly continuous functional since $r, V_n, f, \inf$ are uniformly continuous. By Theorem 12, for any $k \in \mathbb{N}$, there exists a control policy $u^* \in \mathcal{U}$ such that

$$\inf_{x \in X} \left\{ r\left(x, u^*(x)\right) + \gamma V_n(f(x, u^*(x))) \right\} + \frac{1}{k} \geq \inf_{x \in X} \left\{ r\left(x, u(x)\right) + \gamma V_n(f(x, u(x))) \right\}$$

for any control policy u. The difference between (6) and (4) lies in the way the performance mark is defined. In the latter case, $r\left(x, u\right) + \gamma V(f(x, u))$ is optimized in all states x while in the former, the "worst" state is optimized. The optimization problem (6) is thus more mild than the original one, but it is still appropriate for a variety of practical applications.

4.3. Adaptive dynamic programming

ADP is a variant of dynamic programming that is suitable for real–time optimization problems. It may be considered as a reinforcement learning tech-

nique [22] in the sense that it uses an iterative procedure of updating a so–called actor, that produces a control policy, according to a citric that represents the value function. The value function in the framework of ADP is commonly a subject to approximation since exact optimal solutions may be not achieved [15]. In this regard, neural networks are widely used as approximators [25, 26, 6]. For recent surveys on ADP, refer, e. g., to Balakrishnan et al. [2] and Ferrari et al. [13]. It is common to consider ADP in application to discrete–time systems of the form $x_{k+1} := f(x_k, u_k), k \in \mathbb{N}$ or even affine in control: $x_{k+1} := f(x_k) + g(x_k) u_k, k \in \mathbb{N}$. It is also assumed that $f(0) = g(0) = 0$, and that there exists a control policy u such that for all initial conditions $x_0 \in X$, $x_k \to 0$ as $k \to \infty$. ADP usually addresses the following infinite-horizon optimization problem:

$$\inf_u \sum_{l=k}^{\infty} r(x_l, u(x_l)), \forall x_k \in X. \tag{7}$$

Al Tamimi et al. [1] have provided a convergence analysis of an ADP algorithm for affine–in–control systems. Assuming a utility function in the form $r(x, u) := q(x) + u^T R u$ with q being a positive-definite function, and given an arbitrary $\mathcal{C}^{\infty}(X \subset \mathbb{R}^n, \mathbb{R})$ –function $V_0(x)$ such that $\forall x \in X, 0 \leq V(x) \leq Q(x)$, perform the following iterations starting with $i := 0$ for all $x_k \in X$:

$$u_i(x_k) : = \arg\inf_u \{r(x_k, u) + V_i(f(x_k, u))\}, \tag{8}$$

$$V_{i+1}(x_k) : = r(x_k, u_i(x_k)) + V_i(f(x_k, u_i(x))), \tag{9}$$

$$i : = i + 1.$$

Al-Tamimi showed that this algorithm converges to the solution of the Hamilton–Jacobi–Bellman equation

$$V^*(x_k) = \inf_u \{r(x_k, u) + V^*(f(x_k, u))\}, \tag{10}$$

$$u^*(x_k) = \arg\inf_u \{r(x_k, u) + V^*(f(x_k, u))\}, \tag{11}$$

$$\forall x_k \in X \tag{12}$$

as $i \to \infty$. The proof essentially uses the classical monotone convergence theorem that states that a sequence of real numbers converges whenever it is bounded and monotone. Consequently, no estimate on the number of iterations can be given for the prescribed accuracy $||V_i - V^*||$. Another subtle point, that is hard to justify from both the constructive and practical viewpoint, is the assumption that (8) can be solved in terms of a closed-form expression. That is generally impossible. An exception is, e. g., a *linear quadratic regulator* (LQR) that is a solution for linear systems. Liu & Wei [16] introduced a similar proof technique, as in [1], for a *policy iteration* algorithm: start with $i := 0$, and any continuous control policy u_0 such that $u(0) = 0$, the state trajectory $x_k \to 0$ under u_0, and (7) converges, perform the following iterations:

$$V_i(x) : \quad = \quad r(x, u_i(x)) + V_i(f(x, u_i(x))), \tag{13}$$

$$u_i(x_k) : \quad = \quad \arg\inf_u \{r(x_k, u) + V_{i-1}(f(x_k, u))\}, \tag{14}$$

$$\forall x_k \in X \tag{15}$$

$$i : \quad = \quad i + 1.$$

Notice the difference in iteration indices for the value function and the control policy. Again, the proof of convergence uses the monotone convergence theorem and the assumption that (14) can be solved in terms of a closed-form expression. To coup with this problem, Al-Tamimi and Liu suggest to use neural–network based approximators for the value function and the control policy. Unfortunately, no convergence proof has been given for such an approximate setting [16, p. 632]. An alternative approach has been proposed by Heydari [14]. Instead of approximating the control policy, Heydari has shown that the first–order necessary condition for an extremum

$$u_i(x) = -\frac{1}{2}R^{-1}g^T(x)\frac{\partial V_i\left(f(x) + g(x)u_i(x)\right)}{\partial x}, \forall x \in X$$

is a fixed–point equation provided that all the functions in question are $\mathcal{C}^\infty$. By an appropriate choice of the matrix norm of R^{-1} and/or $g(x)$, it can be shown that the mapping $F : \mathcal{C}^\infty(X, \mathbb{R}^m) \to \mathcal{C}^\infty(X, \mathbb{R}^m)$ defined by

$$F[u](x) := -\frac{1}{2}R^{-1}g^T(x)\frac{\partial V_i(f(x) + g(x)u(x))}{\partial x}, \tag{16}$$

is a contraction. The assumption that the control policy at each iteration is a smooth function satisfies our argumentation in Section 4.2 and, provided with a uniform bound and Lipschitz constant, leads to total boundedness of the space of control policies by Corollary 11. However, the first–order condition for an extremum is not sufficient to claim that the infimum of (9) is attained at each iteration. Currently, one can decouple (8) and (9), iterate the value function and then claim existence of an approximate optimal control policy for an alternative performance mark (6). Relaxing the continuity condition by considering measurable functions, and investigating other performance marks, such as Lebesgue integrals, may be of interest for future research.

5. Conclusions

The present work is highlighted in the following points:

- A new constructive proof of the approximate extremum value theorem for function spaces is suggested.

- The methodological approach of the proof takes into account the numerical uncertainty which is related to limitations of real number representations in a computational device.

- The functions forming finite approximations to the respective function spaces are constructed explicitly. In particular, it was shown that the sets of uniformly bounded and uniformly Lipschitz functions on a totally bounded set are totally bounded by explicit constructions of approximating functions.

- As stated in Remark 15, any numerical procedure may in general achieve the result of Theorem 12 at best. That implies that optimality in optimal control problems may in general be achieved only approximately.

- Applications of the theorem to finite-horizon optimal control, dynamic programming and adaptive dynamic programming are addressed. It is

shown that under the stated assumptions, whose practicability is discussed, approximate optimal control policies can be effectively computed up to prescribed accuracy on the approximate optima of the cost functional.

References

[1] Al Tamimi, A., Lewis, F. L., & Abu Khalaf, M. (2008). Discrete-Time Nonlinear HJB Solution Using Approximate Dynamic Programming: Convergence Proof. *IEEE Transactions on Systems, Man, and Cybernetics. Part B (Cybernetics)*, *38*(4), 943–949.

[2] Balakrishnan, S., Ding, J., & Lewis, F. L. (2008). Issues on stability of ADP feedback controllers for dynamical systems. *IEEE Transactions on Systems, Man, and Cybernetics. Part B: Cybernetics*, *38*(4), 913–917.

[3] Banaschewski, B. & Mulvey, C. (1997). A constructive proof of the Stone-Weierstrass theorem. *Journal of Pure and Applied Algebra, 116*(1), 25–40.

[4] Berge, C. (1963). *Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity.* Courier Corporation.

[5] Berger, J., Bridges, D., & Schuster, P. (2006). The fan theorem and unique existence of maxima. *Journal of Symbolic Logic, 71*(02), 713–720.

[6] Bertsekas, D. P. & Tsitsiklis, J. N. (1995). Neuro-dynamic programming: an overview. In *Proceedings of the 34th IEEE Conference on Decision and Control*, volume 1, (pp. 560–564). IEEE.

[7] Bishop, E. (1967). *Foundations of Constructive Analysis*, volume 60. McGraw-Hill New York.

[8] Bishop, E. & Bridges, D. (1985). *Constructive Analysis*, volume 279. Springer Science & Business Media.

[9] Blackwell, D. (1965). Discounted dynamic programming. *The Annals of Mathematical Statistics*, 226–235.

[10] Bridges, D. & Richman, F. (1987). *Varieties of Constructive Mathematics*, volume 97. Cambridge University Press.

[11] Bridges, D. & Vita, L. (2007). *Techniques of Constructive Analysis*. Universitext. Springer New York.

[12] Bridges, D. S. (2007). Constructing local optima on a compact interval. *Arch. Math. Logic*, *46*(2), 149–154.

[13] Ferrari, S., Sarangapani, J., & Lewis, F. L. (2011). Special issue on approximate dynamic programming and reinforcement learning. *Journal of Control Theory and Applications*, *9*(3), 309–309.

[14] Heydari, A. (2014). Revisiting approximate dynamic programming and its convergence. *IEEE Transactions on Cybernetics*, *44*(12), 2733–2743.

[15] Lewis, F. L. & Syrmos, V. L. (1995). *Optimal control*. John Wiley & Sons.

[16] Liu, D. & Wei, Q. (2014). Policy Iteration Adaptive Dynamic Programming Algorithm for Discrete-Time Nonlinear Systems. *IEEE Transactions on Neural Networks and Learning Systems*, *25*(3), 621–634.

[17] Palais, R. S. (2007). A simple proof of the Banach contraction principle. *Journal of Fixed Point Theory and Applications*, *2*(2), 221–223.

[18] Rump, S. (2010). Verification methods: Rigorous results using floating-point arithmetic. *Acta Numerica*, *19*, 287–449.

[19] Schwichtenberg, H. (2012). *Constructive Analysis with Witnesses*. Manuscript.

[20] Sepulchre, R., Jankovic, M., & Kokotovic, P. (2012). *Constructive Nonlinear Control*. Springer Science & Business Media.

[21] Sontag, E. (1989). A "universal" construction of artstein's theorem on nonlinear stabilization. *Systems & Control Letters*, *13*(2), 117–123.

[22] Sutton, R. S. & Barto, A. G. (1998). *Reinforcement learning: An introduction*, volume 1. MIT Press Mambridge.

[23] Tanaka, Y. (2012). On the maximum theorem: A constructive analysis. *Int. J. Comp. and Math. Sciences*, *6*(2), 173–175.

[24] Turing, A. M. (1936). On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, *42*(230–265).

[25] Werbos, P. (1990). A menu of designs for reinforcement learning over time. *Neural networks for control*, 67–95.

[26] Werbos, P. J. (1992). Approximate dynamic programming for real-time control and neural modeling. *Handbook of intelligent control: Neural, fuzzy, and adaptive approaches*, *15*, 493–525.

[27] Ye, F. (2011). *Strict Finitism and the Logic of Mathematical Applications*, volume 355. Springer.

Appendix

Smooth approximations

For the sake of completeness, some technical details of smooth approximation are discussed in this appendix. First, consider the following real–valued non–analytic $\mathcal{C}^\infty$–function on $\mathbb{R}$:

$$
\sigma(t) = \begin{cases} \mathrm{e}^{-1/t^2} & t > 0, \\ \\ 0 & t \leq 0. \end{cases}
$$

Define a $\mathcal{C}^\infty$ bump function $\vartheta : \mathbb{R}^n \to \mathbb{R}, n \in \mathbb{N}$ by $\vartheta(x) := a \cdot \sigma \left(1 - \|x\|^2\right)$ where

$$a := \left(\int_{\mathbb{R}^n} \sigma \left(1 - \|x\|^2\right) dx\right)^{-1}.$$

Then, the support of ϑ lies within the unit ball $\mathcal{B}(0,1)$, i. e. $\{x \in \mathbb{R}^n : \vartheta(x) > 0\} \subset \mathcal{B}(0,1)$. Clearly, $\forall x \in \mathbb{R}^n, \vartheta(x) \geq 0$ and $\int_{\mathbb{R}^n} \vartheta(x) dx = 1$. Now, let $\vartheta_k(x) := k^n \vartheta(kx)$ for $k \in \mathbb{N}$. It follows that $\int_{\mathbb{R}^n} \vartheta_k(x) dx = k^n \int_{\mathbb{R}^n} \frac{1}{k^n} \vartheta(kx) d(kx) = 1$. Let $f : \mathbb{R}^n \to \mathbb{R}$ be an L–Lipschitz function. It may be assumed that $L = 1$ without loss of generality. Clearly, f is locally integrable, i. e. integrable on any compact subset of $\mathbb{R}^n$ since it is Lipschitz continuous on this subset. Since ϑ_k is compactly supported, define a $\mathcal{C}^\infty$ function f_k by convolution as follows:

$$f_k(x) = \int_{\mathbb{R}^n} f(\chi) \vartheta_k(x - \chi) d\chi = \int_{\mathbb{R}^n} f(x - \chi) \vartheta_k(\chi) d\chi =$$
$$k^n \int_{\mathcal{B}(0,1/k)} f(x - \chi) \vartheta(k\chi) d\chi = \int_{\mathcal{B}(0,1)} f\left(x - \frac{\chi}{k}\right) \vartheta(\chi) d\chi.$$

It follows that

$$f_k(x) - f_k(y) = \int_{\mathbb{R}^n} \left(f(x - \chi) - f(y - \chi)\right) \vartheta_k(\chi) d\chi.$$

Since $\forall \chi \in \mathbb{R}^n, \left|f(x - \chi) - f(y - \chi)\right| \vartheta_k(\chi) \leq \|x - y\| \vartheta_k(\chi)$, it holds that:

$$\left|f_k(x) - f_k(y)\right| \leq \int_{\mathbb{R}^n} \left|f(x - \chi) - f(y - \chi)\right| \vartheta_k(\chi) d\chi \leq$$
$$\|x - y\| \int_{\mathbb{R}^n} \vartheta_k(\chi) d\chi = \|x - y\|.$$

Finally, since $f(x) = f(x) \cdot 1 = f(x) \cdot \int_{\mathcal{B}(0,1)} \vartheta(\chi) d\chi = \int_{\mathcal{B}(0,1)} f(x) \vartheta(\chi) d\chi$, it follows that

$$\left| f_k(x) - f(x) \right| = \left| \int\limits_{\mathcal{B}(0,1)} f\left(x - \frac{\chi}{k}\right) \vartheta(\chi)d\chi - f(x) \right| =$$

$$\left| \int\limits_{\mathcal{B}(0,1)} f\left(x - \frac{\chi}{k}\right) \vartheta(\chi)d\chi - \int\limits_{\mathcal{B}(0,1)} f(x)\,\vartheta(\chi)d\chi \right| =$$

$$\left| \int\limits_{\mathcal{B}(0,1)} \left(f\left(x - \frac{\chi}{k}\right) - f(x) \right) \vartheta(\chi)d\chi \right| \le$$

$$\sup_{\chi \in \mathcal{B}(0,1)} \left\| \frac{\chi}{k} \right\| \int\limits_{\mathcal{B}(0,1)} f\vartheta(\chi)d\chi = \frac{1}{k}.$$

Brouwerian counter-examples

The two examples in Section 1, also called Brouwerian counter-examples, demonstrate the computational inability of finding exact optimal control actions and policies in general. The possible computational problems with addressing optimal control may be related to the so-called principles of omniscience [8, p. 11]. One of them, called *limited principle of omniscience* (LPO) states:

Definition 21. (LPO) For any sequence binary $\{a_i\}_i$, the following holds: either $a_i = 0$ for all i, or there is a k with $a_k = 1$.

LPO might be related to the inability of a computer to perform an unbounded search and decide exactly whether a given real number is non-zero or exactly zero which may be in turn related to the Turing's Halting problem [24]. Let $\{a_i\}_i$ be a binary sequence with at least one 1 at some place which is not a priori known. Let the numbers b, c be defined as follows:

$$b = \frac{1}{4} \sum_{i=0}^{\infty} \frac{1}{i+1} a_{2i+1}, c = \frac{1}{4} \sum_{i=1}^{\infty} \frac{1}{i+1} a_{2i}.$$

If it were known that $b = 0$, then one could deduce that all the odd entries of $\{a_i\}$ are zero and, therefore, since one even entry must be 1, there exists an index $2N$ such that $a_{2N} = 1$, i. e.,

$$b = 0 \implies \forall i, \ a_{2i+1} = 0 \implies \exists N. \ a_{2N} = 1.$$

Similarly, if $c = 0$, then some even entry of $\{a_i\}$ must be 1, i. e.,

$$c = 0 \implies \forall i, \ a_{2i} = 0 \implies \exists N. \ a_{2N+1} = 1.$$

In this minimalistic scenario, the appearance of a 1 can be described by the condition $b = 0 \lor c = 0$ which implies LPO. Consider, for instance, the case of the cost function $J(u) = \min\{u^2 + b, (u-1)^2 + c\}$. If an optimal control action u^* could be computed exactly, such that $J(u^*) = \min J$, then either $u^* \geq \frac{1}{3}$ or $u^* \leq \frac{2}{3}$. If $u^* \geq \frac{1}{3}$, then

$$(u^*)^2 + b \geq \frac{1}{9} + b > 0$$

whence $(u^* - 1)^2 + c = 0 \Rightarrow c = 0$. If $u^* \leq \frac{2}{3}$, then

$$(u^* - 1)^2 + c \geq \frac{1}{9} + c > 0$$

whence $(u^*)^2 + b = 0 \Rightarrow b = 0$. Therefore, equivalence to LPO is shown.