
Fast Approximate Dynamic Programming for Infinite-Horizon Markov Decision Processes

Anonymous Author(s)

Affiliation
Address
email

Abstract

1 In this study, we consider the infinite-horizon, discounted cost, optimal control
2 of stochastic nonlinear systems with separable cost and constraints in the state
3 and input variables. Using the linear-time Legendre transform, we propose a
4 novel numerical scheme for implementation of the corresponding value iteration
5 (VI) algorithm in the conjugate domain. Detailed analyses of the convergence,
6 time complexity, and error of the proposed algorithm are provided. In particular,
7 with a discretization of size X and U for the state and input spaces, respectively,
8 the proposed approach reduces the time complexity of each iteration in the VI
9 algorithm from $\mathcal{O}(XU)$ to $\mathcal{O}(X + U)$, by replacing the minimization operation in
10 the primal domain with a simple addition in the conjugate domain.

11 1 Introduction

12 Value iteration (VI) is one of the most basic and wide-spread algorithms employed for tackling
13 problems in reinforcement learning and optimal control [8, 24] formulated as Markov decision
14 processes (MDPs). The VI algorithm simply involves the consecutive applications of the dynamic
15 programming (DP) operator $\mathcal{T}J(x_t) = \min_{u_t} \{C(x_t, u_t) + \gamma \mathbb{E}J(x_{t+1})\}$, where $C(x_t, u_t)$ is the
16 cost of taking the control action u_t at the state x_t . This fixed point iteration is known to converge
17 to the optimal value function for discount factors $\gamma \in (0, 1)$. However, this algorithm suffers from
18 a high computational cost for large scale finite state spaces. For problems with a continuous state
19 space, the DP operation becomes an infinite-dimensional optimization problem, rendering the exact
20 implementation of VI impossible in most cases. A common approach is to incorporate function
21 approximation techniques and compute the output of the DP operator for a finite sample (i.e., a
22 discretization) of the underlying continuous state space. This approximation again suffers from a
23 high computational cost for fine discretizations of the state space, particularly in high-dimensional
24 problems. We refer the reader to [8, 22] for various approximation schemes for VI.

25 For some problems, however, it is possible to partially address this issue by using duality theory,
26 i.e., approaching the minimization problem in the conjugate domain. In particular, as we will see in
27 Section 3, the minimization in the primal domain in the DP operator can be transformed to a simple
28 addition in the dual domain, at the expense of three conjugate transforms. However, proper application
29 of this transformation relies on efficient numerical algorithms for conjugation. Fortunately, such
30 an algorithm, known as linear-time Legendre transform (LLT) has been developed in late 90s [19].
31 Other than the classical application of LLT (and other fast algorithms for conjugate transform) in
32 solving Hamilton-Jacobi equation [1, 12, 13], these algorithms are used in image processing [20],
33 thermodynamics [11], and optimal transport [16].

34 The application of conjugate duality for the DP problem is not new and actually goes back to
35 Bellman [4]. Further applications of this idea for reducing the computational complexity were later
36 explored in [14, 17]. However, surprisingly, the application of LLT for solving discrete-time optimal

37 control problems, has been limited. In particular, in [10], the authors propose the “fast value iteration”
38 algorithm (without a rigorous analysis of the complexity and error of the proposed algorithm) for a
39 particular class of infinite-horizon problems with state-independent stage cost $C(x, u) = C(u)$ and
40 deterministic linear dynamics $x_{t+1} = Ax_t + Bu_t$, where A is a non-negative, monotone, invertible
41 matrix. More recently, in [18], the authors also considered the application of LLT for solving the DP
42 operation in finite-horizon, optimal control of input-affine dynamics. In particular, they introduced
43 the “discrete conjugate DP” (d-CDP) operator, and provided a detailed analysis of its complexity
44 and error. As we will discuss shortly, the current study is an extension of the corresponding d-CDP
45 algorithm that, among other things, considers infinite horizon, discounted cost problems. We note
46 that the algorithms developed in [15, 20] for distance transform can also potentially tackle the optimal
47 control problems similar to the ones of interest in the current study. In particular, these algorithms
48 require the stage cost to be reformulated as a convex function of the “distance” between the current
49 and next states. While this property might arise naturally, it can generally be restrictive, as it is in the
50 problem class considered in this study. Another line of work that is closely related to ours involves
51 utilizing max-plus algebra in solving deterministic, continuous-state, continuous-time, optimal control
52 problems; see, e.g., [2, 21]. These works exploit the compatibility of the DP operation with max-plus
53 operations, and approximate the value function as a max-plus linear combination. Recently, in [3, 5],
54 the authors used this idea to propose an approximate VI algorithm for continuous-state, deterministic
55 MDPs. In this regard, we note that the proposed approach in the current study also involves max-plus
56 linear approximation of the value function as the maximum of affine functions. The key difference is
57 however that by choosing a grid-like (factorized) set of slopes for the linear terms (the basis of the
58 max-plus linear combination), we take advantage of linear time complexity of LLT in computing the
59 constant terms (the coefficients of the max-plus linear combination).

60 **Main contribution.** We focus on an approximate implementation of VI involving discretization of
61 the state-input space for solving the optimal control problem of discrete-time systems with continuous
62 state-input space. Building upon the earlier work [18], we employ conjugate duality to speed-up VI for
63 problems with separable cost (in state and input) and input-affine dynamics. We propose the conjugate
64 VI (CVI) algorithm based on a modified version of the d-CDP operator introduced in [18], and extend
65 the existing results for *infinite-horizon, discounted cost* problems, while taking into account *stochastic*
66 *dynamics* and *numerical approximation of the conjugate of input cost* in our analysis. In particular,
67 (i) we provide sufficient conditions for the convergence of CVI (Theorem 3.9); (ii) we show that CVI
68 can achieve a linear time complexity of $\mathcal{O}(X + U)$ in each iteration (Theorem 3.10), where X (U) is
69 the cardinality of the discrete state (input) space; (iii) we analyze the error of CVI (Theorem 3.11),
70 and use that result to provide specific guidelines on the construction of the discrete dual domain
71 (Section 3.4); (iv) we provide a MATLAB package for the implementation of the proposed algorithm.

72 **Notations.** The standard inner product in \mathbb{R}^n and the corresponding induced 2-norm are denoted
73 by $\langle \cdot, \cdot \rangle$ and $\|\cdot\|_2$, respectively. $\|\cdot\|_\infty$ denotes the infinity norm. We use the superscript d to denote
74 *finite (discrete)* sets (as in \mathbb{X}^d) and *discrete* functions (as in $h^d : \mathbb{X}^d \rightarrow \mathbb{R}$). We use the superscript g
75 to denote *grid-like* finite sets (as in $\mathbb{X}^g = \prod_{i=1}^n \mathbb{X}_i^g$ where $\mathbb{X}_i^g \subset \mathbb{R}$). We also use $\mathbb{X}_{\text{sub}}^g$ to denote the
76 *sub-grid* of \mathbb{X}^g derived by omitting the smallest and the largest elements of \mathbb{X}^g in each dimension. The
77 cardinality of the finite set \mathbb{X}^d or \mathbb{X}^g is denoted by X . We use $\widetilde{h}^d : \mathbb{R}^n \rightarrow \overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$ to denote a
78 *generic extension* of a discrete function h^d , and $\overline{h}^d : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ to denote *multilinear interpolation*
79 *and extrapolation (LERP)* extension of a discrete function with a *grid-like* domain. Let \mathbb{X}, \mathbb{Y} be
80 two arbitrary sets in \mathbb{R}^n . $\text{co}(\mathbb{X})$ is the convex hull of \mathbb{X} . We use $d(\mathbb{X}, \mathbb{Y}) := \inf_{x \in \mathbb{X}, y \in \mathbb{Y}} \|x - y\|_2$
81 to denote the distance between \mathbb{X} and \mathbb{Y} . The one-sided Hausdorff distance *from* \mathbb{X} *to* \mathbb{Y} is defined
82 as $d_{\text{H}}(\mathbb{X}, \mathbb{Y}) := \sup_{x \in \mathbb{X}} \inf_{y \in \mathbb{Y}} \|x - y\|_2$. Let $h : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ be an extended real-valued function
83 with a non-empty effective domain $\text{dom}(h) = \mathbb{X} := \{x \in \mathbb{R}^n : h(x) < \infty\}$. The range of h is
84 denoted by $\text{Rng}(h) := \max_{x \in \mathbb{X}} h(x) - \min_{x \in \mathbb{X}} h(x)$, and the subdifferential of h at a point $x \in \mathbb{X}$
85 is defined as $\partial h(x) := \{y \in \mathbb{R}^n : h(\tilde{x}) \geq h(x) + \langle y, \tilde{x} - x \rangle, \forall \tilde{x} \in \mathbb{X}\}$. We define $\mathbb{L}(h) :=$
86 $\prod_{i=1}^n [L_i^-(h), L_i^+(h)]$, where $L_i^+(h)$ (resp. $L_i^-(h)$) is the maximum (resp. minimum) slope of
87 the function h along the i -th dimension. Note that $\partial h(x) \subseteq \mathbb{L}(h)$ for all $x \in \mathbb{X}$. We report the
88 complexities using the standard big O notations \mathcal{O} and $\tilde{\mathcal{O}}$, where the latter hides the logarithmic
89 factors. In this study, we are mainly concerned with the dependence of the computational complexities
90 *on the size of the finite sets* involved (discretization of the primal and dual domains). In particular, we
91 ignore the possible dependence of the computational complexities on the dimension of the variables,
92 unless they appear in the power of the size of those discrete sets.

93 We note that the extended version of this article, including the technical proofs, is available in the
 94 supplementary material.

95 2 VI in primal domain

96 We are concerned with the infinite-horizon, discounted cost, optimal control problems of the form

$$J_*(x) = \min \mathbb{E}_{w_t} \left[\sum_{t=0}^{\infty} \gamma^t C(x_t, u_t) \middle| x_0 = x \right]$$

s.t. $x_{t+1} = g(x_t, u_t, w_t)$, $x_t \in \mathbb{X}$, $u_t \in \mathbb{U}$, $w_t \sim \mathbb{P}(\mathbb{W})$, $\forall t \in \{0, 1, \dots\}$,

97 where $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^m$, and $w_t \in \mathbb{R}^l$ are the state, input and disturbance variables at time t ,
 98 respectively; $\gamma \in (0, 1)$ is the discount factor; $C : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}$ is the stage cost; $g : \mathbb{R}^n \times \mathbb{R}^m \times$
 99 $\mathbb{R}^l \rightarrow \mathbb{R}^n$ describes the dynamics; $\mathbb{X} \subset \mathbb{R}^n$ and $\mathbb{U} \subset \mathbb{R}^m$ describe the state and input constraints,
 100 respectively; and, $\mathbb{P}(\cdot)$ is the distribution of the disturbance over the support $\mathbb{W} \subset \mathbb{R}^l$. Assuming the
 101 stage cost C is bounded, the optimal value function solves the Bellman equation $J_* = \mathcal{T}J_*$, where \mathcal{T}
 102 is the DP operator (C and J are extended to infinity outside their effective domains) [6, Prop. 1.2.2]

$$\mathcal{T}J(x) := \min_u \{C(x, u) + \gamma \cdot \mathbb{E}_w J(g(x, u, w))\}, \quad \forall x \in \mathbb{X}. \quad (1)$$

103 Indeed, \mathcal{T} is γ -contractive in the infinity-norm [6, Prop. 1.2.4]. This property then gives rise to the VI
 104 algorithm $J_{k+1} = \mathcal{T}J_k$, which converges to J_* as $k \rightarrow \infty$, for arbitrary initialization J_0 . Moreover,
 105 assuming that the composition $J \circ g$ (for each w) and the cost C are jointly convex in the state and
 106 input variables, \mathcal{T} also preserves convexity [7, Prop. 3.3.1].

107 For numerical implementation of VI, we need to address three issues. First, we need to compute
 108 the expectation in (1). In order to simplify the exposition and include the computational cost of this
 109 operation explicitly, we consider disturbances with finite support in this study:

110 **Assumption 2.1** (Disturbance with finite support). *The disturbance w has a finite support $\mathbb{W}^d \subset \mathbb{R}^l$*
 111 *with a given probability mass function (p.m.f.) $p : \mathbb{W}^d \rightarrow [0, 1]$.*

112 Under the preceding assumption, we have $\mathbb{E}_w J(g(x, u, w)) = \sum_{w \in \mathbb{W}^d} p(w) \cdot J(g(x, u, w))$. The
 113 second and more important issue is that the optimization problem (1) is infinite-dimensional for
 114 the continuous state space \mathbb{X} . This renders the exact implementation of VI impossible, except for a
 115 few cases with available closed-form solutions. A common solution to this problem is to deploy a
 116 sample-based approach, accompanied by a function approximation scheme. To be precise, for a finite
 117 subset \mathbb{X}^d of \mathbb{X} , at each iteration $k \geq 0$, we take the discrete function $J_k^d : \mathbb{X}^d \rightarrow \mathbb{R}$ as the input, and
 118 compute the discrete function $J_{k+1}^d = [\mathcal{T} \widetilde{J}_k^d]^d : \mathbb{X}^d \rightarrow \mathbb{R}$, where $\widetilde{J}_k^d : \mathbb{X} \rightarrow \mathbb{R}$ is an extension of J_k^d .
 119 Finally, for each $x \in \mathbb{X}^d$, we have to solve the minimization problem in (1) over the control input.
 120 Here, again, a common approximation is enumeration over a discretization $\mathbb{U}^d \subset \mathbb{U}$.

121 Incorporating these approximations, we end up with the approximate VI algorithm $J_{k+1}^d = \mathcal{T}^d J_k^d$,
 122 characterized by the *discrete DP* (d-DP) operator

$$\mathcal{T}^d J^d(x) := \min_{u \in \mathbb{U}^d} \{C(x, u) + \gamma \cdot \sum_{w \in \mathbb{W}^d} p(w) \cdot \widetilde{J}^d(g(x, u, w))\}, \quad \forall x \in \mathbb{X}^d. \quad (2)$$

123 The convergence of approximate VI described above depends on the properties of the extension
 124 operation $[\cdot]^d$. In particular, if the extension operation is non-expansive (in the infinity-norm), then \mathcal{T}^d
 125 is also γ -contractive. The error of this approximation also depends on the extension operation and
 126 its representative power. We refer the interested reader to [6, 9, 22] for detailed discussions on the
 127 convergence and error of different approximation schemes for VI.

128 The d-DP operator and the corresponding VI algorithm will be our benchmark for evaluating the
 129 performance of the alternative algorithm developed in this study. To this end, we finish this section
 130 with some remarks on the time complexity of the d-DP operation. Let the time complexity of a single
 131 evaluation of the extension operator $[\cdot]^d$ in (2) be of $\mathcal{O}(E)$. Then, the time complexity of the d-DP
 132 operation (2) is of $\mathcal{O}(XUWE)$. In this regard, note that the scheme described above essentially
 133 involves approximating a continuous-state/action MDP with a finite-state/action MDP, and then
 134 applying the VI algorithm. This, in turn, implies the lower bound $\Omega(XU)$ for the time complexity

135 (corresponding to enumeration over $u \in \mathbb{U}^d$ for each $x \in \mathbb{X}^d$). This lower bound is also compatible
 136 with the best existing time complexities in the literature for VI for finite MDPs; see, e.g., [3, 23].
 137 However, as we will see in the next section, for a particular class of problems, it is possible to exploit
 138 the structure of the underlying continuous system in order to achieve a better time complexity in the
 139 corresponding discretized problem.

140 3 Reducing complexity via conjugate duality

141 In this section, we present the class of problems that allows us to employ conjugate duality and
 142 propose an alternative path for solving the corresponding DP operator. We also present the numerical
 143 scheme for implementing the proposed alternative path, and analyze its convergence, complexity,
 144 and error. We note that the proposed algorithm and its analysis are based on the d-CDP algorithm
 145 presented in [18, Sec. 5] for finite-horizon, optimal control of deterministic systems. Here, we extend
 146 those results for infinite-horizon, discounted cost, optimal control of stochastic systems. Moreover,
 147 unlike [18], our analysis includes the case where the conjugate of input cost is not analytically
 148 available and has to be computed numerically; see [18, Assump. 5.1] for more details.

149 3.1 VI in conjugate domain

150 Throughout this section, we assume that the problem data satisfy the following conditions.

151 **Assumption 3.1** (Problem class). *The problem data has the following properties: (i) The disturbance*
 152 *is additive; that is, $g(x, u, w) = f(x, u) + w$, where $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ describes the deterministic*
 153 *dynamics. (ii) The deterministic dynamics is of the form $f(x, u) = f_s(x) + Bu$, where $f_s : \mathbb{R}^n \rightarrow \mathbb{R}^n$*
 154 *is a Lipschitz continuous, possibly nonlinear map describing the “state” dynamics and $B \in \mathbb{R}^{n \times m}$.*
 155 *(iii) The constraint sets $\mathbb{X} \subset \mathbb{R}^n$ and $\mathbb{U} \subset \mathbb{R}^m$ are compact and convex. Moreover, for each $x \in \mathbb{X}$,*
 156 *the set of admissible inputs $\mathbb{U}(x) := \{u \in \mathbb{U} : g(x, u, w) \in \mathbb{X}, \forall w \in \mathbb{W}^d\}$ is nonempty. (iv) The*
 157 *stage cost C is separable in state and input; that is, $C(x, u) = C_s(x) + C_i(u)$, where the state cost*
 158 *$C_s : \mathbb{X} \rightarrow \mathbb{R}$ and the input cost $C_i : \mathbb{U} \rightarrow \mathbb{R}$ are Lipschitz continuous and convex.*

159 Some remarks are in order regarding the preceding assumptions. We first note that the setting of
 160 Assumption 3.1 goes beyond the classical LQR. In particular, it includes nonlinear dynamics, state
 161 and input constraints, and non-quadratic stage costs. Second, the properties laid out in Assumption 3.1
 162 imply that the set of admissible inputs $\mathbb{U}(x)$ is a compact set for each $x \in \mathbb{X}$. This, in turn, implies
 163 that the optimal value in (1) is achieved if $J : \mathbb{X} \rightarrow \mathbb{R}$ is also assumed to be lower semi-continuous.

164 For the problem class of Assumption (3.1), we can use duality theory to present an alternative path
 165 for computing the output of the DP operator. This path forms the basis for the proposed algorithm.
 166 To this end, let us fix $x \in \mathbb{X}$ and consider the following reformulation of the optimization problem (1)

$$\mathcal{T}J(x) = C_s(x) + \min_{u, z} \{C_i(u) + \gamma \cdot \mathbb{E}_w J(z + w) : z = f(x, u)\},$$

167 where we used additivity of disturbance and separability of stage cost in Assumptions 3.1-(i,iv). The
 168 corresponding dual problem then reads as

$$\widehat{\mathcal{T}}J(x) := C_s(x) + \max_y \min_{u, z} \{C_i(u) + \gamma \cdot \mathbb{E}_w J(z + w) + \langle y, f(x, u) - z \rangle\}, \quad (3)$$

169 where $y \in \mathbb{R}^n$ is the dual variable corresponding to the equality constraint. For the dynamics of
 170 Assumption 3.1-(ii), we can obtain the following representation for the dual problem.

171 **Proposition 3.2** (CDP operator). *The dual problem (3) equivalently reads as*

$$\epsilon(x) := \gamma \cdot \mathbb{E}_w J(x + w), \quad x \in \mathbb{X}, \quad (4a)$$

$$\phi(y) := C_i^*(-B^\top y) + \epsilon^*(y), \quad y \in \mathbb{R}^n, \quad (4b)$$

$$\widehat{\mathcal{T}}J(x) = C_s(x) + \phi^*(f_s(x)), \quad x \in \mathbb{X}, \quad (4c)$$

172 where $h^*(y) = \sup_x \{\langle y, x \rangle - h(x)\}$ is the (convex) conjugate of h .

173 Following [18], we call the operator $\widehat{\mathcal{T}}$ in (4) the *conjugate DP* (CDP) operator. We next provide an
 174 alternative representation of the CDP operator that captures the essence of this operation.

175 **Proposition 3.3** (CDP reformulation). *The CDP operator $\widehat{\mathcal{T}}$ equivalently reads as*

$$\widehat{\mathcal{T}}J(x) = C_s(x) + \min_u \{C_i(u) + \gamma \cdot [\mathbb{E}_w J(\cdot + w)]^{**}(f(x, u))\}, \quad (5)$$

176 where $[\cdot]^{**} = [[\cdot]^*]^*$ denotes the biconjugate operation.

177 The preceding result implies that the indirect path through the conjugate domain essentially involves
 178 substituting the (expectation of the) value function by its biconjugate. In particular, it points to a
 179 sufficient condition for zero duality gap.

180 **Corollary 3.4** (Equivalence of \mathcal{T} and $\widehat{\mathcal{T}}$). *If $J : \mathbb{X} \rightarrow \mathbb{R}$ is convex, then $\widehat{\mathcal{T}}J = \mathcal{T}J$.*

181 Hence, $\widehat{\mathcal{T}}$ has the same properties as \mathcal{T} if J is convex. In particular, if $\widehat{\mathcal{T}}$ preserves convexity, then
 182 the *conjugate* VI (CVI) algorithm $J_{k+1} = \widehat{\mathcal{T}}J_k$ converges to the optimal value function J_* , with
 183 arbitrary convex initialization J_0 . For $\widehat{\mathcal{T}}$ to preserve convexity, a sufficient condition is the convexity
 184 of $J \circ f$ (jointly in x and u), given that J is convex (e.g., for linear dynamics $f(x, u) = Ax + Bu$).
 185 We note that, if $\widehat{\mathcal{T}}$ does not preserve convexity, then the alternative path suffers from duality gap.

186 3.2 CVI algorithm

187 The approximate CVI algorithm involves consecutive applications of an approximate implementation
 188 of the CDP operator (4) until some termination condition is satisfied. Algorithm 1 provides the
 189 pseudo-code of this procedure. In particular, we consider solving (4) for a finite set $\mathbb{X}^d \subset \mathbb{X}$, and
 190 terminate the iterations when the difference between two consecutive discrete value functions (in the
 191 infinity-norm) is less than a given constant $\epsilon_t > 0$. Since we are working with a finite subset
 192 of the state space, we can restrict the feasibility condition of Assumption 3.1-(iii) to all $x \in \mathbb{X}^d$:

193 **Assumption 3.5** (Feasible discretization). *We have $\mathbb{U}(x) \neq \emptyset$ for all $x \in \mathbb{X}^d$.*

194 In what follows, we describe the main steps within the initialization and iterations of Algorithm 1.
 195 In particular, the conjugate operations in (4) are handled numerically via the linear-time Legendre
 196 transform (LLT) algorithm [19]. LLT is an efficient algorithm for computing the *discrete* conjugate
 197 function over a finite *grid-like* dual domain. Precisely, to compute the conjugate of the function $h : \mathbb{X} \rightarrow \mathbb{R}$, LLT takes its discretization $h^d : \mathbb{X}^d \rightarrow \mathbb{R}$ as an input, and outputs $h^{d*d} : \mathbb{Y}^g \rightarrow \mathbb{R}$, for the grid-like dual domain \mathbb{Y}^g . We refer the reader to [19] for a detailed description of LLT.

200 The main steps of the proposed approximate implementation of the CDP operator (6) are as follows:
 201 **(i)** For the expectation operation in (4a), by Assumption 2.1, we again have $\mathbb{E}_w J(\cdot + w) = \sum_{w \in \mathbb{W}^d} p(w) \cdot J(\cdot + w)$. Hence, we need to pass the value function $J^d : \mathbb{X}^d \rightarrow \mathbb{R}$ through the
 202 “scaled expectation filter” to obtain $\epsilon^d : \mathbb{X}^d \rightarrow \overline{\mathbb{R}}$ in (6a) as an approximation of ϵ in (4a). Notice that
 204 here we are using an extension $\widetilde{J}^d : \mathbb{X} \rightarrow \mathbb{R}$ of J^d (recall that we only have access to the discrete
 205 value function J^d). **(ii)** In order to compute ϕ in (4b), we need access to two conjugate functions.
 206 First, for ϵ^* , we use the approximation $\epsilon^{d*d} : \mathbb{Y}^g \rightarrow \mathbb{R}$ in (6b), by applying LLT to the data points
 207 $\epsilon^d : \mathbb{X}^d \rightarrow \overline{\mathbb{R}}$ for a properly chosen state dual grid $\mathbb{Y}^g \subset \mathbb{R}^n$. We also need the conjugate C_1^* of the
 208 input cost. If this function is not analytically available, we approximate it as follows: For a properly
 209 chosen input dual grid $\mathbb{V}^g \subset \mathbb{R}^m$, we employ LLT to compute $C_1^{d*d} : \mathbb{V}^g \rightarrow \mathbb{R}$ in (6c), using the data
 210 points $C_1^d : \mathbb{U}^d \rightarrow \mathbb{R}$, where \mathbb{U}^d is a finite subset of \mathbb{U} . With these conjugate functions at hand, we
 211 can now compute $\varphi^d : \mathbb{Y}^g \rightarrow \mathbb{R}$ in (6d), as an approximation of ϕ in (4b). In particular, notice that
 212 we use the LERP extension $\overline{C_1^{d*d}}$ of C_1^{d*d} to approximate C_1^{d*} at the required point $-B^\top y$ for each
 213 $y \in \mathbb{Y}^g$. **(i)** To be able to compute the output according to (4c), we need to perform another conjugate
 214 transform. In particular, we need the value of ϕ^* at $f_s(x)$ for $x \in \mathbb{X}^d$. Here, we use the approximation
 215 $\varphi^{d*d} : \mathbb{Z}^g \rightarrow \mathbb{R}$ in (6e), by applying LLT to the data points $\varphi^d : \mathbb{Y}^g \rightarrow \mathbb{R}$ for a properly chosen grid
 216 $\mathbb{Z}^g \subset \mathbb{R}^n$. Finally, we use the LERP extension $\overline{\varphi^{d*d}}$ of φ^{d*d} to approximate φ^{d*} at the required point
 217 $f_s(x)$ for each $x \in \mathbb{X}^d$, and compute $\widehat{\mathcal{T}}^d J^d$ in (6f) as an approximation of $\widehat{\mathcal{T}}J$ in (4c). With these

Algorithm 1 CVI: Approximate VI in conjugate domain

Input: dynamics $f_s : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $B \in \mathbb{R}^{n \times m}$; finite state space $\mathbb{X}^d \subset \mathbb{X}$; finite input space $\mathbb{U}^d \subset \mathbb{U}$; state cost function $C_s^d : \mathbb{X}^d \rightarrow \mathbb{R}$; input cost function $C_i^d : \mathbb{U}^d \rightarrow \mathbb{R}$; finite disturbance space \mathbb{W}^d and its p.m.f. $p : \mathbb{W}^d \rightarrow [0, 1]$; discount factor γ ; termination bound e_t .

Output: discrete value function $\widehat{J}^d : \mathbb{X}^d \rightarrow \mathbb{R}$.

initialization:

- 1: construct the grids $\mathbb{V}^g, \mathbb{Z}^g$, and \mathbb{Y}^g ;
- 2: use LLT to compute $C_i^{d*d} : \mathbb{V}^g \rightarrow \mathbb{R}$ from $C_i^d : \mathbb{U}^d \rightarrow \mathbb{R}$;
- 3: $J^d(x) \leftarrow 0$ and $J_+^d(x) \leftarrow C_s^d(x) - \min C_i^d$ for $x \in \mathbb{X}^d$;

iteration:

- 4: **while** $\|J_+^d - J^d\|_\infty \geq e_t$ **do**
- 5: $J^d \leftarrow J_+^d$;
- 6: *d-CDP operation:*
- 6: $\varepsilon^d(x) \leftarrow \gamma \cdot \sum_{w \in \mathbb{W}^d} p(w) \cdot \widetilde{J}^d(x+w)$ for $x \in \mathbb{X}^d$; use LLT to compute $\varepsilon^{d*d} : \mathbb{Y}^g \rightarrow \mathbb{R}$ from $\varepsilon^d : \mathbb{X}^d \rightarrow \mathbb{R}$;
- 7: **for each** $y \in \mathbb{Y}^g$ **do**
- 8: use LERP to compute $\overline{C_i^{d*d}}(-B^\top y)$ from $C_i^{d*d} : \mathbb{V}^g \rightarrow \mathbb{R}$; $\varphi^d(y) \leftarrow \overline{C_i^{d*d}}(-B^\top y) + \varepsilon^{d*d}(y)$;
- 9: **end for**
- 10: use LLT to compute $\varphi^{d*d} : \mathbb{Z}^g \rightarrow \mathbb{R}$ from $\varphi^d : \mathbb{Y}^g \rightarrow \mathbb{R}$;
- 11: **for each** $x \in \mathbb{X}^d$ **do**
- 12: use LERP to compute $\overline{\varphi^{d*d}}(f_s(x))$ from $\varphi^{d*d} : \mathbb{Z}^g \rightarrow \mathbb{R}$; $J_+^d(x) \leftarrow C_s^d(x) + \overline{\varphi^{d*d}}(f_s(x))$;
- 13: **end for**
- 14: **end while**
- 15: $\widehat{J}^d \leftarrow J_+^d$.

218 approximations, we can introduce the *discrete* CDP (d-CDP) operator as follows

$$\varepsilon^d(x) := \gamma \cdot \sum_{w \in \mathbb{W}^d} p(w) \cdot \widetilde{J}^d(x+w), \quad x \in \mathbb{X}^d, \quad (6a)$$

$$\varepsilon^{d*d}(y) = \max_{x \in \mathbb{X}^d} \{ \langle x, y \rangle - \varepsilon^d(x) \}, \quad y \in \mathbb{Y}^g, \quad (6b)$$

$$C_i^{d*d}(v) = \max_{u \in \mathbb{U}^d} \{ \langle u, v \rangle - C_i^d(u) \}, \quad v \in \mathbb{V}^g, \quad (6c)$$

$$\varphi^d(y) := \overline{C_i^{d*d}}(-B^\top y) + \varepsilon^{d*d}(y), \quad y \in \mathbb{Y}^g, \quad (6d)$$

$$\varphi^{d*d}(z) = \max_{y \in \mathbb{Y}^g} \{ \langle y, z \rangle - \varphi^d(y) \}, \quad z \in \mathbb{Z}^g, \quad (6e)$$

$$\widehat{\mathcal{T}}^d J^d(x) := C_s^d(x) + \overline{\varphi^{d*d}}(f_s(x)), \quad x \in \mathbb{X}^d. \quad (6f)$$

219 We will discuss the proper construction of the grids $\mathbb{Y}^g, \mathbb{V}^g$, and \mathbb{Z}^g in Section 3.4.

220 3.3 Analysis of CVI algorithm

221 We now provide our main theoretical results concerning the convergence, complexity, and error of
 222 the proposed algorithm. Let us begin with presenting the assumptions to be called in this subsection.

223 **Assumption 3.6** (Grids). *Consider the following properties for the grids in Algorithm 1 (consult the*
 224 *Notations in Section 1): (i) The grid \mathbb{V}^g is constructed such that $\text{co}(\mathbb{V}_{\text{sub}}^g) \supseteq \mathbb{L}(C_i^d)$. (ii) The grid \mathbb{Z}^g*
 225 *is constructed such that $\text{co}(\mathbb{Z}^g) \supseteq f_s(\mathbb{X}^d)$. (iii) The construction of $\mathbb{Y}^g, \mathbb{V}^g$, and \mathbb{Z}^g requires at most*
 226 *$\mathcal{O}(X + U)$ operations. The cardinality of the grids \mathbb{Y}^g and \mathbb{Z}^g (resp. \mathbb{V}^g) in each dimension is the*
 227 *same as that of \mathbb{X}^d (resp. \mathbb{U}^d) in that dimension so that $Y, Z = X$ and $V = U$.*

228 **Assumption 3.7** (Extension operator). *Consider the following properties for the operator $\widetilde{[\cdot]}$ in (6a):*
 229 *(i) The extension operator is non-expansive w.r.t. the infinity norm; (ii) Given a function $J : \mathbb{X} \rightarrow \mathbb{R}$*
 230 *and its discretization $J^d : \mathbb{X}^d \rightarrow \mathbb{R}$, we have $\|J - \widetilde{J}^d\|_\infty \leq e_e$ for some constant $e_e \geq 0$.*

231 **Assumption 3.8** (Dynamics). *Given a convex function $J : \mathbb{X} \rightarrow \mathbb{R}$, the composition $J \circ f$ is jointly*
 232 *convex in the state and input variables, where f describes the deterministic dynamics.*

233 Our first result concerns the contractiveness of the d-CDP operator.

234 **Theorem 3.9** (Convergence). *Let Assumptions 3.6-(ii) and 3.7-(i) hold. Then, the d-CDP operator (6)*
 235 *is γ -contractive w.r.t. the infinity-norm.*

236 The preceding theorem implies that the approximate CVI Algorithm 1 is indeed convergent given that
 237 the required conditions are satisfied. In particular, for deterministic dynamics, $\text{co}(\mathbb{Z}^g) \supseteq f_s(\mathbb{X}^d)$ is
 238 sufficient for Algorithm 1 to be convergent. We next consider the time complexity of our algorithm.

239 **Theorem 3.10** (Complexity). *Let Assumption 3.6-(iii) hold. Also assume that each evaluation of the*
 240 *extension operator $[\cdot]$ in (6a) requires $\mathcal{O}(E)$ operations. Then, the time complexity of initialization*
 241 *and each iteration in Algorithm 1 are of $\mathcal{O}(X + U)$ and $\tilde{\mathcal{O}}(XWE)$, respectively.*

242 The requirements of Assumption 3.6-(iii) will be discussed in Section 3.4. Recall that each iteration
 243 of VI (in primal domain) has a complexity of $\mathcal{O}(XUWE)$, where E denotes the complexity of the
 244 extension operation used in (2). This observation points to a basic characteristic of the proposed
 245 approach: CVI reduces the quadratic complexity of VI to a linear one by replacing the minimization
 246 operation in the primal domain with a simple addition in the conjugate domain. Hence, for problem
 247 class of Assumption 3.1, CVI is expected to lead to a reduction in the computational cost. We
 248 note that CVI, like VI and other approximation schemes that utilize discretization/abstraction of the
 249 continuous state and input spaces, still suffers from the so-called ‘‘curse of dimensionality.’’ This is
 250 because the sizes X and U of the discretizations increase exponentially with the dimensions n and
 251 m of the corresponding space. However, for CVI, this exponential increase is of rate $\max\{m, n\}$,
 252 compared to the rate $m + n$ for VI.

253 **Theorem 3.11** (Error). *Let Assumptions 3.6-(i,ii), 3.7-(i), and 3.8 hold. Consider the true optimal*
 254 *value function $J_* = \mathcal{T}J_* : \mathbb{X} \rightarrow \mathbb{R}$ and its discretization $J_*^d : \mathbb{X}^d \rightarrow \mathbb{R}$, and let Assumption 3.7-(ii)*
 255 *hold for J_* . Also, let $\hat{J}^d : \mathbb{X}^d \rightarrow \mathbb{R}$ be the output of Algorithm 1 with grids $\mathbb{Y}^g, \mathbb{V}^g$, and \mathbb{Z}^g . Then,*

$$\|\hat{J}^d - J_*^d\|_\infty \leq \frac{\gamma(e_e + e_t) + e_d}{1 - \gamma}, \quad (7)$$

256 where $e_d = e_u + e_v + e_x + e_y + e_z$, and

$$\begin{aligned} e_u &= c_u \cdot d_H(\mathbb{U}, \mathbb{U}^d), \quad e_v = c_v \cdot d_H(\text{co}(\mathbb{V}^g), \mathbb{V}^g), \quad e_x = c_x \cdot d_H(\mathbb{X}, \mathbb{X}^d), \\ e_y &= c_y \cdot \max_{x \in \mathbb{X}^d} d(\partial(V_* - C_s)(x), \mathbb{Y}^g), \quad e_z = c_z \cdot d_H(f_s(\mathbb{X}^d), \mathbb{Z}^g), \end{aligned} \quad (8)$$

257 with constants $c_u, c_v, c_x, c_y, c_z > 0$ depending on the problem data.

258 Let us first note that Assumptions 3.6-(i,ii) on the grids \mathbb{V}^g and \mathbb{Z}^g are required for bounding the error
 259 of approximate discrete conjugations using LERP in (6d) and (6f); for more details see the proof
 260 of Theorem 3.11 in Section 5 of the extended manuscript in the supplementary material. Moreover,
 261 Assumption 3.8 implies that the DP and CDP operators preserve convexity, and they both have the
 262 true optimal value function J_* as their fixed point (i.e., the duality gap is zero). Otherwise, the
 263 proposed scheme can suffer from large errors due to dualization. The remaining sources of error in
 264 the proposed approximate implementation of CVI are captured by the three error terms in (7): **(i)** e_e
 265 is due to the approximation of the value function using the extension operator $[\cdot]$; **(ii)** e_t corresponds
 266 to the termination of the algorithm after a finite number of iterations; **(i)** e_d captures the error due to
 267 the discretization of the primal and dual state and input domains.

268 3.4 Construction of the grids

269 In this subsection, we provide specific guidelines for the construction of the grids $\mathbb{Y}^g, \mathbb{V}^g$ and \mathbb{Z}^g . We
 270 note that these discrete sets must be *grid-like* since they form the dual grid for the three conjugate
 271 transforms that are handled using LLT. The presented guidelines aim to minimize the error terms in
 272 (8) while taking into account the properties laid out in Assumption 3.6. In particular, the schemes
 273 described below satisfy the requirements of Assumption 3.6-(iii).

274 **Construction of \mathbb{V}^g .** Assumption 3.6-(i) and the error term e_v in (8) suggest that we find the
 275 smallest input dual grid \mathbb{V}^g such that $\text{co}(\mathbb{V}_{\text{sub}}^g) \supseteq \mathbb{L}(C_i^d)$. This latter condition essentially
 276 means that \mathbb{V}^g must ‘‘more than cover the range of slopes’’ of the function C_i^d ; recall that
 277 $\mathbb{L}(C_i^d) = \prod_{j=1}^m [\mathbb{L}_j^-(C_i^d), \mathbb{L}_j^+(C_i^d)]$, where $\mathbb{L}_j^-(C_i^d)$ (resp. $\mathbb{L}_j^+(C_i^d)$) is the minimum (resp. maxi-
 278 mum) ‘‘slope’’ of C_i^d along the j -th dimension. Hence, we need to compute/approximate $\mathbb{L}_j^\pm(C_i^d)$ for

279 $j = 1, \dots, m$. A conservative approximation for $L_j^-(C_i^d)$ (resp. $L_j^+(C_i^d)$) is $L_j^-(C_i) = \min_{\partial C_i / \partial u_j}$
 280 (resp. $L_j^+(C_i) = \max_{\partial C_i / \partial u_j}$), assuming C_i is differentiable. Alternatively, we can directly use the
 281 discrete input cost C_i^d for computing $L_j^\pm(C_i^d)$. In particular, if the domain $\mathbb{U}^d = \mathbb{U}^g = \prod_{j=1}^m \mathbb{U}_j^g$
 282 of C_i^d is grid-like, we can use the fact that C_i is convex, and take $L_j^-(C_i^d)$ (resp. $L_j^+(C_i^d)$) to be
 283 the minimum first forward difference (resp. maximum last backward difference) of C_i^d along the
 284 j -th dimension (this scheme requires $\mathcal{O}(U)$ operations). Having $L_j^\pm(C_i^d)$ at our disposal, we can
 285 then construct $\mathbb{V}_{\text{sub}}^g = \prod_{j=1}^m \mathbb{V}_{\text{sub}j}^g$ such that, in each dimension j , $\mathbb{V}_{\text{sub}j}^g$ is uniform with the same
 286 cardinality as \mathbb{U}_j^g , and $\text{co}(\mathbb{V}_{\text{sub}j}^g) = [L_j^-(C_i^d), L_j^+(C_i^d)]$. Finally, we construct \mathbb{V}^g by extending $\mathbb{V}_{\text{sub}}^g$
 287 uniformly in each dimension (by adding a smaller and a larger element to $\mathbb{V}_{\text{sub}}^g$ in each dimension).

288 **Construction of \mathbb{Z}^g .** According to Assumption 3.6-(ii), the grid \mathbb{Z}^g must be constructed such that
 289 $\text{co}(\mathbb{Z}^g) \supseteq f_s(\mathbb{X}^d)$. This can be simply done by finding the vertices of the smallest box that contains
 290 the set $f_s(\mathbb{X}^d)$. Those vertices give the range of \mathbb{Z}^g in each dimension. We can then, for example,
 291 take \mathbb{Z}^g to be the uniform grid with the same cardinality as \mathbb{V}^g in each dimension (so that $Z = Y$).
 292 This way, $d_H(f_s(\mathbb{X}^d), \mathbb{Z}^g) \leq d_H(\text{co}(\mathbb{Z}^g), \mathbb{Z}^g)$, and hence e_z in (8) reduces by using finer grids \mathbb{Z}^g .
 293 This construction has a time complexity of $\mathcal{O}(X)$.

294 **Construction of \mathbb{Y}^g .** Construction of the state dual grid \mathbb{Y}^g is more involved. According to Theo-
 295 rem 3.11, we need to choose a grid that minimizes e_y in (8). This can be done by choosing \mathbb{Y}^g such
 296 that $\mathbb{Y}^g \cap \partial(J_\star - C_s) \neq \emptyset$ for all $x \in \mathbb{X}^d$ so that $e_y = 0$. Even if we had access to the optimal value
 297 function J_\star , satisfying such a condition could lead to dual grids $\mathbb{Y}^g \subset \mathbb{R}^n$ of size $\mathcal{O}(X^n)$. Such
 298 a large size violates Assumption 3.6-(iii) on the size of \mathbb{Y}^g , and essentially renders the proposed
 299 algorithm impractical for dimensions $n \geq 2$. A more practical condition is $\text{co}(\mathbb{Y}^g) \cap \partial(J_\star - C_s) \neq \emptyset$
 300 for all $x \in \mathbb{X}^d$ so that $\max_{x \in \mathbb{X}^d} d(\partial(J_\star - C_s)(x), \mathbb{Y}^g) \leq d_H(\text{co}(\mathbb{Y}^g), \mathbb{Y}^g)$, and hence e_y reduces
 301 by using finer grids \mathbb{Y}^g . The latter condition is satisfied if $\text{co}(\mathbb{Y}^g) \supseteq \mathbb{L}(J_\star - C_s)$, i.e., if \mathbb{Y}^g “covers
 302 the range of slopes” of $(J_\star - C_s)$. Hence, we need to approximate the range of slopes of $(J_\star - C_s)$. To
 303 this end, we first use the fact that J_\star is the fixed point of DP operator (1) to approximate $\text{Rng}(J_\star - C_s)$
 304 by $R = \frac{\text{Rng}(C_i^d) + \gamma \cdot \text{Rng}(C_s^d)}{1 - \gamma}$. We then construct the grid $\mathbb{Y}^g = \prod_{i=1}^n \mathbb{Y}_i^g$ such that, for each dimension
 305 i , we have $\pm \alpha R / \Delta_{\mathbb{X}_i^d} \in \text{co}(\mathbb{Y}_i^g)$, where $\Delta_{\mathbb{X}_i^d}$ (with some abuse of notation) denotes the diameter of
 306 the projection of \mathbb{X}^d on the i -th dimension, and $\alpha > 0$ is a scaling factor mainly depending on the
 307 dimension of the state space. This construction has a one-time computational cost of $\mathcal{O}(X + U)$.

308 **Remark 3.12** (Dynamic construction of \mathbb{Y}^g). *Alternatively, we can construct \mathbb{Y}^g dynamically at*
 309 *each iteration in order to minimize the corresponding error in each application of the d-CDP*
 310 *operator given by $e_y = c_y \cdot \max_{x \in \mathbb{X}^d} d(\partial(\mathcal{T}J - C_s)(x), \mathbb{Y}^g)$; see Lemma 5.6 and Proposition 5.8*
 311 *of the extended manuscript in the supplementary material. This means that the construction \mathbb{Y}^g in*
 312 *Algorithm 1 is moved from line 1 to inside the iterations, after line 5. Similar to the static scheme*
 313 *described above, the aim here is to construct \mathbb{Y}^g such that $\text{co}(\mathbb{Y}^g) \supseteq \mathbb{L}(\mathcal{T}J - C_s)$. Since we do*
 314 *not have access $\mathcal{T}J$ at our disposal (it is the output of the current iteration), we can again use the*
 315 *definition of the DP operator (2) to approximate $\text{Rng}(\mathcal{T}J - C_s)$ by $R = \text{Rng}(C_i^d) + \gamma \cdot \text{Rng}(V^d)$*
 316 *where V^d is the output of the previous iteration. We then construct the grid $\mathbb{Y}^g = \prod_{i=1}^n \mathbb{Y}_i^g$ such*
 317 *that, for each dimension i , we again have $\pm \alpha R / \Delta_{\mathbb{X}_i^d} \in \text{co}(\mathbb{Y}_i^g)$. This construction has a one-time*
 318 *computational cost of $\mathcal{O}(U)$ for computing $\text{Rng}(C_i^d)$ and a per iteration computational cost of $\mathcal{O}(U)$*
 319 *for computing $\text{Rng}(V^d)$. Notice, however, that under this dynamic construction, the error bound*
 320 *of Theorem 3.11 does not hold true. More importantly, with a dynamic grid \mathbb{Y}^g that varies in each*
 321 *iteration, there is no guarantee for CVI to converge. However, as we will see in the numerical*
 322 *examples, the proposed scheme leads to (possibly non-monotone) convergence of CVI.*

323 4 Numerical simulations

324 We now showcase the application of the CVI Algorithm 1 in solving the optimal control problem of a
 325 noisy inverted pendulum, and compare its performance with the VI algorithm. The details of these
 326 simulations (and more numerical examples) are provided in Section 4 of the extended manuscript in
 327 the supplementary material. The MATLAB package for the implementation of the proposed CVI
 328 algorithm is provided in the supplementary material. We also note that for the discrete conjugation
 329 (LLT), we used the MATLAB package (in particular, the LLTd routine) provided in [19].

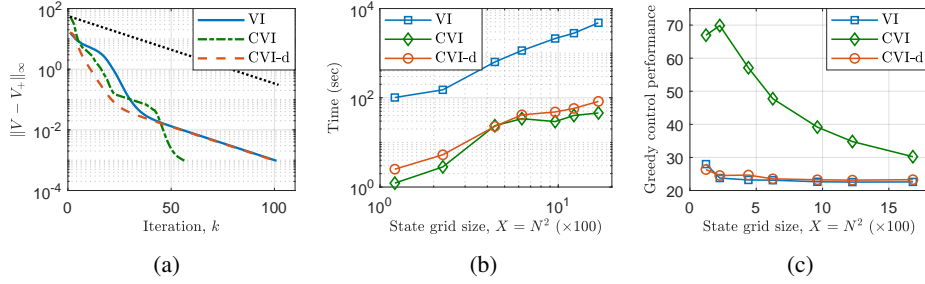


Figure 1: VI vs. CVI: (a) Convergence rate for grid size $X = 41^2$; (b) Running time; (c) Average cost of fifty instances of the control problem with random initial states over $T = 400$ time steps. (d) Difference between outputs of VI and CVI algorithms (over the discrete state space \mathbb{X}^g). The black dotted line in (a) corresponds to exponential convergence with rate $\gamma = 0.95$.

330 We use the setup (model and stage cost) of [18, App. C.2.2], with discount factor $\gamma = 0.95$ and
331 termination bound $e_t = 0.001$. In particular, the state and input costs are both quadratic ($\|\cdot\|_2^2$),
332 and the discrete-time, nonlinear dynamics is of the form $x^+ = f_s(x) + Bu + w$, where where
333 $f_s(x_1, x_2) = [x_1 + \alpha_{12}x_2, \alpha_{21} \sin x_1 + \alpha_{22}x_2]^\top$ and $B = [0, \beta]^\top$ ($\alpha_{12}, \alpha_{21}, \alpha_{22}, \beta \in \mathbb{R}$). State
334 and input constraints are described by $\mathbb{X} = [-\frac{\pi}{3}, \frac{\pi}{3}] \times [-\pi, \pi] \subset \mathbb{R}^2$ and $\mathbb{U} = [-3, 3] \subset \mathbb{R}$. The
335 disturbance has a uniform distribution over the finite support $\mathbb{W}^g = \{0, \pm 0.025\frac{\pi}{3}, \pm 0.05\frac{\pi}{3}\} \times$
336 $\{0, \pm 0.025\pi, \pm 0.05\pi\} \subset \mathbb{R}^2$ of size $W = 5^2$. We use uniform, grid-like discretizations \mathbb{X}^g and \mathbb{U}^g
337 for the state and input spaces such that $\text{co}(\mathbb{X}^g) = [-\frac{\pi}{4}, \frac{\pi}{4}] \times [-\pi, \pi] \subset \mathbb{X}$ and $\text{co}(\mathbb{U}^g) = \mathbb{U}$. This
338 choice of discrete state spaces \mathbb{X}^g particularly satisfies the feasibility condition of Assumption 3.5
339 Also, we use *nearest neighbor* extension (which is non-expansive) for the extension operators in (2)
340 for VI and in (6a) for CVI. The grids $\mathbb{V}^g \subset \mathbb{R}$ and $\mathbb{Z}^g, \mathbb{Y}^g \subset \mathbb{R}^2$ are also constructed uniformly,
341 following the guidelines of Section 3.4 (with $\alpha = 1$). In particular, we also consider the *dynamic*
342 scheme of Remark 3.12 for the construction of \mathbb{Y}^g (hereafter, referred to as CVI-d). Moreover, in
343 each implementation of VI and CVI, all of the involved grids are chosen to be of the same size in
344 each dimension, i.e., $X = Y = Z = N^2$ and $U = V = N$. We are particularly interested in the
345 performance of these algorithms, as the size of the discretizations increases.

346 The results of our numerical simulations are shown in Figure 1. As shown in Figures 1a, both VI and
347 CVI algorithms are indeed convergent with a rate greater than or equal to $\gamma = 0.95$. In particular,
348 CVI terminates in 57 iterations, compared to 101 iterations required for VI to terminate. As expected,
349 this faster convergence, combined with the lower time complexity of CVI in each iteration, lead to a
350 significant reduction in the running time of this algorithm compared to VI. This effect can be clearly
351 seen in Figure 1b, where the run-time of CVI for $N = 41^2$ is an order of magnitude less than that of
352 VI for $N = 11^2$. Since we do not have access to the true optimal value function, in order to evaluate
353 the performance of the outputs of the VI and CVI algorithm, we consider the performance of the
354 greedy policy w.r.t. the discrete value function J^d computed using these algorithms. Figure 1c reports
355 the average cost of fifty instances of the optimal control problem with greedy control actions. As can
356 be seen, the reduction in the running time in CVI comes with an increase in the cost of underlying
357 optimal control problem particularly for coarse discretizations of the state space.

358 Let us now consider the effect of *dynamic* construction of the state dual grid \mathbb{Y}^g . As can be seen in
359 Figure 1a, using a dynamic \mathbb{Y}^g leads to a slower convergence (CVI-d converges in 100 iterations).
360 We note that our simulations show that for the *deterministic* system (with dynamics $x^+ = f_s(x) + Bu$),
361 CVI-d has a similar converge rate as that of CVI. Moreover, for greater values of the scaling factor
362 (e.g., $\alpha = 3$), we observed the non-monotone convergence of CVI-d. See the extended manuscript
363 in the supplementary material for these results. We also note that the relative behaviour of the
364 convergence rates in Figures 1a was also seen for other grid sizes in the discretization scheme.
365 However, thanks to the low time complexity of CVI in each iteration, we see a small increase in the
366 running time of CVI-d compared to CVI; see Figure 1b. More importantly, as depicted in Figure 1c,
367 CVI-d shows almost the same performance as VI when it comes to the quality of the greedy actions
368 generated using the outputs of these algorithms. This is because the dynamic construction of \mathbb{Y}^g
369 in CVI-d uses the available computational power (related to size of the discretization) smartly by
370 finding the smallest grid \mathbb{Y}^g in each iteration, in order to minimize the error of that same iteration.

371 **References**

- 372 [1] Y. Achdou, F. Camilli, and L. Corrias. On numerical approximation of the Hamilton–Jacobi–transport
 373 system arising in high frequency approximations. *Discrete & Continuous Dynamical Systems-Series B*,
 374 19(3), 2014.
- 375 [2] M. Akian, S. Gaubert, and A. Lakhoua. The max-plus finite element method for solving deterministic
 376 optimal control problems: Basic properties and convergence analysis. *SIAM Journal on Control and*
 377 *Optimization*, 47(2):817–848, 2008.
- 378 [3] F. Bach. Max-plus matching pursuit for deterministic Markov decision processes. *arXiv preprint*
 379 *arXiv:1906.08524*, 2019.
- 380 [4] R. Bellman and W. Karush. Mathematical programming and the maximum transform. *Journal of the*
 381 *Society for Industrial and Applied Mathematics*, 10(3):550–567, 1962.
- 382 [5] E. Berthier and F. Bach. Max-plus linear approximations for deterministic continuous-state markov
 383 decision processes. *IEEE Control Systems Letters*, pages 1–1, 2020.
- 384 [6] D. P. Bertsekas. *Dynamic Programming and Optimal Control, Vol. II*. Athena Scientific, Belmont, MA,
 385 3rd edition, 2007.
- 386 [7] D. P. Bertsekas. *Convex Optimization Theory*. Athena Scientific, Belmont, MA, 2009.
- 387 [8] D. P. Bertsekas. *Reinforcement Learning and Optimal Control*. Athena Scientific, Belmont, MA, 2019.
- 388 [9] L. Busoniu, R. Babuska, B. De Schutter, and D. Ernst. *Reinforcement learning and dynamic programming*
 389 *using function approximators*. CRC press, 2017.
- 390 [10] R. Carpio and T. Kamihigashi. Fast value iteration: an application of Legendre–Fenchel duality to a class
 391 of deterministic dynamic programming problems in discrete time. *Journal of Difference Equations and*
 392 *Applications*, 26(2):209–222, 2020.
- 393 [11] L. Contento, A. Ern, and R. Vermiglio. A linear-time approximate convex envelope algorithm using the
 394 double Legendre–Fenchel transform with application to phase separation. *Computational Optimization*
 395 *and Applications*, 60(1):231–261, 2015.
- 396 [12] L. Corrias. Fast Legendre–Fenchel transform and applications to Hamilton–Jacobi equations and conserva-
 397 tion laws. *SIAM Journal on Numerical Analysis*, 33(4):1534–1558, 1996.
- 398 [13] G. Costeseque and J.-P. Lebacque. A variational formulation for higher order macroscopic traffic flow
 399 models: Numerical investigation. *Transportation Research Part B: Methodological*, 70:112 – 133, 2014.
- 400 [14] A. O. Esogbue and C. W. Ahn. Computational experiments with a class of dynamic programming
 401 algorithms of higher dimensions. *Computers & Mathematics with Applications*, 19(11):3 – 23, 1990.
- 402 [15] P. F. Felzenszwalb and D. P. Huttenlocher. Distance transforms of sampled functions. *Theory of computing*,
 403 8(1):415–428, 2012.
- 404 [16] M. Jacobs and F. Léger. A fast approach to optimal transport: the back-and-forth method. *arXiv preprint*
 405 *arXiv:1905.12154*, 2019.
- 406 [17] C. M. Klein and T. L. Morin. Conjugate duality and the curse of dimensionality. *European Journal of*
 407 *Operational Research*, 50(2):220 – 228, 1991.
- 408 [18] M. A. S. Kolarijani and P. Mohajerin Esfahani. Fast approximate dynamic programming for input-affine
 409 dynamics. *preprint arXiv:2008.10362*, 2020.
- 410 [19] Y. Lucet. Faster than the fast Legendre transform, the linear-time Legendre transform. *Numerical*
 411 *Algorithms*, 16(2):171–185, 1997.
- 412 [20] Y. Lucet. New sequential exact Euclidean distance transform algorithms based on convex analysis. *Image*
 413 *and Vision Computing*, 27(1):37 – 44, 2009.
- 414 [21] W. M. McEneaney. Max-plus eigenvector representations for solution of nonlinear H_∞ problems: basic
 415 concepts. *IEEE Transactions on Automatic Control*, 48(7):1150–1163, 2003.
- 416 [22] W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley &
 417 Sons, Hoboken, NJ, 2nd edition, 2011.
- 418 [23] A. Sidford, M. Wang, X. Wu, and Y. Ye. Variance reduced value iteration and faster algorithms for
 419 solving Markov decision processes. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium*
 420 *on Discrete Algorithms*, pages 770–787. SIAM, 2018.
- 421 [24] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.

422 **Checklist**

- 423 1. For all authors...
- 424 (a) Do the main claims made in the abstract and introduction accurately reflect the paper's
425 contributions and scope? [Yes] See Theorems 3.9, 3.10 and 3.11.
- 426 (b) Did you describe the limitations of your work? [Yes] See the discussions following
427 Theorems 3.10 and 3.11.
- 428 (c) Did you discuss any potential negative societal impacts of your work? [N/A]
- 429 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
430 them? [Yes]
- 431 2. If you are including theoretical results...
- 432 (a) Did you state the full set of assumptions of all theoretical results? [Yes] See, in
433 particular, Assumptions 3.6, 3.7, and 3.8.
- 434 (b) Did you include complete proofs of all theoretical results? [Yes] See Section 5 of the
435 extended manuscript in the supplementary material.
- 436 3. If you ran experiments...
- 437 (a) Did you include the code, data, and instructions needed to reproduce the main ex-
438 perimental results (either in the supplemental material or as a URL)? [Yes] See the
439 MATLAB codes provided in the supplementary material.
- 440 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
441 were chosen)? [N/A]
- 442 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
443 ments multiple times)? [N/A]
- 444 (d) Did you include the total amount of compute and the type of resources used (e.g., type
445 of GPUs, internal cluster, or cloud provider)? [Yes] See Section 4 of the extended
446 manuscript in the supplementary material.
- 447 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 448 (a) If your work uses existing assets, did you cite the creators? [Yes] We used the MATLAB
449 package for LLT provided in [19] as mentioned in Section 4.
- 450 (b) Did you mention the license of the assets? [N/A]
- 451 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
452 We are providing a MATLAB package for the implementation of the CVI algorithm
453 introduced in this study. See the MATLAB codes in the supplementary material.
- 454 (d) Did you discuss whether and how consent was obtained from people whose data you're
455 using/curating? [N/A]
- 456 (e) Did you discuss whether the data you are using/curating contains personally identifiable
457 information or offensive content? [N/A]
- 458 5. If you used crowdsourcing or conducted research with human subjects...
- 459 (a) Did you include the full text of instructions given to participants and screenshots, if
460 applicable? [N/A]
- 461 (b) Did you describe any potential participant risks, with links to Institutional Review
462 Board (IRB) approvals, if applicable? [N/A]
- 463 (c) Did you include the estimated hourly wage paid to participants and the total amount
464 spent on participant compensation? [N/A]