Efficient model-based reinforcement learning for approximate online optimal control

Rushikesh Kamalapurkar, Joel A. Rosenfeld, and Warren E. Dixon

Abstract—An the infinite horizon optimal regulation problem is solved online for a deterministic control-affine nonlinear dynamical system using a state following (StaF) kernel method to approximate the value function. Unlike traditional methods that aim to approximate a function over a large compact set, the StaF kernel method aims to approximate a function in a small neighborhood of a state that travels within a compact set. Simulation results demonstrate that stability and approximate optimality of the control system can be achieved with significantly fewer basis functions than may be required for global approximation methods.

I. INTRODUCTION

Reinforcement learning (RL) has become a popular tool for determining online solutions of optimal control problems for systems with finite state and action spaces [1]–[3]. Due to various technical challenges, implementation of RL in systems with continuous state and action spaces has remained an open problem. In recent years, adaptive dynamic programming (ADP) has been successfully used to implement RL in deterministic autonomous control-affine systems to solve optimal control problems via value function approximation [3]-[13]. ADP techniques employ parametric function approximation (typically by employing neural networks (NNs)) to approximate the value function. Implementation of function approximation in ADP is challenging because the controller is void of pre-designed stabilizing feedback and is completely defined by the estimated parameters. Hence, the error between the optimal and the estimated value function is required to decay to a sufficiently small bound sufficiently fast to establish closed-loop stability. The size of the error bound is determined by the selected basis functions, and the convergence rate is determined by richness of the data used for learning.

Sufficiently accurate approximation of the value function over a sufficiently large neighborhood often requires a large number of basis functions, and hence, introduces a large number of unknown parameters. One way to achieve accurate function approximation with fewer unknown parameters is to use some knowledge about the system to determine the basis functions. However, for general nonlinear systems, prior knowledge of the features of the optimal value function is generally not available; hence, a large number of generic basis functions is often the only feasible option.

Fast approximation of the value function over a large neighborhood requires sufficiently rich data to be available for learning. In traditional ADP methods such as [9], [11], [14], richness of data manifests itself as the amount of excitation in the system. In experience replay-based techniques such as [15]-[18], richness of data is quantified by eigenvalues of the recorded history stack. In model-based RL techniques such as [19]-[21], richness of data corresponds to the eigenvalues of a learning matrix. As the dimension of the system and the number of basis functions increases, the required richness of data increases. In traditional ADP methods, the demand for rich data is met by adding excitation signals to the controller, thereby causing undesirable oscillations. Hence, implementation of traditional ADP techniques such as [3]-[14] in high dimensional systems are seldom found in the literature. In experience replay-based ADP methods and in model-based RL, the demand for richer data causes the required amount of data stored in the history stack, and the number of points selected to construct the learning matrix, respectively, to grow exponentially with the dimension of the system. Hence, implementation of datadriven ADP techniques such as [18]-[23] are scarcely found in the literature.

In this paper, a novel model-based RL technique is developed to achieve sufficient excitation without causing undesirable oscillations and expenditure of control effort like traditional ADP techniques and at a lower computational cost than state-of-the-art data-driven ADP techniques. Motivated by the fact that the computational effort required to implement ADP and the data-richness required to achieve convergence both decrease with decreasing number of basis functions, this paper focuses on reduction of the number of basis functions used for value function approximation.

A key contribution of this paper and in our preliminary work in [24] is the observation that online implementation of an ADP-based approximate optimal controller does not require an estimate of the optimal value function over the entire domain of operation of the system. Instead, only an estimate of the slope of the value function evaluated at the current state is required for feedback. Hence, estimation of the value function over a small neighborhood of the current state should be sufficient to implement an ADPbased approximate optimal controller. Since it is reasonable to postulate that approximation of the value function over

Rushikesh Kamalapurkar, Joel A. Rosenfeld, and Warren E. Dixon are with the Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville, FL, USA. Email: {rkamalapurkar, joelar, wdixon}@ufl.edu.

This research is supported in part by NSF award numbers 1161260 and 1217908, ONR grant number N00014-13-1-0151, and a contract with the AFRL Mathematical Modeling and Optimization Institute. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsoring agency.

a local domain would require fewer basis functions than approximation over the entire domain of operation, reduction of the size of the approximation domain is motivated.

Reduction of the size of the approximation domain is achieved via selection of basis functions that travel with the system state (referred to as state-following (StaF) kernels) to achieve accurate approximation of the value function over a small neighborhood of the state. The use of StaF kernels introduces a technical challenge since the ideal values of the unknown parameters corresponding to the StaF kernels are functions of the system state. The Lyapunov-based stability analysis presented in Section IV explicitly incorporates this functional relationship using the result that the ideal weights are continuously differentiable functions of the system state.

Another key contribution of this paper is the observation that model-based RL techniques can be implemented without storing any data if the available model is used to simulate persistent excitation. In other words, an excitation signal added to the simulated system, instead of the actual physical system, can be used to learn the value function. Excitation via simulation is implemented using Bellman error (BE) extrapolation (cf. [19]-[21]); however, instead of a large number of autonomous extrapolation functions employed in results such as [19]-[21], a single time-varying extrapolation function is selected, where the time-variation of the extrapolation function simulates excitation. The use of a single extrapolation point introduces a technical challenge since the BE extrapolation matrix is rank deficient at each time instance. The aforementioned challenge is addressed in Section IV-C by modifying the stability analysis to utilize persistent excitation of the extrapolated regressor matrix. Simulation results including comparisons with state-of-theart model-based RL techniques are included to demonstrate the effectiveness of the developed technique.

II. STAF KERNEL FUNCTIONS

The objective in StaF-based function approximation is to approximate a target function in a region of interest in the neighborhood of a point of interest $x \in \mathbb{R}^n$. In state-of-theart online approximate control, the optimal value function is approximated using a linear-in-the-parameters approximation scheme, and the approximate control law drives the system along the steepest negative gradient of the approximated value function. To compute the controller at the current state, only the gradient of the value function evaluated at the current state is required. Hence, in this application, the target function is the optimal value function, and the point of interest is the system state.

Since the system state evolves through the state-space with time, the region of interest for function approximation also evolves through the state-space. The StaF technique aims to maintain a uniform approximation of the value function over a small region around the current system state so that the gradient of the value function at the current state, and hence the optimal controller at the current state, can be approximated.

To facilitate the theoretical development, this section summarizes key results from our preliminary work in [25], where the theory of reproducing kernel Hilbert spaces (RKHSs) is used to establish continuous differentiability of the ideal weights with respect to the system state, and the postulate that approximation of a function over a small neighborhood requires fewer basis functions is stated and proved.

Let H be a universal RKHS over a compact set $\chi \subset \mathbb{R}^n$ with a continuously differentiable positive definite kernel k: $\underline{\chi} \times \chi \to \mathbb{R}$. Let $\overline{V}^* : \chi \to \mathbb{R}$ be a function such that $\overline{V}^* \in H$. Let $C \triangleq [c_1, c_2, \cdots c_L]^T \in \chi^L$ be a set of distinct centers, and let $\sigma : \chi \times \chi^L \to \mathbb{R}^L$ be defined as $\sigma(x, C) = [k(x, c_1), \cdots, k(x, c_L)]^T$. Then, there exists a unique set of weights W_H such that

$$W_H(C) = \arg\min_{a \in \mathbb{R}^L} \left\| a^T \sigma(\cdot, C) - \overline{V}^* \right\|_H$$

where $\|\cdot\|_{H}$ denotes the Hilbert space norm.

Let $B_r(x) \subset \chi$ denote a closed ball of radius r centered at the current state x. Let $H_{x,r}$ denote the restriction of the Hilbert space H to $B_r(x)$. Then, $H_{x,r}$ is a Hilbert space with the restricted kernel $k_{x,r} : B_r(x) \times B_r(x) \to \mathbb{R}$ defined as $k_{x,r}(y,z) = k(y,z), \forall (y,z) \in B_r(x) \times B_r(x)$. The following result, first stated and proved in [25] is stated here to motivate the use of StaF kernels.

Theorem 1. [25] Let $\epsilon, r > 0$ and let p denote a polynomial that approximates \overline{V}^* within an error ϵ over $B_r(x)$. Let $N(r, x, \epsilon)$ denote the degree of p. Let $k(y, x) = e^{y^T x}$ be the exponential kernel function, which corresponds to a universal RKHS. Then, for each $x \in \chi$, there exists a finite number of centers, $c_1, c_2, ..., c_{M(r,x,\epsilon)} \in B_r(x)$ and weights $w_1, w_2, ..., w_{M(r,x,\epsilon)}$ such that

$$\left\|\overline{V}^*(y) - \sum_{i=1}^{M(r,x,\epsilon)} w_i e^{y^T c_i}\right\|_{B_r(x),\infty} < \epsilon,$$

where $M(r, x, \epsilon) < {\binom{n+N(r, x, \epsilon)+S(r, x, \epsilon)}{N(r, x, \epsilon)+S(r, x, \epsilon)}}$, asymptotically, for some $S(r, x, \epsilon) \in \mathbb{N}$. Moreover, $r, N(r, x, \epsilon)$ and $S(r, x, \epsilon)$ can be bounded uniformly over χ for any fixed ϵ .¹

The Weierstrass theorem indicates that as r decreases, the degree $N(r, x, \epsilon)$ of the polynomial needed to achieve the same error ϵ over $B_r(x)$ decreases [26]. Hence, by Theorem 1, approximation of a function over a smaller domain requires a smaller number of exponential kernels. Furthermore, provided the region of interest is small enough, the number of kernels required to approximate continuous functions with arbitrary accuracy can be reduced to $\binom{n+2}{2}$.

In the StaF approach, the centers are selected to follow the current state x, i.e., the locations of the centers are defined as a function of the system state. Since the system state evolves in time, the ideal weights are not constant. To approximate the ideal weights using gradient-based algorithms, it is essential that the weights change smoothly with respect to the system state. The following result, first stated and proved in [25] establishes continuity of the ideal weights as a function of the centers.

¹The notation $\binom{a}{b}$ denotes the combinatorial operation "a choose b".

Theorem 2. [25] Let the kernel function k be such that the functions $k(\cdot, x)$ are l-times continuously differentiable for all $x \in \chi$. Let $C \triangleq [c_1, c_2, ..., c_L]^T$ be a set of distinct centers such that $c_i \in B_r(x)$, $\forall i = 1, \dots, L$, be an ordered collection of L distinct centers with associated ideal weights

$$W_{H_{x,r}}(C) = \arg \min_{a \in R^M} \left\| \sum_{i=1}^M a_i k(\cdot, c_i) - V(\cdot) \right\|_{H_{x,r}}.$$
 (1)

Then, the function $W_{H_{x,r}}$ is l-times continuously differentiable with respect to each component of C.

Theorem 1 motivates the use of StaF kernels for modelbased RL, and Theorem 2 facilitates implementation of gradient-based update laws to learn the time-varying ideal weights in real-time. In the following, the StaF-based function approximation approach is used to approximately solve an optimal regulation problem online using exact model knowledge via value function approximation. Selection of an optimal regulation problem and the assumption that the system dynamics are known are motivated by ease of exposition. Using a concurrent learning-based adaptive system identifier and the state augmentation technique developed in [20], the technique developed in this paper can be extended to a class of trajectory tracking problems in the presence of uncertainties in the system drift dynamics. Simulation results in Section V-B demonstrate the performance of such an extension.

III. STAF KERNEL FUNCTIONS FOR ONLINE APPROXIMATE OPTIMAL CONTROL

A. Problem Formulation

Consider a control affine nonlinear dynamical system of the form

$$\dot{x}(t) = f(x(t)) + g(x(t))u(t),$$
 (2)

 $t \in \mathbb{R}_{\geq t_0}$,² where t_0 denotes the initial time, $x : \mathbb{R}_{\geq t_0} \to \mathbb{R}^n$ denotes the system state $f : \mathbb{R}^n \to \mathbb{R}^n$ and $g : \mathbb{R}^n \to \mathbb{R}^{n \times m}$ denote the drift dynamics and the control effectiveness, respectively, and $u : \mathbb{R}_{\geq t_0} \to \mathbb{R}^m$ denotes the control input. The functions f and g are assumed to be locally Lipschitz continuous. Furthermore, $f(\mathbf{0}_{n \times 1}) = \mathbf{0}_{n \times 1}$ and $\nabla f : \mathbb{R}^n \to \mathbb{R}^{n \times n}$ is continuous.³ In the following, the notation $\phi^u(t; t_0, x_0)$ denotes the trajectory of the system in (2) under the control signal u with the initial condition $x_0 \in \mathbb{R}^n$ and initial time $t_0 \in \mathbb{R}_{\geq 0}$.

The control objective is to solve the infinite-horizon optimal regulation problem online, i.e., to design a control signal u online to minimize the cost functional

$$J(x,u) \triangleq \int_{t_0}^{\infty} r(x(\tau), u(\tau)) d\tau, \qquad (3)$$

²The notation $\mathbb{R}_{\geq a}$ denotes the interval $[a, \infty)$, and the notation $\mathbb{R}_{>a}$ denotes the interval (a, ∞) .

³The notation $\mathbf{0}_{n \times m}$ denotes an $n \times m$ matrix of zeros.

under the dynamic constraint in (2) while regulating the system state to the origin. In (3), $r : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}_{\geq 0}$ denotes the instantaneous cost defined as

$$r\left(x^{o}, u^{o}\right) \triangleq Q\left(x^{o}\right) + u^{oT} R u^{o},\tag{4}$$

for all $x^o \in \mathbb{R}^n$ and $u^o \in \mathbb{R}^m$, where $Q : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is a positive definite function, and $R \in \mathbb{R}^{m \times m}$ is a constant positive definite symmetric matrix. In (4) and in the reminder of this paper, the notation $(\cdot)^o$ is used to denote an arbitrary variable.

B. Exact Solution

It is well known that since the functions f, g, and Q are stationary (time-invariant) and the time-horizon is infinite, the optimal control input is a stationary state-feedback policy $u(t) = \xi(x(t))$ for some function $\xi : \mathbb{R}^n \to \mathbb{R}^m$. Furthermore, the function that maps each state to the total accumulated cost starting from that state and following a stationary state-feedback policy, i.e., the value function, is also a stationary function. Hence, the optimal value function $V^* : \mathbb{R}^n \to \mathbb{R}_{>0}$ can be expressed as

$$V^{*}\left(x^{o}\right) \triangleq \inf_{u(\tau) \in U \mid \tau \in \mathbb{R}_{\geq t}} \int_{t}^{\infty} r\left(\phi^{u}\left(\tau; t, x^{o}\right), u\left(\tau\right)\right) d\tau, \quad (5)$$

where $U \subset \mathbb{R}^m$ is a compact set. Assuming an optimal controller exists, the optimal value function is characterized by the corresponding HJB equation [27]

$$0 = \min_{u^{o} \in U} \left(\nabla V \left(x^{o} \right) \left(f \left(x^{o} \right) + g \left(x^{o} \right) u^{o} \right) + r \left(x^{o}, u^{o} \right) \right),$$
(6)

with the boundary condition $V(\mathbf{0}_{n\times 1}) = 0$. Provided the HJB in (6) admits a continuously differentiable solution, it constitutes a necessary and sufficient condition for optimality, i.e., if the optimal value function in (5) is continuously differentiable, then it is the unique solution to the HJB in (6) [28]. In (6) and in the following development, the notation $\nabla f(x, y, \cdots)$ denotes the partial derivative of f with respect to the first argument. The optimal control policy $u^* : \mathbb{R}^n \to \mathbb{R}^m$ can be determined from (6) as [27]

$$u^{*}(x^{o}) \triangleq -\frac{1}{2}R^{-1}g^{T}(x^{o})\left(\nabla V^{*}(x^{o})\right)^{T}.$$
 (7)

The HJB in (6) can be expressed in an open-loop form as

$$\nabla V^* (x^o) \left(f(x^o) + g(x^o) u^*(x^o) \right) + r(x^o, u^*(x^o)) = 0.$$
(8)

Using (7), the HJB in (8) can be expressed in a closed-loop form as

$$-\frac{1}{4}\nabla V^{*}(x^{o}) g(x^{o}) R^{-1}g^{T}(x^{o}) (\nabla V^{*}(x^{o}))^{T} + \nabla V^{*}(x^{o}) f(x^{o}) + Q(x^{o}) = 0.$$
(9)

The optimal policy can now be obtained using (7) if the HJB in (9) can be solved for the optimal value function V^* .

C. Value Function Approximation

An analytical solution of the HJB equation is generally infeasible; hence, an approximate solution is sought. In an approximate actor-critic-based solution, the optimal value function V^* is replaced by a StaF kernel-based parametric estimate $\hat{V}: \mathbb{R}^n \times \mathbb{R}^L \to \mathbb{R}$, where ideal values of the unknown parameters, denoted by W, are functions of the state x. Replacing V^* by \hat{V} in (7), an approximation to the optimal policy u^* is obtained as $\hat{u}: \mathbb{R}^n \times \mathbb{R}^L \to \mathbb{R}^m$. The objective of the critic is to learn the parameters W, and the objective of the actor is to implement a stabilizing controller based on the parameters learned by the critic. Motivated by the stability analysis, the actor and the critic maintain separate estimates \hat{W}_a and \hat{W}_c , respectively, of the ideal parameters W. Substituting the estimates \hat{V} and \hat{u} for V^* and u^* in (8), respectively, a residual error $\delta : \mathbb{R}^n \times \mathbb{R}^L \times \mathbb{R}^L \to \mathbb{R}$, called the BE is computed as

$$\delta\left(x^{o}, \hat{W}_{c}, \hat{W}_{a}\right) \triangleq r\left(x^{o}, \hat{u}\left(x^{o}, \hat{W}_{a}\right)\right) + \nabla \hat{V}\left(x^{o}, \hat{W}_{c}\right)\left(f\left(x^{o}\right) + g\left(x^{o}\right)\hat{u}\left(x^{o}, \hat{W}_{a}\right)\right).$$

To solve the optimal control problem, the critic aims to find a set of parameters \hat{W}_c and the actor aims to find a set of parameters \hat{W}_a such that $\delta\left(x^o, \hat{W}_c, \hat{W}_a\right) = 0, \forall x^o \in \mathbb{R}^n$. Since an exact basis for value function approximation is generally not available, an approximate set of parameters that minimizes the BE is sought.

The expression for the optimal policy in (7) indicates that to compute the optimal action when the system is at any given state x^{o} , one only needs to evaluate the gradient ∇V^{*} at x^{o} . Hence, to compute the optimal policy at x^{o} , one only needs to approximate the value function over a small neighborhood around x^{o} .

As established in Theorem 1, the number of basis functions required to approximate the value function is smaller if the approximation space is smaller in the sense of set containment. Hence, in this result, instead of aiming to obtain a uniform approximation of the value function over the entire operating domain, which might require a computationally intractable number of basis functions, the aim is to obtain a uniform approximation of the value function over a small neighborhood around the current system state.

StaF kernels are employed to achieve the aforementioned objective. To facilitate the development, let $\chi \subset \mathbb{R}^n$ be compact and let x^o be in the interior of χ . Then, for all $\epsilon > 0$, there exists a function $\overline{V}^* \in H_{x^o,r}$ such that $\sup_{x^o \in B_r(x^o)} |V^*(x^o) - \overline{V}^*(x^o)| < \epsilon$, where $H_{x^o,r}$ is a restriction of a universal RKHS H, introduced in Section II. In the developed StaF-based method, a small compact set $B_r(x^o)$ around the current state x^o is selected for value function approximation by selecting the centers $C \in B_r(x^o)$ such that $C = c(x^o)$ for some continuously differentiable function \hat{V} and

the approximate policy \hat{u} can then be expressed as

$$\hat{V}\left(x^{o}, \hat{W}_{c}\right) \triangleq \hat{W}_{c}^{T} \sigma\left(x^{o}, c\left(x^{o}\right)\right),$$

$$\hat{u}\left(x^{o}, \hat{W}_{a}\right) \triangleq -\frac{1}{2}R^{-1}g^{T}\left(x^{o}\right) \nabla \sigma\left(x^{o}, c\left(x^{o}\right)\right)^{T} \hat{W}_{a}, \quad (10)$$

where σ denotes the vector of basis functions introduced in Section II.

It should be noted that since the centers of the kernel functions change as the system state changes, the ideal weights also change as the system state changes. The state-dependent nature of the ideal weights differentiates this approach from state-of-the-art ADP methods in the sense that the stability analysis needs to account for changing ideal weights. Based on Theorem 2, it can be established that the ideal weight function $W: \chi \to \mathbb{R}^L$ defined as

$$W\left(x^{o}\right) \triangleq W_{H_{x^{o},r}}\left(c\left(x^{o}\right)\right),$$

where $W_{H_{x^o,r}}$ was introduced in (1), is continuously differentiable provided the functions σ and c are continuously differentiable.

D. Online Learning Based on Simulation of Experience

To learn the ideal parameters online, the critic evaluates a form $\delta_t : \mathbb{R}_{>t_0} \to \mathbb{R}$ of the BE at each time instance t as

$$\delta_{t}\left(t\right) \triangleq \delta\left(x\left(t\right), \hat{W}_{c}\left(t\right), \hat{W}_{a}\left(t\right)\right), \qquad (11)$$

where $\hat{W}_a(t)$ and $\hat{W}_c(t)$ denote the estimates of the actor and the critic weights, respectively, at time t, and the notation x(t) is used to denote the state the system in (2) at time twhen starting from initial time t_0 , initial state x_0 , and under the feedback controller

$$u(t) = \hat{u}\left(x(t), \hat{W}_a(t)\right).$$
(12)

Since (8) constitutes a necessary and sufficient condition for optimality, the BE serves as an indirect measure of how close the critic parameter estimates \hat{W}_c are to their ideal values; hence, in RL literature, each evaluation of the BE is interpreted as gained experience. Since the BE in (11) is evaluated along the system trajectory, the experience gained is along the system trajectory.

Learning based on simulation of experience is achieved by extrapolating the BE to unexplored areas of the state space. The critic selects a set of functions $\{x_i : \mathbb{R}^n \times \mathbb{R}_{\geq t_0} \to \mathbb{R}^n\}_{i=1}^N$ such that each x_i maps the current state x(t) to a point $x_i(x(t), t) \in B_r(x(t))$.

The critic then evaluates a form $\delta_{ti}: \mathbb{R}_{\geq t_0} \to \mathbb{R}$ of the BE for each x_i as

$$\delta_{ti}(t) = \hat{W}_{c}^{T}(t)\,\omega_{i}(t) + r\left(x_{i}(x(t),t),\hat{u}_{i}(t)\right), \quad (13)$$

where

$$\hat{u}_{i}\left(t\right) \triangleq -\frac{1}{2}R^{-1}g^{T}\left(x_{i}\left(x\left(t\right),t\right)\right)$$
$$\cdot \nabla\sigma\left(x_{i}\left(x\left(t\right),t\right),c\left(x\left(t\right)\right)\right)^{T}\hat{W}_{a}\left(t\right),$$

and

$$\begin{split} \omega_{i}\left(t\right) &\triangleq \nabla \sigma\left(x_{i}\left(x\left(t\right),t\right),c\left(x\left(t\right)\right)\right)f\left(x_{i}\left(x\left(t\right),t\right)\right)\\ &-\frac{1}{2}\nabla \sigma\left(x_{i}\left(x\left(t\right),t\right),c\left(x\left(t\right)\right)\right)g\left(x_{i}\left(x\left(t\right),t\right)\right)R^{-1}\cdot\\ &g^{T}\left(x_{i}\left(x\left(t\right),t\right)\right)\nabla \sigma^{T}\left(x_{i}\left(x\left(t\right),t\right),c\left(x\left(t\right)\right)\right)\hat{W}_{a}\left(t\right). \end{split}$$

The critic then uses the BEs from (11) and (13) to improve the estimate $\hat{W}_{c}(t)$ using the recursive least-squares-based update law

$$\dot{\hat{W}}_{c} = -k_{c1}\Gamma\left(t\right)\frac{\omega\left(t\right)}{\rho\left(t\right)}\delta_{t}\left(t\right) - \frac{k_{c2}}{N}\Gamma\left(t\right)\sum_{i=1}^{N}\frac{\omega_{i}\left(t\right)}{\rho_{i}\left(t\right)}\delta_{ti}\left(t\right),$$
(14)

where

$$\begin{split} \omega\left(t\right) &\triangleq \nabla \sigma\left(x\left(t\right), c\left(x\left(t\right)\right)\right) f\left(x\left(t\right)\right) \\ &- \frac{1}{2} \nabla \sigma\left(x\left(t\right), c\left(x\left(t\right)\right)\right) g\left(x\left(t\right)\right) R^{-1} g^{T}\left(x\left(t\right)\right) \\ &\cdot \nabla \sigma^{T}\left(x\left(t\right), c\left(x\left(t\right)\right)\right) \hat{W}_{a}\left(t\right), \end{split}$$

 $\rho_i(t) \triangleq \sqrt{1 + \gamma_1 \omega_i^T(t) \omega_i(t)}, \ \rho(t) \triangleq \sqrt{1 + \gamma_1 \omega^T(t) \omega(t)}, \\
k_{c1}, k_{c2}, \gamma_1 \in \mathbb{R}_{>0}$ are constant learning gains, and $\Gamma(t)$ denotes the least-square learning gain matrix updated according to

$$\dot{\Gamma}(t) = \beta \Gamma(t) - k_{c1} \Gamma(t) \frac{\omega(t) \omega^{T}(t)}{\rho^{2}(t)} \Gamma(t) - \frac{k_{c2}}{N} \Gamma(t) \sum_{i=1}^{N} \frac{\omega_{i}(t) \omega_{i}^{T}(t)}{\rho_{i}^{2}(t)} \Gamma(t), \ \Gamma(t_{0}) = \Gamma_{0}.$$
(15)

In (15), $\beta \in \mathbb{R}_{>0}$ is a constant forgetting factor.

Motivated by a Lyapunov-based stability analysis, the actor improves the estimate $\hat{W}_{a}(t)$ using the update law

$$\hat{\hat{W}}_{a}(t) = -k_{a1} \left(\hat{W}_{a}(t) - \hat{W}_{c}(t) \right) - k_{a2} \hat{W}_{a}(t) + \frac{k_{c1} G_{\sigma}^{T}(t) \hat{W}_{a}(t) \omega(t)^{T}}{4\rho(t)} \hat{W}_{c}(t) + \sum_{i=1}^{N} \frac{k_{c2} G_{\sigma i}^{T}(t) \hat{W}_{a}(t) \omega_{i}^{T}(t)}{4N\rho_{i}(t)} \hat{W}_{c}(t), \quad (16)$$

where $k_{a1}, k_{a2} \in \mathbb{R}_{>0}$ are learning gains,

$$G_{\sigma}(t) \triangleq \nabla \sigma \left(x\left(t \right), c\left(x\left(t \right) \right) \right) g\left(x\left(t \right) \right) R^{-1} g^{T} \left(x\left(t \right) \right) \cdot \nabla \sigma^{T} \left(x\left(t \right), c\left(x\left(t \right) \right) \right),$$

and

$$G_{\sigma i}(t) \triangleq \nabla \sigma \left(x_i(x(t),t), c(x(t)) \right) g \left(x_i(x(t),t) \right) R^{-1} \\ \cdot g^T \left(x_i(x(t),t) \right) \nabla \sigma^T \left(x_i(x(t),t), c(x(t)) \right).$$

IV. ANALYSIS

A. Computational complexity

The computational cost associated with an implementation of the developed method can be computed to be $O\left(N\left(L^3 + mnL + Lm^2 + n^2 + m^2\right)\right)$. Since local approximation is targeted, the StaF kernels result in a reduction in the number of required basis functions (i.e., L). Since the computational cost has a cubic relationship in the number of basis functions, the StaF methodology results in significant computational efficiency. The computational cost does grow linearly with the number of extrapolation points (i.e., N); however, if the points are selected using grid-based methods employed in results such as [19], the number N increases geometrically with respect to the state dimension n. If the extrapolation points are selected to be time varying, then as few as a single point can be sufficient provided the time-trajectory of the point contains enough information to satisfy the subsequent Assumption 1.

In the following, Assumption 1 formalizes the conditions under which the trajectories of the closed-loop system can be shown to be ultimately bounded, and Lemma 1 facilitates the analysis of the closed-loop system when time-varying extrapolation trajectories are utilized.

B. Excitation conditions

For notational brevity, time-dependence of all the signals is suppressed hereafter. Let $B_{\zeta} \subset \mathbb{R}^{n+2L}$ denote a closed ball with radius ζ centered at the origin. Let $\chi \triangleq B_{\zeta} \cap \mathbb{R}^n$. Let the notation $\|(\cdot)\|$ be defined as $\|h\| \triangleq \sup_{\xi \in \chi} \|h(\xi)\|$, for some continuous function $h : \mathbb{R}^n \to \mathbb{R}^k$. To facilitate the subsequent stability analysis, the BEs in (11) and (13) are expressed in terms of the weight estimation errors $\tilde{W}_c \triangleq$ $W - \hat{W}_c$ and $\tilde{W}_a = W - \hat{W}_a$ as

$$\delta_t = -\omega^T \tilde{W}_c + \frac{1}{4} \tilde{W}_a G_\sigma \tilde{W}_a + \Delta(x) ,$$

$$\delta_{ti} = -\omega_i^T \tilde{W}_c + \frac{1}{4} \tilde{W}_a^T G_{\sigma i} \tilde{W}_a + \Delta_i(x) , \qquad (17)$$

where the functions $\Delta, \Delta_i : \mathbb{R}^n \to \mathbb{R}$ are uniformly bounded over χ such that the bounds $\overline{\|\Delta\|}$ and $\overline{\|\Delta_i\|}$ decrease with decreasing $\overline{\|\nabla\epsilon\|}$. Let a candidate Lyapunov function $V_L :$ $\mathbb{R}^{n+2L} \times \mathbb{R}_{>t_0} \to \mathbb{R}$ be defined as

$$V_L(Z,t) \triangleq V^*(x) + \frac{1}{2}\tilde{W}_c^T \Gamma^{-1}(t)\tilde{W}_c + \frac{1}{2}\tilde{W}_a^T \tilde{W}_a,$$

where V^* is the optimal value function, and

$$Z = \left[x^T, \; \tilde{W}_c^T, \; \tilde{W}_a^T \right]^T$$

To facilitate learning, the system states x or the selected functions x_i are assumed to satisfy the following.

Assumption 1. There exists a constant $T \in \mathbb{R}_{>0}$ and such that

$$\underline{c}_{1}I_{L} \leq \int_{t}^{t+T} \left(\frac{\omega\left(\tau\right)\omega^{T}\left(\tau\right)}{\rho^{2}\left(\tau\right)}\right) d\tau, \,\forall t \in \mathbb{R}_{\geq t_{0}},$$
$$\underline{c}_{2}I_{L} \leq \inf_{t \in \mathbb{R}_{\geq t_{0}}} \left(\frac{1}{N}\sum_{i=1}^{N}\frac{\omega_{i}\left(t\right)\omega_{i}^{T}\left(t\right)}{\rho_{i}^{2}\left(t\right)}\right),$$
$$\underline{c}_{3}I_{L} \leq \frac{1}{N}\int_{t}^{t+T} \left(\sum_{i=1}^{N}\frac{\omega_{i}\left(\tau\right)\omega_{i}^{T}\left(\tau\right)}{\rho_{i}^{2}\left(\tau\right)}\right) d\tau, \,\forall t \in \mathbb{R}_{\geq t_{0}},$$

where, at least one of the nonnegative constants $\underline{c}_1, \underline{c}_2$, and \underline{c}_3 is strictly positive.

Remark 1. Unlike typical ADP literature that assumes ω is PE, Assumption 1 only requires *either* the regressor ω or the regressor ω_i to be persistently exciting. The regressor ω is completely determined by the system state x, and the weights \hat{W}_a . Hence, excitation in ω vanishes as the system states and the weights converge. Hence, in general, it is unlikely that $\underline{c}_1 > 0$. However, the regressor ω_i depends on x_i , which can be designed independent of the system state x. Hence, \underline{c}_3 can be made strictly positive if the signal x_i contains enough frequencies, and \underline{c}_2 can be made strictly positive by selecting a sufficient number of extrapolation functions.

C. Boundedness of the least-squares gain under persistent excitation

Intuitively, selection of a single time-varying BE extrapolation function results in virtual excitation. That is, instead of using input-output data from a persistently excited system, the dynamic model is used to simulate persistent excitation to facilitate parameter convergence.

In previous model-based RL results such as [19], stability and convergence of the developed method relied on \underline{c}_2 being strictly positive. In this paper, the computational efficiency of model-based RL is improved by allowing time-varying extrapolation functions that ensure \underline{c}_3 is strictly positive, which can be achieved using a single extrapolation trajectory that contains enough frequencies. The performance of the developed extrapolation method is demonstrated using comparative simulations in Section V-C, where it is demonstrated that the developed method using a single time-varying extrapolation point results in improved computational efficiency when compared to a large number of fixed extrapolation functions.

The following lemma facilitates the stability analysis by establishing upper and lower bounds on the eigenvalues of the least-squares learning gain matrix Γ .

Lemma 1. Provided Assumption 1 holds and $\lambda_{\min} \{\Gamma_0^{-1}\} > 0$, the update law in (15) ensures that the least squares gain matrix satisfies

$$\underline{\Gamma}I_L \le \Gamma(t) \le \overline{\Gamma}I_L, \tag{18}$$

where
$$\overline{\Gamma} = \frac{1}{\min\{k_{c1}\underline{c}_1 + k_{c2}\max\{\underline{c}_2 T, \underline{c}_3\}, \lambda_{\min}\{\Gamma_0^{-1}\}\}e^{-\beta T}}$$
 and $\underline{\Gamma} = \frac{1}{\lambda_{\max}\{\Gamma_0^{-1}\} + \frac{(k_{c1}+k_{c2})}{\beta\gamma_1}}$. Furthermore, $\overline{\Gamma} > 0$.

Proof: The proof closely follows the proof of [29, Corollary 4.3.2]. The update law in (15) implies that $\frac{d}{dt}\Gamma^{-1}(t) = -\beta\Gamma^{-1}(t) + k_{c1}\frac{\omega(t)\omega^{T}(t)}{\rho^{2}(t)} + \frac{k_{c2}}{N}\sum_{i=1}^{N}\frac{\omega_{i}(t)\omega^{T}_{i}(t)}{\rho^{2}(t)}$. Hence,

$$\begin{split} \Gamma^{-1}\left(t\right) &= e^{-\beta t} \Gamma_{0}^{-1} + k_{c1} \int_{0}^{t} e^{-\beta(t-\tau)} \frac{\omega\left(\tau\right) \omega^{T}\left(\tau\right)}{\rho^{2}\left(\tau\right)} d\tau \\ &+ \frac{k_{c2}}{N} \int_{0}^{t} e^{-\beta(t-\tau)} \sum_{i=1}^{N} \frac{\omega_{i}\left(\tau\right) \omega_{i}^{T}\left(\tau\right)}{\rho_{i}^{2}\left(\tau\right)} d\tau. \end{split}$$

To facilitate the proof, let t < T. Then,

$$\Gamma^{-1}(t) \ge e^{-\beta t} \Gamma_0^{-1} \ge e^{-\beta T} \Gamma_0^{-1} \ge \lambda_{\min} \left\{ \Gamma_0^{-1} \right\} e^{-\beta T} I_L.$$

If $t \ge T$, then since the integrands are positive, Γ^{-1} can be bounded as

$$\Gamma^{-1}(t) \ge k_{c1} \int_{t-T}^{t} e^{-\beta(t-\tau)} \frac{\omega(\tau) \omega^{T}(\tau)}{\rho^{2}(\tau)} d\tau$$
$$+ \frac{k_{c2}}{N} \int_{t-T}^{t} e^{-\beta(t-\tau)} \sum_{i=1}^{N} \frac{\omega_{i}(\tau) \omega_{i}^{T}(\tau)}{\rho_{i}^{2}(\tau)} d\tau$$

Hence,

$$\Gamma^{-1}(t) \ge k_{c1}e^{-\beta T} \int_{t-T}^{t} \frac{\omega(\tau)\omega^{T}(\tau)}{\rho^{2}(\tau)} d\tau + \frac{k_{c2}}{N}e^{-\beta T} \int_{t-T}^{t} \sum_{i=1}^{N} \frac{\omega_{i}(\tau)\omega_{i}^{T}(\tau)}{\rho_{i}^{2}(\tau)} d\tau.$$

Using Assumption 1,

$$\frac{1}{N} \int_{t-T}^{t} \sum_{i=1}^{N} \frac{\omega_{i}\left(\tau\right)\omega_{i}^{T}\left(\tau\right)}{\rho_{i}^{2}\left(\tau\right)} d\tau \geq \max\left\{\underline{c}_{2}T, \underline{c}_{3}\right\} I_{L},$$
$$\int_{t-T}^{t} \frac{\omega\left(\tau\right)\omega^{T}\left(\tau\right)}{\rho^{2}\left(\tau\right)} d\tau \geq \underline{c}_{1}I_{L}.$$

Hence a lower bound for Γ^{-1} is obtained as,

$$\Gamma^{-1}(t) \ge \min \left\{ k_{c1}\underline{c}_1 + k_{c2} \max \left\{ \underline{c}_2 T, \underline{c}_3 \right\}, \\ \lambda_{\min}\left\{ \Gamma_0^{-1} \right\} \right\} e^{-\beta T} I_L.$$
(19)

Provided Assumption 1 holds, the lower bound in (19) is strictly positive. Furthermore, using the facts that $\frac{\omega(t)\omega^T(t)}{\rho^2(t)} \leq \frac{1}{\gamma_1}$ and $\frac{\omega_i(t)\omega_i^T(t)}{\rho_i^2(t)} \leq \frac{1}{\gamma_1}$ for all $t \in \mathbb{R}_{\geq t_0}$,

$$\Gamma^{-1}(t) \leq e^{-\beta t} \Gamma_0^{-1} + \int_0^t e^{-\beta(t-\tau)} \left(k_{c1} \frac{1}{\gamma_1} + \frac{k_{c2}}{N} \sum_{i=1}^N \frac{1}{\gamma_1} \right) I_L d\tau,$$

$$\leq \left(\lambda_{\max} \left\{ \Gamma_0^{-1} \right\} + \frac{(k_{c1} + k_{c2})}{\beta \gamma_1} \right) I_L.$$

Since the inverse of the lower and upper bounds on Γ^{-1} are the upper and lower bounds on Γ , respectively, the proof is complete.

D. Main result

Since the optimal value function is positive definite, (18) and [30, Lemma 4.3] can be used to show that the candidate Lyapunov function satisfies the following bounds

$$\underline{v_l}\left(\|Z^o\|\right) \le V_L\left(Z^o, t\right) \le \overline{v_l}\left(\|Z^o\|\right),\tag{20}$$

for all $t \in \mathbb{R}_{\geq t_0}$ and for all $Z^o \in \mathbb{R}^{2+2L}$. In (20), $\underline{v}_l, \overline{v}_l : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ are class \mathcal{K} functions. To facilitate the analysis, let $\underline{c} \in \mathbb{R}_{>0}$ be a constant defined as

$$\underline{c} \triangleq \frac{\beta}{2\overline{\Gamma}k_{c2}} + \frac{\underline{c}_2}{2},\tag{21}$$

and let $\iota \in \mathbb{R}_{>0}$ be a constant defined as

$$\iota \triangleq \frac{3\left(\frac{(k_{c1}+k_{c2})\|\overline{\Delta}\|}{\sqrt{v}} + \frac{\|\overline{\nabla}Wf\|}{\underline{\Gamma}} + \frac{\|\overline{\Gamma}^{-1}G_{W\sigma}W\|}{2}\right)^2}{4k_{c2}\underline{c}} \\ + \frac{1}{(k_{a1}+k_{a2})} \left(\frac{\|\overline{G}_{W\sigma}W\| + \|\overline{G}_{V\sigma}\|}{2} + k_{a2}\|\overline{W}\| + \frac{\|\overline{\nabla}Wf\|}{4\sqrt{v}} + \frac{(k_{c1}+k_{c2})\|\overline{G}_{\sigma}\|\|W\|^2}{4\sqrt{v}}\right)^2 \\ + \frac{1}{2}\|\overline{G}_{V\epsilon}\|.$$

Let $v_l : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ be a class \mathcal{K} function such that

$$v_l(||Z||) \le \frac{Q(x)}{2} + \frac{k_{c2}c}{6} \left\|\tilde{W}_c\right\|^2 + \frac{(k_{a1}+k_{a2})}{8} \left\|\tilde{W}_a\right\|^2.$$

The sufficient conditions for the subsequent Lyapunov-based stability analysis are given by

$$\frac{k_{c2\underline{c}}}{3} \geq \frac{\left(\frac{\|\overline{G}_{W\sigma}\|}{2\underline{\Gamma}} + \frac{(k_{c1}+k_{c2})\|\overline{W}^T G_{\sigma}\|}{4\sqrt{v}} + k_{a1}\right)^2}{(k_{a1}+k_{a2})},$$
$$\frac{(k_{a1}+k_{a2})}{4} \geq \left(\frac{\|\overline{G}_{W\sigma}\|}{2} + \frac{(k_{c1}+k_{c2})\|\overline{W}\|\|\overline{G}_{\sigma}\|}{4\sqrt{v}}\right),$$
$$v_l^{-1}(\iota) < \overline{v_l}^{-1}\left(\underline{v_l}\left(\zeta\right)\right).$$
(22)

Note that the sufficient conditions can be satisfied provided the points for BE extrapolation are selected such that the minimum eigenvalue \underline{c} , introduced in (21) is large enough and that the StaF kernels for value function approximation are selected such that $||\overline{\epsilon}||$ and $||\nabla \overline{\epsilon}||$ are small enough. To improve computational efficiency, the size of the domain around the current state where the StaF kernels provide good approximation of the value function is desired to be small. Smaller approximation domain results in almost identical extrapolated points, which in turn, results in smaller \underline{c} . Hence, the approximation domain cannot be selected to be arbitrarily small and needs to be large enough to meet the sufficient conditions in (22).

Theorem 3. Provided Assumption 1 holds and the sufficient gain conditions in (22) are satisfied, the controller in (12) and the update laws in (14) - (16) ensure that the state x and the weight estimation errors \tilde{W}_c and \tilde{W}_a are ultimately bounded.

Proof: The time-derivative of the Lyapunov function is given by

$$\dot{V}_L = \dot{V}^* + \tilde{W}_c^T \Gamma^{-1} \left(\dot{W} - \dot{\hat{W}}_c \right) + \frac{1}{2} \tilde{W}_c^T \dot{\Gamma}^{-1} \tilde{W}_c + \tilde{W}_a^T \left(\dot{W} - \dot{\hat{W}}_a \right).$$

Using Theorem 2, the time derivative of the ideal weights can be expressed as

$$\dot{W} = \nabla W(x) \left(f(x) + g(x) u \right). \tag{23}$$

Using (14) - (17) and (23), the time derivative of the Lyapunov function is expressed as

$$\begin{split} \dot{V}_{L} &= \nabla V^{*}\left(x\right)\left(f\left(x\right) + g\left(x\right)u\right) \\ &+ \tilde{W}_{c}^{T}\Gamma^{-1}\nabla W\left(x\right)\left(f\left(x\right) + g\left(x\right)u\right) \\ &- \tilde{W}_{c}^{T}\Gamma^{-1}\left(-k_{c1}\Gamma\frac{\omega}{\rho}\left(-\omega^{T}\tilde{W}_{c} + \frac{1}{4}\tilde{W}_{a}G_{\sigma}\tilde{W}_{a} + \Delta(x)\right)\right)\right) \\ &- \tilde{W}_{c}^{T}\Gamma^{-1}\left(-\frac{k_{c2}}{N}\Gamma\sum_{i=1}^{N}\frac{\omega_{i}}{\rho_{i}}\frac{1}{4}\tilde{W}_{a}^{T}G_{\sigma i}\tilde{W}_{a}\right) \\ &- \tilde{W}_{c}^{T}\Gamma^{-1}\left(-\frac{k_{c2}}{N}\Gamma\sum_{i=1}^{N}\frac{\omega_{i}}{\rho_{i}}\left(-\omega_{i}^{T}\tilde{W}_{c} + \Delta_{i}\left(x\right)\right)\right) \\ &- \frac{1}{2}\tilde{W}_{c}^{T}\Gamma^{-1}\left(\beta\Gamma - k_{c1}\Gamma\frac{\omega\omega^{T}}{\rho}\Gamma\right)\Gamma^{-1}\tilde{W}_{c} \\ &- \frac{1}{2}\tilde{W}_{c}^{T}\Gamma^{-1}\left(-\frac{k_{c2}}{N}\Gamma\sum_{i=1}^{N}\frac{\omega_{i}\omega_{i}^{T}}{\rho_{i}}\Gamma\right)\Gamma^{-1}\tilde{W}_{c} \\ &+ \tilde{W}_{a}^{T}\left(\nabla W\left(x\right)\left(f\left(x\right) + g\left(x\right)u\right) - \dot{\tilde{W}}_{a}\right). \end{split}$$

Provided the sufficient conditions in (22) hold, the time derivative of the candidate Lyapunov function can be bounded as

$$\dot{V}_L \le -v_l \left(\|Z\| \right), \quad \forall \zeta > \|Z\| > v_l^{-1} \left(\iota \right).$$
 (24)

Using (20), (22), and (24), [30, Theorem 4.18] can be invoked to conclude that Z is ultimately bounded, in the sense that $\limsup_{t\to\infty} \|Z(t)\| \leq \underline{v_l}^{-1}(\overline{v_l}(\iota))$.

Since $Z \in \mathcal{L}_{\infty}$, x, \tilde{W}_a , and $\tilde{W}_c \in \mathcal{L}_{\infty}$. Since $x \in \mathcal{L}_{\infty}$ and since W is a continuous function of x, $W \circ x \in \mathcal{L}_{\infty}$. Hence, \hat{W}_a and $\hat{W}_c \in \infty$, which implies $u \in \mathcal{L}_{\infty}$.

V. SIMULATION

A. Optimal regulation problem with exact model knowledge

1) Simulation parameters: To demonstrate the effectiveness of the StaF kernels, simulations are performed on a two-state nonlinear dynamical system. The system dynamics are given by (2), where $x^o = [x_1^o, x_2^o]^T$,

$$f(x^{o}) = \begin{bmatrix} -x_{1}^{o} + x_{2}^{o} \\ -\frac{1}{2}x_{1}^{o} - \frac{1}{2}x_{2}^{o}\left(\cos\left(2x_{1}^{o}\right) + 2\right)^{2} \end{bmatrix},$$
$$g(x^{o}) = \begin{bmatrix} 0 \\ \cos\left(2x_{1}^{o}\right) + 2 \end{bmatrix}.$$
(25)

The control objective is to minimize the cost

$$\int_{0}^{\infty} \left(x^{T}\left(\tau\right) x\left(\tau\right) + u^{2}\left(\tau\right) \right) d\tau.$$
(26)

The system in (25) and the cost in (26) are selected because the corresponding optimal control problem has a known analytical solution. The optimal value function is $V^*(x^o) = \frac{1}{2}x_1^{o2} + x_2^{o2}$, and the optimal control policy is $u^*(x^o) = -(\cos(2x_1^o) + 2)x_2^o$ (cf. [9]).

To apply the developed technique to this problem, the value function is approximated using three exponential StaF kernels, i.e., $\sigma(x^o, C) =$



Figure 1. State trajectories generated using StaF kernel-based ADP.

 $\begin{bmatrix} \sigma_1 \left(x^o, c_1 \right), & \sigma_2 \left(x^o, c_2 \right), & \sigma_3 \left(x^o, c_3 \right) \end{bmatrix}^T. \text{ The kernels are selected to be } \sigma_i \left(x^o, c_i \right) = e^{x^{o^T}c_i} - 1, & i = 1, \cdots, 3. \\ \text{The centers } c_i \text{ are selected to be on the vertices of a shrinking equilateral triangle around the current state, i.e., <math>c_i = x^o + d_i \left(x^o \right), & i = 1, \cdots, 3, \text{ where } d_1 \left(x^o \right) = 0.7\nu^o \left(x^o \right) \cdot [0, 1]^T, & d_2 \left(x^o \right) = 0.7\nu^o \left(x^o \right) \cdot [0.87, -0.5]^T, \\ \text{and } d_3 \left(x^o \right) = 0.7\nu^o \left(x^o \right) \cdot [-0.87, -0.5]^T, \text{ and } \\ \nu^o \left(x^o \right) \triangleq \left(\frac{x^{o^T}x^o + 0.01}{1 + \gamma_2 x^o Tx^o} \right) \text{ denotes the shrinking function, } \\ \text{where } \gamma_2 \in \mathbb{R}_{>0} \text{ is a constant normalization gain. The point for BE extrapolation is selected at random from a uniform distribution over a <math display="inline">2.1\nu^o \left(x \left(t \right) \right) \times 2.1\nu^o \left(x \left(t \right) \right) \text{ square centered at the current state } x \left(t \right) \text{ so that the function } x_i \text{ is of the form } x_i \left(x^o, t \right) = x^o + a_i \left(t \right) \text{ for some } a_i \left(t \right) \in \mathbb{R}^2. \\ \end{bmatrix}$

The system is initialized at $t_0 = 0$ and the initial conditions⁴

$$\begin{aligned} x\left(0\right) &= [-1, 1]^{T}, \ W_{c}\left(0\right) = 0.4 \times \mathbf{1}_{3 \times 1}, \\ \Gamma\left(0\right) &= 500I_{3}, \ \hat{W}_{a}\left(0\right) = 0.7\hat{W}_{c}\left(0\right), \end{aligned}$$

and the learning gains are selected as

$$k_{c1} = 0.001, \ k_{c2} = 0.25, \ k_{a1} = 1.2, \ k_{a2} = 0.01,$$

 $\beta = 0.003, \ \gamma_1 = 0.05, \ \gamma_2 = 1.$

2) *Results:* Figure 1 shows that the developed StaF-based controller drives the system states to the origin while maintaining system stability. Figure 2 shows the implemented control signal compared with the optimal control signal. It is clear that the implemented control converges to the optimal controller. Figure 3 shows that the weight estimates for the StaF-based value function and policy approximation remain bounded and converge as the state converges to the origin. Since the ideal values of the weights are unknown, the weights can not directly be compared with their ideal values. However, since the optimal solution is known, the value function estimate corresponding to the weights in Figure 3



Figure 2. Control trajectory generated using StaF kernel-based ADP compared with the optimal control trajectory.



Figure 3. Trajectories of the estimates of the unknown parameters in the value function generated using StaF kernel-based ADP. The ideal weights are unknown and time-varying; hence, the obtained weights can not be compared with their ideal weights.

can be compared to the optimal value function at each time t. Figure 5 shows that the error between the optimal and the estimated value functions rapidly decays to zero.

B. Optimal tracking problem with parametric uncertainties in the drift dynamics

1) Simulation parameters: Similar to [20], the developed StaF-based RL technique is extended to solve optimal tracking problems with parametric uncertainties in the drift dynamics. The drift dynamics in the two-state nonlinear dynamical system in (25) are assumed to be linearly param-

⁴The notation I_n denotes a $n \times n$ identity matrix and $\mathbf{1}_{n \times m}$ and $\mathbf{0}_{n \times m}$ denote $n \times m$ matrices of ones and zeros, respectively.



Figure 4. Trajectories of the estimates of the unknown parameters in the policy generated using StaF kernel-based ADP. The ideal weights are unknown and time-varying; hence, the obtained weights can not be compared with their ideal weights.



Figure 5. The error between the optimal and the estimated value function.

eterized as

$$f\left(x^{o}\right) = \underbrace{\left[\begin{array}{ccc} \theta_{1} & \theta_{2} & \theta_{3} \\ \theta_{4} & \theta_{5} & \theta_{6} \end{array}\right]}_{\theta^{T}} \underbrace{\left[\begin{array}{ccc} x_{1}^{o} \\ x_{2}^{o} \\ x_{2}^{o} \left(\cos\left(2x_{1}^{o}\right)+2\right) \end{array}\right]}_{\sigma_{\theta}\left(x^{o}\right)},$$

where $\theta \in \mathbb{R}^{3 \times 2}$ is the matrix of unknown parameters and σ_{θ} is the known vector of basis functions. The ideal values of the unknown parameters are $\theta_1 = -1$, $\theta_2 = 1$, $\theta_3 = 0$, $\theta_4 = -0.5$, $\theta_5 = 0$, and $\theta_6 = -0.5$. Let $\hat{\theta}$ denote an estimate of the unknown matrix θ . The control objective is to drive the estimate $\hat{\theta}$ to the ideal matrix θ , and to drive the state x to follow a desired trajectory x_d . The desired trajectory is

selected to be solution of the initial value problem

$$\dot{x}_d(t) = \begin{bmatrix} -1 & 1\\ -2 & 1 \end{bmatrix} x_d(t), \quad x_d(0) = \begin{bmatrix} 0\\ 1 \end{bmatrix}, \quad (27)$$

and the cost functional is selected to be $\begin{pmatrix} e^{T}(t) \operatorname{diag}(10, 10) e(t) + (\mu(t))^{2} \\ 0 &= x(t) - x_{d}(t), \quad \mu(t) = \\ x_{d}(t)) \left(\begin{bmatrix} -1 & 1 \\ -2 & 1 \end{bmatrix} x_{d}(t) - f(x_{d}(t)) \right), \quad \text{and}$ \int_0^∞ where e(t)u(t)_ $g^{+}(x^{o})$ $g^{+}\left(x_{d}\left(t\right)\right)$ denotes the pseudoinverse of $g(x^o)$.

The value function is a function of the concatenated state $\zeta \triangleq \begin{bmatrix} e^T & x_d^T \end{bmatrix}^T \in \mathbb{R}^4$. The value function is approximated using five exponential StaF kernels given by $\sigma_i(\zeta^o, C)$, where the five centers are selected according to $c_i = \zeta^o + d_i(\zeta^o)$ to form a regular five dimensional simplex around the current state with $\nu^o(\zeta^o) \equiv 1$. Learning gains for system identification and value function approximation are selected as

$$k_{c1} = 0.001, \ k_{c2} = 2, \ k_{a1} = 2, \ k_{a2} = 0.001,$$
$$\beta = 0.01, \ \gamma_1 = 0.1, \ \gamma_2 = 1, \ k = 500,$$
$$\Gamma_{\theta} = I_3, \ \Gamma(0) = 50I_5, \ k_{\theta} = 20,$$

To implement BE extrapolation, a single state trajectory ζ_i is selected as ζ_i (ζ^o, t) = $\zeta^o + a_i$ (t), where a_i (t) is sampled at each t from a uniform distribution over the a 2.1 × 2.1 × 2.1 × 2.1 × 2.1 hypercube centered at the origin. The history stack required for concurrent learning contains ten points, and is recorded online using a singular value maximizing algorithm (cf. [17]), and the required state derivatives are computed using a fifth order Savitzky-Golay smoothing filter (cf. [31]).

The initial values for the state and the state estimate are selected to be $x(0) = [0,0]^T$ and $\hat{x}(0) = [0,0]^T$, respectively. The initial values for the NN weights for the value function, the policy, and the drift dynamics are selected to be $0.025 \times \mathbf{1}_{5\times 1}$, $0.025 \times \mathbf{1}_{5\times 1}$, and $\mathbf{0}_{3\times 2}$, respectively. Since the system in (25) has no stable equilibria, the initial policy $\hat{\mu}(\zeta, \mathbf{0}_{3\times 2})$ is not stabilizing. The stabilization demonstrated in Figure 6 is achieved via fast simultaneous learning of the system dynamics and the value function.

2) Results: Figures 6 and 7 demonstrate that the controller remains bounded and the tracking error is regulated to the origin. The NN weights are functions of the system state ζ . Since ζ converges to a periodic orbit, the NN weights also converge to a periodic orbit (within the bounds of the excitation introduced by the BE extrapolation signal), as demonstrated in Figures 8 and 9. Figure 10 demonstrates that the unknown parameters in the drift dynamics, represented by solid lines, converge to their ideal values, represented by dashed lines.

C. Comparison

The developed technique is compared with the modelbased RL method developed in [19] for regulation and [20] for tracking, respectively. The simulations are performed in MATLAB[®] Simulink[®] at 1000 Hz on the same machine. The simulations run for 100 seconds of simulated time. Since the objective is to compare computational efficiency of the

Problem description	Regulation	(2-state system)	Regulation (3-state system)		Tracking (4-state system)	
Controller	StaF	Controller in [19]	StaF	Controller in [19]	StaF	Controller in [20]
Running time (seconds)	6.5	17	9.5	62	12	260
Total cost	2.8	1.8	9.3	12.3	3.9	3.4
RMS steady-state error	2.5×10^{-6}	0	4.3×10^{-6}	4.5×10^{-6}	3.5×10^{-4}	2.5×10^{-4}

Table I	
a i	

SIMULATION RESULTS FOR 2, 3 AND 4 DIMENSIONAL NONLINEAR SYSTEMS.



Figure 6. Tracking error trajectories generated using the proposed method for the trajectory tracking problem.



Figure 7. Control signal generated using the proposed method for the trajectory tracking problem.

model-based RL method, exact knowledge of the system model is used. Table I shows that the developed controller requires significantly fewer computational resources than the controllers from [19] and [20]. Furthermore, as the system dimension increases, the developed controller significantly outperforms the controllers from [19] and [20] in terms of computational efficiency.



Figure 8. Policy weight trajectories generated using the proposed method for the trajectory tracking problem. The weights do not converge to a steadystate value because the ideal weights are functions of the time-varying system state. Since an analytical solution of the optimal tracking problem is not available, weights cannot be compared against their ideal values

Since the optimal solution for the regulation problem is known to be quadratic, the model-based RL method from [19] is implemented using three quadratic basis functions. Since the basis used is exact, the method from [19] yields a smaller steady-state error than the developed method, which uses three inexact, but generic StaF kernels. For the 3-state regulation problem and the tracking problem, the methods from [19] and [20] are implemented using polynomial basis functions selected based on a trial-and-error approach. The developed technique is implemented using generic StaF kernels. In this case, since the optimal solution is unknown, both the methods use inexact basis functions, resulting in similar steady-state errors.

The two main advantages of StaF kernels are that they are universal, in the sense that they can be used to approximate a large class of value functions, and that they target local approximation, resulting in a smaller number of required basis functions. However, the StaF kernels trade optimality for universality and computational efficiency. The kernels are inexact, and the weight estimates need to be continually adjusted based on the system trajectory. Hence, as shown in Table I, the developed technique results in a higher total cost than state-of-the-art model-based RL techniques.



Figure 9. Value function weight trajectories generated using the proposed method for the trajectory tracking problem. The weights do not converge to a steady-state value because the ideal weights are functions of the time-varying system state. Since an analytical solution of the optimal tracking problem is not available, weights cannot be compared against their ideal values



Figure 10. Trajectories of the unknown parameters in the system drift dynamics for the trajectory tracking problem. The dotted lines represent the true values of the parameters.

VI. CONCLUSION

In this paper an infinite horizon optimal control problem is solved using a new approximation methodology called the StaF kernel method. Motivated by the fact that a smaller number of basis functions is required to approximate functions on smaller domains, the StaF kernel method aims to maintain good approximation of the value function over a small neighborhood of the current state. Computational efficiency of model-based RL is improved by allowing selection of fewer time-varying extrapolation trajectories instead of a large number of autonomous extrapolation functions. Simulation results are presented that solve the infinite horizon optimal regulation and tracking problems online for a two state system using only three and five basis functions, respectively, via the StaF kernel method.

State-of-the-art solutions to solve infinite horizon optimal control problems online aim to approximate the value function over the entire operating domain. Since the approximate optimal policy is completely determined by the value function estimate, state-of-the-art solutions generate, often at an intractable computational cost, policies that are valid over the entire state space. Since the StaF kernel method aims at maintaining local approximation of the value function around the current system state, the StaF kernel method lacks memory, in the sense that the information about the ideal weights over a region of interest is lost when the state leaves the region of interest. Thus, unlike existing techniques, the StaF method trades global optimality for computational efficiency to generate a policy that is near-optimal only over a small neighborhood of the origin. A memory-based modification to the StaF technique that retains and reuses past information is a subject for future research.

REFERENCES

- R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduc*tion. Cambridge, MA, USA: MIT Press, 1998.
- [2] D. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2007.
- [3] P. Mehta and S. Meyn, "Q-learning and pontryagin's minimum principle," in Proc. IEEE Conf. Decis. Control, Dec. 2009, pp. 3598–3605.
- [4] K. Doya, "Reinforcement learning in continuous time and space," *Neural Comput.*, vol. 12, no. 1, pp. 219–245, 2000.
- [5] R. Padhi, N. Unnikrishnan, X. Wang, and S. Balakrishnan, "A single network adaptive critic (SNAC) architecture for optimal control synthesis for a class of nonlinear systems," *Neural Netw.*, vol. 19, no. 10, pp. 1648–1660, 2006.
- [6] A. Al-Tamimi, F. L. Lewis, and M. Abu-Khalaf, "Discrete-time nonlinear HJB solution using approximate dynamic programming: Convergence proof," *IEEE Trans. Syst. Man Cybern. Part B Cybern.*, vol. 38, pp. 943–949, 2008.
- [7] F. L. Lewis and D. Vrabie, "Reinforcement learning and adaptive dynamic programming for feedback control," *IEEE Circuits Syst. Mag.*, vol. 9, no. 3, pp. 32–50, 2009.
- [8] T. Dierks, B. Thumati, and S. Jagannathan, "Optimal control of unknown affine nonlinear discrete-time systems using offline-trained neural networks with proof of convergence," *Neural Netw.*, vol. 22, no. 5-6, pp. 851–860, 2009.
- [9] K. Vamvoudakis and F. Lewis, "Online actor-critic algorithm to solve the continuous-time infinite horizon optimal control problem," *Automatica*, vol. 46, no. 5, pp. 878–888, 2010.
- [10] H. Zhang, L. Cui, X. Zhang, and Y. Luo, "Data-driven robust approximate optimal tracking control for unknown general nonlinear systems using adaptive dynamic programming method," *IEEE Trans. Neural Netw.*, vol. 22, no. 12, pp. 2226–2236, 2011.
- [11] S. Bhasin, R. Kamalapurkar, M. Johnson, K. Vamvoudakis, F. L. Lewis, and W. Dixon, "A novel actor-critic-identifier architecture for approximate optimal control of uncertain nonlinear systems," *Automatica*, vol. 49, no. 1, pp. 89–92, 2013.
- [12] H. Zhang, L. Cui, and Y. Luo, "Near-optimal control for nonzero-sum differential games of continuous-time nonlinear systems using singlenetwork adp," *IEEE Trans. Cybern.*, vol. 43, no. 1, pp. 206–216, 2013.
- [13] H. Zhang, D. Liu, Y. Luo, and D. Wang, Adaptive Dynamic Programming for Control Algorithms and Stability, ser. Communications and Control Engineering. London: Springer-Verlag, 2013.
- [14] K. Vamvoudakis and F. Lewis, "Online synchronous policy iteration method for optimal control," in *Recent Advances in Intelligent Control Systems*, W. Yu, Ed. Springer, 2009, pp. 357–374.
- [15] G. Chowdhary, "Concurrent learning adaptive control for convergence without persistencey of excitation," Ph.D. dissertation, Georgia Institute of Technology, December 2010.

- [16] G. Chowdhary and E. Johnson, "A singular value maximizing data recording algorithm for concurrent learning," in *Proc. American Con*trol Conf., 2011, pp. 3547–3552.
- [17] G. Chowdhary, T. Yucelen, M. Mühlegg, and E. N. Johnson, "Concurrent learning adaptive control of linear systems with exponentially convergent bounds," *Int. J. Adapt. Control Signal Process.*, vol. 27, no. 4, pp. 280–301, 2013.
- [18] H. Modares, F. L. Lewis, and M.-B. Naghibi-Sistani, "Integral reinforcement learning and experience replay for adaptive optimal control of partially-unknown constrained-input continuous-time systems," *Automatica*, vol. 50, no. 1, pp. 193–202, 2014.
- [19] R. Kamalapurkar, P. Walters, and W. E. Dixon, "Concurrent learningbased approximate optimal regulation," in *Proc. IEEE Conf. Decis. Control*, Florence, IT, Dec. 2013, pp. 6256–6261.
- [20] R. Kamalapurkar, L. Andrews, P. Walters, and W. E. Dixon, "Modelbased reinforcement learning for infinite-horizon approximate optimal tracking," in *Proc. IEEE Conf. Decis. Control*, 2014, pp. 5083–5088.
- [21] R. Kamalapurkar, J. Klotz, and W. Dixon, "Concurrent learning-based online approximate feedback Nash equilibrium solution of N -player nonzero-sum differential games," *IEEE/CAA J. Autom. Sin.*, vol. 1, no. 3, pp. 239–247, July 2014.
- [22] B. Luo, H.-N. Wu, T. Huang, and D. Liu, "Data-based approximate policy iteration for affine nonlinear continuous-time optimal control design," *Automatica*, 2014.
- [23] X. Yang, D. Liu, and Q. Wei, "Online approximate optimal control for affine non-linear systems with unknown internal dynamics using adaptive dynamic programming," *IET Control Theory Appl.*, vol. 8, no. 16, pp. 1676–1688, 2014.
- [24] R. Kamalapurkar, J. A. Rosenfeld, and W. E. Dixon, "State following (StaF) kernel functions for function approximation part II: Adaptive dynamic programming," in *Proc. Am. Control Conf.*, 2015, to appear (see also arXiv:1502.02609).
- [25] J. A. Rosenfeld, R. Kamalapurkar, and W. E. Dixon, "State following (StaF) kernel functions for function approximation part I: Theory and motivation," in *Proc. Am. Control Conf.*, 2015, to appear (see also arXiv:1503.04854).
- [26] G. G. Lorentz, *Bernstein polynomials*, 2nd ed. Chelsea Publishing Co., New York, 1986.
- [27] D. Kirk, Optimal Control Theory: An Introduction. Dover, 2004.
- [28] D. Liberzon, Calculus of variations and optimal control theory: a concise introduction. Princeton University Press, 2012.
- [29] P. Ioannou and J. Sun, Robust Adaptive Control. Prentice Hall, 1996.
- [30] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2002.
- [31] A. Savitzky and M. J. E. Golay, "Smoothing and differentiation of data by simplified least squares procedures." *Anal. Chem.*, vol. 36, no. 8, pp. 1627–1639, 1964.