European Journal of Operational Research 289 (2021) 683–695

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Addressing state space multicollinearity in solving an ozone pollution dynamic control problem

Bancha Ariyajunya[a], Ying Chen[b,*], Victoria C. P. Chen[c], Seoung Bum Kim[d], Jay Rosenberger[e]

[a] Faculty of Engineering, Burapha University, Chonburi, Thailand
[b] Department of Management Science and Engineering, Harbin Institute of Technology, China
[c] Department of Industrial, Manufacturing, and Systems Engineering, The University of Texas at Arlington, U.S.A.
[d] School of Industrial Management Engineering, Korea University, Seoul, Korea
[e] Department of Industrial, Manufacturing, and Systems Engineering, The University of Texas at Arlington, U.S.A.

ARTICLE INFO

Article history:
Received 23 September 2019
Accepted 6 July 2020
Available online 12 July 2020

Keywords:
Ozone pollution
Computer experiments
Multicollinearity
Statistical modeling
Approximate dynamic programming

ABSTRACT

High ground-level ozone concentrations constitute a serious air quality problem in many metropolitan regions. In this paper, we study a stochastic dynamic programming (SDP) formulation of the Atlanta metropolitan ozone pollution problem that seeks to reduce ozone via reductions of nitrogen oxides. The initial SDP formulation involves a 524-dimensional continuous state space, including ozone concentrations that are highly correlated. In prior work, a design and analysis of computer experiments (DACE) based approximate dynamic programming (ADP) solution method was able to conduct dimensionality reduction and value function approximation to enable a computationally-tractable numerical solution. However, this prior work did not address state space multicollinearity. In statistical modeling, high multicollinearity is well-known to adversely affect the generalizability of the constructed model. This issue is relevant whenever an empirical model is trained on data, but is largely ignored in the ADP literature. We propose approaches for addressing the multicollinearity in the Atlanta case study and demonstrate that if high multicollinearity is ignored, the resulting empirical models provide misleading information within the ADP algorithm. Because many SDP applications involve multicollinear continuous state spaces, the lessons learned in our research can guide the development of ADP approaches for a wide variety of SDP problems.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Ozone exists naturally in the atmosphere. In the stratosphere, high concentrations of ozone protect the earth's surface from harmful ultraviolet radiation emitted by the sun. Much lower ozone concentrations are found in the troposphere and are mainly due to the occurrence of photochemical reactions and stratospheric intrusions. However, in the troposphere, especially in metropolitan areas characterized by substantial anthropogenic emissions, photochemical reactions involving anthropogenic nitrogen oxides (NO_x) and volatile organic compounds (VOCs) significantly increase ground-level ozone concentrations. High ozone levels and other air pollutants are harmful to both the natural ecosystem and humans (U.S. Environmental Protection Agency, EPA, 2018). Thus, there is the need for regulations to control the emission of NO_x and/or VOCs to reduce ground-level ozone concentrations.

In this study, we consider the issue of a state space multicollinearity for an ozone pollution problem from Atlanta, Georgia. The stochastic dynamic programming (SDP) formulation of this ozone pollution problem involves a 524-dimensional continuous state space (Yang, Chen, Chang, Murphy & Tsai, 2007), containing ozone concentrations and precursor NO_x emissions occurring from 4:00 AM through 7:00 PM across the Atlanta metropolitan region. The cycle of ozone formation is daily, rising during the day and falling overnight. As commented in Ariyajunya, Chen, Chen and Kim (2017), the ozone concentrations are correlated spatially and temporally. Ozone formation is limited by the presence of precursor NO_x and/or VOC emissions. Atlanta, is "NO_x limited," which means that lowering ozone concentrations can be best achieved by reducing NO_x emissions (Chameides, Lindsay, Richardson & Kiang, 1988). Hence, the decision variables are reductions in NO_x emissions in different locations and time periods across the Atlanta metropolitan region. The goal is to minimize the expected cost

* Corresponding author.
E-mail addresses: bancha@buu.ac.th (B. Ariyajunya), yingchen@hit.edu.cn (Y. Chen), vchen@uta.edu (V.C.P. Chen), sbkim1@korea.ac.kr (S.B. Kim), jrosenbe@uta.edu (J. Rosenberger).

https://doi.org/10.1016/j.ejor.2020.07.014
0377-2217/© 2020 Elsevier B.V. All rights reserved.

of reducing NO_x emissions at each stage, so as to maintain the hourly-averaged ozone level at or below the EPA limit.

A design and analysis of computer experiments (DACE) based approximate dynamic programming (ADP) method has been employed to achieve a computationally-tractable numerical solution to this high-dimensional SDP problem (Yang et al., 2007; Yang, Chen, Chang, Sattler & Wen, 2009). In particular, Yang et al. (2009) focused on using ADP to solve the ozone pollution control problem with the state transition metamodel developed using stepwise regression in Yang et al. (2007). However, neither addressed the issue of state space multicollinearity. In fact, Ariyajunya et al. (2017) revealed that 3 of the 16 regression models in Yang et al. (2007) had very high variance inflation factors (VIFs), where high VIFs occur when there is high multicollinearity. Ariyajunya et al. (2017) demonstrated the presence of high state space multicollinearity in the ozone pollution SDP problem and proposed methods to orthogonalize the multicollinear state space, but did not demonstrate the impact of multicollinearity on the ADP solution.

In ADP value function approximation, the state space is analogous to the predictor or input variable space in statistics, and the value function is the unknown response surface. High multicollinearity among predictor variables leads to the situation in which there are infinitely many model forms with seemingly similar prediction accuracy, and there is no easy algorithm to choose among them (Kutner, Nachtsheim, Neter & Li, 2005: Ch 7). Consequently, an empirical model fit over a multicollinear space is subject to poor generalizability, meaning that the model form is sensitive to small changes in the data. While prediction accuracy might appear good within the multicollinear space, the wide variation in model forms can be extreme, with completely different sets of significant predictor variables, and estimated model coefficients that vary by orders of magnitude or switch signs, consequently affecting generalizability.

If the state variables are not multicollinear, then most any flexible statistical modeling method (or machine learning algorithm) can be used to generate a reasonable value function approximation. If multicollinear state space data are used to fit a value function approximation, then one is unlikely, given the infinite number of similarly accurate models, to uncover the appropriate relationships between the state variables and the value function. This, in turn, will provide misleading information to the subsequent stagewise optimizations that depend on the approximate value function. DACE based ADP is unique in ADP because it ensures an uncorrelated (or orthogonal) representation of the state space by using design of experiments, so as to provide better data for the value function approximation. Design of experiments is the gold standard for studying causal relationships (Box, Hunter & Hunter, 1978), with the most well-known being the clinical trial. However, state space multicollinearity must be addressed in order to take proper advantage of a DACE approach. The strength of our ADP research lies it in its foundation in well-established statistical approaches. In the next section, we provide background on the DACE based ADP approach. In Section 3, background is given on the ozone pollutions SDP problem. Section 4 presents development of state transition metamodels to study state space multicollinearity, and SDP results using these state transition metamodels are compared in Section 5. Finally, Section 6 presents the concluding recommendations.

2. Background and contribution

Eq. (1) is a recursive formulation of the finite-horizon continuous-state SDP problem (Bellman, 1957; Bertsekas, 2005):

$$
\begin{aligned}
&V_t(\boldsymbol{x}_t) = \min_{\boldsymbol{u}_t} E\{c_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t) + V_{t+1}(\boldsymbol{x}_{t+1})\} \\
&s.t. \quad \boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t), \text{for } t = 1, \ldots, T-1 \\
&\qquad \boldsymbol{u}_t \in \Gamma_t, \text{for } t = 1, \ldots, T \\
&\text{where } V_T(\boldsymbol{x}_T) = c_T(\boldsymbol{x}_T). \qquad (1)
\end{aligned}
$$

At the beginning of stage t, the state vector is $\boldsymbol{x}_t \in R^n$, the decision vector is $\boldsymbol{u}_t \in R^m$, and the vector of random variables is $\boldsymbol{\varepsilon}_t \in R^l$. We denote the future value function (FVF) by $V_t(\cdot)$, where the stagewise optimization minimizes the expectation of the stagewise cost function $c_t(\cdot)$ with the FVF for the future stage $t+1$, and the minimization is taken over $\boldsymbol{u}_t$, subject to the constraints Γ_t and the state transition function $f_t(\cdot)$. Exact solutions are only possible for small-scale problems under restrictive assumptions, hence, ADP and reinforcement learning (RL) algorithms (e.g. Powell 2011, Bertsekas, 2005) arose to enable numerical solutions, in particular, to address the curse of dimensionality. However, as commented in Powell (2011), for high-dimensional problems with continuous state and decision spaces, most existing ADP and RL algorithms still cannot perform well. To specifically address the challenges of continuous spaces, Chen, Ruppert and Shoemaker (1999), introduced a statistical perspective that we now refer to as a DACE based ADP approach. This approach has enabled numerical solutions to several high-dimensional continuous-state SDP problems (see Cervellera, Chen & Wen, 2006; Chen, 1999; Tsai, Chen, Beck & Chen, 2004; Yang et al., 2009). A brief introduction to DACE is provided in Section 2.1, followed by the DACE based ADP algorithm in Section 2.2, and our contribution in Section 2.3.

2.1. Design and analysis of computer experiments (DACE)

The concept of DACE was introduced by Sacks, Welch, Mitchell and Wynn (1989) as an approach to collect and analyze data from a deterministic simulation model. The approach is similar to response surface methodology (Box et al., 1978). Computer experiments constitute the use of design of experiments to control input settings for computer model runs. In most DACE applications, the computer model is a simulation model, and in Sacks et al. (1989), the simulation was assumed to be deterministic. Sacks et al. (1989) modeled the deterministic response output as the realization of a stochastic process, thereby providing a statistical basis for choosing the inputs for efficient prediction.

Chen, Tsui, Barton and Meckesheimer (2006) provided a review of DACE literature, extending the computer model to essentially any computer algorithm that defines an input and output relationship. Specifically, Chen et al. (2006) defined two basic tasks in DACE that must be conducted: (i) data collection; and (ii) statistical modeling. The data collection task is controlled by design of experiments, commonly space-filling designs, such as Latin hypercubes, orthogonal arrays and number-theoretic methods. The intention of the data collection task is not to merely sample data from a complex computer model. Rather, the intention is to control data collection, so as to improve the ability of the statistical modeling task to uncover a high quality approximation with minimal data. As for statistical modeling, this task is not merely the fitting of a pre-specified model to data. Rather, this task utilizes statistical modeling methods that are adaptively constructed to fit the data, such as regression trees and multivariate adaptive regression splines (MARS). Chen et al. (1999) utilized orthogonal array experimental designs and MARS, and Yang et al. (2009) employed a number-theoretic method, specifically called a low-discrepancy sequence, and MARS for the Atlanta ozone pollution case study.

2.2. DACE based ADP approach

In ADP, the stagewise optimization is the computer model, state space sampling is analogous to data collection, and FVF

1. For each stage t: use DoE to sample N points from the state space $\{x_{jt}\}_{t=1}^{N}$.
2. In last stage T:
   (a) For each discretized point x_{jT}, $j = 1, \ldots, N$, solve
   $$V_T(x_{jT}) = \min_{u_{jT}} E\{c_T(x_{jT}, u_{jT}, \varepsilon_{jT})\}$$
   (b) Approximate $V_T(x_T)$ with $\hat{V}_T(x_T)$ for all $x_T \in R^n$, by applying a statistical modeling approach to the data for V_T from step 2(a).
3. In each stage $t = T-1, \ldots, 1$:
   (a) For each sampled state point x_{jt}, $j = 1, \ldots, N$, solve
   $$\tilde{V}_t(x_{jt}) = \min_{u_{jt}} E\{c_t(x_{jt}, u_{jt}, \varepsilon_{jt}) + \hat{V}_{t+1}(f(x_{jt}, u_{jt}, \varepsilon_{jt})\}$$
   (b) Approximate $\tilde{V}_t(x_t)$ with $\hat{V}_t(x_t)$ for all $x_T \in R^n$, as in step 2(b).

Fig. 1. General DACE based ADP algorithm for continuous-state SDP problems (Chen et al., 1999).

approximation is analogous to statistical modeling. How one controls and represents the input data is actually more important than how the output data are modeled. In observational data studies, there is no control over data collection, but in computer experiments, there is absolute control over data collection. DACE based ADP, presented in Fig. 1, fully understands this concept, while the majority of ADP literature have yet to recognize this and continue to rely on Monte-Carlo or grid-based techniques for state space sampling. By fully controlling the sampling of the continuous state space via design of experiments (DoE in Fig. 1), DACE based ADP is able to overcome the curse of dimensionality and tackle very high-dimensional SDP problems. The statistical modeling task that approximates the FVF in DACE based ADP can be addressed with most any flexible statistical modeling method, although some methods may perform better with specific experimental designs (e.g., Cervellera, Wen & Chen, 2007). In general, it is assumed that prior knowledge of any special structure for the FVF is not available.

In Fig. 1, the output of each stagewise optimization, assuming a global optimum, is the point on the FVF corresponding to each DoE state point. Upon completion of the entire set of computer model runs, statistical modeling is implemented to approximate the FVF. The resulting solution to the algorithm in Fig. 1 is the SDP control policy, represented by the approximated FVFs, $\hat{V}_t(\cdot)$. The SDP control policy can then be implemented forward in real time by re-optimizing the approximated FVFs (Chen et al., 1999; Tejada-Guibert, Johnson & Stedinger, 1993). While the ADP algorithm is typically computationally-intensive, the re-optimization is very quick, enabling implementation of the SDP control policy in real time.

2.3. Contribution

While DACE based ADP effectively avoids multicollinearity, a key aspect is the appropriate representation of the state space prior to implementing a DACE approach. In this study, to demonstrate the effect of multicollinearity, we chose the high-dimensional Atlanta ozone pollution SDP problem, for which only the DACE based ADP algorithm has been successful in achieving a numerical solution (Yang et al., 2009). The baseline DACE based solution approach in this study technically avoids multicollinearity within the ADP algorithm, which is already an improvement over more common ADP approaches that do not take advantage of DACE. In order to fully address the multicollinearity issue, it is necessary to consider how the state space is represented from two aspects:

(i) If the naturally occurring state space is multicollinear, then a near-orthogonal DACE based representation is not appropriate. Example systems include energy, where market price and meteorology are correlated over time and space; the natural environment, where air quality or water quality observations are correlated over time and space; and health care, where patient characteristics are correlated due to commonly occurring medical conditions. However, it is also not appropriate to simply conduct state space sampling over the multicollinear state space. Hence, alternative approaches are needed, which motivates the second aspect.

(ii) The state transition model represents how the state changes. For many real world problems, the state transitions are unknown and must be estimated. In this case, there are two ways to handle multicollinearity. The traditional statistical modeling approach is to conduct a model-search process that enables a low-VIF state transition metamodel (Kutner et al., 2005). With such a model, the state space can appropriately be represented by near-orthogonal experimental designs in a DACE based ADP method. The other approach is to transform the multicollinear space into an orthogonal space (e.g., Ariyajunya et al., 2017; Kim, 2009), and then construct the state transition metamodel over the orthogonalized space.

Based on the two aspects presented above, in order to demonstrate the drawbacks of ignoring the multicollinearity, three types of metamodels are developed. The first type ignores the inherent multicollinearity in order to generate a high-VIF metamodel that will demonstrate the adverse effects of multicollinearity. The second type eliminates predictors from the regression model, so as to yield a low-VIF metamodel, which should be more generalizable and stable than high-VIF models. However, it is not uncommon that predictors that can improve the accuracy of a statistical model will need to be eliminated in order to avoid high VIFs. In this situation, an accurate low-VIF model will not be possible to obtain. Therefore, the third type tries to keep all the important features in the state space and eliminate multicollinearity by orthogonalizing the state space through feature extraction techniques. A challenge for this third type is that it requires the transition model to be reconstructed in the orthogonalized state space, so that the associated Bellman equation and corresponding backward ADP solution algorithm can appropriately utilize the orthogonalized state space. Since the state space is transformed to an orthogonalized space in this third type, the direct relationships with the original state variables are no longer what is modeled. Specifically, the orthogonalized state space employs linear combinations of the original state variables. While sometimes linear combinations uncover latent meanings, in general, they are not necessarily meaningful. As a consequence, the relationships constructed over this orthogonalized state space often do not provide insight into the relationships with the original state variables. For the Atlanta ozone pollution case study, we demonstrate one approach for reconstructing the transition metamodel for the orthogonalized state. However, future work will need to continue to study this empirical state transition modeling challenge.

After developing these three types of metamodels, we follow Yang et al. (2009) to apply them to conditions based on one of the worst cases in Atlanta ozone pollution history. In summary, the major contributions are as follows:

(a) The issue of multicollinearity in high-dimensional continuous-state SDP is demonstrated to encourage other ADP researchers to recognize the potential drawbacks of ignoring multicollinearity.
(b) Two options for state transition metamodeling, low-VIF and orthogonalized, are presented to develop an appropriate DACE based ADP approach for a multicollinear, continuous state space.

3. Atlanta ozone pollution control problem

The Atlanta ground-level ozone pollution SDP problem in Yang et al. (2009) includes highly correlated ozone concentration state variables at different monitoring stations and different time periods. As described in Section 1, NO_x and VOC emissions are the main precursors for high ground-level ozone concentrations in urban areas. However, Atlanta is "NO_x-limited," which means that it would be ineffective to target VOC emissions in a control strategy. Thus, the focus is on only NO_x emissions in this case study. To control NO_x emissions, it is necessary to control both point and non-point sources. Power plants and other heavy industries are categorized as point sources of NO_x emissions, while other sources, such as automobiles and small industries are treated as non-point sources. Air quality researchers utilize complex 3-D air chemistry photochemical models to study emission control strategies. Following the work of Yang et al. (2009), we utilize the Urban Airshed Model (UAM, U.S. EPA, 1990) to simulate Atlanta ozone pollution for a 160×160 sq. km region over the metropolitan area. The region contains a total of 102 point sources, and non-point sources are represented on a 5×5 grid over the region. Ozone concentrations are monitored at four Photochemical Assessment Monitoring Stations located at Conyers, South Dekalb, Tucker, and Yorkville. To reduce the ozone concentrations, emission controls are applied to specific point sources and non-point sources in the 5×5 grid and over specific time periods. Because ozone concentrations increase during the daytime when the sun is shining, only time periods from 4:00 AM to 7:00 PM are considered as potential time periods for reducing emissions.

To apply SDP to pollution control, five 3-hour time periods are defined: 4:00 AM to just before 7:00 AM (time period 0), 7:00 AM to just before 10:00 AM (time period 1), 10:00 AM to just before 1:00 PM (time period 2), 1:00 PM to just before 4:00 PM (time period 3), and 4:00 PM to just before 7:00 PM (time period 4). Time period 0 is considered as the initialization stage, and the SDP controls are applied to time periods 1 through 4. Following the SDP formulation in Eq. (1), state and decision variables are defined as follows. Decision variables ($\boldsymbol{u}_t$) are the fractions of the nominal emissions to be reduced and are component-wise defined as $u_t^i = (M_t^i - E_t^i)/M_t^i$, where M_t^i is the nominal base case emission at (point or non-point) source i in time period t, and E_t^i is the corresponding reduced emission, such that $M_t^i - E_t^i$ is the reduction in the emissions. State variables ($\boldsymbol{x}_t$) at the beginning of SDP stage t consist of maximum hourly-averaged ozone concentrations at the four monitoring stations in time periods t – 1 and earlier, and decisions on emission reductions from time periods t – 1 and earlier. Because the last time period has the longest history, it has the highest dimension, and consequently defines the dimensionality of the SDP problem. Specifically, SDP stage 4 consists of 524 state variables, covering 4 monitoring stations, 25 (5×5) grid squares for non-point sources, and 102 point sources over 4 time periods:

4 time periods $\times$ (4 stations + 25 non-point sources + 102 point sources) = 524.

The decision variables provide targeted information on where and when to reduce emissions, but assessing direct control interventions is not the purpose of the present study. Yang et al. (2009) provided a discussion in Section 4.4 on implementation considerations, and related work studying potential control strategies in State Implementation Plans, such as emission reductions at brick kilns, traffic signal improvement, or carpooling programs, has been published by Sule, Chen and Sattler (2011) and Hsu, Rosenberger, Sule, Sattler and Chen (2014).

The UAM was used both for conducting data collection to build the state transition metamodel (Yang et al., 2007) and for evaluating the SDP control policies (Yang et al., 2009), including statistical models for maximum hourly-averaged ozone O_t^S at station S in time period t. For $t = 1, 2, 3$, the models for O_t^S are part of the estimated SDP state transition function in Eq. (1), which depends on state, decision, and random vectors. For $t = 4$, O_t^S is a final outcome of the evolution of ozone at station S over the day, but can still be modeled like a transition function. Let the vector $\boldsymbol{O}_t$ hold the ozone concentrations O_t^S for all four monitoring stations. For SDP stage t, the state $\boldsymbol{x}_t$ includes the past history on both ozone $\{\boldsymbol{O}_1, \boldsymbol{O}_2, \ldots, \boldsymbol{O}_{t-1}\}$ and emission reduction decisions $\{\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_{t-1}\}$, the decision vector is $\boldsymbol{u}_t$, and the statistical metamodel form assuming an additive random error ε can be written as in Eq. (2):

$$O_t^S = f_t^S(\boldsymbol{x}_t, \boldsymbol{u}_t, \varepsilon) = f_t^S(\boldsymbol{O}_1, \boldsymbol{O}_2 \ldots, \boldsymbol{O}_{t-1}, \boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_t) + \varepsilon, \tag{2}$$

where $f_t^S(\cdot)$ is the function to be approximated for station S in time period t. As shown in Fig. 2, ozone levels $\boldsymbol{O}_1, \boldsymbol{O}_2, \ldots, \boldsymbol{O}_{t-1}$ can involve high multicollinearity (Ariyajunya et al., 2017). As noted in Section 2.3, the state transitions are unknown and are estimated via a statistical metamodel using data from the Atlanta UAM. Technically, the UAM may be used as a state transition function to predict the state values at next stage. However, the direct application of the UAM in SDP is impractical because of the high computational cost. Yang et al. (2007) employed stepwise regression for state transition metamodeling, and variants of the regression approach were explored by Ariyajunya et al. (2017) to provide guidance on addressing the multicollinearity. Details on state transition metamodeling for the current study are given in Section 4.

The main goal of the Atlanta ozone pollution problem is to maintain the ozone level to satisfy the EPA standard. Yang et al. (2009) used the one-hour EPA ozone standard of 0.125 ppm (parts per million), although the standard has since decreased (see https://www.epa.gov/criteria-air-pollutants). The stagewise objective function from Yang et al. (2009), specified in Eq. (3), is divided into two parts, namely, the emission reduction cost function using $c_e(\cdot)$ in Eq. (4) and an ozone penalty cost function using $c_{max}(\cdot)$ in Eq. (5):

$$c_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t) = \alpha \sum_{u_t^i \in \boldsymbol{u}_t} W_t^i c_e\left(u_t^i\right) + \beta \sum_S c_{max}\left(O_t^S\right), \tag{3}$$

$$c_e(u) = \begin{cases} 0, & u \le 0 \\ 4u^3 - 4u^4, & 0 < u < 0.5 \\ u - 0.25, & u \ge 0.5 \end{cases} \tag{4}$$

$$c_{max}(x) = \begin{cases} 0, & x \le 0.118 \\ 2.5 \times 10^{11}(x - 0.118)^3 & \\ -6.25 \times 10^{13}(x - 0.118)^4, & 0.118 < x < 0.12 \\ 10^6(x - 0.119), & x \ge 0.12 \end{cases} \tag{5}$$

The index t in Eq. (3) denotes the four SDP stages. Both $c_e(\cdot)$ and $c_{max}(\cdot)$ follow smooth convex functional forms that have continuous second derivatives to facilitate the use of a fast sequential quadratic programming optimization routine from the NAG library (NAG, 1991). The cost function $c_e(\cdot)$ in Eq. (4) follows a slowly rising polynomial curve when the fraction of reduction is less than 0.5 and then rises linearly for higher fractions. This form is consistent with practice in that it is relatively easy to achieve small emission reductions, but higher reductions involve higher investment. The fraction of reduction is also constrained within [0,1]. In Eq. (3), W_t^i is a weighting factor for the emission source i, where a source with higher nominal base case emissions has proportionally higher W_t^i, such that the per unit cost of reduction is equivalent across all emission sources. The ozone penalty cost in Eq. (5) applies no penalty if the estimated maximum hourly-averaged ozone stays below 0.118 ppm. While the standard is 0.125 ppm, a more conservative cutoff was chosen to allow a margin for error. Between 0.118 and 0.120, a slowly rising penalty is applied, while above 0.120,

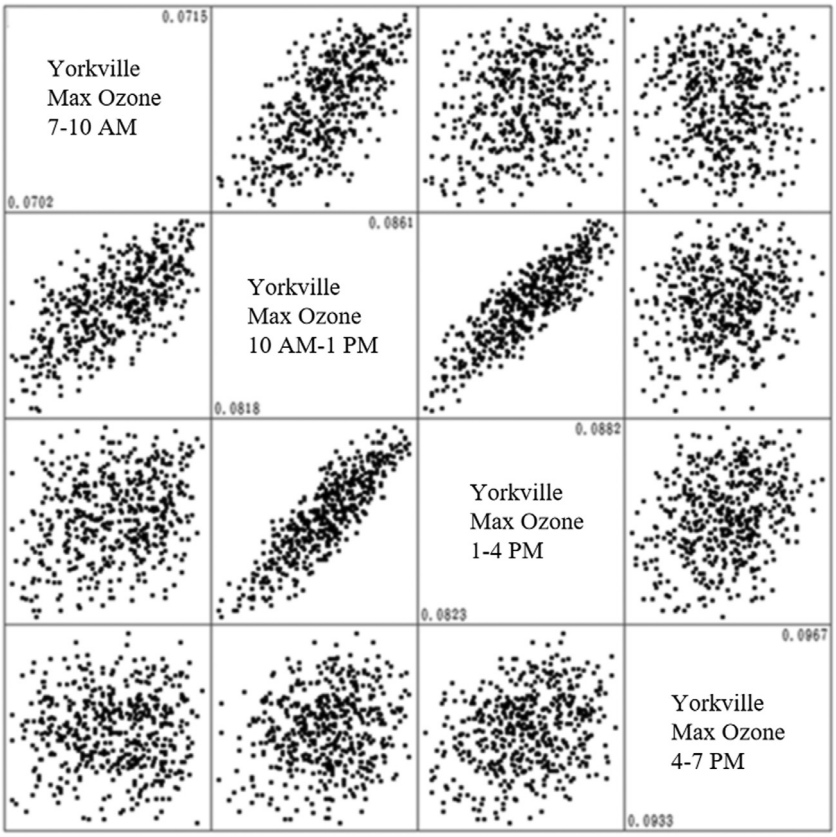

Fig. 2. Scatterplots showing correlations between the state variables of four time periods in Yorkville (Ariyajunya et al., 2017).

the penalty rises quickly, so that the optimization prioritizes reducing ozone. In order to balance the emission reduction cost and the ozone penalty, the α and β values are chosen such that the ozone penalty dominates the emission reduction cost when the maximum hourly-averaged ozone exceeds the EPA standard. To numerically solve for the SDP control policy, we implement the DACE based ADP approach based on Yang et al. (2009) using the state transition metamodels developed in Section 4. ADP implementation details are given in Section 5 with the computational results.

4. Development of state transition functions

When Yang et al. (2007) exploited statistical modeling methods to estimate the state transition functions, the multicollinearity issue in the state space was ignored. In this study, we used VIFs, as in Ariyajunya et al. (2017), to examine the degree of the multicollinearity of the state space. They defined a metamodel with VIFs < 4 as having low multicollinearity, and a metamodel with maximum VIF > 10 as having high multicollinearity. A VIF of 10 means that the variance of that model coefficient estimator is 10 times higher than it should be, resulting in an unstable model.

Three state transition metamodels are developed in this section, the high-VIF metamodel (Section 4.1), the low-VIF metamodel (Section 4.2), and the orthogonalized metamodel (Section 4.3). Recall that the purpose of the high-VIF metamodel is to illustrate the situation when high multicollinearity is ignored, while the other two metamodel types are approaches for addressing state space multicollinearity in ADP. Our process for estimating the state transition metamodels follows the approach of Yang et al. (2007), which is summarized in Fig. 3.

The initialization, data collection, and mining phases were previously conducted (Yang et al., 2007), so our descriptions for each metamodel type focus on the modeling phase. Following the mining phase, 87 of the 102 point sources are eliminated, while all other state variables are kept to move on to the modeling phase. Additional dimension reduction occurs in the modeling phase, depending on which variables are selected for the final state transition metamodel. All three metamodel types are constructed using the same data collected from the Atlanta UAM, and start with the same set of variables following the mining phase. The summary count of variables in post-mining phase and post-modeling phase for each of the three metamodel types is given in Table 1. For listings of the decision and state variables and their estimated coefficients for the three metamodel types, refer to the Appendix in Ariyajunya (2012).

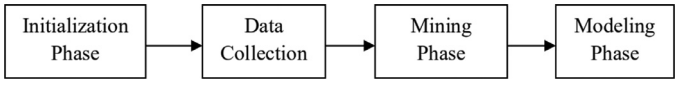

Fig. 3. State transition metamodeling process.

4.1. High-VIF metamodels

The high-VIF state transition metamodel is deliberately developed to represent the situation when high multicollinearity is ignored. To build high-VIF metamodels, only the emission decision variables $(\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_t)$ are considered for elimination in the mining phase, while all ozone levels $(\boldsymbol{O}_1, \boldsymbol{O}_2, \ldots, \boldsymbol{O}_{t-1})$ are maintained for the modeling phase. Table 2 provides summary information on the high-VIF state transition metamodel. In particular, it can be seen that the maximum VIFs are above 10 for all models, except for the Conyers maximum ozone level during 7–10 AM. The maximum VIF for the Yorkville maximum ozone level during 1–4 PM indicates that a model coefficient estimator's variance is over 1300

Table 1
Summary count of the state and decision variables for the different metamodel types.

Stage	Variables	Number of variables post-mining phase	State transition metamodel type		
			High-VIF	Low-VIF	Orthogonalized
Stage 1	Decision variables	40	29	17	29
	State variables	44	34	16	25
	Total	84	63	33	54
Stage 2	Decision variables	40	31	9	28
	State variables	88	59	23	23
	Total	128	90	32	51
Stage 3	Decision variables	40	30	9	25
	State variables	132	82	21	14
	Total	172	112	30	39
Stage 4	Decision variables	40	12	3	7
	State variables	176	92	19	9
	Total	216	104	22	16

Table 2
Summary of the high-VIF ozone state transition metamodel.

Maximum ozone model	Model R^2	Root MSE	Maximum VIF
Conyer 7–10 AM	0.2682	0.0007	1.0859
S.Dekalb 7–10 AM	0.9864	0.0006	17.8589
Tucker 7–10 AM	0.9612	0.0012	62.0089
Yorkville 7–10 AM	0.9945	<0.0001	62.1610
Conyer 10 AM-1 PM	0.9937	0.0003	30.2103
S.Dekalb 10 AM-1 PM	0.2642	0.0056	69.7131
Tucker 10 AM-1 PM	0.6370	0.0027	69.4942
Yorkville 10 AM-1 PM	0.9993	<0.0001	163.6091
Conyer 1–4 PM	0.9846	0.0006	76.9823
S.Dekalb 1–4 PM	0.9920	0.0010	75.9784
Tucker 1–4 PM	0.9747	0.0010	74.8509
Yorkville 1–4 PM	0.9994	<0.0001	1366.4426
Conyer 4–7 PM	0.9847	0.0013	95.7745
S.Dekalb 4–7 PM	0.9930	0.0009	86.5118
Tucker 4–7 PM	0.9891	0.0008	92.4630
Yorkville 4–7 PM	0.9994	<0.0001	374.8161

Table 3
Summary of the low-VIF ozone state transition metamodel.

Maximum ozone model	Model R^2	Root MSE	Maximum VIF
Conyer 7–10 AM	0.2646	0.0007	1.0149
S.Dekalb 7–10 AM	0.9855	0.0007	1.0075
Tucker 7–10 AM	0.9607	0.0013	1.0161
Yorkville 7–10 AM	0.9942	<0.0001	1.0158
Conyer 10 AM–1 PM	0.9935	0.0003	1.2438
S.Dekalb 10 AM–1 PM	0.1954	0.0058	1.0306
Tucker 10 AM–1 PM	0.6080	0.0028	1.0282
Yorkville 10 AM–1 PM	0.9992	<0.0001	1.0128
Conyer 1–4 PM	0.9808	0.0007	1.0179
S.Dekalb 1–4 PM	0.9692	0.0019	1.0297
Tucker 1–4 PM	0.9536	0.0014	1.3483
Yorkville 1–4 PM	0.9990	0.0000	2.3313
Conyer 4–7 PM	0.9625	0.0019	3.4408
S.Dekalb 4–7 PM	0.9801	0.0014	1.5708
Tucker 4–7 PM	0.9308	0.0019	1.5675
Yorkville 4–7 PM	0.9624	0.0001	1.0166

times higher than it should be, but the Model R^2 and Root MSE indicate a high quality fit to the data. This high-VIF metamodel demonstrates that quality of fit alone is insufficient for assessing empirical models when multicollinearity is present.

4.2. Low-VIF metamodels

To ensure the creation of low-VIF metamodels, after the mining phase in Fig. 3, we add a multicollinearity assessment procedure using VIFs to remove predictor variables that contribute to high VIFs, re-fit stepwise regression and re-evaluate the metamodel until all VIFs are below 4. In the modeling phase, all maintained predictor variables must be statistically significant at a significance level of 0.05. It must be noted that the process to identify a low-VIF metamodel is a time-consuming manual search process to find that "best" low-VIF empirical model, and may not always yield a quality model. The summary of the low-VIF state transition metamodel is presented in Table 3. Most of the VIFs are near 1.0, indicating nearly no variance inflation, despite the presence of high multicollinearity among the *full* set of predictor variables. In other words, there is very little multicollinearity among the *selected* set of predictor variables.

To illustrate the difference between low-VIF and high-VIF metamodels, Fig. 4 compares the estimated model coefficients for predicting the maximum hourly-averaged ozone at the Yorkville monitoring station in the four time periods. According to Tables 2 and 3, the fit is excellent for all these models, but in Fig. 4 it is clear that the high-VIF model is very different from the low-VIF model. This demonstrates the challenge of identifying good models in the presence of multicollinearity because many models have a similar fit to the data. For Yorkville, high quality low-VIF models were possible, and these models are seen to have fewer predictors. Further, the high-VIF models exhibit coefficients with the incorrect sign, namely a negative relationship between ozone variables. Blind machine learning algorithms that do not intelligently address multicollinearity can encounter similar issues.

4.3. Orthogonalized metamodels

An alternative approach to addressing multicollinearity is to orthogonalize the state space. To conduct this, an orthogonalization phase is added after the mining phase, as shown in Fig. 5. Ariyajunya et al. (2017) explored several modeling approaches for the orthogonalization of the state space, and their comparison, which was only conducted for the last stage of the SDP problem, recommended using stepwise regression and partial least squares (PLS). We refer to this as stepwise-PLS in the orthogonalization phase. Compared to the low-VIF modeling procedure, the process in Fig. 5 is completely automated. However, as previously mentioned, the direct relationships with the original state variables after state space transformation are no longer what is modeled, which brings a new challenge to the ADP solution approach. Note that only the state space is orthogonalized while the current stage decision variables are maintained in their original form for the stagewise optimization.

As described in Kim (2009), PLS performs fitting and transformation simultaneously during the modeling process. Consequently, PLS cannot be directly used to build the state transition metamodel

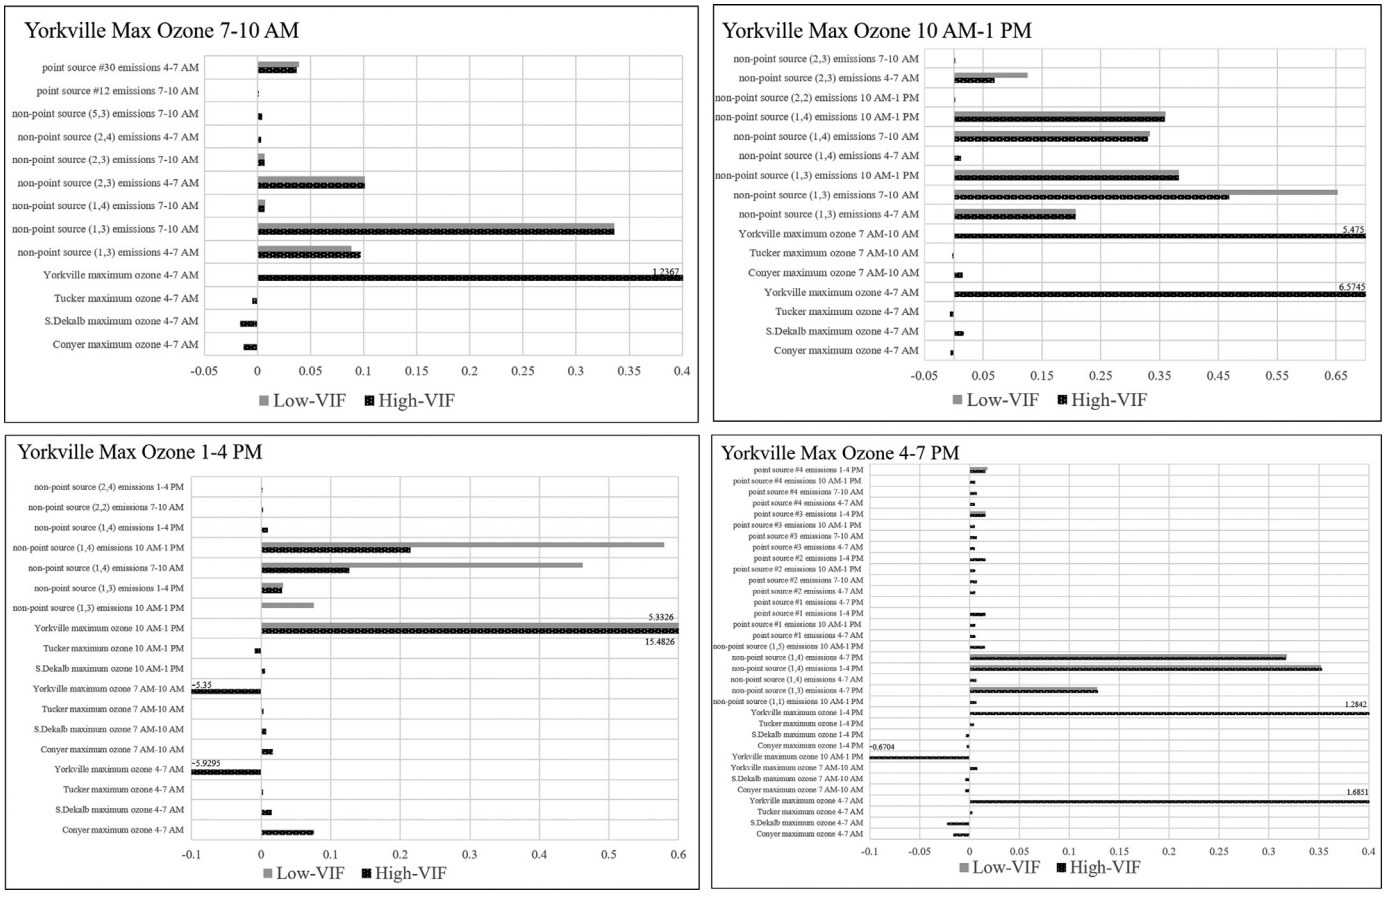

Fig. 4. Coefficients for the Yorkville low-VIF and high-VIF models in the four time periods.

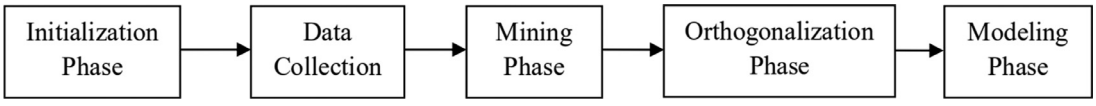

Fig. 5. Process for development of orthogonalized state transition metamodels.

because the decision variables are not orthogonalized. To incorporate the orthogonalized state variables from PLS into a backward ADP algorithm, the modeling phase in Fig. 5 employs stepwise regression that is divided into two parts. First, in each SDP stage t, the ozone variables $\boldsymbol{O}_t$ are modeled as a function ($g1_t$) of the orthogonalized state variables $\boldsymbol{z}_t$ in Eq. (6), and the orthogonalized future state, $\boldsymbol{z}_{t+1}$ is also be modeled as a function ($f1_t^z$) of $\boldsymbol{z}_t$ in Eq. (7):

$$\boldsymbol{O}_t = g1_t(\boldsymbol{z}_t) + \boldsymbol{\varepsilon} \tag{6}$$

$$\boldsymbol{z}_{t+1} = f1_t^z(\boldsymbol{z}_t) + \boldsymbol{\varepsilon} \tag{7}$$

Second, the deviations $\{\boldsymbol{O}_t - \widehat{g1}_t(\boldsymbol{z}_t)\}$ are modeled as a function ($g2_t$) of the decision vector $\boldsymbol{u}_t$ in Eq. (8). Similarly, the deviations $\{\boldsymbol{z}_{t+1} - \widehat{f1}_t^z(\boldsymbol{z}_t)\}$ are modeled as a function ($f2_t^z$) in Eq. (9):

$$\boldsymbol{O}_t - \widehat{g1}_t(\boldsymbol{z}_t) = g2_t(\boldsymbol{u}_t) + \boldsymbol{\varepsilon} \tag{8}$$

$$\boldsymbol{z}_{t+1} - \widehat{f1}_t^z(\boldsymbol{z}_t) = f2_t^z(\boldsymbol{u}_t) + \boldsymbol{\varepsilon} \tag{9}$$

Finally, the approximated $\hat{\boldsymbol{O}}_t$ and $\hat{\boldsymbol{z}}_{t+1}$ combine the two parts in Eqs. (10) and (11):

$$\hat{\boldsymbol{O}}_t = \widehat{g1}_t(\boldsymbol{z}_t) + \widehat{g2}_t(\boldsymbol{u}_t) \tag{10}$$

$$\hat{\boldsymbol{z}}_{t+1} = \widehat{f1}_t^z(\boldsymbol{z}_t) + \widehat{f2}_t^z(\boldsymbol{z}_t) \tag{11}$$

Table 4
Summary of the stepwise-PLS ozone state transition metamodel.

Maximum Ozone Model	Model R^2	Root MSE	Part 1 Max VIF	Part 2 Max VIF
Conyer 7–10 AM	0.2759	0.0007	1	1.0040
S.Dekalb 7–10 AM	0.9184	0.0016	1	1.0033
Tucker 7–10 AM	0.9149	0.0019	1	1.0025
Yorkville 7–10 AM	0.9551	<0.0001	1	1.0071
Conyer 10 AM–1 PM	0.9348	0.0011	1	1.0009
S.Dekalb 10 AM–1 PM	0.2847	0.0055	1	1.0088
Tucker 10 AM–1 PM	0.6400	0.0027	1	1.0001
Yorkville 10 AM–1 PM	0.9690	0.0002	1	1.0048
Conyer 1–4 PM	0.9430	0.0012	1	1.0043
S.Dekalb 1–4 PM	0.9318	0.0029	1	1.0007
Tucker 1–4 PM	0.9156	0.0019	1	1.0049
Yorkville 1–4 PM	0.9735	0.0002	1	1.0000
Conyer 4–7 PM	0.9795	0.0014	1	1.0001
S.Dekalb 4–7 PM	0.9858	0.0012	1	1.0074
Tucker 4–7 PM	0.9469	0.0017	1	1.0171
Yorkville 4–7 PM	0.9231	0.0002	1	1.0016

A summary of the results of the stepwise-PLS metamodel for each stage is presented in Table 4. To verify the success of the PLS orthogonalization, it is seen that all VIFs from Part 1 show no variance inflation, which indicates that the state space

Table 5
DACE based ADP implementation details for all runs.

DoE for state space discretization	2000-point Sobol sequence
Ozone threshold	0.125 ppm (modeled in penalty functions)
Negative coefficients in ozone models	Truncated to zero
MARS approximation algorithm	MARS ASR-II
Maximum basis functions for MARS	2000
Maximum order of interaction in MARS	2
Number of candidate knots per dimension	35
Non-linear optimization library	NAG Fortran Mark 15
Optimization starting points for stages 1 and 2	Midpoint, lower bound, and 10 random points
Optimization starting points for stages 3 and 4	Midpoint and lower bound
Running environment	Workstation with dual 2.6 G AMD Atlon processors and 3 gigabyte memory Cent OS 4.9 gcc version 3.4.6 20,060,404 (Red Hat 3. 4. 6–9)

multicollinearity has been eliminated. In Part 2, the emission decision variables are not orthogonalized; however, we still see that the DACE approach maintains VIFs that are very close to 1.0, indicating nearly no multicollinearity or near orthogonality in Part 2.

5. Computational results

In order to solve the Atlanta ozone pollution, high-dimensional, continuous-state SDP problem over a multicollinear state space, the DACE based ADP algorithm described in Section 2.2 is employed with each of the three state transition metamodel types described in Section 4 (high-VIF, low-VIF, and stepwise-PLS). The resulting SDP control policies to control NO_x emission reduction decisions are then forward simulated within the Atlanta UAM, so as to compare their performance. In Section 5.1, details are provided on the DACE based ADP process for the Atlanta ozone pollution SDP problem. In Section 5.2, comparisons using the Atlanta UAM are presented. Finally, in Section 5.3, the accuracy of the state transition metamodels is verified.

5.1. DACE based ADP for the Atlanta ozone pollution SDP problem

The solution of the ozone pollution problem using the DACE based ADP method begins at the last stage and moves backward until all the stages have been solved, as illustrated in Fig. 1. Following Yang et al. (2009), a 2000-point low-discrepancy sequence (Sobol, 1967) is used as the experimental design to discretize the state space, and MARS is used to approximate the FVFs. At each discretization point, a commercial sequential quadratic programming algorithm (NAG E04) is used to obtain an optimal solution that corresponds to a point on the FVF. Unlike in the work of Yang et al. (2009), where the MARS approximations are allowed to have negative values, the negative MARS values are truncated to zero in the present study, because a negative cost is unrealistic.

To reduce the possibility of achieving local optima, multiple starting points are used. However, while the use of many starting points increases the chance of approaching a global optimum, for computational reasons, only two starting points (the midpoint and the lower bound) and an additional ten random points between the two are employed in the present study. Further, previous numerical experiments have demonstrated that the use of multiple starting points tends to produce better overall results compared with the use of only one (the middle point), especially in time periods 1 and 2. Details on our DACE based ADP implementation are given in Table 5. Note that negative coefficients in the ozone metamodels associated with the emission decision variables are truncated to zero to avoid undesirable relationships that would cause the optimization to increase emissions to reduce ozone. These undesirable relationships can occur when NO_x levels are high, as explained in Yang et al. (2007). As a consequence, by truncating these negative coefficients to zero, the accuracy of the developed metamodels is degraded.

Table 6 summarizes the running time and MARS approximations for the FVFs in each stage. The MARS algorithm running time grows linearly with the number of input variables, but grows as a cubic polynomial with the number of basis functions (Friedman, 1991). In general, a more complex function requires more MARS basis functions. However, it is suspected that the more complex MARS FVF approximations for stages 1 and 2 are due to the presence of more local optima in the FVF data used to fit MARS than due to the underlying true FVF having a significantly more complex structure. Please refer to Ariyajunya (2012: Chapter 5) for some figures of the MARS FVF approximation functions. It is also seen in Table 6 that the ADP solution time is much greater when using the stepwise-PLS state transition metamodel. The computational time for the ADP solution is primarily subject to the nonlinearity of the stagewise optimization. In the case of the stepwise-PLS state transition metamodel, the direct relationships with the original state variables are no longer what is modeled. Instead, the orthogonalized state variables possibly induce more complex and nonlinear relationships that, in turn, create a more complex nonlinear optimization. As seen in Eq. (1), the state transition function is an integral part of the SDP optimization. Overall, the computational times are mostly provided as informational and do not provide judgment of the quality of an SDP control policy. Further, the ADP algorithm is executed off-line prior to implementation, so the long computational times in Table 6 do not reflect an inability to employ the resulting SDP control policies in real time. The re-optimization required to implement the SDP control policy executes within seconds. In the next section, the SDP control policies for the high-VIF, low-VIF, and stepwise-PLS cases, represented by the MARS FVF approximations, are compared via results from a forward simulation that mimics real-time implementation.

5.2. Simulation results

The re-optimization mentioned in Section 2.2 is used to simulate an SDP control policy forward in real time to generate the specific sequence of NO_x emission reductions (point and non-point sources) over the time periods. Each sequence of NO_x emission reductions is implemented in the Atlanta UAM to generate the resulting maximum hourly-averaged ozone. If the SDP control policy performs well, then the UAM results should be similar to the results from the forward simulation using re-optimization. In this section, we study 50 hypothetical scenarios that were generated by randomly varying the initial state conditions in the 4:00–7:00 AM time period from the actual ozone episode during July 29-August 1, 1987, which is one of the worst on record. Please refer to Appendix E in Ariyajunya (2012) for the 50 initial random state vectors.

Table 6
Number of MARS basis functions and the running times.

Model	Stage	No. of state variables	No. of decision variables	No. of basis functions selected by MARS	Fitting MARS (hh:mm:ss)	ADP solution (hh:mm:ss)	Total running time (hh:mm:ss)
High-VIF	Stage 1	34	29	1296	30:01:51	0:59:24	31:01:23
	Stage 2	59	31	300	1:01:09	4:40:24	5:41:33
	Stage 3	82	30	227	0:54:26	0:23:25	1:17:51
	Stage 4	92	12	182	1:57:53	0:02:24	2:00:17
				Total time	33:55:19	6:05:45	40:01:04
Low-VIF	Stage 1	16	17	394	0:53:31	1:07:57	2:01:28
	Stage 2	23	9	1853	50:09:49	0:16:44	50:26:33
	Stage 3	21	9	104	0:02:47	0:05:21	0:08:08
	Stage 4	19	3	90	0:02:30	0:00:32	0:03:02
				Total time	51:08:37	1:30:34	52:39:11
Stepwise-PLS	Stage 1	25	29	1354	24:27:20	29:49:16	54:16:36
	Stage 2	23	28	964	7:57:58	27:54:13	35:52:11
	Stage 3	14	25	215	0:07:14	4:32:07	4:39:21
	Stage 4	9	7	72	0:00:32	0:03:58	0:04:30
				Total time	32:33:04	62:19:34	94:52:38

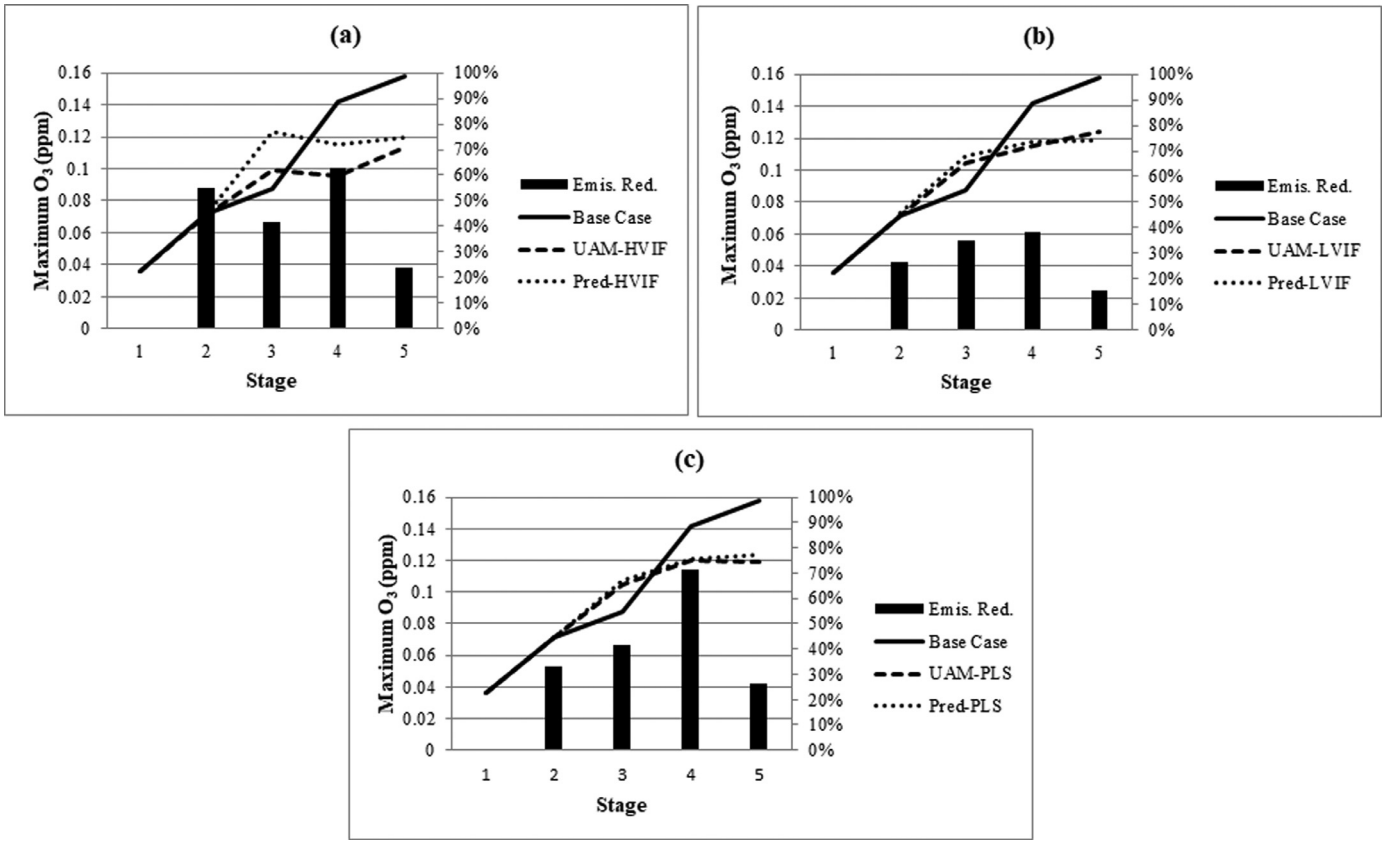

Fig. 6. Hypothetical scenarios (50): average of the maximum hourly-averaged ozone levels from the nominal base case and from the ADP solutions, both simulated via the UAM and directly predicted using the three state transition metamodels, (a) high-VIF (HVIF), (b) low-VIF (LVIF), and (c) stepwise-PLS (PLS), along with the corresponding emission reductions (Emis. Red.).

The average emission reductions required by the SDP control policies are presented in Fig. 6 and summarized in Table 7. Fig. 6 compares the trajectory for maximum hourly-averaged ozone using the nominal base case historical data and the average of ozone levels over 50 hypothetical scenarios simulated by the Atlanta UAM and predicted directly using the state transition metamodel. It can be seen in Fig. 6 that the nominal base case ozone level rises significantly above the EPA ozone standard of 0.125 ppm by the last time period. The three re-optimized SDP control policies are clearly able to lower ozone levels. However, it can be seen for the SDP control policy using the high-VIF metamodel, that the UAM-simulated ozone levels are noticeably below those predicted by the high-VIF metamodel. By contrast, the SDP control policies using the low-VIF and stepwise-PLS metamodels that account for the inherent state space multicollinearity yield UAM-simulated ozone levels that are very close to the predicted ones. As a consequence, the SDP re-optimization using the FVF approximation derived via the high-VIF metamodel selects much higher emission reductions. Since emission control strategies can be costly, it is not desirable to require more emission reduction than necessary. Alternately, if the high-VIF metamodel had been similarly inaccurate, but in the opposite direction, under-estimating ozone instead of

Table 7
Hypothetical scenarios (50): average (Avg.) emission reductions (gm-mol) required by the SDP control policies, percent of the total nominal base case emissions, and corresponding average costs based on the weighted sum of emission reductions for the high-VIF, low-VIF, and stepwise-PLS state transition metamodels.

	High-VIF			Low-VIF			Stepwise-PLS		
Time period	Avg. Reduction	%	Avg. Cost	Avg. Reduction	%	Avg. Cost	Avg. Reduction	%	Avg. Cost
1	1667,499	55.1	45,864	811,482	26.8	19,822	1001,731	33.1	15,922
2	1060,932	41.5	25,274	900,083	35.2	24,375	1065,100	41.6	22,803
3	1630,610	62.7	32,337	997,836	38.4	27,692	1850,742	71.1	37,815
4	735,351	23.8	35,883	473,427	15.3	27,612	813,879	26.4	36,057
Daily	5094,392	45.2	139,358	3182,828	28.2	99,501	4731,453	42.0	112,597

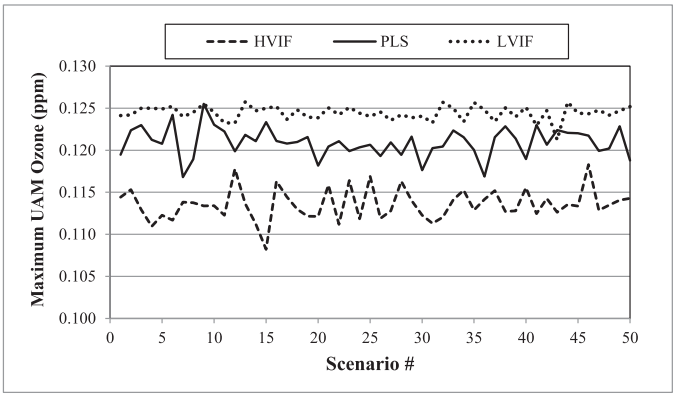

Fig. 7. Maximum UAM hourly-averaged ozone for all 50 hypothetical scenarios using each state transition metamodel type: high-VIF, low-VIF, and stepwise-PLS.

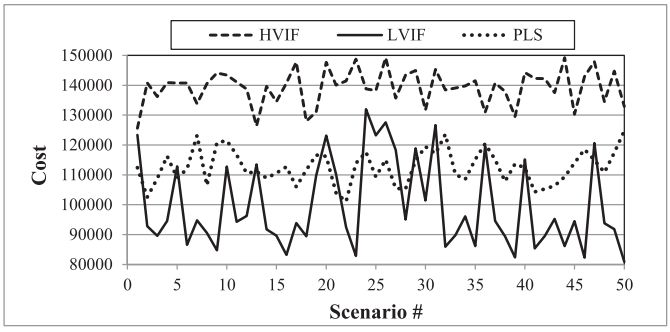

Fig. 8. Weighted sum costs for all 50 hypothetical scenarios using each state transition metamodel type: high-VIF, low-VIF, and stepwise-PLS.

over-estimating, then the ADP solution would have recommended lower emission reductions, while risking exceedingly high ozone levels.

Fig. 7 shows the maximum hourly-averaged ozone levels from the Atlanta UAM for each of the 50 hypothetical scenarios. Because the high-VIF state transition metamodel severely overestimates ozone, as seen in Fig. 6(a), higher emission reductions are selected and the resulting UAM-simulated ozone levels are unnecessarily far below the EPA standard of 0.125 ppm in Fig. 7. The stepwise-PLS metamodel leads to emission controls with the best performance in Fig. 7, showing ozone levels at or just below the EPA standard. While the low-VIF metamodel achieves an average ozone level below the EPA standard in Fig. 6(b), the ozone levels across the 50 scenarios in Fig. 7 show several instances (10 out of 50) of exceeding the EPA standard. As reported by EPA (2019), such exceedances can increase the frequency of asthma attacks and aggravate lung diseases, which correlates with increased hospital admissions and costs to the healthcare system.

Table 7 provides a summary of the averages over the 50 scenarios for the emission reductions, percentages, and weighted sum costs required by the ADP solutions for each of the three metamodels in each of the time periods. The high-VIF metamodel leads to the highest total emission reduction and the corresponding highest weighted sum cost, due to the overestimation of the ozone levels seen in Fig. 6(a). The low-VIF metamodel leads to the lowest total emission reduction and the lowest weighted sum cost. Using the stepwise-PLS metamodel, the total emission reduction is about 48% higher than using the low-VIF metamodel, but the weighted sum cost is still only about 13% higher. In other words, while using the stepwise-PLS metamodel leads to much higher emission reductions, they are located at lower-weighted emission sources than those resulting from using the low-VIF metamodel. Using the high-VIF metamodel, the total emission reduction is about 60% higher than using the low-VIF metamodel, and the weighted sum cost is 40% higher. Fig. 8 shows the weighted sum costs for each of the 50 hypothetical scenarios. It is clear that the high-VIF method results in higher weighted sum costs, uniformly across all the scenarios. The costs for the low-VIF and stepwise-PLS methods exhibit some overlap. However, while the low-VIF method leads to lower costs, on average, the stepwise-PLS method leads to lower variance on the costs, indicating the stepwise-PLS method yields a less risky (more stable) SDP control policy. Markowitz (1952) pioneered the concept of higher variance being associated with a riskier solution that is undesirable in optimization. In addition, Fig. 7 shows that the low-VIF metamodel leads to several scenarios with ozone levels above the EPA standard, indicating that the low-VIF emission controls are insufficient in these scenarios. Hence, if we only consider the emission cost objective, the low-VIF metamodel leads to the best performing SDP control policy. However, if we consider both the emission cost and air quality objectives, the stepwise-PLS metamodel leads to the best performing SDP control policy with ozone levels at or below the EPA standard, achieved with low weighted sum cost. Overall, both the low-VIF and stepwise-PLS metamodels lead to superior SDP control policies compared to the high-VIF case that ignores multicollinearity, demonstrating the negative impact that multicollinearity can have on SDP control policies and the success of both low-VIF and orthogonalization approaches in mitigating this impact.

The Atlanta UAM 40 $\times$ 40 grid covering the metropolitan region is aggregated into a 5 $\times$ 5 grid, as shown in Figs. 9 to 12. We refer to these grid squares with the notation (i, j), for i =1, 2, 3, 4, 5 and j=1, 2, 3, 4, 5, where grid square (5, 1) is at the bottom right. With the initial state conditions from the historical base case on July 31 as an example, Figs. 9 through 12 spatially illustrate the magnitude and location of the emission reductions for the SDP control policies using each of the three metamodels in each of the four SDP stages, where point source emission reductions are combined with non-point source emission reductions based on location. The different sources of emissions are weighted, as explained earlier in Eq. (3), such that higher nominal emissions lead to higher weights in the SDP cost function. The highest nominal emissions are in the center of the region, grid square (3, 3). The second highest

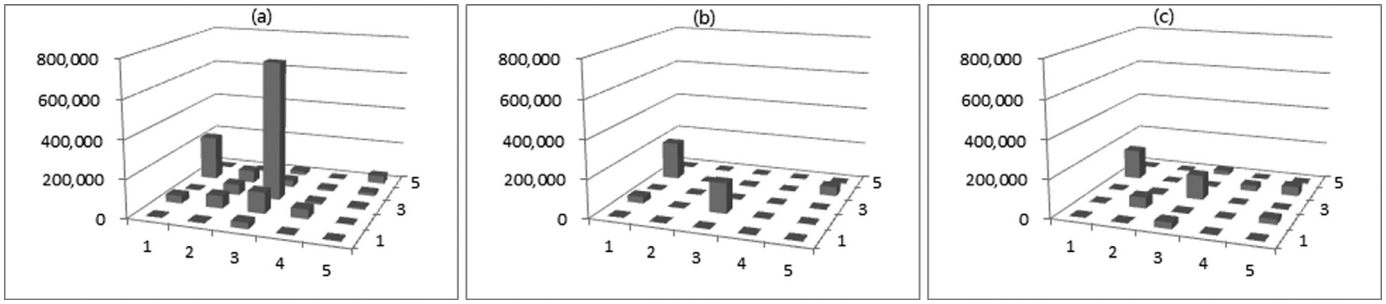

Fig. 9. Emission reductions during 7:00 AM – 10:00 AM within each grid square for the real Atlanta scenario using (a) high-VIF, (b) low-VIF, and (c) stepwise-PLS metamodels.

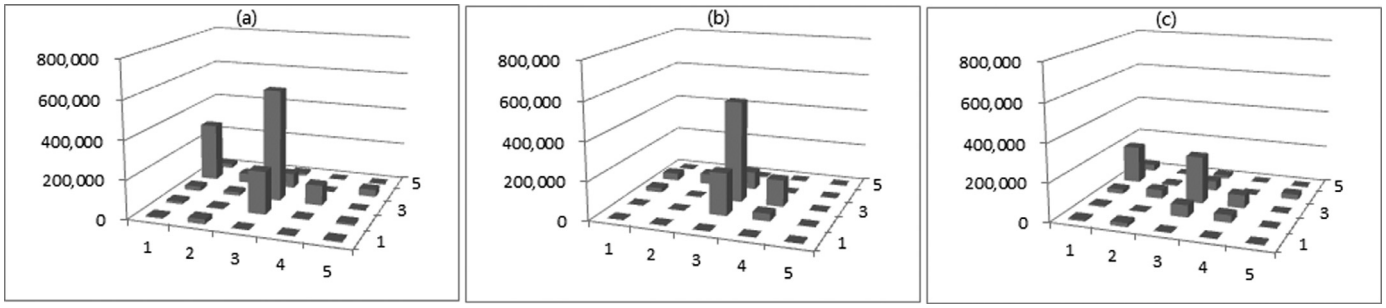

Fig. 10. Emission reductions during 10:00 AM – 1:00 PM within each grid square for the real Atlanta scenario using (a) high-VIF, (b) low-VIF, and (c) stepwise-PLS metamodels.

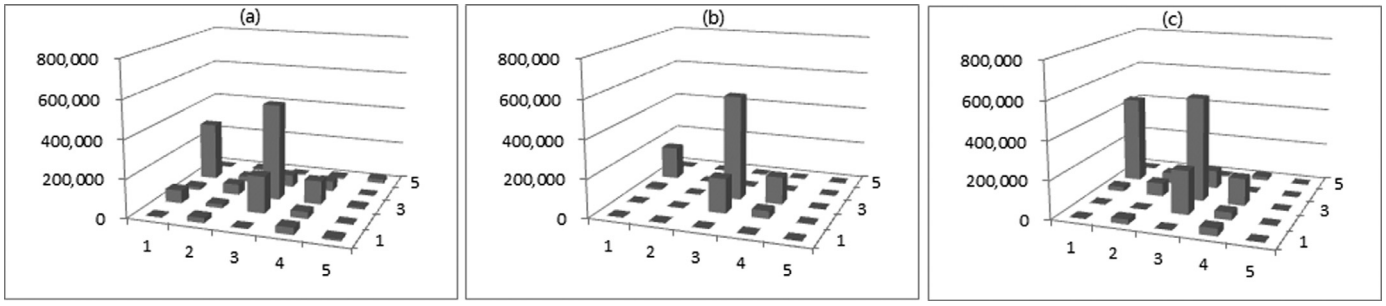

Fig. 11. Emission reductions during 1:00 PM – 4:00 PM within each grid square for the real Atlanta scenario using (a) high-VIF, (b) low-VIF, and (c) stepwise-PLS metamodels.

source, at about 1/4 of the highest, is one space in front of center, grid square (3, 2). The next two highest sources, each at about 1/2 of the second highest, are grid square (4, 3) to the right of center and a point source located on the left edge in grid square (1, 4). It can be seen in Fig. 9 that the high-VIF metamodel yields a high emission reduction in the most expensive center grid square, while the low-VIF and stepwise-PLS metamodels do not yield noticeable emission reductions in the center. In the second time period, covering the lunch hour, Fig. 10 shows all three metamodels yield emission reductions in the four expensive grid squares. However, the low-VIF and stepwise-PLS metamodels yield generally lower emission reductions, compared to the high-VIF metamodel, In the third time period, covering mid-afternoon, the three metamodels yield similar emission reduction spatial patterns in Fig. 11. Finally, in the last time period, covering the evening rush hour, all three metamodels focus emission reductions in the center, but the low-VIF metamodel has the lowest amount, as shown in Fig. 12.

5.3. Verification of metamodels

The most desirable SDP control policy for the Atlanta ozone pollution SDP problem is that which requires the least emission reductions to maintain the maximum ozone level within the EPA standard. However, the numerical results of the SDP process are affected by the accuracy of the state transition metamodels. To assess accuracy, the predicted ozone levels using the metamodels are compared to the UAM-simulated ozone levels using the SDP control policies obtained for the 50 hypothetical scenarios in Section 5.2. Recall, as stated in Section 5.1, that the accuracy of the developed metamodels is degraded due to truncation of negative coefficients. The average percent errors with respect to each monitoring station are shown in Fig. 13(a) and with respect to each SDP stage are presented in Fig. 13(b). It is seen that the high-VIF metamodel has the highest average percent errors in most cases, indicating the worst performance overall. In general, the metamodels overestimate ozone levels, on average, although the low-VIF metamodel does underestimate ozone at the stage 4 by an average percent error of 1.47%. Underestimation is riskier in controlling ozone since it could lead to insufficient emission control, as seen in Fig. 7. The stepwise-PLS metamodel has clearly larger average percent error than the low-VIF metamodel in three cases, South Dekalb, Tucker, and stage 1. In stage 1, the stepwise-PLS metamodel has a higher percent error than the high-VIF metamodel. In summary, both the low-VIF and stepwise-PLS metamodels are much more stable than the high-VIF metamodel.

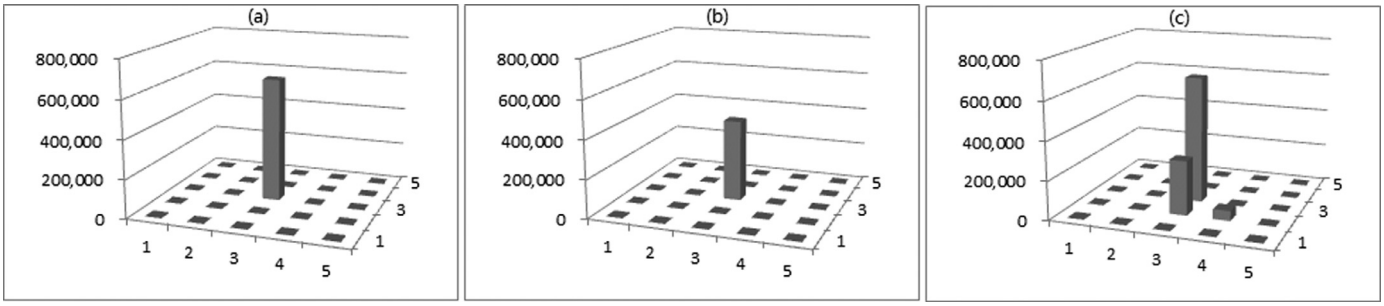

Fig. 12. Emission reductions during 4:00 PM – 7:00 PM within each grid square for the real Atlanta scenario using (a) high-VIF, (b) low-VIF, and (c) stepwise-PLS metamodels.

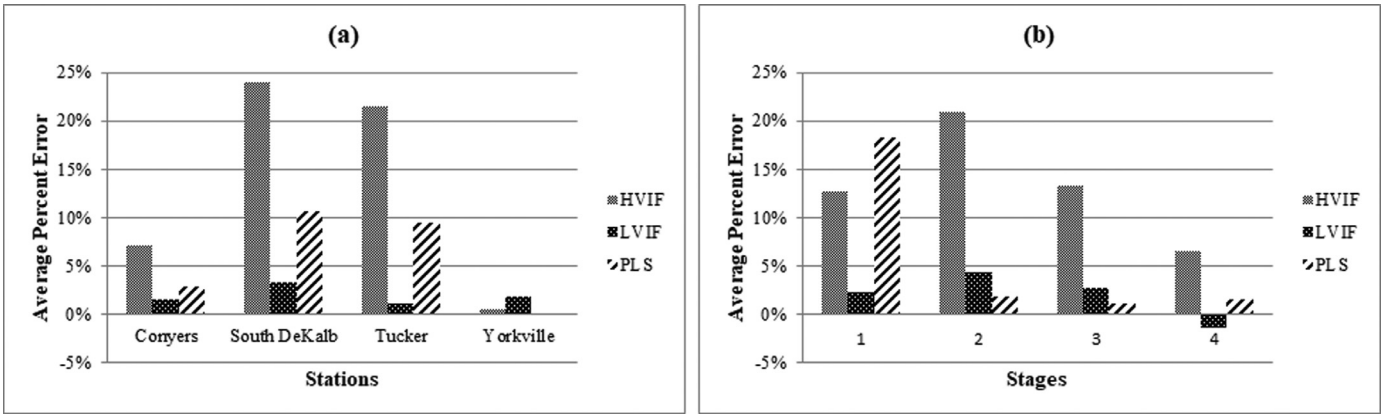

Fig. 13. Average percent error between the metamodel predictions and UAM-simulated ozone for the 50 hypothetical scenarios by (a) stations and (b) stages.

6. Concluding remarks

In this study, we focus on the issue of high state space multicollinearity for a ground-level ozone pollution SDP problem for Atlanta, Georgia. The drawbacks of ignoring high state space multicollinearity in ADP are demonstrated, and two state transition modeling approaches are presented, combined with DACE based ADP, to appropriately address the state space multicollinearity. An important goal of our study is to raise the awareness that multicollinear state spaces require care when modeling in ADP.

For the high-dimensional, continuous-state Atlanta ozone pollution SDP problem, we employed regression techniques to construct state transition metamodels of the Atlanta UAM over the multicollinear state space, where the high-VIF metamodel ignores the multicollinearity, the low-VIF metamodel follows best practices for handling multicollinearity, and the stepwise-PLS model transforms the multicollineaar state space to an orthogonal one. The high-VIF metamodel leads to clearly inferior SDP performance compared to the low-VIF and stepwise-PLS metamodels. While the low-VIF metamodel produces the most cost-efficient ADP solutions (Fig. 8), there are several cases in which the emission reductions are insufficient, leading to UAM-simulated ozone levels that exceeds the EPA standard by a small amount (Fig. 7). The stepwise-PLS metamodel produces ADP solutions with weighted sum cost about 13% higher on average than using the low-VIF metamodel and UAM-simulated ozone levels that comply with the EPA standard (Fig. 7). From the human health aspect, this is important since the exceedance of the EPA ozone level is known to be harmful for the vulnerable groups, such as children, the elderly, and patients with lung diseases.

With regard to recommendations and future work, both the low-VIF approach and the stepwise-PLS approach are viable for handling state space multicollinearity. The disadvantage of the low-VIF approach is that best practices are still a rather manual process and are not guaranteed to achieve a high quality low-VIF metamodel. If one is experienced with building low-VIF metamodels, and a high quality low-VIF metamodel can be constructed, then we would deem the low-VIF approach to be more straightforward to implement for SDP. While an orthogonalization approach can be automated, the disadvantage is the challenge of modeling the state transition relationships over the orthogonalized state space, which may be more complex and nonlinear than the relationships over the original state space. However, the orthogonalization approach guarantees an orthogonal (uncorrelated) representation of the state space and performed well for the Atlanta case study. Future work is needed to further study the challenge of state transition modeling. In this work, we are able to make use of the Atlanta UAM simulation model to provide data for state transition modeling, but in some applications, real world data may be used, which creates additional modeling challenges (LeBoulluec et al., 2018).

Acknowledgments

This research was partially supported by the National Science Foundation (ECCS-0801802 and CMMI-1434401), the Dallas-Fort Worth International Airport and National Natural Science Foundation of China (Grant no. 91846301) and China Postdoctoral Science Foundation (Grant no. 2020M670917).

References

Ariyajunya, B. (2012). *Adaptive dynamic programming for high-dimensional multi-collinear state spaces*. University of Texas at Arlington.
Ariyajunya, B., Chen, Y., Chen, V. C. P., & Kim, S. B. (2017). Data mining for state space orthogonalization in adaptive dynamic programming. *Expert System with Applications, 76*, 49–58.
Bellman, R. (1957). *Dynamic programming*. Princeton: Princeton University Press.
Bertsekas, D. (2005). *Dynamic programming and optimal control*. Belmont, MA: Athena Scientific.
Box, G. E. P., Hunter, W. G., & Hunter, J. S. (1978). *Statistics for experimenters: An introduction to design, data analysis, and model building*. New York: Wiley.
Cervellera, C., Chen, V. C. P., & Wen, A. (2006). Optimization of a large-scale water reservoir network by stochastic dynamic programming with efficient state space discretization. *European Journal of Operational Research, 171*, 1139–1151.

Cervellera, C., Wen, A., & Chen, V. C. P. (2007). Neural network and regression spline value function approximations for stochastic dynamic programming. *Computer & Operations Research, 34*, 70–90.
Chameides, W. L., Lindsay, R. W., Richardson, J., & Kiang, C. S. (1988). The role of biogenic hydrocarbons in urban photochemical smog: Atlanta as a case-study. *Science, 241*(4872), 1473–1475.
Chen, V. C. P. (1999). Application of orthogonal arrays and MARS to inventory forecasting stochastic dynamic programs. *Computational Statistics and Data Analysis, 30*(3), 317–341.
Chen, V. C. P., Ruppert, D., & Shoemaker, C. A. (1999). Applying experimental design and regression splines to high-dimensional continuous-state stochastic dynamic programming. *Operations Research, 47*, 38–53.
Chen, V. C. P., Tsui, K. L., Barton, R., & Meckesheimer, M. (2006). A review on design, modeling, and applications of computer experiments. *IIE Transactions, 38*, 273–291.
Friedman, J. H. (1991). Multivariate adaptive regression splines (with discussion). *Annals of Statistics, 19*, 1–67.
Hsu, W., Rosenberger, J. M., Sule, N. V., Sattler, M. L., & Chen, V. C. P. (2014). Mixed Integer Programming Models for Selecting Ground-Level Ozone Control Strategies. *Environmental Modeling and Assessment, 19*(6), 503–514.
Kim, S. B. (2009). Feature extraction/selection in high-dimensional spectral data. *Encyclopedia of Data Warehousing and Mining, 2*, 863–869.
Kutner, M. H., Nachtsheim, C. J., Neter, J., & Li, W. (2005). *Applied linear statistical models*. Boston: McGraw-Hill Irwin.
LeBoulluec, A. K., Ohol, N., Zeng, L., L. , Chen, V. C. P., Rosenberger, J. M., & Gatchel, R. J. (2018). Handling time-varying confounding in state transition models for dynamic optimization of adaptive interdisciplinary pain management. *IISE Transactions on Healthcare Systems Engineering, 8*(1), 83–92.
Markowitz, H. M. (1952). Portfolio Selection. *Journal of Finance, 7*(1), 77–91.
NAG Inc. (1991). *The nag FORTRAN library manual (Mark 15)*. Oxford, UK: NAG Inc.
Sacks, J., Welch, W. J., Mitchell, T. J., & Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical Science, 4*(4), 409–423.
Sobol, I. M. (1967). The distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics, 7*, 86–112.
Sule, N., Chen, V. C. P., & Sattler, M. L. (2011). A Decision-making framework for assessing control strategies for ground level ozone. *Atmospheric Environment, 45*, 4996–5004.
Tejada-Guibert, J. A., Johnson, S. A., & Stedinger, J. R. (1993). Comparison of two approaches for implementing multi-reservoir operating policies derived using dynamic programming. *Water Resources Research, 29*, 3969–3980.
Tsai, J. C. C., Chen, V. C. P., Beck, M. B., & Chen, J. (2004). Stochastic dynamic programming formulation for a wastewater treatment decision-making framework. *Annals of Operations Research, Special Issue on Applied Optimization under Uncertainty, 132*, 207–221.
U.S. Environmental Protection Agency. Accessed June 6, 2018, https://www.epa.gov/clean-air-act-overview/air-pollution-current-and-future-challenges.
U.S. EPA. (1990). *User's Guides for the Urban Airshed Model*. EPA-450/4-90-007A-E, Washington, D.C.
Yang, Z., Chen, V. C. P., Chang, M. E., Murphy, T. E., & Tsai, J. C. C. (2007). Mining and modeling for a metropolitan Atlanta ozone pollution decision-making framework. *IIE Transactions, 39*, 607–615.
Yang, Z., Chen, V. C. P., Chang, M. E., Sattler, M. L., & Wen, A. (2009). A decision-making framework for ozone pollution control. *Operations Research, 57*, 484–498.