Addressing State Space Multicollinearity for an Ozone Pollution Dynamic Control Problem

Bancha Ariyajunya

Faculty of Engineering, Burapha University, Chonburi, Thailand, bancha@buu.ac.th

Ying Chen*

Department of Management Science and Engineering, School of Management, Harbin Institute of Technology, Harbin, China, yingchen@hit.edu.cn

Victoria C. P. Chen

Department of Industrial, Manufacturing, and Systems Engineering, The University of Texas at Arlington, U.S.A., vchen@uta.edu

Seoung Bum Kim

School of Industrial Management Engineering, Korea University, Seoul, Korea, sbkim1@korea.ac.kr

Jay Rosenberger

Department of Industrial, Manufacturing, and Systems Engineering, The University of Texas at Arlington, U.S.A., jrosenbe@uta.edu

*Corresponding author. Phone number: (+86)-451-86414009, Email: yingchen@hit.edu.cn

Department of Management Science and Engineering, School of Management, Harbin Institute of Technology, Harbin, China, 150000.

Abstract

High ground-level ozone concentrations constitute a major problem in the United States. In this study, we considered the case of Atlanta, Georgia State, where the control of nitrogen oxides (NO_x) emission is necessary to reduce ozone pollution. Stochastic dynamic programming (SDP) was used to examine the existing NO_x emission control strategies with respect to time stage and location. A design and analysis of a computer experiments (DACE)-based SDP method was used to obtain a computationally tractable solution involving both metamodeling to estimate the state transitions, and value function approximation to generate a solution policy. However, the state variables of the considered ozone pollution control problem included spatially and temporally correlated observations. Hence, the state transition models constructed in the presence of multicollinearity might be unstable, with resultant effects on the SDP control policy. This necessitated the consideration of multicollinearity in the state space. Three types of state transition metamodels were developed to demonstrate the potential drawbacks of ignoring multicollinearity and the SDP policy improvement that can be achieved by addressing it in the state space.

Subject classifications*:* environment: pollution; dynamic programming/optimal control: Markov-infinite state; statistics: correlation.

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, mainly because of the occurrence of photochemical reactions and stratospheric intrusions. However, in the troposphere, especially in metropolitan areas characterized by substantial anthropogenic emissions, photochemical reactions involving nitrogen oxides (NO_x) and volatile organic compounds (VOCs) significantly increase the ground-level ozone concentration. This high ozone level and other air pollutants are harmful to both the natural ecosystem and humans (U.S.

Environmental Protection Agency 2018). There is thus the need for regulations for controlling the emission of NO_x and/or VOCs to reduce ground-level ozone concentration.

Dynamic programming (DP) (Bellman 1957), which is a mathematical programming method, has been used to solve multistage optimization problems in engineering, economics, social science, and many other fields (e.g., Bertsekas 2005; Adda and Cooper 2003; Burkhauser et al. 2004). The objective of the present study was to develop optimal strategies for controlling ground-level ozone pollution. We focused on finite-horizon stochastic DP (SDP), in which the objective is to sequentially minimize the expected cost over several time stages and under prescribed constraints. However, the ozone pollution problem is high-dimensional and the state and decision variables of it are continuous, whereas most current SDP/reinforcement learning (RL) algorithms still suffer the so-called "curse of dimensionality." Although a finite-grid method could be used to discretize a state space and then interpolate the optimal value function between grid points (Johnson et al. 1993), in a high-dimensional problem, a straightforward grid of points corresponding to a full factorial experimental design would grow exponentially with increasing number of state variables. Recognizing this issue, Chen et al. (1999) used a design of experiments (DoE) method to implement a more efficient state space discretization process. Efficient experimental designs have also been applied to continuous-state problems, such as orthogonal arrays (OAs), Latin hypercubes, and low-discrepancy sequences (Chen et al. 1999, Chen 1999, Cervellera et al. 2006, Cervellera et al. 2007, Yang et al. 2009, Fan et al. 2013). This approach is generally referred to as the design and analysis of computer experiments (DACE) (Sacks et al. 1989, Chen et al. 2006).

Yang et al. (2009) was the first to use DACE-based SDP to study emission control strategies with regard to the ozone pollution problem in Atlanta, Georgia State, United States of America. Their work was based on the DACE-based state transition modeling process developed by Yang et al. (2007). However, they ignored multicollinearity in the state space. To our knowledge, there has been no SDP or related RL study that considered multicollinearity in continuous state spaces. The state variables of the ozone pollution problem include spatially and temporally correlated observations, because the NO_x and/or VOCs emission locations are close to each other and the system dynamics late in the day depend on the

system states earlier in the day. Ariyajunya et al. (2017) were the first to examine multicollinearity issues in SDP state spaces, when they investigated the use of data mining techniques such as feature extraction and feature selection to mitigate multicollinearity in state transition metamodeling. They demonstrated the consequences of ignoring multicollinearity in the case of Atlanta, including the variance inflation factors (VIFs) reaching 56, where VIF > 10 is considered a serious problem in statistical modeling, because high multicollinearity is known to lead to unstable high variance parameter estimates (Kutner et al. 2005).

In the present study, we developed three types of state transition metamodels and numerically solved the corresponding SDP problems. The first type of state transition metamodels ignores the inherent multicollinearity and can be used to develop high-VIF metamodels. It is likely to produce poor results for SDP problems with unchecked high multicollinearity. The second type of state transition metamodels addresses multicollinearity using classical regression analysis techniques to yield low VIFs. Although it is simpler to implement this type of metamodels in SDP, it is not always possible to obtain them. The third type of state transition metamodels is used to orthogonalize the state space, guaranteeing zero multicollinearity. However, the associated Bellman equation and corresponding backward SDP solution algorithm must be modified to utilize the orthogonalized state space. All three types of state transition metamodels were applied to the Atlanta ozone pollution SDP problem and their results were compared. The major contributions of the study are as follows:

(a) Restructuring the Bellman equation and corresponding solution algorithm in terms of the latent variables described by Ariyajunya et al. (2017).
(b) Development of three types of state transition metamodels and applying them to the SDP process.
(c) Simulation of computational experiments that demonstrate the performance of the value functions obtained from three test scenarios.
(d) Verification of the three developed state transition metamodels.

In section 2, the DACE-based SDP method is reviewed and a brief background of the data-mining techniques used in this study is presented. Section 3 introduces the Atlanta ozone pollution control

problem. In section 4, the general procedures for developing the state transition function are described and tests are conducted on three cases of the state transition metamodeling of the Atlanta ozone pollution problem. Section 5 describes the implementation of the proposed state transition functions in an SDP process and application to a real Atlanta air quality scenario and 50 other random hypothetical scenarios. The accuracies of the state transition metamodels are also verified. Finally, section 6 presents the conclusions of the study and the scope of future work.

2. Background

2.1. Stochastic Dynamic Programming

Equation (1) is the formulation of continuous state SDP for T discrete time stages presented by Chen et al. (1999) and Yang et al. (2009). 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 are $\boldsymbol{\varepsilon}_t \in R^l$. The known stagewise cost functions are denoted by $c_t : R^{n+m+l} \to R^1$ and depend on the state, decision, and random variables at the beginning of stage t. The objective is to minimize the expected cost over T discrete stages, subject to certain constraints, namely, $\Gamma_t \subset R^{n+m}$ and the state transition given by $f_t(\cdot)$. The state transition function defines the transition of the state variables from the current stage ($\boldsymbol{x}_t$) to the next stage ($\boldsymbol{x}_{t+1}$).

$$\begin{aligned} &\min_{\boldsymbol{u}_1,\ldots,\boldsymbol{u}_T} E\left\{\sum_{t=1}^{T} c_t\left(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t\right)\right\} \\ &\text{s.t. } \boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t), \text{ for } t = 1,\ldots,T-1 \\ &\text{and } (\boldsymbol{x}_t, \boldsymbol{u}_t) \in \Gamma_t, \text{ for } t = 1,\ldots,T. \end{aligned} \tag{1}$$

Given the state $\boldsymbol{x}_t$ of the system at the beginning of stage t, we solve for the future value function (FVF) ($V_t(\boldsymbol{x}_t)$) using the recursion in Equation (2) (Bellman 1957, Bertsekas 2005):

$$
\begin{aligned}
& V_t(x_t) = \min_{u_t} E\left\{c_t\left(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t\right) + V_{t+1}(\boldsymbol{x}_{t+1})\right\} \\
& \text{s.t.} \quad \boldsymbol{x}_{t+1} = f_t\left(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t\right), \text{for } t = 1,..., T-1 \\
& \qquad \boldsymbol{u}_t \in \Gamma_t, \text{ for } t = 1,..., T \\
& \text{where } V_T(\boldsymbol{x}_T) = c_T(\boldsymbol{x}_T).
\end{aligned} \tag{2}
$$

2.2. DACE-based SDP

Chen et al. (1999) used experimental design to discretize the state space and statistical modeling to approximate the FVF. This DACE-based SDP solution method is described in Figure 1. From a statistical perspective, the unknown FVF is the response surface of interest, and the state variables are the predictor variables. Experimental design is a statistical method for organizing a data collection to enable the desired data analysis (Mason et al. 2003). In DACE-based SDP, the experimental design selects state points over the state space. The computer experiment is used to conduct stagewise optimization, as expressed by Equation (2), for each of the designed state points to obtain the optimized objective value on the FVF. The points on the FVF represent the response values of the computer experiment. A statistical model is then constructed to estimate the FVF based on these data, with the input (predictor) variable values obtained from the experimental design, while the output (response) variables are obtained by optimization using a computer model. Specifically, an OA of strength 3 is used as the experimental design, and the multivariate adaptive regression splines (MARS) method is used to approximate the FVFs of Chen et al. (1999).

<insert Figure 1>

2.3. Data Mining in Computer Experiments for Optimization

Ariyajunya et al. (2017) used a combination of data mining techniques such as feature selection and feature extraction to develop new state transition models of the Atlanta ozone pollution control problem considered by Yang et al. (2009). Specifically, they used feature selection techniques such as stepwise regression, regression trees (Breiman et al. 1984), and a multiple-testing procedure based on the false discovery rate (FDR) (Benjamini and Hochberg 1995) to downsize the problem by identifying the

important subset of the original features. Feature extraction techniques such as principal component analysis (PCA) and partial least squares (PLS) (Wold et al. 2001) were used to create new orthogonal features based on transformations of the original features, obtaining useful information for modeling (Kim 2009). As an illustration for the Atlanta ozone pollution problem, the original correlated state space is presented in Figure 2(a), and the orthogonalized state space in Figure 2(b). However, not every technique combination is efficient for this particular problem. Through the application of four different evaluation criteria, Ariyajunya et al. (2017) demonstrated that the combination of stepwise regression and PLS (stepwise-PLS) produced better overall results. This combination was thus used in the present study to directly reduce the dimension of the ozone pollution problem and orthogonalize the state space.

<insert Figure 2>

3. Atlanta Ozone Pollution Control Problem

The Atlanta ground-level ozone pollution problem in Yang et al. (2009) includes the ozone state variables at different monitoring stations and different time stages, which are highly correlated. As described in section 1, NO_x and VOCs are the main causes of high ground-level ozone concentration in urban and rural 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 was on only NO_x emissions in this case study. To control NO_x, 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. Yang et al. (2009) aggregated the 40 × 40 grid of the Atlanta Urban Airshed Model (UAM), which covers a 160 × 160 km region in the metropolitan area, into a 5 × 5 grid representing non-point source emissions. The region contains a total of 102 point sources. The ozone concentrations are monitored by four Photochemical Assessment Monitoring Stations (PAMS) located at Conyers, S. Dekalb, Tucker, and Yorkville.

The present study objective for the Atlanta ozone pollution case was to prevent the hourly averaged ozone concentrations from exceeding the U.S. EPA standard limit, which was 0.125 parts per million (ppm) at the time of the study (it has since been decreased; see http://www.epa.gov/air/criteria.html). To reduce the ozone concentrations, emission controls were applied to specific areas and time stages. Because the ozone concentration increases during the daytime when the sun is shining, only time stages from 4:00 AM to 7:00 PM were considered as potential time stages for reducing emissions. To apply SDP to the pollution control, five 3-h time stages were defined: 4:00 AM to just before 7:00 AM (time stage 0), 7:00 AM to just before 10:00 AM (time stage 1), 10:00 AM to just before 1:00 PM (time stage 2), 1:00 PM to just before 4:00 PM (time stage 3), and 4:00 PM to just before 7:00 PM (time stage 4). Time stage 0 was considered as the initialization stage and the SDP stages were based on time stages 1 through 4.

Regarding the use of SDP to solve this ozone pollution control problem, at the beginning of time stage *t*, the known state variables describe the status of all the factors that may affect the ozone concentrations. A series of decisions must be made in time stages 1 through 4 to minimize the total cost of achieving the EPA ozone goals. According to the SDP formulation in Equation (1), the state and decision variables of the Atlanta ozone pollution case can be defined as follows. The state variables ($\boldsymbol{x}_t$) at time stage *t* include all the historical information on the ozone concentrations and NO_x emissions at various spatial locations across the metropolitan Atlanta area. In other words, the initial set of potential state variables at time stage t includes information related to the ozone air chemistry during time stages 0 to *t*-1. The decision variables ($\boldsymbol{u}_t$) are the actions to be performed at time stage t to control the amounts of emissions at various locations and times over the course of the entire day.

According to Yang et al. (2009), the objective in each stage is established using the following criteria: (1) if the ozone levels cannot satisfy the EPA standard, then the control policy should minimize the ozone levels; (2) if the ozone levels can satisfy the EPA standard, then the expected cost of the control policy should be minimized. Instead of using strict constraints, Yang et al. (2009) used a penalty approach to prioritize satisfying the EPA standard. They divided the objective function into two parts, namely, the

emission reduction cost function and the penalty cost function. The above criteria and objective function proposed by Yang et al. (2009) were utilized in the present study (see Yang et al. (2009) for the details).

4. Development of State Transition Function

In the SDP formulation in Equation (2), the state transition function is that which describes the evolution of the system state between the current time stage and the next one. In the case of the present ozone problem, at each stage, an air quality model such as the Atlanta UAM is required to evaluate the emission action strategies to determine the resulting ozone concentrations based on the state and decision variables. The UAM may be used as a state transition function to predict the initial ozone state variables of the next stage. However, the direct application of the UAM when using the SDP for ozone concentration and state transition calculations is impractical because of the high computational cost and requirement of a very large amount of input data. There is thus the need to use a more efficient state transition model as a surrogate for the UAM in SDP. In the SDP literature, the state transition model is typically stationary and known (e.g., Chen et al. 1999, Fan et al. 2013). In the present study, SDP was used to address the more challenging case of non-stationary state transitions and to estimate the unknown system dynamics via data mining techniques. In addition to the following state transition function,

$$\boldsymbol{x}_{t+1} = f_t(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t) \quad , \qquad (3)$$

which applies to the other historical state variables at time stage t+1, the identity transitions defined by Yang et al. (2007) were also employed in this study.

Section 4.1 below describes a general procedure for developing state transition functions, and section 4.2 presents three test cases of state transition metamodels for the Atlanta ozone pollution scenario.

4.1. Procedure of State Transition Modeling

It was assumed in this study that the state transition functions were unknown. It was therefore necessary to estimate the functions using real data acquired from the system itself, or data obtained by a simulation of

the system dynamics. Even if the Atlanta UAM were available, it might be too computationally impractical to be directly employed within an optimization. In the use of the DACE-based SDP approach, this issue is further complicated by the presence of multicollinearity over the state space. The degree of multicollinearity can be measured using the VIFs obtained by regression modeling (Kutner et al. 2005). In this study, we defined a case with VIFs < 4 as low multicollinearity, a case with maximum VIF > 10 as high multicollinearity, and a case with 4 < VIFs < 10 as mid-level multicollinearity. To distinguish the effect of high multicollinearity from that of low multicollinearity, the cases with mid-level multicollinearity was not considered in this study. In theory, high VIFs indicate that the variances of the parameter estimators are inflated by high multicollinearity, resulting in undesirable models. An appropriate process should therefore be adopted in developing the state transition model. The following method was previously used by the authors (Yang et al. 2007)

<insert Figure 3>

Firstly, the initialization phase is used to identify the stages, state variables, and decision variables of the system, including the modeling space (the boundaries of the state and decision variables). The data collection phase is then used to collect data on the system dynamics and performance as it evolves through the time stages. In the mining phase, feature selection data mining techniques is used to eliminate the state and decision variables that clearly do not influence the state transitions. The modeling phase is used to construct the statistical prediction models of the future state outputs (previously expressed by Equation (3)), and the uncertainty modeled through statistical analysis is combined with the prediction models to incorporate random disturbances in the state transition. To improve the modeling accuracy, the modeling phase may involve additional data collection for the variables selected in the mining phase.

Considering that the dimensionality of the original state variables is 524, the primary goal of the above process is to reduce this dimensionality of the problem. Regarding the high multicollinearity in the state space, after the general feature selection procedure of the mining phase, further and careful variable selection may potentially yield regression models with low VIFs. An alternative approach developed in the present study involved modification of the representation of the state variables by transforming them

into an orthogonalized set ($\boldsymbol{z}_t$) using a feature extraction technique, after the feature selection in the mining phase. Thus, compared with Equation (3), the stochastic state transition function for $\boldsymbol{z}_t$ is as follows:

$$\boldsymbol{z}_{t+1} = g_t(\boldsymbol{z}_t, \boldsymbol{u}_t, \boldsymbol{\varepsilon}_t) \quad , \qquad (4)$$

where g_t is the state transition function of the orthogonalized variables $\boldsymbol{z}_t$ at time stage *t*. The decision variables are not orthogonalized because they are not part of the experimental design process. Besides, for optimization purposes, it is more practical to maintain the decision variables in their original form.

4.2. Three Test Cases of the Atlanta Ozone Pollution Problem

To examine the impact of multicollinearity on the SDP solution, three test cases of the Atlanta ozone pollution problem were considered. The differences between the cases were in the methods of addressing multicollinearity. In the high-VIF case, multicollinearity was not addressed and high-VIF metamodels were used. In the low-VIF case, multicollinearity was addressed by using carefully crafted regression models to obtain low-VIF metamodels. In the orthogonalized case, feature extraction was used to orthogonalize the state space after the feature selection procedure. The stepwise-PLS approach, which performs better than the other data mining techniques presented by Ariyajunya et al. (2017), was used for the mining phase in the orthogonalized case. In addition to the transition models used in the high- and low-VIF cases, identity transition (Yang et al. 2007) was also applied. Identity transition could not be applied to the orthogonalized case because of the orthogonalized state variables. Instead, prediction models developed with the aid of orthogonalized data were used to determine the values of the ozone variables. The high-VIF, low-VIF, and orthogonalized test cases are described in sections 4.2.1, 4.2.2, and 4.2.3 below, respectively.

4.2.1. High-VIF Test Case

The high-VIF state transition metamodels were deliberately developed to represent a worst-case outcome in the presence of high multicollinearity. The high-VIF models forced all the ozone state variables in the

model, while the stepwise regression for feature selection was only used to select additional emission variables.

<insert Table 1>

Table 1 summarizes the variables of the high-VIF ozone models. Based on the notations of Yang et al. (2009), the first two letters indicate the PAMS sites (cy = Conyers, sk = South DeKalb, tk = Tucker, and yk = Yorkville); the second two letters indicate the ozone level (M3 = maximum ozone); and the last two letters indicate the time stage (p1 = stage 1, p2 = stage 2, p3 = stage 3, and p4 = stage 4). For example, "cyM3p1" denotes the maximum ozone (M3) level at the Conyers site at time stage 1. It is necessary to consider all previous state variables in an ozone model. For example, the high-VIF metamodel for the Yorkville site at time stage 1 (ykM3p1) contains all the four ozone state variables for time stage 0 (cyM3p0, skM3p0, tkM3p0, and ykM3p0). Stepwise regression is then used to select 13 additional emission variables to give a total of 17 variables. The highest VIF value for the ykM3p1 model is 62.1101, which is statistically very high.

4.2.2. Low-VIF Test Case

The low-VIF test case was based on the work of Yang et al. (2007). However, an examination of the 16 models used in this previous study clearly revealed that three of them had very high VIFs. As noted in section 4.1, however, further careful variable selection after the general feature selection may be used to obtain low-VIF metamodels. The low-VIF metamodels were obtained in this study by adding a model correction procedure that removed some variables and refitted and re-evaluated the models until VIFs < 4. In addition, if some predictors in the models were not significant, the models were further corrected through stepwise regression to select only the statistically significant predictors using a significance level of 0.05. The results of the low-VIF metamodels are presented in Table 2.

<insert Table 2>

4.2.3. Orthogonalized Test Case

An alternative approach to addressing multicollinearity in the state space was introduced in section 4.1. The approach involves the addition of a feature extraction procedure to the mining phase. In this study, we used the stepwise-PLS technique recommended by Ariyajunya et al. (2017), in which stepwise regression is used to select the important variables and PLS is used to orthogonalize the state space in the mining phase. Following is the procedure for using this technique to obtain the state transition model in the mining phase:

1. Stepwise regression is used to select the statistically significant state and decision variables.
2. At each time stage *t*, PLS is used to orthogonalize the selected state variables while the decision variables are maintained.
3. The orthogonal state transition model for transitioning $\boldsymbol{z}_t$ to $\boldsymbol{z}_{t+1}$ is constructed for time stage *T*-1 to 1.

   (a). $\boldsymbol{z}_{t+1}$ is modeled as a function of the orthogonal state variables $\boldsymbol{z}_t$.

   (b). The residuals from 3(a) are modeled as functions of the decision variables ($\boldsymbol{u}_t$).

In addition to the orthogonalized state transition metamodels, the prediction models of the ozone variables were created from the last time stage *T* to time stage 1. The first step of this process was the use of the orthogonalized state variables ($\boldsymbol{z}_t$) to model the ozone variables ($\boldsymbol{O}_t$). This was followed by using the original decision variables to model the function of the residuals of the first step.

Appendix B in the work of Ariyajunya (2012) presents an equation for obtaining the ozone variables $\boldsymbol{O}_t$ for time stages 1–4, as well as the orthogonal state transition model. A summary of the results of the stepwise-PLS metamodels for each stage is also presented in Chapter 4 of the same work.

5. Computational Results

Table 3 summarizes the numbers of state variables and decision variables obtained by the three types of state transition metamodels described above. The results for the high-VIF metamodels in the table include the majority of both the state and decision variables. All the variables for the low-VIF metamodels are

subsets of the variables for the high-VIF and stepwise-PLS metamodels. The dimension of an SDP problem is determined by the maximum number of state variables for all the stages; the SDP dimensions when using the low-VIF, high-VIF, and stepwise-PLS metamodels are thus 23, 92, and 25, respectively, which indicate complexity of the ozone pollution problem.

<insert Table 3>

The application of SDP to the present ozone pollution control problem is described in section 5.1 below. After using the DACE-based SDP method to approximate the FVFs, simulation of a forward SDP "real-time" re-optimization process was used to determine the expected value of a real-time policy. The computational results of the optimal control policy for one real Atlanta scenario and 50 random hypothetical scenarios are presented in section 5.3. The verification of each of the utilized transition models is described in section 5.4.

5.1. Constructing FVFs of Atlanta Ozone Problem

The solution of the ozone problem using the DACE-based SDP method begins at the last stage and moves backward until all the stages have been solved, as illustrated in Figure 1. Following Yang et al. (2009), a low-discrepancy sequence developed by Sobol (1967) was used to discretize the state space, and MARS was used to approximate the FVFs of the control problem. The 2000 designed points of the Sobol sequence were generated to represent the state space. At each designed point, nonlinear programming was used to obtain an optimal solution, and a commercial optimization library (NAG E04) was then employed as the optimization module for achieving the SDP solution of the Atlanta ozone problem. The three different methods for modeling the state transition function described in section 4 were respectively implemented. Unlike in the work of Yang et al. (2009), where the MARS approximations were allowed to have negative values, the negative MARS values were truncated to zero in the present study, because a negative cost is unrealistic.

<insert Table 4>

To reduce the possibility of only a local optima being achieved, multiple starting points are usually used in the implementation of an optimization module. However, while the use of many starting points also increases the chance of approaching a global optimal cost, for computational reasons, only two starting points (the midpoint and the lower bound) and an additional ten random points between the two were 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 stages 1 and 2. The addition of ten other random points, as was done in the present SDP implementation, generally tends to reduce the cost. The details of the present application of SDP to the Atlanta ozone pollution problem are given in Table 4. Table 5 summarizes the running time and MARS approximations for the FVFs in each stage. As can be observed, the low-VIF procedure requires the longest time for the MARS approximation of the FVFs, but the shortest time for the SDP solution. The high-VIF procedure requires the shortest total time, and the stepwise-PLS procedure the longest. The FVF approximations were subsequently used to determine the optimal policies for the "real-time" forward simulation, as described in the next section.

<insert Table 5>

5.2. Forward Re-optimization for Optimal Control Policy

The forward re-optimization technique was used to simulate the real-time optimal control policy (Tejada-Guibert et al. 1993) in the present study. The simulation procedure is expressed in Figure 4. Given the initial state vector for stage 1 ($\boldsymbol{x}_1$), the process was sequentially implemented to solve the optimal control policy ($\boldsymbol{u}_t$) until all the stages were solved. For the present ozone problem, the optimal control policy for the specified initial emissions and ozone levels before time stage 1 determined the reduction of the emissions at specific locations and times over the course of the day.

<insert Figure 4>

5.3. Comparison of Results

The results of using the three different transition models for the SDP implementation in this study were evaluated by considering two test cases of the initial state vector. The first test case involved a real Atlanta scenario, with July 31, 1987 being the first day of the ozone episode. The second test case involved 50 random hypothetical initial scenarios.

5.3.1. Real Atlanta Scenario

The initial emissions and ozone levels at the beginning of the first time period were obtained from the nominal values recorded on July 31, 1987, and the re-optimization algorithm was used to determine the optimal control decisions. The Atlanta UAM was used to simulate the optimized policy to calculate the ozone level. The application of the obtained optimized policy to the UAM thus afforded the best representation of the actual ozone level. Table 7 presents the emission reductions achieved by the optimal control policies for this real scenario using different metamodels. Based on the objective function of the Atlanta pollution problem introduced in section 3, the optimal control policy obtained by the low-VIF metamodels requires the lowest daily total emission reduction of 27.66%, followed by that obtained by the stepwise-PLS metamodels (36.63%), and lastly that obtained by the high-VIF metamodels (47.03%). All three policies generally require the highest emission reduction in time stages 2 and 3 (covering 10 AM to 4 PM). The maximum ozone level trajectories obtained by the transition metamodels and the UAM are shown in Figure 5. The primary *y*-axis on the left side represents the maximum ozone level, while the secondary *y*-axis on the right side represents the percentage emission reduction. The "BASE CASE" line represents the maximum ozone level when no control action is implemented. The "UAM-LVIF," "UAM-HVIF," and "UAM-PLS" lines represent the actual ozone levels obtained by simulation using the UAM together with the optimal control policies obtained by the low-VIF, high-VIF, and Stepwise-PLS methods, respectively. The "LVIF," "HVIF," and "PLS" lines represent the ozone levels predicted by the low-VIF, high-VIF, and stepwise-PLS metamodels, separately.

It can further be seen from Figure 5 that the high-VIF metamodels are the least accurate and always overestimate the maximum ozone level. Consequently, they require greater emission reduction than

necessary, and hence a higher NO_x emission reduction cost. The low-VIF metamodels seem to have the best performance, although they slightly underestimate the ozone levels, which may result in lower emission reductions than necessary in time stages 3 and 4. The stepwise-PLS metamodels perform better than the High-VIF metamodels, but result in a small exceedance of the EPA limitation at stage 4, and this may prompt stringent emission reduction policies in time stage 3. However, because all the actual ozone levels in the low-VIF case are within the EPA limit of 0.125 ppm, the low-VIF metamodels, which require the least emission reductions, are the first recommendation for the development of control strategies for this particular case.

<insert Table 6>

<insert Figure 5>

5.3.2. Fifty Hypothetical Scenarios

The above real scenario represents only one specific situation. To demonstrate the applicability of SDP to different situations, 50 hypothetical scenarios were generated and tested. The initial emissions for these hypothetical scenarios were randomly generated based on the emission ranges of the above real scenario; values within the ranges were used as inputs to the UAM to obtain the initial ozone levels of the hypothetical scenarios. The optimal control policies for the hypothetical scenarios were obtained by the same procedure used for the real scenario. The average emission reductions required by the optimal control policies obtained by the different metamodels are presented in Figure 6 and summarized in Table 7. The emission reduction requirements for the 50 scenarios can be seen to be comparable to those for the real scenario.

<insert Figure 6>

<insert Table 7>

The high-VIF metamodels still have the worst performance for the 50 hypothetical scenarios, with the stepwise-PLS and low-VIF metamodels exhibiting comparable performances. The major difference

between the low-VIF and the stepwise-PLS metamodels occurs in time stage 4, where the maximum ozone level is underestimated by the low-VIF metamodels but overestimated by the stepwise-PLS metamodels. Nevertheless, the maximum ozone level determined by the low-VIF metamodels for time stage 4 is still within the EPA limit of 0.125 ppm, indicating that the low-VIF metamodels once again outperform the others overall. Different from the real scenario in section 5.3.1, the actual ozone levels using stepwise-PLS metamodels at stage 4 are not exceeded the EPA limitation, which denotes overall, stepwise-PLS metamodels have good performance.

5.4. Verification of Metamodels

The most desirable optimal policy for the Atlanta ozone problem is that which requires the least emission reductions to maintain the maximum ozone level within EPA standards. However, the computational results of the SDP process are affected by the accuracy of the state transition metamodels. The best optimal policy obtained by SDP alone may thus be insufficient to achieve the best overall results. There is thus the need to assess the accuracy of the metamodels. For this purpose, the optimal control policies obtained for the 50 hypothetical scenarios in section 5.3.2 were simulated using the UAM to determine the maximum ozone levels. The results were then compared with the maximum ozone levels predicted by the metamodels. The deviations between the two with respect to the monitoring station and time stage are presented in Figs. 7(a) and 7(b), respectively.

Figure 7 once again shows that the high-VIF metamodels have the worst performance, because they produce the largest average deviations relative to the UAM results with respect to both station and time stage. Compared with the performance of the high-VIF metamodels, the low-VIF and stepwise-PLS metamodels yield much better control of ozone pollution. However, at stage 1, the stepwise-PLS metamodels produce larger deviations than the other two metamodel types; at stage 4, the low-VIF metamodels underestimate the ozone levels with an average deviation of -0.00183, indicating the requirement for more actions for emission reduction in time stage 2 or 3. Overall, taking time stage and

station into consideration, the low-VIF metamodels are the most accurate, which is consistent with the results of the previous experiments.

<insert Figure 7>

6. Conclusion

In this study, we focus on developing the more practical and more efficient emission control policy for cities such as Atlanta, GA, dealing with harmful ground-level ozone concentrations. This is a major subject of interest to many modern cities and governments faced with the need to implement appropriate environmental protection regulations. However, the state variables of the ozone pollution problem such as Atlanta case, at different monitoring stations and different time stages are known to be highly correlated. In statistics, multicollinear state variables produce unstable models with high VIFs. To address this issue, we developed three types of statistical metamodels, namely, high-VIF, low-VIF, and stepwise-PLS metamodels, as surrogates of the UAM, using different approaches to consider the multicollinearity in the state space. The three types of models were used to implement SDP for one real Atlanta scenario and 50 other random hypothetical initial scenarios. The results showed that the high-VIF metamodels were less accurate than the other models and always overestimated the maximum ozone level. The optimal policies obtained by it tended to require excessive emission reductions. The SDP optimal policies based on the low-VIF metamodels tended to require the least emission reductions. Furthermore verification also revealed that the low-VIF metamodels afforded the most accurate predictions. Although the maximum ozone levels predicted by the low-VIF metamodels in time stage period 4 in Figures 5–7 represent slight underestimations on average, they would not significantly affect environmental policies because the resultant actual ozone levels are within the EPA limit of 0.125 ppm.

However, whereas the low-VIF metamodels produce better results, their metamodeling process attempts to avoid combinations of variables that are highly correlated. This is likely to yield a false impression of the excluded variables being unimportant. In contrast, the stepwise-PLS approach maintains

the important state variables, even if they are correlated. Moreover, the results of the stepwise-PLS metamodels are almost comparable to those of the low-VIF metamodels. In terms of the modeling process, low-VIF metamodeling also requires special effort, while the stepwise-PLS approach can be automated. Nevertheless, the results of the considered Atlanta ozone pollution control problem with multicollinear state spaces suggest that the low-VIF metamodels should be the first choice. An orthogonalization-type method such as the use of stepwise-PLS metamodels should be considered when accurate low-VIF metamodels cannot be constructed, or when an automated process is desired.

We would note in closing that the MARS-approximated FVFs in the present study contained non-convexity, and there is thus room for further study toward using a convex version of MARS to approximate the FVFs. The use of a monotonic ozone transition function without negative coefficients should also be explored; all the negative coefficients associated with the decision variables were truncated to zero in this study, degrading the accuracy of the developed models. In addition, the stepwise-PLS approach should be further investigated toward using it to achieve more accurate metamodels, thus avoiding the special effort required to create the low-VIF metamodels.

Acknowledgements

This research was partially supported by the National Science Foundation (ECCS-0801802).

References

Adda J, Cooper RW (2003) *Dynamic Economics* (MIT Press, Cambridge).

Ariyajunya B (2012) Adaptive dynamic programming for high-dimensional multicollinear state spaces. Ph.D. dissertation, University of Texas at Arlington.

Ariyajunya B, Chen Y, Chen VCP, Kim SB (2017) Data mining for state space orthogonalization in adaptive dynamic programming. *Expert Syst. Appl.* 76:49–58.

Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate-a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B: Methodol.* 57:289–300.

Bellman R (1957) *Dynamic Programming* (Princeton University Press, Princeton).

Bertsekas D (2005) *Dynamic Programming and Optimal Control* (Athena Scientific, Belmont, MA).

Breiman L, Friedman JH, Olshen RA, Stone CJ (1984) *Classification and Regression Trees* (Chapman & Hall New York, NY).

Burkhauser RV, Butler JS, Gumus G (2004) Dynamic programming model estimates of Social Security Disability Insurance application timing. *J. Appl. Econom.* 19:671–685.

Cervellera C, Chen VCP, Wen A (2006) Optimization of a large-scale water reservoir network by stochastic dynamic programming with efficient state space discretization. *Eur. J. Oper. Res.* 171:1139–1151.

Cervellera C, Wen A, Chen VCP (2007) Neural network and regression spline value function approximations for stochastic dynamic programming. *Comput. Oper. Res.* 34(1):70–90.

Chen VCP, Tsui K-L, Barton RR, Meckesheimer M (2006) Design, modeling, and applications of computer experiments. *IIE Trans.* 38:273–291.

Chen VCP, Ruppert D, Shoemaker CA (1999) Applying experimental design and regression splines to high-dimensional continuous-state stochastic dynamic programming. *Oper. Res.* 47:38–53.

Chen VCP (1999) Application of orthogonal arrays and MARS to inventory forecasting stochastic dynamic programs. *Comput. Stat. Data Anal.* 30(3)317–341.

Fan H-Y, Tarun PK, Chen, VCP (2013). Adaptive value function approximation for continuous-state stochastic dynamic programming. *Comput. Oper. Res.* 40:1076–1084.

Hocking RR (1976) The analysis and selection of variables in linear regression. *Biometrics* 32:1–49.

Johnson SA, Stedinger JR, Shoemaker CA, Li Y, Tejada-Guibert JA (1993) Numerical solution of continuous-state dynamic programs using linear and spline interpolation. *Oper. Res.* 41:484–500.

Kim SB (2009) Feature extraction/selection in high-dimensional spectral data. Wang J, ed. *Encyclopedia of Data Warehousing and Mining, Vol. 2* (Information Science Reference, Hershey, PA), 863–869.

Kutner MH, Nachtsheim CJ, Neter J, Li W (2005) *Applied Linear Statistical Models* (McGraw-Hill Irwin, Boston).

Mason R, Gunst RF, Hess JL (2003) Statistical Design and Analysis of Experiments, With Applications to Engineering and Science, 2nd edition (Wiley, Hoboken, NJ).

Seinfeld JH, Kyan CP (1971) Determination of optimal air pollution control strategies. *Socio-Econom. Plan. Sci.* 5:173–190.

Sobol IM (1967) The distribution of points in a cube and the approximate evaluation of integrals. *USSR Comput. Math. Math. Phys.* 7:86–112.

NAG Inc. (1991) *The NAG Fortran Library Manual (Mark 15)* (NAG Inc., Oxford, UK).

Tejada-Guibert JA, Johnson SA, Stedinger JR (1993) Comparison of two approaches for implementing multi-reservoir operating policies derived using dynamic programming. *Water Resour. Res.* 29:3969–3980.

U.S. Environmental Protection Agency. Accessed June 6, 2018, https://www.epa.gov/clean-air-act-overview/air-pollution-current-and-future-challenges,.

Wold S, Sjöström M, Eriksson L (2001) PLS-regression: a basic tool of chemometrics. *Chemom. Intell. Lab. Syst.* 58:109–130.

Yang Z, Chen VCP, Chang ME, Sattler ML, Wen A (2009) A decision-making framework for ozone pollution control. *Oper. Res.* 57:484–498.

Yang Z, Chen VCP, Chang ME, Murphy TE, Tsai JCC (2007) Mining and modeling for a metropolitan Atlanta ozone pollution decision-making framework. *IIE Trans.* 39:607–615.

Tables

Table 1. Summary of the high-VIF ozone state transition functions

Maximum Ozone Model	Number of Forced Variables	Number of Selected Variables	Number of Total Variables in Model	Model R-Square	Root MSE	Maximum VIF
cyM3p1	4	7	11	0.2682	0.0007	1.0859
skM3p1	4	11	15	0.9864	0.0006	17.8589
tkM3p1	4	5	9	0.9612	0.0012	62.0089
ykM3p1	4	13	17	0.9945	<0.0001	62.1610
cyM3p2	8	9	17	0.9937	0.0003	30.2103
skM3p2	8	10	18	0.2642	0.0056	69.7131
tkM3p2	8	10	18	0.6370	0.0027	69.4942
ykM3p2	8	20	28	0.9993	<0.0001	163.6091
cyM3p3	12	14	26	0.9846	0.0006	76.9823
skM3p3	12	28	40	0.9920	0.0010	75.9784
tkM3p3	12	17	29	0.9747	0.0010	74.8509
ykM3p3	12	15	27	0.9994	<0.0001	1366.4426
cyM3p4	16	24	40	0.9847	0.0013	95.7745
skM3p4	16	26	42	0.9930	0.0009	86.5118
tkM3p4	16	43	59	0.9891	0.0008	92.4630
ykM3p4	16	32	48	0.9994	<0.0001	374.8161

Table 2. Summary of the low-VIF ozone state transition functions

Maximum Ozone Model	Model R-Square	Root MSE	Maximum *p*-value	Maximum VIF	Number of Variables in Model
cyM3p1	0.2646	0.0007	0.0224	1.0149	7
skM3p1	0.9855	0.0007	0.0072	1.0075	7
tkM3p1	0.9607	0.0013	0.0259	1.0161	5
ykM3p1	0.9942	<0.0001	0.0027	1.0158	7
cyM3p2	0.9935	0.0003	0.0056	1.2438	7
skM3p2	0.1954	0.0058	0.0364	1.0306	7
tkM3p2	0.6080	0.0028	0.0052	1.0282	6
ykM3p2	0.9992	<0.0001	<0.0001	1.0128	6
cyM3p3	0.9808	0.0007	<0.0001	1.0179	7
skM3p3	0.9692	0.0019	<0.0001	1.0297	7
tkM3p3	0.9536	0.0014	<0.0001	1.3483	7
ykM3p3	0.9990	0.0000	<0.0001	2.3313	5
cyM3p4	0.9625	0.0019	0.0021	3.4408	7
skM3p4	0.9801	0.0014	<0.0001	1.5708	7
tkM3p4	0.9308	0.0019	<0.0001	1.5675	7
ykM3p4	0.9624	0.0001	<0.0001	1.0166	7

Table 3. Summary of the numbers of state and decision variables for the different model types

Stage	Variables	Number of Initial Variables	State Transition Model Type		
			Low-VIF	High-VIF	Stepwise-PLS
Stage 1	Number of decision variables	40	17	29	29
	Number of state variables	44	16	34	25
	Total number of variables	84	33	63	54
Stage 2	Number of decision variables	40	9	31	28
	Number of state variables	88	23	59	23
	Total number of variables	128	32	90	51
Stage 3	Number of decision variables	40	9	30	25
	Number of state variables	132	21	82	14
	Total number of variables	172	30	112	39
Stage 4	Number of decision variables	40	3	12	7
	Number of state variables	176	19	92	9
	Total number of variables	216	22	104	16

Table 4. SDP implementation details for all the runs

DoE for state spaces 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 knots	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 GB memory Cent OS 4.9 gcc version 3.4.6 20060404 (Red Hat 3. 4. 6-9)

Table 5. 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)	SDP Solution (hh:mm:ss)	Total Running Time (hh:mm:ss)
Low-VIF	Stage 1	16	17	394	0:53:31	1:07:57	2:01:28
Low-VIF	Stage 2	23	9	1853	50:09:49	0:16:44	50:26:33
Low-VIF	Stage 3	21	9	104	0:02:47	0:05:21	0:08:08
Low-VIF	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
High-VIF	Stage 1	34	29	1296	30:01:51	0:59:24	31:01:23
High-VIF	Stage 2	59	31	300	1:01:09	4:40:24	5:41:33
High-VIF	Stage 3	82	30	227	0:54:26	0:23:25	1:17:51
High-VIF	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
Stepwise-PLS	Stage 1	25	29	1354	24:27:20	29:49:16	54:16:36
Stepwise-PLS	Stage 2	23	28	964	7:57:58	27:54:13	35:52:11
Stepwise-PLS	Stage 3	14	25	215	0:07:14	4:32:07	4:39:21
Stepwise-PLS	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

Table 6. Emission reductions required by the determined optimal policies for the real Atlanta scenario

	Low-VIF		High-VIF		Stepwise-PLS	
Real Atlanta Scenario	Emission Reduction (gm-mol)	% Reduction	Emission Reduction (gm-mol)	% Reduction	Emission Reduction (gm-mol)	% Reduction
Stage 1	446,941.4	14.77%	1,531,936.0	50.63%	520.937.7	17.22%
Stage 2	1,147,042.2	44.81%	1,535,720.7	60.00%	862,915.9	33.71%
Stage 3	1,101,894.7	42.35%	1,594,932.6	61.30%	1,754,379.6	67.42%
Stage 4	422,009.0	13.68%	639,188.0	20.71%	990,565.7	32.10%
Daily Total	3,117,887.4	27.66%	5,301,777.3	47.03%	4,128,798.9	36.63%

Table 7. Average emission reductions required by the determined optimal policies for the 50 hypothetical scenarios

50 Hypothetical Scenarios	Low-VIF		High-VIF		Stepwise-PLS	
	Average Emission Reduction (gm-mol)	% Average Reduction	Average Emission Reduction (gm-mol)	% Average Reduction	Average Emission Reduction (gm-mol)	% Average Reduction
Stage1	811,481.7	26.82%	1,667,498.8	55.11%	1,001,731.1	33.11%
Stage 2	900,083.1	35.16%	1,060,932.1	41.45%	1,065,100.0	41.61%
Stage 3	997,836.3	38.35%	1,630,609.9	62.67%	1,850742.3	71.13%
Stage 4	473,427.0	15.34%	735,351.2	23.83%	813,879.3	26.37%
Daily Total	3,182,828.0	28.23%	5,094,392.0	45.19%	4,731,452.7	41.97%

Figures

1. For each stage *t*: use DoE to sample *N* points from the state space $\{\boldsymbol{x}_{jt}\}_{t=1}^{N}$.
2. In last stage *T*:
   (a) For each discretized point x_{jT}, j=1…*N*, solve
   $$V_T(x_{jT}) = \min_{u_{jT}} E\{c_T(x_{jT}, u_{jT}, \varepsilon_j)\},$$
   (b) Approximate $V_T(x_T)$ with $\hat{V}_T(x_T)$ for all $x_T \in R^n$, by applying a statistical regression approach such as MARS 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
   $$\tilde{V}_t(x_{jt}) = \min_{u_{jt}} E\{c_t(x_{jt}, u_{jt}, \varepsilon_j) + \hat{V}_{t+1}(f(x_{jt}, u_{jt}, \varepsilon_j))\}$$
   (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).

Figure 1. General algorithm for solving continuous-state SDP problems (Chen et al. 1999).

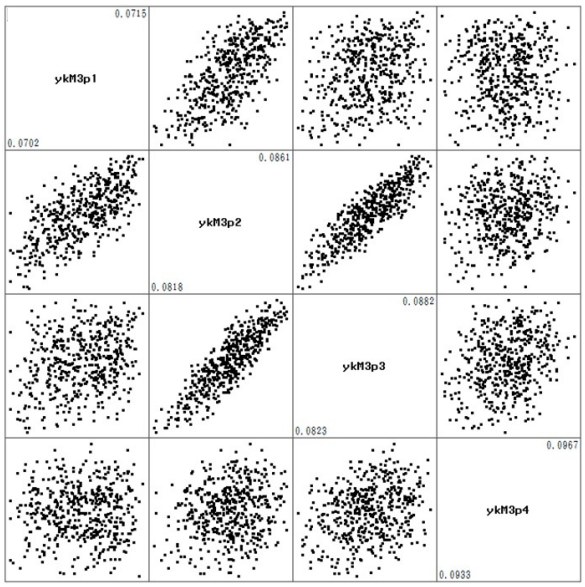

(a)

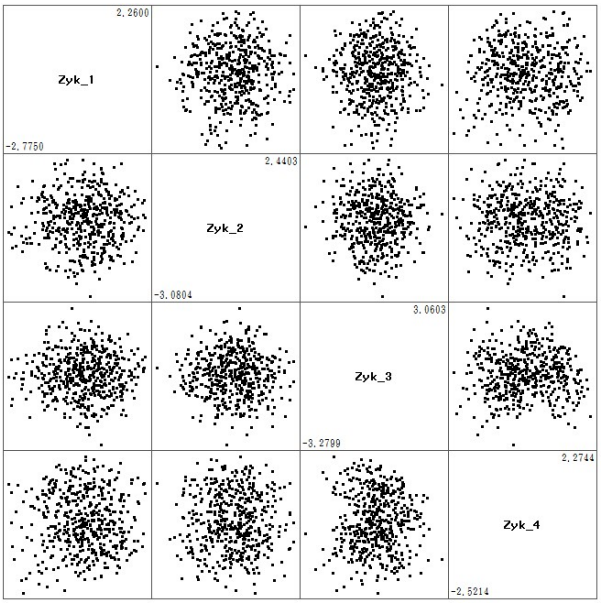

(b)

Figure 2. Relationship between state variables in the (a) original multicollinear state space and (b) orthogonalized state space (Ariyajunya et al. 2017).

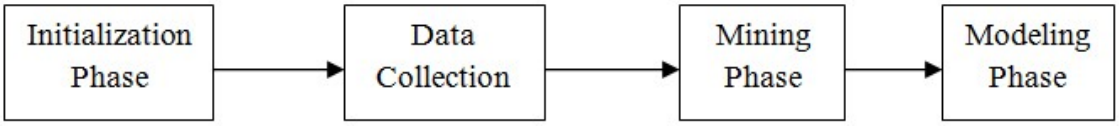

Figure 3. State transition modeling process (Yang et al. 2007).

1. For stage $t = 1,\ldots,T\text{-}1$
   a. Solve $\tilde{V}_t(x_t) = \min_{u_t} E\{c_t(x_t, u_t, \varepsilon_t) + \hat{V}_{t-1}(f_t(x_t, u_t, \varepsilon_t))\}$ for u_t;
   b. Calculate $x_{t+1} = f_t(x_t, u_t, \varepsilon_t)$
2. For stage T, solve $\tilde{V}_T(x_T) = \min_{u_t} E\{c_T(x_T, u_T, \varepsilon_T)\}$ for u_T.

Figure 4. Forward simulation procedure for solving for the optimal control policy (Yang et al. 2009).

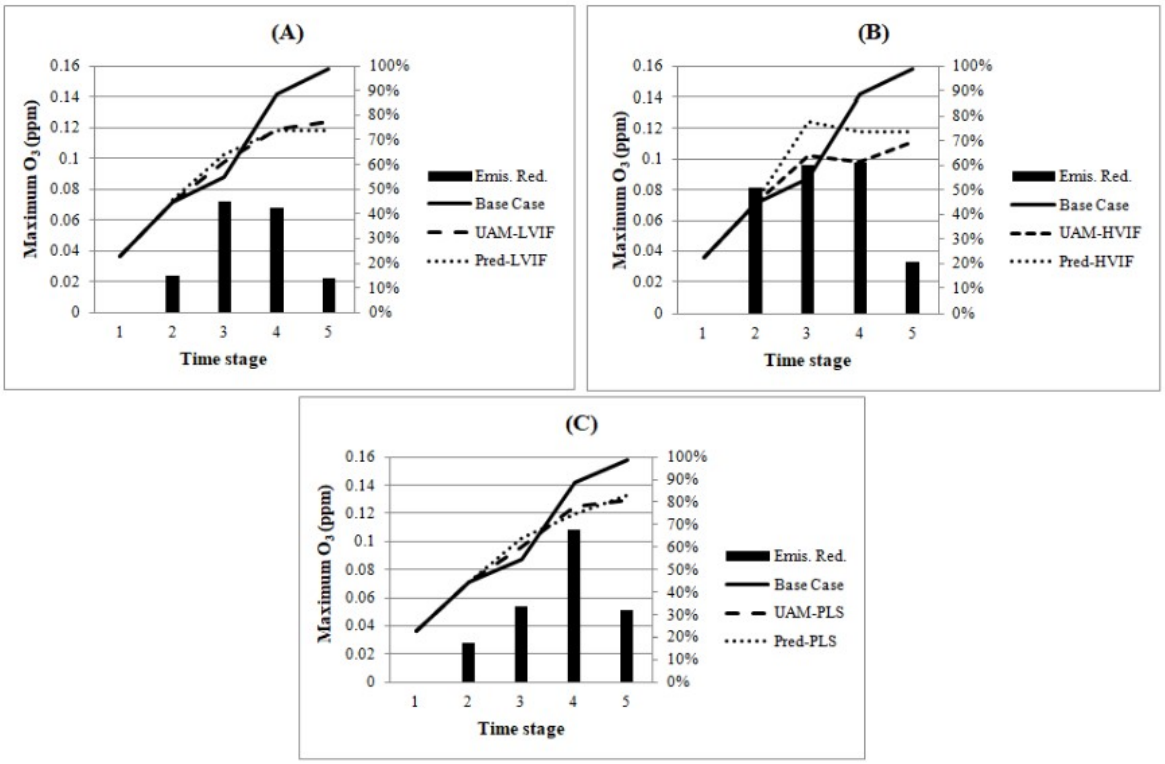

Figure 5. Maximum ozone levels and emission reductions (Emis. Red.) for the real Atlanta scenario required by the optimal control policies determined using (A) low-VIF, (B) high-VIF, and (C) stepwise-PLS metamodels.

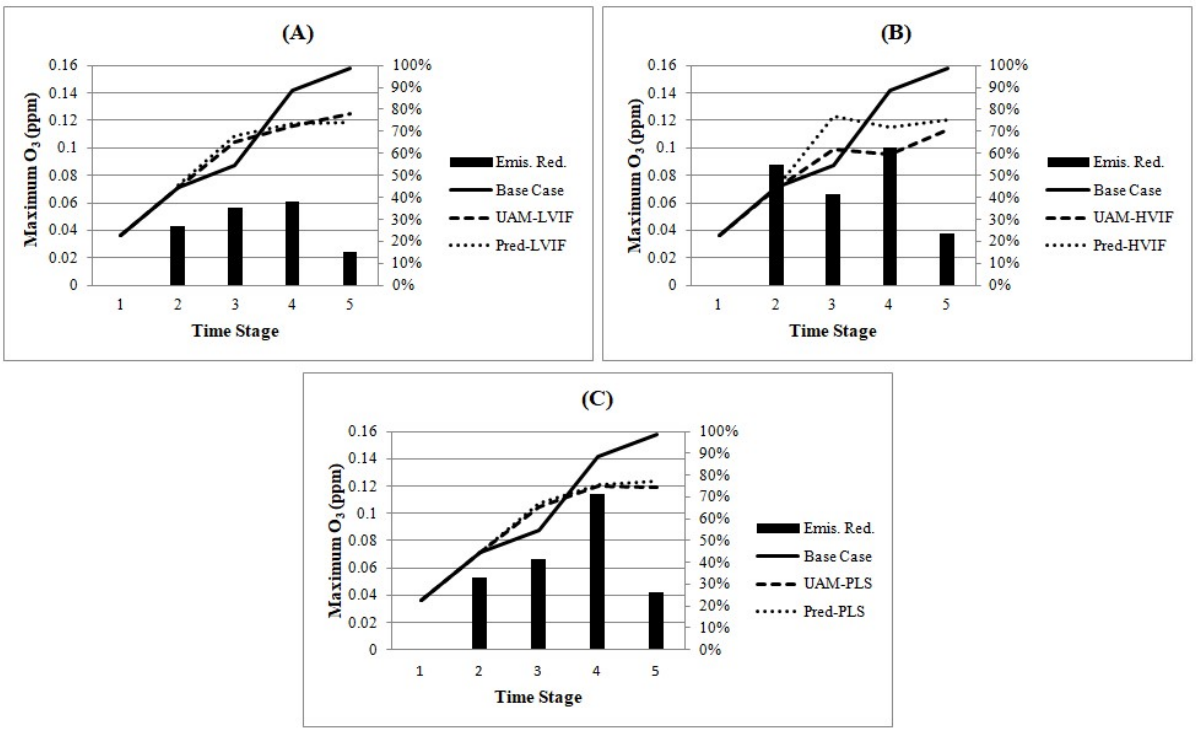

Figure 6. Maximum ozone levels and emission reductions (Emis. Red.) for the 50 hypothetical scenarios required by the optimal control policies determined using (A) low-VIF, (B) high-VIF, and (C) stepwise-PLS metamodels.

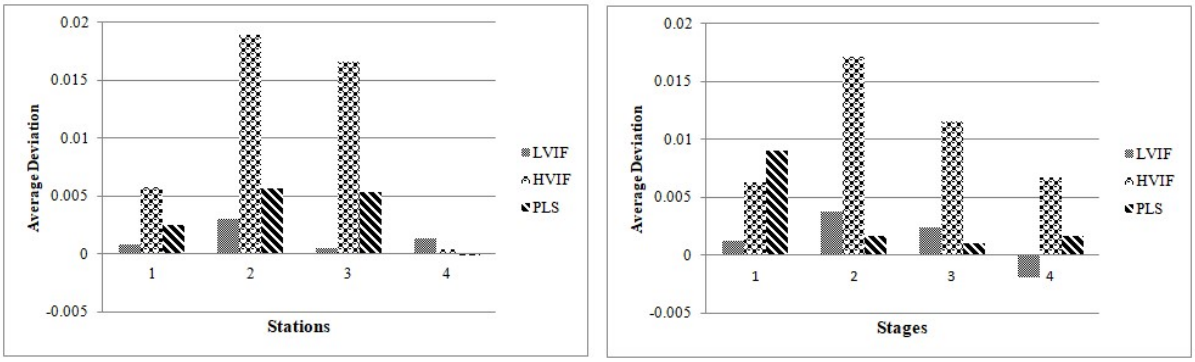

Figure 7. Average deviations between the metamodel and UAM predictions for the 50 hypothetical scenarios by (a) stations and (b) stages.