PERGAMON

Atmospheric Environment 35 (2001) 617–630

ATMOSPHERIC ENVIRONMENT

www.elsevier.com/locate/atmosenv

A review of statistical methods for the meteorological adjustment of tropospheric ozone

Mary Lou Thompson[a], Joel Reynolds[a,1], Lawrence H. Cox[b,*], Peter Guttorp[a], Paul D. Sampson[a]

[a]National Research Center for Statistics and the Environment, University of Washington, USA

[b]National Exposure Research Laboratory (MD-75), United States Environmental Protection Agency, Research Triangle Park, NC 27711, USA

Received 9 June 1999; accepted 4 May 2000

Abstract

A variety of statistical methods for meteorological adjustment of ozone have been proposed in the literature over the last decade for purposes of forecasting, estimating ozone time trends, or investigating underlying mechanisms from an empirical perspective. The methods can be broadly classified into regression, extreme value, and space–time methods. We present a critical review of these methods, beginning with a summary of what meteorological and ozone monitoring data have been considered and how they have been used for statistical analysis. We give particular attention to the question of trend estimation, and compare selected methods in an application to ozone time series from the Chicago area. We conclude that a number of approaches make useful contributions to the field, but that no one method is most appropriate for all purposes and all meteorological scenarios. Methodological issues such as the need for regional-scale analysis, the nonlinear dependence of ozone on meteorology, and extreme value analysis for trends are addressed. A comprehensive and reliable methodology for space–time extreme value analysis is attractive but lacking. Published by Elsevier Science Ltd.

Keywords: Regression; Extreme values; Time series; Spatial statistics; Environmetrics

1. Introduction

The meteorological adjustment of tropospheric ozone can be achieved by statistical modeling of the association between ozone concentration and meteorological variables. The last decade has seen a growing diversity of statistical literature on the subject with the application of a wide range of statistical methodologies, the use of widely differing data, and with adjustment being considered for different policy objectives. This review provides a summary and critical evaluation of this literature. We categorize models under three broad statistical approaches: regression-based modeling, extreme value approaches, and space–time models. Using data from the Chicago area, we compare those methods that we regard as having the most merit to derive meteorologically adjusted ozone and investigate time trend.

The main objectives for meteorological adjustment of surface ozone measurements include: (a) obtaining air quality forecasts, (b) investigating and estimating ozone time trends, and (c) increasing scientific understanding of the underlying mechanisms. The objective can influence the choice of both appropriate statistical methods and relevant data. Regression-based and extreme value methods are aimed primarily at forecasting or trend estimation, and to a lesser degree at elucidating underlying mechanisms. The early literature on meteorological adjustment for trend estimation is summarized in NRC (1991, Chapter 2). Space–time modeling has so far received little attention in the literature but could, in principle, address all three adjustment goals, with the

* Corresponding author.

E-mail address: cox.larry@epamail.epa.gov (L.H. Cox).

[1] Currently at Alaska Department of Fish and Game, Alaska, USA.

1352-2310/00/$ - see front matter. Published by Elsevier Science Ltd.
PII: S1352-2310(00)00261-2

disadvantage of increased complexity of modeling and data collection.

Forecasting extreme ozone events in order to provide public health warnings may focus the analysis on investigating those observations exceeding a threshold and their association with readily predicted meteorology. In contrast, assessing time trends may involve modeling all available daily surface ozone observations and their relationship to any relevant and available surface or upper air meteorological covariates. Estimation of a trend will most frequently be undertaken with a view to assessing the effect of changes in emissions. Models developed to provide insight into the conditions conducive to ozone production might demand measurement of precursors as well as synoptic-scale meteorology. This has largely been the focus of photochemical modeling research, while the statistical literature has focused primarily on objectives (a) and (b) above. There is room for greater contributions of statistical methods in combination with photochemical process models, but statistically valid assessments of forecasts and trends still can be made using empirically based statistical models that recognize the fundamental chemical and physical processes underlying tropospheric ozone, although not attempting detailed representation of these processes.

Another context for assessing the relationship between ozone and meteorology, but which we do not specifically consider, is the analysis of ozone and meteorology jointly in the determination of potential ozone-related health effects (Stieb et al., 1996; Kuenzli et al., 1997). The statistical issues in the analysis of health effects are different from those considered here because ozone is a predictor rather than an outcome in the modeling and the impact of meteorology would be considered in the context of epidemiological confounding, such as temperature affecting both ozone production and hospital admissions.

In Section 2, we summarize the types of data that have been used in the literature on the statistical adjustment of ozone for meteorology. Section 3 contains a description and critical review of statistical methodologies to be found in this literature. In Section 4, we compare the application of those methodologies that we have evaluated positively to data from the Chicago area and in Section 5 we summarize our evaluation and recommendations.

2. Data considerations

Even within the context of a particular modeling objective, the data used in the literature vary widely both in terms of the variables considered (ozone summaries chosen, surface and/or upper air meteorological variables included) and in terms of observation scales in space (single monitor, network of monitors) and time (hourly, daily, seasonal, annual). Different sets of meteorological variables are used depending on local and synoptic meteorology and data availability. We briefly summarize the temporal and spatial scales of ozone observations commonly used, as well as the meteorological variables investigated and eventually incorporated into the analyses.

2.1. Ozone measures

2.1.1. Time scales

The time scales for summary measures of ozone range from 5 min to daily summaries, most commonly daily 1 h maximum levels (Table 1). Daily maximum 1 h average concentration levels are the focus of the current US National Ambient Air Quality Standards (NAAQS) (EPA, 1998a) because of expected health effects of exceedances. Unfortunately, health effects are not sufficiently well understood, so the choice of ozone summary measures for analysis cannot be fully resolved from this perspective. Daily summaries are also the basis of most recent statistical assessments of trend, which we believe to be appropriate in view of the time scales of meteorological impact on ozone, which is on the order of days. Finer (hourly) scales are most relevant for process modeling, short-term predictions, and photochemical model evaluation. Changes in emissions, such as would result from air pollution control strategies, can result in different atmospheric chemistry, and hence different relationships between ozone and meteorology. If such changes occur on time scales much shorter or much larger than those used for summary measures, serious difficulties in meteorological adjustment of ozone would result. In this evaluation, we assume that no major structural changes occur on the time scales considered.

Ozone data have most often been modeled in terms of original concentration scales (parts per million), although in some cases transformations, such as square root or logarithm, are used. The appropriate transformation will depend on the temporal and spatial scale of the particular analysis, with greater spatial and temporal averaging domains generally resulting in more nearly Gaussian distributions. The different investigations of Chicago ozone time series reported below (Section 4) include analyses of original, square root, and log-transformed daily ozone measures. The scale of the ozone response has clear implications for the way that meteorological measures are most appropriately combined in explanatory regression models. These issues do not appear to have been addressed adequately in the literature.

2.1.2. Spatial scales

There is also considerable variety in the spatial scales represented in analyses of ozone measurements (Table 1). While several studies use data from a single location, many consider a regional network of ozone monitors. Analyzing data from a regional network requires a decision on how to address the spatial ozone field. Among

Table 1
Summary of ozone measures

Time scale	
5 min	Abdul-Wahab et al., Poissant et al., van Ooy and Carroll
Daily 30 min max	Spichtinger et al.
Hourly average	Carroll et al., Cox and Chu, Fiore et al., Galbally et al., Katsoulis
Daily average	Feister and Balzer
Daily 1 h max	Bloomfield et al., Burrows et al., Davis et al., Eder et al., Flaum et al., Gao et al., Huang and Smith, Milanchus et al., Joe et al., Korsog and Wolff, McKendry, Milanchus et al., Niu, Porter et al., Pryor et al., Rao et al., Reynolds et al., Smith and Huang, Stoeckenius and Hudischewskyj, Xu et al.
Daily 8 h max	Porter et al.
Length of record	
Single year	Abdul-Wahab et al., Poissant et al., van Ooy and Carroll
Multiple years	Bloomfield et al., Burrows et al., Carroll et al., Cox and Chu, Davis et al., Eder et al., Feister and Balzer, Fiore et al., Flaum et al., Galbally et al., Gao et al., Huang and Smith, Joe et al., Katsoulis, Korsog and Wolff, McKendry, Niu, Porter et al., Pryor et al., Rao et al., Reynolds et al., Smith and Huang, Spichtinger et al., Stoeckenius and Hudischewskyj, Xu et al.
Sites	
Single site	Abdul-Wahab et al., McKendry, Poissant et al., Pryor et al.
Multiple sites, modeled separately	Burrows et al., Cox and Chu, Feister and Balzer, Fiore et al., Flaum et al., Galbally et al., Joe et al., Katsoulis, Korsog and Wolff, Milanchus et al., Rao et al., Smith and Huang, Spichtinger et al., van Ooy and Carroll, Xu et al.
Multiple sites, univariate summary	Bloomfield et al., Davis et al., Eder et al., Gao et al., Huang and Smith, Niu, Reynolds et al., Smith and Huang, Stoeckenius and Hudischewskyj
Multiple sites, modeled jointly	Carroll et al., Porter et al.
Transformation	
None	Burrows et al., Davis et al., Feister and Balzer, Fiore et al., Gao et al., Joe et al., Katsoulis, McKendry, Niu, Poissant et al., Pryor et al., Reynolds et al. (1998), Spichtinger et al., van Ooy and Carroll
Logarithm	Abdul-Wahab et al., Bloomfield et al., Flaum et al., Korsog and Wolff, Porter et al., Rao et al., Xu et al.
Square root	Carroll et al., Reynolds et al. (1999)
Distribution	Cox and Chu, Galbally et al.

the ozone network analyses reviewed here, most model each monitor independently, some model a derived univariate network summary, and one models the multivariate spatial field. The design of an ozone monitoring network is an important and often overlooked issue, in which NAAQS compliance-based networks are designed to find the maximum of a random field and subsequent analysis ought to take this into account. None of the analyses to date do this.

Separate modeling of the association between each ozone monitor and local meteorology is the simplest and most common approach to analysis. However, this approach ignores any information on regional dynamics of meteorology and ozone available in the analysis of a network of ozone monitoring sites, and may therefore result in a statistically less powerful and possibly misleading analysis for purposes such as the assessment of regional trend. Analysis of the full network response through space–time modeling resides at the other end of the spectrum, being more complex both theoretically and computationally, yet potentially more flexible in its ability to capture regional associations between ozone and meteorology. Modeling a univariate summary of the ozone network retains the simplicity and wealth of tools available for univariate responses, but requires choosing the appropriate summary.

Two broad solutions to the question of choice of spatial summary appear in the literature. The first involves selecting a simple network summary. Examples include average over the network of the site-specific daily 1 h maximum (Stoeckenius and Hudischewskyj, 1990; Eder et al., 1994) and the network maximum of the daily 1 h maxima, (Niu, 1996; Smith and Huang, 1993; Stoeckenius and Hudischewskyj, 1990). Alternatively, network summaries have been derived using multivariate dimension reduction techniques. For example, principal components analysis of a network of ozone monitors provided a univariate network summary for the modeling of Bloomfield et al. (1993a, b), and the subsequent reanalysis by Gao et al. (1996). Bloomfield et al.

(1996) (Davis et al., 1998a), and subsequently Davis et al. (1998b, c), used a median polish algorithm applied to a two-way representation of ozone data in terms of site and day to compute a "daily network typical value".

Another approach leading to one or more univariate summaries uses the singular value decomposition (SVD) of the cross-covariance matrix between measures from a network of ozone monitors and a network of meteorological monitors. The resulting ozone and meteorology summaries are those that are most highly associated in terms of explained covariance (Reynolds et al., 1998). In application of these methods to data from the Chicago region, the ozone network summaries were approximately proportional to network averages, allowing simple interpretations of the subsequent analysis (see Section 4).

The modeling of ozone network data requires a balance between interpretation, simplicity of approach, and incorporation of regional-scale response information for increased statistical power. In our view a multivariate statistical technique should be used for the explicit definition of a regional-scale response, rather than a relatively arbitrary selection of a network mean or median that does not take into account the various purposes and empirical variability of individual monitors. A simple summary such as a network average might be appropriate if consideration of the relationships within the network (by, say, PCA) indicates exchangeability of monitors. We suggest consideration of the SVD methodology for its explicit construction of summaries that take into account the association between ozone and meteorology.

2.2. Meteorology and meteorological variables

2.2.1. The effects of meteorology on tropospheric ozone

The fundamental production vehicle for ozone is photochemistry. In addition, reactions with volatile organic compounds (VOCs) are important in the formation of ozone. Periods of high ozone concentrations are observed with slow-moving, high-pressure weather systems that result in sunshine, high temperatures, and stagnant air (NRC, 1991, Chapter 4). These systems typically result in a stable trophospheric layer less conducive to convective mixing with a temperature inversion that helps contain precursor pollutants (NO_x and VOCs) in the troposphere. Winds associated with high-pressure systems are typically light, thus increasing the chance that pollutants will accumulate in the atmospheric boundary layer. And, finally, warm cloudless conditions associated with these systems are favorable to photochemical production of ozone.

A detailed specification of the influence on tropospheric ozone of the primary measures of meteorology, including spatio-temporal characterizations of wind, pressure, clouds, sunshine, etc., is exceedingly complex and not the intent of the statistical models discussed here. Nonetheless, understanding of these factors is important in considering which meteorological variables to incorporate in statistical models and how to consider them.

2.2.2. Meteorological variables for statistical modeling

The choice of relevant meteorological variables depends on the purpose of the analysis, regional differences in meteorology and emission patterns, and data availability. Forecasting restricts the analysis to readily predicted (or lagged) meteorology and, in the context of forecasting extreme events, focuses on just those meteorological variables indicative of such events. Alternatively, assessing trends and other long-term developments necessitates study of ozone production under a variety of atmospheric conditions. Meteorological variables included in the reviewed literature are summarized in Table 2. For each application, we list only those variables that are explicitly mentioned in the corresponding paper. Temperature, wind speed and direction are included in most models.

Cox and Chu (1993, 1996) considered some 100 meteorological variables and found maximum surface temperature, wind speed, relative humidity, mixing height, and opaque cloud cover, as well as wind speed by temperature interaction, to be significant meteorological predictors over most major metropolitan areas of the United States. However, regional differences in meteorology and emissions patterns make more precise statements of dominant variables tenuous. At perhaps the other extreme, Pryor et al. (1995), consider only pressure and geopotential height, and use these as a basis for defining atmospheric circulation modes by principal components analysis, arguing that "a wide range of meteorological variables are implicitly contained within each mode".

Upper air measurements have been included in a number of studies. They have proved useful for predicting the conditions underlying extreme events (Burrows et al., 1995; Pryor et al., 1995) as such events have been shown to be associated with specific synoptic weather patterns (e.g., Eder et al., 1994; McKendry, 1994; Davis et al., 1998b, c).

Another issue in the analysis of ozone and meteorology concerns the relative locations of the ozone and meteorology monitors. While airports often provide high-quality meteorological surface and upper air observations, measurements from meteorological monitors not co-located with air-quality monitors are frequently used, as well as analyzed regional meteorological data resulting from a variety of sources including radar, satellites and balloon measurements. Analyses that we have carried out in Washington State (Reynolds et al., 1998) and Chicago (Reynolds et al., 1999) indicate that airport meteorological data may provide the most reliable meteorological measures for associations with ozone over

Table 2
Available and included meteorology[a]

Meteorological variable	Chicago analyses						Regression									
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Surface temperature	I	I	I	I	I	I	I	I	I	I	I	I	I	I		I
Wind speed	I	I	I	A	I	I	I	I	A		I		A	I		E
Wind direction	I	I	I	A	I	I	A	I	A				A	I		E
Humidity	I	I	I	I	I	I	I	A	A			I	A	I		
Pressure	A	A	I			A	A		A			I	A		I	
Radiation	I	A		IE	I	A	I	A	I			A	A			
Upper temperature	A	A				E			A				I			I
Upper wind speed	I	I				E							A			
Upper wind direction	A	A				E							A			
Geopotential height	A	A			I	E			A				A		I	I

	Regression			CART		Extreme value				F	S	Other		
Meteorological variable	17	18	19	20	21	22	23	24	25	26	27	28	29	30
Surface temperature	I	I	I	I	I	I	I	I	I	I	I	I		I
Wind speed	I	A	I	I	I	I	I		I		A	I		I
Wind direction		A	I	A	I		I		A		A	I		I
Humidity	I	I	I	I		I	I					I		I
Pressure		A	A	I								I	I	
Radiation	I	I	I	A	I	I	I			I	I			A
Upper temperature			A	I	A	I						I		
Upper wind speed			I	A								I		
Upper wind direction			I	I							I	I		
Geopotential height		A		I	I	I	A					I	I	

[a]I – incorporated into final model; A – available but not incorporated into final model; E – meteorology estimated from deterministic models. Analyses: 1 – Bloomfield et al. (1996), 2 – Gao et al. (1996), 3 – Huang and Smith (1999), 4 – Milanchus et al. (1998), 5 – Niu (1996), 6 – Reynolds et al. (1999), 7 – Smith and Huang (1993), 8 – Abdul-Wahab et al. (1996), 9 – Feister and Balzer (1991), 10 – Fiore et al. (1998), 11 – Galbally et al. (1986), 12 – Katsoulis (1996), 13 – Korsog and Wolff (1991), 14 – Poissant et al. (1996), 15 – Pryor et al. (1995), 16 – Reynolds et al. (1998), 17 – Spichtinger et al. (1996), 18 – Xu et al. (1996), 19 – Davis et al. (1998b, c), 20 – Burrows et al. (1995), 21 – Stoeckenius and Hudischewskyj (1990), 22 – Cox and Chu (1993), 23 – Cox and Chu (1996), 24 – Joe et al. (1996), 25 – Smith and Shively (1995), 26 – Rao et al. (1997), 27 – Carroll et al. (1997), 28 – Eder et al. (1994), 29 – McKendry (1994), 30 – van Ooy and Carroll (1995).

broad regions. While more work needs to be done to determine the implications of using different summary spatial scales for meteorological variables, there remains the practical statistical question of how best to use those data that are currently available.

3. Methods

The statistical approaches to ozone adjustment for meteorology may be divided into three broad areas, with considerable variety within each: regression-based modeling, extreme value approaches, and space–time models. A common focus is the reduction of unexplained variability in ozone through meteorological adjustment. Rao et al. (1997) argue that process changes due to policy or climate changes may be very small and difficult to detect unless they are separated from weather and seasonality. In Section 4, we will compare those approaches which we believe have merit in an application to data from the Chicago area.

3.1. Regression-based methods

The majority of approaches to meteorological adjustment of ozone are in some sense regression-based, with widely varying degrees of complexity. These approaches model average behavior and their structure can be considered in three broad categories: linear regression, regression trees, and nonlinear regression. Within each category there are further distinctions such as the method of introducing meteorological variables (directly or via dimension reduction) and the incorporation of temporal and spatial dependence. The class of regression models dictates the availability of software for model fitting and diagnostics, the ability to model complex interactions among the meteorological variables, and the familiarity and ease of interpretation for the policy-maker.

3.1.1. Models based on linear regression

The most familiar of the methodologies employed in the literature is linear regression. All linear regression models are open to the criticism that the underlying chemical and physical processes are unlikely to be linear and additive. The assumed simple linear and additive associations between the variables are unlikely to aid understanding of processes driving the relationships, in that they are inadequate to capture interactions and nonlinearities in the ozone response. In its simplest form multiple linear regression modeling is used to link ozone measurements to contemporaneous meteorological measures (e.g., Feister and Balzer, 1991; Korsog and Wolff, 1991; Abdul-Wahab et al., 1996; Katsoulis, 1996; Fiore et al., 1998). Next in complexity is time series modeling which incorporates lagged relationships and a correlation structure (usually a simple AR(1)) for the residuals (e.g. Galbally et al., 1986). However, incorporation of lagged ozone measurements is not suited to the estimation of time trends.

Both the ozone and meteorological measures are typically highly multivariate and this poses serious challenges for scientifically meaningful linear regression analyses. The reduction of a multivariate network of ozone measurements to a single summary was addressed above (Section 2.1.2). The most common methods considered for reducing the dimensionality of the ozone predictors include factor analysis (Spichtinger et al., 1996) and varimax principal component analysis (Pryor et al., 1995; Poissant et al., 1996). The singular-value decomposition (SVD) or canonical covariance analysis (Reynolds et al., 1998) is the one method proposed that jointly determines linear combinations of the predictors (meteorology) and responses (ozone) for simplified interpretation and analysis. These approaches enable handling large numbers of potentially multicollinear covariates and reduce the complexity of possible interactions to be considered. However, all of them, if applied naively, may overlook nonlinear associations. The methods should be most useful and meaningful when motivated by a scientific model that proposes underlying meteorological factors as explanatory variables; this is the underlying rationale for the SVD analysis. But even in this case, as mentioned in Section 2.1.1 above, care must be taken in considering whether to examine nonlinear transformations of meteorological variables before combining them into a composite (as in the transformations of a network of temperature measurements by Reynolds et al., 1998) or to consider nonlinear transformations of simple (linear) additive composites as in projection pursuit regression methods (which have not been applied in this setting, to our knowledge). The same concern arises in forming ozone network summaries (see Section 2.1.2).

Finally, there are approaches that involve some temporal filtering of both ozone and meteorological variables before the resulting filtered variables are modeled by linear regression (Flaum et al., 1996; Porter et al., 1996; Eskridge et al., 1997; Rao et al., 1997; Milanchus et al., 1998). Details are provided in Section 3.3.

3.1.2. Tree-based and stratified models

The possibility that the association between ozone and meteorology may be different in different regimes has led to methods of analysis for identification of different meteorological regimes and subsequent analysis stratified by regime. These stratifications may bypass some of the complexities of formal modeling of the separation of the regimes yet provide a structured approach to non-additivities. In addition, the identification of the regimes (strata) may aid our scientific understanding of the underlying processes.

Both classification and regression trees (CART; Breiman et al., 1984) and cluster analysis have been used to define meteorological regimes (Huang and Smith, 1999; Davis et al., 1998b, c). Huang and Smith (1999) fitted separate linear trend models to the observations in each terminal node of a tree, investigating ozone trends within each cluster of meteorological conditions identified by CART. Meteorological clusters conducive to higher ozone tended to display stronger (downward) temporal trends. Their analysis provided insight into the dependence of time trends on different meteorological conditions. The adjustment for meteorology considered here was limited to identification of the clusters. A related approach by Davis et al. (1998b, c) took this a step further by including formal modeling of ozone-meteorology associations within a cluster, but computed the clusters using meteorological data alone without consideration of ozone concentrations.

Davis et al. (1998b, c) analyzed 11 ozone monitoring sites for Houston, Texas, for the period 1981–1992 using the same form of ozone network summary as that of Bloomfield et al. (1996) for Chicago. Their meteorological variables were derived from a single site. First, the data matrix of meteorological observations (on seven variables recorded every 3 h, for a total of 56 variables) was subjected to a singular-value decomposition, from which six component scores were computed to summarize daily meteorology. Cluster analysis applied to these component scores separated the records into seven clusters showing significantly different ozone concentrations. The relationship between ozone and meteorological variables was then modeled separately in each cluster using stepwise selection and generalized additive models

(Hastie and Tibshirani, 1990). Seasonal effects were essentially accommodated in the cluster definitions. This approach has the advantage of allowing different associations between ozone and meteorology within different meteorological regimes, and thus different, possibly more sensitive, estimates of ozone trends within regimes.

CART's capacity to identify the meteorological conditions most commonly associated with high ozone events has also led to good performance as a forecasting tool (Burrows et al., 1995; Huang and Smith, 1999). Stoeckenius and Hudischewskyj (1990), in an alternative application, use CART to implicitly capture the effect of changing precursor levels through time. Unfortunately, the final analysis of temporal trends in their derived estimates of expected number of ozone exceedances does not account for the varying uncertainty in the estimates.

3.1.3. Models based on nonlinear regression

Bloomfield et al. (1996), (Davis et al., 1998a), argue correctly that statistical linear models "have difficulty capturing the complex relationships between the meteorological variables and ozone". They develop a parametric nonlinear model for data from 45 monitoring stations in the Chicago region over 1981–1991 represented in the AIRS database (EPA, 1998b). These authors model the daily median (across sites) of the daily site-specific maximum 1 h average ozone values, using nonlinear least squares. Parametric forms for trend and the relationship between contemporaneous and lagged ozone, surface temperature, relative humidity, surface wind speed and 700 hPa wind speed are identified in stages by exploratory graphical displays and non-parametric modeling. Seasonal terms are modeled via a short Fourier series. In estimating the standard errors of the fitted coefficients, the authors acknowledge the existence of serial autocorrelation in the model residuals and make appropriate adjustments, using the methods of Gallant (1987).

Bloomfield et al. (1996) provide a useful and thorough approach. Their choice of the network median as regional ozone summary was guided by a principal components analysis that provided essentially an average as the first component. The median was chosen rather than the average "in order to be less sensitive to individual extreme ozone concentrations". The exploratory multivariate graphical displays used have the advantage of revealing interactions in the relationships between the various meteorological variables in their association with ozone. For instance, they found that the scale of the ozone association with temperature depends on relative humidity. The graphical displays can motivate various nonlinear functional forms for the ozone model. Their analysis of a network summary based on untransformed ozone concentrations undoubtedly influenced the multiplicative nonlinear form of the final model. Model residuals had a distinctly long upper tail, in addition to serial correlation and heteroscedasticity, all of which led to a jackknife approach to compute meaningful standard errors for the parameter estimates. Modeling approaches that consider transformations computed to provide more nearly Gaussian ozone summaries will result in notably different model forms.

In summary, regression models are useful for modeling average behavior. Complex, stratified, nonlinear regression models are needed to approximate the true underlying mechanisms. Even then, if the interest is in extreme values, these models may not be sufficient. Smith and Huang (1993) implemented the model developed by Bloomfield et al. but with particular interest in the extreme behavior and found that the fit in the tail of the distribution is inadequate. Thus, while the model may describe average behavior well, it is not a good fit for applications assessing acute health effects due to pollutant extremes.

3.2. Extreme value approaches

The inherent averaging in regression analyses often makes fitted models poor predictors of extreme values (NRC, 1991, p. 61). An alternative approach, particularly useful in the context of modeling threshold exceedances, is to use extreme value theory (e.g. Gumbel, 1958; Leadbetter et al., 1983).

Cox and Chu (1993, 1996) applied a Weibull hazard model to an ozone value Y, namely

$$P(Y > y) = \exp(-(y/\sigma)\lambda)$$

in which the scale parameter σ is allowed to depend on meteorological variables $M_1, \ldots, M_n$ (about 100 candidate meteorological measures were used) in the form

$$\sigma = \exp(\Sigma\beta_j M_j + \zeta T),$$

where T is the year. An advantage with this approach is that it is not confined to a single threshold, as is the case with that of Smith and Huang (1993) described below. The model was fitted to daily maximum hourly average ozone values for 43 Metropolitan Statistical Areas (MSAs) during 1981–1993 using the average from all available station hours. The parameters were estimated using maximum likelihood, assuming independence between days. In order to take account of the daily dependence not included in the parameter estimation, standard errors were computed using a block bootstrap approach with 3-day blocks, yielding standard errors 30–40% larger than those suggested by standard likelihood asymptotics.

Niu (1996) extended the Cox and Chu approach by explicit modeling of the dependence in the errors using a heteroscedastic ARMA model with innovations scale parameter σ having the structure given in the previous paragraph. In addition, they used a nonlinear additive model to express the time-varying mean dependence on atmospheric variables. This model, while using a large number of parameters, did improve on the Cox and Chu results for prediction of percentiles.

Joe et al. (1996), as part of the Lower Fraser Valley (British Columbia, Canada) ozone monitoring project, looked for trends in high quantiles of ozone at two sites by examining days with maximum daily temperature above and below 20°C. The character of the trends was similar in both temperature ranges, although the values were, of course, larger in the higher temperature range. In order to test for trend, Joe et al. used isotonic regression (Barlow et al., 1972), i.e., testing the null hypothesis of equal means (or population quantiles) against a hypothesis of monotonically increasing (or decreasing) levels across years. For estimation of population quantiles (here the 80th and 85th) they assumed lognormality and applied standard parametric likelihood methods. Dependence was modeled using an AR(1)-process.

Smith and Huang (1993) developed a logit model for threshold exceedance, in which the probability p_i of exceeding the threshold on day i is given by

$$\log(p_i/1 - p_i) = \Sigma\beta_j M_{ij}.$$

The data consist of a sequence of ones and zeros, corresponding to exceedance or non-exceedance, respectively, and the likelihood was computed assuming independence between days. In fitting the model to one of the most extreme Chicago sites, using the same data as Bloomfield et al. (1993a, 1996), they found significant coefficients for year, seasonal effect, temperature, specific humidity, wind speed, and temperature–wind speed interaction. The fit was improved by introducing an indicator for whether the previous day was an exceedance, in effect turning the model into a first-order Markov chain. Further lags were not needed.

Smith and Huang also modeled the amount of excess over the threshold, but using a generalization of the generalized Pareto distribution (GPD) of Pickands (1975), which occurs as the limiting distribution for high-level exceedance distributions under very general circumstances. Using the framework of Davison and Smith (1990) they considered models in which the logarithm of the scale parameter in the GPD is linearly related to the meteorological covariates, in a vein similar to the Cox and Chu (1993, 1996) modeling of the Weibull scale parameter. The results indicated significant year effect, temperature, wind speed, and temperature–wind speed interaction, much as in the analysis of exceedances. Smith (1989) considered an extreme value analysis of ozone without including covariates, and found the model to have exponential tails. In contrast, the inclusion of covariates resulted in a model with shorter than exponential tails. Thus, the inclusion of covariates was important for the choice of exceedance distribution.

One advantage of the extreme value approach is that one can easily compute a forecast probability of exceedance for a day, given its (forecasted) meteorology and current-day maximum ozone. This enables local governments to release reliable ozone alert forecasts. In addition, an index of "bad ozone years" can be computed by summing exceedance probabilities over days to compute the expected number of exceedances, based only on meteorology. Such an index would be useful in determining which violations of the ozone standard are amenable to air pollution control strategies. The "area over threshold", viz., the sum over exceedance days of the amount by which the regulatory level is exceeded, may also be of interest as an indicator of ambient ozone exposure.

3.3. Time series filtering

Rao, Zurbenko, and colleagues have published a series of papers in recent years on meteorological adjustment of ozone data for the assessment of ozone trends and management programs (Rao et al., 1992, 1995, 1997; Flaum et al., 1996; Porter et al., 1996; Eskridge et al., 1997; Milanchus et al., 1998). The Rao–Zurbenko approach aims to separate ozone time series into three components: a synoptic-scale component attributable to weather and short-term fluctuations in precursor emissions, a seasonal scale component reflecting variation in the solar angle, and a long-term component manifesting effects of changes in climate, policy, and/or economics. We focus on two of the most recent publications of the Rao–Zurbenko methodology: Rao et al. (1997) and Milanchus et al. (1998). They considered the model

$$X(t) = b(t) + W(t),$$

where $b(t)$ is a baseline component, consisting of the sum of a long-term (trend) and a seasonal variation component while $W(t)$ is short-term (weather) variation. The authors analyzed log-transformed ozone, so that the model accounts for multiplicative effects of weather on the baseline. The same type of decomposition is applied to meteorological time series such as temperature.

The authors separated the baseline component from the short-term variation using a computationally simple iterative application of a moving average (the KZ filter). Define

$$X_t^{(i+1)} \equiv \frac{1}{m} \sum_{j=-k}^{k} X_{t+j}^{(i)},$$

where $X_t^{(0)} = X_t$ and $m = 2k + 1$. Then the KZ filter can be written as $KZ_{mp}(X_t) = X_t^{(p)}$, where p is the number of iterations. Eskridge et al. (1997) shows that this calculation approximately filters all periods of less than $m \approx p^{1/2}$ days. They note that other methods of decomposition, in particular using wavelets, could also be used. Thus, general issues that arise in the application and interpretation of decomposition-based analyses are not necessarily specific to the KZ filter.

Rao, Zurbenko, and colleagues applied this filter to both log-daily max ozone, $O_{kz}(t) = KZ_{mp}(O_t)$, and to one or more meteorological variables, beginning with maximum daily temperature, $T_{kz}(t) = KZ_{mp}(t)$, with $m = 29$ days and $p = 3$ iterations. Simple linear regressions of the filtered ozone series on the filtered temperature series showed that the relationship is stronger when the filtered temperature series is lagged; for example, a lag of 16 days was selected for Chicago. Milanchus et al. (1998) explained this time lag by the relationship between solar angle, which peaks in late June, and the maximum surface temperature, which peaks in July. However, the estimated phase shift changes greatly when the model is expanded to include an additional meteorological covariate, specific humidity. For example, the optimal time lag changes from 16 days to 7 days when specific humidity was added to the model for Chicago, and from 23 days to 44 days in a model for Cliffside Park, NJ. This casts doubt on the interpretation of the phase shift parameter as scientifically meaningful.

Both the KZ-filtered series and the resulting short-term series are regressed on pairs of similarly decomposed meteorological variables. For the KZ-filtered meteorology, the authors switch from maximum temperature to a computed estimate of solar radiation (for which no phase lag appears necessary) and specific humidity. The short-term series covariates are maximum temperature and dew point depression.

The sum of the residuals from these two regressions is assumed to reveal changes in ozone attributable to changes in emissions. The KZ filter is applied again (with $m = 365$ and $p = 3$) in order to look at trends in these meteorologically adjusted residuals.

To relate this approach to simple (undecomposed) linear regression, consider the consequences of decomposing ozone $O_t = O_{\mathrm{lt}} + O_{\mathrm{st}}$ into a long-term and a short-term component, with a corresponding decomposition $M_t = M_{\mathrm{lt}} + M_{\mathrm{st}}$ for meteorology, here assumed one-dimensional for simplicity. The Rao–Zurbenko approach considers the linear models

$$O_{\mathrm{lt}} = \alpha_0 + \alpha_1 M_{\mathrm{lt}} + \varepsilon_{\mathrm{lt}},$$

$$O_{\mathrm{st}} = \beta_0 + \beta_1 M_{\mathrm{st}} + \varepsilon_{\mathrm{st}}$$

yielding the resulting composite ozone model

$$O_t = \alpha_0 + \beta_0 + \alpha_1 M_t + (\beta_1 - \alpha_1) M_{\mathrm{st}} + \varepsilon_t$$

indicating that this approach, unless the short-term and long-term regression slopes are identical, results in a linear model that takes into account short-term meteorology in a more complicated fashion than a standard linear regression model. With simple linear regression the component M_{st} would be part of the error term. Linear relationships may, of course, not capture the complexity of the ozone-meteorology association. The general idea of separating out the meteorological association with ozone for different time scales is a valuable one. Modern wavelet approaches (see, e.g., Foufoula-Georgiou and Kumar, 1994) would be an improvement over the rather simplistic Rao–Zurbenko filtering approach.

3.4. Spatio-temporal modeling

Carroll et al. (1997) used a spatially homogeneous and temporally stationary space–time model to study ozone exposure in Texas. Their purpose was modeling for spatio-temporal prediction, not meteorological adjustment for estimation of trend. The data came from 11 stations in the Houston area, and consisted of hourly measurements of ambient ozone between 1980 and 1993. At each monitoring site, temperature, wind speed and wind direction were also measured. The modeling used a square root transformation, and involved a deterministic trend, depending on time and temperature,

$$O_t^{1/2} = f_t(M_t) + Z_t,$$

where Z_t is an error term with space–time dependence structure. The relationship with temperature was a quadratic polynomial; the authors did not use wind data, since the resulting predictions of 1993 data had higher variability than predictions without the wind data. The mean function was estimated using ordinary least-squares, due to computational problems in performing the appropriate generalized least-squares estimation. The authors claimed that because of the large size of the data set, the loss of efficiency in using ordinary least squares should not be very important. The validity of this claim is not apparent to us.

In order to predict the spatial ozone field away from the monitoring stations, Carroll et al. used a kriging technique, with a space–time covariance function of a form that is not necessarily valid (Cressie, 1997). In discussion of the paper by Carroll et al., Cressie (1997) and Stein and Fang (1997) criticize the model for not incorporating a spatial trend component. The authors respond that it would be difficult to predict this part of the trend.

While this analysis is one of the very few in the literature that explicitly incorporates spatial and temporal dependence, as well as accounting for meteorology, further development of the approach is needed for more accurate representation of the spatio-temporal structure of hourly ozone data. In particular, it is important to use valid space–time covariance structures, to incorporate meteorological variables that affect the covariance structure (such as winds), and to develop computationally feasible approaches to moderately large data sets. The potential advantage is the ability to include atmospheric science explicitly in the modeling, which may lead to improved understanding of the processes involved. The approach of Wikle et al. (1998), using a Bayesian hierarchical space–time model, shows particular promise in this respect.

3.5. Model assessment

In any statistical modeling, model assessment is necessary. Attention to model assessment has varied in the literature reviewed here, and some issues such as adjustment for autocorrelation have frequently been overlooked. We do not intend here to discuss and critique model assessment as it has been carried out. Rather, we want to emphasize two particular issues of some statistical sophistication: variable selection and trend estimation.

3.5.1. Variable selection

One important issue that arises in the context of model assessment is the choice of explanatory variables to include in the model. Most assessments of variables to use in a meteorological adjustment tend to be stepwise rather than a consideration of all possible subsets, the latter being computationally challenging with large numbers of variables. Stepwise selection lacks a clear global model selection criterion, and has the problem that a variable which is eliminated early may have important interactions with other variables, which are masked by variables that are later dropped from the model (Weisberg, 1985).

A different approach, developed by Raftery et al. (1997) for linear models, and employed by Clyde (1999) in the context of health effects of particulate matter air pollution, is a fully Bayesian approach in which one sets down prior probabilities for including the various variables and then computes the resulting model uncertainty in terms of posterior probabilities for a large array of models. Proponents of Bayesian model averaging argue that better predictions are obtained and that standard errors of coefficient estimates and model predictions more appropriately account for model selection, which is largely unaccounted for in traditional regression analysis. To our knowledge, this approach has not yet been used in the present context.

3.5.2. Modeling trend

The common overall goal of the ozone modeling considered in this review may be characterized as "minimizing the unexplained variation in ozone". The approach to achieving this goal will, however, differ depending on the purpose of the analysis. In analyses aimed at forecasting, for instance, a major predictor of ozone is ozone at one or more previous time points (e.g. Galbally et al., 1986; Feister and Balzer, 1991). Analysis aimed at trend estimation would, on the other hand, be confounded by incorporation of lagged ozone values. In consideration of ozone trend estimation, the relevant trend is an adjusted trend, i.e., one that is not accounted for by meteorology.

Two justifiable approaches to trend estimation are possible. One may attempt to assess the magnitude of a trend whose form has been hypothesized by, e.g., chemical and mathematical modeling. Alternatively, trend estimation may be data driven, in that one estimates the trend structure observed in the data non-parametrically. Much of the trend estimation reviewed here falls into neither of these categories, but simply estimates a linear trend, without any justification. It is not difficult to imagine a situation where the presence of significant long-term structure could be masked by assuming a linear form.

Some models developed to estimate trend include an explicit trend term in the overall model while others assess trend in the ozone residuals, after adjustment for meteorology, seasonality and other known sources of variation. The latter approach may lead to optimistic standard error estimates for the trend as the covariability of the trend estimate with other coefficients in the model is ignored. Alternatives to assuming a functional form for testing trend are isotonic regression (Joe et al., 1996), non-parametric regression on year within a general additive model framework (Reynolds et al., 1999), using Kendall's tau on the time series of meteorologically adjusted expected ozone exceedances (Stoeckenius and Hudischewskyj, 1990), or applying rank correlation to investigate monotonic time trends in the residuals (Reynolds et al., 1998).

4. Comparison of selected methods using Chicago data

We now turn to a comparison of the leading methods for trend estimation in meteorologically adjusted ozone as applied to data from Chicago. The choice of the Chicago data set is motivated by the fact that several authors have analyzed these data. Trend analyses for the region have returned differing conclusions.

An SVD analysis of square root-transformed ozone concentrations by Reynolds et al. (1999) led to estimation of a non-monotonic trend over the years 1981–1991 which was not statistically significant. Bloomfield et al. (1996) (Davis et al., 1998a), considering a network median (untransformed) ozone level over the same years, estimated a negative trend, also not statistically significant. The authors remark that the fairly wide confidence interval on the trend coefficient implies that "even strong trends may not be detected as statistically significant results for some years". Niu (1996), using similar data for the period 1982–1993, but carrying out an extreme value analysis, obtained a significant negative trend.

Huang and Smith (1999), using the same network summary as Bloomfield et al. (1996), demonstrated that the strength of temporal trends can differ across meteorological conditions. They considered 15 terminal nodes or meteorological clusters produced by a regression tree analysis. The clusters most conducive to high ozone levels were also the smallest: three of the clusters were represented by fewer than 35 days and three further clusters had fewer than 100 observations. Nonetheless, trends were predominantly negative and they were increasingly so in meteorological conditions conducive to high ozone events.

Davis et al. (1998b) reported that re-estimation of the nonlinear model developed by Bloomfield et al. (1996) using data for the Houston region led to "unsatisfactory results". They suggested that the existence of widely disparate meteorological regimes for the Gulf Coast area favored the use of their cluster analysis approach that aims to identify unique regimes. Without certain knowledge of the complexity of the meteorology for Chicago vis-à-vis Houston, we decided nonetheless to apply the methodology of Davis et al. to the Chicago data of Bloomfield et al., but using the network ozone summary computed by Reynolds et al. (1998). Details of our calculations differed in a few ways from those of Davis et al. The meteorological data were already in the form of daily summaries (rather than 3 h data as in Houston). We found that their average linkage followed by k-means clustering was unsuccessful, but that the model-based clustering methods of Banfield and Raftery (1993) did identify seemingly useful regimes with substantially different ozone distributions. Days of high ozone fell into two of the five clusters identified and retained for subsequent analysis. Within-cluster nonlinear additive models fitted by stepwise procedures resulted in only slightly better fits then a global nonlinear additive model. Formal testing of residual time trends within clusters was not carried out, but visual inspection of the residual time series suggested no trends, positive or negative. This result contrasts with the findings of Huang and Smith (1999) based on a regression-tree approach (explicitly considering ozone) to define the meteorological clusters. The much larger number of meteorological clusters of Huang and Smith and their analysis of a network summary based on untransformed (skewed) ozone concentrations contribute to the resulting contrast.

Milanchus et al. (1998) estimated a non-significant positive trend at a single Chicago location using log daily maximum ozone concentrations for the period 1984–1995. The discrepancy with the trend estimates based on regional summaries points to the difficulties with describing regional behavior using a single site. In fact, application of the Rao–Zurbenko approach to the 1981–1991 network summary obtained in the SVD analysis results in a trend estimate comparable to that of Reynolds et al. (1999). Interestingly, a simple linear regression of ozone on meteorology applied to the single site yields residuals with very similar structure and variability to that of Milanchus et al. (1998). In this instance there does not appear to be any advantage to applying the decomposition approach, although the implicit assumption that short-term meteorology is uncorrelated with long-term ozone (and vice versa) seems justified in the Chicago data.

Smith and Huang (1993), modeling probability of exceedance of 120 ppb at three Chicago stations, estimated negative linear trends for the period 1981–1991, two of which are statistically significant. Using a threshold of 100 ppb, all three stations displayed a negative linear trend, but only one was statistically significant. Further analysis of this station revealed a significant negative quadratic trend in exceedances (at 120 ppb), a result that also held for the network maximum. In both cases, a slight increase in ozone occurred up to 1984 or 1985, followed by a decrease.

The results of Smith and Huang (1993) are different from the network summary trend analyses because they were considering trends in extreme events rather than in the average ozone baseline. While shifts in average level are sufficient to produce shifts in exceedance frequency, they are not necessary for such shifts to occur. The lack of significant trends when a lower threshold was used to define exceedance is consistent with the lack of significant trends in the average ozone level.

In summary, while there is some agreement that ozone in Chicago has tended to decrease over the period considered (with the exception of the positive trend found at a single site by Milanchus et al., 1998), there is considerable variability in the way trends are modeled and in the assessment of their statistical significance. Comparison of trend estimates between methods using the Chicago data is made more difficult by the lack of strong signal in average ozone levels.

5. Discussion

The philosophy that has guided this assessment of the literature is that ideally the statistical methodology should be process-driven, by which we mean that it should be guided by an understanding of the underlying physical mechanisms. Each of the methods compared in Section 4 makes some steps in this direction and each offers some useful contributions to the meteorological adjustment of ozone. Bloomfield et al. (1996) attempt to incorporate scientific understanding of the associations between ozone and meteorology in their nonlinear regression model. SVD (Reynolds et al., 1999) provides a sensible way of forming regional summaries for ozone and meteorology. The Rao–Zurbenko approach (Milanchus et al., 1998) allows for different associations between ozone and meteorology on different time scales. Huang and Smith (1999) distinguish trends under different meteorological conditions. The cluster analysis approach of Davis et al. (1998b, c) addresses a similar aim. In the application to Chicago region ozone data, the regression tree approach of Huang and Smith led to the identification of 15 meteorological regimes whereas our application of the Davis et al. approach led to only five meteorological clusters.

There are disadvantages to meteorological adjustment of regional summaries, and space–time models of the association between ozone and meteorology would be preferable. Research on relevant space–time modeling should explicitly acknowledge the fact that compliance monitoring networks have been designed (in part) towards finding large values of the underlying ozone field. Extreme value techniques (Smith and Huang, 1993; Niu, 1996), in our view, provide the right approach with a view to standards violations. Effects based on an assessment of mean levels are not as sensitive to tail behavior in the data, as is evidenced in the Chicago analyses. An important research area is the development of space–time models for ozone extreme values.

With regard to trend estimation, the question arises as to whether a reasonable parametric form for trend can be gleaned from physical considerations. If so, statistical methods for trend estimation are well established. On the other hand, recent techniques of non-parametric function estimation, such as generalized additive models (Hastie and Tibshirani, 1990) or wavelet approaches (Foufoula-Georgiou and Kumar, 1994), have a potential to produce not only trend estimates (which could be linear if the data warrant it), but also simultaneous confidence bands for these estimates. Such confidence bands can be used to test particular parametric models.

Summarizing nonlinear trend beyond a graphical display is complicated, as it cannot be captured by single measures, such as "percent per decade". Yet, some form of summary measure is necessary to describe the trend to the public. Time-dependent summaries may be needed over the periods in which nonlinear trends are monotone.

While, as discussed above, several useful approaches to meteorological adjustment of ozone have been developed, we believe that no single approach has distinguished itself as uniformly most appropriate. The choice of methodology will depend on the purpose of the analysis (which may, for instance, be to assess changes in mean level, changes in extremes, or forecasting of these) and the meteorological complexity of ozone formation in a given region. Emission control strategies may change both the amounts and the relative contributions of precursors, thereby changing the atmospheric chemistry of ozone production. The methods considered here all require the assumption that the meteorological dependence structure remains relatively constant over the modeled period.

The issues raised here apply more generally to meteorological adjustment of other types of air pollution, where, however, the underlying physical mechanisms will be different. Communication between statisticians and atmospheric scientists is an essential part of the model building process.

Acknowledgements

The US Environmental Protection Agency through its Office of Research and Development partially funded and collaborated in the research described here under CR 825173-01-0 to the National Research Center on Statistics and the Environment, University of Washington. It has been subjected to Agency review and approved for publication. Mention of trade names or commercial products does not constitute an endorsement or recommendation for use. The authors wish to thank two anonymous reviewers for comments on an earlier draft and to acknowledge the computational assistance of David Caccia.

References

Abdul-Wahab, S., Bouhamra, W., Ettouney, H., Sowerby, B., Crittenden, B.D., 1996. Predicting ozone levels: a statistical model for predicting ozone levels. Environmental Science and Pollution Research 3, 195–204.

Banfield, J.D., Raftery, A.E., 1993. Model-based Gaussian and non-Gaussian clustering. Biometrics 49, 803–822.

Barlow, R.E., Bartholomew, D.J., Bremner, J.M., Brunck, H.D., 1972. Statistical Inference Under Order Restrictions. Wiley, New York.

Bloomfield, P.J., Royle, J.A., Steinberg, L.J., Yang, Q., 1996. Accounting for meteorological effects in measuring urban ozone levels and trends. Atmospheric Environment 30 (17), 3067–3077.

Bloomfield, P.J., Royle, J.A., Yang, Q., 1993a. Accounting for meteorological effects in measuring urban ozone levels and trends. National Institute of Statistical Sciences Technical Report #1.

Bloomfield, P.J., Royle, J.A., Yang, Q., 1993b. Rural ozone and meteorology: analysis and comparison with urban ozone. National Institute of Statistical Sciences Technical Report #3.

Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J., 1984. Classification and Regression Trees. Wadsworth and Brooks/Cole, Monterey.

Burrows, W.R., Benjamin, M., Beauchamp, S., Lord, E.R., McCollor, D., Thomson, B., 1995. CART decision tree statistical analysis and prediction of summer season maximum surface ozone for the Vancouver, Montreal, and Atlantic Regions of Canada. Journal of Applied Meteorology 34, 1848–1862.

Carroll, R.J., Chen, R., George, E.I., Li, T.H., Newton, H.J., Schmiediche, H., Wang, N., 1997. Ozone exposure and population density in Harris County. Texas. Journal of the American Statistical Association 92, 392–404.

Clyde, M., 1999. Bayesian model averaging and model search strategies (with discussion). Bayesian Statistics 6, 157–185.

Cox, W., Chu, S., 1993. Meteorologically adjusted ozone trends in urban areas: a probabilistic approach. Atmospheric Environment 27B, 425–434.

Cox, W., Chu, S., 1996. Assessment of interannual ozone variation in urban areas from a climatological perspective. Atmospheric Environment 30, 2615–2625.

Cressie, N., 1997. Comment. Journal of the American Statistical Association 92, 411–413.

Davis, J.M., Eder, B.K., Bloomfield, P., 1998a. Modeling ozone in the Chicago urban area. In: Nychka, D., Piegorsch, W.W., Cox, L.H. (Eds.), Case Studies in Environmental Statistics, Lecture Notes in Statistics, Vol. 132. Springer, New York, pp. 5–26.

Davis, J.M., Eder, B.K., Bloomfield, P., 1998b. Regional and temporal models for ozone along the Gulf Coast. In: Nychka, D., Piegorsch, W.W., Cox, L.H. (Eds.), Case Studies in Environmental Statistics, Lecture Notes in Statistics, Vol. 132. Springer, New York, pp. 27–50.

Davis, J.M., Eder, B.K., Nychka, D., Yang, Q., 1998c. Modeling the effects of meteorology on ozone in Houston using cluster analysis and generalized additive models. Atmospheric Environment 32, 2505–2520.

Davison, A.C., Smith, R.L., 1990. Models for exceedances over high thresholds (with discussion). Journal of the Royal Statistical Society B 52, 393–442.

Eder, B.K., Davis, J.M., Bloomfield, P.J., 1994. An automated classification scheme designed to better elucidate the dependence of ozone on meteorology. Journal of Applied Meteorology 33, 1182–1199.

EPA (Environmental Protection Agency), 1998a. EPA's updated air quality standards for smog (ozone) and particulate matter. Available at the web site http://ttnwww.rtpnc.epa.gov/naaqsfin/.

EPA (Environmental Protection Agency), 1998b. Aerometric Information Retrieval System. Available at the web site http://www.epa.gov/airsdata/info.htm.

Eskridge, R.E., Ku, J.Y., Rao, S.T., Porter, P.S., Zurbenko, I.G., 1997. Separating different scales of motion in time series of meteorological variables. Bulletin of the American Meteorological Society 78, 1473–1483.

Feister, U., Balzer, K., 1991. Surface ozone and meteorological predictors on a subregional scale. Atmospheric Environment 25, 1781–1790.

Fiore, A.M., Jacob, D.J., Logan, J.A., Yin, J.H., 1998. Long-term trends in ground level ozone over the contiguous United States, 1980–1995. Journal of Geophysical Research 103, 1471–1480.

Flaum, J.B., Rao, S.T., Zurbenko, I.G., 1996. Moderating the influence of meteorological conditions on ambient ozone concentrations. Journal of the Air and Waste Management Association 46, 35–46.

Foufoula-Georgiou, E., Kumar, P. (Eds.), 1994. Wavelets in Geophysics. Academic Press, Boston.

Galbally, I.E., Miller, A.J., Hoy, R.D., Ahmet, S., Joynt, R.C., Attwood, D., 1986. Surface ozone at rural sites in the Latrobe Valley and Cape Grim, Australia. Atmospheric Environment 20, 2403–2422.

Gallant, A.R., 1987. Nonlinear Statistical Models. Wiley, New York.

Gao, F., Sacks, J., Welch, W.J., 1996. Predicting urban ozone levels and trends with semiparametric modeling. Journal of Agricultural, Biological, and Environmental Statistics 1, 404–425.

Gumbel, E.J., 1958. Statistics and Extremes. Columbia University Press, New York.

Hastie, T.J., Tibshirani, R.J., 1990. Generalized Additive Models. Chapman & Hall, London.

Huang, L., Smith, R.L., 1999. Meteorologically-dependent trends in urban ozone. Environmetrics 10, 103–118.

Joe, H., Steyn, D.G., Susko, E., 1996. Analysis of trends in tropospheric ozone in the lower Fraser Valley. British Columbia. Atmospheric Environment 30, 3413–3421.

Katsoulis, B.D., 1996. The relationship between synoptic, mesoscale and microscale meteorological parameters during poor air quality events in Athens, Greece. Science of the Total Enironment 181, 13–24.

Korsog, P.E., Wolff, G.T., 1991. An examination of urban ozone trends in the northeastern US (1973–1983) using a robust statistical method. Atmospheric Environment 25, 47–57.

Kuenzli, N., Lurmann, F., Segal, M., Ngo, L., Balmes, J., Tager, I.B., 1997. Association between lifetime ambient ozone exposure and pulmonary function in college freshmen—results from a pilot study. Environmental Research 72, 8–23.

Leadbetter, M.R., Lindgren, G., Rootzén, H., 1983. Extremes and Related Properties of Random Sequences and Series. Springer, New York.

McKendry, I.G., 1994. Synoptic circulation and summertime ground-level ozone concentrations at Vancouver, British Columbia. Journal of Applied Meteorology 33, 627–641.

Milanchus, M.L., Rao, T.S., Zurbenko, I.G., 1998. Evaluating the effectiveness of ozone management efforts in the presence of meteorological variability. Journal of the Air and Waste Management Association 48, 201–215.

NRC (National Research Council), 1991. Rethinking the Ozone Problem in Urban and Regional Air Pollution. National Academy Press, Washington.

Niu, X.F., 1996. Nonlinear additive models for environmental time series with applications to ground-level ozone data analysis. Journal of the American Statistical Association 91, 1310–1321.

Pickands, J., 1975. Statistical inference using extreme order statistics. Annals of Statistics 3, 119–131.

Poissant, L., Bottenheim, J.W., Roussel, P., Reid, N.W., Niki, H., 1996. Multivariate analysis of a 1992 SONTOS data subset. Atmospheric Environment 30 (12), 2133–2144.

Porter, P.S., Rao, S.T., Zurbenko, I.G., Zalewsky, E., Henry, R.F., Ku, J.Y., 1996. Statistical characteristics of spectrally-decomposed ambient ozone time series data. In Reports on the OTAG Web Site http://capita.wustl.edu/otag/Reports/Reports.html.

Pryor, S.C., McKendry, I.G., Steyn, D.G., 1995. Synoptic-scale meteorological variability and surface ozone concentrations in Vancouver, British Columbia. Journal of Applied Meteorology 34, 1824–1833.

Raftery, A.E., Madigan, D., Hoeting, J.A., 1997. Bayesian model averaging for linear regression models. Journal of the American Statistical Association 92, 179–191.

Rao, S.T., Sistia, G., Henry, R., 1992. Statistical analysis of trends in urban ozone air quality. Journal of the Air and Waste Management Association 42, 1204–1211.

Rao, S.T., Zalewsky, E., Zurbenko, I.G., 1995. Determining temporal and spatial variations in ozone air quality. Journal of the Air and Waste Management Association 45, 57–61.

Rao, S.T., Zurbenko, I.G., Neagu, R., Porter, P.S., Ku, J.Y., Henry, R.F., 1997. Space and time scales in ambient ozone data. Bulletin of the American Meteorological Society 78, 2153–2166.

Reynolds, J.H., Das, B., Sampson, P.D., Guttorp, P., 1998. Meteorological adjustment of Western Washington and Northwest Oregon surface ozone observations with investigation of trends. National Research Center for Statistics and the Environment Technical Report 15. Available at http:// www.nrcse.washington.edu/research/reports/papers/trs15_doe/trs15_doe.pdf.

Reynolds, J.H., Caccia, D., Sampson, P.D., Guttorp, P., 1999. Meteorological adjustment of Chicago, Illinois, regional surface ozone observations with investigation of trends. National Research Center for Statistics and the Environment Technical Report 25. Available at http://www.nrcse.washington.edu/research/reports/papers/trs25_chicago/trs25_chicago.pdf.

Smith, R.L., 1989. Extreme value analysis of environmental time series: an application to trend detection in ground-level ozone (with discussion). Statistical Science 4, 367–393.

Smith, R.L., Huang, L., 1993. Modeling high threshold exceedances of urban ozone. National Institute for Statistical Science Technical Report #6.

Smith, R.L., Shively, T.S., 1995. Point process approach to modeling trends in tropospheric ozone based on exceedances of a high threshold. Atmospheric Environment 29 (23), 3489–3499.

Spichtinger, N., Winterhalter, M., Fabian, P., 1996. Ozone and Grosswetterlagen: analysis for the Munich metropolitan area. Environmental Science and Pollution Research 3, 145–152.

Stein, M.L., Fang, D., 1997. Comment. Journal of the American Statistical Association 92, 408–411.

Stieb, D.M., Burnett, R.T., Beveridge, R.C., Brook, J.R., 1996. Association between ozone and asthma emergency department visits in Saint John, New Brunswick, Canada. Environmental Health Perspectives 104, 1354–1360.

Stoeckenius, T.E., Hudischewskyj, A.B., 1990. Adjustment of ozone trends for meteorological variation. SYAPP-90/008, Systems Applications Inc., San Rafael, CA.

van Ooy, D.J., Carroll, J.J., 1995. The spatial variation of ozone climatology on the western slope of the Sierra Nevada. Atmospheric Environment 29 (11), 1319–1330.

Weisberg, S., 1985. In: Applied Linear Regression, 2nd Edition. Wiley, New York, pp. 214–215.

Wikle, C.K., Berliner, L.M., Cressie, N., 1998. Hierarchical Bayesian space–time models. Environmental and Ecological Statistics 5, 117–154.

Xu, D., Yap, D., Taylor, P.A., 1996. Meteorologically adjusted ground level ozone trends in Ontario. Atmospheric Environment 30, 1117–1124.