

KRIGING WITH ESTIMATED PARAMETERS

Richard L. Smith

Department of Statistics and Operations Research

University of North Carolina

Chapel Hill

UNC Statistics and Operations Research Colloquium

January 10 2005

Acknowledgements to:

Zhengyuan Zhu

Elizabeth Shamseldin

Stas Kolenikov

Joe Ibrahim

References

For a preliminary version of a paper on this talk is based, see:

Smith, R.L. and Zhu, Z. (2004), Asymptotic theory for kriging with estimated parameters and its application to network design.
<http://www.stat.unc.edu/postscript/rs/supp5.pdf>

I Motivating Example: Interpolation of Air Pollution Data

II Universal Kriging, REML Estimation and Bayesian Spatial Statistics

III Existing Results on “Kriging With Estimated Parameters”

IV The Approach Based on Second-Order Asymptotics

V Computations and Applications

I. Motivating Example: Interpolation of Air Pollution Data

The problems of spatial interpolation and network design are introduced through an example using the EPA fine particulates ($PM_{2.5}$) network. The analysis is taken from Smith, Kolenikov and Cox (2003).

Current EPA standard for PM_{2.5}:

- Twenty-four hour average PM_{2.5} not to exceed 65 $\mu\text{g}/\text{m}^3$ for a three-year average of annual 98th percentiles at any population-oriented monitoring site in a monitoring area.
- Three-year annual average PM_{2.5} not to exceed 15 $\mu\text{g}/\text{m}^3$ concentrations from a single community-oriented monitoring site or the spatial average of eligible community-oriented monitoring sites in a monitoring area.

The data (compiled from a larger data set) consisted of weekly average PM_{2.5} levels during 1999 at 74 EPA stations in NC, SC, GA. We fitted a model of the form

$$y_{xt} = f(x) + g(t) + \eta_{xt}$$

in which y_{xt} is the square root of PM_{2.5} in location x in week t , $f(x)$ and $g(t)$ are fixed functions of time and space respectively (both represented as linear regression functions), and η_{xt} is a random error (independent in time but dependent in space).

The model parameters were estimated by maximum likelihood, and kriging (described later) was used to interpolate η_{xt} at locations x off the network. Based on that, a map was constructed of the interpolated annual average for each point in the study region, together with an estimated root mean square prediction error.

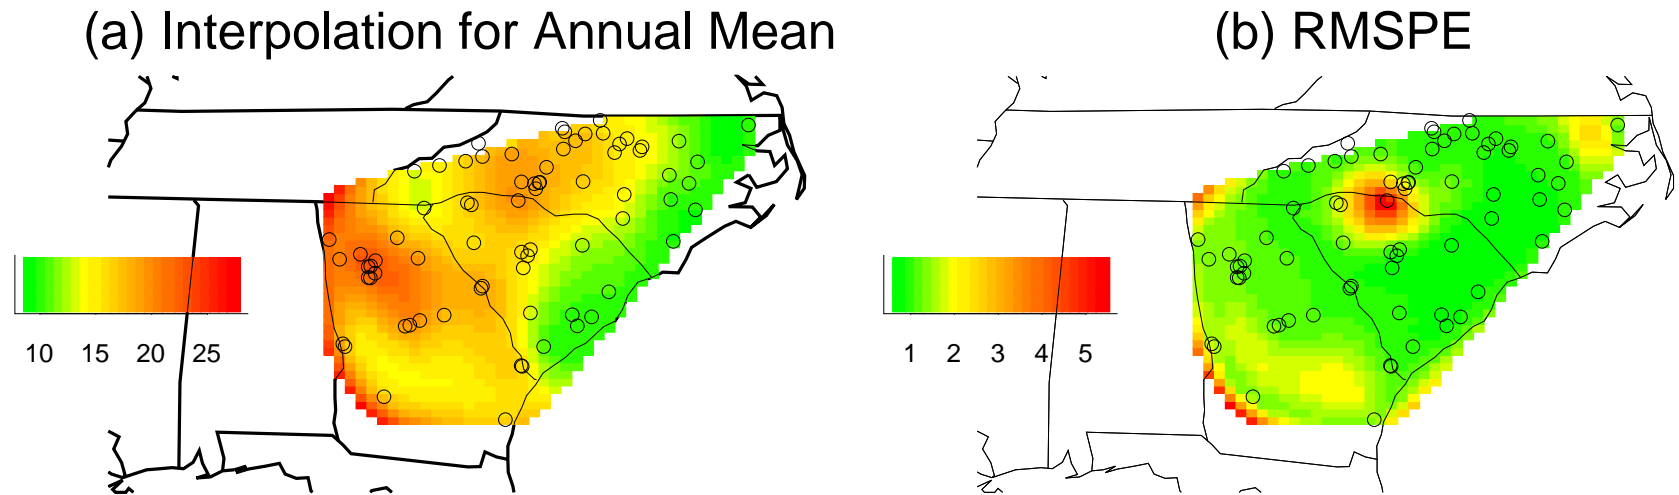


Fig. 1. (a) Interpolated surface for annual mean PM_{2.5} in $\mu\text{g}/\text{m}^3$, and monitor locations (circles). (b) Root mean squared prediction error for the surface in (a). Adapted from Smith, Kolenikov and Cox (2003).

Questions for this talk:

There are many questions along the lines of what is the right model for this data set, but in this talk, I focus on two more theoretical aspects:

1. *Computation of RMSPEs.* The calculations above used the conventional formulae for kriging prediction errors that you can find in any text on geostatistics, but these ignore model estimation errors. Is this appropriate, and how could we improve on this approach?
2. *Design of the network.* The EPA is constantly adding or deleting stations as it attempts to provide better coverage or to reduce costs. How could it do this most efficiently?

The two questions are linked because accurate determination of kriging errors is critical in assessing the design.

II. Universal Kriging, REML Estimation and Bayesian Spatial Statistics

We assume data follow a *Gaussian random field* with mean and covariance functions represented as functions of finite-dimensional parameters.

Define the prediction problem as

$$\begin{pmatrix} Y \\ Y_0 \end{pmatrix} \sim N \left[\begin{pmatrix} X\beta \\ x_0^T\beta \end{pmatrix}, \begin{pmatrix} V & w^T \\ w & v_0 \end{pmatrix} \right] \quad (1)$$

where Y is an n -dimensional vector of observations, Y_0 is some unobserved quantity we want to predict, X and x_0 are known regressors, and β is a p -dimensional vector of unknown regression coefficients. For the moment, we assume V , w and v_0 are known.

The most widely used spatial models (stationary and isotropic) assume the covariance between components Y_i and Y_j is a function of the (scalar) distance between them, $C_\theta(d_{ij})$. An example is the *exponential power model*

$$C_\theta(d) = \sigma^2 \exp \left\{ - \left(\frac{d}{\rho} \right)^\kappa \right\},$$

where $\theta = (\kappa, \sigma^2, \rho)$ with $0 < \kappa \leq 2$, $\sigma^2 > 0$, $\rho > 0$.

The PM_{2.5} data analysis actually assumed

$$\text{Var}\{Y_i - Y_j\} = \theta_1 + \theta_2 d_{ij}^{\theta_3} \quad (\theta_1 > 0, \theta_2 > 0, 0 < \theta_3 \leq 2).$$

This is of *intrinsically stationary* form and $C_\theta(d)$ does not exist, though the model can easily be transformed into one for which the methods of this talk apply.

The key assumption is: the covariances are unknown in practice, but expressed as functions of finitely many parameters θ .

Universal Kriging

Assume model (1) where the covariances V , w , v_0 are known but β is unknown. The classical formulation of *universal kriging* asks for a predictor $\hat{Y}_0 = \lambda^T Y$ that minimizes $\sigma_0^2 = E \{ (Y_0 - \hat{Y}_0)^2 \}$ subject to the unbiasedness condition $E \{ Y_0 - \hat{Y}_0 \} = 0$.

The classical solution:

$$\begin{aligned}\lambda &= w^T V^{-1} + (x_0 - X^T V^{-1} w)^T (X^T V^{-1} X)^{-1} X^T V^{-1}, \\ \sigma_0^2 &= v_0 - w^T V^{-1} w + (x_0 - X^T V^{-1} w)^T (X^T V^{-1} X)^{-1} (x_0 - X^T V^{-1} w).\end{aligned}$$

Maximum Likelihood and REML Estimation

Model of form

$$Y \sim N[X\beta, V(\theta)].$$

A classical method of estimation is the *method of maximum likelihood* (MLE), in which the parameters β and θ are chosen to maximize the joint density of Y given β and θ (the likelihood function).

In practice, once θ is specified, the MLE $\hat{\beta}$ can be calculated by elementary algebra (the generalized least squares estimator). Therefore, in practice MLE is computed by maximizing the *profile log likelihood*

$$-\frac{1}{2} \log |V(\theta)| - \frac{G^2(\theta)}{2}$$

where $G^2 = Y^T W Y$, $W = V^{-1} - V^{-1} X^T (X V^{-1} X^T)^{-1} X V^{-1}$, is the generalized residual sum of squares.

An alternative to MLE is to use the *restricted likelihood function*. As defined originally by Patterson and Thompson (1971), if X is $n \times p$, we define a $(n - p)$ -dimensional vector of contrasts $A^T Y$, where A is $n \times (n - p)$ has rank $n - p$ and satisfies $A^T X = 0$. The restricted likelihood is the density of $A^T Y \sim N[0, A^T V A]$.

In practice this is equivalent to maximizing the function

$$\ell_n(\theta) = -\frac{1}{2} \log |V(\theta)| - \frac{1}{2} \log |X^T V(\theta)^{-1} X| - \frac{G^2(\theta)}{2}.$$

The resulting estimator is called REML. It is usually considered superior to MLE, though the two estimators are equivalent to first-order asymptotics.

Bayesian Reformulation

Suppose (β, θ) have a joint prior density of the form $\pi(\theta)d\beta d\theta$ (constant in β).

The Bayesian predictive density of Y_0 given Y is

$$p(Y_0 | Y) = \frac{\int \int f(Y, Y_0 | \beta, \theta) \pi(\theta) d\beta d\theta}{\int \int f(Y | \beta, \theta) \pi(\theta) d\beta d\theta}.$$

After some algebraic manipulation, this may be rewritten

$$p(Y_0 | Y) = \frac{\int e^{\ell_n(\theta)} \psi(Y_0 | Y, \theta) \pi(\theta) d\theta}{\int e^{\ell_n(\theta)} \pi(\theta) d\theta} \quad (2)$$

where

$$\psi(Y_0 | Y, \theta) = \frac{1}{\sqrt{2\pi}\sigma_0(\theta)} \exp \left\{ -\frac{1}{2} \left(\frac{Y_0 - \lambda(\theta)^T Y}{\sigma_0(\theta)} \right)^2 \right\}.$$

Two forms of predictive density

The REML estimator $\hat{\theta}$ is the value of θ that maximizes $\ell_n(\theta)$.

The conventional kriging formula uses the predictive density

$$\hat{\psi}(Y_0 | Y) = \psi(Y_0 | Y, \hat{\theta})$$

also known as *estimative density* (Aitchison) or *plug-in rule*.

In contrast to this, we write (2) as $\tilde{\psi}(Y_0 | Y)$, which Aitchison called the *predictive density*.

In subsequent discussion, we shall mostly use the predictive distribution function, i.e. redefine $\psi(z|Y, \theta) = \Phi\left(\frac{z - \lambda(\theta)^T Y}{\sigma_0(\theta)}\right)$ where $\Phi(\cdot)$ is the standard normal distribution function.

The first central question of this talk is then: *Is $\tilde{\psi}$ superior to $\hat{\psi}$?* (and if so, how is this influenced by the choice of prior?)

Designing Monitor Networks

Large literature, many different approaches.

Recent work has focussed on contrast between two types of criterion:

- *Estimative* — e.g. choose the design to maximize the determinant of the Fisher information matrix of θ
- *Predictive* — focus on a specific Y_0 , find a design to minimize σ_0 . Note that this ignores the estimation of θ , in effect assuming θ known.

Zhu and Stein (2004) discussed the idea of using Bayesian prediction intervals as a basis for network design, arguing that Bayesian intervals take account of the uncertainty of θ and therefore should be superior to the predictive approach. However, they rejected this as too computationally intensive, and instead proposed a two-stage criterion (more later).

Direct Bayesian Approach

- For any data set, use MCMC to construct the Bayesian predictive distribution
- For any given design, run the Bayesian analysis on simulated data sets to determine the expected length of Bayesian prediction intervals
- Use an optimization algorithm (e.g. simulated annealing) to find the optimal design

Direct implementation of this approach requires a lot of Monte Carlo simulation. The second major theme of this talk is to propose an approximate approach that reduces the first two steps to a simple algebraic formula for the expected length of a Bayesian prediction interval.

III Existing Results on “Kriging With Estimated Parameters”

Suppose we apply universal kriging to predict Y_0 by $\lambda^T Y$, but estimate $\hat{\lambda} = \lambda(\hat{\theta})$ where $\hat{\theta}$ is the MLE or REMLE.

Harville and Jeske (1992) and Zimmerman and Cressie (1992) proposed the following correction to the mean squared prediction error:

$$V_1 = E \left\{ (Y_0 - \hat{\lambda}^T Y)^2 \right\} \approx \sigma_0^2 + \text{tr} \left\{ \mathcal{I}^{-1} \left(\frac{\partial \lambda}{\partial \theta} \right)^T V \left(\frac{\partial \lambda}{\partial \theta} \right) \right\}$$

where \mathcal{I} is the observed information matrix for θ . This formula corrects for the error in specifying the kriging weights λ .

The derivation of this formula assumed that $\hat{\theta} - \theta$ was independent of $Y_0 - \lambda^T Y$. Abt (1999) derived an improved formula without this assumption, but noted that in practice, the improvement made little difference to the result.

However, in calculating a prediction *interval* for Y_0 , it is also necessary to consider the effect of σ_0^2 being unknown. Stein (1999) and Zhu and Stein (2004) defined

$$V_2 = \left(\frac{\partial \sigma_0^2}{\partial \theta} \right)^T \mathcal{I}^{-1} \left(\frac{\partial \sigma_0^2}{\partial \theta} \right)$$

as a measure of the uncertainty in σ_0^2 , and they suggested that some linear combination of V_1 and $\frac{V_2}{\sigma_0^2}$ would best measure the overall uncertainty. In particular, they suggested

$$V_3 = V_1 + \frac{1}{2} \cdot \frac{V_2}{\sigma_0^2}$$

as a suitable combined criterion. However, it's not clear exactly why this particular linear combination is appropriate.

Bayesian Approaches Based on Reference Priors

For a model with parameters (β, σ^2, ρ) where σ^2 is the marginal variance and ρ is the range parameter, Berger, De Oliveira and Sansó (2001) derived the “reference prior” in the form

$$\pi(\beta, \sigma^2, \rho) \propto \frac{1}{\sigma^2} \left[\text{tr} \left(W \frac{\partial V}{\partial \rho} W \frac{\partial V}{\partial \rho} \right) - \frac{1}{n-p} \left\{ \text{tr} \left(W \frac{\partial V}{\partial \rho} \right) \right\}^2 \right]^{1/2}. \quad (3)$$

They compare this with alternative definitions of reference prior and Jeffreys prior. However, they don't point out that (3) is actually the same as the Jeffreys prior derived from the restricted likelihood.

Paulo (2003) extended their result to a general multi-parameter spatial likelihood.

Both papers used simulation to demonstrate that the method would produce good frequentist coverage probability.

IV The Approach Based on Second-Order Asymptotics

Long history —

- *Frequentist Asymptotics for Prediction* — Cox (1975), Barndorff-Nielsen and Cox (1996), Hall, Peng and Tajvidi (1999),...
- *Predictive Likelihood* — Lauritzen (1974), Hinkley (1979), Butler (1986), Davison (1986), Bjørnstad (1990),....
- *Decision Theoretic Approaches* — Aitchison (1975), Harris (1989), Komaki (1996), Smith (1999)
- *Matching Bayesian and Frequentist Inference* — Welch and Peers (1963),....., Datta and Mukerjee (2004 Springer-Verlag Monograph). See in particular, Datta, Mukerjee, M. Ghosh and Sweeting (2000, *Annals of Statistics*) for a “matching prior” approach to predictive inference.

With scattered exceptions, all of this literature applies only to the case of *independent* observations.

Notation

Define

$$\tilde{\psi}(z | Y) = \frac{\int e^{\ell_n(\theta)+Q(\theta)}\psi(z | Y, \theta)d\theta}{\int e^{\ell_n(\theta)+Q(\theta)}d\theta} \quad (4)$$

where $e^{\ell_n(\theta)}$ is the restricted likelihood of θ , $Q(\theta) = \log \pi(\theta)$ and $\psi(z | Y, \theta) = \Phi\left(\frac{z-\lambda(\theta)^T Y}{\sigma_0(\theta)}\right)$. Also let $\tilde{\psi}^{-1}$ be inverse function, i.e. $\tilde{\psi}^{-1}(P | Y)$ is the value of z for which $\tilde{\psi}(z | Y) = P$.

For $P \in (0, 1)$ define

$$\begin{aligned} z_P(Y | \theta) &= \lambda(\theta)^T Y + \sigma_0(\theta)\Phi^{-1}(P), \\ \hat{z}_P(Y) &= \hat{\lambda}^T Y + \hat{\sigma}_0\Phi^{-1}(P), \\ \tilde{z}_P(Y) &= \tilde{\psi}^{-1}(P | Y). \end{aligned}$$

For an estimator z_P^* (could be \hat{z}_P or \tilde{z}_P) we would like to calculate

$$E \{ \psi(z_P^*(Y) | Y, \theta) - \psi(z_P(Y | \theta) | Y, \theta) \} \quad (5)$$

and

$$E \{ z_P^*(Y) - z_P(Y | \theta) \} \quad (6)$$

(5) is called the coverage probability bias (CPB). (6) leads to the expected length of a prediction interval (our proposed design criterion) because for a $100(P_2 - P_1)\%$ interval,

$$\begin{aligned} & E \{ z_{P_2}^*(Y) - z_{P_1}^*(Y) \} \\ = & E \{ z_{P_2} - z_{P_1} \} + E \{ z_{P_2}^* - z_{P_2} \} - E \{ z_{P_1}^* - z_{P_1} \} \\ = & \sigma_0 \{ \Phi^{-1}(P_2) - \Phi^{-1}(P_1) \} + E \{ z_{P_2}^* - z_{P_2} \} - E \{ z_{P_1}^* - z_{P_1} \} \end{aligned}$$

Define $U_i = \frac{\partial \ell_n(\theta)}{\partial \theta^i}$, $U_{ij} = \frac{\partial^2 \ell_n(\theta)}{\partial \theta^i \partial \theta^j}$, $U_{ijk} = \frac{\partial^3 \ell_n(\theta)}{\partial \theta^i \partial \theta^j \partial \theta^k}$.

The matrix with entries U_{ij} has an inverse with entries U^{ij} .

Other quantities $Q(\theta) = \log \pi(\theta)$, $\lambda(\theta)$, $\sigma_0(\theta)$. Suffixes denote partial differentiation, e.g. $Q_i = \frac{\partial Q}{\partial \theta^i}$, $\sigma_{0ij} = \frac{\partial^2 \sigma_0}{\partial \theta^i \partial \theta^j}$. Let

$$\begin{aligned} U_i &= n^{1/2} Z_i, \\ U_{ij} &= n^{1/2} Z_{ij} + n \kappa_{ij}, \\ U_{ijk} &= n^{1/2} Z_{ijk} + n \kappa_{ijk}, \end{aligned}$$

and define also $\kappa_{i,j} = n^{-1} E \{U_i U_j\} = -\kappa_{ij}$, $\kappa_{ij,k} = n^{-1} E \{U_{ij} U_k\}$. Suppose inverse of $\{\kappa_{i,j}\}$ matrix has entries $\{\kappa^{i,j}\}$. We assume all the Z quantities are $O_p(1)$ and all the κ quantities are $O(1)$ as $n \rightarrow \infty$ and we employ the summation convention.

Step 1: Taylor expansion of $\hat{\psi}$

$$\begin{aligned}\hat{\psi} - \psi &= n^{-1/2} \kappa^{i,j} Z_i \psi_j + n^{-1} (\kappa^{i,j} \kappa^{k,\ell} Z_{ik} Z_j \psi_\ell \\ &\quad + \frac{1}{2} \kappa^{i,r} \kappa^{j,s} \kappa^{k,t} \kappa_{ijk} Z_r Z_s \psi_t \\ &\quad + \frac{1}{2} \kappa^{i,j} \kappa^{k,\ell} Z_i Z_k \psi_{j\ell}) + O_p(n^{-3/2}).\end{aligned}$$

Follows well-known references on higher-order asymptotics of MLE, e.g. McCullagh (1987), Barndorff-Nielsen and Cox (1994).

Step 2: From $\hat{\psi}$ to $\tilde{\psi}$

Using a Laplace approximation to the Bayesian integral, Lindley (1980) showed that

$$\tilde{\psi} - \hat{\psi} = \frac{1}{2} \hat{U}_{ijk} \hat{\psi}_\ell \hat{U}^{ij} \hat{U}^{kl} - \frac{1}{2} (\hat{\psi}_{ij} + 2\hat{\psi}_i \hat{Q}_j) \hat{U}^{ij} + O_p(n^{-2}).$$

where the hats indicate that U_{ijk}, ψ_ℓ , etc., are evaluated at the REMLE $\theta = \hat{\theta}$.

Hence

$$\tilde{\psi} - \hat{\psi} = \frac{1}{2n} \left\{ \kappa_{ijk} \kappa^{i,j} \kappa^{k,\ell} \psi_\ell + (\psi_{ij} + 2\psi_i Q_j) \kappa^{i,j} \right\} + O_p(n^{-3/2}).$$

Together, these approximations give an expansions of either $\hat{\psi} - \psi$ or $\tilde{\psi} - \psi$ in powers of $n^{-1/2}$.

Step 3: From a distribution function to its inverse (Cox 1975, Smith 1999)

Suppose $\psi^*(z)$ (could be $\hat{\psi}$ or $\tilde{\psi}$) is an estimator of the conditional prediction distribution function $\psi(z) = \psi(z | Y, \theta)$ that has an expansion

$$\psi^*(z) = \psi(z) + n^{-1/2}R + n^{-1}S + o_p(n^{-1}).$$

Define predictive quantile z_P^* by $\psi^*(z_P^*) = P$. Then

$$z_P^* - z_P = -n^{-1/2} \frac{R}{\psi'} - n^{-1} \left(\frac{RR'}{\psi'^2} - \frac{R^2}{\psi'^3} - \frac{S}{\psi'} \right) + o_p(n^{-1}), \quad (7)$$

$$\psi(z_P^*) - \psi(z_P) = -n^{-1/2}R - n^{-1} \left(\frac{RR'}{\psi'} - S \right) + o_p(n^{-1}). \quad (8)$$

Here primes denote differentiation with respect to z . By taking expectations in (7) and (8) respectively, we derive expressions for the expected length of a prediction interval and the coverage probability bias (CPB).

Step 4: Evaluate the expectations

Side comment: Up to this point, the calculations are the same as in the independent case. In particular, by taking expectations in (8), it should be possible to re-derive the results in Datta, Mukerjee, Ghosh and Sweeting (2000) (though they used a different method)

However, for dependent observations, these expectations are much harder to evaluate.

As an example of the terms to be evaluated, consider

$$\begin{aligned}
 E \left\{ \frac{RR'}{\psi'} \right\} &= E \left\{ \frac{\kappa^{i,j} Z_i \psi_j \kappa^{k,l} Z_k \psi'_l}{\psi'} \right\} \\
 &= \kappa^{i,j} \kappa^{k,l} \phi(\Phi^{-1}(P)) \times \\
 E \left[Z_i Z_k \left\{ \frac{\lambda_j^T Y}{\sigma_0} + \frac{\sigma_{0j}}{\sigma_0} \Phi^{-1}(P) \right\} \left\{ \frac{\sigma_{0l}}{\sigma_0} - \left(\frac{\lambda_l^T Y}{\sigma_0} + \frac{\sigma_{0l}}{\sigma_0} \Phi^{-1}(P) \right) \Phi^{-1}(P) \right\} \right]
 \end{aligned}$$

This requires evaluating moments of Y 's of up to sixth order, but it is still an explicit calculation!

If $\epsilon_i, \epsilon_j, \dots$ are jointly normal random variables with mean 0 and covariances $\{\sigma_{ij}\}$, then

$$\begin{aligned} E\{\epsilon_i \epsilon_j\} &= \sigma_{ij}, \\ E\{\epsilon_i \epsilon_j \epsilon_k \epsilon_\ell\} &= \sigma_{ij} \sigma_{kl} + \sigma_{ik} \sigma_{j\ell} + \sigma_{il} \sigma_{jk}, \\ E\{\epsilon_i \epsilon_j \epsilon_k \epsilon_\ell \epsilon_r \epsilon_s\} &= \sigma_{ij} \sigma_{kl} \sigma_{rs} [15] \end{aligned}$$

where [15] denotes that the same term is repeated and summed 15 times according to all the possible ways of dividing the six indices into three blocks of size two.

All odd-order moments are 0.

Results

$$\begin{aligned}
& nE \{ \psi(\hat{z}_P(Y) \mid Y, \theta) - \psi(z_P(Y \mid \theta) \mid Y, \theta) \} \\
& \sim \phi(\Phi^{-1}(P))\Phi^{-1}(P) \left[-\frac{1}{2}\Phi^{-1}(P)^2 \kappa^{i,j} \frac{\sigma_{0i}\sigma_{0j}}{\sigma_0^2} \right. \\
& \quad \left. + \kappa^{i,j} \kappa^{k,\ell} \left(\kappa_{jk,\ell} + \frac{1}{2}\kappa_{jkl} \right) \frac{\sigma_{0i}}{\sigma_0} + \frac{1}{2}\kappa^{i,j} \left\{ \frac{\sigma_{0ij}}{\sigma_0} - \frac{\lambda_i^T V \lambda_j}{\sigma_0^2} \right\} \right. \\
& \quad \left. - \frac{1}{2}\kappa^{i,k} \kappa^{j,\ell} \cdot \frac{1}{n\sigma_0^2} \left(\lambda_i^T V \frac{\partial W}{\partial \theta^k} V \frac{\partial W}{\partial \theta^\ell} V \lambda_j + \lambda_i^T V \frac{\partial W}{\partial \theta^\ell} V \frac{\partial W}{\partial \theta^k} V \lambda_j \right) \right],
\end{aligned}$$

$$\begin{aligned}
& nE \{ \psi(\tilde{z}_P(Y) \mid Y, \theta) - \psi(z_P(Y \mid \theta) \mid Y, \theta) \} \\
& \sim \phi(\Phi^{-1}(P))\Phi^{-1}(P) \left[\kappa^{i,j} \kappa^{k,\ell} \left(\kappa_{jk,\ell} + \kappa_{jkl} \right) \frac{\sigma_{0i}}{\sigma_0} \right. \\
& \quad \left. - \kappa^{i,j} \left(\frac{\sigma_{0i}\sigma_{0j}}{\sigma_0^2} - \frac{\sigma_{0ij}}{\sigma_0} \right) + \kappa^{i,j} \frac{\sigma_{0i}}{\sigma_0} Q_j \right. \\
& \quad \left. - \frac{1}{2}\kappa^{i,k} \kappa^{j,\ell} \cdot \frac{1}{n\sigma_0^2} \left(\lambda_i^T V \frac{\partial W}{\partial \theta^k} V \frac{\partial W}{\partial \theta^\ell} V \lambda_j + \lambda_i^T V \frac{\partial W}{\partial \theta^\ell} V \frac{\partial W}{\partial \theta^k} V \lambda_j \right) \right].
\end{aligned}$$

$$nE \{ \hat{z}_P - z_P \} \approx \Phi^{-1}(P) \left\{ \kappa^{i,j} \kappa^{k,\ell} \sigma_{0\ell} \left(\kappa_{ik,j} + \frac{1}{2} \kappa_{ijk} \right) + \frac{1}{2} \kappa^{i,j} \sigma_{0ij} \right\}$$

$$\begin{aligned} nE \{ \tilde{z}_P - z_P \} \approx & \Phi^{-1}(P) \left\{ \kappa^{i,j} \kappa^{k,\ell} \sigma_{0\ell} (\kappa_{ik,j} + \kappa_{ijk}) \right. \\ & + \kappa^{i,j} \left(\sigma_{0ij} - \frac{\sigma_{0i} \sigma_{0j}}{\sigma_0} \right) + \kappa^{i,j} Q_j \sigma_{0i} \\ & \left. + \frac{1}{2} \Phi^{-1}(P)^2 \kappa^{i,j} \frac{\sigma_{0i} \sigma_{0j}}{\sigma_0} + \frac{1}{2} \kappa^{i,j} \frac{\lambda_i^T V \lambda_j}{\sigma_0} \right\}. \end{aligned}$$

These results imply the existence of a “matching prior” for which the second-order CPB is 0. However we can also manipulate the asymptotic expressions to obtain a direct estimate of z_P with the same property:

$$\begin{aligned}
z_P^\dagger = & \hat{z}_P - n^{-1} \Phi^{-1}(P) \left\{ \hat{\kappa}^{i,j} \hat{\kappa}^{k,\ell} \hat{\sigma}_{0\ell} \left(\hat{\kappa}_{ik,j} + \frac{1}{2} \hat{\kappa}_{ijk} \right) \right. \\
& + \frac{1}{2} \hat{\kappa}^{i,j} \left(\hat{\sigma}_{0ij} - \frac{\hat{\sigma}_{0i} \hat{\sigma}_{0j}}{\hat{\sigma}_0} \Phi^{-1}(P)^2 \right) - \frac{1}{2\hat{\sigma}_0} \hat{\kappa}^{i,j} \hat{\lambda}_i^T \hat{V} \hat{\lambda}_j \\
& \left. - \frac{1}{2n\hat{\sigma}_0} \hat{\kappa}^{i,j} \hat{\kappa}^{k,\ell} \left(\hat{\lambda}_j^T \hat{V} \frac{\partial \hat{W}}{\partial \theta^i} \hat{V} \frac{\partial \hat{W}}{\partial \theta^k} \hat{V} \hat{\lambda}_\ell + \hat{\lambda}_j^T \hat{V} \frac{\partial \hat{W}}{\partial \theta^k} \hat{V} \frac{\partial \hat{W}}{\partial \theta^i} \hat{V} \hat{\lambda}_\ell \right) \right\}.
\end{aligned}$$

Application to Network Design

Suppose we use an estimator of z_P whose second-order CPB is 0 (e.g. either the Bayes estimator with matching prior, or z_P^\dagger). In either case we have

$$nE \left\{ z_P^\dagger - z_P \right\} \approx \frac{1}{2\sigma_0} \Phi^{-1}(P) \left\{ \kappa^{i,j} \lambda_i^T V \lambda_j + \Phi^{-1}(P)^2 \kappa^{i,j} \sigma_{0i} \sigma_{0j} \right. \\ \left. + \kappa^{i,k} \kappa^{j,\ell} \left(\lambda_i^T V \frac{\partial W}{\partial \theta^k} V \frac{\partial W}{\partial \theta^\ell} V \lambda_j + \lambda_i^T V \frac{\partial W}{\partial \theta^\ell} V \frac{\partial W}{\partial \theta^k} V \lambda_j \right) \right\}.$$

The second line is Abt's refinement of the Harville-Jeske-Zimmerman-Cressie correction and will be ignored in subsequent discussion.

Use this to construct a two-sided prediction interval, with tail probability $1 - P$ in each tail. The approximate expected length of this prediction interval is

$$2\Phi^{-1}(P)\sqrt{\sigma_0^2 + n^{-1}\kappa^{i,j}\lambda_i^T V \lambda_j + n^{-1}\Phi^{-1}(P)^2 \kappa^{i,j}\sigma_{0i}\sigma_{0j}}.$$

In the notation of Zhu and Stein (2004), the quantity under the square root sign is

$$V_4 = V_1 + \frac{\Phi^{-1}(P)^2}{4} \cdot \frac{V_2}{\sigma_0^2}.$$

Recall their own criterion was $V_3 = V_1 + \frac{1}{2} \cdot \frac{V_2}{\sigma_0^2}$.

Two Design Criteria

$$V_3 = V_1 + \frac{1}{2} \cdot \frac{V_2}{\sigma_0^2} \quad (\text{Zhu and Stein})$$

$$V_4 = V_1 + \frac{\Phi^{-1}(P)^2}{4} \cdot \frac{V_2}{\sigma_0^2} \quad (\text{this talk})$$

The present formula V_4 has the unusual feature that the design might depend on the desired coverage probability of a prediction interval.

It is also tied directly to two specific methods of constructing a prediction interval whose second-order coverage probability bias is 0, whereas previous approaches have not shown how to construct such an interval.

Summary

1. The second-order coverage probability bias of the Bayes estimator of z_P is smaller than that of the plug-in estimator in the limit as $P \rightarrow 0$ or 1 , regardless of the prior.
2. For the Bayesian predictive distribution there is a matching prior, i.e. one for which the second-order CPB of \tilde{z}_P is 0.
3. However we can also achieve the same second-order properties directly, using the estimator z_P^\dagger .
4. For any of these estimators of predictive quantiles, we have an approximation for the expected length of a prediction interval, and this can be used as a design criterion.
5. In the case of an estimate whose second-order CPB is 0, we obtain a design criterion very similar to that of Zhu and Stein, but adapted to a specific construction of a prediction interval.

V. Computations and Applications

First we show the results of a small simulation experiment, to test the coverage probability formula for REML.

A set of 20 “prediction locations” was chosen. These were predicted based on respectively 20, 35, 50 observation locations. For each sample size and prediction location, true coverage probability was estimated by simulation (10,000 replications). Also computed was the theoretical coverage probability using the formulae in this talk.

Fig. 2 shows the prediction/observation locations and Fig. 3 shows the simulated v. theoretical coverage probabilities.

For each of the three sample sizes, the correlation between simulated coverage probability and theoretical coverage probability was at least 94%.




Fig. 2. Locations for observations (red circles) and predictions (blue crosses). One far away prediction location is omitted.

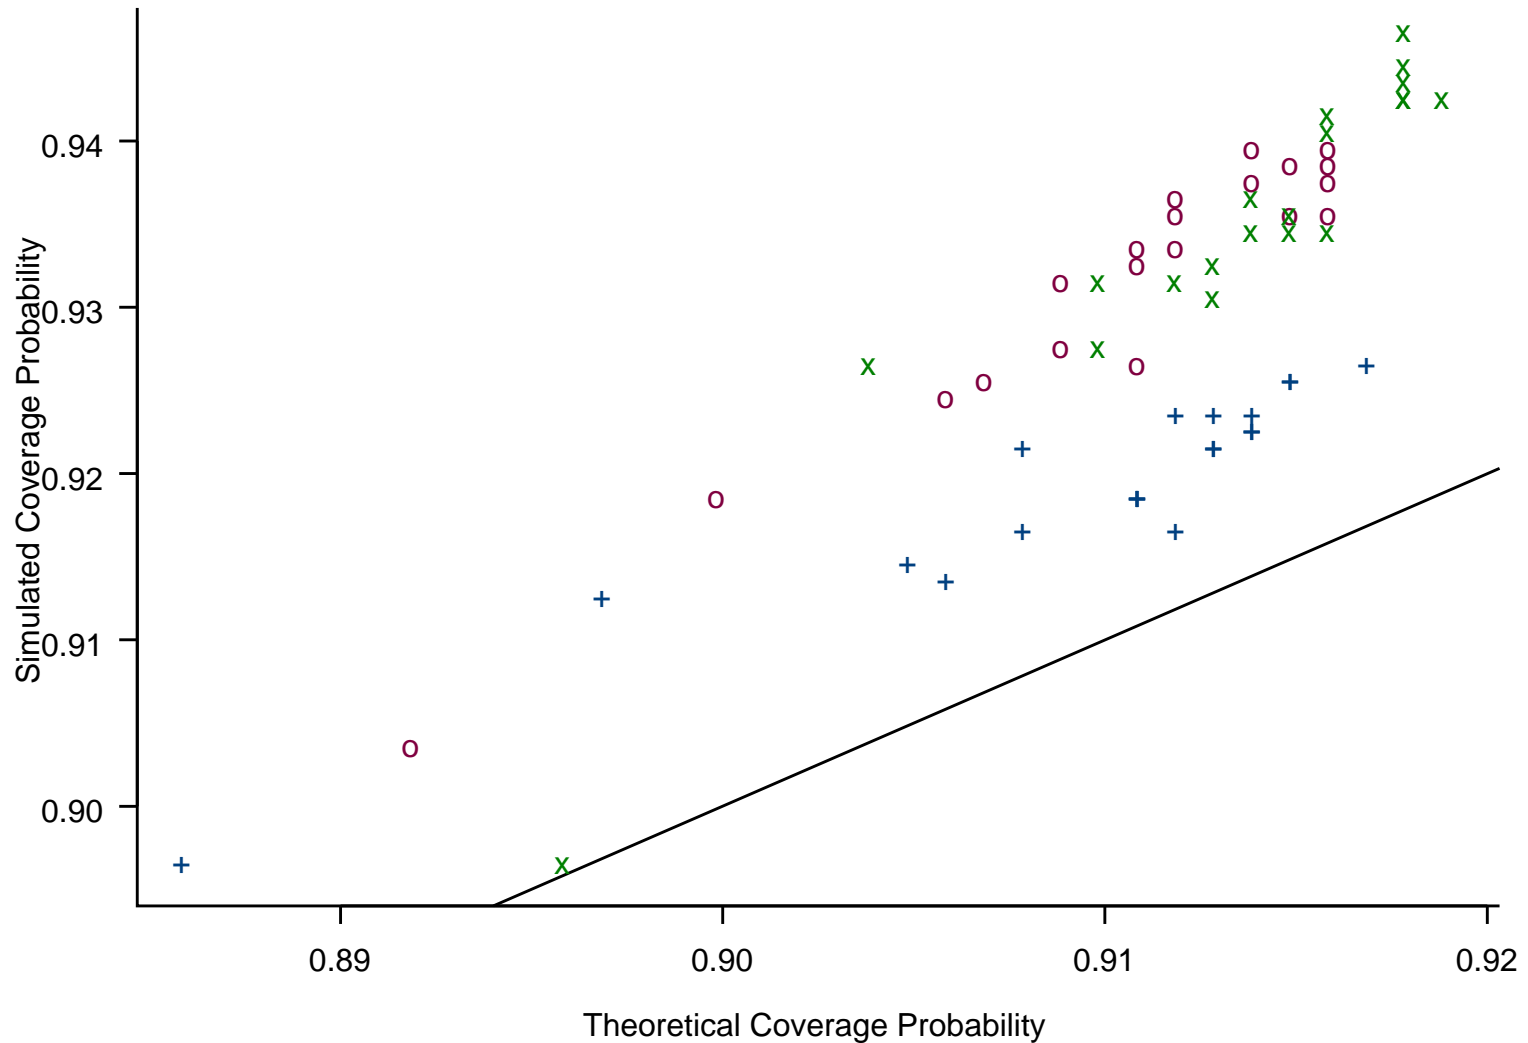


Fig. 3. Simulated v. theoretical coverage probabilities for sample sizes 20 (blue crosses), 35 (red circles) and 50 (green crosses). Based on 10000 simulations fitted by REML.

Example based on the 38 PM_{2.5} monitors in North Carolina

Assume the objective is to estimate population-weighted daily average. Daily data from 2000. Assume individual days' data are independent replications of the model

$$\text{Cov}(y_i, y_j) = \begin{cases} \theta_1^2 & \text{if } i = j, \\ \theta_3 \theta_1^2 e^{-d_{ij}/\theta_2} & \text{if } i \neq j, \end{cases}$$

with y_i, y_j the PM_{2.5} at locations i and j , d_{ij} is distance (units of 100 km.), and we estimated $\theta_1 = 6.495$, $\theta_2 = 4.019$, $\theta_3 = .9423$. Treat this as the true model, but assume $\theta_1, \theta_2, \theta_3$ would have to be re-estimated on any given day.

Population-weighted averages were calculated using data from the 2000 U.S. census for the 809 zip code tabulation areas (ZCTA) in North Carolina. Select 38 ZCTA out of 809 to place the monitoring station to give most accurate prediction of the total population PM2.5 exposure defined as

$$y_0 = \sum_i p_i y_i,$$

where p_i is the population at the i 'th ZCTA, and y_i is the PM2.5 level there. V_1 and V_4 with two-sided tail probabilities $P = 0.1, 0.01, 0.001$ are used as design criteria, and a simulated annealing algorithm is used to find the designs given in Figure 4.

Optimal Designs Under Four Criteria

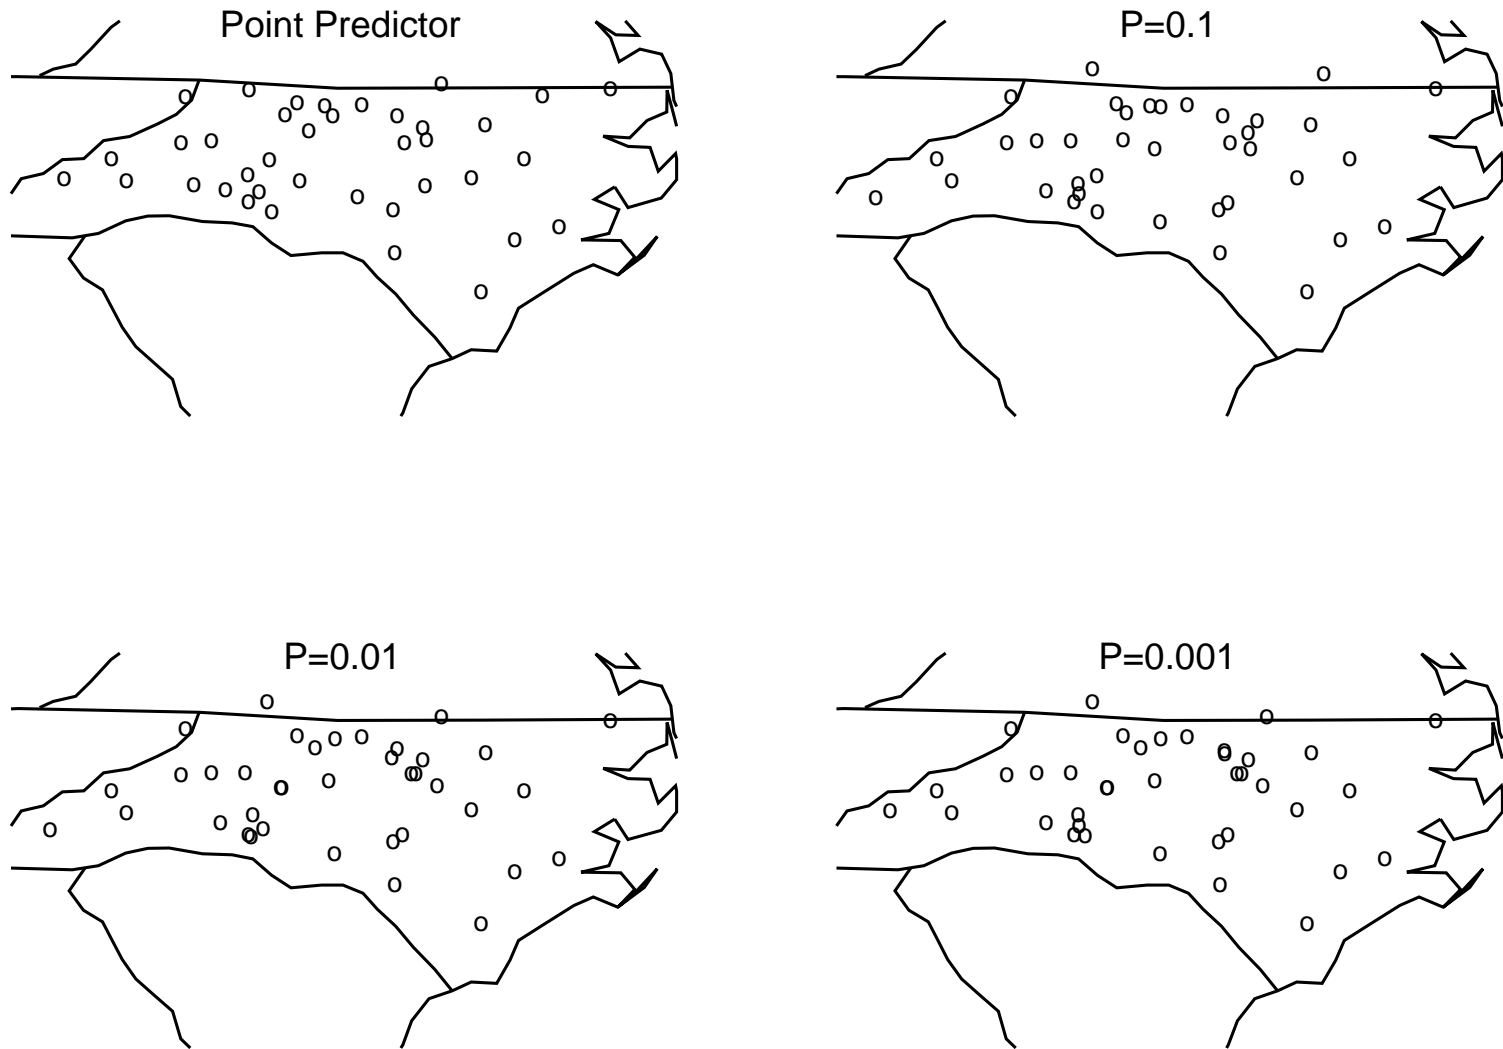


Fig. 4. Four designs selected using criteria of this talk (calculations due to Zhengyuan Zhu)

All four designs tend to place monitors in regions of high population density (as does the current EPA network, Fig. 1) but it is noticeable that the criterion V_4 , especially for smaller P , tends to favor a network with clusters of nearby monitors, reflecting the role such clusters play in ensuring good estimation of model parameters.

Finally we show the actual EPA map of nonattainment areas.

Attainment and Nonattainment Areas in the U.S. PM_{2.5} Standards

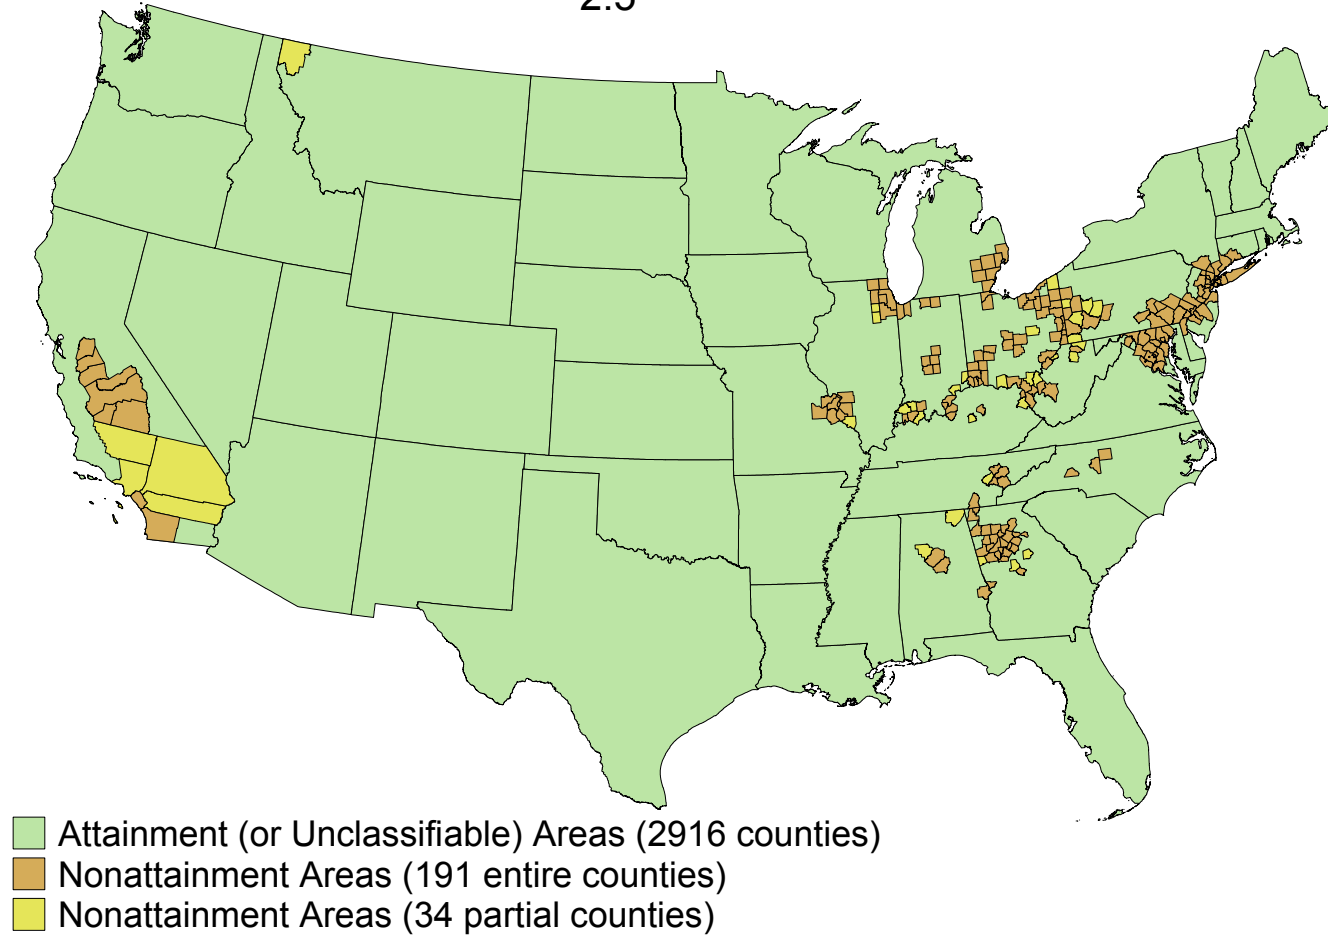


Fig. 5. Current nonattainment areas (Source: EPA website, 12/18/2004).