

1 Introduction

An important step in modeling a spatial processes is choosing the form of the covariance function. The form of the covariance function will have an effect on kriging as well as parameter estimates and the associated uncertainty. Misspecification of the second order properties can lead to misleading inference. Sherman (2011, pg. 87-90) and Guan et al. (2004) summarize the implications of incorrectly specifying isotropy properties through numerical examples. In order to choose an appropriate model, a statistician must first assess the nature of the spatial variation of his or her data. To this end, a number of graphical techniques are available, such as directional sample variograms or rose diagrams to check for anisotropy (Matheron, 1961; Isaaks and Srivastava, 1989, pg. 149-154). Spatial statisticians may have intuition about the interpretation and reliability of these diagnostics, but a less experienced user may desire a more objective evaluation in the form of a hypothesis test. In this paper, we use two real data examples to illustrate how the nonparametric hypothesis tests available in the R package **spTest** (Weller, 2015) can be used to assess isotropy properties in spatially-referenced data. The examples also demonstrate how graphical techniques and hypothesis tests can, and should, be used in a complementary role. The remainder of the paper is organized as follows: Section 2 establishes notation and definitions; Section 3 describes the functionality of the **spTest** package; Section 4 demonstrates how to use the functions in **spTest** in conjunction with graphical techniques on two different data sets; Section 5 concludes the paper with a discussion.

2 Notation

2.1 Spatial statistics

In Section 2.1 we briefly review key definitions required for describing and performing tests of isotropy. For additional background, see Weller and Hoeting (2015) or Schabenberger and Gotway (2004). Let $\{Y(\mathbf{s}) : \mathbf{s} \in \mathcal{D} \subseteq \mathbb{R}^2\}$ be a second order stationary random field (RF). Let $\{\mathbf{s}_1, \dots, \mathbf{s}_n\} \subset \mathcal{D}$ be the finite set of locations at which the random process is observed, providing the random vector $\mathbf{Y} = (Y(\mathbf{s}_1), \dots, Y(\mathbf{s}_n))^T$. The sampling locations may follow one of several spatial sampling designs, for example, gridded locations, randomly and uniformly distributed locations, a cluster design, or any other general design. Note that there is a distinction between a lattice process and a geostatistical process observed on a grid (e.g., Fuentes and Reich, 2010; Schabenberger and Gotway, 2004, pg. 6-10), but we do not explore this distinction and will use the term grid throughout.

For a spatial lag $\mathbf{h} = (h_1, h_2)$, the semivariogram function describes dependence between observations, \mathbf{Y} , at spatial locations separated by lag \mathbf{h} and is defined as

$$\gamma(\mathbf{h}) = \frac{1}{2} \text{Var}(Y(\mathbf{s} + \mathbf{h}) - Y(\mathbf{s})). \quad (2.1)$$

The classical moment-based estimator of the semivariogram (Matheron, 1962) is

$$\hat{\gamma}(\mathbf{h}) = \frac{1}{2|\mathcal{D}(\mathbf{h})|} \sum [Y(\mathbf{s}) - Y(\mathbf{s} + \mathbf{h})]^2, \quad (2.2)$$

where the sum is over $\mathcal{D}(\mathbf{h}) = \{\mathbf{s} : \mathbf{s}, \mathbf{s} + \mathbf{h} \in \mathcal{D}\}$ and $|\mathcal{D}(\mathbf{h})|$ is the number of elements in $\mathcal{D}(\mathbf{h})$. The set $\mathcal{D}(\mathbf{h})$ is the set of sampling location pairs that are separated by spatial lag \mathbf{h} . The covariance function is defined by

$$C(\mathbf{h}) = \text{Cov}(Y(\mathbf{s}), Y(\mathbf{s} + \mathbf{h})) \quad (2.3)$$

and is an alternative to the semivariogram for describing spatial dependence. When it exists, the covariance function, $C(\mathbf{h})$, is related to the semivariogram by $C(\mathbf{h}) = \lim_{\|\mathbf{v}\| \rightarrow \infty} \gamma(\mathbf{v}) - \gamma(\mathbf{h})$ if the limit exists.

It is often of interest to infer the effect of covariates on the process, deduce dependence structure, and/or predict Y with quantifiable uncertainty at new locations. To achieve these goals, the practitioner must specify the distributional properties of the spatial process. A common assumption is that the finite-dimensional joint distribution is multivariate normal (MVN), in which case we call the RF a Gaussian random field (GRF). Since we are interested in the second order properties, hereafter we assume that the mean of the RF is zero.

A common simplifying assumption about the spatial dependence structure is that it is isotropic.

Definition 2.1.1. *A second-order stationary spatial process is isotropic if the semivariogram, $\gamma(\mathbf{h})$, [or equivalently, the covariance function $C(\mathbf{h})$] of the spatial process depends on the lag vector \mathbf{h} only through its Euclidean length, $\|\mathbf{h}\|$, i.e., $\gamma(\mathbf{h}) = \gamma_0(\|\mathbf{h}\|)$ for some function $\gamma_0(\cdot)$ of a univariate argument.*

Isotropy implies that the dependence between any two observations depends only on the Euclidean distance between their sampling locations and not on their relative orientation. A spatial process that is not isotropic is called anisotropic. The methods in **spTest** are designed to assist in determining whether or not the assumption of isotropy holds. Namely, the functions in **spTest** implement non-parametric hypothesis tests of isotropy, which avoid the choice of a parametric form for the covariance (semivariogram) function.

There are two important modifications to the estimator in Equation 2.2 that are pertinent to the methods described in this paper. First, for non-gridded sampling locations, the estimator needs to be modified to account for the fact that very few pairs of locations will be separated by a specific spatial lag, \mathbf{h} . One solution to this challenge is to specify a distance tolerance, ϵ , such that lags having length $\|\mathbf{h}\| \pm \epsilon$ are included in estimating the semivariance at lag \mathbf{h} . Second, directional sample semivariograms can be estimated by only using observations that are separated by spatial lags in a specific direction. For example, to investigate potential anisotropy, we can compare sample semivariograms between the horizontal and vertical directions. For non-gridded sampling locations, very few pairs of locations will lie at a specific distance and directional lag, so we need to allow for both a distance and a directional tolerance when estimating the semivariance. A common method for doing this is by using a product kernel smoother that smoothes over both the horizontal (h_1) and vertical (h_2) components of the spatial lag $\mathbf{h} = (h_1, h_2)$.

Spatial RFs and their second order properties can also be expressed in the spectral (or frequency) domain using Fourier transforms. The spectral representation of RFs and their second order properties provides alternative methods for testing second order properties. For brevity we focus only on the methods in the spatial domain and refer the interested reader to Weller and Hoeting (2015)

or Fuentes and Reich (2010). We note that that, in addition to the methods from the spatial domain, the spectral methods from Lu and Zimmerman (2005) are also implemented in **spTest**.

2.2 Nonparametric tests of isotropy

Lu (1994) and Lu and Zimmerman (2001) pioneered a popular approach to testing second-order properties when they used the asymptotic joint normality of the sample semivariogram computed at different spatial lags to evaluate symmetry and isotropy properties. The subsequent works of Guan et al. (2004, 2007) and Maity and Sherman (2012) built upon these ideas and are the primary methods programmed in **spTest**. Here we give an overview of the tests in Guan et al. (2004) and Maity and Sherman (2012).

Under the null hypothesis that the RF is isotropic, it follows that the values of $\gamma(\cdot)$ evaluated at any two spatial lags that have the same norm are equal, independent of the direction of the lags. For example, under the assumption of isotropy, $\gamma((-3, 0)) = \gamma((\sqrt{3}, \sqrt{3}))$. To completely specify the null hypothesis of isotropy, theoretically, one would need to compare semivariogram values for an infinite set of lags. In practice a small number of lags are specified. Then it is possible to test a hypothesis consisting of a set of linear contrasts of the form

$$H_0 : \mathbf{A}\gamma(\cdot) = \mathbf{0} \quad (2.4)$$

as a proxy for the null hypothesis of isotropy, where \mathbf{A} is a full row rank matrix (Lu and Zimmerman, 2001). For example, a set of lags, denoted $\mathbf{\Lambda}$, commonly used in practice on gridded sampling locations with unit spacing is

$$\mathbf{\Lambda} = \{\mathbf{h}_1 = (1, 0), \mathbf{h}_2 = (0, 1), \mathbf{h}_3 = (1, 1), \mathbf{h}_4 = (-1, 1)\}, \quad (2.5)$$

and the corresponding \mathbf{A} matrix under $H_0 : \mathbf{A}\gamma(\mathbf{\Lambda}) = \mathbf{0}$ is

$$\mathbf{A} = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix}. \quad (2.6)$$

In other words, by using Equations 2.5 and 2.6 under the hypothesis 2.4, we contrast the semivariogram values at lags $\mathbf{h}_1 = (1, 0)$ and $\mathbf{h}_2 = (0, 1)$, and likewise, contrast semivariogram values at lags $\mathbf{h}_3 = (1, 1)$ and $\mathbf{h}_4 = (-1, 1)$. An important step in detecting anisotropy is the choice of lags, $\mathbf{\Lambda}$, as the test results will only hold for the set of lags considered (Guan et al., 2004). While this choice is somewhat subjective, there are several considerations for determining the set of lags. First, in terms of Euclidean distance, short lags should be chosen since estimates of the semivariance at long lags will be less reliable than estimates at short lags because they are based on fewer pairs of observations and hence more variable. Second, pairs of orthogonal lags should be contrasted since, for an anisotropic process, the directions of strongest and weakest spatial correlation will typically be orthogonal. For other considerations and more detailed guidelines regarding the choice of lags, see Weller and Hoeting (2015), Lu and Zimmerman (2001), and Guan et al. (2004).

The tests in Guan et al. (2004) and Maity and Sherman (2012) involve estimating either the semivariogram 2.1 or covariogram 2.3 and evaluating the estimator at the set of chosen lags. Denoting the set of point estimates of the

variogram/covariogram at the chosen lags as $\hat{\mathbf{G}}_n$, the true values as \mathbf{G} , and normalizing constant a_n , a central result for all three methods is that

$$a_n(\hat{\mathbf{G}}_n - \mathbf{G}) \xrightarrow[n \rightarrow \infty]{d} MVN(\mathbf{0}, \mathbf{\Sigma}), \quad (2.7)$$

under increasing domain asymptotics and mild moment and mixing conditions on the RF. The test statistic is a quadratic form

$$TS = b_n^2(\mathbf{A}\hat{\mathbf{G}}_n)^\top (\mathbf{A}\hat{\mathbf{\Sigma}}\mathbf{A}^\top)^{-1}(\mathbf{A}\hat{\mathbf{G}}_n), \quad (2.8)$$

where $\hat{\mathbf{\Sigma}}$ is an estimate of the asymptotic variance-covariance matrix and b_n is a normalizing constant. A p value can be obtained from the asymptotic χ^2 distribution with degrees of freedom given by the row rank of \mathbf{A} . Important differences in these works regard the distribution of sampling locations, shape of the sampling domain, estimation of \mathbf{G} and $\mathbf{\Sigma}$, and assumptions on the dependence structure of the random field (see Weller and Hoeting (2015) for more details).

3 Nonparametric tests implemented in spTest

The R package **spTest** includes functions for implementing the tests developed in Guan et al. (2004), Lu and Zimmerman (2005), and Maity and Sherman (2012). The primary functions in **spTest** for implementing these tests are **LuTest**, **MaityTest**, **GuanTestGrid**, and **GuanTestUnif**. For example, the test from Guan et al. (2004) for data observed at non-gridded, but uniformly distributed, sampling locations is implemented in the function **GuanTestUnif**, which takes the following arguments:

```
GuanTestUnif(spdata, lagmat, A, df, h = 0.7, kernel = "norm",
  truncation = 1.5, xlims, ylims, grid.spacing = c(1, 1),
  window.dims = c(2, 2), subblock.h, sig.est.finite = T).
```

There are several necessary inputs. The matrix **spdata** includes the coordinates of sampling locations and the corresponding data values. The spatial lags that will be used to estimate the semivariance, denoted \mathbf{A} , are specified in the matrix **lagmat**. The matrix **A** provides the contrasts of the semivariance estimates, and its row rank is indicated by the parameter **df** (the degrees of freedom for the asymptotic χ^2 distribution). The values **h** and **kernel** provide the bandwidth (smoothing) parameter and form of the kernel smoother, respectively, used to smooth over spatial lags when estimating the semivariance. If a normal smoothing kernel is used, then the **truncation** parameter indicates where to truncate the normal kernel (i.e., zero weight for values larger than this value). The parameters **xlims** and **ylims** give the horizontal and vertical limits of the sampling region (a rectangular sampling region is assumed). A grid is laid over the sampling region according the width and height specified by **grid.spacing**. The dimensions of the moving windows, given in the units of the underlying grid, are determined by the values in **window.dims**. The bandwidth used to estimate the semivariance on the subblocks of data is indicated by **subblock.h** and a finite sample adjustment to the estimate of the asymptotic variance-covariance matrix is made by setting **sig.est.finte**

= T. For more information about the different arguments and guidelines on how to choose them, see Weller and Hoeting (2015), the **spTest** manual (<https://cran.r-project.org/web/packages/spTest/index.html>), and the original works.

4 Applications: Using **spTest** to check for anisotropy

We demonstrate the functions in **spTest** on two data sets: the first containing gridded sampling locations, the second non-gridded sampling locations. For more details on the functions and more examples using simulated data, see the official **spTest** manual. The **spTest** package can be used independently of other packages built for analyzing spatial data, but it works nicely with two other packages loaded into R: **fields** (Nychka et al., 2014) and **geoR** (Ribeiro Jr. and Diggle, 2001). The package **fields** is automatically installed with **spTest**, while **geoR** is not. We also load the **splines** package (R Core Team, 2015), which we'll use to fit mean functions.

```
> library("spTest")
> library("fields")
> library("geoR")
> library("splines")
```

For the two examples given below, we'll use graphical diagnostics and the hypothesis tests implemented in **spTest** to determine whether or not an assumption of isotropy is reasonable for spatially-referenced data. The general strategy will be to first do exploratory data analysis (EDA) of the original data and create a model for the mean of the spatial process using appropriate covariates. After estimating a model for the mean, we'll extract residuals and again use EDA to check for remaining spatial dependence and utilize graphical diagnostics and hypothesis tests to investigate potential anisotropy. We note that, for brevity, we have not included the full version of EDA code and plots; instead, we include only the most relevant to demonstrating the functionality of the **spTest** package. The complete version of the code is available on [github](#).

4.1 Gridded sampling locations

The gridded data used in this section come from the North American Regional Climate Change Assessment Program [NARCCAP] (Mearns et al., 2009). The data set **WRFG** in **spTest** includes coordinates and a 24-year average of yearly average temperatures from runs of the Weather Research and Forecasting - Grell configuration (WRFG) regional climate model (RCM) using boundary conditions from the National Centers for Environmental Prediction (NCEP). The original data and the R code used to create the yearly averages are available online. The data set contains both latitude and longitude coordinates and map projection coordinates that specify the regular grid for 14,606 grid boxes along with average temperature at each grid box. We can display a heat map of all of the data using the `image.plot` function from the **fields** package (see Figure 1a). Due to computational considerations and because the methods in **spTest** assume stationarity, for our analysis we'll use a 20×20 subset of the grid boxes over the central United States (see Figure 1b).




Figure 1: Heat maps of average temperatures of WRFG-NCEP output from NARCCAP. Figure 1a displays the average of WRFG-NCEP yearly average temperature over 24 years with coordinates in latitude and longitude. Figure 1b displays the same data but on the gridded map projection (easting/northing) coordinates, and the black points indicate the subset of the data used for this example.

```
> data("WRFG")
> image.plot(WRFG$lon - 360, WRFG$lat, WRFG$WRFG.NCEP.tas, xlab = "Longitude",
+           ylab = "Latitude", main = "Mean WRFG-NCEP Temperatures")
> world(add = T)

> coords <- expand.grid(WRFG$xc, WRFG$yc)
> sub <- which(coords[, 1] > 2950000 & coords[, 1] < 4e+06 & coords[,
+           2] > 1200000 & coords[, 2] < 2250000)
> image.plot(WRFG$xc, WRFG$yc, WRFG$WRFG.NCEP.tas, xlab = "Easting",
+           ylab = "Northing", main = "Map Projection Coordinates")
> points(coords[sub, ], pch = 16, cex = 0.35)
> my.coords <- coords[sub, ]
> colnames(my.coords) <- c("xc", "yc")
> tas <- WRFG$WRFG.NCEP.tas
> tas <- c(tas)
> tas <- tas[sub]
> mydata <- cbind(my.coords[, 1], my.coords[, 2], tas)
```

To investigate potential anisotropy in the relevant subset of these data, we can examine two graphical diagnostics: a heat map and directional sample semi-variograms. We'll use the function `variog4` from the **geoR** package to estimate directional semivariograms to visually assess isotropy properties.

```
> x.coord <- unique(my.coords[, 1])
> y.coord <- unique(my.coords[, 2])
> nx <- length(x.coord)
```

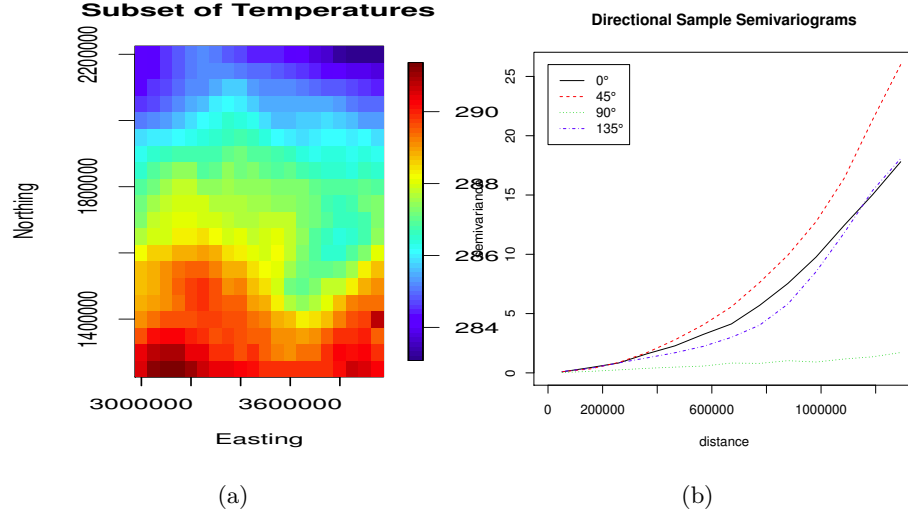


Figure 2: Graphical assessments of isotropy in the 20×20 subset of WRFG temperature data. Because northing (latitude) coordinates have not been accounted for, the heat map (2a) indicates that the dependence between observations is stronger in the east-west direction than the north-south direction. The directional dependence is also evidenced by the differences between the directional sample semivariograms (2b).

```
> ny <- length(y.coord)
> tas.mat <- matrix(mydata[, 3], nrow = nx, ncol = ny, byrow = F)
> par(mfrow = c(1, 2), oma = c(0, 0, 2, 0))
> image.plot(x.coord, y.coord, tas.mat, ylab = "Northing", xlab = "Easting",
+           main = "Subset of Temperatures")

> geodat <- as.geodata(mydata)
> svar4 <- variog4(geodat)
> plot(svar4)
> title("Directional Sample Semivariograms")
```

The heat map in Figure 2a indicates that the spatial process is anisotropic, having a stronger spatial dependence in the horizontal direction than the vertical direction. Intuitively, northing (latitude) is an important factor in determining average temperature, and we need to include its effect in a model for these data. We also notice non-linear trends in temperature as a function of easting (longitude) in Figure 2a. Thus, the anisotropy can be attributed, at least in part, to the fact that we have not modeled important covariates related to the process. The directional sample semivariograms in Figure 2b reaffirm the notion that the data exhibit anisotropy as the 90° sample semivariogram appears much different than the other three. Before modeling the effects of northing and easting coordinates, we'll use the `GuanTestGrid` function from **spTest** to affirm our understanding that these data exhibit anisotropy.

While some of the nonparametric methods for testing isotropy do not require Gaussian data, necessary conditions for the asymptotic properties of the test to hold are typically met when the data are Gaussian. A histogram (not shown) of the relevant subset of WRFG temperature data indicates that a Gaussian assumption is not unreasonable. To implement the nonparametric test in Guan et al. (2004) via the function `GuanTestGrid`, we need to specify the spatial lags that will be used to test for differences in the semivariogram. Based on the sample semivariograms in Figure 2b, we begin to see a difference in the semivariograms at a distance of approximately 200,000. Since the distance between sampling locations is 50,000, we should use spatial lags with lengths of approximately $200,000/50,000 = 4$ units. For this test we choose the lag set

$$\mathbf{A} = \{\mathbf{h}_1 = (4, 0), \mathbf{h}_2 = (0, 4), \mathbf{h}_3 = (3, 3), \mathbf{h}_4 = (-3, 3)\},$$

and we use the matrix \mathbf{A} in Equation 2.6. To create subblocks of data used to estimate $\mathbf{\Sigma}$, we choose a moving window with size 2×2 . The moving window dimensions should be chosen so that the window has the same shape (i.e., square or rectangle) and orientation as the entire data set. To maximize the amount of data used to estimate $\mathbf{\Sigma}$, the dimensions of the window should evenly divide the number of columns and rows, respectively, of the entire region. For this example there are 20 rows and columns; hence, window dimensions of 2×2 or 4×4 are ideal choices. Using a window size of 2×2 , we run the test via the following code.

```
> my.delta <- 50000
> mylags <- rbind(c(4, 0), c(0, 4), c(3, 3), c(-3, 3))
> myA <- rbind(c(1, -1, 0, 0), c(0, 0, 1, -1))
> tr <- GuanTestGrid(spdata = mydata, delta = my.delta, lagmat = mylags,
+   A = myA, window.dims = c(2, 2), pt.est.edge = TRUE, sig.est.edge = TRUE,
+   sig.est.finite = TRUE, df = 2)
> tr$pvalue.finite

[1] 0
```

As expected, the results of the hypothesis test (p value < 0.05) indicate that the data exhibit anisotropy. We note that the p value is given as exactly 0 due to a finite sample correction (Guan et al., 2004). The function `GuanTestGrid`, and other functions in `spTest`, return a p value for the test and information used in computing the p value, such as the point estimates ($\hat{\mathbf{G}}_n$), estimates of the asymptotic variance-covariance matrix ($\hat{\mathbf{\Sigma}}$), the number of subblocks used to estimate $\mathbf{\Sigma}$, and other information about the estimation process. Here we note that the point estimates for the directional semivariance are slightly different between the functions from the `spTest` and `geoR` packages due to different kernel methods used in estimation.

As mentioned earlier, we need to model the effects of northing and easting (latitude and longitude) coordinates on average temperature. We fit temperature as a nonparametric additive function of both the northing and easting coordinates via least-squares using cubic splines. The cubic splines can be specified using the function `ns` from the `splines` package and the least squares fit is computed via the `lm` function.

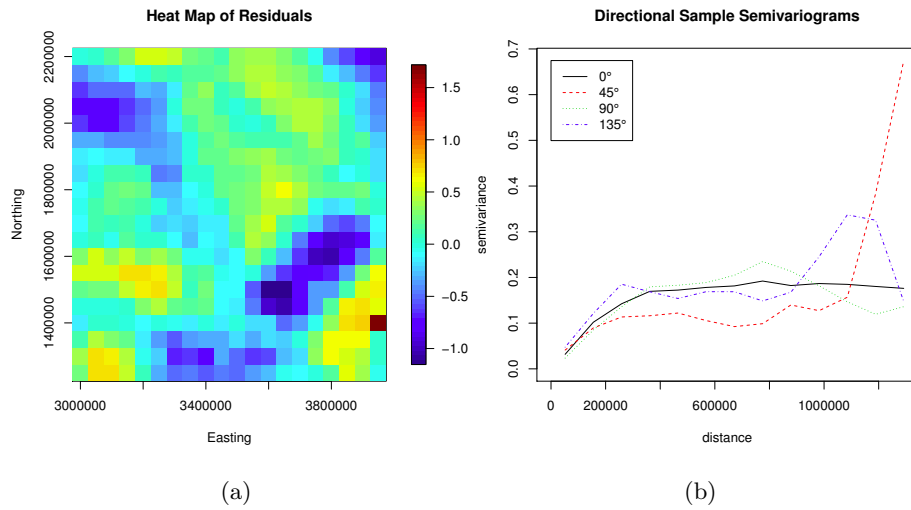


Figure 3: Graphical assessments of isotropy in the residuals. The appearance of elongated areas of similar residual values in the heat map (3a) indicates that the process may be anisotropic. The directional semivariograms (3b) do not appear to exhibit differences, indicating that the process is isotropic.

```
> x1 <- my.coords[, 1]
> x2 <- my.coords[, 2]
> m1 <- lm(tas ~ ns(x1, df = 3) + ns(x2, df = 3))
> summary(m1)
```

After removing the mean effects of the coordinates, we can check for any remaining (unaccounted for) spatial dependence and evidence of anisotropy in the residuals using graphical diagnostics and a hypothesis test. A histogram of the residuals (not shown) indicates that a Gaussian assumption is reasonable.

```
> resid <- m1$resid
> resid.mat <- matrix(resid, nrow = nx, ncol = ny, byrow = F)
> image.plot(x.coord, y.coord, resid.mat, xlab = "Easting", ylab = "Northing",
+           main = "Heat Map of Residuals")

> resid.dat <- cbind(mydata[, 1:2], resid)
> geodat.resid <- as.geodata(resid.dat)
> plot(variogram4(geodat.resid))
> title("Directional Sample Semivariograms")
```

The clusters of similar values in the heat map of Figure 3a, and the increase, followed by a leveling off, of the semivariance values as distance increases in the directional sample semivariograms in Figure 3b indicate that the residuals are still spatially dependent. However, the plots in Figure 3 do not clearly illustrate whether or not the residuals exhibit anisotropy. There appears to be directional dependence along the NW to SE direction in the northern parts of the heatmap (Figure 3a). The directional sample semivariograms do not appear to be different until the distance is greater than 200000. Semivariogram

estimates at large distances can be unreliable because there are fewer pairs of sampling locations to estimate this value than at short distances. Likewise, directional semivariograms are less reliable than a uni-directional semivariogram because fewer pairs of sampling locations are used at each distance for directional estimation. The unreliability of the semivariograms at the larger distances, coupled with the lack of a measure of uncertainty, make it difficult to determine whether or not an assumption of isotropy is reasonable. In hopes of gaining more insight into the isotropy properties, we perform a nonparametric hypothesis test of isotropy using the residuals with the same choices for Λ , \mathbf{A} , and the window dimensions.

```
> tr <- GuanTestGrid(spdata = resid.dat, delta = my.delta, lagmat = mylags,
+   A = myA, window.dims = c(2, 2), pt.est.edge = TRUE, sig.est.edge = TRUE,
+   sig.est.finite = TRUE, df = 2)
> tr$pvalue.finite

[1] 0.06111111
```


Here the residuals do not provide evidence for anisotropy (p value > 0.05). These results suggest that it may be appropriate to choose an isotropic covariance function to model the residuals. However, it is important to note that we have not included the effect of other potentially influential covariates such as elevation or water cover in the model for temperature. Additionally, although we examined a 20×20 subset of the data, the grid boxes still cover a large geographic region of the U.S., and thus an assumption of stationarity, which is needed for the asymptotic properties of the hypothesis test to hold, may not be reasonable.

4.2 Non-gridded sampling locations

The non-gridded data set used in this section describes monthly surface meteorology for the state of Colorado and comes from the National Center for Atmospheric Research (NCAR). The data are available in the package **fields** or online (<http://www.image.ucar.edu/GSP/Data/US.monthly.met/CO.html>). For this example, our variable of interest is the log of the 30-year average of average yearly precipitation at 344 station locations during the time period 1968-1997.

Like the temperature data, our goal will be to model the mean effect of covariates and check for spatial dependence and potential anisotropy in the residuals. Because the sampling locations cover a much smaller region than the subset of WRFG temperatures, we choose to use the latitude and longitude coordinates for this example. We can create a heat map of the log precipitation values and the elevation of the stations using the function `quilt.plot` from the **fields** package.

```
> data(COmonthlyMet)
> dim(CO.ppt)
> sub30 <- CO.ppt[74:103, , ]
> dim(sub30)
> nstations <- 376
> years <- 1968:1997
> nyears <- length(years)
```



(a) Station locations and $\log(\text{precip})$ (b) Station locations and elevation.

Figure 4: Heat maps showing the locations of the stations in a region of Colorado along with the log of average precipitation (4a) and elevation (4b) at each station.

```
> yr.avg <- matrix(data = NA, nrow = nstations, ncol = nyears)
> for (i in 1:nyears) {
+   yr.dat <- sub30[i, , ]
+   yr.avg[, i] <- apply(yr.dat, 2, mean, na.rm = T)
+ }
> avg30 <- apply(yr.avg, 1, mean, na.rm = T)
> quilt.plot(CO.loc, log(avg30), xlab = "Longitude", ylab = "Latitude",
+   main = "Quilt Plot of log(precip)")
> US(add = T)

> quilt.plot(CO.loc, CO.elev, xlab = "Longitude", ylab = "Latitude",
+   main = "Quilt Plot of Elevation")
> US(add = T)
```

Colorado has two distinct geographic regions: the mountainous region in the west and the plains region in the east. Figure 4b illustrates these two regions, and we can begin to notice a possible relationship between elevation and average precipitation. We explore the potential relationship between precipitation and elevation using scatter plots (see Figure 5a).

```
> plot(CO.elev, log(avg30), xlab = "Elevation", ylab = "log(precip)",
+   main = "Scatter of log(precip) vs. Elevation")
> m1 <- lm(log(avg30) ~ ns(CO.elev, df = 3))
> summary(m1)
> fits <- m1$fitted.values
> bad <- is.na(avg30)
> x <- CO.elev[which(!bad)]
> ord <- order(x)
> x <- sort(x)
```




Figure 5: The plot in Figure 5a displays the scatter plot and the estimate of the nonparametric relationship between log(precipitation) and elevation. Figure 5b displays a histogram of the residuals from the model relating log(precipitation) to elevation.

```
> fits <- fits[ord]
> lines(x, fits, lwd = 3, col = "red")

> resid <- m1$resid
> hist(resid, freq = F, main = "Histogram of Residuals")
> lines(density(resid), lwd = 2)
```

We fit a cubic smoothing spline via least squares to model the relationship between log(precipitation) and elevation. The estimate is shown in Figure 5a, and the histogram of residuals in Figure 5b indicates that a Gaussian assumption is reasonable. We will use the residuals from this model to check for remaining spatial dependence and potential anisotropy. We use `variog4` to estimate directional sample semivariograms.

```
> precip.resid <- cbind(CO.loc[which(!bad), ], CO.loc[which(!bad),
+   ], 2], resid)
> geodat <- as.geodata(precip.resid)
> svar4 <- variog4(geodat)
> plot(svar4, legend = F)
> legend("bottomright", legend = c(expression(0 * degree), expression(45 *
+   degree), expression(90 * degree), expression(135 * degree)),
+   col = 1:4, lty = 1:4)
> title("Directional Sample Semivariograms")
```

The increase, followed by a leveling off, of the semivariance values as distance increases in Figure 6 indicates that there is spatial dependence remaining in the data. We notice that the 0° semivariogram appears to be slightly different than the other three, but there is no attached measure of uncertainty. We again turn




Figure 6: Directional sample semivariograms using residuals from the model relating $\log(\text{precipitation})$ to elevation.

to a nonparametric hypothesis test of isotropy to assist in determining if an assumption of isotropy is reasonable.

There are two testing procedures for non-gridded available in **spTest**: Guan et al. (2004) and Maity and Sherman (2012). To choose between these two, we need to decide whether or not it is reasonable to assume that sampling locations are uniformly distributed on the sampling domain. The methods for non-gridded data from Guan et al. (2004) rely on the assumption that sampling locations are uniformly distributed while the Maity and Sherman (2012) methods can be used on any general sampling design. To check this assumption, we can turn to methods from the spatial point process literature to perform a test of complete spatial randomness (CSR), i.e., a uniform spatial distribution, for the sampling locations. Methods for testing CSR are available in the R package **spatstat** (Baddeley and Turner, 2004). For brevity, we do not include those results here, and we will proceed assuming that CSR holds for these sampling locations.

For both Guan et al. (2004) and Maity and Sherman (2012), we need to choose the lag set, $\mathbf{\Lambda}$, and the contrast matrix, \mathbf{A} . The large jumps and decrease in the semivariance values in Figure 6 indicate that semivariogram estimates become unreliable beyond a distance of 2; thus, we should choose lags having Euclidean distance less than this distance. We choose the lag set

$$\mathbf{\Lambda} = \{\mathbf{h}_1 = (3/4, 0), \mathbf{h}_2 = (0, 3/4), \mathbf{h}_3 = (3/4, 3/4), \mathbf{h}_4 = (-3/4, 3/4)\},$$

and again we use the matrix \mathbf{A} in Equation 2.6.

```
> mylags <- 0.75 * rbind(c(1, 0), c(0, 1), c(1, 1), c(-1, 1))
```

```
> myA <- rbind(c(1, -1, 0, 0), c(0, 0, 1, -1))
```

The next step is to implement the methods from Guan et al. (2004) and Maity and Sherman (2012) is to determine the size of the moving windows and the block size, respectively, that will be used in estimating the asymptotic variance-covariance matrix. To make this determination, we need to place a grid over the sampling region. The choice of grid will depend on the density and design of the sampling locations, but, ideally, should divide the region so that uniformly distributed sampling locations are approximately evenly divided between the cells. We recommend using visualizations of different grid choices to assist in choosing a grid. For our example, we set the limits of the sampling region to be slightly wider than the observed minimum and maximum coordinates in both the x and y directions, and we divide both the x -axis and y -axis into 12 parts. The result of this grid is plotted in Figure 7.

```
> quilt.plot(precip.resid[, 1:2], precip.resid[, 3], xlab = "Longitude",
+           ylab = "Latitude")
> title("Quilt Plot of Residuals and Grid Used for Subsampling")
> min(precip.resid[, 1])

[1] -109.483

> max(precip.resid[, 1])

[1] -101.02

> min(precip.resid[, 2])

[1] 36.512

> max(precip.resid[, 2])

[1] 41.467

> my.xlims <- c(-109.5, -101)
> my.ylims <- c(36.5, 41.5)
> xlen <- my.xlims[2] - my.xlims[1]
> ylen <- my.ylims[2] - my.ylims[1]
> my.grid.spacing <- c(xlen/12, ylen/12)
> xgrid <- seq(my.xlims[1], my.xlims[2], by = my.grid.spacing[1])
> ygrid <- seq(my.ylims[1], my.ylims[2], by = my.grid.spacing[2])
> abline(v = xgrid, lty = 2)
> abline(h = ygrid, lty = 2)
```

Now that we have a grid on the sampling region, for the test in Guan et al. (2004), we need to choose the dimensions of the moving window and the smoothing (bandwidth) parameters. The p value of the hypothesis test will be sensitive to the choice of moving window dimensions. For this example, the dimensions of the sampling region are $8.5^\circ \times 5^\circ$ (width \times height). Thus, we should create a rectangular moving window having a slightly larger width than height. Our window dimensions should also be compatible with the grid placed on the sampling region. Recalling that we placed a 12×12 grid on the region, a good

Quilt Plot of Residuals and Grid Used for Subsampling




Figure 7: The grid placed on the sampling region used for spatial resampling. Because the sampling locations are not gridded, it can be difficult to assess isotropy properties via a heat map.

choice for window dimensions is 4×3 since these values both evenly divide 12 and create the desired rectangular orientation. Another potential choice for window dimensions is 2×1 , but this may not be a good choice because very few sampling locations will be available in each subblock. The smoothing parameter should be chosen based on the number and density of the sampling locations with larger values inducing higher levels of smoothing, i.e., averaging over all directions. In our experience, smoothing parameter values between 0.25 and 0.9 tend to produce reasonable results; however, the p value of the hypothesis test will change with the bandwidth. For this example, we choose a bandwidth of 0.6 for smoothing over lags over the entire domain, and a bandwidth of 0.7 for smoothing over lags on the subblocks of data created by the moving window.

Likewise, for the test in Maity and Sherman (2012), we need to choose the block size used to implement the grid-based block bootstrap [GBBB] (Lahiri and Zhu, 2006). Like the moving windows, we use a block size of 4×3 . Finally, we need to choose the number of bootstrap resamples that will be used to estimate Σ . We recommend using at least 50 bootstrap samples; however, more than 200 bootstrap samples comes with an increase in computation cost. We choose 100 bootstrap samples for our example, and we note that the number of bootstraps does not affect the precision of the p value, which is computed via the asymptotic χ^2 distribution.

When moving window or block dimensions are not compatible with the number of rows or columns of gridded sampling locations or the dimensions of grid laid on the sampling region for non-gridded locations, the functions in **spTest** will print a warning message since they do not currently handle partial (incom-

plete) blocks. Likewise, if the window dimensions are too small, the function `GuanTestUnif` will discard subblocks that do not have enough sampling locations and print a warning message. The problem of choosing an optimal bandwidth and optimal window/block size remains an open question. See Weller and Hoeting (2015) and the original works for guidelines on choosing these values.

Having determined values for the different options, we can now run the hypothesis tests.

```
> myh <- 0.6
> myh.sb <- 0.7
> tr.guan <- GuanTestUnif(spdata = precip.resid, lagmat = mylags,
+   A = myA, df = 2, h = myh, kernel = "norm", truncation = 1.5,
+   xlims = my.xlims, ylims = my.ylims, grid.spacing = my.grid.spacing,
+   window.dims = c(4, 3), subblock.h = myh.sb)
> tr.guan$pvalue.chisq

[1] 0.2267888

> tr.maity <- MaityTest(spdata = precip.resid, lagmat = mylags,
+   A = myA, df = 2, xlims = my.xlims, ylims = my.ylims, grid.spacing = my.grid.spacing,
+   block.dims = c(4, 3), nBoot = 100, kappa = 1)
> tr.maity$pvalue.chisq

[1] 0.1945388
```

For both tests, the residuals do not provide evidence in favor of anisotropy. Thus, in developing a spatial model for the residuals, it may be reasonable to assume isotropy.

5 Discussion

Choosing a covariance function is an important step in modeling spatially-referenced data and a variety of choices for the covariance function are available, e.g., anisotropy, nonstationarity, parametric forms. The R package **spTest** implements several nonparametric tests for checking isotropy properties which avoid specifying a parametric form for the covariance function. One concern regarding the methods in **spTest** is that they tend to have low power when the anisotropy is weak (see, e.g., Weller and Hoeting, 2015).

We have demonstrated how graphical techniques and the functions in **spTest** can be used in a complementary role to check for anisotropy. Future work includes extending the functionality of **spTest** to allow implementing the tests on non-rectangular sampling domains. Additionally, computational efficiency can be improved by programming functions in C++.

Acknowledgements

Weller's work was supported by the National Science Foundation Research Network on Statistics in the Atmospheric and Ocean Sciences (STATMOS) through grant DMS-1106862. We wish to thank the North American Regional Climate Change Assessment Program (NARCCAP) for providing the data used in this

paper. NARCCAP is funded by the National Science Foundation (NSF), the U.S. Department of Energy (DoE), the National Oceanic and Atmospheric Administration (NOAA), and the U.S. Environmental Protection Agency Office of Research and Development (EPA).

References

- Baddeley, A. J. and Turner, R. (2004). Spatstat: An r package for analyzing spatial point patterns.
- Fuentes, M. and Reich, B. (2010). Spectral domain. *Handbook of Spatial Statistics*, pages 57–77.
- Guan, Y., Sherman, M., and Calvin, J. A. (2004). A nonparametric test for spatial isotropy using subsampling. *Journal of the American Statistical Association*, 99(467):810–821.
- Guan, Y., Sherman, M., and Calvin, J. A. (2007). On asymptotic properties of the mark variogram estimator of a marked point process. *Journal of statistical planning and inference*, 137(1):148–161.
- Isaaks, E. H. and Srivastava, R. M. (1989). *Applied geostatistics*, volume 2. Oxford University Press New York.
- Lahiri, S. and Zhu, J. (2006). Resampling methods for spatial regression models under a class of stochastic designs. *The Annals of Statistics*, 34(4):1774–1813.
- Lu, H. and Zimmerman, D. (2001). Testing for isotropy and other directional symmetry properties of spatial correlation. *preprint*.
- Lu, H.-C. (1994). *On the distributions of the sample covariogram and semi-variogram and their use in testing for isotropy*. PhD thesis, University of Iowa.
- Lu, N. and Zimmerman, D. L. (2005). Testing for directional symmetry in spatial dependence using the periodogram. *Journal of Statistical Planning and Inference*, 129(1):369–385.
- Maity, A. and Sherman, M. (2012). Testing for spatial isotropy under general designs. *Journal of statistical planning and inference*, 142(5):1081–1091.
- Matheron, G. (1961). Precision of exploring a stratified formation by boreholes with rigid spacing-application to a bauxite deposit. In *International Symposium of Mining Research, University of Missouri*, volume 1, pages 407–22.
- Matheron, G. (1962). *Traité de géostatistique appliquée*. Editions Technip.
- Mearns, L. O., Gutowski, W., Jones, R., Leung, R., McGinnis, S., Nunes, A., and Qian, Y. (2009). A regional climate change assessment program for north america. *Eos, Transactions American Geophysical Union*, 90(36):311–311.
- Nychka, D., Furrer, R., and Sain, S. (2014). *fields: Tools for spatial data*. R package version 7.1.
- R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ribeiro Jr., P. and Diggle, P. (2001). geoR: a package for geostatistical analysis. *R-NEWS*, 1(2):15–18.

- Schabenberger, O. and Gotway, C. A. (2004). *Statistical methods for spatial data analysis*. CRC Press.
- Sherman, M. (2011). *Spatial statistics and spatio-temporal data: covariance functions and directional properties*. John Wiley & Sons.
- Weller, Z. D. (2015). *spTest: Nonparametric Hypothesis Tests of Isotropy and Symmetry*. R package version 0.2.0.
- Weller, Z. D. and Hoeting, J. (2015). A review of nonparametric hypothesis tests of isotropy properties in spatial data. *submitted to Statistical Science, available on arXiv*.