

UC Berkeley

UC Berkeley Previously Published Works

Title

A local nearest-neighbor convex-hull construction of home ranges and utilization distributions

Permalink

<https://escholarship.org/uc/item/2dd9c1mh>

Journal

Ecography, 27(4)

ISSN

0906-7590

Authors

Getz, Wayne M
Wilmers, C C

Publication Date

2004-08-01

Peer reviewed

A local nearest-neighbor convex-hull construction of home ranges and utilization distributions

Wayne M. Getz and Christopher C. Wilmers

Getz, W. M. and Wilmers, C. C. 2004. A local nearest-neighbor convex-hull construction of home ranges and utilization distributions. – *Ecography* 27: 489–505.

We describe a new method for estimating the area of home ranges and constructing utilization distributions (UDs) from spatial data. We compare our method with bivariate kernel and α -hull methods, using both randomly distributed and highly aggregated data to test the accuracy of area estimates and UD isopleth construction. The data variously contain holes, corners, and corridors linking high use areas. Our method is based on taking the union of the minimum convex polygons (MCP) associated with the $k - 1$ nearest neighbors of each point in the data and, as such, has one free parameter k . We propose a “minimum spurious hole covering” (MSHC) rule for selecting k and interpret its application in terms of type I and type II statistical errors. Our MSHC rule provides estimates within 12% of true area values for all 5 data sets, while kernel methods are worse in all cases: in one case overestimating area by a factor of 10 and in another case underestimating area by a factor of 50. Our method also constructs much better estimates for the density isopleths of the UD than kernel methods. The α -hull method does not lead directly to the construction of isopleths and also does not always include all points in the constructed home range. Finally we demonstrate that kernel methods, unlike our method and the α -hull method, does not converge to the true area represented by the data as the number of data points increase.

W. M. Getz (getz@nature.berkeley.edu) and C. C. Wilmers, Dept of Environmental Sciences, Policy and Management, Univ. of California Berkeley, CA 94720, USA.

The construction of space use maps from points representing distributions of animals or plants in space or time are critical in addressing a range of questions in ecology from the behavioral to the landscape level. Ecologists are generally interested in building two types of such maps: home range maps (Burt 1943) that delineate the spatial extent or outside boundary of an animals movement, and utilization distributions (UDs) (Jennrich and Turner 1969, Ford and Krumme 1979) that represent the density of space used by animals.

The simplest method for constructing home ranges is the minimum convex polygon (MCP). This method is still widely employed (Meulman and Klomp 1999, Baker 2001, Creel and Creel 2002, Rurik and Macdonald 2003) despite recent recognition that it provides an extremely

poor fit to data when the home range of an animal or the distribution of a population is strongly non-convex (Burgman and Fox 2003). In search of a better method, Burgman and Fox (2003) propose using α -hull constructions, which involve producing Delauney triangulations of the data and then removing all sides that are α times longer than the median of the original sides. Like the MCP, this method does not explicitly reveal high and low density use areas or clusters of points in cores. Also, in applications to real data, α -hull constructions leave some points hanging outside the area they bound, resulting in area estimates of home ranges that are often too conservative.

In the context of statistical errors, a home range or UD map can be regarded as a hypothesis about the

Accepted 9 February 2004

Copyright © ECOGRAPHY 2004
ISSN 0906-7590

expected space use of an organism and is subject to both type I (excluding valid areas) and type II (including invalid areas) errors. In methods with adjustable parameters one can trade-off these errors, where the optimal trade depends on the consequence of each type of error. Thus, if one is looking at the association of the UD of an animal population with background vegetation types, the balance may be tipped in favor of type I over type II errors (i.e. reducing false associations of animals with particular vegetation types). On the other hand, if one is looking for areas in a landscape that contain hidden factors causing some disease in a population, then type II errors may be more serious than type I errors in generating a list of putative factors common to all areas (i.e. we would not want to omit any area that might hold some clue to the cause of the disease).

Obviously, a method that produces both smaller type I and type II errors than another is preferred, provided the method is not computationally difficult to implement. Various user-friendly software packages, particularly those that include spatial mapping utilities, are available for implementing some of the more computationally challenging methods (e.g. see Lawson and Rodgers, 1997 for a comparison of five packages), such as kernel methods (Silverman 1986, Worton 1989). Kernel methods construct UDs by taking weighted sums of local parametric distributions (e.g. bivariate normal kernels) centered on each point in the data set being modeled (Silverman 1986, Worton 1989). Hence they perform well in constructing multimodal UDs for data generated as the sum of several bivariate normal distributions (Seaman and Powell 1996). The simplest of the kernel methods is the fixed method: it uses the same smoothing parameter value h at each point (this value determines the relative peakedness of the local distributions). A "best" value for h can be found by minimizing the mean-integrated-square-error of the UD fitted to the data as a function of h (Worton 1989). Adaptive kernel methods require additional computations to implement: they modify the value of h from point to point, based on local densities of points. In theory, adaptive methods should perform even better than fixed methods in characterizing the tails of the UD, but in practice this is not always true. Also, kernel methods are known to perform rather poorly when it comes to estimating areas of home ranges from data (Lawson and Rodgers 1997, Casaer et al. 1999, Ostro et al. 1999).

In this paper, we demonstrate that kernel methods perform poorly when fitted to distributions arising in landscapes that have distinct boundaries determined by geographic or physiographic features such as cliffs, rivers, or abrupt changes in soil types leading to abrupt changes in vegetation or other ecological determinants. Our method performs much better than kernel methods

in fitting UDs to home ranges with distinct boundaries and better than the α -hull methods in incorporating all points into the home range. We do not compare our method to grid or rectangular methods because these methods appear to have no advantages over ours. In particular, they are sensitive to the size of the underlying paving units and they use ad-hoc criteria to fill in holes after paving has been completed (e.g. Ostro et al. 1999 propose filling all holes $\leq 1\%$ of the area of associated MCP with the data – also see Plotnick et al. 1993, 1996, Dale et al. 2002). We have also not compared our method to those based on spatial statistics (such as the "local index of correlation association" and "spatial analysis by distance indices" reviewed by Dale et al. 2002) or on cluster analysis (e.g. Kenward et al. 2001, Plotkin et al. 2002) because the implementation of these is more complicated even than adaptive kernel methods, and they have not been widely applied.

Our method is direct and easily implemented. It involves constructing a UD from the union of convex hulls associated with each point and its $k-1$ nearest neighbors. We refer to this union as a k -NNCH (k nearest neighbor convex hull) covering, while the subcovering obtained from a union of the smallest of these convex hulls covering $x\%$ of points provides for the construction of the $x\%$ isopleth (e.g. the decile isopleths: 10%, 20%...100%). In the first part of this paper, we describe the method and then use it to map the UDs associated with computer-generated data that has sharp boundaries, multinuclear cores (cf. Kenward et al. 2001) and corridors. We then demonstrate that our algorithm performs better than kernel methods in identifying these features and in estimating area. Finally, we discuss where our method is superior to the α -hull method in constructing home ranges and utilization distributions (as characterized by the isopleths associated the density of points used to construct home ranges).

Methods

A k -NNCH covering for constructing UDs

Given a set of specified points the method begins by constructing the convex hull associated with each point and its $(k-1)$ nearest neighbors. We refer to the area covered by the union of all these convex hulls as a k -NNCH covering. We then order the hulls from the smallest to the largest. By progressively taking the union of these from the smallest upwards, until $x\%$ of points are included (with some rounding error), we construct the areas whose boundaries represent the $x\%$ isopleth of the densest set of points in our k -NNCH covering. (See Appendix for technical details.)

Kernel methods

Both fixed and adaptive bivariate normal kernel methods were coded in MATLAB using algorithms described in Worton (1989). MATLAB routines were then used to draw isopleths at the $p\%$ of the kernel density function. We used the $p = 0.99$ as the outer boundary for area calculations, although some studies rather use $p = 0.95$. This choice does not affect our conclusions regarding the poor estimation performance of kernel methods because kernel methods both greatly over and under estimate the areas involved. See Appendix for details regarding implementation of the reference or the least-squares cross-validated smoothing parameter values h_{REF} (which is optimal when the data are bivariate – see Silverman 1986) and h_{LSCV} respectively for both the fixed and adaptive kernel methods.

The α -hull method for constructing UD

Following the method of Burgman and Fox (2003), we constructed a Delauney tessellation to bound the data. We then calculated the mean length of all connections in this tessellation and removed those that were α times greater than this mean for specific values of α . Finally, we added up the area of the remaining triangles to obtain our estimate of the area. Based on Burgman and Fox's (2003) finding that $\alpha = 3$ is the most robust integer value of α with regard to sampling artifacts, we focused our analysis on this value and, for purposes of comparison, on twice this value ($\alpha = 6$). We also explored other values of α to get a sense of how the α -hull method performs as a function of α . Currently, no rule (such as the MSHC rule we propose below for selecting k for our algorithm) has been proposed for selecting an appropriate or "best" value for α : a value that is bound to differ for different sets of data.

Computer-generated data sets

We generated the five data sets below using Monte Carlo methods (Ripley 1987). The data are designed to test how well the methods perform at different ends of the data spectrum (random vs highly aggregated data), on contrasting shapes (donuts, squares, and multicore constructs), and identifying high use and odd-shaped boundaries (e.g. edges of lakes or land used on only one side of a the confluence of a river and one of its tributaries). Specifically, our idealized data sets are: 1) Random square (RS) (Fig. 6A): 1089 points were placed at random on the unit square. (Area = 1 arbitrary unit.) 2) Aggregated square (AS) (Fig. 5A): 1089 points were randomly assigned $x - y$ coordinates on the unit square. These coordinates were then cubed leading to increasingly higher densities of points having lower (x, y)

values (i.e. strongly aggregating around the axes, especially the origin). (Area = 1 arbitrary unit). 3) Random donut (RD) (Fig. 4A): We distributed 1089 points at random on a donut that has an inner radius of 1 and an outer radius of 5. The radius of each point was obtained from the equation $r = 1 + 4\sqrt{\xi}$ where ξ is a random variable rectangularly distributed on $[0, 1]$ and an angle between 0 and 2π was assigned at random. (Area = 75.4 arbitrary units). 4) Aggregated donut (AD) (Fig. 1A): We distributed 1089 points, as in 3 above, except in this case we used the formula $r = 1 + 4\xi^3$. This results in an extremely strong clustering around the inner boundary of the donut. (Area = 75.4 arbitrary units). 5) Multicore (MC) (Fig. 8A): This data set was constructed by placing less dense versions of the above 4 data sets at corners of a 25×25 unit quadrant and then connecting them with corridors. (Area is approximately 320 arbitrary – see Appendix for details.)

The MSHC rule for selecting k

For relatively low values of k the resulting k -NNCH coverings contain a number of holes that disappear with increasing k . For areas with known topologies (squares, donuts, etc.) the "minimum spurious hole covering" (MSHC) rule is to select the smallest value of k that produces a covering that has the same topology as the given set. If the topology of the space associated with the data is not known, we can guess its genus (number of holes) by identifying relatively large physical features, such as lakes, mountain peaks, or inhospitable habitats at comparable scales. We expect these objects to produce real holes in the data. Of course, real holes at scales that are relatively small compared with the size of the home range may well be missed. Differences between real and spurious holes in k -NNCH coverings of data sets should also be evident in plots of the number of holes in a particular k -NNCH covering against the value of k : the covering of spurious holes should correspond to a leveling off of the resulting graph. Only experience with the method, however, will reveal appropriate methods for deciding when this leveling off has been achieved. In our case, we know the topology of the data; and we use k^* to denote the value obtained using our MSHC rule.

Results

The AD data (Fig. 1 A) and various k -NNCH coverings ($k = 2, 6, 10, k^* = 17$, and $k = 301$, where the latter is the smallest k that covers the permanent hole in the center) are plotted in Fig. 1. UD and shaded deciles are drawn in Fig. 2 for the 17-NNCH and 50-NNCH coverings (panels A and B) for the h_{REF} and h_{LSCV} fixed and




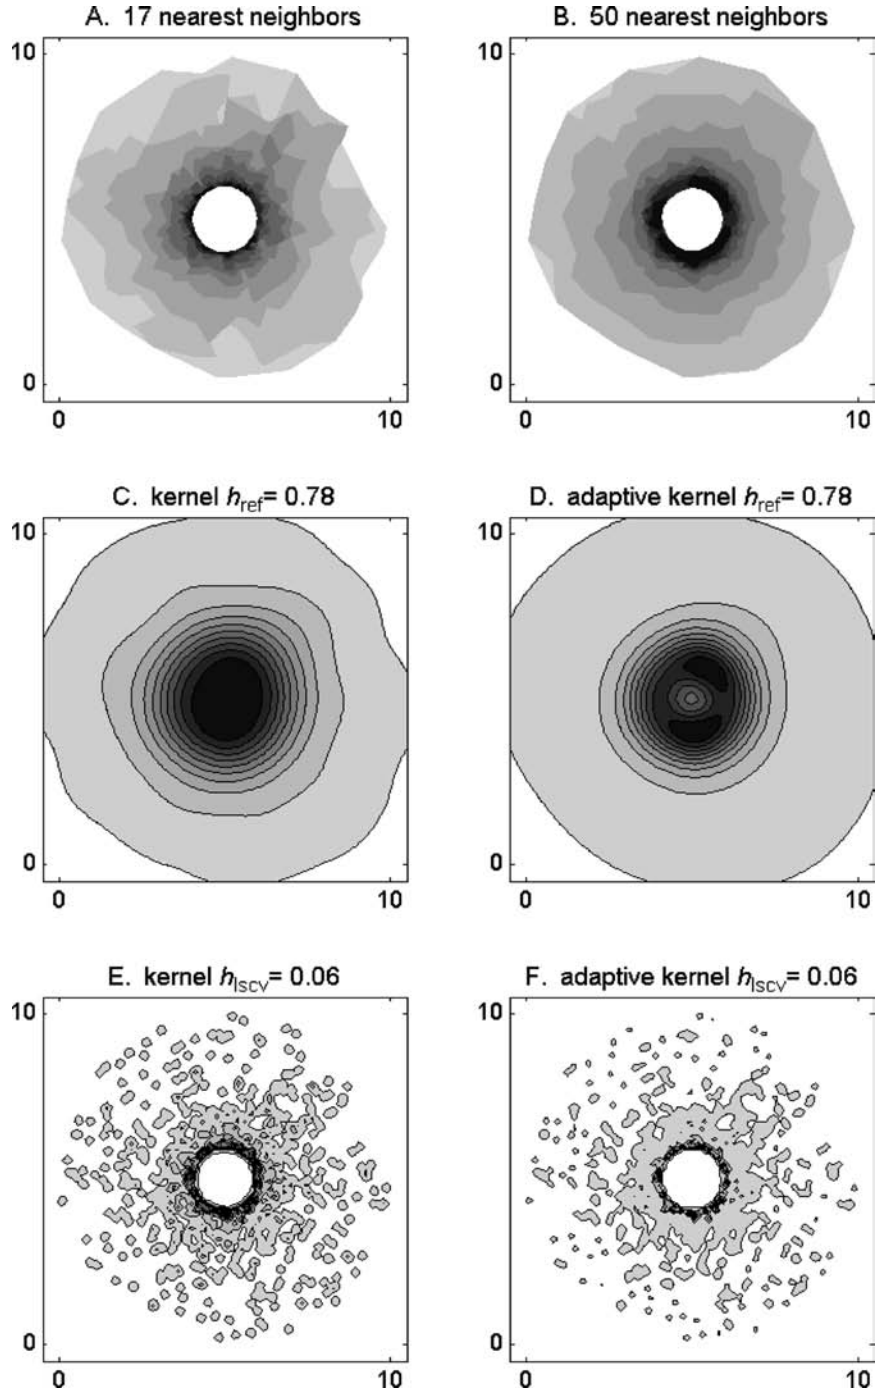
Fig. 1. k -NNCH coverings (recall that k is the number of nearest neighbors used to construct local minimum convex polygons) of the (A) AD data are illustrated for the cases (B) $k = 2$, (C) $k = 6$, (D) $k = 10$, (E) $k^* = 17$ and (F) $k = 301$. (The donut hole is first filled when $k = 301$ and the covering becomes the minimum convex polygon $k = 1089$, the total number of points in the data set.)

adaptive kernel methods (panels C–F). The areas associated with some of these constructions are plotted (Fig. 3A) for decile isopleths (kernel constructions) with the number of points and associated densities covered by each decile interval plotted in Figs 3B and C.

Decile shadings of the RD (Fig. 4A), AS (Fig. 5A) and RS (Fig. 6A) data are plotted for the 5-NNCH,

k^* -NNCH, and 50-NNCH coverings (panels B–D in Figs 4–7), and for both fixed and adaptive kernel methods using h_{REF} (panels E–F in Figs 4–7) and h_{LSCV} (panels G–H in Figs 4–7) smoothing parameter values. For the RS, we also include the area, number and density plots associated with these decile intervals (Figs 7A–C). Finally, for MC data (Fig. 8A), we plot

Fig. 2. Decile-shaded k -NNCH coverings of the AD (see Fig. 1A, data generated in an area of ca 75 units) are illustrated for the cases (A) $k^* = 17$ (area = 66 units) and (B) $k = 50$ (area = 68 units). Decile isopleths are plotted for distributions obtained using the REF smoothing parameters for the fixed (C) $h_{\text{REF}} = 0.78$ (area_{99%} = 99 units) and adaptive (D) $h_{\text{REF}} = 0.78$ (area_{99%} = 107 units) kernel methods and using the LSCV smoothing parameter value for the fixed (E) $h_{\text{LSCV}} = 0.058$ (area_{99%} = 27 units) and adaptive (F) $h_{\text{LSCV}} = 0.058$ (area_{99%} = 21 units) kernel methods.



decile shadings of the 5-NNCH (Fig. 8B), k^* -NNCH ($k^* = 17$, Fig. 8C) and 50-NNCH (Fig. 8D) coverings. For comparison we plot decile shadings of the fixed and adaptive kernel distributions for this data for the h_{REF} (Fig. 8E–F) and h_{LSCV} (Fig. 8I–J) cases.

To examine how well the methods converge to the area associated with the AD data (Fig. 1A), we sub-sampled five sets for each of a 30-point, 100-point and 300-point assessment of the performance of our method (Table 1). The UDs obtained from the k^* -NNCH covering and




Fig. 3. The (A) area, (B) number of points, and (C) corresponding density (number of points divided by area) included in each decile partition are graphed for each of the two k-NNCH coverings and two adaptive kernel distributions of the AD data plotted in Fig. 2.

h_{LSCV} adaptive kernel method are illustrated in Fig. 9 for one of the five 100 and 300 data point subsets (Fig. 9A–B). Area estimates averaged over the five different sets for each of the three cases are given in Table 1.

The 3-hull coverings of all five data sets are illustrated in Fig. 10A–E. The comparison of areas estimated by these coverings, as well as 6-hull coverings, with those of selected k-NNCH coverings and kernel methods are tabulated in Table 2.

Discussion

Minimum convex polygon (MCP) and kernel methods are currently the mainstay of the home range construction literature. The reason could be the ease of calculating areas from MCPs and the existence of software packages for implementation of kernel methods including the more complicated adaptive kernel method (e.g. CALHOME, RANGES IV and V, and TRACKER – see Lawson and Rodgers 1997 for a review).

Our k-NNCH covering method is a simple extension of MCP to a union of a set of local MCPs. As such, our method is easy to understand and relatively easy to implement. The primary challenge in producing a

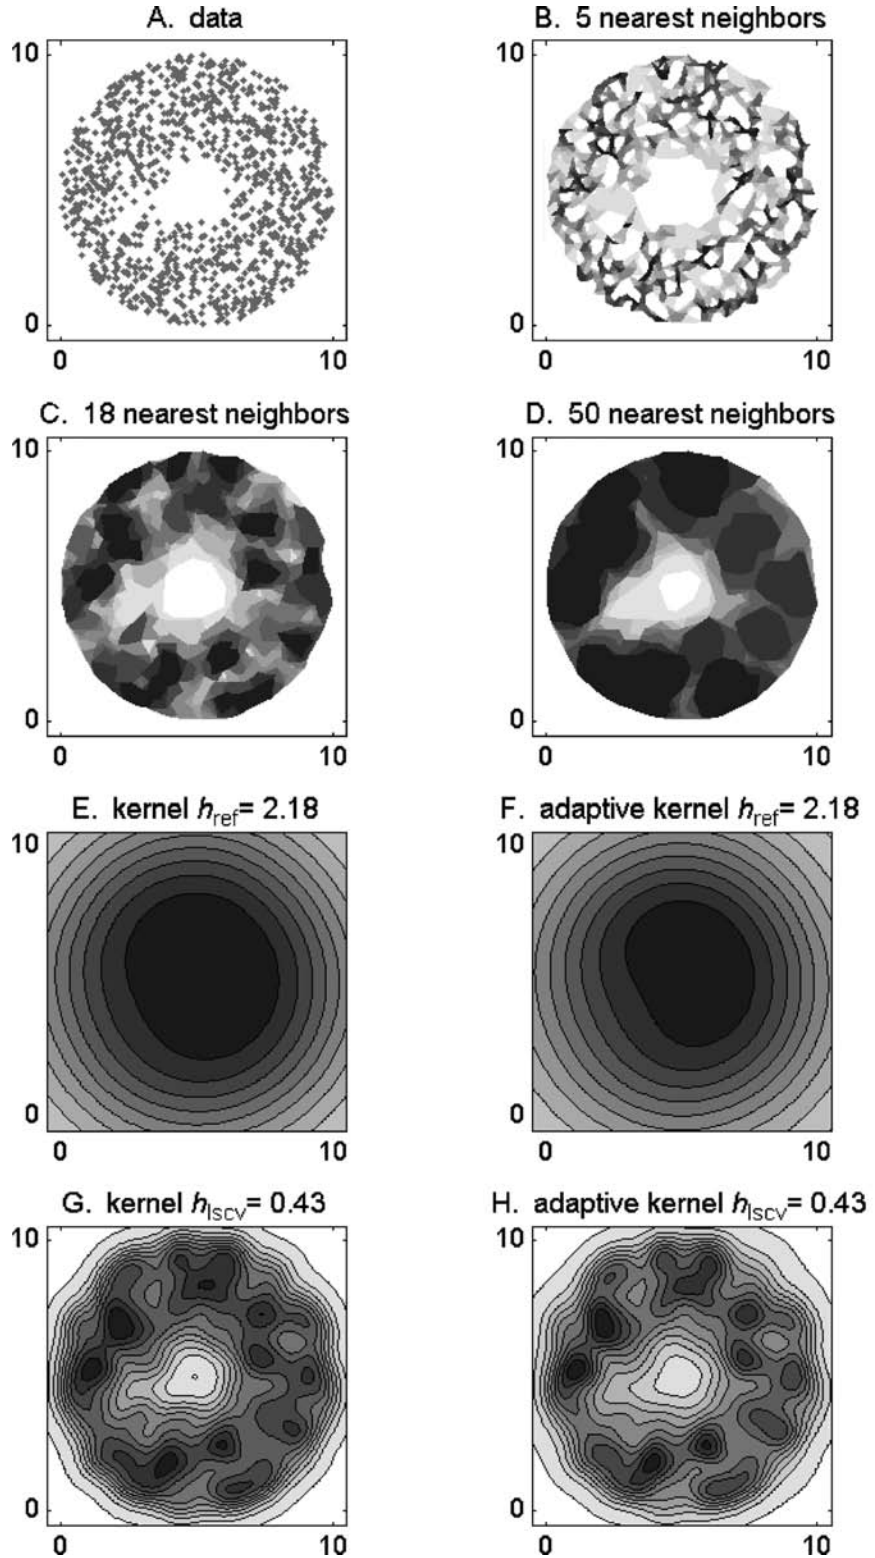
k-NNCH covering is deciding for a particular set of data what the “best” value for k might be. The best value for k should clearly equal or exceed k^* , as evident from Fig. 1B–E. For the first four data sets (AD, RD, AS and RS), however, $k = 50$ provides slightly better area estimates than k^* (Table 2). This is not the case for the fifth data set (MC: Table 2): the value of k producing the best area estimate is likely to vary for different data sets. Selecting the best value for k could be based on minimizing changes in area as a function of k , but the question remains open until more experience is gained using our approach. The question, however, appears to be much less pressing than that of finding the best value of the smoothing parameter h for kernel methods because of the vast range of area estimates obtained for different values of h (cf. panels E–J in Fig. 8). By contrast, comparisons of k-NNCH area estimates indicate very little difference between area estimates using k^* and the ad-hoc value $k = 50$ for 4 of the 5 data sets (Table 2).

Also of consideration in selecting a value for k is the issue, as discussed in the introduction, of the relative importance of avoiding type I vs type II errors. Errors are unavoidable and the smaller the data set the greater the error rate should be (although, this sensible requirement is not always true for kernel methods – Table 1). Relatively large smoothing parameter values for the fixed and adaptive kernel methods (i.e. $h_{REF} = 0.78$) may avoid type II errors (the donut is completely covered – see Fig. 2C, D) but produce large type I errors (at least 32% and 43% respectively of the area are misidentified) through the inclusion of regions that lie beyond the outer circumference of the AD (Table 2). Further the fixed kernel method misidentifies the AD hole as the most heavily utilized part of the home range (Fig. 2D).

At the other extreme, for relatively small values of the smoothing parameter (i.e. $h_{LSCV} = 0.058$), both the fixed and adaptive kernel methods do well at minimizing type II errors, but only at considerable expense with regard to type I errors and extensive fragmentation of the identified area (Fig. 2E, F). In particular, these methods underestimate the area of the aggregated donut by 64% and 72% respectively (Table 2). The α -hull method performs hardly better for the case $\alpha = 3$ in underestimating the area of the AD by 49%, although the underestimate for the case $\alpha = 6$ is much improved at 16% (Table 2). On the other hand, our k-NNCH method performs well over a large range of k values, underestimating the area of the AD data by 12% for the 17-NNCH covering and only 9% for the 50-CH covering (Table 2).

Comparisons of home ranges constructed using k-NNCH coverings and kernel methods for both the AD (Fig. 2) and AS (Fig. 4) data sets indicate how much better the former are than the latter when the data includes heavily used boundaries and intersections of

Fig. 4. Decile-shaded k-NNCH coverings of the (A) RD data (generated in an area of approximately 75 units) are plotted for the cases (B) $k = 5$ (area = 44 units), (C) $k^* = 17$ (area = 72 units), and (D) $k = 50$ (area = 75 units). Decile isopleths are plotted for distributions obtained the following kernel methods: (E) fixed, $h_{REF} = 2.2$ (area_{99%} = 382 units); (F) adaptive, $h_{REF} = 2.2$ (area_{99%} = 382); (G) fixed, $h_{LSCV} = 0.44$ (area_{99%} = 110 units); and (H) adaptive $h_{LSCV} = 0.44$ (area_{99%} = 115 units).







Fig. 5. Decile-shaded k-NNCH coverings of the (A) AS data (generated in an area of 1 unit) are plotted for the cases (B) $k^* = 5$ (area = 0.51 units), (C) $k = 29$ (area = 0.95 units), and (D) $k = 50$ (area = 0.95 units). Decile isopleths are plotted for distributions obtained using the following kernel methods: (E) fixed, $h_{REF} = 0.025$ (area_{99%} = 0.73 units); (F) adaptive, $h_{REF} = 0.025$ (area_{99%} = 0.29 units); (G) fixed, $h_{LSCV} = 0.0037$ (area_{99%} = 0.13 units); and (H) adaptive, $h_{LSCV} = 0.0037$ (area_{99%} = 0.015 units).

Fig. 6. Decile-shaded k-NNCH coverings of the (A) RS data (generated in an area of 1 unit) are plotted for the cases (B) $k = 5$ (area = 0.56 units), (C) $k^* = 22$ (area = 0.97), and (D) $k = 50$ (area = 0.98 units). Decile isopleths are plotted for distributions obtained using the following kernel methods: (E) fixed, $h_{\text{REF}} = 0.026$ (area_{99%} = 1.21 units); (F) adaptive, $h_{\text{REF}} = 0.026$ (area_{99%} = 1.23 units); (G) fixed, $h_{\text{LSCV}} = 0.057$ (area_{99%} = 1.52 units); and (H) adaptive, $h_{\text{LSCV}} = 0.057$ (area_{99%} = 1.61 units).



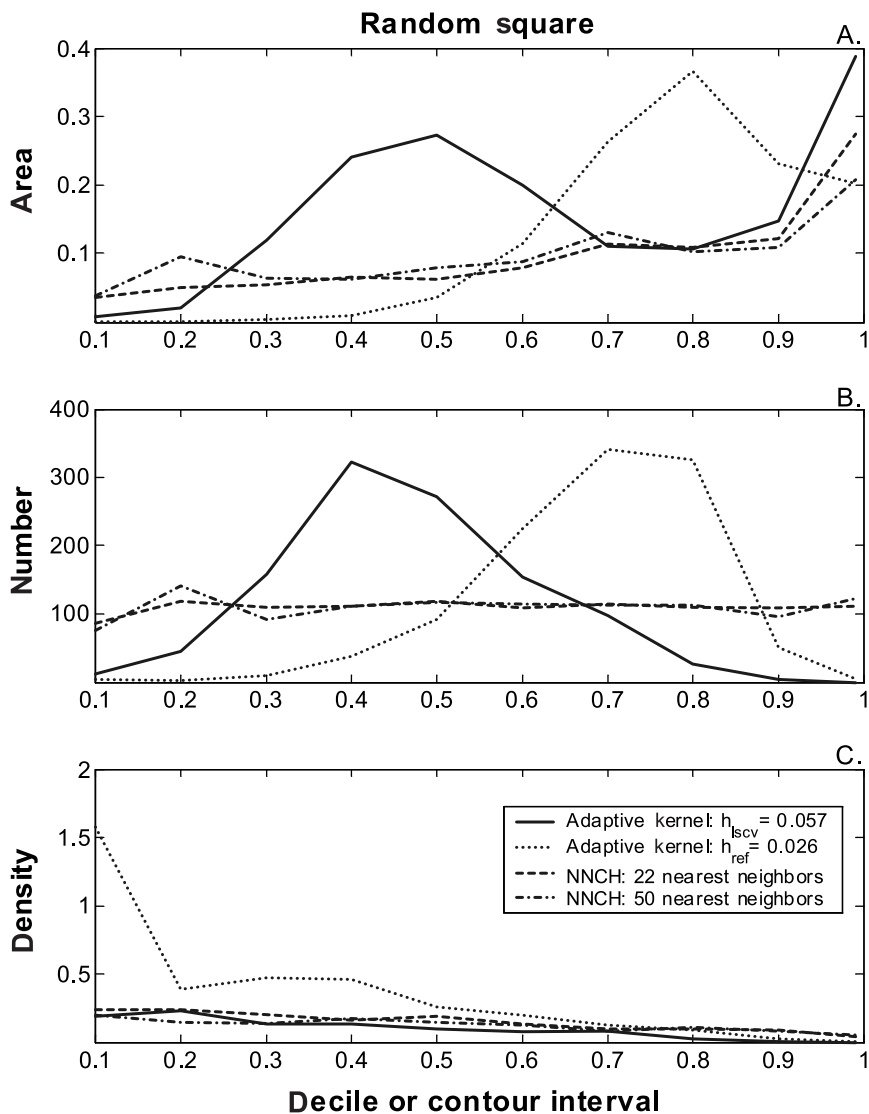


Fig. 7. The (A) area, (B) number of points, and (C) corresponding density (number of points divided by area) included in each decile partition are graphed for the larger two k -NNCH coverings and two adaptive kernel distributions of the RS data plotted in Fig. 6.

such boundaries (corners). For example, all three k -NNCH coverings (Fig. 5B–D) clearly identify the high-density (lower-left) and two medium-density (upper-left and lower-right) corners of the aggregated square (AS). Only the low-density (upper-right) corner is not detected, and then only because no data point falls close enough to this corner to permit identification under any method. On the other hand, kernel methods by design are unable to trace out corners. For the relatively large smoothing parameter value $h_{REF} = 0.025$ corners are obscured (Fig. 5E–F), and for the much smaller smoothing parameter value $h_{LSCV} = 0.0037$ the area is extraordinarily fragmented (Fig. 5G–H). Further, the estimated areas are off by orders

of magnitude: 87% and 98% underestimates respectively for the fixed and adaptive kernel methods (Table 2).

Kernel methods perform better on non-aggregated than aggregated data, but still have problems with corners and donut holes. For the random square (RS) data (Fig. 6), kernel methods smear out the corners and, surprisingly, the algorithmically complicated h_{LSCV} smoothing parameter construction (Fig. 6G–H) overestimates the area of the square by more than twice that of the much simpler h_{REF} case (Fig. 6E–F). Also surprisingly, in both cases the adaptive kernel method performs marginally worse than the fixed kernel method in estimating area (see Table 2). By contrast, provided k is sufficiently large to cover all of the spurious holes, our

Fig. 8. Decile-shaded k-NNCH coverings of the (A) MC data (generated in an area of approximately 320 units are plotted for the cases (B) $k = 5$ (area = 199) units; (C) $k^* = 17$ (area = 347 units) and (D) $k = 50$ (area = 449 units). Decile isopleths are plotted for distributions obtained the following kernel methods: (E) fixed, $h_{REF} = 28$ (area_{99%} = 3459 units), (F) adaptive, $h_{ref} = 28$ (area_{99%} = 3459 units); (G) fixed $h = 2.8$ (area_{99%} = 1429 units); (H) adaptive $h = 2.8$, area_{99%} = 1445 units); (I) fixed, $h_{LSCV} = 0.12$ (area_{99%} = 228 units); and (J) adaptive, $h_{LSCV} = 0.12$ (area_{99%} = 215 units).

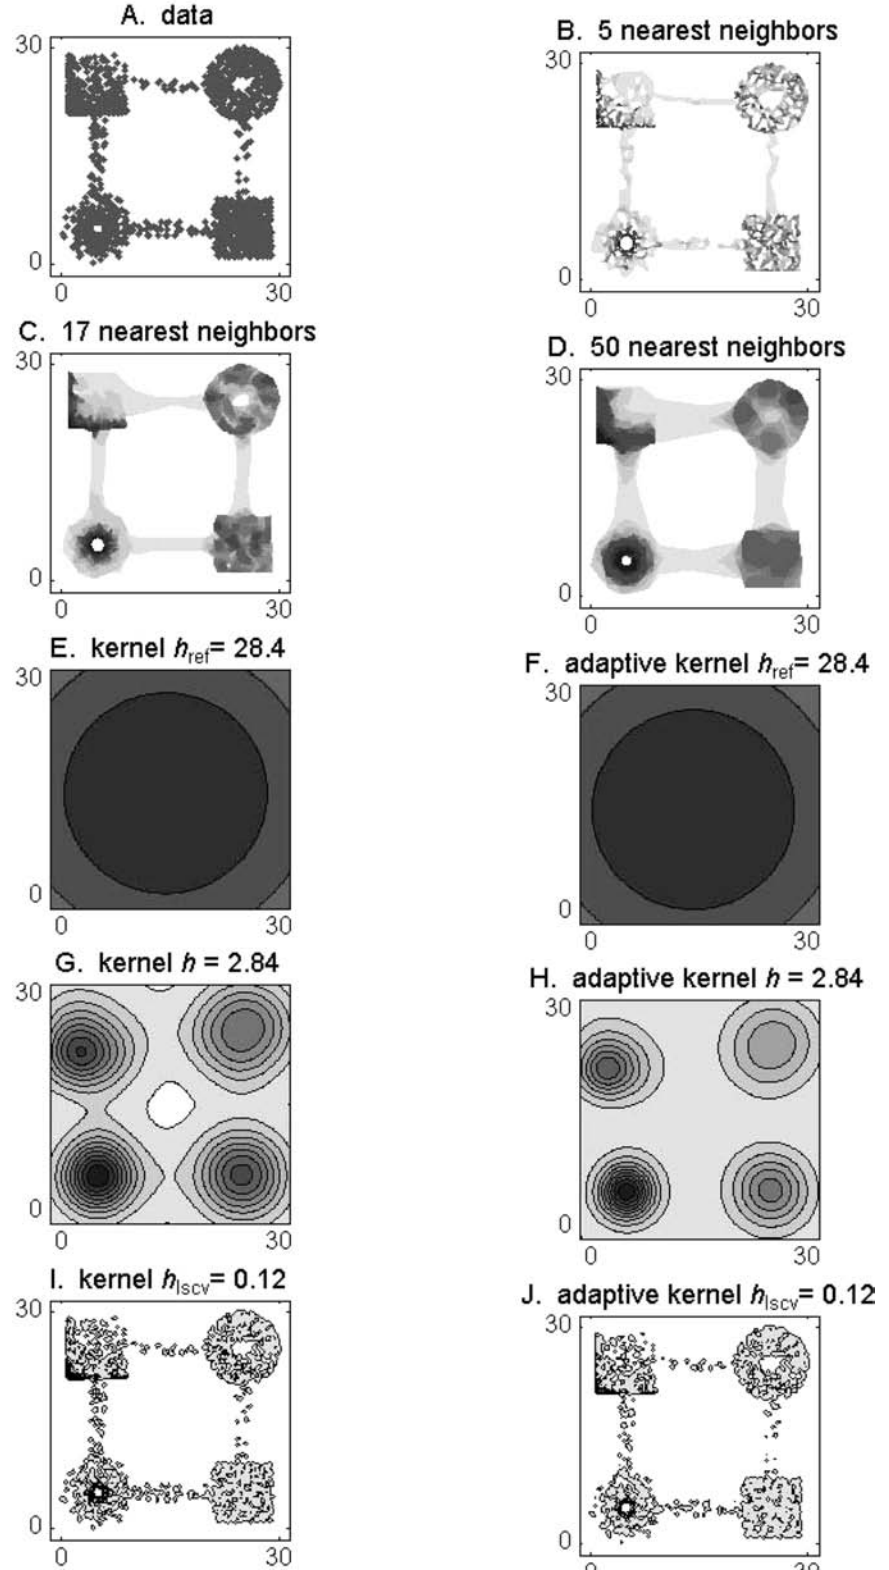


Table 1. Estimates of area from k-NNCH coverings and the 99th percentile of LSCV-optimized adaptive kernel distributions obtained using each of five 30-point, 100-point and five 300-point randomly sampled subsets of the 1089 points in the AD data (Fig. 1A).

Sample size	k-NNCH		Adaptive kernel	
	k*	Area	h_{LSCV}	Area _{99%}
30	10	28.9	0.43	27.0
	9	26.9	0.60	28.6
	10	32.7	0.62	25.3
	9	26.5	0.49	27.4
	10	27.5	0.13	10.0
Mean (SD)		28.4 (2.6)		23.6 (7.7)
Percent error ¹		-62%		-69%
100	11	48.2	0.20	34.1
	12	35.3	0.19	29.0
	9	42.3	0.21	35.0
	14	42.9	0.21	33.6
	12	39.0	0.30	50.8
Mean (SD)		41.6 (4.8)		36.5 (8.3)
Percent error ¹		-45%		-52%
300	13	54.3	0.10	23.4
	16	61.2	0.09	22.3
	16	58.3	0.09	17.7
	12	60.6	0.12	31.5
	15	56.3	0.11	25.4
Mean (SD)		58.1 (2.9)		24.1 (5.0)
Percent error ¹		-23%		-68%
1089	17	66.0	0.06	20.8
Percent error ¹		-12%		-72%

¹The percentage of the known actual value that would have to be added or subtracted (negative numbers) to this value to obtain the estimated value.

k-NNCH coverings accurately maps out the home range and its associated distributions of points (higher and lower densities areas arise at random). In particular, k-NNCH coverings underestimates the area of the square by 3% for $k^* = 22$ and by 2% when $k = 50$ (Table 2). (Note because the points always fall within the defined unit square, the actual area represented by the points is always < 1 , so the best method should always give a slight underestimate.) The α -hull method does comparatively well in underestimating the area of the square by 5% when $\alpha = 3$ and only 2% when $\alpha = 6$.

For the random donut (RD) data (Fig. 4A), our k-NNCH method continues to provided good estimates of the area, underestimating it by 5% for the k^* -NNCH ($k^* = 18$) covering and by only 1% for selected 50-NNCH covering. Kernel methods, on the other hand fail to locate the hole in all case (Fig. 4E-H). Further, kernel methods provide very poor estimates of the RD area using h_{REF} , overestimating it by 409% in the case of both the fixed and adaptive kernel methods (Table 2). Even the "optimized" h_{LSCV} parameter performs poorly, overestimating the area by 48% for the fixed and 53% for the adaptive kernel methods (Table 2). Again, the α -hull method does well in underestimating the area by 8% when $\alpha = 3$ and 4% when $\alpha = 6$, which we can compare


with 5% and 1% underestimates for the k^* -NNCH and 50-NNCH constructions respectively (Table 2).

Multimodal data also challenges the construction of UD. Although kernel methods are regularly used to fit distributions to multimodal data, Casear et al. (1999) have demonstrated that the Thiessen method, employing a simple Dirichelet tessellations of the data, is superior to kernel methods in identifying core usage areas. The Thiessen method itself provides an estimate of area equal to MCP, which is generally very poor (Bergman and Fox 2003). From Fig. 8, it is clear that kernel methods perform very poorly in mapping out the home range distribution of the MC data (Fig. 8A). In the case of $h_{REF} = 28.4$, the fixed and adaptive kernel methods completely fail to identify high use areas (Fig. 8E-F); and they overestimate the area by an order of magnitude (Table 2). In the case of $h_{LSCV} = 0.12$, the fixed and adaptive kernel methods yield highly fragmented home ranges (Fig. 8I-J); and they underestimate the area by close to 33% (Table 2). For the completely ad-hoc intermediate case $h = 2.84$ the core areas are identified without unduly fragmenting the home range (Fig. 8G-H), but even then very poor representations are obtained of the shape and size of the core areas and associated corridors.

Our k-NNCH coverings capture very well the shape of the core areas and the corridors (Fig. 8B-C) associated with the MC data. Additionally, the k^* -NNCH covering identifies both donut holes and only overestimates the area by 8% (Table 2). The more arbitrary 50-NNCH covering does not do quite as well: it covers one of the donut holes and overestimates the area by 41%. For $\alpha = 3$, though, the α -hull method, provides an area estimate matching the 8% performance of the k^* -NNCH covering, except it provides an under rather than an over estimate. The 3-hull method, however, does not identify corridors as well as k-NNCH coverings (Fig. 10E), yielding one fragmented corridor and leaving two of the remaining three corridors linked by lines rather than area segments.

A critical weakness of kernel methods is that unlike hull methods (both the α -hull and k-NNCH constructions) they do not provide convergent area estimates with increasing number of points. As demonstrated by Seaman et al. (1999), this holds even for bivariate normal data. The problem is much worse for aggregated data sets, such as AD (Fig. 1A). Rather than reaching an asymptote, the area estimates get worse in the case of the h_{LSCV} adaptive kernel method. Specifically, for subsamples of 30, 100, 300, and the full 1089 points the method underestimates the area by 69%, 52%, 68% and 72% respectively (Table 1 - Fig. 9B, D, and F) which should be compared with the converging sequence 62%, 45%, 23% and 12% for the k^* -NNCH coverings (Table 1 - Fig. 9A, C, and E) and the converging sequence.

Fig. 9. k^* -NNCH and adaptive kernel constructions of UD_s using 100-point (A) $k^* = 11$, (B) $h_{LSCV} = 0.33$; 300-point (C) $k^* = 14$, (D) $h_{LSCV} = 0.15$; and 1000-point (E) $k^* = 16$, (F) $h_{LSCV} = 0.06$ sub samples of the AD data (Fig. 1A). (See Table 1 for information on area estimates.)



Beyond the questions of the accuracy and convergence of area estimates, and of identifying high-density regions in multimodal data, is the question of the accuracy of the density isopleths themselves. For example, a plot of decile isopleths against the number of points actually bounded by those isopleths should be flat. This is nearly the case for the k -NNCH constructions plotted in Figs 3B and 7B (AD and RS data), although the lines are flatter for the smaller than larger values of k because of rounding errors (the union of groups of $k \ll N$ points into precise decile intervals of size $N/10$ produces smaller rounding errors for smaller values of k). The number-of-

points plotted per decile isopleth is not at all flat for the kernel UD_s. Specifically, for the adaptive $h_{REF} = 0.78$ UD constructed from the AD data, the tails (the first, second, and last deciles of the distribution) contain at least twice as many of points as they should, thereby producing erroneous area (Fig. 3A and C) and density estimates of the associated UD.

Errors associated with the adaptive $h_{LSCV} = 0.06$ UD are even more severe with hardly any points included in the first seven decile intervals and most of the points in the last decile interval (Fig. 3B) resulting to nonsensical area and density plots (Fig. 3A and C). For the RS data,

α -hull




Fig. 10. For the case $\alpha = 3$, α -hull coverings of the 5 data sets (A) AD (area 38 units), (B) RD (area 69 units), (C) AS (area 0.50 units), (D) RS, (area 0.95 units) and (E) MC (area 296 units).

Table 2. Areas associated with the generated data and percent error¹ in estimates obtained by k-NNCH, α -hull, and 99th% of kernel distributions.

Data (Area ²)	k-NNCH		Fixed kernel		Adaptive kernel		α -CH	
	k*	k = 50	h _{REF}	h _{LSCV}	h _{REF}	h _{LSCV}	$\alpha = 3$	$\alpha = 6$
AD (75.4)	-12% (17)	-9%	32%	-64%	43%	-72%	-49%	-16%
RD (75.4)	-5% (18)	-1%	409%	48%	409%	53%	-8%	-4%
AS (1)	-5% (29)	-5%	-27%	-87%	-71%	-98%	-50%	-12%
RS (1)	-3% (22)	-2%	21%	52%	23%	61%	-5%	-3%
MC (320)	8% (17)	41%	981%	-29%	981%	-33%	-8%	11%

¹Negative values are underestimates, positive overestimates (see footnote to Table 1).

²This is the area in arbitrary units (approximated to two significant digits) of the set bounding the data that is generated using the rules described in the Methods. Units are unspecified or, equivalently, dimensionless.

the adaptive kernel UD's now completely underestimate the number of points in the tail for both the $h_{REF} = 0.026$ and $h_{LSCV} = 0.077$ constructions (Fig. 7B: most of the points are covered by the seventh and eighth decile intervals in the former case and fourth and fifth decile intervals in the latter case). Again, the inability of these kernel methods to demarcate decile intervals of points with any reasonable accuracy translates into hopelessly erroneous area (Fig. 7B) and density (Fig. 7C) plots.

Conclusion

The construction of unbiased high resolution UD's ultimately depends on the quantity and quality of the data available, and issues such as serial correlations (De Solla et al. 1999) and sampling errors affect all methods to a greater or lesser degree. Modern radio telemetry, however, provides data in much greater quantities and of much higher quality than ever before. Thus our k-NNCH covering, which converges on the true distribution as the quality and quantity of data increases, provides a superior alternative to methods such as kernel methods, which do not converge. Further, we have demonstrated that k-NNCH provide much better fits than kernel methods across a spectrum of distributions of data, from uniform to highly aggregated, and multimodal.

Kernel methods perform particularly poorly on aggregated and clustered data. Also, they were unable to clearly demarcate boundaries and tended to fill in real holes. We are certainly not the first to recognize this problem. Creel and Creel (2002), p. 37, for example, in their application of the adaptive kernel module of the CALHOME program (Kie et al. 1994) to construct utilization distributions from GPS data on the movement of wild dogs in Africa state "...[we] modified the shapes of several home ranges to exclude areas that could not be used (lakes) ... [by] overlaying the home range contours onto a base map of the study area and cutting out the unusable areas by manual onscreen digitizing." Further, the poor performance of kernel methods in estimating home range areas is well docu-

mented (Lawson and Rodgers 1997, Ostro et al. 1999), as is the problem of non convergence of kernel methods with increasing sample size to some unbiased area estimate (Casaer et al. 1999).

Yet kernel methods continue to be widely used. The reason for this might be that other relatively simple methods, such as MCP and α -hulls, do not produce density isopleths; even though an algorithm can be devised to construct density isopleths associated with a given α -hull construction. Our k-NNCH does not have this deficiency and leads directly to the construction of density isopleths. It appears to provide very good area estimates for challenging data sets and converges to the true area as the number of data points increase. Although α -hull methods, also satisfy this latter property, they suffer from the deficiency of not always including all points within or on the boundary of the constructed area (i.e. some points may not be included at all or they may be joined to an area by a line segment). In short, k-NNCH coverings provide a general approach to home range and UD construction that is superior to existing kernel and hull methods, particularly when the data reflects the existence of real boundaries, is multimodal, and topologically complex.

Acknowledgements – This research was funded in part by the United States National Science Foundation Ecology of Infectious Disease Grant DEB-0090323 and a James S. McDonnell Foundation Complex Systems Award to WMG and by an Environmental Protection Agency STAR Fellowship to CCW. We thank Paul Cross and Sadie Ryan for valuable discussion and comments during the preparation of this paper.

References

- Baker, J. 2001. Population density and home range estimates for the eastern bristlebird at Jervis Bay, south-eastern Australia. – *Corella* 25: 62–67.
- Burgman, M. A. and Fox, J. C. 2003. Bias in species range estimates from minimum convex polygons: implications for conservation and options for improved planning. – *Anim. Conserv.* 6: 19–28.
- Burt, W. H. 1943. Territoriality and home range concepts as applied to mammals. – *J. Mammal.* 24: 346–352.
- Casaer, J. et al. 1999. Analysing space use patterns by Thiessen polygon and triangulated irregular network interpolation: a

- non-parametric method for processing telemetric animal fixes. – *Int. J. Geogra. Inf. Sci.* 13: 499–511.
- Creel, S. and Creel, N. M. 2002. The African wild dog: behavior, ecology, and conservation. – Princeton Univ. Press.
- Dale, M. R. T. et al. 2002. Conceptual and mathematical relationships among methods for spatial analysis. – *Ecography* 25: 558–577.
- De Solla, S. R., Bonduriansky, R. and Brooks, R. J. 1999. Eliminating autocorrelation reduces biological relevance of home range estimates. – *J. Anim. Ecol.* 68: 221–234.
- Ford, R. G. and Krumme, D. W. 1979. The analysis of space use patterns. – *J. Theor. Biol.* 76: 125–157.
- Jennrich, R. I. and Turner, F. B. 1969. Measurement of non-circular home range. – *J. Theor. Biol.* 22: 227–237.
- Kenward, R. E. et al. 2001. Density and linkage estimators of home range: nearest-neighbor clustering defines multinuclear cores. – *Ecology* 82: 1905–1920.
- Kie, J. G., Baldwin, J. A. and Evans, C. J. 1994. CALHOME home range analysis program user's manual. – US Forest Service Pacific Southwest Research Station, Fresno.
- Lawson, E. J. G. and Rodgers, A. R. 1997. Differences in home-range size computed in commonly used software programs. – *Wildl. Soc. Bull.* 25: 721–729.
- Meulman, E. P. and Klomp, N. I. 1999. Is the home range of the heath mouse *Pseudomys shortridgei* an anomaly in the *Pseudomys* genus? – *Victorian Nat.* 116: 196–201.
- Ostro, L. E. T. et al. 1999. A geographic information system method for estimating home range size. – *J. Wildl. Manage.* 63: 748–755.
- Plotnick, R. E., Gardner, R. H. and O'Neill, R. V. 1993. Lacunarity indices as measures of landscape texture. – *Landscape Ecol.* 8: 201–211.
- Plotkin, J. B., Chave, J. and Ashton, P. S. 2002. Cluster analysis of spatial patterns in Malaysian tree species. – *Am. Nat.* 160: 629–644.
- Plotnick, R. E. et al. 1996. Lacunarity analysis: a general technique for the analysis of spatial patterns. – *Physical Rev. E* 53: 5461–5468.
- Ripley, B. D. 1987. Stochastic simulation. – Wiley.
- Rurik, L. and Macdonald, D. W. 2003. Home range and habitat use of the kit fox (*Vulpes macrotis*) in a prairie dog (*Cynomys ludovicianus*) complex. – *J. Zool.* 259: 1–5.
- Seaman, D. E. et al. 1999. Effects of sample size on kernel home range estimates. – *J. Wildl. Manage.* 63: 739–747.
- Seaman, D. E. and Powell, R. A. 1996. An evaluation of the accuracy of kernel density estimators for home range analysis. – *Ecology* 77: 2075–2085.
- Silverman, B. W. 1986. Density estimation for statistics and data analysis. – Chapman and Hall.
- Worton, B. J. 1989. Kernel methods for estimating the utilization distribution in home-range studies. – *Ecology* 70: 164–168.
- Worton, B. J. 1995. A convex hull-based estimator of home range size. – *Biometrics* 51: 1206–1215.

Appendix

A k-NNCH covering for constructing UD

Given a set of specified points (vectors) $Z^n = \{z_i = (x_i, y_i) | i = 1, \dots, n\}$ the method begins by constructing the convex hull associated with each point z_i and its $(k - 1)$ nearest neighbors (i.e. its k-NNCH). The area covered by each k-NNCH is then calculated and the points z_i are sorted and renumbered according to the area of the associated local convex hull. This results in a

list $L_0 = \{(z_i; c_i^k; a_i) | i = 1, \dots, n\}$, where c_i^k is the name for the convex hull associated with z_i and its $(k - 1)$ nearest neighbors, and a_i the area of c_i^k with indices reordered such that $a_1 \leq a_2 \leq \dots \leq a_n$. Then defining the unions $C_i^k = \cup_{j=1}^i c_j^k$, the list L_0 is extended to obtain $L_E = \{(z_i; c_i^k; a_i; C_i^k; A_i; N_i) | i = 1, \dots, n\}$, where A_i is the area of C_i^k and N_i is the number of points z_i associated with C_i^k (a certain number will lie on the boundary, defining the boundary elements and the rest will be in the interior).

The list L_E is used to construct percentiles of points contained in the UD as follows. The densest area containing at most $p\%$ of the points is C_i^k where i is the largest integer for which $N_i \leq np/100$. If we denote this value of i by i_p , then, for a selected set of values $0 < p_1 < p_2 < \dots < p_m = 100$ we can construct a corresponding nested set of regions $\{C_{i_{p_1}}^k, C_{i_{p_2}}^k, \dots, C_{i_{p_{m-1}}}^k, C_{i_{100}}^k\}$ each with area $\{A_{i_{p_1}}^k, A_{i_{p_2}}^k, \dots, A_{i_{p_{m-1}}}^k, A_{i_{100}}^k\}$ to represent the UD. These areas can be represented graphically and values tabulated provide a visual and quantitative characterization of the UD. Of particular interest are the densities (here we define $p_{i-1} = 0$ and $A_{i_{p_0}}^k = 0$) $d_{p_i}^k = (A_{i_{p_i}}^k - A_{i_{p_{i-1}}}^k) / n(p_i - p_{i-1})$, $i = 1, \dots, m$, which necessarily are non-increasing with i .

The algorithm was implemented using vector-based methods in MATLAB v6.12 (Mathworks).

Kernel methods

The standard choice for the smoothing parameter h , which is known to be optimal for bivariate data (Silverman 1986, Worton 1995), is the “reference” value $h_{REF} = sn^{-1/6}$, where $s = ((s_x^2 + s_y^2)/2)^{1/2}$ and s_x^2 and s_y^2 are the variances respectively of the x , and y locations of the data. For non-bivariate data, the least-squares cross-validated smoothing parameter value, denoted h_{LSCV} , is regarded as “best” or “optimal” for the fixed kernel method: it is the value of h that minimizes the cross-validation function $CV(h)$ (the formula for this function is described elsewhere – see Silverman 1986, Worton 1995, Seaman and Powell 1996). We found the minimum, as suggested in Worton (1995), by plotting $CV(h)$ over the range $0.1h_{REF} < h < 1.5h_{REF}$ to the desired degree of accuracy, extending the range if, initially, an internal minimum was not obtained. For the adaptive kernel method, following Silverman's recommendation that h_{LSCV} is a reasonable choice for the global smoothing parameter useful form for minimizing (also see Seaman and Powell 1996) we also used h_{LSCV} as our global smoothing parameter and locally modified as prescribed for adaptive kernel methods by Silverman (1986), also see Worton 1995, Seaman and Powell 1996).

We also calculated the areas enclosed by a sequence of isopleths p_i , $i = 1, \dots, m - 1$, in the same way as we do for our k-NNCH algorithm, and then calculate the

densities $d_{p_i}^f$ (fixed method) and $d_{p_i}^a$ (adaptive method), $i = 1, \dots, m - 1$, using the number of points with in the area bounded by each isopleth divided by the area itself.

Multicore computer-generated data sets

An AD (cf. data set 4) containing 500 hundred rather than 1089 points (i.e. area = 75.4 with density of $500/75.4 = 6.63$ points per unit area) was connected by the rectangular corridor $[9 \leq x \leq 21] \times [4 \leq y \leq 6]$ containing 50 points (i.e. area = 24 with density of $50/24 = 2.08$ points per unit area) to an RS (cf. data set 1) containing 500 points (i.e. area is 64 with density of $500/64 = 7.81$ points per unit area). This AD was also connected by the 50 point $[4 \leq x \leq 6] \times [9 \leq y \leq 21]$

corridor (i.e. area = 75.4 with density of 2.08 individuals per unit area) to a 500 point AS (cf. data set 2) located at $[1 \leq x \leq 9] \times [21 \leq y \leq 29]$ (i.e. area = 64 with density of 7.81 individuals per unit area). Finally, both rectangles were connected by the 25-point corridors $[9 \leq x \leq 21] \times [24 \leq y \leq 26]$ (i.e. density is $25/24 = 1.04$ points per unit area) and $[24 \leq x \leq 26] \times [9 \leq y \leq 21]$ (i.e. area of 12 with a density of $50/12 = 1.04$ individuals per unit area) to a 500 point RD (cf. data set 3) centered at (25,25) (i.e. area of 75.5 with a density of 6.63 points per unit area). Taking into account that the corridors are not flush with the donuts, but overlap by approximately 1.9 units of area, the total area of MC is approximately $(2 \times 64 + 2 \times 75.4 + 4 \times 12) - 4 \times 1.9 \approx 320$ units and the mean density of points is $2150/320 = 6.7$ points per unit area.