

U.S. Department of Commerce
National Oceanic and Atmospheric Administration
National Weather Service
National Centers for Environmental Prediction
5830 University Research Court
College Park, MD 20740-3818

Office Note 494

<https://doi.org/10.7289/V5/ON-NCEP-494>

HILBERT CURVES ISOMETRICALLY FILLING A SPHERICAL SHELL, AND
THEIR APPLICATION TO THE ESTIMATION OF SPATIAL DATA DENSITY

R. James Purser*
IM Systems Group, Rockville, Maryland

May 11, 2018

THIS IS AN UNREVIEWED MANUSCRIPT, PRIMARILY INTENDED FOR INFORMAL
EXCHANGE OF INFORMATION AMONG THE NCEP STAFF MEMBERS

* email: jim.purser@noaa.gov

Abstract

We review ways of constructing space-filling Hilbert curves that isometrically occupy rectangular or spherical regions in either two or three dimensions, together with a proposal for using the projection and sorting of observational data locations on such curves to facilitate the rough but efficient estimation of the spatial density of these data. For data whose spatial distributions tend to exhibit sporadic clumping, these density estimates can be used to trigger the progressive down-weighting within the assimilation of those data whose perceived density has exceeded a threshold related to the effective resolution of the assimilation. In this way, we are able to partially and approximately account for some of the effects that strongly correlated representativeness errors in the data at short distances should have on the proper assignment of measurement weights where the data are excessively crowded, as occurs, for example, in aircraft reports near busy terminal hubs. The advantage of using space-filling curves, as opposed to direct spatial estimation of data density on a regular fine grid, is that we can avoid the wasted computations that the latter method would entail in large expanses of the domain where the data are too sparse to be of concern. Since our proposed method is, in effect, essentially a stochastic one, a superior estimate of the density, with improved isotropy of the implied spatial averaging of neighboring data, can be obtained by averages of repeated trials in which the particular orientations of the frameworks used to construct the space-filling curves are randomized. Strategies for systematic randomization are suggested.

1. INTRODUCTION

In a space of dimension two or more equipped with a coordinate system, a curve parameterized by a segment of the real line can be thought of as a continuous function evaluating to a unique coordinate set at each value of the parameter in that segment. It was first noted by Peano (1890) that, by this definition, a continuous curve can actually pass through every point of a finite region of the space, and a modification of his construction by Hilbert (1891) produced the simplest geometrical definition, and the one of most practical utility, amongst the related family of space-filling curves. The most obvious computational applications concern the problem of systematically bringing serial order to an unstructured distribution of discrete points in space in situations where it is desirable that the images of the points that end up relatively near to one another on the space-filling curve are also guaranteed to be close to one another in the original ‘physical’ space. Nice examples of the practical value of employing this tool in computer science related to configuring processors for adaptively-resolving fluid dynamics models were given by Behrens and Zimmermann (2000) and by Dennis (2003). In the context of meteorological (or more generally, geophysical) data assimilation, the utility of such a curve combined with an efficient sorting algorithm is that it allows a sporadic distribution in space, such as observation locations, to be rendered into a sorted one-dimensional series of real-numbered locations along the parameterized Hilbert curve.

A space-filling curve can be made ‘isometric’ in the sense that the measure of length of a segment of the curve is exactly proportional to the corresponding measure of area or volume of

the space that the segment of the curve occupies. Observations that map to points clustered close together on the curve necessarily also lie close together in space and, although the converse is not strictly true, we can usually expect that a pair of points that are close to one another in space will most probably not lie very far apart on the Hilbert curve. Thus, if we take a sparsely-sampled subset of the observations as they map to a Hilbert curve, then for the most part (a few exceptions) they will tend to be sparsely spread in the spatial domain, which gives us a convenient means to generate the kinds of samples of subsets of the data we wish to use as the validating subsets of cross-validation procedures. This is how the two-dimensional Hilbert curve is presently used in NOAA/NCEP’s operational Real-Time Mesoscale Analysis (RTMA, e.g., Pondeca et al., 2011) to facilitate the testing and tuning of statistics via cross-validation. In a previous note, Purser et al., (2009, henceforth denoted ON460) it was also suggested that data thinning and super-obbing could be undertaken with the help of this isometric space-filling curve.

In the present note we take these applications further and develop a scheme whereby a crude estimate of the spatial density of a type of local measurement, such as aircraft data, can be inferred by first estimating the density of the data along the Hilbert curve on which the data are projected. The motivation for wanting such estimates relates to the problems associated with the over-representation of opportunistic data which, if not addressed and dealt with via thinning or down-weighting, can lead to serious ill-conditioning problems within the variational assimilation algorithms. Again, the characteristic size of the volume associated with the local density estimation maps directly (via the isometry property) to the characteristic linear parameter distance of the smoother applied on the Hilbert curve to smooth the ‘indicator functions’ (unit delta functions) at each measurement location. The crudeness of the estimate of spatial density is, again, the consequence of the fact that not all close pairs of data in space map to close pairs on the curve, but this imperfection can be mitigated in a statistical sense by repeating the procedure multiple times with differently oriented Hilbert curves, and averaging all the individual density estimates. We will explore the practical advantages of such stochastic ‘randomization’ later in this note.

In order to make this method practical we need to be able to perform the smoothing operation directly to the Hilbert parameter data images themselves and not rely on using any kind of regular numerical grid. Fortunately, just such a smoother was recently developed for another application (time smoothing); it is the B-spline smoother described (for the low-order versions at least) in Purser and Zhu (2016, henceforth PZ16), although, as we find out, it is only the trivial lowest order versions of this family of filters that need to be considered in practice.

In the next section we review the construction of space-filling curves in the domains germane to the problems of meteorological data assimilation, in both two and three dimensions, and in both regional rectangular domains where the sphericity can be neglected and over the entire globe where it must be accommodated.

2. SPACE-FILLING HILBERT CURVES

Peano (1890) first realized that a continuous curve could fill a space of more than one dimension and gave the first construction. A simpler form was soon proposed by Hilbert, and it is this form that finds the most practical applications of such curves (of which there are now




Figure 1. Generic refinement of a portion of the Hilbert curve in a square, the arrows showing the exit corners, and pointing to the next neighbor, for each sub-square of this stage of refinement. Since the entry and exit points of the parent square are at opposite ends of a common edge of it, and this property is inherited by the daughter sub-squares, the path taken by the whole curve is uniquely defined by similar refinements at each recursive stage of subdivision, where the numerals indicate the corresponding quaternary digit in the expression of the curve parameter.

many variants, for example see the book by Sagan, 1994). We give only a cursory review here as the topic is largely covered in ON460 or in the many available general references that define and describe the standard forms of space-filling curves. We shall refer to the geometrical regions in which each Hilbert curve is constructed as its ‘Frame’. The two-dimensional (horizontal) frames we review are the square and sphere; the three dimensional ones are the cube and the thin and thick spherical shells.

(a) *Square*

In the case of the two-dimensional domain comprising the unit square (the simplest example), the Hilbert curve is defined recursively through rules that govern the successive stages of binary refinement of the linear dimensions, quaternary refinement of each square into its four sub-squares, and the corresponding quaternary refinement of the associated segments of the parameter space of the Hilbert curve itself. The point of ingress, or beginning, and the point of egress or finish, are, at every stage of the recursive construction, at adjacent corners of the square or a sub-square. The whole curve is parameterized from, say, zero to one, but each of its quarter-ranges covers each of the four principal half-scale sub-squares, in the order indicated in Fig. 1, and arranged such that the entry and exit points of each sub-square are also at their corners and have an edge of that sub-square in common. The arrows in Fig. 1 indicate the corner at which each sub-square’s portion of the curve exits the square, and points towards the neighboring sub-square about to be entered next as the ‘Hilbert parameter’ increases. The numerals in each square can be thought of as the first quaternary (base-4) digit of the parameter of that portion of the curve occupying the sub-square in question. At each stage of refinement, the linear sizes of the sub-squares are halved and the length of the corresponding segment of the curve decreases by a factor of four, covering a parameter subrange that corresponds to a new quaternary digit. It is easy to see that, once the process is begun, the subsequent refinements are uniquely determined.

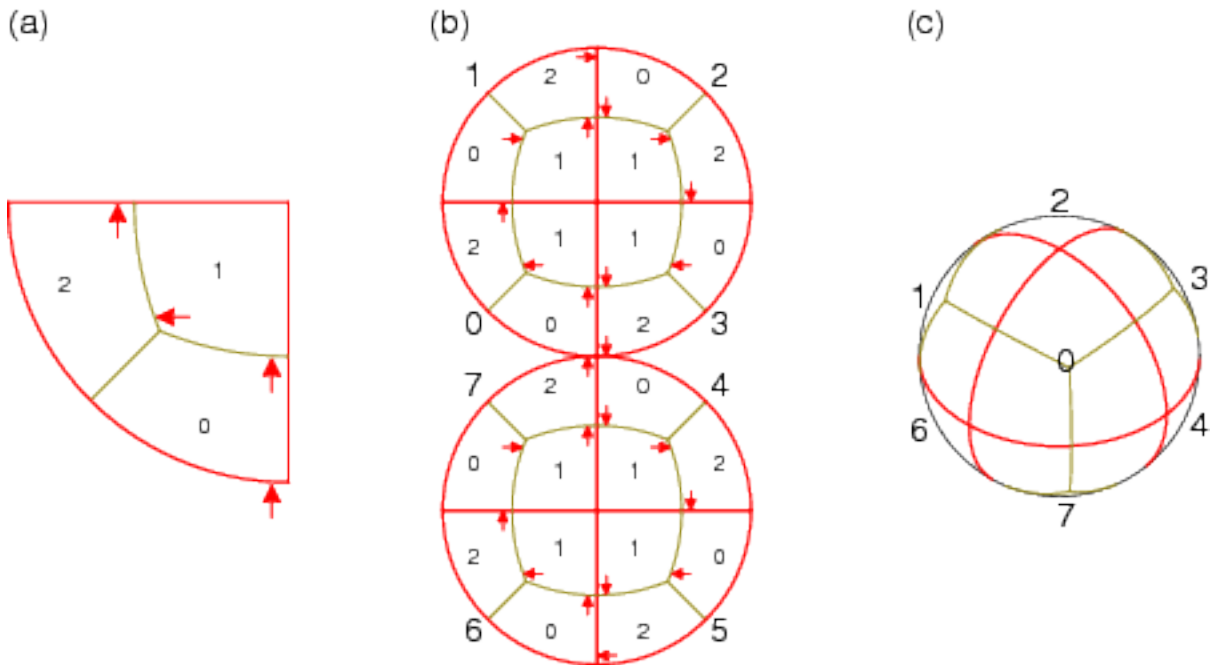


Figure 2. (a) Refinement of a representative spherical octant into its constituent triplet of quadrilaterals, whose further refinement follows the standard Hilbert curve construction in the square once the quadrilateral is suitably mapped to such a square. (b) A map showing the eight spherical octants numbered in the order in which the space-filling Hilbert curve visits them. Note that the shape of the Hilbert curve's path interior to each octant (neglecting the direction of progression and the points of ingress and egress) is perfectly mirror-symmetric about all three principal Cartesian planes when the octants are symmetrically wrapped around the sphere as shown in (c).

(b) *Sphere*

As shown in detail in ON460, the surface of a sphere can be divided first into spherical octants and then each octant further subdivided into three quadrilaterals. Fig. 2a shows one route through the three quadrilaterals of the octant labeled '0' in the map (b) displaying the arrangement where all the octants for a sphere (c) are labeled in a sequence that not only allows the Hilbert curve to fill the spherical surface, but does so in such a way that the interior portions of the curve in each octant are exact mirror-images of each other by reflection through the three principal Cartesian planes. This mirror symmetry is a valuable organizing principle to adhere to as we generalize the curve to the third dimension.

(c) *Cube*

For the unit cube, Fig. 3 shows one way to thread the eight sub-cubes, maintaining the rule that ingress and egress from each are at opposite ends of a shared edge, as well as preserving the less obvious property that a reversal of the direction along the path is geometrically a mirror-image of the original path. In order that the refinement rule for the cube be unambiguous, the path through each sub-cube must be a *proper* (unreflected) rotation of a half-sized copy of the parent cube's path that it forms a part of. (This also restricts the way we are permitted to order the sequence of eight sub-cubes in the definition of the refinement.)

sufficient number of copies of the first two layers in panel (c) of that figure.

(e) *Isometry on the sphere and aspect ratios considerations*

The geometrical method guaranteeing the maintenance of the principle of ‘isometry’ on the spherical surface is fully described, for the purely two-dimensional (2D) Hilbert curve in ON460 and, for the three-dimensional (3D) extension, the only remaining question concerns how we choose to relate the effective scales in the vertical to the effective scale in the horizontal directions.

This is resolved by prescribing, at each altitude, z , an aspect ratio, $\alpha(z) = \Delta x / \Delta z$ that we consider equitable in a comparison of characteristic dimensions of a perturbation. In atmospheric data assimilation at synoptic scales, dynamically balanced perturbations tend to exhibit a ratio of horizontal to vertical scales in roughly the same ratio of the Brunt-Väisälä frequency to the characteristic rotation rate or Coriolis frequency, in which case, $\alpha \approx 200$ is plausible in the troposphere, larger values being perhaps more appropriate in the stratosphere. Either a horizontal, or thin-shell, Hilbert curve over the spherical Earth of area, $4\pi R_e^2$ ($R_e \approx 6.371 \times 10^6$ m being the radius) corresponds, according to our custom, to 24 units of the Hilbert parameter, which implies that the horizontal scale of the Hilbert parameter unit is, $X = \sqrt{\pi/6}R_e \approx 4.61 \times 10^6$ m. The effective thickness, in Hilbert length parameter units, of a target extent of the atmosphere going from a lower altitude (which may be negative) of z_0 up to an altitude of z_1 , is

$$Z = \frac{1}{X} \int_{z_0}^{z_1} \alpha(z) dz. \quad (2.1)$$

If we suppose a constant value of $\alpha = 200$ and a vertical atmospheric extent in our target regions of 18000m, then the effective thickness of our shell is about $Z \approx 0.78$ units which, being smaller than unity, counts as a ‘thin’ shell in this example.

We should remain aware that, in order to map the Hilbert curve to the surface of a sphere in an isometric manner, the geometrical distortions incurred by the quadrilateral horizontal elements in this construction will tend to effectively broaden the distributions corresponding to Fig. 5 very slightly, but this effect is not assumed to be large enough to invalidate the methods of density estimation we are proposing in this note.

(f) *Choice of number bases*

How should we expand the leading digits of spherical Hilbert curves of these types? In general, since the coarsest subdivision is into octants, the natural choice for the leading digit is one of base-8, followed in the case of the thin-shell expansion by base-3. We adopt the convention that these provide the whole units part of the expansion, with the subsequent refinements of the principal quadrilaterals (or hexahedra) supplying the fractional part of the standard parameter definition. For the thick-shell case, the first two digits are followed by a digit of whichever base corresponds to the number of stacked principal hexahedra and *this* then completes the expansion of the ‘units’ part of the parameter. Thus, for the thin-shell examples, the expansions involve two digits left of the ‘point’; for thick-shell examples, the expansion involves three digits of mixed-base left of the ‘point’ and only octal digits right of it. In practice it is rarely expected that the thick-shell examples will be needed for ordinary applications in

the assimilation of data for weather models, but space weather applications may need to use the thick-shell versions.

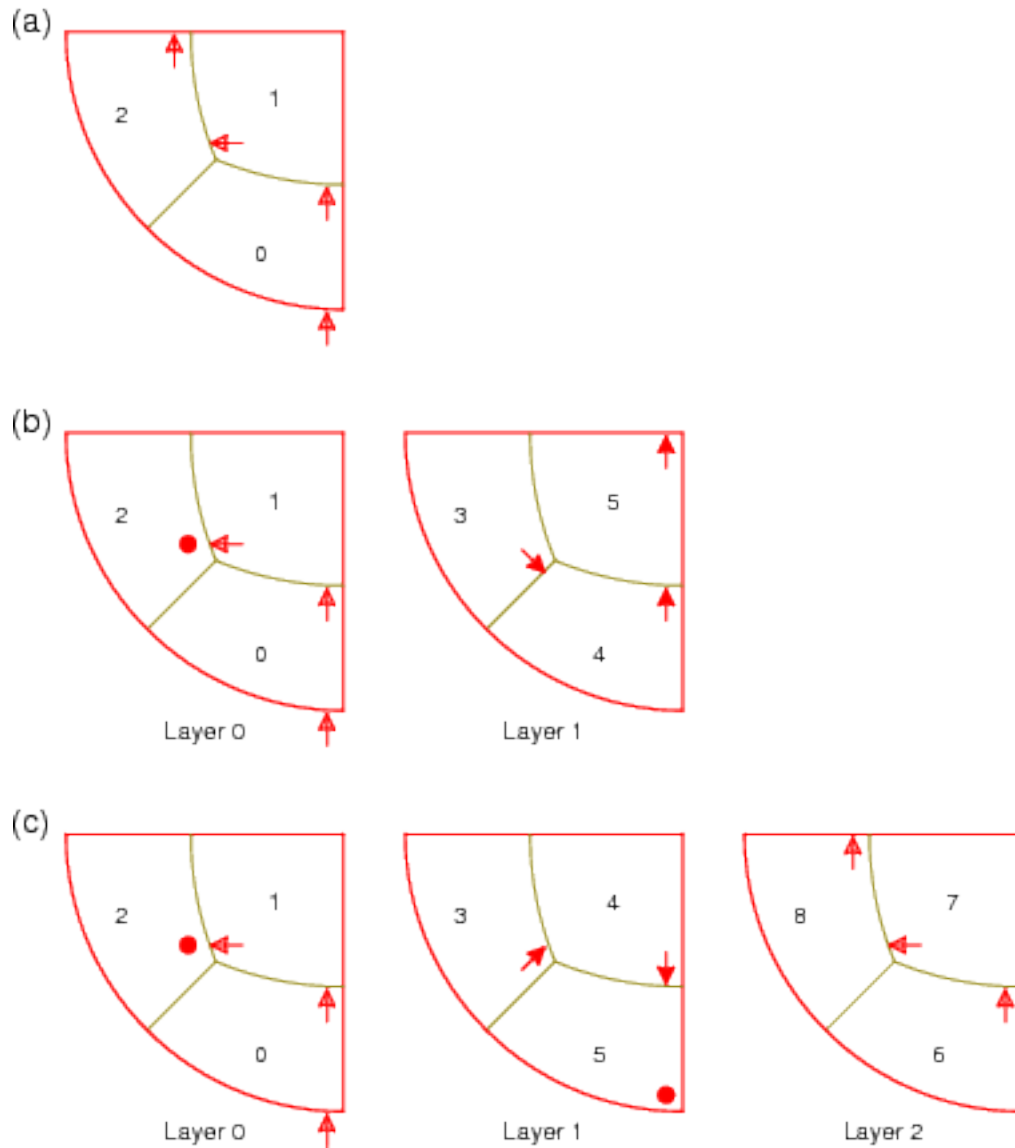


Figure 4. Refinement of a representative spherical octant shells into its constituent single (a) or multiple (b) and (c) layers of hexahedra, whose further refinement follows the standard Hilbert curve construction in the cube once the quadrilateral footprint is suitably mapped to a square.

3. ESTIMATING SPATIAL DENSITY

A generalized function comprising a unit impulse, or delta function, at each observation location, whether it be in physical space, or projected onto the Hilbert curve, will be called the ‘indicator function’. Smoothing the spatial indicator function using a homogeneous and

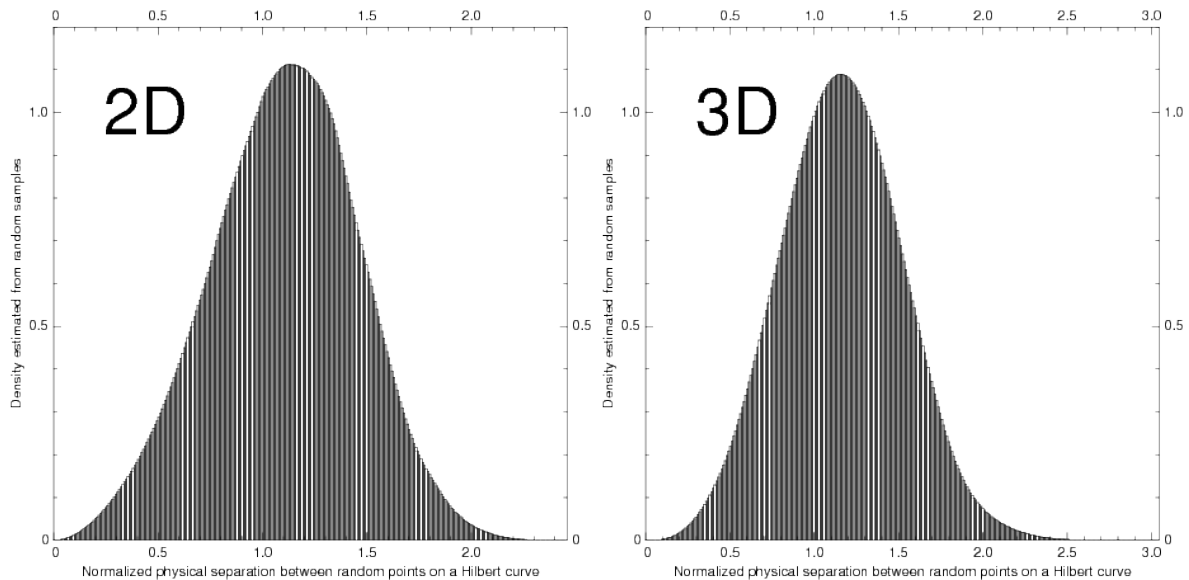


Figure 5. Two histograms that show the distributions of physical separation in two and three dimensions respectively for randomly positioned standardized intervals of the Hilbert curve parameters. In the two-dimensional case (left panel), a sharp cut-off occurs at a separation of $6^{1/2}$; in the three-dimensional case (right panel) the cut-off occurs at $28^{1/3}$.




Figure 6. The extreme standardized separation possible in a Hilbert curve based on a square or cubic frame can be understood by considering configurations, as depicted, where a succession of sub-squares or sub-cubes within the Hilbert curve construction line up with all their entry and exits on a common line, and the sizes form the double-sided series of progressive halving towards the ends as shown. In such configurations, the section of the Hilbert curve between the entry and exit arrows shown fills all the intervening squares and cubes. The spatial separation of these ends divided by the square-root (for squares) or cube-root (for cubes) of the filled space in between defines the standardized separation in both cases, and is $6^{1/2}$ for the 2D case; $28^{1/3}$ for the 3D case, as we see from summing the infinite geometrical series involved.

isotropic filter whose kernel integrates to unity produces a new spatial distribution whose value at each point represents the particular version of the locally-averaged data density we associated with that smoothing, or averaging, filter. The continuity property of the Hilbert curve means that a one-dimensional smoothing filter applied to the indicator function on this curve also produces an effective spatial density estimate of the data, but now the implied smoothing kernel in physical space is highly inhomogeneous and ragged, or ‘fractal’ in appearance. Nevertheless, for some purposes, it might be good enough, and if the computational effort involved in computing it is very modest (compared to the cost of performing the smoothing of the indicator

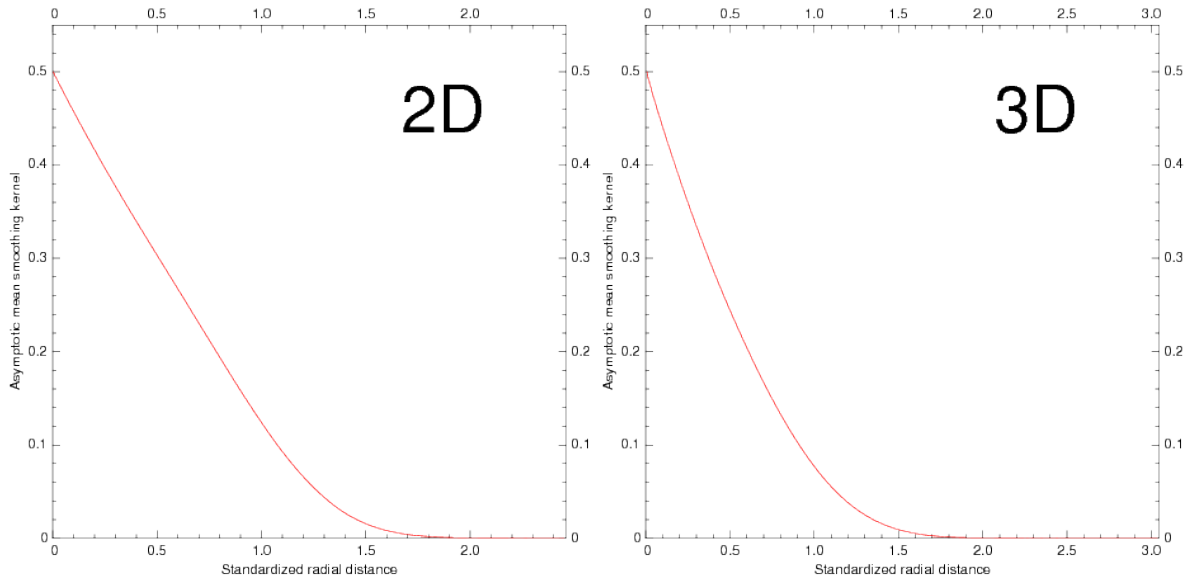


Figure 7. Radial distributions of the smoothing kernels implied by the histogram distributions of Fig. 5 when a box-car distribution of unit half-width is used to smooth the indicator function of data projected to the Hilbert curves.




Figure 8. Radial distributions of the smoothing kernels implied by the histogram distributions of Fig. 5 when a triangle distribution of unit half-width is used to smooth the indicator function of data projected to the Hilbert curves.

function on a sufficiently fine spatial grid, for example), then it is worth choosing the Hilbert curve method, which needs no grid whatsoever — just the sorted linked list of the Hilbert curve parameters of the data.

To give an idea of how faithful an approximation to a well-behaved (homogeneous, isotropic)

spatial smoother a simple Hilbert curve smoother might be, we can examine the shape of the frequency distribution of spatial separation corresponding to a unit separation in the Hilbert parameter when averaged uniformly over all orientations of the frame and, uniformly in the logarithmic sense, over all scales of the frame within the range representing a factor of two of the scaling in physical space, which is sufficiently representative given the Hilbert curve's intrinsic self-similarity at this factor of rescaling. The desired frequency distribution can be estimated empirically by Monte Carlo sampling (Hammersley and Handscomb, 1964) of random intervals of various small sizes in the interiors of square-framed and cube-framed Hilbert curves to generate histograms. Hence we obtain the frequency distributions of these standardized separations in two and three dimensions. For 5×10^9 samples, these are shown in two dimensions (2D) and three dimensions (3D) in the left and right panels of Fig. 5. By standardization, we mean that, for the case of the unit square, a Hilbert curve interval of Δs that corresponds to a spatial separation of Δx implies a standardized spatial separation of $\Delta x/\sqrt{\Delta s}$. Likewise, for the cube-framed Hilbert curve, we need to divide the actual physical separation by the cube-root of the Hilbert parameter interval. Clearly, there has to be some dispersion in these distributions, but the plots of Fig. 5 show that at least the amount of dispersion is not too large – both histograms showing relatively narrow peaks, which is encouraging.

We can take the distributions of Fig. 5, interpreted as probability densities, $B_0(\chi)$ of the respective 2D and 3D standardized separations, χ , and derive from them the implied average radial distributions of the smoothing kernels that correspond to simple local averages in the Hilbert parameters. Consider the case where the smoothing in Hilbert parameter space (neglecting end-parameter adjustments) is the convolution between the smoothing kernel function $S(s)$ with finite support $[-\hat{s}, \hat{s}]$ and the indicator function, $f(s)$:

$$\bar{f}(s) = \int_{-\hat{s}}^{\hat{s}} S(s') f(s - s') ds', \quad (3.1)$$

which, when the kernel is symmetric, becomes:

$$\bar{f}(s) = \int_0^{\hat{s}} S(s') [f(s - s') + f(s + s')] ds'. \quad (3.2)$$

Assuming $B_0(\chi)$ applies at all scales, then, given that the separation in Hilbert parameter is a general value, s , the conditional probability it implies for a spatial separation, r , is:

$$B(r|s) = \begin{cases} \frac{1}{s^{1/2}} B_0\left(\frac{x}{s^{1/2}}\right) & : \text{ in 2D} \\ \frac{1}{s^{1/3}} B_0\left(\frac{x}{s^{1/3}}\right) & : \text{ in 3D} \end{cases}. \quad (3.3)$$

The support of $B_0(\chi)$ is the open interval, $(0, \hat{\chi})$, with $\hat{\chi} = 6^{1/2}$ in 2D; $\hat{\chi} = 28^{1/3}$ in 3D, the reason for these values being explained in the caption of Fig. 6. The averaged implied radial kernel distribution in physical space corresponding to the symmetric kernel, S , is given by the integral:

$$R_{(2D)}(x) = 2 \int_0^{6^{1/2}} B_0(\chi) S\left(x^2/\chi^2\right) \frac{d\chi}{(\pi\chi^2)}, \quad (3.4)$$

in 2D, and:

$$R_{(3D)}(x) = 2 \int_0^{28^{1/3}} B_0(\chi) S\left(x^3/\chi^3\right) \frac{d\chi}{\left(\frac{4\pi\chi^3}{3}\right)}, \quad (3.5)$$

in 3D, where the factor of two before each integral comes from the two symmetric halves of the kernel, S , and the separations defined by (3.3) are taken to be equally likely in all possible directions.

The simplest conceptual and computational example of a smoother in the Hilbert parameter space is the uniform-profile ‘box-car’ moving average of half-width Δs , for which,

$$S(s) = \begin{cases} \frac{1}{2\Delta s} & : |s| < \Delta s \\ 0. & : |s| > \Delta s \end{cases}. \quad (3.6)$$

Applied to the indicator function, the evaluation of this smoother at each data point amounts simply to a count of this datum and the other data within a parameter distance, Δs , all divided by the interval width, $2\Delta s$. This can be computed very rapidly for all the data once they have been ordered into a linked list. Fig. 7 shows the corresponding shapes of these radial profiles in 2D and 3D. Another relatively simple smoother, which can also be computed efficiently, is the ‘triangle’ smoother. If we scale its kernel to also fit inside the half-width Δs , it is defined:

$$S(s) = \begin{cases} \frac{1-s/\Delta s}{\Delta s} & : |s| < \Delta s \\ 0. & : |s| > \Delta s \end{cases}. \quad (3.7)$$

The implied radial distributions of averaged kernels in 2D and 3D for this case are shown in Fig 8. Even allowing for the fact that the second moment in the s -domain of the triangle smoother is only a half of that of the box-car smoother of the same Δs , as we have defined them, the effect of the triangle smoother is to further exacerbate the acuteness of the peaks implied in the averaged spatial kernels’ radial distributions. Thus, while it is in principle quite straightforward to further generalize these smoothing kernels to higher-degree piecewise polynomials for $S(s)$ based on the family of uniform-grid B-spline basis functions (e.g., de Boer 1978), which become progressively more Gaussian in shape (but only in s -space) as the chosen degree increases, the corresponding further changes in the qualitative shapes of the implied spatial kernels’ radial profiles are probably not in the desired direction. Instead, if our goal is to formulate Hilbert curve-based local averages whose radial kernels exhibit flatter, more plateau-like profiles, it might be better to combine two (or more) of the simple box-car Hilbert-space smoothers but with differing scales and weighted both positively and negatively. In Fig. 9 we show how a pair of moving averages on the Hilbert curve, with characteristic scales differing there by a factor of two, can be combined with appropriate positive and negative weights to make the corresponding averaged kernels in the spatial domain have the radial profiles with the relatively flat peaks, as the two panels show.

Any of these filters should be computationally efficient to apply but, owing to the stochastic nature of this data density estimation technique, the results from a single Hilbert curve will be

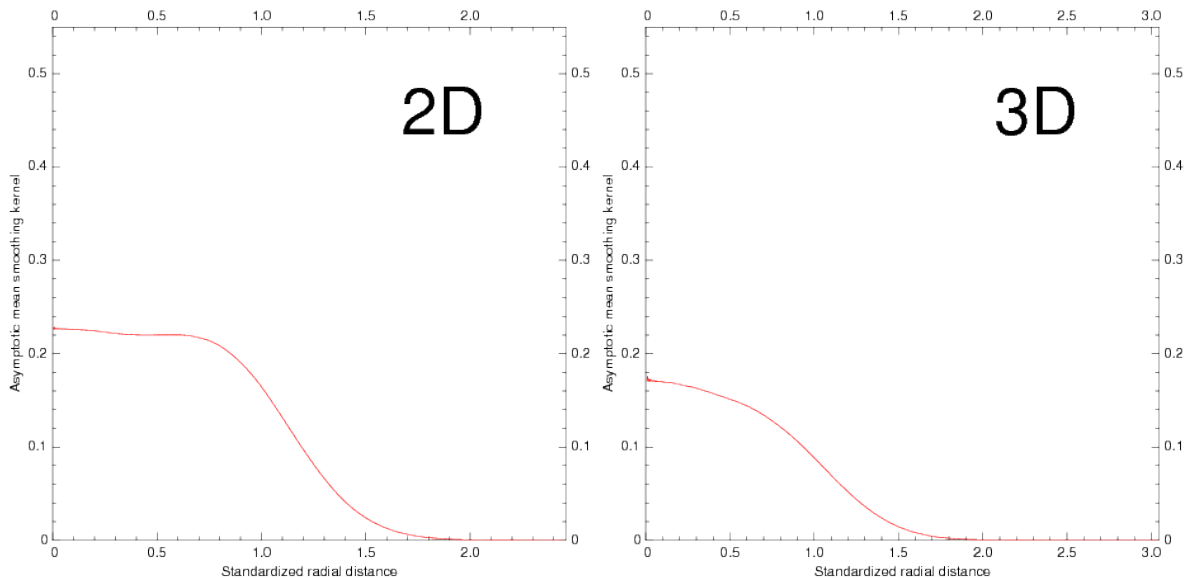


Figure 9. Radial distributions of the smoothing kernels implied by the histogram distributions of Fig. 5 when pairs of box-car distributions, with scales in the Hilbert domain differing by a factor of two, are combined with opposite signs and weights chosen to make the implied radial distribution of the averaged kernels in physical space have the form of ‘plateaux’ instead of the acute peaks implied in the case of the single box-car.

subject to a large amount of sampling randomness. This is especially true in the cases, such as those illustrated in Fig. 9, where the smoothing combines two weighted box-car averages with opposite signs. But since the method is expected to be computationally cheap to apply, we can repeat the estimation several times with differently oriented constructions of the Hilbert curves in each instance, and simply average the resulting independent estimates of data density. The following section indicates how the Hilbert curves can be ‘randomized’ in systematic ways to ensure that the estimates that come out of the process truly are diverse and are unlikely to possess any correlated sampling errors.

4. SYSTEMATIC ‘RANDOMIZATION’ OF HILBERT CURVE CONSTRUCTIONS

Consider a simulated data array, very crudely mimicking the locations one might expect of aircraft reports, constructed to exhibit clumps (terminal hubs) and concentrated tracks (flight paths) between them, such as that shown in Fig. 10a, and an estimate of the local density of these data that one might obtain by smoothing the indicator functions of these data using a simple Gaussian smoother, resulting in the contours shown in that figure. In practice, we are only interested in the estimated values of the local data density at the locations of the data themselves, so that we can use this information to adjust, by judicious down-weighting, the data diagnosed to be within an excessively dense, over-sampled region. While the Gaussian shape for a smoothing kernel is extremely convenient to choose when the smoothing is done by conventional filtering operations on a regular Cartesian grid, we have seen in the previous section that, in the case of estimates that emerge from the use of the Hilbert curve projection and smoothing technique, the shapes of the averages of orientation-averaged kernels implied by

an ensemble of Hilbert curves and simple moving-average box-car or triangle filters are more sharply peaked than the Gaussian, and are of finite support (contained within a small finite radius), so we would not expect the density estimates from the two methods to match exactly. Nevertheless, we should expect that there is a good correlation discernible between the two methods of estimation, even though there will be a certain amount of stochastic scatter in the estimates obtained with the Hilbert curve method. For a single Hilbert curve, Fig. 10b shows a scatter plot of the Hilbert curve density estimates against the Gaussian bell smoother estimates, showing that there is indeed a striking correlation, as well as the expected stochastic scatter.

We randomize the Hilbert curve framework by rotating the basic square, made large enough to contain the data domain at all angles of rotation, where the angles of rotation are exact subdivisions of 90° (since a rotation by 90° just brings the principal directions of the square framework back into alignment again). The axis of this rotation should not be the dead center of the square since, by the manner in which the square is subdivided to construct the Hilbert curve, this point is a very special one with the property that nearby points on either side of it are essentially *always* going to be found far apart on the Hilbert curve. For this reason, we choose a pivot point for the frame rotations located at a point without a terminating binary expansion of its coordinates, such as the point, $(1/3, 1/3)$. Fig. 10c shows the very significant reduction in the scatter we obtain with 10 randomizations where, again, the estimates of density from the Hilbert curve are compared with the direct estimates that employ a Gaussian filter on a fine grid. Lastly, in Fig. 10d we make the corresponding comparison between the average of the same 10 randomizations with an alternative benchmark comprising the average of 1000 randomizations, showing that there is a significant further sharpening of the correlation compared with panel (c) where the benchmark was with the two-dimensional Gaussian filter. As discussed in the previous section, we can attribute this further narrowing of the scatter to the fact that the effective averaged radial profile of the smoothing kernel of the Hilbert curve method is *not* a Gaussian shape; if we insist on using the Gaussian smoother as the benchmark, the scatter should therefore *not* be expected to converge to a line in the asymptotic limit of averaging ever larger numbers of randomized rotations of the Hilbert curve.

In three dimensions, we can carry out a corresponding exercise of comparing the Hilbert curve-based density estimate with that obtained from a Gaussian filter. In this case, we contour only the vertically-accumulated density for the randomly generated data shown in Fig. 11a. The scatter-plot comparison of panel (b) shows that, if we randomize the frame of the 3D Hilbert curve by 10 rotations only in the horizontal plane, as was adequate for the 2D Hilbert curve, we do not achieve the same disciplined concentration about a straight line, but rather an untidy scatter that persists even for much larger numbers of randomized rotations. However, if we randomize the frame of the Hilbert curves by rotations that are fully three dimensional, we begin to see the points gather more approximately towards a line (Fig. 11c). Lastly, panel (d) shows that, once again, a tighter scatter is obtained when the density of the average from 10 three-dimensionally rotated Hilbert curves is compared against the more compatible benchmark obtained from an average of 1000 randomly rotated Hilbert curves. One important but unavoidable difference in the three dimensional case is that there is not a way to ‘randomize’ the rotations three-dimensionally in a perfectly uniform distribution of orientations; the randomizations of panels (c) and (d) use truly independent random orientations, which is not optimal since, by chance, the resulting orientations are not all guaranteed to be significantly different

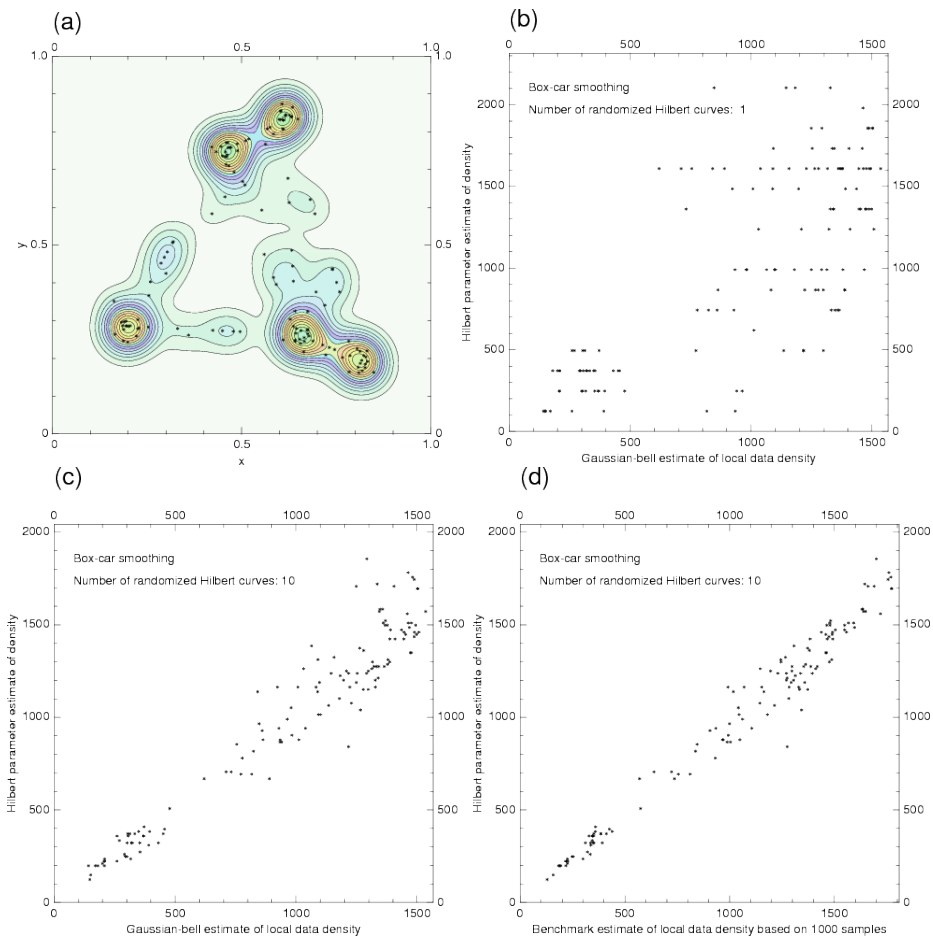


Figure 10. Idealized simulation of the process of estimating the density of data in 2D that can be contained within all members of a set of mutually rotated Hilbert curves constructed within square frames. Panel (a) shows the disposition of the clustered data and contours of local density estimated using a Gaussian-kernel spatial smoother directly. When the estimates of density at roughly the same characteristic horizontal scale of averaging are obtained by a single Hilbert curve projection with box-car smoothing there of the indicator function, we obtain an estimate of density at each data point that correlates roughly with the corresponding estimate that the Gaussian smoother delivers (b). By averaging 10 such Hilbert curve estimates from frames rotated every 9° , the corresponding scatter starts to coalesce towards a line (c), but not as much as it does when the benchmark is taken to be the average of 1000 rotated Hilbert curve estimates (d) instead of the Gaussian filter, since the orientation-averages Hilbert curve estimates do not imply an exactly Gaussian radial profile for the kernel in physical space.

from each other. Also, it is not actually feasible to swivel and tilt the frame of the Hilbert curve generated within a spherical shell in a fully three-dimensional way, but this example does at least suggest that, in the spherical shell example, we should find some way to ensure that the randomization of the frame exercises perturbations of it in all three dimensions at each location.

For the spherical shell, the randomization can be done by combining, for each realization, a pseudo-random rigid rotation of the spherical surface with a globally-uniform vertical displacement that is different for each realization. In practice, it seems easiest to do this vertical

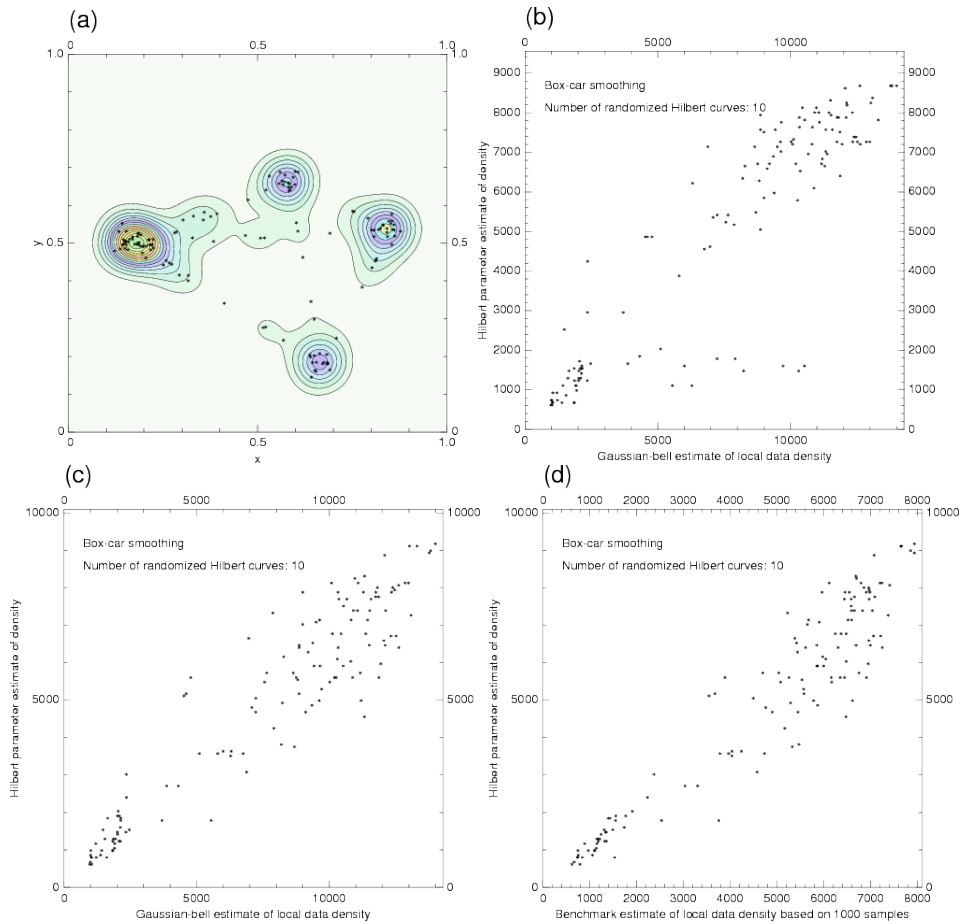


Figure 11. A simulation of data density estimation in a limited 3D domain contained within the common interior of rotated cubes. The vertically integrated (projected) data density obtained from a Gaussian smoother on a fine three-dimensional grid, and the clustered data themselves, are shown in panel (a). Panel (b) shows, using a scatter plot, how the estimates of data density averaged from 10 Hilbert curve frames are rotated only in the horizontal plane, are compared against the benchmark defined by the Gaussian smoother, and panel (c) shows the improved match obtained when the randomization is done by fully three-dimensional rotations of the Hilbert curve frame. The effect of the mismatch between the implied radial profile of the averaged Hilbert curve kernels and the Gaussian profile shows up as the small further improvement (d) in the fit to the alternative benchmark constructed in a compatible manner by a large ensemble (1000) of the same Hilbert curves rotated randomly in three dimensions.

contribution of the randomization by slightly expanding the thickness (in Hilbert scale units) of the shell so that the target region fits within it with some slack. Then for each realization, this target region can be transformed at slightly different vertical displacements to ensure that the horizontal binary partitions in the Hilbert curve construction occur at different physical altitudes for each realization. The amounts of these vertical shifts should not have terminating binary expansions in the Hilbert scale units. Fig. 12 illustrates the effectiveness of this strategy for a simulated clumped distribution of data at different altitudes within a thin spherical shell, smoothed in the horizontal and vertical with a Gaussian smoother (panel (a)) to provide an initial benchmark. In (b) we see how the average of 10 randomized (in the horizontal and vertical) Hilbert curve projection estimates of data density compare with the Gaussian-based

benchmark. A fairer comparison in (c) shows the corresponding scatter we obtain when the benchmark estimate is replaced by the average of 1000 randomized compatible Hilbert curve estimates, with the tighter scatter indicating much greater consistency, as we should expect.

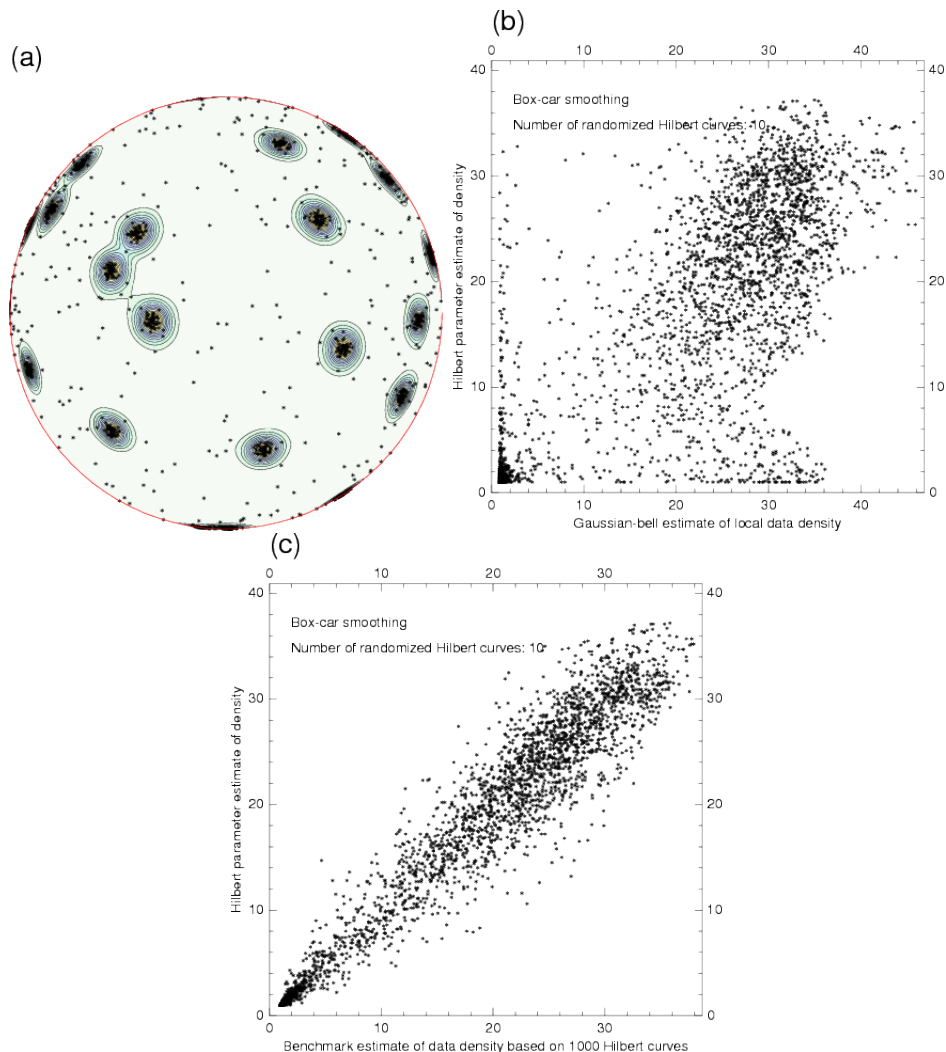


Figure 12. A simulated clustered distribution of data in a thin shell wrapping the sphere (a) together with an associated estimate of data density obtained by smoothing the indicator function of the data with a grid-based numerical filter whose kernel very closely approximates the Gaussian profile in both the horizontal and vertical. The contours show only the vertically integrated density over the obverse surface. Panel (b) shows a comparison between the Hilbert curve estimates of data density, using 10 randomizations, and the estimates that come from the Gaussian-bell smoother used to construct the contours in (a). The two methods have implied smoothing kernels of different shapes and presumably respond to the upper boundaries and lower boundaries of the shell in a different way, so it is not surprising that there is a relatively large scatter here. Panel (c) confirms that, when the 10-randomization estimates are compared with a compatible benchmark formed using 1000 randomizations of the Hilbert curve, the results exhibit the expected consistency with small residual scatter.

For the horizontal component of the shell rotations, we can expect to do better than simply choosing the orientations randomly. Recognizing that the global construction of the Hilbert curve is, at its coarsest scale, based on a cubical/octahedral decomposition of the spherical

surface, we restrict our search for suitable rigid rotations of the spherical frame that transformed that inscribed cube defining the frame to distinguishable inequivalent images. Also, for a small ensemble of the rotations we strive to ensure not only that each pair of rotations is inequivalent, but that, collectively, the set of chosen orientations are in some sense as far away from each other as possible. In our previous note, Purser (2017, henceforth ON489), we showed that a quaternionic analysis of the spaces of orientations of any symmetric polyhedra, including the cube that is relevant here, provides a valuable way of looking at problems of this kind.

A general proper rotation of any object in 3D space about a fixed point has a natural identification with an antipodal pair of unit quaternions. Such a quaternion can be thought of as a unit 4-vector with components indexed from 0 to 3, the zeroth component being the ‘real’ part and the other components being ‘imaginary’. The null rotation of an object is identified by the identity unit quaternion, $(1, 0, 0, 0)$, together with (as always for these rotations) its negative, $(-1, 0, 0, 0)$. Consider the rotation of the object from its standard (identity) orientation to a new one obtained by rotating it by an angle of θ in the right-handed sense about the axis pointing in the direction of the 3D Cartesian unit-vector, (x, y, z) . Let $c = \cos(\theta/2)$ and $s = \sin(\theta/2)$. Then the equivalent representation of this rotation by a quaternion pair q and $-q$ has one member defined:

$$\mathbf{q} = (c, sx, sy, sz). \tag{4.1}$$

As an algebraic system, quaternions are equipped with the ordinary vector rules of addition together with rules for noncommutative multiplication which, in the case of unit quaternions, and apart from the aforementioned duplication, match the rules of compounding the corresponding proper rotations (i.e., by 3×3 orthogonal matrices) of rigid 3D objects. (An explicit definition of the quaternion multiplative rules is provided in ON489.) It is easier to find the angle between two general 3D spatial orientations by referring to their quaternion representations than by their more obvious representations as orthogonal 3×3 matrices; the angular distance between the quaternions (along a great-circle of the 3-sphere of unit quaternions) corresponds to half the angular separation between the two spatial orientations they represent.




Figure 13. Schematic view of the Euclidian (tangent-hyperplane) radial projection of the Voronoi cell associated with the central representative quaternion that belongs to the binary-octahedral group representation. 48 such congruent cells completely tile the 3-sphere of unit quaternions, each face coinciding with a left-handed twisted image of its opposite face in the way that causes the labels, A–F, to correspond.

The 24 proper rotations that map the cube in its standard orientation into itself form a

group – the *octahedral group*. A double-cover of this group, the *binary-octahedral group*, is represented by the 48 associated unit quaternions of these special rotations. The geometry of this configuration is described in detail in ON489. We construct the Voronoi decomposition of the 3-sphere hypersurface of the unit quaternions associated with these 48 binary-octahedral quaternions according to the rule that a quaternion, \mathbf{q} , belongs to the Voronoi cell of \mathbf{q}_a when it is closer to \mathbf{q}_a than to any of the other 47 representatives of the group, or at least, no farther away, if we include the boundary of the cell in the definition of ‘inclusion’. The 48 congruent cells thus defined have the geometrical shape, when projected onto the 3-sphere’s tangent-hyperplane at their central quaternion, \mathbf{q}_a , of a truncated cube – the shape of the Archimedean solid with cubic/octahedral symmetry comprising six regular octagons and eight equilateral triangles, as shown in the sketch of Fig. 13 taken from ON489. Any orientation of the cube can now be obtained by rotating it from its standard orientation by a rotation operator that corresponds to a unit quaternion that belongs to the interior, or a nonredundant half of the boundary, of the Voronoi cell that has the identity as its center – the *principal Voronoi cell*. But that particular quaternion has 47 other images distributed amongst each of the other Voronoi cells that represent the equivalent cube orientation (these images are obtained from the principal one by right-multiplying it by the quaternion elements of the binary-octahedral group). So the problem of finding the most diverse set of N inequivalent rotations of the cube boils down to finding, within, or on a nonredundant half of the boundary of, the principal Voronoi cell, the N unit quaternions that maximize the minimum distance between each other and to their images in the above sense. Geometrically, this is a kind of optimal sphere-packing problem, but in the curved space defined by the portion of the unit 3-sphere occupied by the Voronoi cell and with the implied topological connectivity by which opposite octahedral faces map to each other (with a left-handed rotation of 45°) and opposite triangular faces likewise map to each other (but with a left-handed twist of 60°) in the manner one can deduce from the labeled vertices of Fig. 13. For $N \leq 5$ the solutions to this puzzle can be given analytically and reasonably simply, and are provided in the appendix. Analytic solutions also exist for the special cases, $N = 7$ and $N = 13$, but for other N there do not seem to be solutions amenable to easy algebraic expression. However, as discussed in the appendix, we can formulate larger sets of distinct rotations of the cube, up to $N = 273$, by compounding the rotations in the special analytic families that have $N = 3$, $N = 7$ and $N = 13$ (although these are not generally optimal). These sets of up to $N = 273$ rotations of the Cartesian frames of 3D Hilbert curves serve as the nearest equivalents to the uniform planar rotations we use for the 2D Hilbert curves. They enable the ‘randomization’ to be performed in a formulaic manner; they conveniently avoid requiring the generation of truly random realizations via the more conventional ‘Monte Carlo’ pseudo-random number method (Hammersley and Handscomb, 1964), which would have the unattractive feature of not generally giving reproducible results over different computing platforms (since the pseudo-random numbers generated would not be expected to be identical).

5. DISCUSSION AND CONCLUSIONS

We have described how space-filling curves can serve to help with the efficient estimation of data density in both two and three dimensions, in both limited and spherical domains. The proposed method is intended for use when the atmospheric data whose density is being

estimated tend to exhibit a marked degree of tight clustering, such as the global distribution of aircraft reports (tightly clustered around major airport hubs), satellite cloud-motion winds (clustered in regions where the appropriate clouds are found), and mesonet data (clustered in heavily-populated urban areas). For data with a more uniform coverage, other methods that employ filtering on a regular grid would be expected to be more efficient.

The cost of projecting data to a Hilbert curve scales in proportion to the number, M of data, although the proportionality factor itself scales as the logarithm of the desired resolution of the implied Hilbert curve. The cost of sorting (by a binary-sorting procedure) M data rises as $M \log M$. The cost of smoothing the indicator function for M sorted data is about M . The attractive feature of this method is that there is no cost incurred by the existence of large expanses of empty space devoid of data, which occurs for many data types (including those we have mentioned). Since the final smoothing of the sorted data locations is relatively cheap, it should be inexpensive to refine the simple moving-average filters based on the triangle smoother and, especially, the very simple box-car smoother, and synthesize from these building blocks a filter with a more interesting profile. However, the random scatter we observed when the filter estimates were plotted in comparison with the more robust benchmark values, will tend to be exacerbated for most compound filters that employ a mixture of positive and negative combinations of the simpler building blocks, so this is something to take into consideration.

ACKNOWLEDGMENTS

This note has benefitted from several stimulating discussions with Drs. Andrew Collard, John Derber, Xiujuan Su and Yanqiu Zhu, and the internal reviews of Manuel Pondeva and Jun Wang.

APPENDIX A

Optimal sets of rotations of a cubic framework

If we adopt the criterion that the orientations of the N cubic frameworks in which the Hilbert curve is constructed be separated (in angle) from each other as much as possible, then for sets of N up to five, as well as for $N = 7$ and $N = 13$, the optimal solutions possess some symmetries and their quaternionic representations have tractable analytic solutions. We list these representations here, relative to the standard orientation of one (the first in each list) of these cubes. In each case, the quaternion representation of each cube's orientation lie within (or on the surface of) the Voronoi fundamental region of the 3-sphere of unit quaternions centered on the identity. The shape of that region is a truncated cube (six octagonal and eight triangular faces, all edges equal), as discussed and illustrated in ON489.

For the case of $N = 1$, the representative unit quaternion is trivially the identity itself:

$$\mathbf{q}_{(1)} = [1, 0, 0, 0]. \tag{A.1}$$

For the case $N = 2$, the identity is augmented with one of the similar representations of the most distant point from it that still belongs to the fundamental region – that is one of the 24 vertices of that Voronoi polyhedron. An example of such a unit quaternion is:

$$\mathbf{q}_{(2)} = \left[\frac{1}{2} + 1/\sqrt{8}, -1/\sqrt{8}, -\frac{1}{2} + 1/\sqrt{8}, 1/\sqrt{8} \right], \tag{A.2}$$

and, in the case where $N = 3$, we need only augment this point by a third unit quaternion,

$$\mathbf{q}_{(2)} = \left[\frac{1}{2} + 1/\sqrt{8}, 1/\sqrt{8}, -\frac{1}{2} + 1/\sqrt{8}, -1/\sqrt{8} \right], \quad (\text{A.3})$$

since the three cube orientations implied by these three quaternions are equally distant from each other in terms of the angular distance between them. These three implied cubes are generated by rotations of 120° about any one of the four axes defined in the Cartesian frame of the standard cube: $[\pm\sqrt{2/3}, 0, \sqrt{1/3}]$ or $[0, \pm\sqrt{2/3}, \sqrt{1/3}]$.

For the case, $N = 4$, we define:

$$a = (4 + 6\sqrt{2})/14 \quad (\text{A.4a})$$

$$b = (4 - \sqrt{2})/14 \quad (\text{A.4b})$$

and, in addition to the identity, the three unit quaternions:

$$\mathbf{q}_{(2)} = [a, 2b, -b, b] \quad (\text{A.5a})$$

$$\mathbf{q}_{(3)} = [a, b, 2b, -b] \quad (\text{A.5b})$$

$$\mathbf{q}_{(4)} = [a, -b, b, 2b]. \quad (\text{A.5c})$$

This quartet forms a tetrahedral arrangement. This is made obvious by rotating to the equivalent more symmetric, presentation:

$$\mathbf{q}_{(1)} = [a', b', b', b'] \quad (\text{A.6a})$$

$$\mathbf{q}_{(2)} = [a', b', -b', -b'] \quad (\text{A.6b})$$

$$\mathbf{q}_{(3)} = [a', -b', b', -b'] \quad (\text{A.6c})$$

$$\mathbf{q}_{(4)} = [a', -b', -b', b'], \quad (\text{A.6d})$$

where:

$$a' = \sqrt{(13 + 9\sqrt{2})/28} \quad (\text{A.5.7a})$$

$$b' = \sqrt{(5 - 3\sqrt{2})/28}. \quad (\text{A.5.7b})$$

The case, $N = 5$, has a highly symmetric, but as it turns out, a suboptimal, solution that was discussed in ON489; the five distinct ways that cubes can be inscribed in a regular dodecahedron while sharing the latter's vertices. In the quaternion notation, this icosahedrally-symmetric case has the non-identity quaternions defined in the fundamental region:

$$\mathbf{q}_{(2)} = [c, d, d, d] \quad (\text{A.8a})$$

$$\mathbf{q}_{(3)} = [c, d, -d, -d] \quad (\text{A.8b})$$

$$\mathbf{q}_{(4)} = [c, -d, d, -d] \quad (\text{A.8c})$$

$$\mathbf{q}_{(5)} = [c, -d, -d, d], \quad (\text{A.8d})$$

where

$$c = (\phi + 1)/\sqrt{8} \quad (\text{A.9a})$$

$$d = (\phi - 1)/\sqrt{8}, \quad (\text{A.9b})$$

where $\phi = (\sqrt{5} + 1)/2$ is the ‘Golden ratio’. But these cubes are closer to one another (in the rotational sense) than they need to be. The optimal arrangement, which is not as symmetrical, is one whose quaternionic representation (in addition to the usual identity) is of the form

$$\mathbf{q}_{(2)} = [e, f, -g, h] \quad (\text{A.10a})$$

$$\mathbf{q}_{(3)} = [e, -j, k, l] \quad (\text{A.10b})$$

$$\mathbf{q}_{(4)} = [e, k, j, -l] \quad (\text{A.10c})$$

$$\mathbf{q}_{(5)} = [e, g, -h, -f], \quad (\text{A.10d})$$

where, in terms of ϕ , $\sigma = 1/\phi$, and $\chi = \sqrt{2} - 1$, the new coefficients are defined:

$$e = 1/(2\sigma + \chi - \sqrt{4\sigma + 2\chi - 3}) \quad (\text{A.11a})$$

$$f = \chi e \quad (\text{A.11b})$$

$$g = -1 + (-1 + \sqrt{2}\phi)e \quad (\text{A.11c})$$

$$h = -\phi + (1 + \sqrt{2}\sigma)e \quad (\text{A.11d})$$

$$j = \sigma/\sqrt{2} + (-1 + \sqrt{2}\sigma)e \quad (\text{A.11e})$$

$$k = 1/\sqrt{2} + (-1 + \sqrt{2}\chi\sigma)e \quad (\text{A.11f})$$

$$l = \phi/\sqrt{2} - e. \quad (\text{A.11g})$$

In this case, a geometrical characterization of the frames implied by the \mathbf{q} is that they form a cycle obtained by successive rotations of 72° about an axis defined by the (unnormalized) Cartesian vector,

$$v = [e + f - g - h, e - f - g + h, e + f + g + h]. \quad (\text{A.12})$$

In the case where $N = 7$ we start by defining the quantities:

$$K_1 = 2 \cos(\pi/7) \quad (\text{A.13a})$$

$$K_2 = 2 \cos(2\pi/7) \quad (\text{A.13b})$$

$$K_3 = 2 \cos(3\pi/7), \quad (\text{A.13c})$$

and define the two coefficients, a and b , of the quadratic equation:

$$aw^2 - 2bw + 1 = 0, \quad (\text{A.14})$$

by:

$$a = 4(-5 + 6K_1 - 2K_2 + \sqrt{2}(K_1 - K_2 - 2)) \quad (\text{A.15a})$$

$$b = 1/K_3 + \sqrt{2}K_3, \quad (\text{A.15b})$$

Taking the larger root,

$$w = \frac{b + \sqrt{b^2 - a}}{a} \approx 0.914163, \quad (\text{A.16})$$

we define x , y and z by:

$$x = \sqrt{2}w - K_2/2 \approx 0.669333 \quad (\text{A.17a})$$

$$y = (K_1 - 2)\sqrt{2}w + K_2/2 \approx 0.367430 \quad (\text{A.17b})$$

$$z = (2 + (1 - K_1)\sqrt{2})w + K_2/2 \approx 0.168074. \quad (\text{A.17c})$$

The desired set of seven cube rotations form a symmetric cycle, with adjacent members separated by a relative rotation of $2\pi/7$ about an oblique axis, and are encoded by the usual identity, together with following quaternions:

$$\mathbf{q}_{(2)} = [K_1/2, -z/K_1, -x/K_1, y/K_1] \quad (\text{A.18a})$$

$$\mathbf{q}_{(3)} = [w, -(y+z)/\sqrt{2}, (K_2/2 - x)/\sqrt{2}, (y-z)/\sqrt{2}] \quad (\text{A.18b})$$

$$\mathbf{q}_{(4)} = [w, K_2(x-y)/\sqrt{2}, -(K_2x - K_3/2)/\sqrt{2}, -(K_2z + K_3/2)/\sqrt{2}] \quad (\text{A.18c})$$

$$\mathbf{q}_{(5)} = [w, K_2(x-y)/\sqrt{2}, -(K_2z + K_3/2)/\sqrt{2}, -(K_2z - K_3/2)/\sqrt{2}] \quad (\text{A.18d})$$

$$\mathbf{q}_{(6)} = [w, (z-y)/\sqrt{2}, -(K_2/2 - x)/\sqrt{2}, -(y+z)/\sqrt{2}] \quad (\text{A.18e})$$

$$\mathbf{q}_{(7)} = [K_1/2, z/K_1, x/K_1, -y/K_1]. \quad (\text{A.18f})$$

The case $N = 13$ yields an optimal configurations that retains the full symmetry of the canonical cube. First we define the following intermediate terms:

$$T_1 = 29\sqrt{2} - 1 \quad (\text{A.19a})$$

$$T_2 = (11 + 9\sqrt{2} + 2\sqrt{T_1})/41 \quad (\text{A.19b})$$

$$T_3 = (30 - 9\sqrt{2} - 2\sqrt{T_1})/82, \quad (\text{A.19c})$$

Now take,

$$\tau = \sqrt{T_2} \quad (\text{A.20a})$$

$$v = \sqrt{T_3}, \quad (\text{A.20b})$$

and obtain the set consisting of the identity, $\mathbf{q}_{(1)}$, together with the remaining twelve, $\mathbf{q}_{(2)}$ — $\mathbf{q}_{(13)}$, forming the set:

$$[\tau, 0, \pm v, \pm v] \cup [\tau, \pm v, 0, \pm v] \cup [\tau, \pm v, \pm v, 0]. \quad (\text{A.21})$$

Remarks

The solutions given above for the cases $N = 3$, $N = 7$ and $N = 13$ are particularly efficient at distancing the implied member cubes from one another. But since these quaternion sets, regarded as groups of geometrical rotations, share no coinciding subgroup representations other than the trivial identity, we can compound pairs or triplets, one member from each set, to form new sets of distinct, though sub-optimal, implied rotations of the cube – as many as 273 such quaternion products. Using the superscript indices to denote the sets from which each factor is taken, we find the good results are obtained, for sets numbering more than 13, but not exceeding 91, by the product quaternions:

$$\mathbf{q} \in \left\{ \mathbf{q}^{(7)} \times \mathbf{q}^{(13)} \right\} \quad (\text{A.22})$$

and for sets numbering more than 91, and not exceeding 273,

$$\mathbf{q} = \in \left\{ \mathbf{q}^{(13)} \times \mathbf{q}^{(3)} \times \mathbf{q}^{(7)} \right\}. \quad (\text{A.23})$$

REFERENCES

- | | | |
|--|------|--|
| Behrens, J., and J. Zimmermann | 2000 | Paralleling an unstructured grid generator with a space-filling curve approach. <i>Lecture Notes in Computer Science</i> Springer, Berlin, pp. 815–823. |
| de Boor, C. | 1978 | <i>A Practical Guide to Splines</i> . Springer, New York. 392 pp. |
| Dennis, J. M. | 2003 | Partitioning with space-filling curves on the cubed-sphere. In: <i>Proceedings of the International Parallel and Distributed Processing Symposium., April 2003</i> , IEEE. |
| Hammersley, J. M., and D. C. Handscomb | 1964 | <i>Monte Carlo Methods</i> . Methuen and Co., London. 178 pp. |
| Hilbert, D. | 1891 | Über die stetige Abbildung einer Linie auf ein Flächenstück. <i>Math. Ann.</i> , 38 , 459–460. |
| Peano, G. | 1890 | Sur une courbe, qui remplit toute une aire plane. <i>Math. Ann.</i> , 36 , 157–160. doi: 10.1007/BF01199438. |
| De Pondeca, M. S. F. V., and Coauthors | 2011 | The Real-Time Mesoscale Analysis at NOAAs National Centers for Environmental Prediction: Current Status and Development. <i>Wea. Forecasting</i> , 26 , 593-612. |
| Purser, R. J. | 2017 | Sets of optimally diversified polyhedral orientations. NOAA/NCEP Office Note 489. |
| Purser, R. J., M. de Pondeca, and S.-Y. Park | 2009 | Construction of a Hilbert curve on the sphere with an isometric parametrization of area. NOAA/NCEP Office Note 460. |
| Purser, R. J., and Y. Zhu | 2016 | Comparison of finite differencing, time smoothing and spline fitting algorithms for estimating airspeed metadata from coarse-resolution aircraft position and wind reports. NOAA/NCEP Office Note 485. |
| Sagan, H. | 1994 | <i>Space-filling curves</i> Springer-Verlag. |