A note on multivariate Gauss-Hermite quadrature

Peter Jäckel*

16th May 2005

1 Introduction

Gaussian quadratures are an ingenious way to approximate the integral of an unknown function $f(x)$ over a specified domain $\mathcal{D}$ with a known weighting kernel $\psi(x)$. If the function $f(x)$ is well approximated by a polynomial of order $2m-1$, then a quadrature with n nodes suffices for a good estimate of the integral, i.e.

$$\int_{\mathcal{D}} f(x)\psi(x)\,\mathrm{d}x \;\approx\; \sum_{i=1}^{m} w_i f(x_i)\;. \tag{1}$$

The nodes x_i and weights w_i are uniquely determined by the choice of the domain $\mathcal{D}$ and the weighting kernel $\psi(x)$. In fact, one may go as far as to say that the choice of the domain and the kernel *defines* a quadrature. In particular, the location of the nodes x_i are given by the roots of the polynomial of order m in the sequence of orthonormal polynomials $\{\pi_j\}$ generated by the metric $\langle\pi_j|\pi_k\rangle := \int_{\mathcal{D}} \pi_j(x)\pi_k(x)\psi(x)\,\mathrm{d}x = \delta_{jk}$, and the weights w_i can be computed from a linear system once the roots are known. The mathematics of quadrature methods is well understood and described in most textbooks on numerical analysis [PTVF92].

In the case of the integration domain to be the entire real axis, and the integration kernel given by the density of a standard normal distribution, the associate quadrature scheme is known under the name *Gauss-Hermite* since the involved orthogonal polynomials turn out to be Hermite polynomials. Gauss-Hermite quadrature is of fundamental importance in many areas of applied mathematics that uses statistical representations, e.g. financial mathematics and actuarial sciences. Reliable routines for the calculation of the roots and weights are readily available [PTVF92] and most mathematical software packages provide means for one-dimensional Gauss-Hermite quadrature calculations.

*Global head of credit, hybrid, commodity, and inflation derivative analytics, ABN AMRO, 250 Bishopsgate, London EC2M 4AA

2 Multivariate Gaussian quadrature

In contrast to one-dimensional Gauss-Hermite quadrature, already at the level of two dimensions little is published on Gaussian quadrature with a multivariate normal metric

$$\varphi(\boldsymbol{r}, C) \;=\; \frac{\mathrm{e}^{-\frac{1}{2}\boldsymbol{r}^\top\cdot C^{-1}\cdot\boldsymbol{r}}}{(2\pi)^{d/2}\sqrt{|C|}} \tag{2}$$

with the correlation matrix C, apart from the obvious recommendation to decompose it into a sequence of nested one-dimensional quadratures.

This is where we start having a large amount of ambiguity: the set of nested one-dimensional integrations effectively amounts to a multivariate quadrature within which the integrand is evaluated at a discrete set of points and the computed function values are summed with weightings given by products of one-dimensioal weights. Since the alignment of the one-dimensional integrations is totally arbitrary, we effectively have a multitude of integration schemes at our disposal. To explain this in simple terms, assume that we are to compute a two-dimensional Gaussian integral $\iint f(x,y)\varphi(x)\varphi(y)\,\mathrm{d}x\,\mathrm{d}y$ with zero correlation between the two standard normal variates x and y. We can construct a two-dimensional Gaussian quadrature scheme with $m\times m$ function evaluations from a one-dimensional Gauss-Hermite quadrature with roots $\{z_k\}$ and weights $\{w_k\}$ for $k = 1..m$

by the summation

$$\sum_{i=1,j=1}^{m,m} w_i w_j f(z_i, z_j) \ . \tag{3}$$

This scheme essentially samples the function f on a set of points in a rectangular configuration as shown in figure 1 for $m = 15$ where contourlines of the

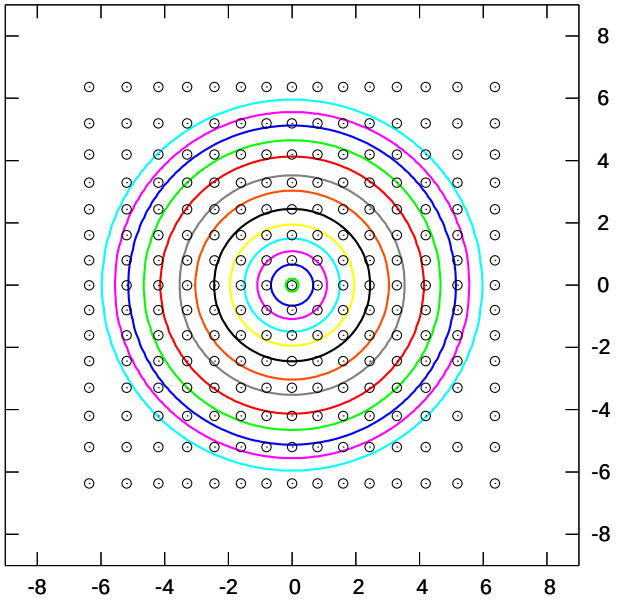

Figure 1: Sampling from a simple combination of one-dimensional Gauss-Hermite quadratures results in a rectangular set of function evaluation points.

distribution have been added for comparison. Since we have zero correlation between x and y, we could have chosen any arbitrary planar rotation of the the set $\{(z_i, z_j)\} \to \{(z'_i, z'_j)\}$ with

$$\boldsymbol{z}' \ = \ R \cdot \boldsymbol{z} \tag{4}$$

with R being a rotation operator, and we will come back to this observation later.

Another aspect of the weighted sampling scheme depicted in figure 1 is that the sheer appearance of the sampling set has one obvious unpleasant feature: on the diagonal, sampling points go much further in terms of distance to the origin than along the axes. Since the function values in those corner regions end up being weighted with extremely low weights, they contribute practically nothing at all to the total integral value. As a consequence, the quadrature scheme effectively wastes precious calculation time in the corners. This can be remedied to some extent, however, by the aid of the age-old straightforward trick known as *pruning*. We may simply set a threshold, say

$$\theta_m \ := \ \frac{w_1 \cdot w_{\left[\frac{m+1}{2}\right]}}{m} \ , \tag{5}$$

and drop all points that would be weighted with a net weight below θ_m:

$$\sum_{i=1,j=1}^{m,m} \mathbf{1}_{\{w_i w_j > \theta_m\}} \cdot w_i w_j f(z_i, z_j) \ . \tag{6}$$

The effect of pruning can be seen in figure 2.

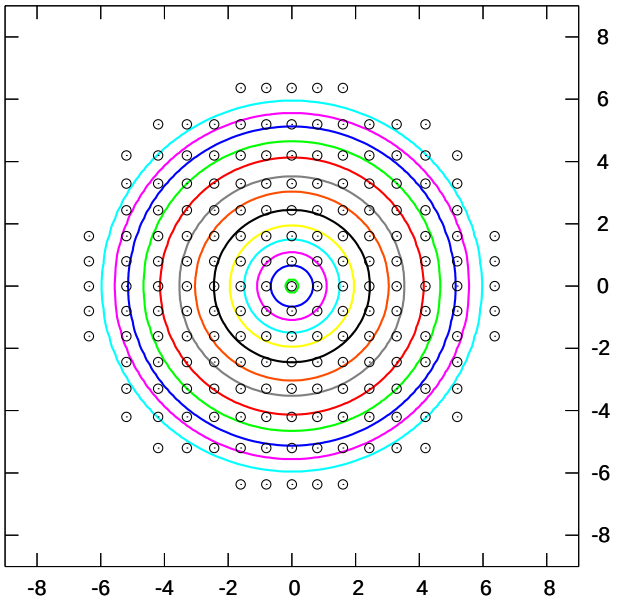

Figure 2: Pruning.

So far, we ignored correlation. In practice, we usually have non-zero correlation requirements. It may appear obvious to the reader how to account for correlation, though, here again, we have ambiguity as to how to proceed. An immediately intuitive method to incorporate correlation of magnitude ρ would be to use

$$\sum_{i=1,j=1}^{m,m} w_i w_j f(z_i, \rho z_i + \rho' z_j) \tag{7}$$

with $\rho' := \sqrt{1-\rho^2}$. What we are implicitly doing with this scheme is that we are transforming from uncorrelated to correlated variates by the aid of a Cholesky decomposition:

$$\boldsymbol{r} \ = \ L \cdot \boldsymbol{z} \tag{8}$$

The lower triangular matrix L satisfies $L \cdot L^\top = C$. A surprising side-effect of this algorithm is that we are

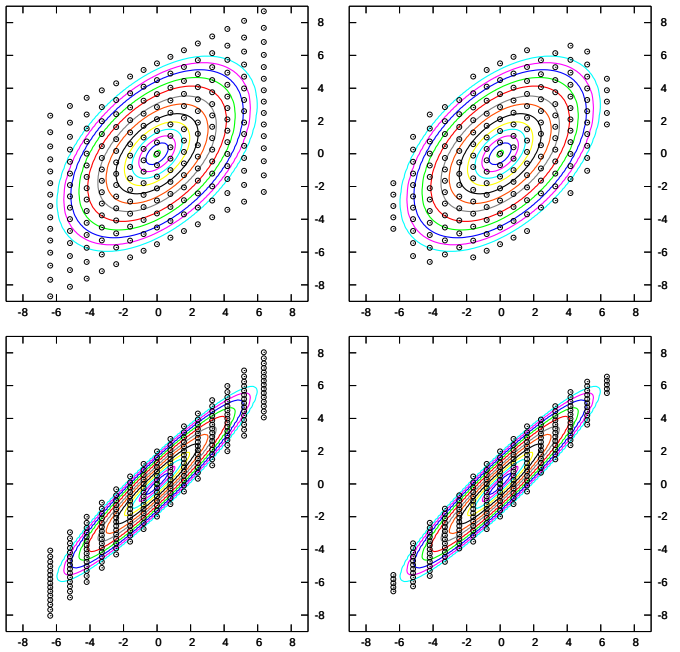

Figure 3: Correlation incorporation by the aid of Cholesky decomposition. Top left: $\rho = 0.5$. Top right: $\rho = 0.5$ (pruned). Bottom left: $\rho = 0.95$. Bottom right: $\rho = 0.95$ (pruned).

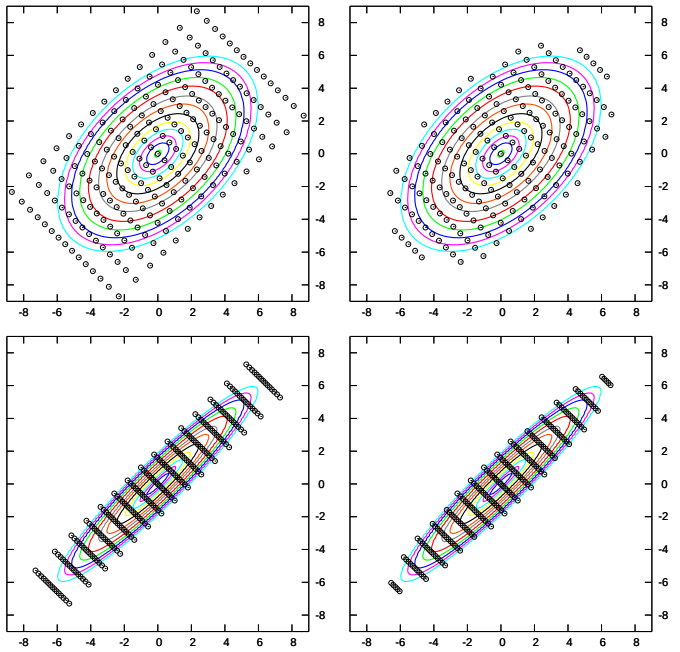

Figure 4: Correlation incorporation by the aid of spectral decomposition. Top left: $\rho = 0.5$. Top right: $\rho = 0.5$ (pruned). Bottom left: $\rho = 0.95$. Bottom right: $\rho = 0.95$ (pruned).

not sampling symmetrically with respect to the principal axes which is in practice usually, though arguably, a desirable feature to have for a multivariate quadrature scheme[1]. The effect of Cholesky decomposition on the sampling set is shown in figure 3.

An alternative to Cholesky decomposition, and one that is definitely to be preferred when Monte Carlo methods with low-discrepancy number generators are used, is to decompose spectrally [Jäc02]: find the eigensystem S and diagonal eigenvalue matrix Λ of C such that

$$C \;=\; S \cdot \Lambda \cdot S^{\top} \;, \tag{9}$$

set

$$A \;:=\; S \cdot \sqrt{\Lambda} \;, \tag{10}$$

and sample from the correlated set of sampling points

$$\boldsymbol{r} \;=\; A \cdot \boldsymbol{z} \;. \tag{11}$$

This scheme takes care of alignment with principal axes and the resulting sampling set is shown in figure 4.

We now return to the issue of arbitrary rotation. We can see in figure 4 that the automatic alignment with principal axes incurred by spectral incorporation of correlation effectively rotates the original square of sampling points showin in figure 1 such that the axes of the square become aligned with the principal axes. Out of all possible sections through the initial square, however, it is probably a different one that we would have preferred to become aligned with the principal axes, namely the one that already enhances sampling along the respective direction by virtue of the geometry of a square: the direction along the diagonals of the square. In two dimensions, this can be achieved by rotating[2] the set of uncorrelated standard normal variates on a square grid by 45° prior to the principal axes transformation. In higher dimensions, this can be done analogously: transform the initial set of d-dimensional rectangular grid coordinates by a sequence of $d-1$ planar rotations of 45° degrees (the net rotation of which may be denoted as R), and then transform to principal axes by setting

$$\boldsymbol{r} \;=\; A \cdot R \cdot \boldsymbol{z} \;. \tag{12}$$

In two dimensions, this scheme can be reduced to the formula

$$\sum_{i=1,j=1}^{m,m} \mathbf{1}_{\{w_i w_j > \theta_m\}} \cdot w_i w_j f(az_i + bz_j, bz_i + az_j) \tag{13}$$

with

[1] This is a purely empirical observation.

[2] It doesn't matter whether clockwise or counter-clockwise: both ways lead to the desired effect.

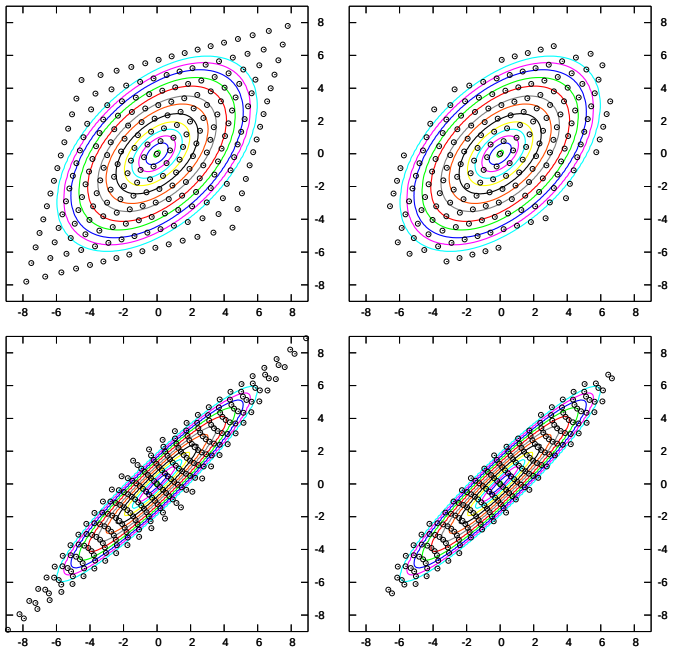

Figure 5: Principal axes rotation with spectral decomposition. Top left: $\rho = 0.5$. Top right: $\rho = 0.5$ (pruned). Bottom left: $\rho = 0.95$. Bottom right: $\rho = 0.95$ (pruned).

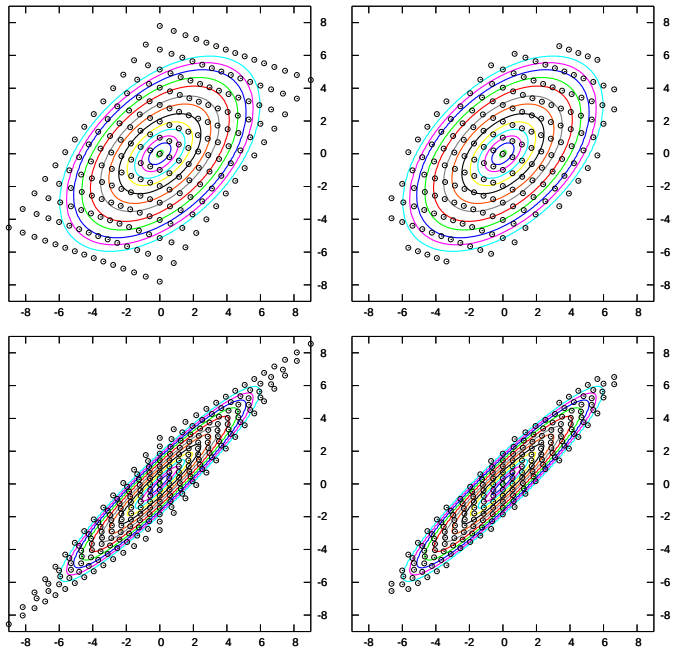

Figure 6: Principal axes rotation with Cholesky decomposition. Top left: $\rho = 0.5$. Top right: $\rho = 0.5$ (pruned). Bottom left: $\rho = 0.95$. Bottom right: $\rho = 0.95$ (pruned).

$$a := \frac{\sqrt{1+\rho} + \sqrt{1-\rho}}{2} \tag{14}$$

$$b := \frac{\sqrt{1+\rho} - \sqrt{1-\rho}}{2} \,. \tag{15}$$

The effect of this scheme is shown in figure 5.

One could, of course, also apply the upfront rotation to the Cholesky decomposition scheme, and this results in the sampling set shown in figure 6

As we can see, the rotation-Cholesky scheme samples over a set of points that appear to be be just slightly skewed away from the principal axes. In contrast, the rotation-spectral scheme preserves symmetry with respect to the principal axes. This may be one of the reasons why I found empirically the two-dimensional Gaussian quadrature scheme given by equation (13) to be the most robust in practical applications among those presented above.

3 Polar coordinates

In the previous section, we discussed the various possible choices how a multi-dimensional zero correlation Gauss-Hermite quadrature scheme can be adapted to account for correlation. The underlying scheme for zero correlation was given by a sequence of nested one-dimensional standard Gauss-Hermite quadratures which resulted effectively in a rectangular sampling grid. For two-dimensional problems, an alternative to this approach can be found by transforming the original integration problem to polar coordinates:

$$\iint f(x,y) \frac{\mathrm{e}^{-\frac{1}{2}\left(x^2+y^2\right)}}{2\pi} \,\mathrm{d}x\,\mathrm{d}y \;= \tag{16}$$

$$\int_0^1 \int_0^\infty f\left(r\cos\left[2\pi u\right], r\sin\left[2\pi u\right]\right) \,\mathrm{e}^{-\frac{1}{2}r^2} r \,\mathrm{d}r\,\mathrm{d}u$$

At this point, it is tempting to continue with

$$\int_0^1 \int_0^\infty f\left(r\cos\left[2\pi u\right], r\sin\left[2\pi u\right]\right) \,\mathrm{e}^{-\frac{1}{2}r^2} r \,\mathrm{d}r\,\mathrm{d}u \;= \tag{17}$$

$$\int_0^1 \int_0^\infty f\left(\sqrt{v}\cos\left[2\pi u\right], \sqrt{v}\sin\left[2\pi u\right]\right) \,\mathrm{e}^{-v} \,\mathrm{d}v\,\mathrm{d}u$$

and use Gauss-Laguerre quadrature for the inner integral. Hold your horses right there! This will not work since, chances are, $f(\cdot,\cdot)$ is a reasonably smooth function near the origin. This is because all Gaussian quadrature scheme are designed to work well for functions that are well or reasonably well approximated by polynomials. The square root function introduced

in transformation (17), however, has no polynomial representation at all! Thus, no matter how high we may go in the Gauss-Laguerre quadrature order, we will never be able to capture the fact that the square root function is not differentiable at the origin, and as a consequence, a two-dimensional Gauss-Hermite quadrature based on transformation (17) is very inefficient and unreliable.

Still, the idea of the polar coordinate transformation (16) may not be all that bad after all. The main crux is that the roots and weights for a quadrature on the domain $r \in (0, \infty)$ with metric $\mathrm{e}^{-\frac{1}{2}r^2} r$ are not readily available. However, if we manage to compute the roots $\{r_j\}$ and weights $\{\omega_j\}$ for $j = 1..m$ for the m-th order quadrature

$$\int_0^\infty f(r)\, \mathrm{e}^{-\frac{1}{2}r^2} r \, \mathrm{d}r \;\approx\; \sum_{j=1}^{m} \omega_j f(r_j) \;, \qquad (18)$$

then the polar coordinate approach to two-dimensional Gauss-Hermite quadrature (16) can be used still! The required mathematics for the calculation of the roots $\{r_j\}$ and weights $\{\omega_j\}$ for any one-dimensional quadrature defined by its integration domain and weighting kernel is described excellently in [PTVF92]. For the particular case (18), the roots and weights have been computed and are tabulated up to order 40 in [Jäc05].

The last point to be addressed is the outer integral on the right hand side of equation (16). A naïve attempt would be to tackle this one by the aid of Gauss-Legendre quadrature. However, since in practice the function $f\left(r\cos\left[2\pi u\right], r\sin\left[2\pi u\right]\right)$ is almost surely not of polynomial form in u for constant r, this is also unlikely to lead to satisfactory results. Without further knowledge, and to avoid the risk of introducing a lateral bias, one may need to resort to simply equally weighted sampling in u, i.e. $u_j := {}^{(j-1)}/_{m}$ with equal weights ${}^{1}/_{m}$. An example for this is shown in figure 7.

It remains to be mentioned that, since three-dimensional polar coordinates are also well tried and tested in mathematics and numerical analysis, it is conceivable that the approach shown for two-dimensionsal polar coordinates can be extended to three dimensions.

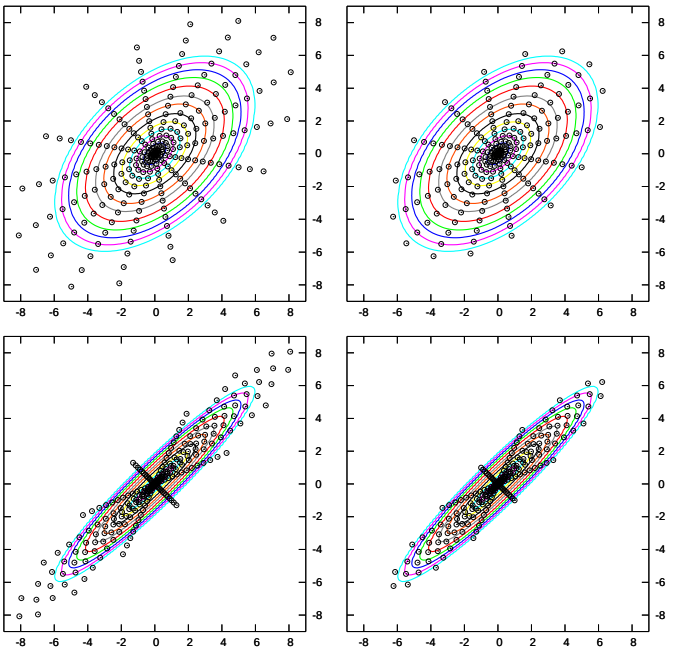

Figure 7: 16×16 two-dimensional Gauss-Hermite quadrature roots set based on the polar coordinate transformation (16). Top left: $\rho = 0.5$. Top right: $\rho = 0.5$ (pruned). Bottom left: $\rho = 0.95$. Bottom right: $\rho = 0.95$ (pruned).

References

[Jäc02] P. Jäckel. *Monte Carlo methods in finance*. John Wiley and Sons, February 2002.

[Jäc05] P. Jäckel. Roots and weights for radial Gauss-Hermite quadrature. www.jaeckel.org/RootsAndWeightsForRadialGaussHermiteQuadrcpp, May 2005.

[PTVF92] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1992. www.library.cornell.edu/nr/cbookcpdf.html.