Data Compression by Geometric Quantization

Nkem-Amin Khumbah and Edward J. Wegman
Center for Computational Statistics
George Mason University

Abstract

In this paper, we propose a nonparametric method for data quantization so as to reduce massive data sets to more manageable sizes. We investigate the probabilistic foundation and demonstrate statistical results for the quantization process. We discuss optimal geometric quantization procedures and discuss the computational and storage complexity of these procedures.

1 Introduction

The conjunction of large data set size, high dimensional data sets, and high-order algorithmic complexity makes exploratory data analysis and data mining infeasible for a number of data sets. Wegman (1995) argues that there are two significant thresholds when exploring data sets. One, he argues occurs around 10^6 to 10^7 bytes when issues of computational feasibility, data transfer, and human visual resolution become major issues. Of course, since 1995, computer chip speed has improved by roughly two orders of magnitude, and network communication speed and data transfer speed from hard drives has improved by at least one order of magnitude, but regrettably the human visual system has probably even degraded somewhat as populations age. Even so, computer screen resolution has actually only improved marginally. The rough figure of 10^6 to 10^7 bytes is still a reasonable threshold for interactive, visual-exploratory analysis including visual data mining. The second threshold Wegman (1995) argues occurs when data is sufficiently massive that it can no longer be stored on hard drives and must be stored in magnetic tape silos. In 1995, this threshold occurred at about one terabyte. This threshold too has moved upwards since 1995 and it is now feasible for a PC computer to have a terabyte of on-line storage installed in a single machine. However, while this threshold described in Wegman (1995) has moved up by perhaps two orders of magnitude, it is clear that data acquisition capability has also moved by similar orders of magnitude, so that the overall two-threshold premise is basically valid for the foreseeable future with only slight adjustments in the exact placement of the thresholds.

It is therefore of interest to be able to compress data in such a way that inference made about the data remain essentially unchanged, but that the data are amenable to issues of computational speed, transfer speed, storage capability, and capability for visualization. In short we would like to have the ability to compress data to around 10^6 to 10^7 observations without losing substantial inference capabilities. This is essentially a process of binning data. The phrase *quantization* is used in electrical engineering literature and carries much more of a sense of what we mean than the word binning typically used within the statistics community. In essence we are thinking of binning at a very fine scale consistent with sample size, say a million bins for 10^{12} observations. For comparison, digital images are usually quantized at something on the order of 2^{24} color levels, which is well beyond the capability of the human visual system to distinguish distinct colors. Similarly, CD quality audio is typically quantized at 2^{16} bits, which for all intents is beyond the capability of the human auditory system to distinguish.

The electrical engineering literature has discussed vector quantization, which is a process of forming clusters and then associating with each cluster a representor of that cluster. We note that Braverman (2001, 2002) has made very effective use of vector quantization for certain classes of satellite data for statistical purposes. However, vector quantization has the drawback that most clustering algorithms are $O(n^2)$ complexity so that a data set with 10^{12} observations would require $O(10^{24})$ computations, which even with teraflop computers is not feasible. We are proposing a nonparametric, geometry-based quantization with $O(n)$ computational complexity. In what follows, we adopt the convention that n refers to data set size, d refers to dimension of the data, and k will refer to the number of bins or tiles into which the data is quantized. We begin by laying out the probabilistic framework.

2 Quantization and Probability Spaces

The basic idea of quantization is to divide the underlying space into a finite number of representors. This can be done either in a Euclidean d-space or in the underlying probability space depending which is more convenient for developing theory or practice. In either case we want the theory to be consistent. We describe here the basic structure. First we consider an abstract probability space described by $(\Omega, \mathfrak{A}, P)$, where Ω is an abstract set, $\mathfrak{A}$ is a σ-field, and P is a probability measure. Consider a mapping $Q_k : \Omega \longrightarrow \Omega$, which is a many-to-one mapping and maps a measurable subset $A \in \Omega$ into a single point, $Q(A)$. Q_k will be our quantizer map. Later, we will suppress the subscript k. We will eventually choose only a finite number, k, of $A \in \mathfrak{A}$ according to a geometric criterion. Let $\Omega_k = \{b_j\}$ be the finite collection of images of Q_k. The family $\mathfrak{A}_k$ is given by

$$\mathfrak{A}_k = \{Q_k(A) : A \in \mathfrak{A}\}$$

so that for each $B \in \mathfrak{A}_k$, $Q_k^{-1}(B) \subset \Omega$ is measurable. For each $B \subset \mathfrak{A}_k$,

$B = \underset{j \in I_B}{\cup} \{b_j\}$, with b_j the singleton images of some $A_j \in \mathfrak{A}$ and I_B and index set for B. Thus

$$Q_k^{-1}(B) = Q_k^{-1}(\underset{j \in I_B}{\cup} \{b_j\}) = \underset{j \in I_B}{\cup} A_j \in \mathfrak{A}.$$

Thus $\mathfrak{A}_k$ contains countable (actually finite) unions. Similarly, for each $b_j \in \Omega_k$, we define $p_j = P(Q_k^{-1}(b_j))$ and let P_k be the finite probability measure associated with measurable sets in $\mathfrak{A}_k$. The points b_j are the representors of the set A_j in the original probability space. For $B \in \mathfrak{A}_k$, $B = \underset{j \in I_B}{\cup} \{b_j\}$, we define

$$P_k(B) = P(Q_k(B)) = \sum_{j \in I_B} P(Q_k(b_j)) = \sum_{j \in I_B} p_j.$$

Lemma 1: $\mathfrak{A}_k$ is a field of sets.

Proof: Suppose first that $\{A_j\}$ is a partition of Ω where each $A_j \in \mathfrak{A}$. Then $\Omega_k = Q_k(\Omega)$. Then $\Omega_k \in \mathfrak{A}_k$ since by definition $\Omega \in \mathfrak{A}$. Next suppose $B_1, B_2 \in \mathfrak{A}_k$. Then $\exists\ A_1, A_2 \in \mathfrak{A}$ such that $B_1 = Q_k(A_1)$ and $B_2 = Q_k(A_2)$ and $B_1 \cup B_2 = Q_k(A_1) \cup Q_k(A_2)$. Let $b \in B_1 \cup B_2$. Then either $b \in B_1$ or $b \in B_2$ or both. If $b \in B_1 \Rightarrow b \in Q_k(A_1)\, \exists\, \omega \in A_1$ such that $b = Q_k(\omega) \Rightarrow \omega \in A_1 \cup A_2 \Rightarrow b \in Q_k(A_1 \cup A_2)$. Similarly if $b \in B_2$. Thus

$$Q_k(A_1) \cup Q_k(A_2) \subset Q_k(A_1 \cup A_2).$$

Conversely, suppose $b \in Q_k(A_1 \cup A_2)$. Then $\exists\, \omega \in A_1 \cup A_2$ such that $b = Q_k(\omega)$. If $\omega \in A_1$, then $b \in Q_k(A_1)$. If $\omega \in Q_k(A_2)$, then $b \in Q_k(A_2)$. In either case, $b \in Q_k(A_1) \cup Q_k(A_2)$. Thus

$$B_1 \cup B_2 = Q_k(A_1) \cup Q_k(A_2) = Q_k(A_1 \cup A_2) \subset \mathfrak{A}_k.$$

Thus, by induction, $\mathfrak{A}_k$ is closed under countable (finite) unions. Now consider B_1 and B_2 and suppose $b \in B_1 - B_2$. Then there exists $A_1 \in \mathfrak{A}$ such that $B_1 = Q_k(A_1)$ and $A_2 \in \mathfrak{A}$ such that $B_2 = Q_k(A_2)$. Since $b \in B_1$, for some $\omega \in A_1$, $b = Q_k(\omega)$. Suppose for that ω, $\omega \in A_2$. Then $b \in Q_k(A_2) \Rightarrow$by above $b \in Q_k(A_1 \cup A_2) = B_1 \cup B_2$. This is contradiction $b \in B_1 - B_2$. Thus $Q_k(A_1) - Q_k(A_2) = Q_k(A_1 - A_2)$. This $\mathfrak{A}_k$ is closed under complementation. Finally, $\phi = \Omega_k - \Omega_k = Q(\Omega - \Omega) = Q_k(\phi)$.

It is straightforward to show that P_k is a finite probability measure. Now suppose Y is a random vector mapping $(\Omega, \mathfrak{A}, P)$ into $\mathcal{R}^d$, d-dimensional Euclidean space. Let F be the induced probability distribution. We will have in mind partitioning $\mathcal{R}^d$ into a finite number of sets, which we will call S_j, $j = 1, \ldots, k$. The S_j will be a tessellation of $\mathcal{R}^d$ and for each S_j we will choose a representor. In particular, $q_k(\boldsymbol{x}) = y_j$ if $\boldsymbol{x} \in S_j$. Here we denote the vector $\boldsymbol{x}$ by a bold $\boldsymbol{x}$. Let $W = \{y_1, \ldots, y_k\}$ the quantized space. Equivalently, $q_k^{-1}(y_j) = \{\boldsymbol{x} \in \mathcal{R}^d : q_k(\boldsymbol{x}) = y_j\}$, $j = 1, \ldots, k$. As above there is an induced probability $P(y_j) = \int_{S_j} dF(\boldsymbol{x})$. The geometric tessellation on $\mathcal{R}^d$ induces a partition on Ω as follows: $A_j = Y^{-1} \circ q_k^{-1}(y_j)$ which in turn creates the finite probability space $(\Omega_k, \mathfrak{A}_k, P_k)$ described above. This in turn induces a finite random variable $Y_k : \Omega_k \longrightarrow W$. In short the following diagram commutes.

$$\begin{array}{ccc} (\Omega, \mathfrak{A}, P) & \xrightarrow{Q_k} & (\Omega_K, \mathfrak{A}_k, P_k) \\ \downarrow Y & & \downarrow Y_k \\ \mathcal{R}^d & \xrightarrow{q_k} & W. \end{array}$$

The choice of representor y_j could reasonably be any element of S_j. However, an excellent choice is $y_j = \int_{S_j} \boldsymbol{x} \, dF(\boldsymbol{x}) = E(Y|Q_k = y_j)$. We are now able to suppress the subscript k and rewrite this equation as $E(Y|Q) = Q$. That is the expected value of a quantized random vector is just the quantized vector itself. Another way of thinking is that within a tile of our tessellation, we choose the representor to be the mean value of the conditional distribution in the tile. This feature is called self-consistency, which has some important implications for quantization. For more details on self consistency, see Tarpey and Flury (1996). This idea was exploited in Braverman (2000, 2002) for vector quantization.

3 Self Consistency and its Implications

Clearly, we have as an easy result $E(Y) = E_Q E(Y|Q) = E(Q)$. Thus if we have a sample $Y_1, Y_2, \ldots, Y_n$ and $\widehat{\theta}(Y_1, \ldots, Y_n)$ is a linear unbiased estimator of a parameter θ, the same estimator based on the quantized version, say $E[\widehat{\theta}|Q]$ will also be a linear unbiased estimator.

Theorem 1:

1. $E(Y) = E(Q)$
2. If $\widehat{\theta}$ is a linear unbiased estimator of θ, then so is $E(\widehat{\theta}|Q)$
3. If h is a convex function, then $E(h(Q)) \leq E(h(Y))$. In particular $E(Q^2) \leq E(Y^2)$ so that $var(Q) \leq var(Y)$
4. $E[Q(Q - Y)] = 0$
5. $cov(Y - Q) = cov(Y) - cov(Q)$
6. $E(Y - P)^2 \geq E(Y - Q)^2$ where P is any other quantizer.

Proof:

We have just shown 1 and 2 above. To see 3, recall from Jensen's Inequality $E(h(Y)) \geq hE(Y)$. Thus $E[h(Q)] = E[h(E(Y|Q)] \leq E[E(h(Y)|Q] = E(h(Y))$. Because $h(y) = y^2$ is a convex function, it follows that $E(Q^2) \leq E(Y^2)$. Moreover, because $E(Q) = E(Y)$, it follows that $var(Q) \leq var(Y)$.

To see 4, consider

$$E(Q(Q-Y)) = E_Q E(Q(Q-Y)|Q)$$
$$= E_Q E(Q^2|Q) - E_Q E(QY|Q)$$
$$= E_Q(Q^2) - E_Q[QE(Y|Q)]$$
$$= E_Q(Q^2) - E_Q(Q^2)$$
$$= 0.$$

To see 5,

$$cov(Y-Q) = E[(Y-Q)'(Y-Q)]$$
$$= E(Y-Q)'Y - E[(Y-Q)'Q]$$
$$= E(Y'Y) - E(Y'Q)$$
$$= E(Y'Y) - E_Q E(Y'Q|Q)$$
$$= E(Y'Y) - E_Q[E(Y'|Q)Q]$$
$$= E(Y'Y) - E(Q^2).$$

Because $[E(Y)]^2 = [E(Q)]^2$, it follows that $cov(Y-Q) = cov(Y) - cov(Q)$.

Finally to see 6, consider

$$E[Y-P]^2 = E[Y-Q+Q-P]^2$$
$$= E[Y-Q]^2 + E[Q-P]^2 + 2E(Y-Q)(Q-P)$$
$$= E[Y-Q]^2 + E[Q-P]^2 + 2\{E[(Y-Q)Q\} - E[(Y-Q)P]\}$$
$$= E[Y-Q]^2 + E[Q-P]^2 - 2E[(Y-Q)P]$$
$$= E[Y-Q]^2 + E[Q-P]^2 - 2E_P(P[E(Y-Q)]|P)$$
$$= E[Y-Q]^2 + E[Q-P]^2$$
$$\geq E[Y-Q]^2.$$

Result 6 indicates that the optimal strategy for quantizing in terms of reducing mean square error is to use the conditional expectation of the random vector for a given tile as the representor of the tile.

One slightly troubling result of this theorem is that quantization is variance reducing. Ideally, we would prefer that the variance structure remained identical between the quantized data and the original data. Experimental results shown in Figure 1 suggest that indeed for data set sizes we are considering, there need be little concern.

4 Some Notes on Convergence

An important fact to recognize is that quantization turns continuous spaces on Ω into a discrete space Ω_k. This makes convergence substantially easier to deal with. Consider a sequence of random variables $Y_i : \Omega_k \longrightarrow W$. Then we have the following result.

Theorem 2: Suppose Y_i converges to Y in any of the following modes of convergence. Then it converges in all the other modes as well. That is the following modes of convergence are equivalent.

(a) Convergence in probability
(b) Convergence almost surely
(c) Convergence in p-norm
(d) Almost uniform convergence
(e) Uniform convergence

Proof:

Standard results show that (e) $\Rightarrow$ (d) $\Rightarrow$ (b) $\Rightarrow$ (a).

Show that (a) $\Rightarrow$ (b). Suppose $Y_i \longrightarrow Y$ as $i \longrightarrow \infty$ in probability. Then for every $\epsilon > 0$,

$$\underset{i \longrightarrow \infty}{lim} P_k(|Y_i - Y| \leq \epsilon) = 1.$$

Let $\Omega_k = \{\omega_j\}_{j=1}^k$. We may assume the ω_j are all distinct. Let us consider the corresponding $y_j \in W$. Because y_j is the representor of tile S_j, there is an induced metric on W. Suppose $d(y_j, y_m)$ is that metric. We may assume that the y_j like the ω_j are distinct. Choose $\epsilon > 0$ such that $\epsilon < min\, d(y_j, y_m)$, $j, m = 1, \ldots, k$. Then

$$\underset{i \longrightarrow \infty}{lim} P_k[|Y_i - Y| \leq min\, d(y_j, y_m)] = 1.$$

This implies there exists an N_ϵ such that for every $i > N_\epsilon$, $P_k(Y_i = Y) = 1$. Thus (a)$\Rightarrow$(b).

Show that (b)$\Rightarrow$(e). Suppose $Y_i \longrightarrow Y$ pointwise almost surely. Then given $\epsilon > 0$ and for every $j = 1, \ldots, k$, $\exists\, N_j$ such that whenever $i > N_j$, $|Y_i(\omega_j) - Y(\omega_j)| < \epsilon$. Take $N = \underset{1 \leq j \leq k}{max} N_j$. Then whenever $i \geq N$

$$|Y_i(\omega_j) - Y(\omega_j)| \leq \underset{1 \leq j \leq k}{max} |Y_i(\omega_j) - Y(\omega_j)| \leq \epsilon.$$

Thus (b)$\Rightarrow$(e)

Show that (c)$\Rightarrow$(a). Suppose $Y_i \longrightarrow Y$ in p-norm. Thus $\underset{i \longrightarrow \infty}{lim} \sum_{j=1}^k |Y_i(\omega_j) - Y(\omega_j)|^p p_j = 0$. Then for any $\epsilon > 0$, let $A_i = \{\omega_j : |Y_i(\omega_j) - Y(\omega_j)|^p > \epsilon\}$. Then

$$\textstyle\sum_{j \in A_i} |Y_i(\omega_j) - Y(\omega_j)|^p p_j \geq \epsilon P_k(A_i).$$

Hence

$$\textstyle\sum_{j=1}^k |Y_i(\omega_j) - Y(\omega_j)|^p p_j =$$

$= \sum_{j \in A_i} |Y_i(\omega_j) - Y(\omega_j)|^p p_j + \sum_{j \in \Omega_k - A_i} |Y_i(\omega_j) - Y(\omega_j)|^p p_j$

$\geq \epsilon P_k(A_i) + \sum_{j \in \Omega_k - A_i} |Y_i(\omega_j) - Y(\omega_j)|^p p_j$

$\geq \epsilon P_k(A_i)$.

Thus $0 \leq \underset{i \longrightarrow \infty}{lim} \epsilon P_k(A_i) \leq \underset{i \longrightarrow \infty}{lim} \sum_{j=1}^{k} |Y_i(\omega_j) - Y(\omega_j)|^p p_j = 0$. Thus we have (c)⇒(a).

Finally show (b)⇒(c). Suppose $Y_i \longrightarrow Y$ almost surely. Then given $\epsilon > 0$, $\exists N_\epsilon$ such that whenever $i > N_\epsilon$,

$|\mathrm{Y}_i(\omega_j) - Y(\omega_j)| < \epsilon$ for every $\omega_j \in \Omega_k$.

Then

$\sum_{j=1}^{k} |\mathrm{Y}_i(\omega_j) - Y(\omega_j)|^p p_j \leq |\epsilon|^p \sum_{j=1}^{k} p_j = |\epsilon|^p$.

Thus $Y_i \longrightarrow Y$ in p-norm.

5 Minimizing Distortion and Geometry

We will focus on distortion as the mean square error difference between the unquantized data and the quantized data. We have previously shown that by choosing $E(Y|Q) = Q$, we minimize $E(Y - Q)^2$, i.e. the distortion. That is, we choose the representor of a tile to be the average values within the tile. However, the shape of the tile impacts the distortion as well. Clearly if the tile is a sphere and the distribution is symmetric within the sphere, then the geometric centroid and the conditional expectation will coincide, and the ideal representor would simply be the geometric centroid. In general this is not the case. However it is shown by Gersho (1979) and Gish and Pierce (1968), that in $\mathcal{L}_2$-vector spaces, the geometric shape that minimizes distortion is a sphere. Unfortunately, the sphere does not tessellate any Euclidean space except $\mathcal{R}^1$, in which case it is a straight line segment. Ideally, geometric quantization has the following constraints:

1. We must tessellate the space with tiles that are a spherical as possible,

2. The tiles should preferably be congruent,

3. The tessellation must be space filling so as not to miss observations.

As a measure of sphericity, we can take the dimensionless second moment. Consider a regular polytope in d-dimensions. The (hyper)-volume and second moment can be expressed in terms of the volume and second moments of its faces (the part of the polytope that lies in one of its enclosing hyperplanes),

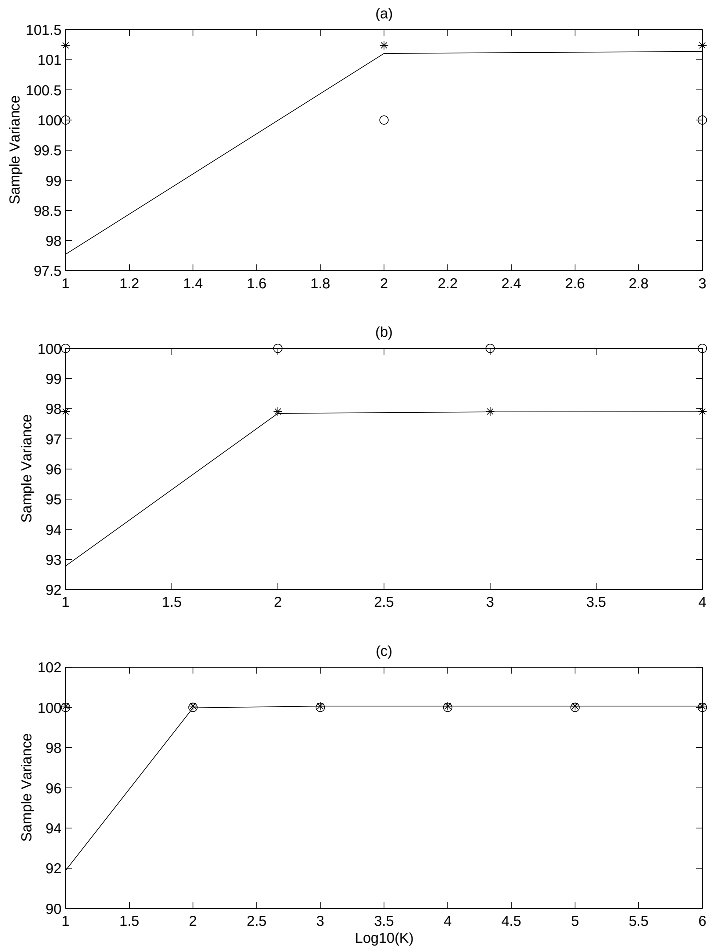

Figure 1: Comparison of true variance (given by circle), estimated variance (given by asterisks) and quantized estimated variance (given by solid line). The three panels starting at top have $n = 10^3$, $n = 10^4$, and $n = 10^6$

Table 1: Dimensionless Second Moments for Several Regular 3-D Polytopes. Those marked with an * tessellate 3-space.

Tetrahedron*	.1040042...
Cube*	.0833333...
Octahedron	.0825482...
Hexagonal Prism*	.0812227...
Rhombic Dodecahedron*	.0787451...
Truncated Octahedron*	.0785433...
Dodecahedron	.0781285...
Icosahedron	.0778185...
Sphere	.0769670...

then in terms of its $(d-2)$-dimensional faces and so on. See Sloane and Conway (1999) for details.

Suppose S is a d-dimensional polytope with N_1 congruent faces, $F_1^1, \ldots, F_1^{N_1}$; N_2 congruent faces, $F_2^1, \ldots, F_2^{N_2}$ and so on. Suppose that S contains a point, A, such that all the generalized pyramids AF_1^j, $j = 1, \ldots, N_1$ are congruent, and AF_2^j, $j = 1, \ldots, N_2$ are congruent, and so on. Let $a_i \in F_i$ be the foot of the perpendicular from A to F_i, $V_{d-1}(i)$ be the volume of the face F_i, $U_{d-1}(i)$ be the second moment of F_i about a_i and define $h_i = \|A - a_i\|$. Then the volume and second moment of S about A are given in Sloane and Conway (1999) by

$$vol(S) = \sum_i \frac{N_i h_i}{d} V_{d-1}(i)$$

and

$$U(S) = \sum_i \frac{N_i h_i}{d+2} \left[h_i^2 V_{d-1}(i) + U_{d-1}(i) \right].$$

As mentioned above, in 1-dimensional space the only polytope of interest is a straight line, which also happens to be a one-dimensional sphere. In two dimensions, the only regular polytopes that tessellate 2-space are equilateral triangles, squares, and hexagons. Of these, hexagons have the smallest dimensionless second moment, and from distortion considerations, this would be the tessellation of choice. In three dimensions, perhaps the most interesting, the regular polytopes that can tessellate three space are tetrahedrons, cubes, hexagonal prisms, rhombic dodecahedrons, and truncated octahedrons.

As can be see from Table 1, the truncated octahedron would be the polytope of choice for tessellating 3-space from the perspective of minimizing distortion. In four dimensions, 4-simplices, hypercubes and 24-cells tessellate 4-dimensional space, with the 24-cell having the minimal dimensionless second moment. Many of these polytopes are illustrated in Wegman (2001) with animations on a CD.

6 Computational and Storage Complexity of Geometric Quantization

Of course, the main point of the geometric quantization is to reduce the data set size effectively so that computationally complex algorithms may be used, so that data transfer in reasonable time frames is feasible, and so that visualization methods may be used. All of this depends on the quantizing algorithms to be computationally feasible. As mentioned above the engineering concept of vector quantization relies on clustering algorithms, which, generally, are $O(n^2)$. Fortunately, the geometric quantizing algorithms are $O(n)$.

Consider first the one-dimensional case. Suppose the data are located within the set $[a, b]$ and that we desire to quantize the data into k tiles. Then the algorithm

$$j = fixed[k * (x_i - a)/(b - a)]$$

calculates j the index of the tile into which x_i falls. The $fixed$ is an operation which changes a floating point number into a fixed point number. The term $(b-a)$ need only be calculated once. Thus there is one subtraction, one division, one multiplication and one $fixed$ operation for each observation. In total for a data set of size n there are $4n + 1 = O(n)$ operations required. Notice also that the storage requirement drops from n to $3k$. For each of the k quantized representors, we must store the location of the tile, the number of observations in the tile, and the value of the representor. Generally we want to choose $k \ll n$. Our operational target is $k = 10^6$.

In the two-dimensional case, if we tessellate with hexagons, each hexagon has three pairs of parallel sides. Thus the tile must be triply indexed and the above algorithm applied in each of three directions. Thus the computational complexity for hexagons is $12n + 3 = O(n)$ and the storage complexity remains $3k$. Of course, if we tessellate with squares or rectangles, there are only two pairs of parallel sides, so that the tiles need be only doubly indexed. Thus the computational complexity for squares or rectangles is $8n + 2 = O(n)$ with the storage requirement remaining at $3k$.

In the three-dimensional case, the most spherical polytope is the truncated octahedron. This polytope has 3 pairs of square sides and 4 pairs of hexagonal sides. Thus tiles when tessellated using truncated octahedrons must be indexed by 7 indices so that the one-dimensional algorithm must be applied 7 times. This yields a computational complexity of $28n + 7 = O(n)$ and storage requirements of $3k$. Of course if cubes are used to tessellate 3-space, they need only be triply indexed so the computational complexity is $12n + 3 = O(n)$.

Of course in general if d-dimensional hypercubes are used to tessellate d-space, there are d orthogonal directions so the general form for the computational complexity will be $(4n + 1)d = O(n)$ with storage requirements always remaining at $3k$. Unfortunately, while the computational complexity is always $O(n)$, the number of tiles, assuming a constant tile size per orthogonal direction, grows exponentially with dimension. Thus for high dimensions the number of tiles can easily be greater than the number of observations making quantization a useless exercise. Geometric quantization is probably not useful beyond about 5 dimensions unless the partition in each direction is severely limited. Indeed, within the computer science literature, the so-called datacube is essentially a hyper-rectangular tiling with each variable limited to a small number of categories, usually 3 or 4.

One extremely positive aspect of this geometric quantization is that the tiles are determined independent of data. Thus geometric quantization can be used for streaming data, i.e. as a data point comes in, it can be placed into a tile and the data for that tile recursively updated. The mean value for a tile can be recursively computed.

7 Discussion and Conclusions

In this paper we have introduced the idea of quantizing massive data sets in order to compress data sets to a feasible size from several points of view. This quantization is essentially a nonparametric procedure. We have laid out the probabilistic foundations and exploited the self-consistency property to show that most statistical properties for quantized data remain unchanged from what they would have been with unquantized data. Quantization of course reduces the underlying probability space to a finite space. This has great advantages from a convergence perspective. These results hold no matter the mode of quantization. However, in order to minimize distortion, i.e. mean square error, the optimal geometry involves tessellating space with the most spherical possible tessellating polytopes. We have identified these in dimensions 1 through 4 and shown that computational complexity in general is $O(n)$. We also commented that geometric quantization is probably not practical beyond 5 dimensions unless extremely coarse tiles are used.

Practically speaking, it is probably better to use hypercubes rather than the distortion optimal polytope because of reduced complexity and greater interpretability. Alternate geometric tessellations are possible. One procedure might be to select some data points at random and create a Delauney tessellation. This procedure was used for multivariate density estimation in Hearne and Wegman (1991, 1992, 1994). The rate of growth of the number of tiles with random tessellations has been shown experimentally to be slower than with regular tessellations. However, there are much more severe departures from spherical shapes. An alternate procedure being currently explored is to use density estimates to determine local modes and then use the local modes for Voronoi tessellations. Because kernel density estimates are also $O(n)$ computationally, this approach to forming clusters may lead to an alternate to distance-based clustering that would make vector quantization feasible for large data sets.

Acknowledgments

This paper is based on an invited talk prepared by Dr. Wegman for the International Conference on Current Advances and Trends in Nonparametric Statistics held July 15-19, 2002 in Crete, Greece. Due to an illness in the family, Dr. Wegman was unable to attend and gratefully acknowledges Dr. Don Faxon's assistance by presenting the paper on Dr. Wegman's behalf. We also gratefully acknowledge the kind invitation and the flexibility of the organizers in adapting to the substitute speaker. This paper is based in part on the Ph.D. dissertation of Nkem-Amin Khumbah (Khumbah, 2000) under the direction of Dr. Wegman. Dr. Wegman's work was funded by the Office of Naval Research under contract DAAD19-99-1-0314 administered by the Army Research Office, by the Air Force Office of Scientific Research under contract F49620-01-1-0274 and contract DAAD19-01-1-0464, the latter also administered by the Army Research Office and finally by the Defense Advanced Research Projects Agency through cooperative agreement 8105-48267 with the Johns Hopkins University.

References

Braverman, Amy (2000), "Compressing massive geophysical data using quantization," *Computing Science and Statistics (CD)*, 32, 28-36

Braverman, Amy (2002), "Compressing massive geophysical datasets using vector quantization," *Journal of Computational and Graphical Statistics*, 11(1), 44-62

Gersho, Allen (1979), "Asymptotically optimal block quantization," *IEEE Transactions on Information Theory*, IT-25(4), 373-380

Gish, H. and Pierce, John N. (1968), "Asymptotically efficient quantizing," *IEEE Transactions on Information Theory*, IT-12(5), 678-681, 683

Hearne, Leonard B. and Wegman, Edward J. (1991), "Adaptive probability density estimation in lower dimensions using random tessellations," *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, 241-245

Hearne, Leonard B. and Wegman, Edward J. (1992), "Maximum entropy density estimation using random tessellations," *Computing Science and Statistics*, 24, 483-487

Hearne, Leonard B. and Wegman, Edward J. (1994), "Fast multidimensional density estimation based on random-width bins," *Computing Science and Statistics*, 26, 150-155

Khumbah, Nkem-Amin N. (2000), *Mathematical Quantization for Massive Data*

Sets, Ph.D. dissertation, School of Information Technology and Engineering, George Mason University

Sloane, N. J. A. and Conway, J. H. (1999), *Sphere Packings, Lattices and Groups*, (Third Edition) Springer-Verlag: New York

Tarpey, T. and Flury, B. (1996), "Self-consistency: A fundamental concept in statistics," *Statistical Science*, 11, 229-243

Wegman, Edward J. (1995), "Huge data sets and the frontiers of computational feasibility," *Journal of Computational and Graphical Statistics*, 4(4), 281-295

Wegman, Edward J. (2001), "Data compression by quantization," *Computing Science and Statistics(CD)*, 33, \Interface2001\I2001HTML\EWegman.htm