Assume the scalar quantizer has point density λ_1.

The cell $S_{\underline{x}}$ of the product quantizer containing $\underline{x}$ is a rectangle:

$$S_{\underline{x}} = \frac{1}{2^R\lambda_1(x_1)} \times \frac{1}{2^R\lambda_1(x_2)} \times \frac{1}{2^R\lambda_1(x_3)} \times \cdots \frac{1}{2^R\lambda_1(x_k)}$$

with volume

$$|S_{\underline{x}}| = \frac{1}{2^{kR}} \times \frac{1}{\lambda_1(x_1)\lambda_1(x_2)\ldots\lambda_1(x_k)}$$

The "product" point density is

$$\lambda_{pr,k}(\underline{x}) = \frac{1}{2^{kR}|S_{\underline{x}}|} = \lambda_1(x_1)\,\lambda_1(x_2)\ldots\lambda_k(x_k) \qquad \text{(It's a product!)}$$

The "product" inertial profile is

$$m_{pr,k}(\underline{x}) = \frac{1}{12}\,\frac{\frac{1}{k}\sum_{i=1}^{k}\frac{1}{\lambda_1(x_i)^2}}{\left(\prod_{i=1}^{k}\frac{1}{\lambda_1(x_i)^2}\right)^{1/k}} = m_1^*\,\frac{\frac{1}{k}\sum_{i=1}^{k}\frac{1}{\lambda_1(x_i)^2}}{\left(\prod_{i=1}^{k}\frac{1}{\lambda_1(x_i)^2}\right)^{1/k}}$$

Consider "loss" of the scalar quantizer (i.e. the product quantizer) relative to the VQ:

$$L = \frac{D(Q_1)}{D(Q_k)} = \frac{D(Q_{pr,k})}{D(Q_k)}$$

Applying Bennett's integral to both terms. Let

$$B(k, m, \lambda, f) = \int \frac{m(\underline{x})}{\lambda^{2/k}(\underline{x})}\, f(\underline{x})\, d\underline{x} = \text{Bennett's integral}$$

Then

$$L = \frac{D(Q_1)}{D(Q_k)} = \frac{D(Q_{pr,k})}{D(Q_k)} \cong \frac{B(k, m_{pr,k}, \lambda_{pr,k}, f)}{B(k, m_k, \lambda_k, f)}$$

$$= \frac{B(k, m_{pr,k}, \lambda_{pr,k}, f)}{B(k, m_k, \lambda_{pr,k}, f)} \times \frac{B(k, m_k, \lambda_{pr,k}, f)}{B(k, m_k, \lambda_k, f)}$$

$= \text{cell shape loss} \times \text{point density loss}$

$= L_{ce} \times L_{pt}$

Now assume k is large and Q_k is optimal, so that

$$m_k(\underline{x}) \cong m_\infty^* = \frac{1}{2\pi e} \quad \text{and} \quad \lambda_k(\underline{x}) \cong \lambda_k^*(\underline{x}) = c\, f^{k/(k+2)}(\underline{x}) \cong f(\underline{x})$$

Then the loss of the scalar quantizer relative to the optimal high dim'l VQ is

$$L = \frac{m_1^*}{m_\infty^*} \times \frac{\int \frac{1}{k} \sum_{i=1}^{k} \frac{1}{\lambda_1^2(x_i)} f(\underline{x})\, d\underline{x}}{\int \frac{1}{(\prod_{i=1}^{k} \lambda_1^2(x_i))^{1/k}} f(\underline{x})\, d\underline{x}} \times \frac{B(k, m_k^*, \lambda_{pr,k}, f)}{B(k, m_k^*, \lambda_k^*, f)}$$

$$= \text{cubic loss} \times \text{oblongitis loss} \times \text{point density loss}$$

$$= L_{cu} \times L_{ob} \times L_{pt}$$

where we have factored the cell shape loss L_{ce} into the product of a "cubic loss" L_{cu} and an "oblongitis loss".

Z-51

Now assume source is IID.

Consider the choice of $\lambda_1(x)$ to minimize the loss, i.e. to minimize $L_{ob} \times L_{pt}$.

- Choosing $\lambda_1(x)$ to be a constant on the set where $f_1(x)$ is large makes $L_{ob} \cong 1$. However, L_{pt} becomes very large.
- On the other hand choosing $\lambda_1(x) = f_1(x)$ causes

$$\lambda_{pr,k}(\underline{x}) \cong \prod_{i=1}^{k} f_1(x_i) = f(\underline{x}) \cong c\, f_k^{k/(k+2)}(\underline{x}) = \lambda_k^*(\underline{x}) \qquad (k \text{ large} \Rightarrow k/(k+2) \cong 1)$$

so that $L_{pt} = 1$. In other words scalar quantization can produce the optimal point density! This fact is often overlooked, because for IID Gaussian case, the product quantizer looks like it has a "cubical" point density, when it actually has a spherical one. Unfortunately, however, for this λ_1, However, $L_{ob} = \infty$.

Z-52

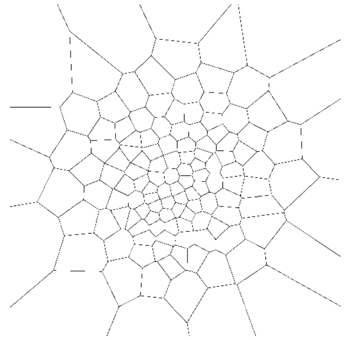

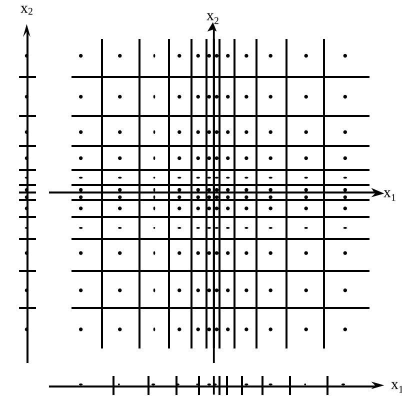

The point density that minimizes $L_{ob} \times L_{pt}$ is the compromise

$$\lambda_1(x) = c\, f_1(x)^{1/3}$$

This makes $\lambda_1(x)$ "flatter" (more uniform) in regions where $f_1(x)$ is large than the previous choice of $\lambda_1(x) = f(x)$. Therefore, the rectangular cells are more nearly cubical in the important region where $f_1(x)$ is large, so there is less oblongitis loss.

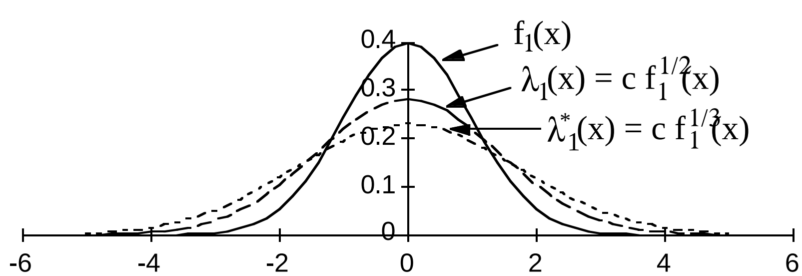

The Resulting Losses

	L_{cu}	L_{ob}	$L_{ce} = L_{cu}L_{ob}$	L_{pt}	$L_{ob}L_{pt}$	$L = L_{cu}L_{ob}L_{pt}$
	cubic	oblongitis	cell shape = cubic & obl.	point density	"shape" loss = pt. dens. & obl.	total
	m_1^*/m_∞^*				β_1/β_∞	$m_1^*\beta_1/m_\infty^*\beta_\infty$
Gaussian	$2\pi e/12$	$\sqrt{3}\,e^{-1/3}$	$4^{-1}3^{-1/2}\pi e^{2/3}$	$3\,e^{-2/3}$	$3\sqrt{3}/e$	$\pi\sqrt{3}/2$
dB	1.533	0.938	2.471	1.876	2.814	4.347
Laplacian	$2\pi e/12$	$3\,e^{-2/3}$	$2^{-1}\pi e^{1/3}$	$9\,e^{-4/3}$	$27/e^2$	$9\pi/2e$
dB	1.533	1.876	3.409	3.752	5.628	7.161

The oblongitis and point density losses are larger for Laplacian than for the Gaussian density, because the Laplacian's sharper peak at the origin and heavier tail means that a good scalar quantizer must be more nonuniform. This causes more oblongitis, which in turn causes more compromising of the optimal point density in order to reduce the oblongitis.

In summary, for an IID source the shortcomings of scalar quantization relative to high-dimensional VQ are

(a) The cubic loss L_{cu} = 1.53 dB, which is a measure of its inability to produce cells with smaller nmi than a cube.

(b) The lack of sufficient degrees of freedom to simultaneously attain good inertial profile and good point density.

Shortcomings of Optimal k'-dimensional VQ

One can similarly decompose the loss of optimal k'-dimensional VQ ($k' \geq 2$) relative to high-dimensional VQ into point density, oblongitis, and "space-filling" losses by comparing the characteristics of an optimal k'-dimensional VQ to that of an optimal high-dimensional VQ. To make the comparison, one considers the point density and inertial profile of the product quantizer formed by using the k'-dimensional VQ many times. The point density and oblongitis losses are then defined in the same way as before. The "space-filling" loss, which is

$$L_{sp} = \frac{m^*_{k'}}{m^*_\infty} ,$$

generalizes the "cubic loss" we considered for $k'=1$. It is called the "space-filling" loss because it represents the loss due to the product quantizers inability to fill space with cells whose NMI is better than that which induces m^*_k. (For k=1 it's a "cubic loss", for k=2 it's a "hexagonal loss"

As k' increases, one finds:

(a) The "space filling" loss decreases to 1 (0 dB).

(b) There are more degrees of freedom, so less compromise is needed between the k'-dimensional point density that minimizes oblongitis and that which minimizes point density loss. Consequently, when optimized, the oblongitis & point density losses are smaller, and decrease to 1 (0 dB)

Example: IID Gaussian Source

$$L_{ob} = \left(\frac{k+2}{k}\right)^{k/2} e^{-k/(k+2)}, \quad L_{pt} = \frac{k+2}{k} e^{-2/(k+2)}, \quad L_{pt} L_{ob} = L_{shape} = \frac{\beta_k}{\beta_\infty} = \frac{1}{e}\left(\frac{k+2}{k}\right)^{(k+2)/2}$$

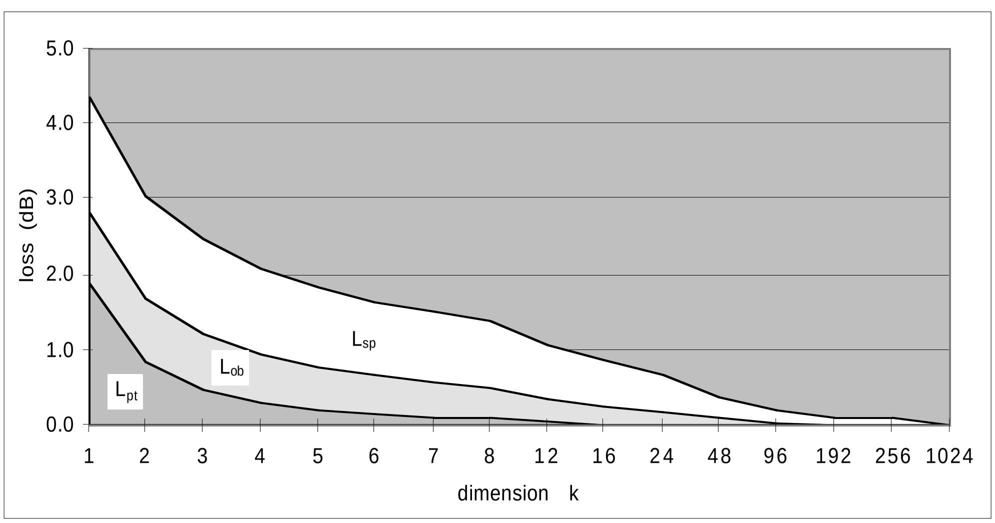

k' dimen.	L_{sp} space filling $m^*_{k'}/m^*_\infty$	L_{ob} oblongitis	$L_{ce} = L_c L_{ob}$ cell shape = ideal & obl.	L_{pt} point density	$L_{ob} L_{pt}$ "shape" loss = pt. dens. & obl. β_{k}/β_∞	$L = L_{cc} L_{ob} L_{pt}$ total $m^*_{k'}\beta_{k'}/m^*_\infty\beta_\infty$
1	1.533	0.938	2.471	1.876	2.814	4.347
2	1.366	0.839	2.205	0.839	1.678	3.043
3	1.276	0.722	1.998	0.481	1.203	2.479
4	1.138	0.627	1.765	0.313	0.940	2.078
5	1.055	0.551	1.606	0.220	0.772	1.827
6	0.986	0.491	1.477	0.164	0.655	1.640
7	0.927	0.442	1.369	0.126	0.569	1.495
8	0.876	0.402	1.278	0.101	0.503	1.378
12	0.725	0.294	1.019	0.049	0.343	1.068
16	0.624	0.232	0.855	0.029	0.261	0.884
24	0.494	0.163	0.657	0.014	0.176	0.671
48	0.277	0.086	0.363	0.004	0.089	0.366
96	0.169	0.044	0.213	0.001	0.045	0.214
192	0.100	0.022	0.122	0.000	0.023	0.122
256	0.080	0.017	0.097	0.000	0.017	0.097
1024	0.000	0.004	0.004	0.000	0.004	0.004
∞	0.000	0.000	0.000	0.000	0.000	0.000

Memory Loss

For sources with memory (e.g. autoregressive), scalar and low-dimensional quantization suffer an additional loss, namely the inability to exploit or fully exploit the memory, i.e the dependence between source samples.

Both oblongitis and point density losses can be factored into two terms, one of which expresses the quantizers inability to fully exploit the source memory.