Geometric Approach to Sampling and Communication

Emil Saucan
Technion

Joint work with

Eli Appleboim and Yehoshua Y. Zeevi.

German-Israel Workshop for Vision and Image Sciences

Haifa, November 4, 2008.

Geometric Approach to Sampling and Communication

or...

Geometric Approach to Sampling and Communication

or...

In Shannon's Long Shadow...

Why?...

Let's start with the following

Quiz: From what classical textbook in Calculus is this image taken?

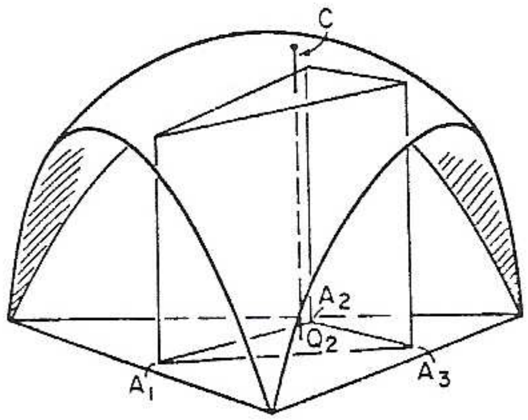

(a) Spivak

(b) Apostol

(c) Loomis

(d) Your favorite textbook

(e) My favorite textbook

Answer:

Neither! – It is from Shannon's paper *Some Geometrical Results in Channel Capacity**!...

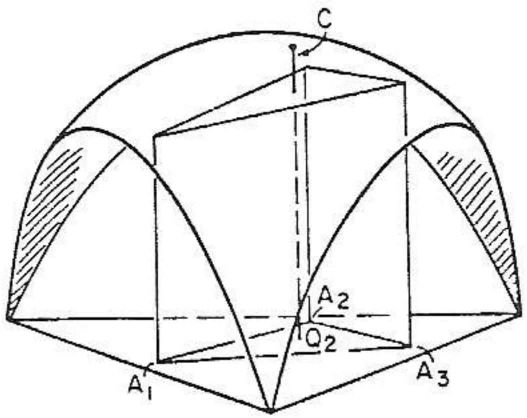

Fig. 1 Construction for channel capacity with three input and three output letters

Moreover, the geometric thinking doesn't stop at this picture, nor is it restricted to this paper...

*Nachrichtentechnische Zeit, vol. 10, 259-264, 1957.

... it is the "moving force" in the seminal (and by now classical) *Communication in the presence of noise**

Table 1

Communication System	*Geometrical Entity*
The set of possible signals	A space of $2TW$ dimensions
A particular signal	A point in the space
Distortion in the channel	A warping of the space
Noise in the channel	A region of uncertainty about each point
The average power of the signal	$(2TW)^{-1}$ times the square of the distance from the origin to the point
The set of signals of power P	The set of points in a sphere of radius $\sqrt{2TW\ P}$
The set of possible messages	A space of $2T_1W_1$ dimensions
The set of actual messages distinguishable by the destination	A space of D dimensions obtained by regarding all equivalent messages as one point, and deleting messages which the source could not produce
A message	A point in this space
The transmitter	A mapping of the message space into the signal space
The receiver	A mapping of the signal space into the message space

*Proceedings of the IRE, vol. 37, no. 1, 1949, 10-21.

In the same geometric spirit we have obtained earlier*

Theorem 1 *Let* $\Sigma^n, n \geq 2$ *be a connected, not necessarily compact, smooth manifold, with finitely many compact boundary components. Then there exists a sampling scheme of* Σ^n*, with a density* $\mathcal{D} = \mathcal{D}(p) = \mathcal{D}\left(\frac{1}{k(p)}\right)$*, where* $k(p) = \max\{|k_1|, \ldots, |k_n|\}$*, and where* $k_1, \ldots, k_n$ *are the principal curvatures of* Σ^n*, at the point* $p \in \Sigma^n$*.*

Corollary 2 *Let* $\Sigma^n, \mathcal{D}$ *be as above. If there exists* $k_0 > 0$*, such that* $k(p) \leq k_0$*, for all* $p \in \Sigma^n$*, then there exists a sampling of* Σ^n *of finite density everywhere. In particular, if* Σ^n *is compact, then there exists a sampling of* Σ^n *having uniformly bounded density.*

*Journal of Mathematical Imaging and Vision, **30**(1), 2008, 105-123. See also Leibon & Letsher, D. Proceedings of the sixteenth annual symposium on Computational geometry, 341 - 349, 2000.

We do not stress here

- The 1-dimensional case (and the comparison with the classical version of the Shannon Sampling Theorem)[*,†]

nor

- The intrinsic capability for Sparse Sampling[‡]

and we certainly do not explore the intricacies of the proof; we just mention a few essential steps with a view to their use for the building of our "dictionary":

[*]Journal of Mathematical Imaging and Vision, **30**(1), 2008, 105-123.

[†]CCIT Report #707 November 2008 EE Pub No. 1664.

[‡]Geometric Sampling for Signals with Applications to Images, SampTA 07 – Sampling Theory and Applications, 2008.

• Use Nash's Embedding Theorem to isometrically embed M^n in some $\mathbb{R}^N$, for some N large enough.

• Produce a point set $A \subseteq M^n$, that is maximal with respect a certain density condition (given by the curvature).

• Employ the density condition above to construct the finite cell complex "cut out of M" by the N-dimensional Voronoi complex, whose (closed) cells are given by:

$$\bar{c}_k = \bar{c}_k^{\nu} = \{x \in \mathbb{R}^{\nu} \,|\, d_{eucl}(a_k, x) \leq d_{eucl}(a_i, x),\ a_i \in A\,,\ a_i \neq a_k\}\,.$$

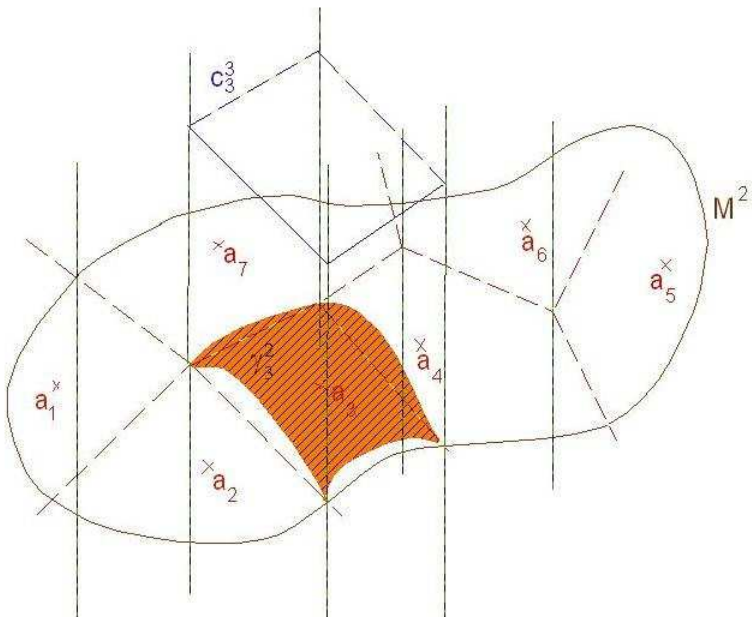

Let us muse a bit on the on the first point:

• While for surfaces in $\mathbb{R}^3$ (hence for BW images) the embedding is given (hence no problem there), but...

• Finding the precise embedding dimension N is excessively hard (even for very simple manifolds!....)

Moreover

• The canonical N is prohibitively high: $N = 17$ for a general surface (and even after further refinements due to Gromov and Günther the lowest embedding dimension for surfaces is $N = 10$!...)

However, let us despair not, but rather search for advantage and

Applications

- Pulse Code Modulation (PCM) for Images

Advantage: sampling points associated with relevant geometric features (via curvature), not chosen randomly via the Nyquist rate.

As a consequence the sampling is adaptive (and, indeed, compressive).

- Vector Quantization for Images

Immediate from Step 3 above.

Gives estimates for the error in terms of length and angle distortion when passing from the cell complex $\{\bar{\gamma}_k^n\}$ to the Euclidean cell complex $\{\bar{c}_k^n\}$ having the same set of vertices as $\{\bar{\gamma}_k^n\}$.

More precisely, if $M = M^n$ is a manifold without boundary, then (locally):

$$\frac{3}{4}d_M(x,y) \leq d_{eucl}(\bar{x},\bar{y}) \leq \frac{5}{3}d_M(x,y)\,;$$

where d_{eucl}, d_M denote the Euclidean and intrinsic metric (on M) respectively, and where $x, y \in M$ and $\bar{x}, \bar{y}$ are their preimages on the piecewise-flat complex.

If $M = M^n$ has a boundary ∂M, then:

$$\frac{3}{4}d_M(x,y) - f(\theta)\eta_\partial \leq d_{eucl}(\bar{x},\bar{y}) \leq \frac{5}{3}d_M(x,y) + f(\theta)\eta_\partial\,;$$

where $f(\theta)$ is a constant depending on the geometry of M and ∂M.

As noted, the embedding dimension is prohibitive (for practical use) in most cases.

But, instead of despairing, we can ask ourself the following

Question 3 *Can we overcome the "curse of dimensionality"?*

Better yet

Question 4 *Can we turn the "curse of dimensionality" to our advantage?*

The answer is "**Yes**" and resides in Zador's Theorem that states that there is an inherent advantage in using higher dimensional quantizers, more precisely one can reduce the *average mean squared error per dimension*

$$\mathcal{E} = \frac{1}{N}\int_{\mathbb{R}^N} d_{eucl}(x, p_i)p(x)dx \, ,$$

p_i being the *code point* closest to x and $p(x)$ denoting the *probability density function* of x, can be reduced by making avail of higher dimensional quantizers.

For embedded manifolds $p(x) = p_1(x)\chi_M$, we obtain:

$$\mathcal{E} = \frac{1}{N}\int_{M^n} d_{eucl}(x, p_i)p_1(x)dx \, .$$

It follows from Zador's Theorem that if the main issue is **accuracy**, not **simplicity**, then 1-dimensional coding algorithms* perform far worse than higher dimensional ones.

Unfortunately the proof of Zador's Theorem is nonconstructive, hence no **optimal** embedding dimension has been established.[†]

Our geometric coding method provides a **natural** and **constructive** high dimension for the quantization of M^n – the embedding dimension N.

*such as the classical Ziv-Lempel algorithm

[†]An upper limit for the coding dimension must exist – otherwise one could just code the whole data as one N-dimensional vector (albeit of unpractically high dimension)!...

Let us continue a bit with the "dictionary":

It is natural to extend the classical definitions of average power in the signal:

$$P = \frac{1}{T}\int_0^T f^2(t)dt\,,$$

and the rate of the code:

$$R = \frac{1}{T}\log_2 N\,,$$

in the context of lattices with fundamental cell λ, where N represents the number of code points, in the following manner:

$$P = \frac{1}{\mathrm{Vol}(\Lambda)}\int_\lambda f^2(t)dt = \frac{1}{N_1\mathrm{Vol}(\lambda)}\int_\lambda f^2(t)dt\,,$$

and

$$R = \frac{1}{\mathrm{Vol}(\Lambda)}\log_2 N = \frac{1}{N_1\mathrm{Vol}(\lambda)}\log_2 N,$$

respectively, N_1 being the number of cells.

Similarly, one can adapt the classical definition of the channel capacity:

$$C = \lim_{T\to\infty} R = \lim_{T\to\infty} \frac{\log_2 N}{T},$$

to become

$$C = \lim_{T\to\infty} \frac{\log_2 N}{\mathrm{Vol}(\Lambda)} = \lim_{T\to\infty} \frac{1}{N_1 \mathrm{Vol}(\lambda)} \log_2 N.$$

Since N and N_1 are related by $N_1 = \alpha(N)$, where α is the growth function of the manifold, the expression of C becomes:

$$C = \lim_{T\to\infty} \frac{1}{\mathrm{Vol}(\lambda)} \frac{\log_2 N}{\alpha(N)}.$$

Note that by putting $1/T = 1/\text{Vol}(M)$, the definitions above apply for any sampling scheme of any manifold of finite volume, not just for lattices.* In this case N and N_1 represent the number of vertices, respective simplices, of the triangulation.

The interpretation of frequency considered above does not extend, however, to general signals. For a proper generalization we have to look into the **geometric** analogue of W.

We start with the basics: 1-dimensional signals...

*For lattices. geometric measures such as $\text{diam}(\lambda)$ (or, alternatively, the length of the longest edge) and its volume $\text{Vol}(\lambda)$ replace the duration of "standard" signal.

But, for 1-dimensional (with normalized energy) signals, W equals the curvature rate $k/2$, were k represents the maximal absolute curvature of the curve*. This, and our more general Sampling Theorem,†, naturally leads us naturally to the following definition of W for general "geometric signals":

Definition 5 *Let $M = M^n$ be an n-dimensional manifold $n \geq 2$. $W = W_M = 1/k_M$, where $k_M = \max k_i$ and $k_i, i = 1, \ldots, n$ are the principal curvatures of M.*

*Journal of Mathematical Imaging and Vision

†and since a stick point in two directions...

But what about **Shannon's Second Theorem** (and the **Channel Coding Problem**)?...

Recall that in the classical context a received signal is represented by a vector $X = F + Y$, where $F = (f_1, \ldots, f_N)$ is the transmitted signal, and $Y = (y_1, \ldots, y_n)$ represents the noise, whose components y_i are independent Gaussian random variables, of mean 0 and average power σ^2.

Theorem 6 (Shannon's Second Theorem) *For any rate R not exceeding the capacity C_0:*

$$C_0 = W \log_2\left(1 + \frac{P}{\sigma^2}\right),$$

where W represents the frequency, there exists a sufficiently large T, such that there exists a code of rate R and average power $\leq P$, and such that the probability of a decoding error is arbitrarily small. Conversely, it is not possible to obtain arbitrarily small errors for rates $R > C_0$.

Before we try and formulate a Geometric Version, let us take a look at Shannon's view on this problem:*

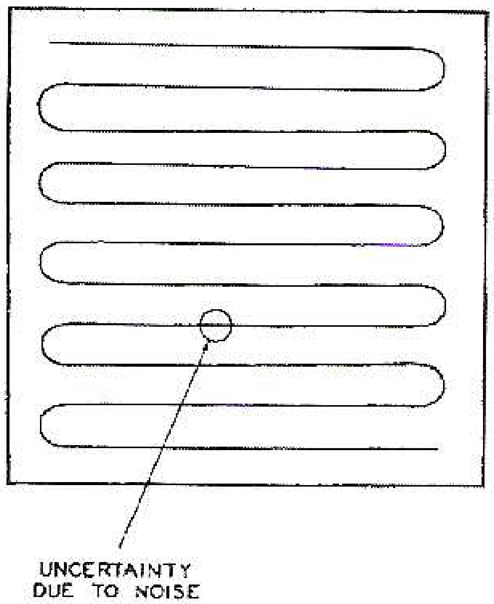

Fig. 4. Efficient mapping of a line into a square.

**Communication in the presence of noise*, Proceedings of the IRE, vol. 37, no. 1, 1949, 10-21.

In the case of geometric signals, F is given by the sampling (code) points on the manifolds and, since the mean equals 0, the noisy transmitted signal $F+Y$ lies in the tube* $\mathrm{Tub}_\sigma(M)$:

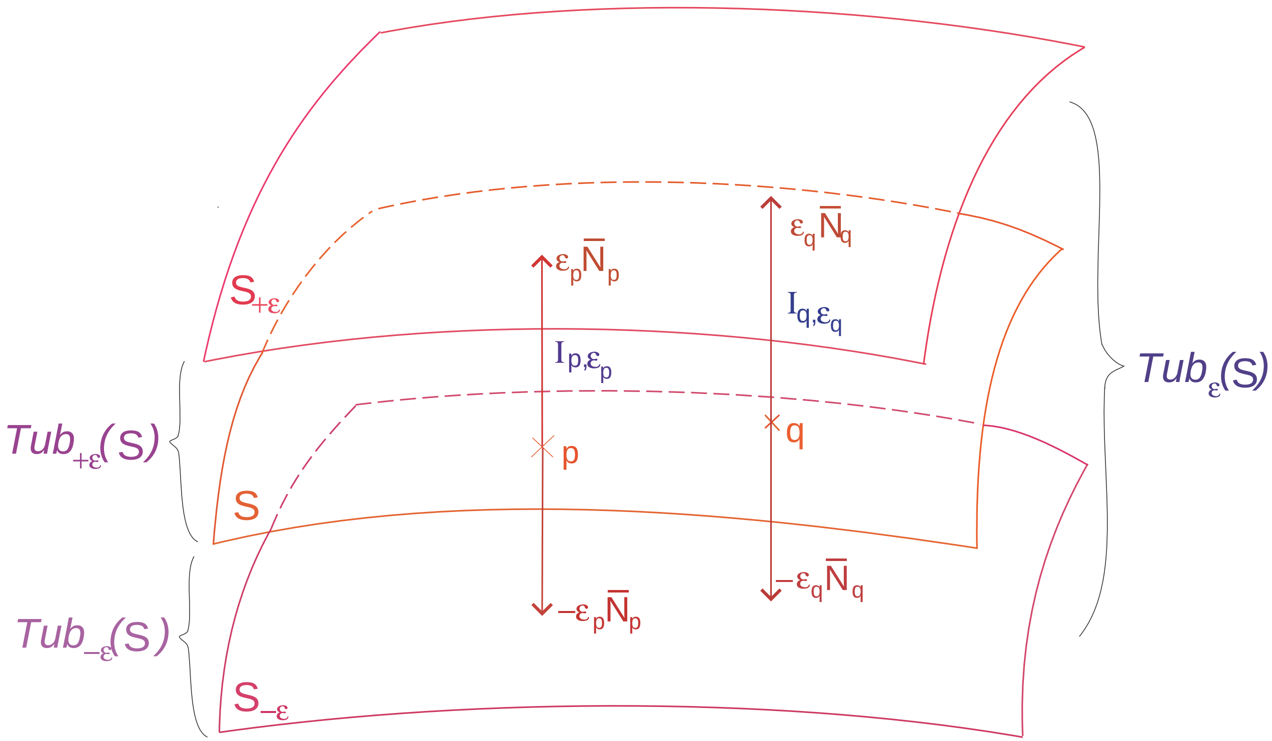

*a standard tool in (Geometric) Statistics!...

Remark 7 *The existence of tubes is assured locally, for any regular, orientable surface, and globally for regular, compact, orientable surfaces.*

Remark 8 *Note the following facts regarding the regularity of $S_{\pm\varepsilon}$:*

- *If S is convex, then $S_{\pm\varepsilon}$ are piecewise $\mathcal{C}^{1,1}$ surfaces (i.e., they admit parameterizations with continuous and bounded derivatives), for all $\varepsilon > 0$.*

- *If S is a smooth enough surface with a boundary (that is, at least piecewise $\mathcal{C}^2$), then $S_{\pm\varepsilon}$ are piecewise $\mathcal{C}^2$ surfaces, for all small enough ε.*

- *For* **any** *compact set $S \subset \mathbb{R}^3$, $S_{\pm\varepsilon}$ are Lipschitz surfaces for a.e. ε.*

In the geometric setting, $\sigma \equiv \varepsilon$ can be taken, of course, to be the maximal Euclidean deviation. However, a better deviation measure is, at least for compact manifolds, the Haussdorf Distance (between M and $\partial \mathrm{Tub}_{\sigma}^{-}(M)$, $\partial \mathrm{Tub}_{\sigma}^{+}(M)$):

Definition 9 *Let* (X, d) *be a metric space and let* $A, B \subseteq (X, d)$*. The Hausdorff distance between* A *and* B *is defined as:*

$$d_H(A, B) = \max\{\sup_{a \in A} d(a, B), \sup_{b \in B} d(b, A)\}\,.$$

For non-compact manifolds one has to consider the more general Gromov-Hausdorff distance.

Since both the distance between M and $\partial\text{Tub}_\sigma^-(M)$, $\partial\text{Tub}_\sigma^+(M)$ and the deviations of their curvature measures are arbitrarily small for small enough σ, we can state the following Qualitative

Theorem 10 (Shannon's Second Theorem, Geometric version) *Let M^n be a smooth geometric signal (manifold) and let σ be small enough, such that $\text{Tub}_\sigma(M)$ is a submanifold of $\mathbb{R}^{n+1}$. Then, given any noisy signal $M + Y$, such that the average noise power σ_Y is at most σ, there exists a sampling of $M + Y$ with an arbitrarily small probability of resultant decoding error.*

A few brief remarks:

- The analogue of the capacity in the context of the geometric approach to codes is $C_0 = C_0(n, \sigma, r)$, where r represents the differentiability class of M.

- For compact manifolds the sampling scheme necessitates $O(N)$ points, $N = N_M$. For noncompact manifolds it requires $N_{M+Y} = O(N_M^n)$ sampling points.

- A Quantitative version is in preparation.

More details and considerations on

- Implementation by Wavelets,

- The Uncertainty Principle,

and

- Applicability to Infinite Dimensional Signals,

can be found in our preprint Technion CCIT Report #707 (EE Pub. # 1664 November 2008)

Thanks for your attention!