Effects of Data Geometry in Early Deep Learning

Saket Tiwari

Department of Computer Science
Brown University
Advisor: Professor George Konidaris

Motivation

- Deep Neural Networks (DNNs) can approximate arbitrarily complex functions.

Motivation

- Deep Neural Networks (DNNs) can approximate arbitrarily complex functions.
- What makes them so effective?

Motivation

- Deep Neural Networks (DNNs) can approximate arbitrarily complex functions.
- What makes them so effective?
- How does data geometry affect DNNs ability to approximate these functions?

Motivation

- Deep Neural Networks (DNNs) can approximate arbitrarily complex functions.
- What makes them so effective?
- How does data geometry affect DNNs ability to approximate these functions?
- Can we derive theoretical results that can be corroborated empirically, to understand DNNs approximation capacity?

Linear Regions as a Proxy for Approximation Capacity

We consider DNNs with piece wise linear activation function (e.g. ReLU).

$$\sigma(x) = \max(0, x)$$

DNNs divide the input space into pieces with it computing a single linear function on each piece.

Definition

The set $A \in \mathbb{R}^{n_{\text{in}}}$ is defined to be a *linear region* of a neural network F if it is the maximal connected such set where F behaves as a linear function.

Counting the number of linear regions tells us "how non linear a function is?"

How many Linear Regions?

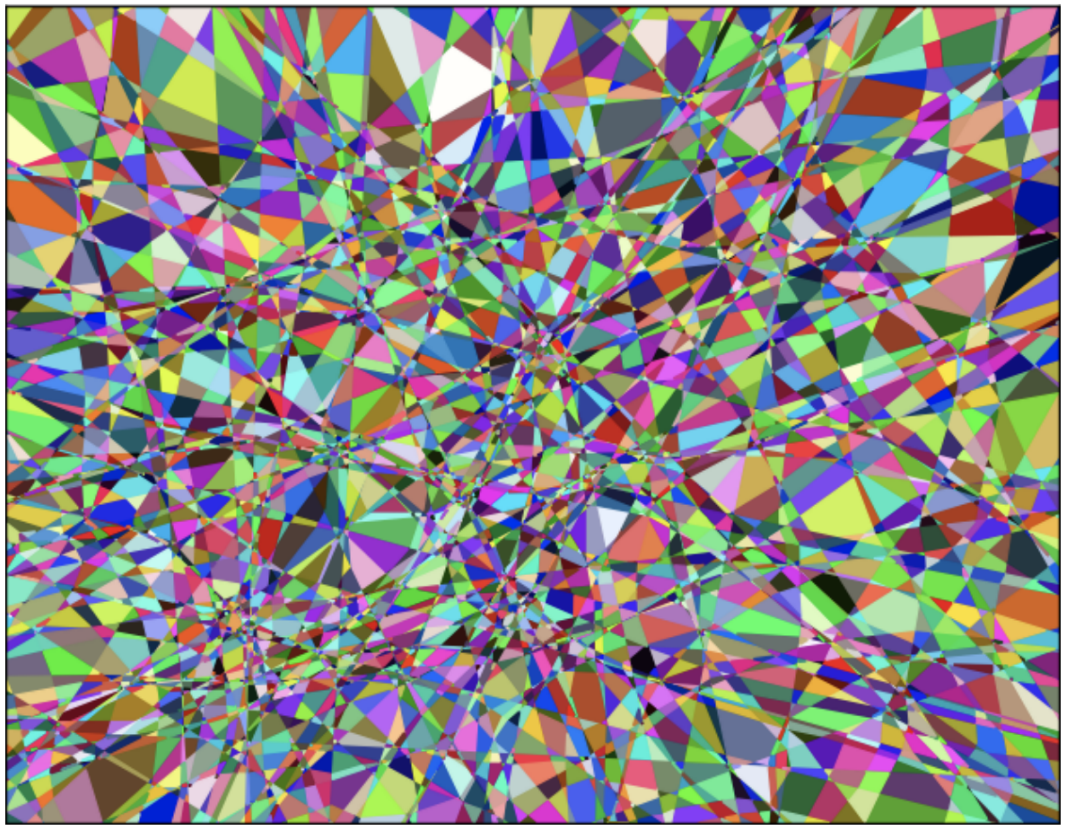

Figure: A two-dimensional slice through the 784-dimensional input space of vectorized MNIST, as represented by a fully-connected ReLU network with three hidden layers of width 64 each (Figure by Hanin and Rolnick, 2019).

What is Data Geometry?

The Manifold Hypothesis

Natural data forms lower-dimensional manifolds in its embedding space.

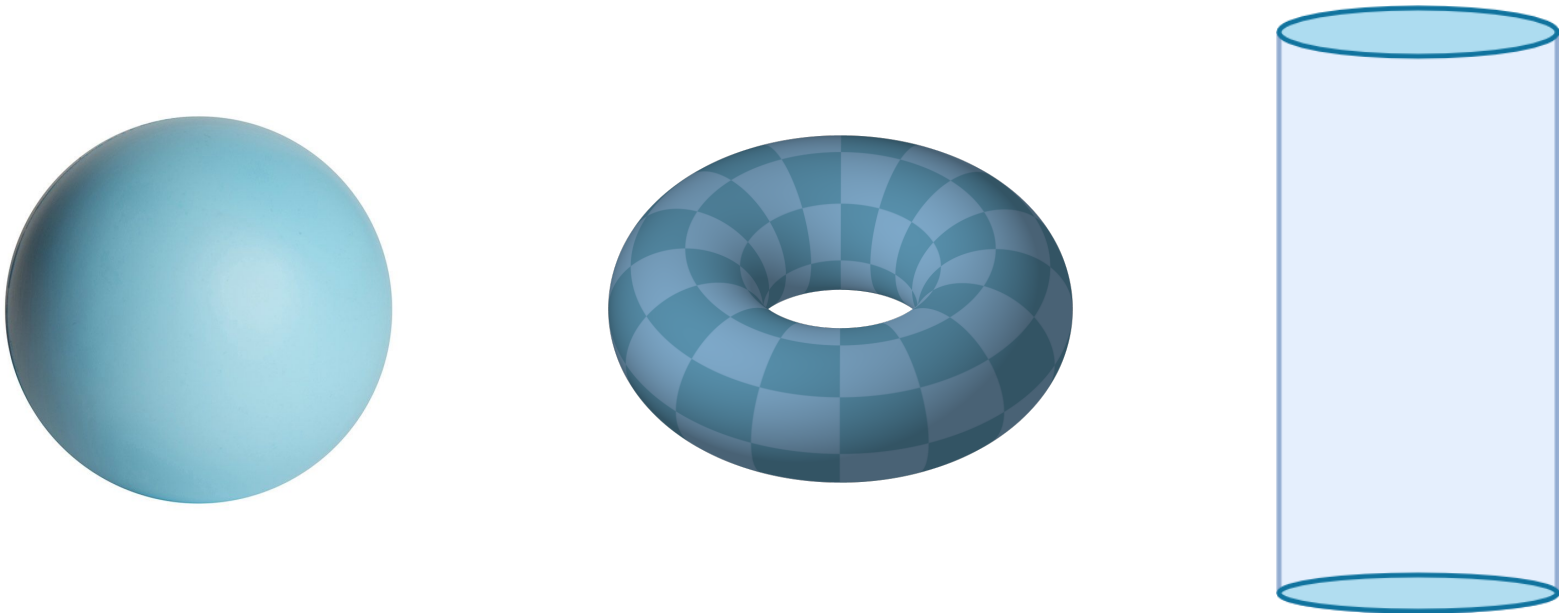

Figure: 2D Manifolds embedded in 3D spaces: Sphere, Torus, and Cylinder

The Manifold Hypothesis

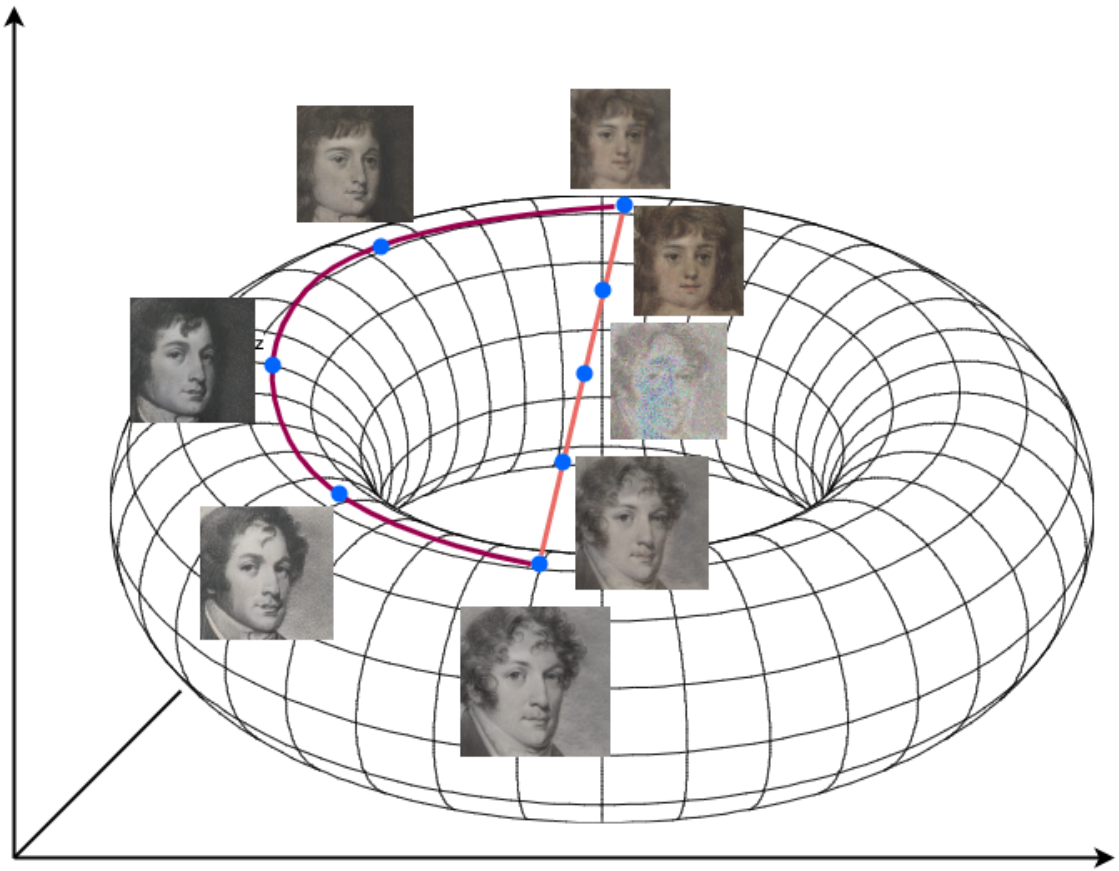

Figure: A visualization of how the 2D surface, here represented by a 2-torus, is embedded in a larger input space, R^3. Suppose each point corresponds to an image of the face on this 2-torus. We can chart two curves: one straight line cutting across the 3D space and another curve that stays on the torus.

The Manifold Correction

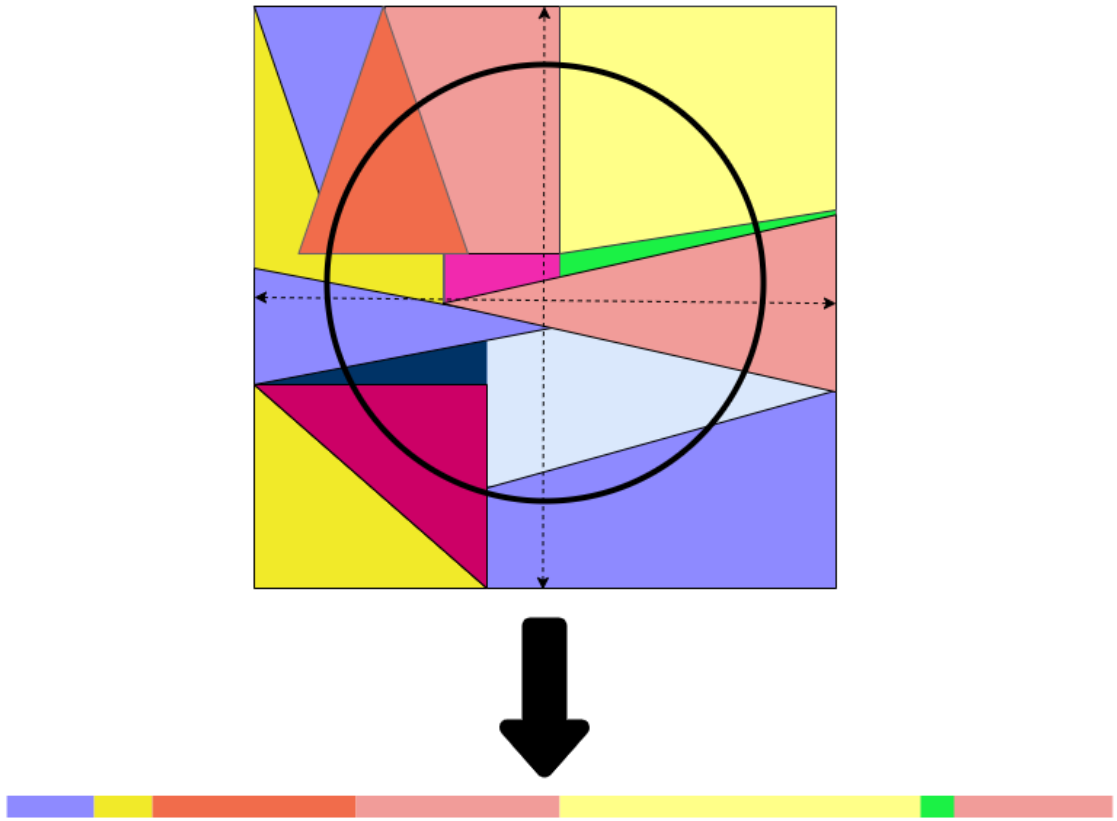

Figure: Linear regions of a DNN over a one-dimensional manifold embedded in a two-dimensional space. Each color uniquely represents a linear region.

Our Theoretical Results: Stronger Bounds

Theorem

Theorem 1: *For data sampled uniformly from a compact and measurable m dimensional manifold M:*

$$\mathbb{E}\left[\frac{vol_{m-1}(\textbf{Boundary regions in manifold } M)}{vol_m(M)}\right] \leq 2\#neurons C_{grad} C_{bias} C_M,$$

where C_{grad} depends on the gradient of the individual neurons w.r.t the input, C_M the geometry of M, and C_{bias} on the distribution of biases.

Our Theoretical Results: Stronger Bounds

Theorem

Theorem 2: *For any point, x, chosen randomly from M, we have:*

$$\mathbb{E}[\mathit{distance}_M(x, \mathcal{B}_F \cap M)] \geq \frac{C_{M,\kappa}}{C_{grad} C_{bias} C_M \#\mathit{neurons}},$$

where $C_{M,\kappa}$ depends on the scalar curvature, the input dimension and the dimensionality of the manifold M. The function $\mathit{distance}_M$ is the distance on the manifold M.

Empirical Observations: Supervised Learning

Consider two different settings where data sampled from two different one-dimensional manifolds with different geometries embedded in $\mathbb{R}^2$: Circle and Tractrix.

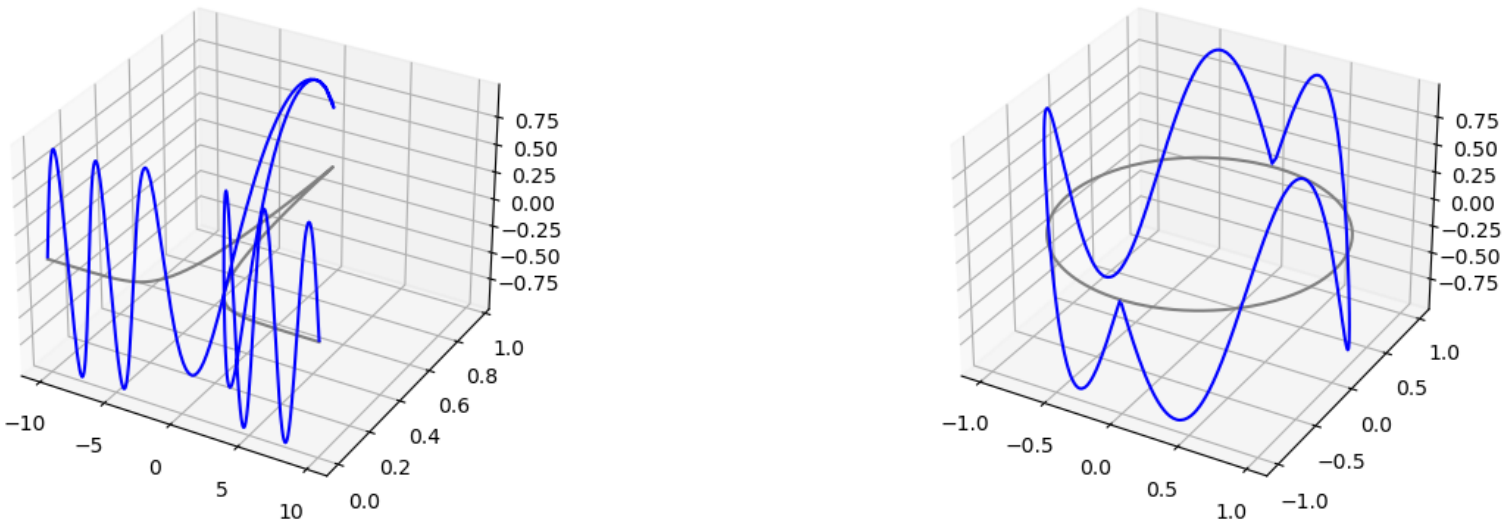

Figure: 1-D data manifolds embedded in 2D space: Circle and Tractrix

Empirical Results

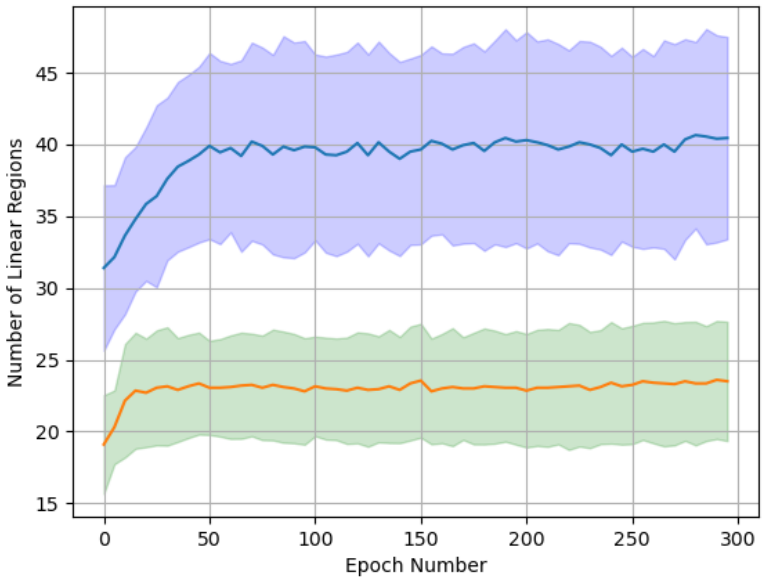

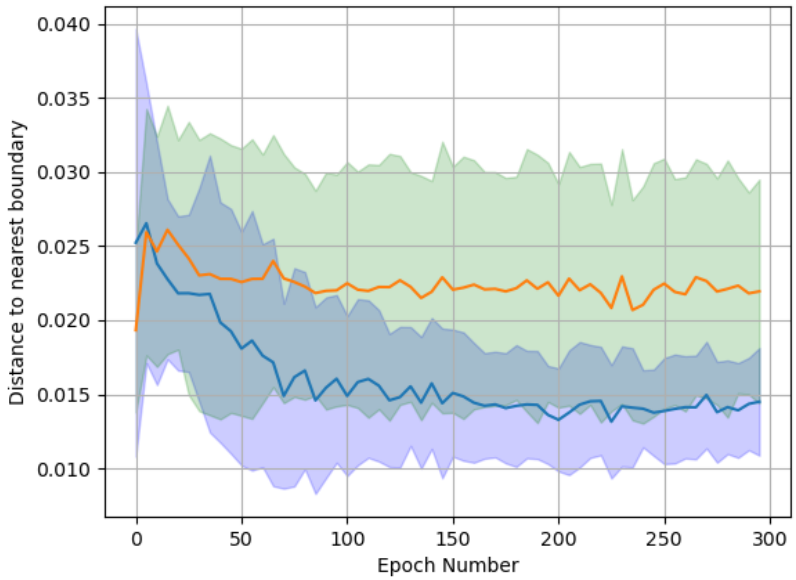

Graphs for Tractrix (blue) and Sphere (orange). We see that the number of linear regions stays close to the number of neurons (26). The difference in number of regions is attributed to different geometries.

Empirical Results

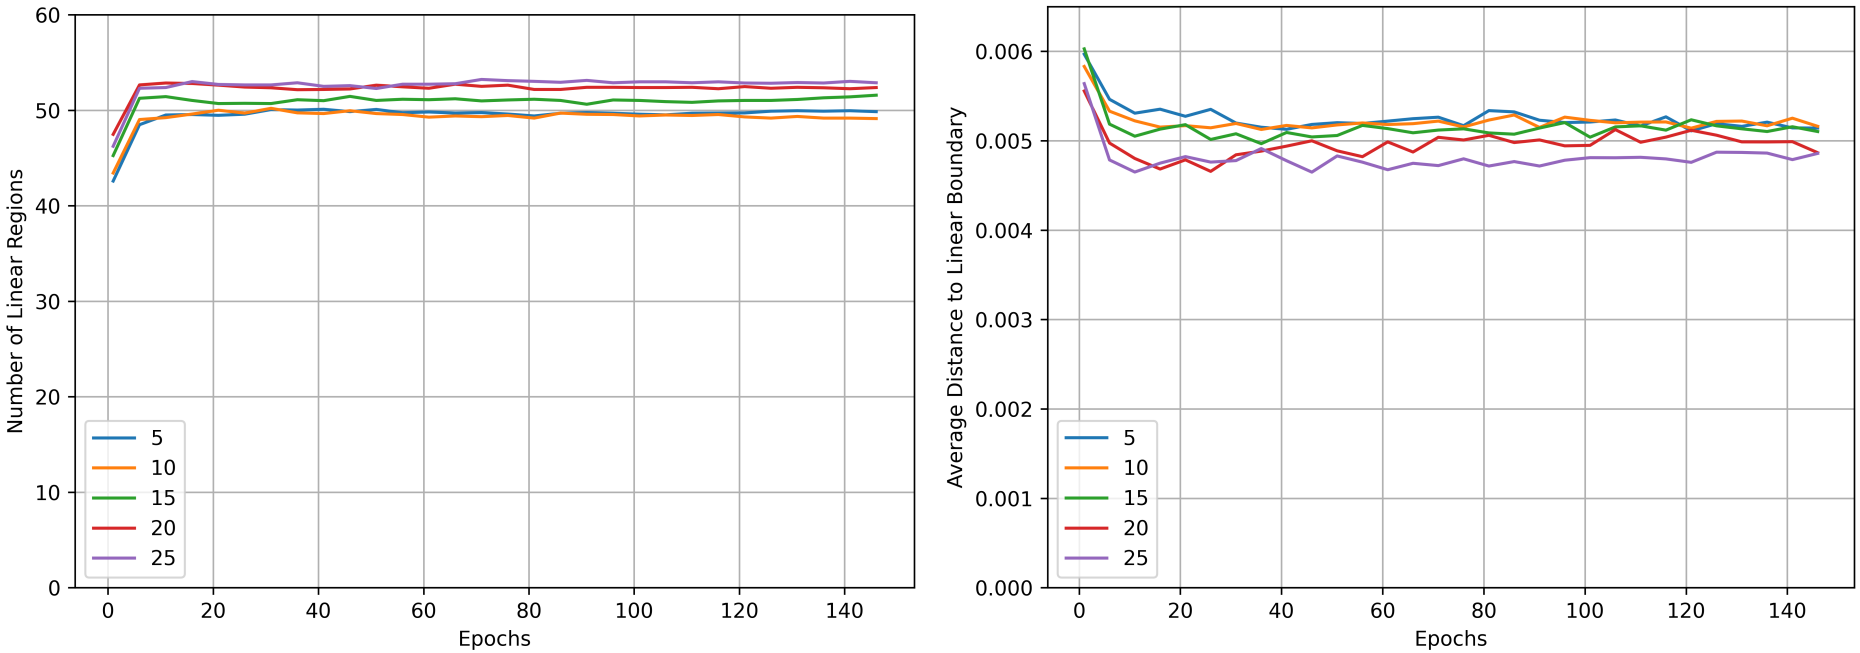

Figure: As we increase n_{in}, for a sphere, while keeping m constant at 1 we notice that the number of linear regions and the average distance to the linear boundary stay almost constant.

Figure: Shirana Shahbazi, [Composition-45-2011]