

Machine Learning, Spring 2019: Exercise Sheet 2

Problem 1 (a conceptual challenge, not a mathematical challenge). Consider the set of all photographic color (RGB) images that have been posted on the web, and assume they have been normalized to a size of 1000 by 1000 pixels (times 3 for the three color channels). Color intensity values are between 0 (black) and 1 (brightest). An image thus normalized therefore is a point in the three-million-dimensional unit hypercube $I = [0, 1]^{3,000,000}$. But most points in I , if you would print their corresponding image, will not look like a photograph of some real-world scene at all – they will just be “colored noise” pictures. The real world and the physics of taking photographic images impose many constraints on points in I – only points in I which satisfy these constraints will appear like a photograph of a real-world scene. Mathematically speaking, these constraints confine the “real photo-like” points to a relatively low-dimensional manifold R within I . Your task: develop an educated guess on what is the dimension of this manifold R . – This is not an easy problem, because you have to think about all kinds of regularities and factors which govern the appearance of real-world-like images. Some of these regularities will be derived from low-level geometry (like “neighboring pixels have a tendency to have similar color values”), others will be derived from just how our world is (“if there are two nearby patches that appear like eyes, then very likely there is a patch in the image close by that looks like a nose”). Any such regularity trims down the three million dimensions of I toward the dimension of R . In summary, how low-dimensional do you think is the “Flickr data manifold”? Try to argue for a *lower bound*!

Lower bounds *can* be argued; but arguing for a “correct” manifold dimension is ill-defined. As explained in the LNs, real image data points will jitter slightly around the “ideal” manifold, and this jitter can be in any direction in $[0, 1]^{3,000,000}$, so the true dimension of the manifold is the full dimension 3,000,000. Thus, a more technically correct statement of the problem is, “argue for an estimate of a manifold dimension such that any real-world photograph x , when projected to the nearest point x' in the manifold, leads to a pair x, x' which cannot be distinguished by the human eye.”

Problem 2 Consider the set $G = \{(x, y) \in \mathbb{R}^2 \mid 0 \leq x \leq 1, y = x^2\}$. This is the set of all points belonging to the *graph* of the square function within the unit square. It is a 1-dimensional manifold, of the familiar parabola shape, embedded in the 2-dimensional interval $[0, 1] \times [0, 1]$. Give a map $f: G \rightarrow \mathbb{R}$, which “flattens” the parabola, similar to what you see is happening in panel **a** in Figure 4 in the lecture notes. Specifically, f should preserve distance along the parabola. That is, if two points $(x, y), (x', y')$ have a distance d , measured along the parabola, then the distance between $f(x, y)$ and $f(x', y')$ should also be d .

Hint: this problem leads to evaluating an integral. At https://en.wikipedia.org/wiki/List_of_integrals_of_irrational_functions you will find the integral formula that solves the integral that you will need.