1 Horizontal and Vertical Lines

Take the set X to be the unit square $[0,1]^2$, and a hypothesis class $\mathcal{H}$ all of the horizontal and vertical lines that intersect this square. This has finite VC dimension of 2, as you can't shatter 3 points. We attempt to construct a fractal witness set: a fractal subset of the square, with Minkowski dimension not much bigger than 1, that has non-trivial intersection with every line in $\mathcal{H}$. Of course, we could choose the diagonal which would give us an exactly line.

We do so recursively. We divide the unit square into an $n \times n$ grid of pixels, find a witness set of pixels, then divide those pixels up to find another witness set, and so on. We eventually take the intersection.

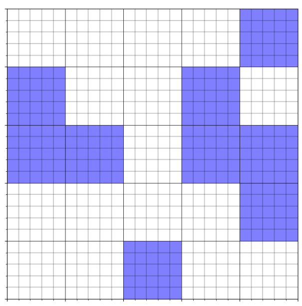
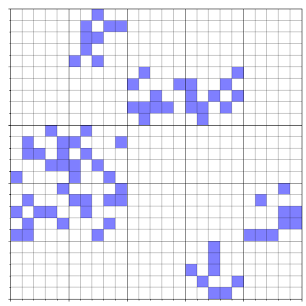

At the k^{th} step (starting at stage 0), we have a square of side length $\left(\frac{1}{n}\right)^k$. We divide this square into a grid of n^2 pixels, each side-length $\left(\frac{1}{n}\right)^{k+1}$. The pixels in our grid is a new set A. For every horizontal line that intersects the square, it intersects exactly n pixels of A, an entire row; a column for a vertical line. These rows and columns of n points each form our new hypothesis class/range space R.

Our witness set definition states that for $\epsilon > 0$, then $N \subseteq A$ is an ϵ-net if for any $r \in R$ (a row or column) such that $|r \cap A| > \epsilon |A|$. Since each $|r \cap A| = \frac{1}{n}$, so if $\epsilon < \frac{1}{n}$, then any epsilon net will contain at least one pixel in every row and column.

Our witness set theorem states that there exists a witness set of size, for some constants,

$$\frac{C_0}{\epsilon} \log \frac{C_1}{\epsilon}$$

Precisely, the probability that a Witness set of this size exists is non-zero, so one exists.

So we know that, for our set A, there is a witness set of the rows and columns of size

$$c_0 n \log(c_1 n)$$

So we have $c_0 n \log(c_1 n)$ pixels that are sure to intersect all of the horizontal and vertical lines that go through this particular square. We then take all of the chosen pixels to be the start of our next level: divide each of them up into their own grid, and complete the same process. Since no row or column will ever be lost– we are always choosing some small pixel in every row and column, then these sets will intersect non-trivially with every horizontal and vertical line. This gives us a fractal witness set.

Starting at the unit square $X = E_0$, we divide into n^2 side-length $\frac{1}{n}$ pixels, and take $c_0 n \log(c_1 n)$ to be E_1. In each of those, we divide into n^2 side-length $\frac{1}{n^2}$ pixels, and take $c_0 n \log(c_1 n)$, giving us

$$(c_0 n \log(c_1 n))^2$$

pixels of side-length $\frac{1}{n^2}$ that intersect every horizontal and vertical line, these form E_2. Continuing down, at step k we have E_k made up of

$$(c_0 n \log(c_1 n))^k$$

pixels of side-length $\frac{1}{n^k}$ that form a witness set. We take the intersection of all of these; equivalent to taking E_∞, we get a witness set $E = \bigcap_{k=0}^{\infty} E_k$.

Now to find the Minkowski of E, we find an upper and lower bound on the covering number $N_\delta(E)$. Since for every k, $E \subset E_k$, so we know $N_\delta(E) \leq N_\delta(E_k)$. When $\frac{2}{n^k} < \delta$, then one δ ball can cover a pixel of length $\frac{1}{n^k}$ (but $\delta \leq \frac{2}{n^{k-1}}$ ensures that it wouldn't cover not a pixel at one level up/larger length). Using one ball for each pixel, we can exhibit a covering by covering each square. Then an upper bound on the covering number is the number of squares, $(c_0 n \log(c_1 n))^k$.

Since $k - 1 \leq \log_n(\frac{2}{\delta})$ and $x^{\log_n(y)} = y^{\log_n(x)}$

$$\begin{aligned}
N_\delta(E) &\leq (c_0 n \log(c_1 n))^k \\
&\leq (c_0 n \log(c_1 n)) \, (c_0 n \log(c_1 n))^{\log_n(\frac{2}{\delta})} \\
&\leq (c_0 n \log(c_1 n)) \left(\frac{2}{\delta}\right)^{\log_n(c_0 n \log(c_1 n))} \\
&= (c_0 n \log(c_1 n)) \left(\frac{2}{\delta}\right)^{\frac{\log(c_0)+\log(n)+\log(\log(c_1 n))}{\log(n)}} \\
&= (c_2 n \log(c_1 n)) \left(\frac{1}{\delta}\right)^{1+\frac{\log(c_0)+\log(\log(c_1 n))}{\log(n)}}
\end{aligned}$$

This is what we want: n is chosen before δ, but we can make it as large as we want such that the exponent of $\frac{1}{\delta}$ is less than $1 + \alpha$ for any α.

Similarly for the lower bound: every pixel of E_k will contain some part of the witness set, so to cover E, we will have to cover something in every pixel of E_k. When $\frac{1}{n^{k+1}} \leq \delta < \frac{1}{n^k}$, then one δ-ball can cover at most 4 pixels of E_k (and not at the next level). So the covering number has to contain at least $\frac{1}{4}(c_0 n \log(c_1 n))^k$ balls. This gives the same exponent as above.

So the dimension of this witness set, for a fixed n, is

$$1 + \frac{\log(c_0) + \log(\log(c_1 n))}{\log(n)}$$

Which means we can get a fractal witness set with dimension as close as we want to 1 by taking n large enough because $\lim_{n\to\infty} \frac{\log(c_0)+\log(\log(c_1 n))}{\log(n)} = 0$.

2 Diagonal Lines

The goal now is to complete the same process for diagonal lines in the unit square. Thus, take the set X to be the unit square $[0,1]^2$, and a hypothesis class $\mathcal{H}$ all of lines that intersect this square. We are hoping to construct a witness set for all of the lines that intersect the square with "sizable" length- at least some length. For now, we'll say that lines have to intersect at least half of the width or height of the square.

Re-run the same process as before. Divide the square into an $n \times n$ grid of pixels. The range in our discretized grid corresponding to any line are the pixels, each squares of side-length $\frac{1}{n}$, that the line goes through at least half of their width or height. By the nature of how our lines and squares line up, we know that

Lemma. *For a square divided into a $k \times k$ grid of pixels, any line that intersects over half of the square's width or height will intersect over half of a pixel's width or height for at least $\frac{k}{2}$ pixels of the grid.*

Proof. TODO. Dividing into cases of slope $|m| \leq 1$ or $|m| \geq 1$, and thinking about columns or rows respectively. $\square$

Thus, the lines we are considering will each have a corresponding range of at least $\frac{n}{2}$, or a $\frac{1}{2n}$ proportion of the n^2 pixels in the grid. And, as the lines intersect these over half of these smaller square's width or height, we can repeat the same process on them. We do, however, need to know that this range space has finite VC dimension:

Lemma. *The above range space made up of the grid and pixels with sizable intersection has VC dimension that is bounded as the grid gets finer, $n \to \infty$.*

Proof. TODO. Probably true for the similar reasons line's have finite VC dimension. $\square$

When $\epsilon < \frac{1}{2n}$, then $|r \cap A| > \frac{n}{2} > \epsilon|A|$ for all of the ranges corresponding to our lines, so we know there is a witness set of size

$$c_0 n \log(c_1 n)$$

that intersects all of these (the 2 gets absorbed in the constant). That would be the set E_1, which we iterate this process on.

We then can complete the same process as above to show that we can get this witness set to have dimension as close to 1 as we would like.