A Study of Points in Hypercubes

Siddartha Devic, Rocky Cardenas, Yusuf Ahsanullah, Francisco Villegas, Michelle Sferrazza

UTD & UTRGV

December 2nd, 2019

Table of Contents

Motivation: The Simple Case

Consider the following problem from a previous homework.

Problem

Consider any five points within the unit square. Show that one can always find two of the points at distance at most $\frac{\sqrt{2}}{2}$ apart.

Solution: Divide the unit square into quadrants. Each quadrant is a square and has side length of $\frac{1}{2}$. By the pigeonhole principle, two points must lie within the same quadrant.

Next, given two points in a single quadrant, the farthest they could be apart is if each point was placed at opposite diagonals, which is $\frac{\sqrt{2}}{2}$ by the Pythagorean theorem.

Unit square

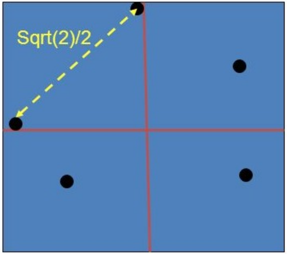

Figure: Unit square with five points lying inside

Potential Generalizations and Variants

We can extend the problem in multiple ways.

1. Consider k points in a n-dimensional hypercube
2. Consider a distance d to be given. How many points do we need to guarantee that two will be at most d apart (in the hypercube)?
3. Consider randomly sampling k points in the hypercube. Can we bound d, the distance between the closest pair of points?

n-tants

Definition

***n*-tant**: An n-tant is the slice of space in $\mathbb{R}^n$ defined by an n-tuple $\{+, -\}^n$, where each index $k \in (1, n)$ in the n-tuple determines which side of zero the slice will lie along dimension k.

A Natural Order

- For convenience, we induce a natural ordering based on the sign of the coordinates. Although we focus on the finite dimensional case, the countably infinite case follows similarly.
- We let the first n-tant be the set of all points where each coordinate is positive. The second n-tant lies where the sign of the first coordinate within the n-tuple is restricted to negative values while the other coordinates remain positive.
- The third through n'th are defined in similar manner, flipping the sign of the $i^t h$ component to negative, where $i \in [1, n]$.
- The $n + 1$'th through $n + 1 + \binom{n}{2}$'th n-tants are defined by flipping every combination of two indices two negative values, in the natural ordering. Similarly for the following n-tants, we simply iterate through all $\binom{n}{k}, k \in [1, n]$ n-tants by permuting signs in the natural order.

A Natural Order Contd

As an example, we list the eight 3-tants in our natural ordering:

1. $(+, +, +)$
2. $(-, +, +)$
3. $(+, -, +)$
4. $(+, +, -)$
5. $(-, -, +)$
6. $(-, +, -)$
7. $(+, -, -)$
8. $(-, -, -)$

Our 3-tant labeling method is different from the conventional octant labeling (and similarly, 2-tant order is not the same as quadrant), but will be convenient in arbitrary finite dimension space like R^n.
Lastly, notice that every point in space falls either in the interior one of these n-tants or on the boundary. If the point lies on the boundary of multiple n-tants, we introduce the convention of including the point in the n-tant of the lowest index number.

n-tants and the Euclidean Metric

Theorem

Given $2^n + 1$ points in a hypercube $H \in \mathbb{R}^n$ with side lengths 2, there exist a pair of points that are at most $\sqrt{n}$ apart.

proof: Suppose we have $2^n + 1$ points distributed through the interior of H. Without loss of generality, center H at the origin.
Divide the hypercube into 2^n n-tants. By the pigeonhole principle, two points must lie within the same n-tant.
The maximal distance between 2 points in the same n-tant is achieved by maximizing the sum of the squares the differences in each component.
Without loss of generality, let us operate in the first n-tant, where every coordinate component is restricted to [0,1]. The distance clearly achieves its max at $\sqrt{(1-0)^2 + \cdots + (1-0)^2}$, simplifying to $\sqrt{n}$.

an upper bound on the distance

We see then that this application of the pigeonhole principle has given us an upper bound, and that we can always find at least 2 points within $\sqrt{n}$ given $2^n + 1$ points on our 2-hypercube.
(side length 2 was selected for convenience of calculation, but we claim without proof that this choice was arbitrary, since we can apply some substitution $y = 2x$, where x was our unit length).
Question: Is there some smaller distance that we are still guaranteed to find 2 points within?
No, at least for $n < 5$.

Theorem

The minimal distance between two points among $2^n + 1$ points distributed within some hypercube H is exactly $\sqrt{n}$ for $n < 5$.

Exactly $\sqrt{n}$

Theorem

The minimal distance between two points among $2^n + 1$ points distributed within some hypercube H is exactly $\sqrt{n}$ for $n < 5$.

proof: Suppose we had some distance less than $\sqrt{n}$ that we claimed as possible, I will call it $\sqrt{n} - \epsilon$, where $\epsilon > 0$.

Consider the configuration that puts 1 point in the center of H, and then puts the others on the 2^n spots near the corners at distance exactly $\sqrt{n} - \epsilon/2$ from the center. **Notice the** $\epsilon/2$.

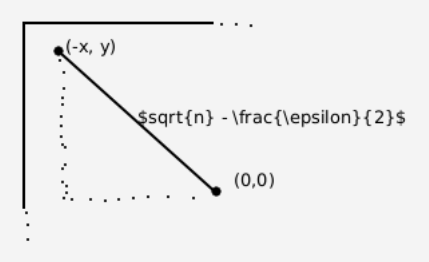

Euclidean metric contd.

In 2d, we have the following when attempting to solve for the coordinates:

$$2x^2 = (\sqrt{n} - \frac{\epsilon}{2})^2$$

In the hypercube, we must solve for x in the generalization.

$$\begin{aligned} nx^2 &= (\sqrt{n} - \frac{\epsilon}{2})^2 \\ \sqrt{nx^2} &= \sqrt{n} - \frac{\epsilon}{2} \\ x &= 1 - \frac{\epsilon}{2\sqrt{n}} \end{aligned}$$

Formally speaking, subtract $\epsilon/(2\sqrt{n})$ from the each component of the corner's coordinates (subtract if w.l.o.g. we are operating in the first n-tant, but at each corner we are adding/subtracting this distance with the sign that keeps it in the hypercube).

Using n-tants

We can nicely describe each of these placed points using our previously defined n-tants.

Enumerate over each n-tant represented by N_k. Then, the corresponding location for the point within its respective n-tant (ignoring the center point of H) is given by:

$$P_k = N_k \circ \mathbb{1}(1 - \frac{\epsilon}{2\sqrt{n}})$$

Where $\circ$ is the Hadamard element-wise product, and $\mathbb{1}(1 - \frac{\epsilon}{2\sqrt{n}})$ indicates a vector where each component is equal to the value $1 - \frac{\epsilon}{2\sqrt{n}}$.

Euclidean metric contd.

Place each $P_k, k \in [1, 2^n]$ points at these near-corner locations, and P_0 at the center of H. We find that no 2 are within distance $\sqrt{n} - \epsilon$ (so long as the $n < 5$). We conclude $\sqrt{n}$ is indeed the best that we can do to complete the proof of minimality for $n < 5$.
For the unit-hypercube, this becomes $\sqrt{n}/2$.
Question: Why $n < 5$?
Answer: Unfortunately, we don't know. However, consider $n = 5$. If we put each of our points in the configuration we described, then there exists points at $(1, 1, 1, 1, 1)$ and $(1, 1, 1, 1, -1)$.
The euclidean distance between these two points is 2, which is lower than the $\sqrt{5}$ bound offered by our theorem.
For $n \geq 5$, we have that there is always two points on the same face of a square in $\mathbb{R}^n$ that are at most 2 apart.
However, we could not figure out how to scale this with dimension, and it is clearly not the exact bound we proved for $n < 5$.

Let's Go Further

We have described we a possible limiting case configuration of points, where the minimal distance between any 2 point was 2. Finding the exact answer is not simple.(as described above with the $n = 5$ example.

We leave proving this geometrically as an exercise for Matthew.
Question: When considering a number of points given as $x^n + 1$ distributed within a hypercube H, is it "often" the case that the pigeon hole solution is a strict equality, and not just an upper bound (like it was for $2^n + 1$ points for $n < 5$)?
No. Consider the question from the second midterm (rip), that states given $10 = 3^2 + 1$ points in the unit square, we can find 2 a distance of $\sqrt{2}/3$ apart.
To solve, we divide each side of the square into thirds and observe (by pigeonhole) that 2 points are in 1 of the 9 sections. As it happens, this answer is NOT as good as we can do.

Circle Packing

An equivalent problem to the minimal distance we may find 2 points within is to pack 10 circles with their centers in the interior of the unit square with maximal radius, and the distance between the 2 points is analogous to twice the radius of the circles packed. It is commonly stated as packing n unit circles into a square of minimal size.

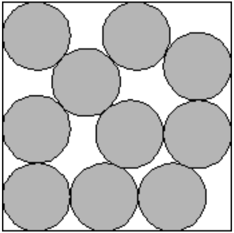

- Has a solution that is less than $\frac{\sqrt{2}}{3}$, but not so easy to find
- Corresponds to the roots of a polynomial that is at least of degree 18.

Circle Packing

From OEIS:

N	d_min
2	$\sqrt{2}$
3	$\sqrt{6}-\sqrt{2}$
4	1
5	$\frac{\sqrt{2}}{2}$
6	$\frac{\sqrt{13}}{6}$
7	$4-2\sqrt{3}$
8	$\sqrt{2-\sqrt{3}}$
9	$\frac{1}{2}$
10	min poly of deg 18

"The smallest square ten unit circles will fit into has side length s $= 2 + 2/\text{d} = 6.74744152...$ and the maximum radius of ten non-overlapping circles in the unit square is 1 / s $= 0.14820432...$"

Circle packing's formulations illustrated

Number of circles (n)	Square size (side length (L))	d_n [1]	Number density (n/L^2)	Figure
1	2	∞	0.25	
2	$2+\sqrt{2}$ ≈ 3.414...	$\sqrt{2}$ ≈ 1.414...	0.172...	
3	$2+\frac{\sqrt{2}}{2}+\frac{\sqrt{6}}{2}$ ≈ 3.931...	$\sqrt{6}-\sqrt{2}$ ≈ 1.035...	0.194...	
4	4	1	0.25	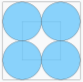
5	$2+2\sqrt{2}$ ≈ 4.828...	$\frac{\sqrt{2}}{2}$ ≈ 0.707...	0.215...	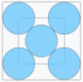

credit: wikipedia - circle packing in a square

Circle packing illustrated contd.

6	$2+\frac{12}{\sqrt{13}}$ ≈ 5.328...	$\frac{\sqrt{13}}{6}$ ≈ 0.601...	0.211...	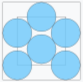
7	$4+\sqrt{3}$ ≈ 5.732...	$4-2\sqrt{3}$ ≈ 0.536...	0.213...	
8	$2+\sqrt{2}+\sqrt{6}$ ≈ 5.863...	$\frac{\sqrt{6}-\sqrt{2}}{2}$ ≈ 0.518...	0.233...	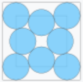
9	6	0.5	0.25	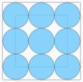
10	6.747...	0.421... OEIS: A281065	0.220...	

credit: wikipedi a- circle packing in a square

Circle Packing

Note that our problem is exactly a formulation of the circle packing problem, generalized to n dimensions.
Therefore, circle packing solutions are an *upper bound* to our problem, since many circle packing configurations are not proven to be optimal, merely the best that we have found.
Furthermore, consider points placed on the boundary of H. The corresponding spheres clearly are not completely contained in H. Now, consider the larger hypercube H' that *does* completely contain all the spheres in H. Most likely, there is a way to pack more circles in H' than circles "packed" in H to solve our original problem.

Revisiting $3^n + 1$ Points

Next, let's revisit $3^n + 1$ points.
Applying our pigeonhole argument will only give us an upper bound in the case of $3^n + 1$ points within $H \in \mathbb{R}^n$.
We do so anyway:

Proof

Let $H \in \mathbb{R}^n$ be a hypercube of side length 3. Then in each 3rd division we have a maximal distance of $\sqrt{n}$ by the same argument as before. Without loss of generality, look at the diagonal along the first compartment. Then $\sqrt{n}/3$ is our upper bound for the unit hypercube given $3^n + 1$ points.

For an exact answer, this becomes a "hypersphere packing" question, and unfortunately the answer is not so clear. In fact, it is a special case of Hilbert's 18th problem (or equivalently, the Kepler conjecture) which was formally proven in 2013 using computers [Wikipedia].

A Generalization on the Upper Bound

Theorem

Given $x^n + 1$ points in the hypercube, we can always find 2 a distance of $\sqrt{n}/x$ apart.

proof: Let $H \in \mathbb{R}^n$ be a hypercube with side lengths x. We may divide each edge into x segments of length 1, drawing "lines" through the entire cube that are orthogonal to the origin edge.
Repeat this along every edge, and we have partitioned the cube into x^n different compartments.

Formal Compartments

A slightly "formal" proof is the following combinatorial argument:

- Center the cube with a corner at the origin, and consider the edges leaving the origin. (For a square, 2 such edges a exist, for a cube, 3, . . . , n for H)
- These edges are a basis in $\mathbb{R}^n$
- **Claim**: These compartments form a partition of H, since 2 different compartments sharing an edge of the cube must have a divider in between them by construction, or be the same compartment
- Thus each chopped orthogonal edge contributes x pieces towards the number of compartments, and the exact location of each compartment may be specified by which section of each partitioned edge it corresponds to
- In $\mathbb{R}^n$, we have n such edges, and x choices for each edge. Therefore, the number of compartments is x^n.

A Generalization on the Upper Bound Contd.

The previous argument confirms that indeed we have x^n compartments, and thus 2 of the $x^n + 1$ points are in the same compartment.
Since we chose each compartments' side length to be 1, by making the hypercube have side length x, we know the maximal distance in 1 compartment is $\sqrt{n}$ by our previous results.
Then by scaling back to a unit hypercube, the maximal distance between two points d is upper bounded by the pigeonhole argument s.t. $d \leq \frac{\sqrt{n}}{x}$.

A glimpse into why the pigeonhole argument fails in general

- First, consider $n = 2$. We claim 2 of the $x^2 + 1$ points on the square must be closer than the pigeonhole argument asserts for $x > 2$.
- We reason that the vertices of the x^2 compartments are the optimal places for points to be assigned, since it is "clearly" best to put points at the corners, where the "packing sphere" around them takes up the least volume in the square. (We believe this is equivalent to Kepler conjecture statement)

We want to pack them as close as possible to avoid wasting space, and these points are precisely the lattice points separating compartments. (the natural packing in 3d asserted by the Kepler conjecture but in 2d) then we have $(x+1)^2$ vertices, and $x^2 + 1$ points. We color the vertices with a chess board coloring, since adjacent vertices have a distance of 1 between them, and are thus too close to be selected.

Failure of pigeonhole contd

- We may only then select half of the vertices (or half $+ 1$) in our selection, since half are one color. But then since $x^2 + 1 > ((x+1)^2)/2 + 1$ for $x > 2$ where we have rounded up to integer values if necessary, we conclude that we run out vertices (the RHS) for our points (the LHS).
- in general this argument gets more and more complicated, and it is not obvious whether the number of available vertices grows faster or slower than the number of vertices eliminated each time a point is placed, since points on 1 face of the cube stay the same distance while the distance of interest, $\sqrt{n}$ keeps growing.

Independent Sampling from Hypercubes

1. Consider randomly sampling points on the interval $(0, 1)$.
2. Let's say we don't want to sample points within $\epsilon > 0$ of the edges 0 or 1.
3. Then we sample from the "interior" with probability $\frac{1-2\epsilon}{1}$.

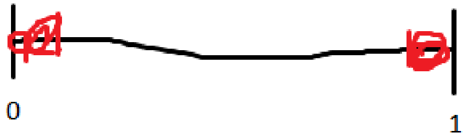

Independent Sampling from Hypercubes

Definition

k-**hypercube**: Let a k-hypercube $\in \mathbb{R}^n$ be the hypercube in $\mathbb{R}^n$ with (each) side of length k.

For each single dimension, we can sample from the "interior" along that dimension with probability $\frac{k-2\epsilon}{k}$.
$\epsilon \in (0, k)$ determines the size of our interior relative to the side of an edge in the cube.
Then we have the following result:

$$\lim_{n\to\infty} \left(\frac{k-2\epsilon}{k}\right)^n = 0.$$

This is true since $\left(\frac{k-2\epsilon}{k}\right) < 1$, and any number of magnitude less than one to the n'th power will go to 0 in the power limit.

Implications of Independent Sampling

With probability 1, every sampled point in an n-dimensional will be on the/a boundary in the limit as n goes to infinity.

Theorem

For a finite dimensional hypercube $H \in \mathbb{R}^n$, given n independently sampled points, there will be a pair that will be $d = 1$ from each other (in expectation).

proof: Given n independently sampled points in a finite dimensional hypercube, with probability 1, there will be at least two points on boundaries (in expectation).
Select any two boundary points. Given that they are not in the same n-tant, the distance between them will be lower bounded by k. e.g. (1,1,1,1,1) and (1,1,1,1,-1), $k = 2$.

Implications of Independent Sampling

- Given a random sampling of points, in expectation the min distance between two will be lower bounded by k.
- Probability of two points being in the same n-tant is 0 in expectation, given countably infinite dimension n.

Question: How does this relate to the earlier work on point-square packing?
L_2 norm for distance between points in high dimensional space can become meaningless, as everything tends to the same distance apart.

The Curse of Dimensionality

Curse of Dimensionality

As $n \to \infty$, within a set of points in $\mathbb{R}^n$, the ratio

$$\frac{min\{D_{p_i,p_j}\}}{max\{D_{p_i,p_j}\}}$$

, where $\{D_{p_i,p_j}\}$ is the set of distances w.r.t. the L_2 norm between all pairs of points, goes to 1.

proof: Proof/motivation given in "When is Nearest Neighbor meaningful?" [Beyer et al., ICDT 99], outside scope of presentation.

- Relating back to the original problem, this implies that by solving the problem in high dimensional space, we can get exact bounds (in expectation) of the distance between *any* pair of points.
- Packing may not be well defined with our choice of norm.

Summary

1. We first formalized and generalized the point-square problem to hypercubes.
2. We proved tight bounds for $n < 5$ in the $2^n + 1$ number of points case.
3. We conjectured a constant bound of 2 for $2^n + 1, n \in \mathbb{R}^n$.
4. We showed a reduction to a recently closed problem about circle packing, where the reduction provides a loose bound to our original problem.
5. We proved a loose bound on a distribution of points of number $x^n + 1$ within $H \in \mathbb{R}^n$.
6. We provided some intuition on the failure of the pigeonhole principle in these higher dimensions.
7. Finally, we showed a small aside on sampling points from hypercubes, and it's relation to the original problem.

Potential Future Directions

1. Consider other norms, L_1, L_p, L_∞. This corresponds to different packing problems.
2. Consider general convex shapes, or restricted shapes based on SA, volume.
3. ... Tighten optimality bounds for $2^n + 1$ case. Is this an open problem?

Conclusion

Circle packing is hard.
Questions?