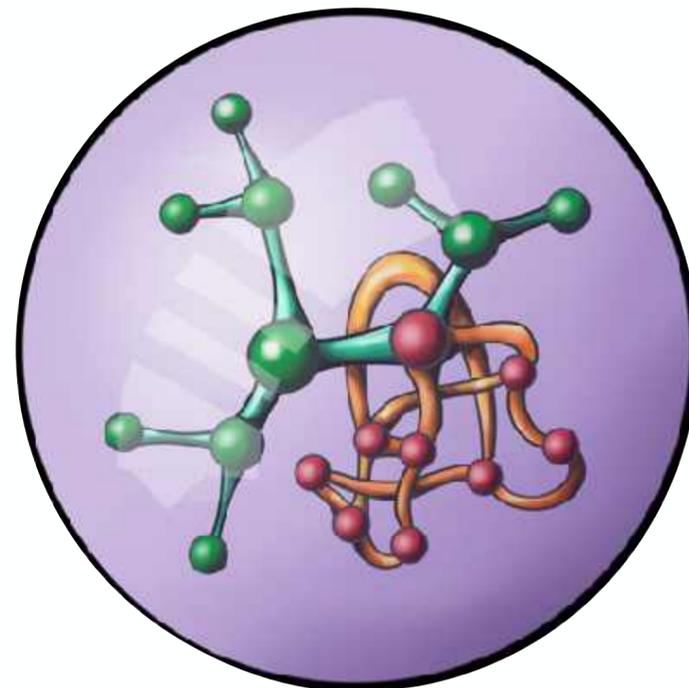
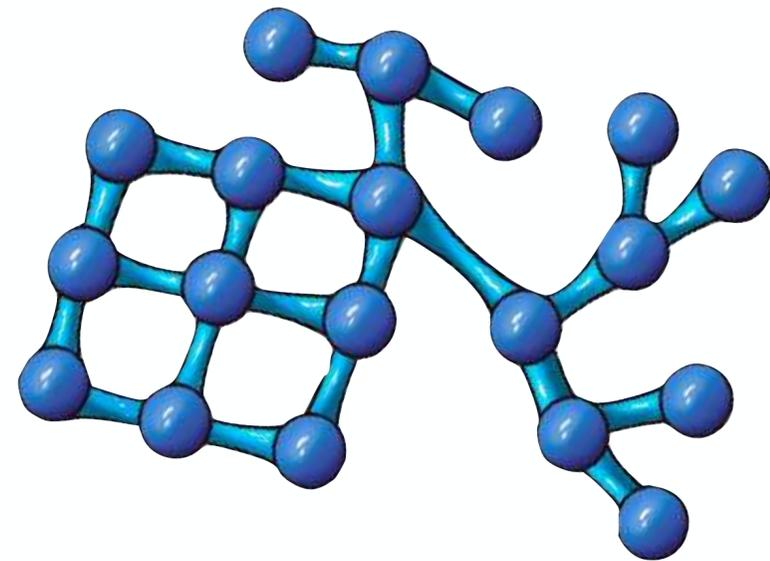
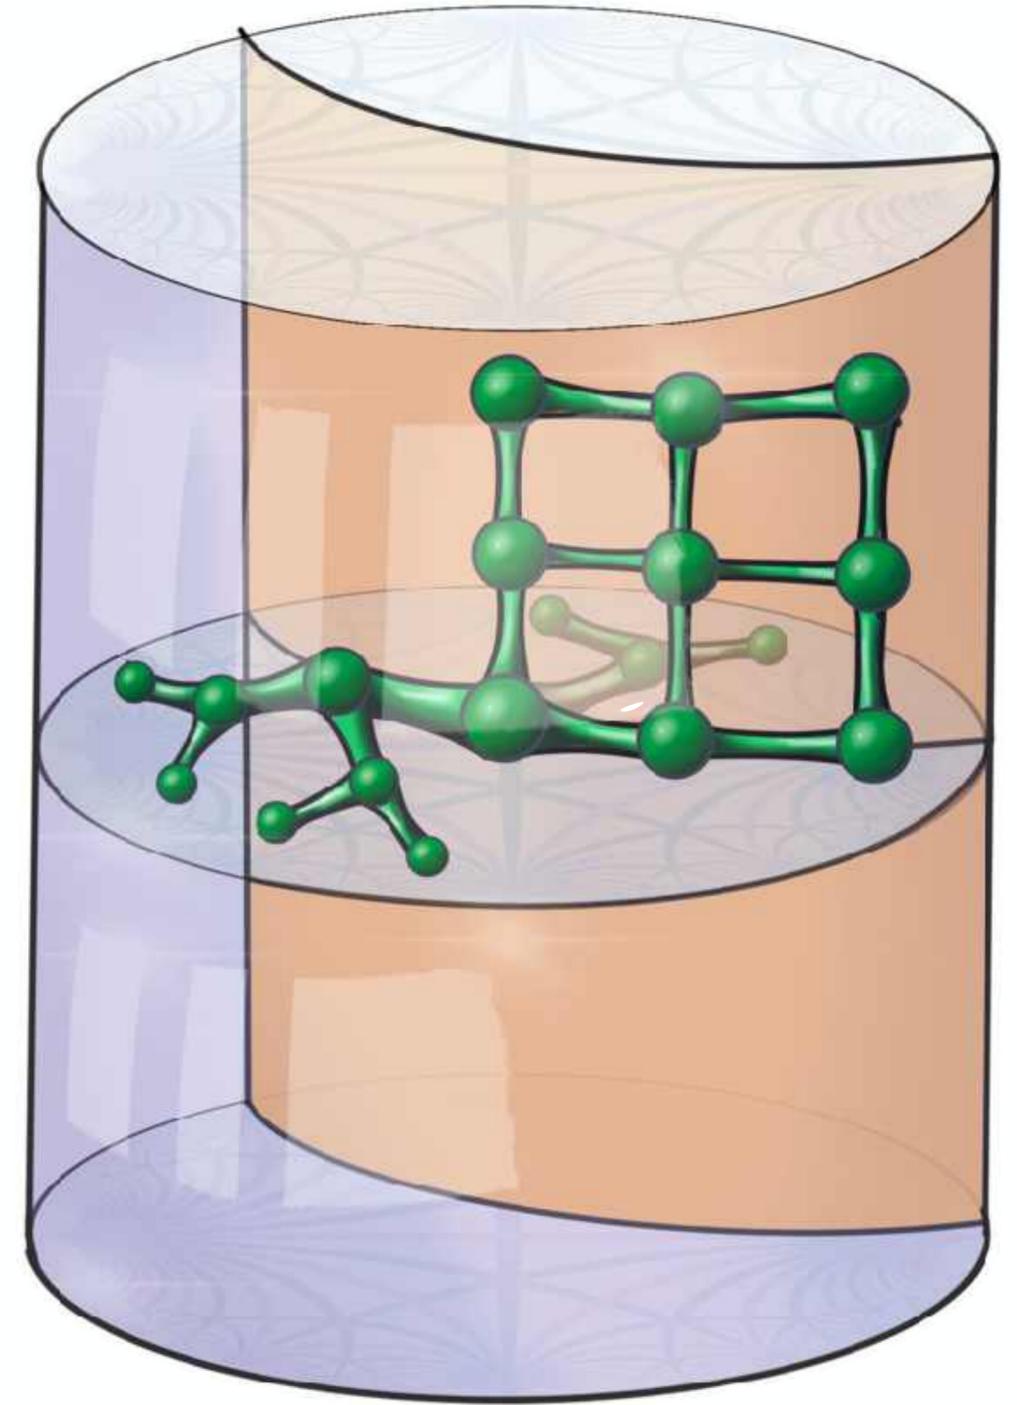
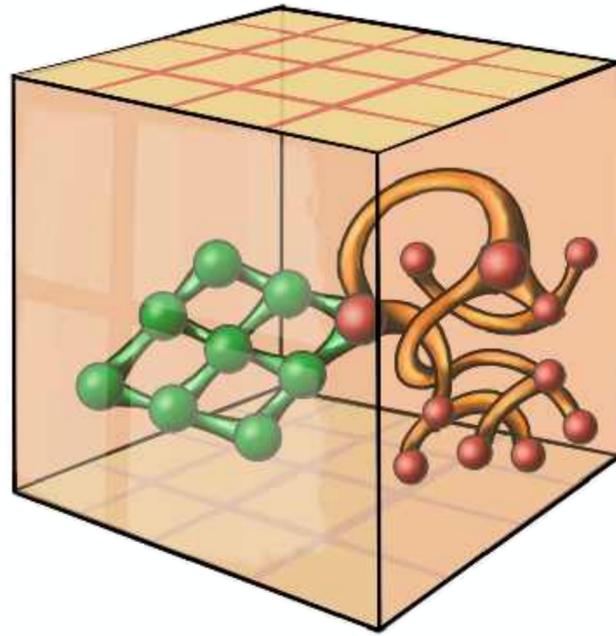


Modern Geometry &
Machine Learning



Leveraging the geometry of Riemannian symmetric spaces for graph embeddings



Federico
Lopez



Beatrice
Pozzetti



Michael
Strube



Steve
Trettel



Anna
Wienhard



Max
Reistenberg



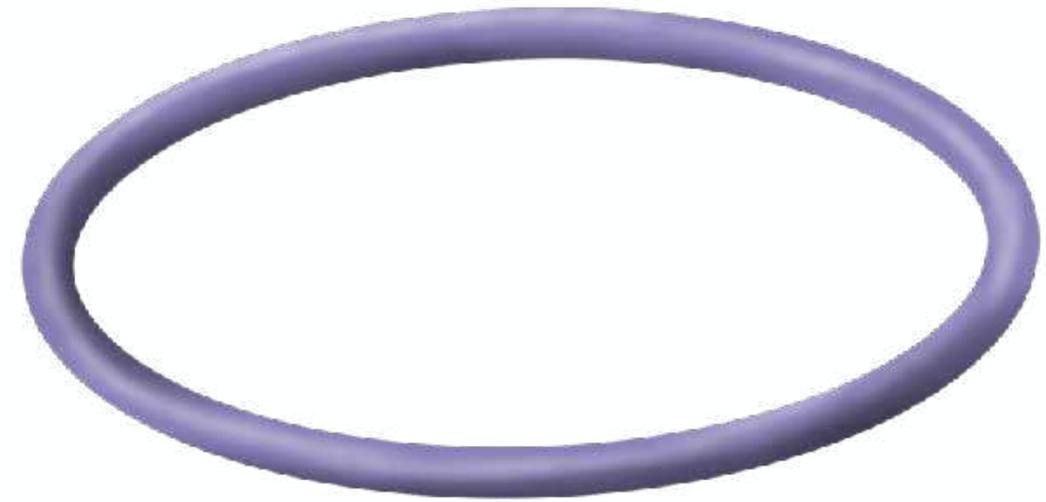
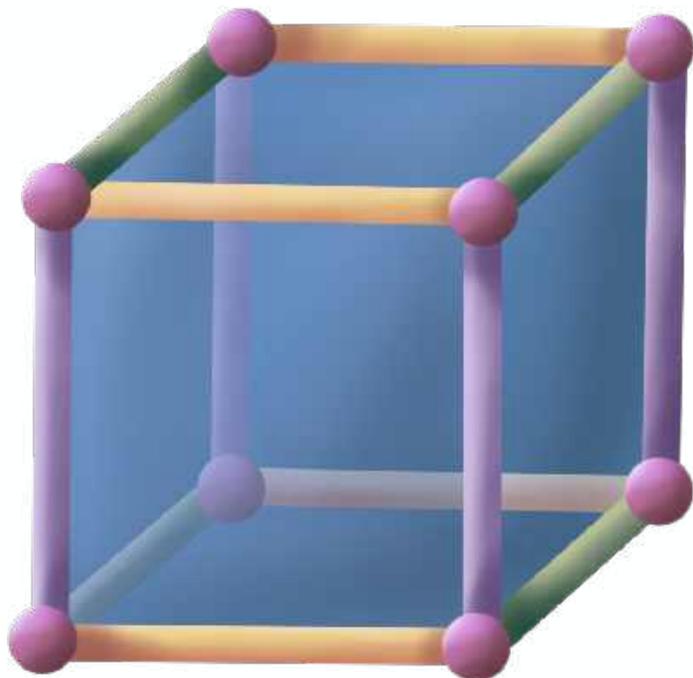
Diaaeldin
Taha



Wei
Zhao

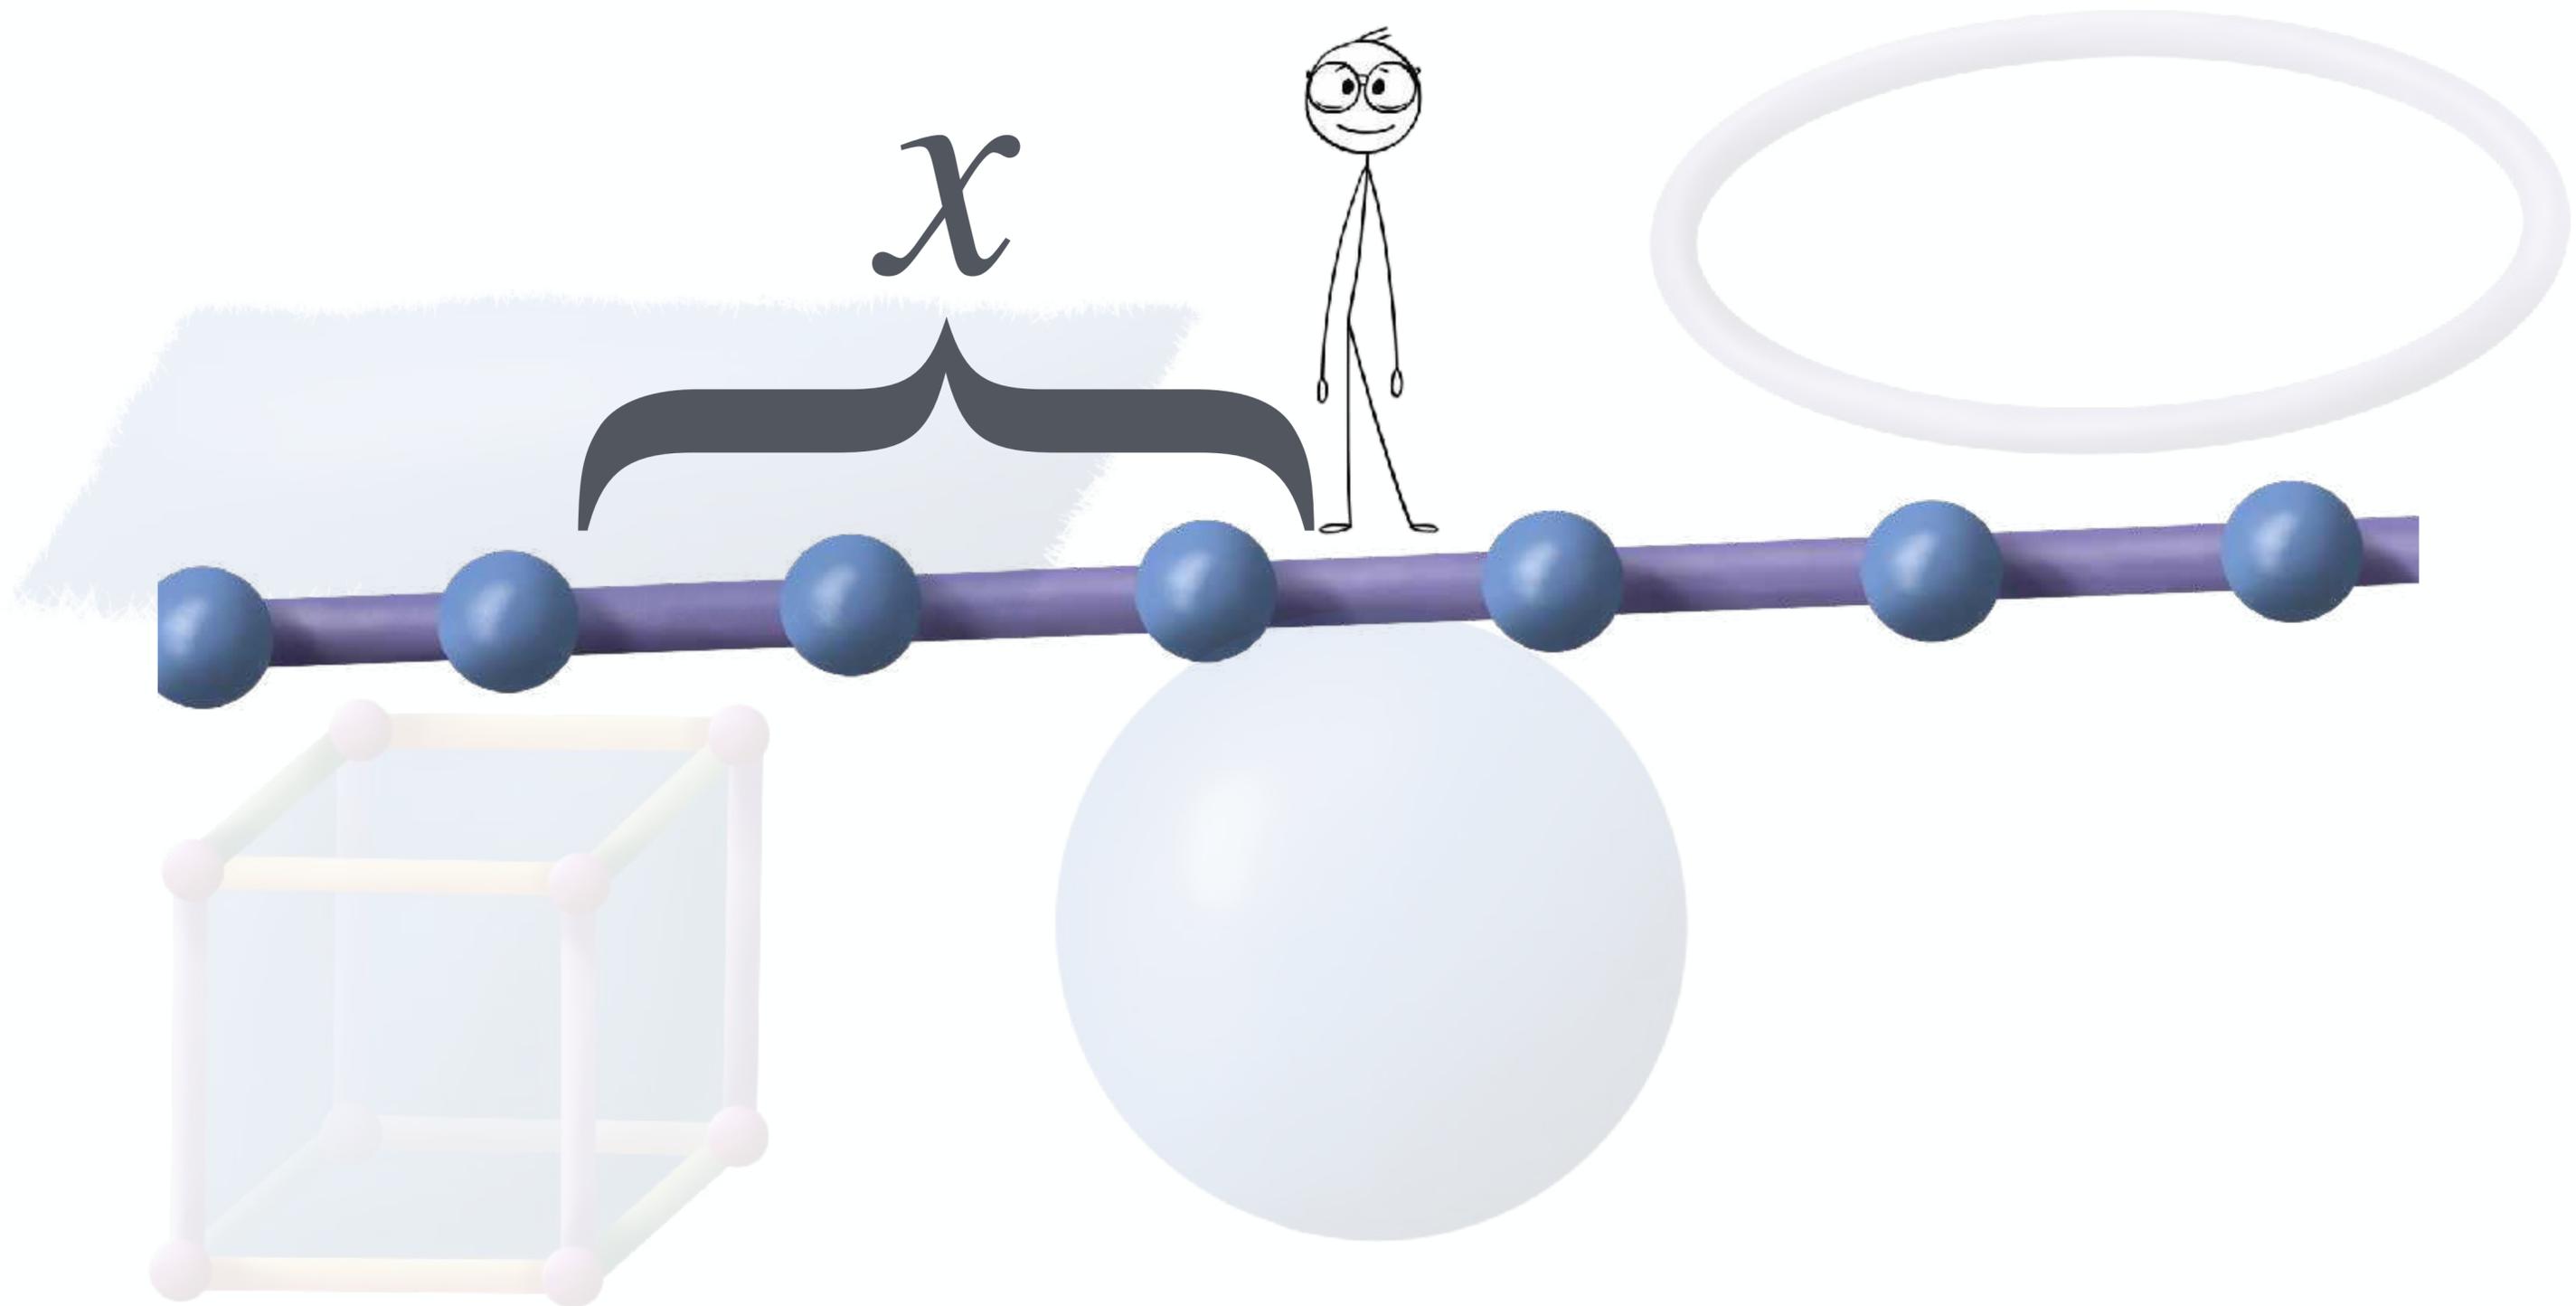
The Ideas of Modern Geometry

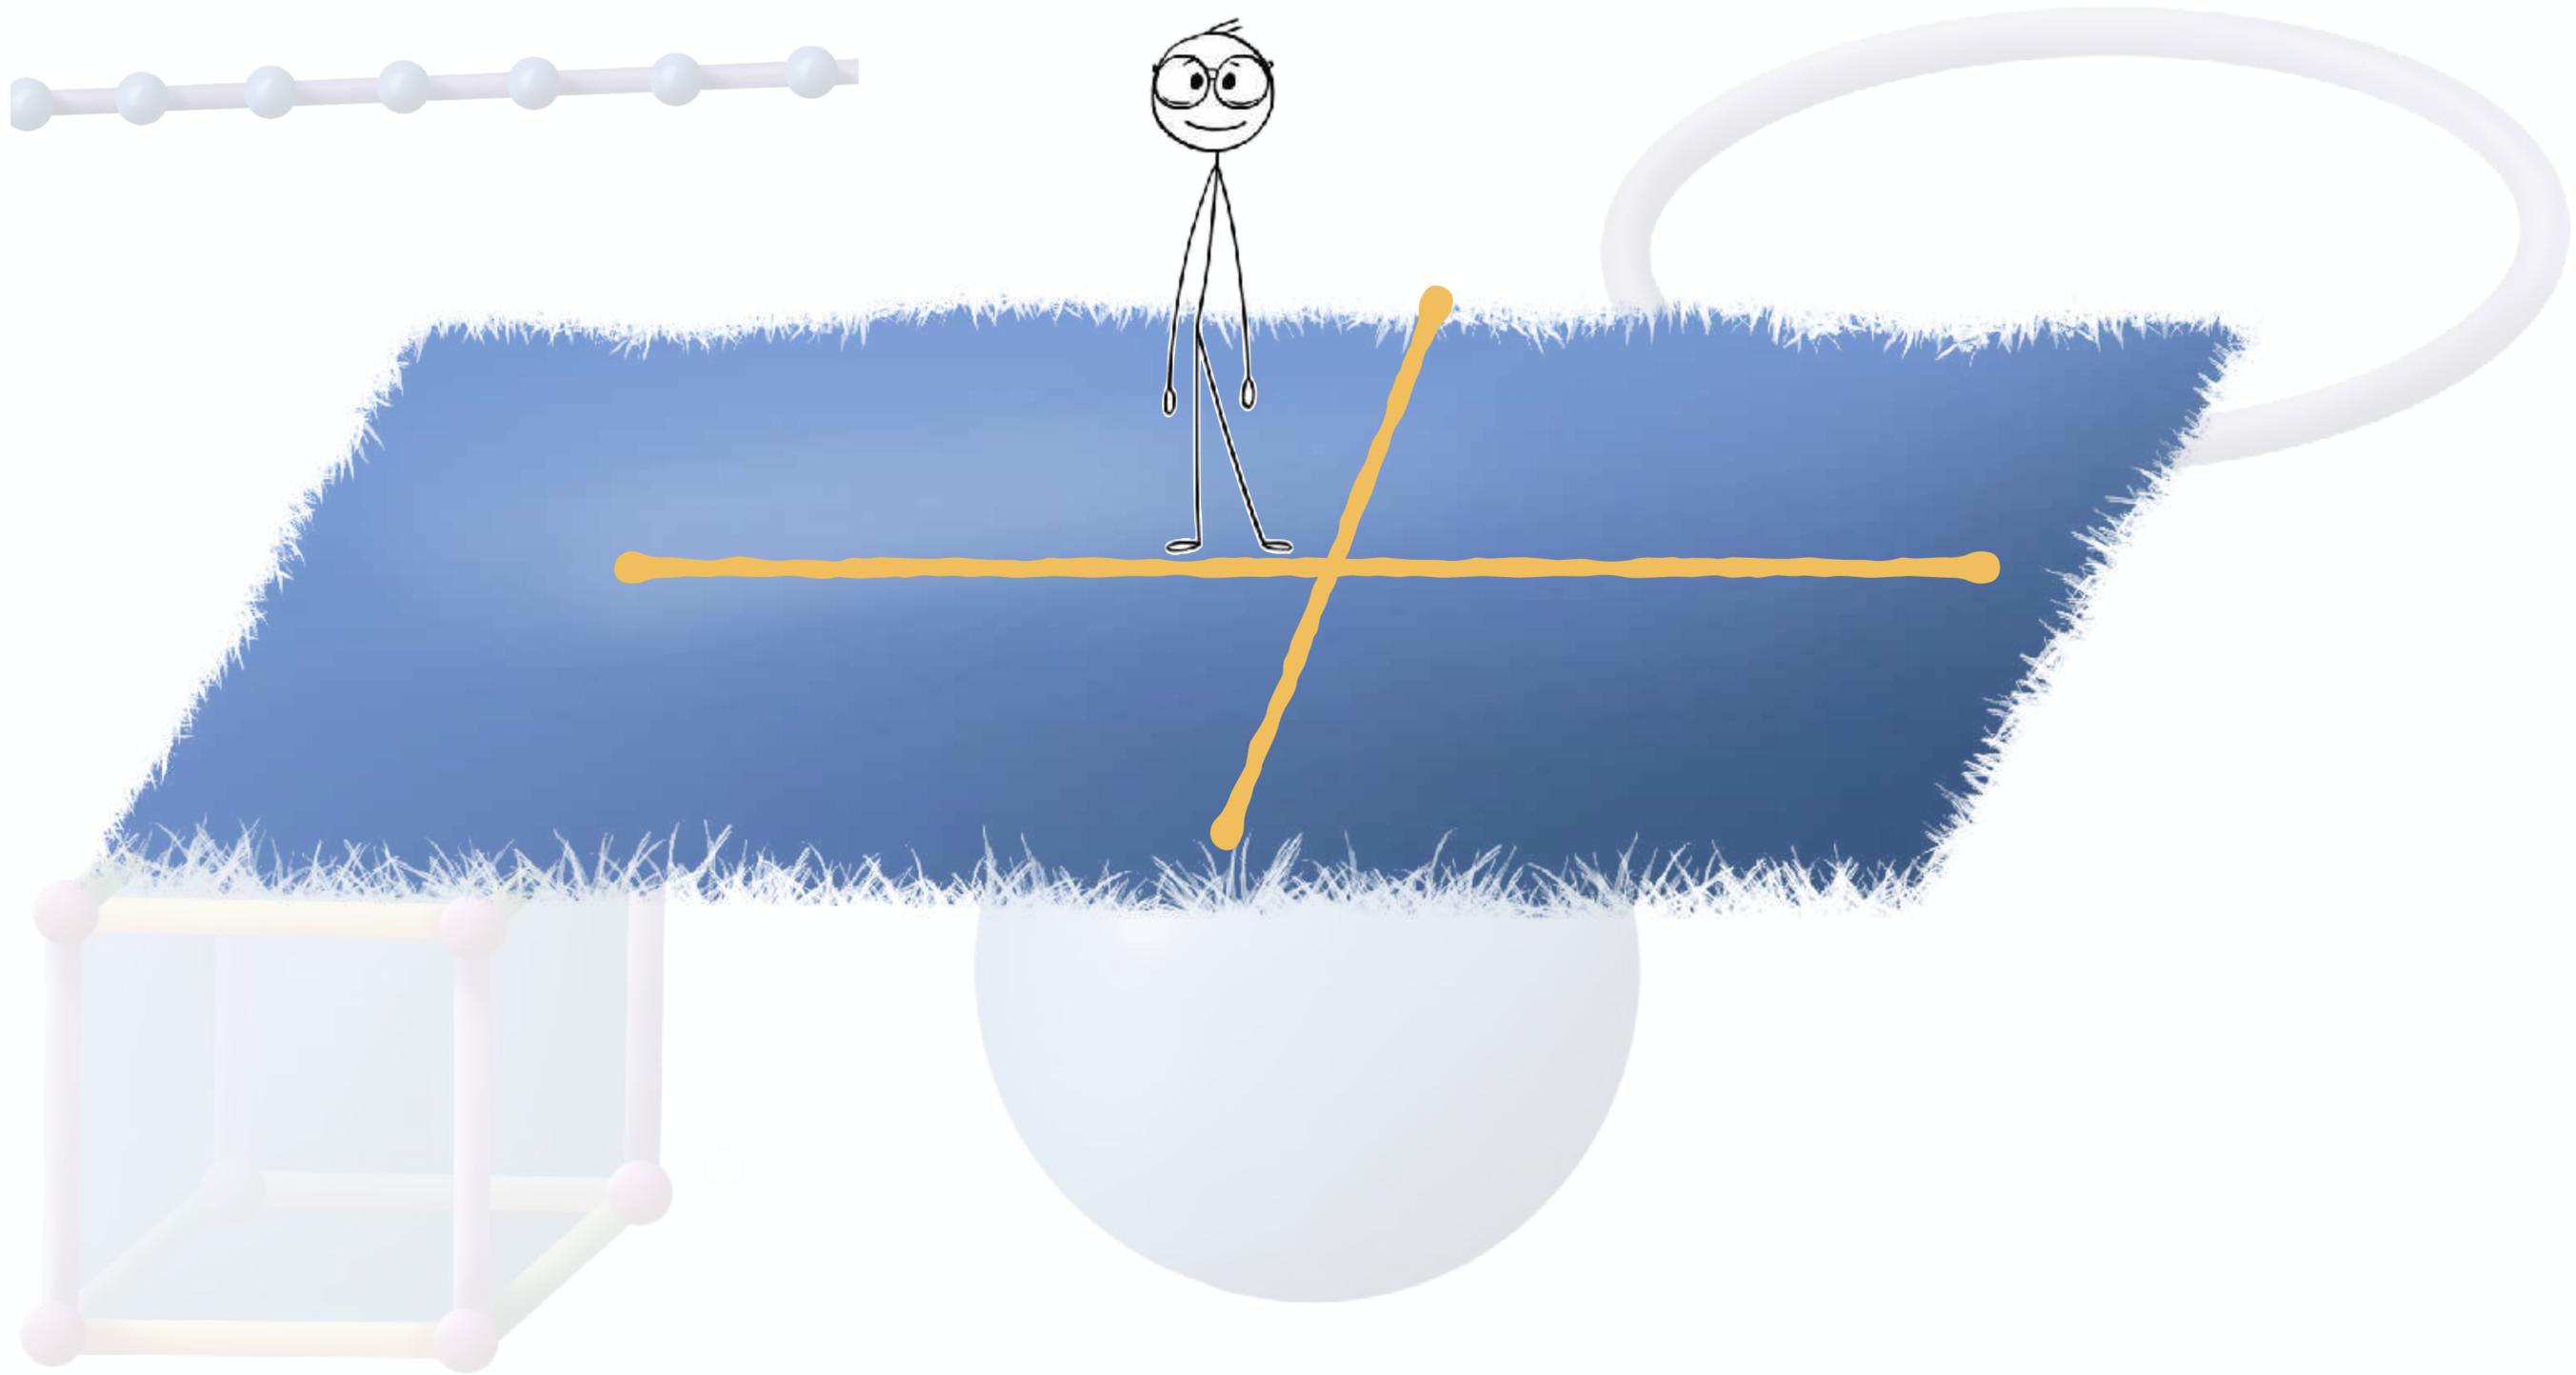
High Dimensions & Curved Spaces

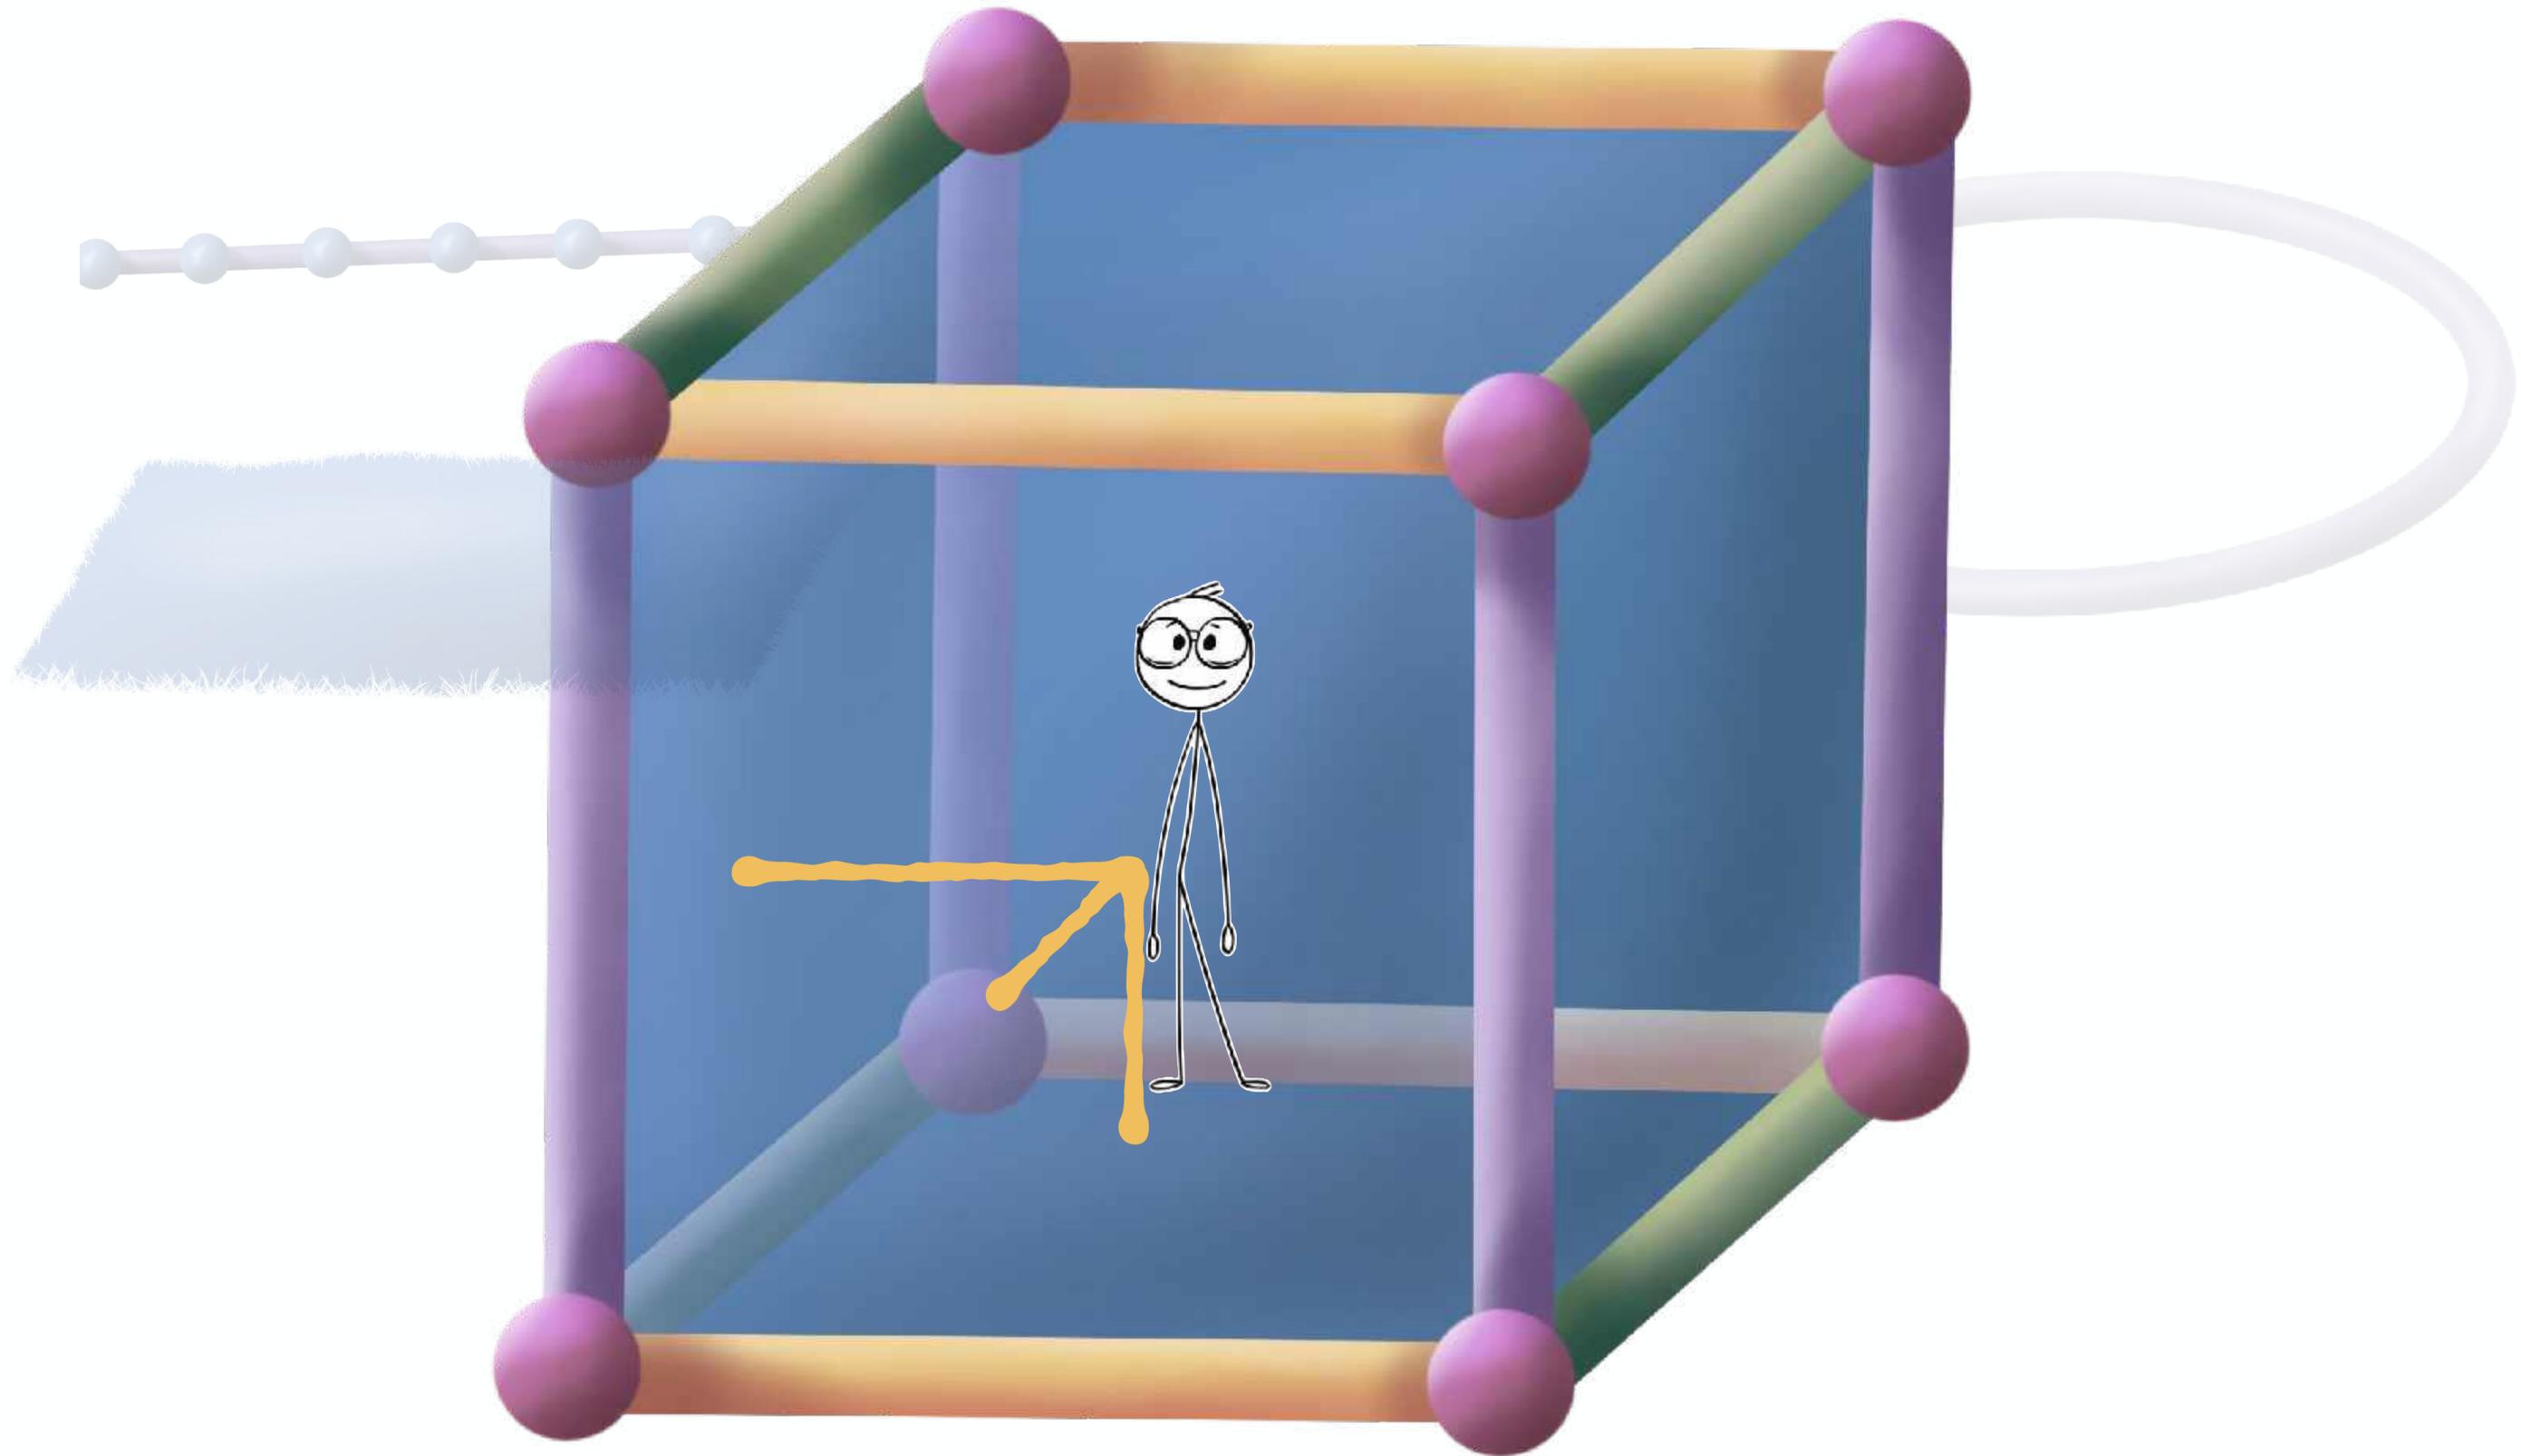


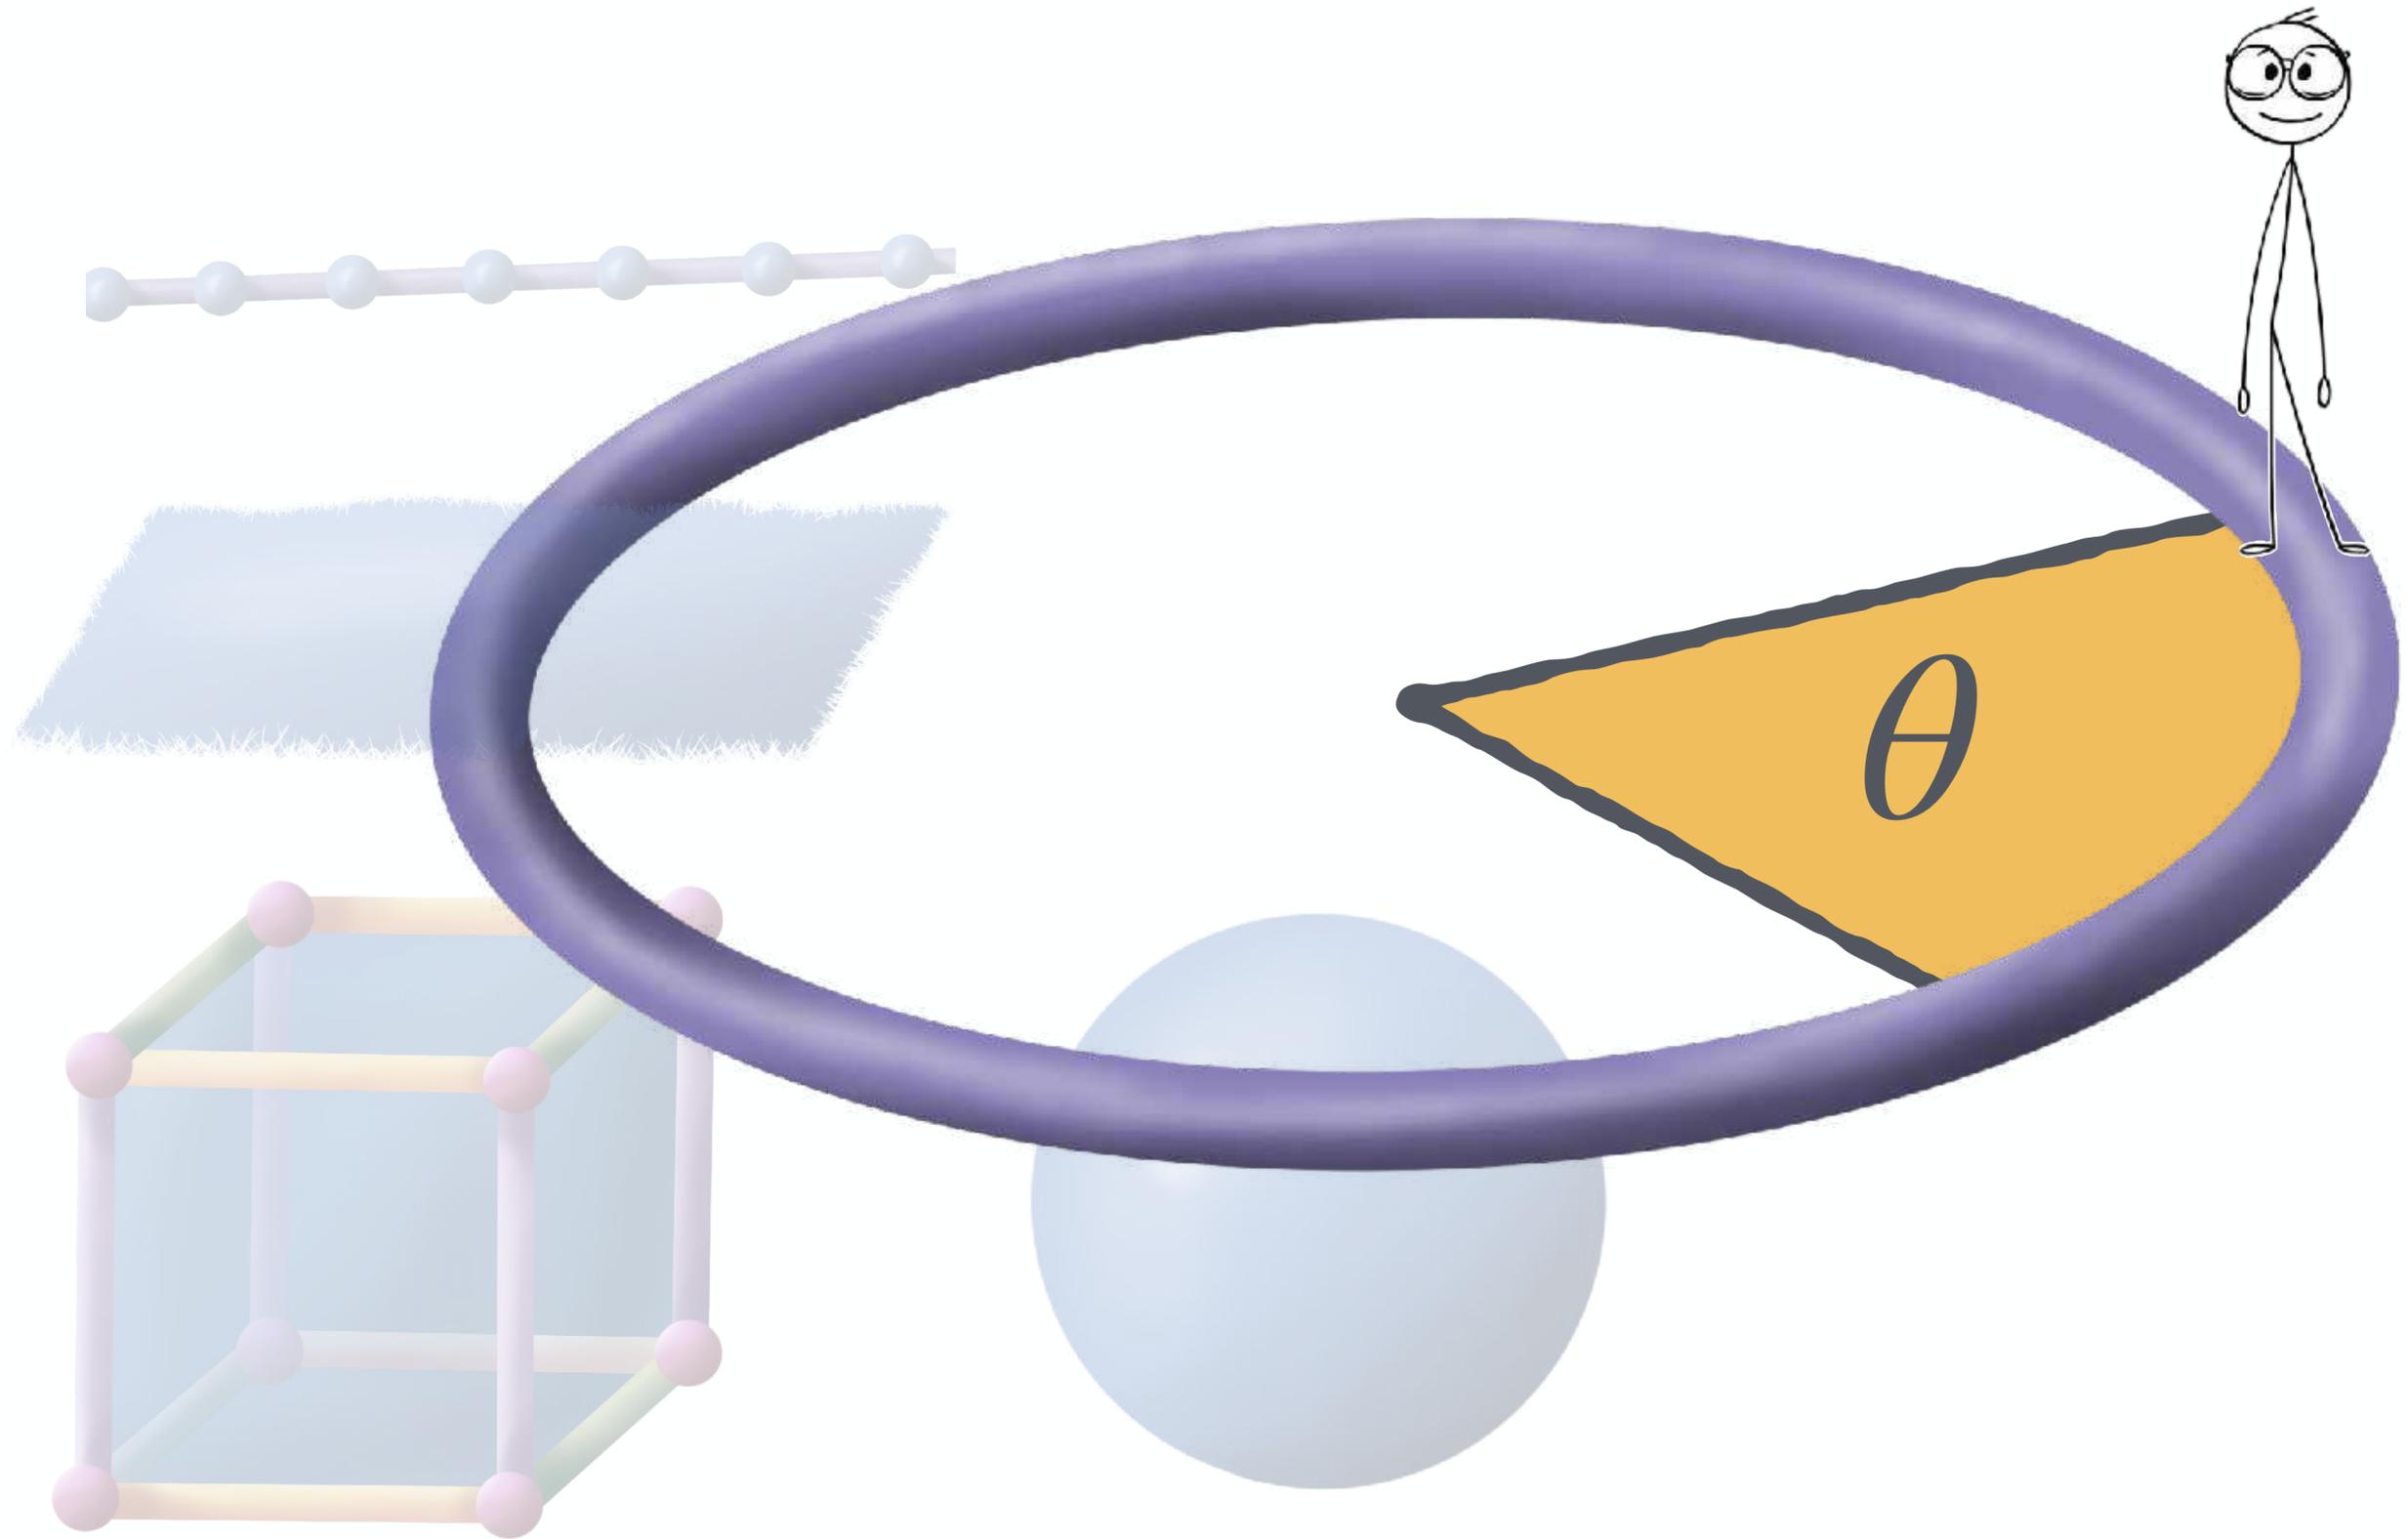
What is dimension?

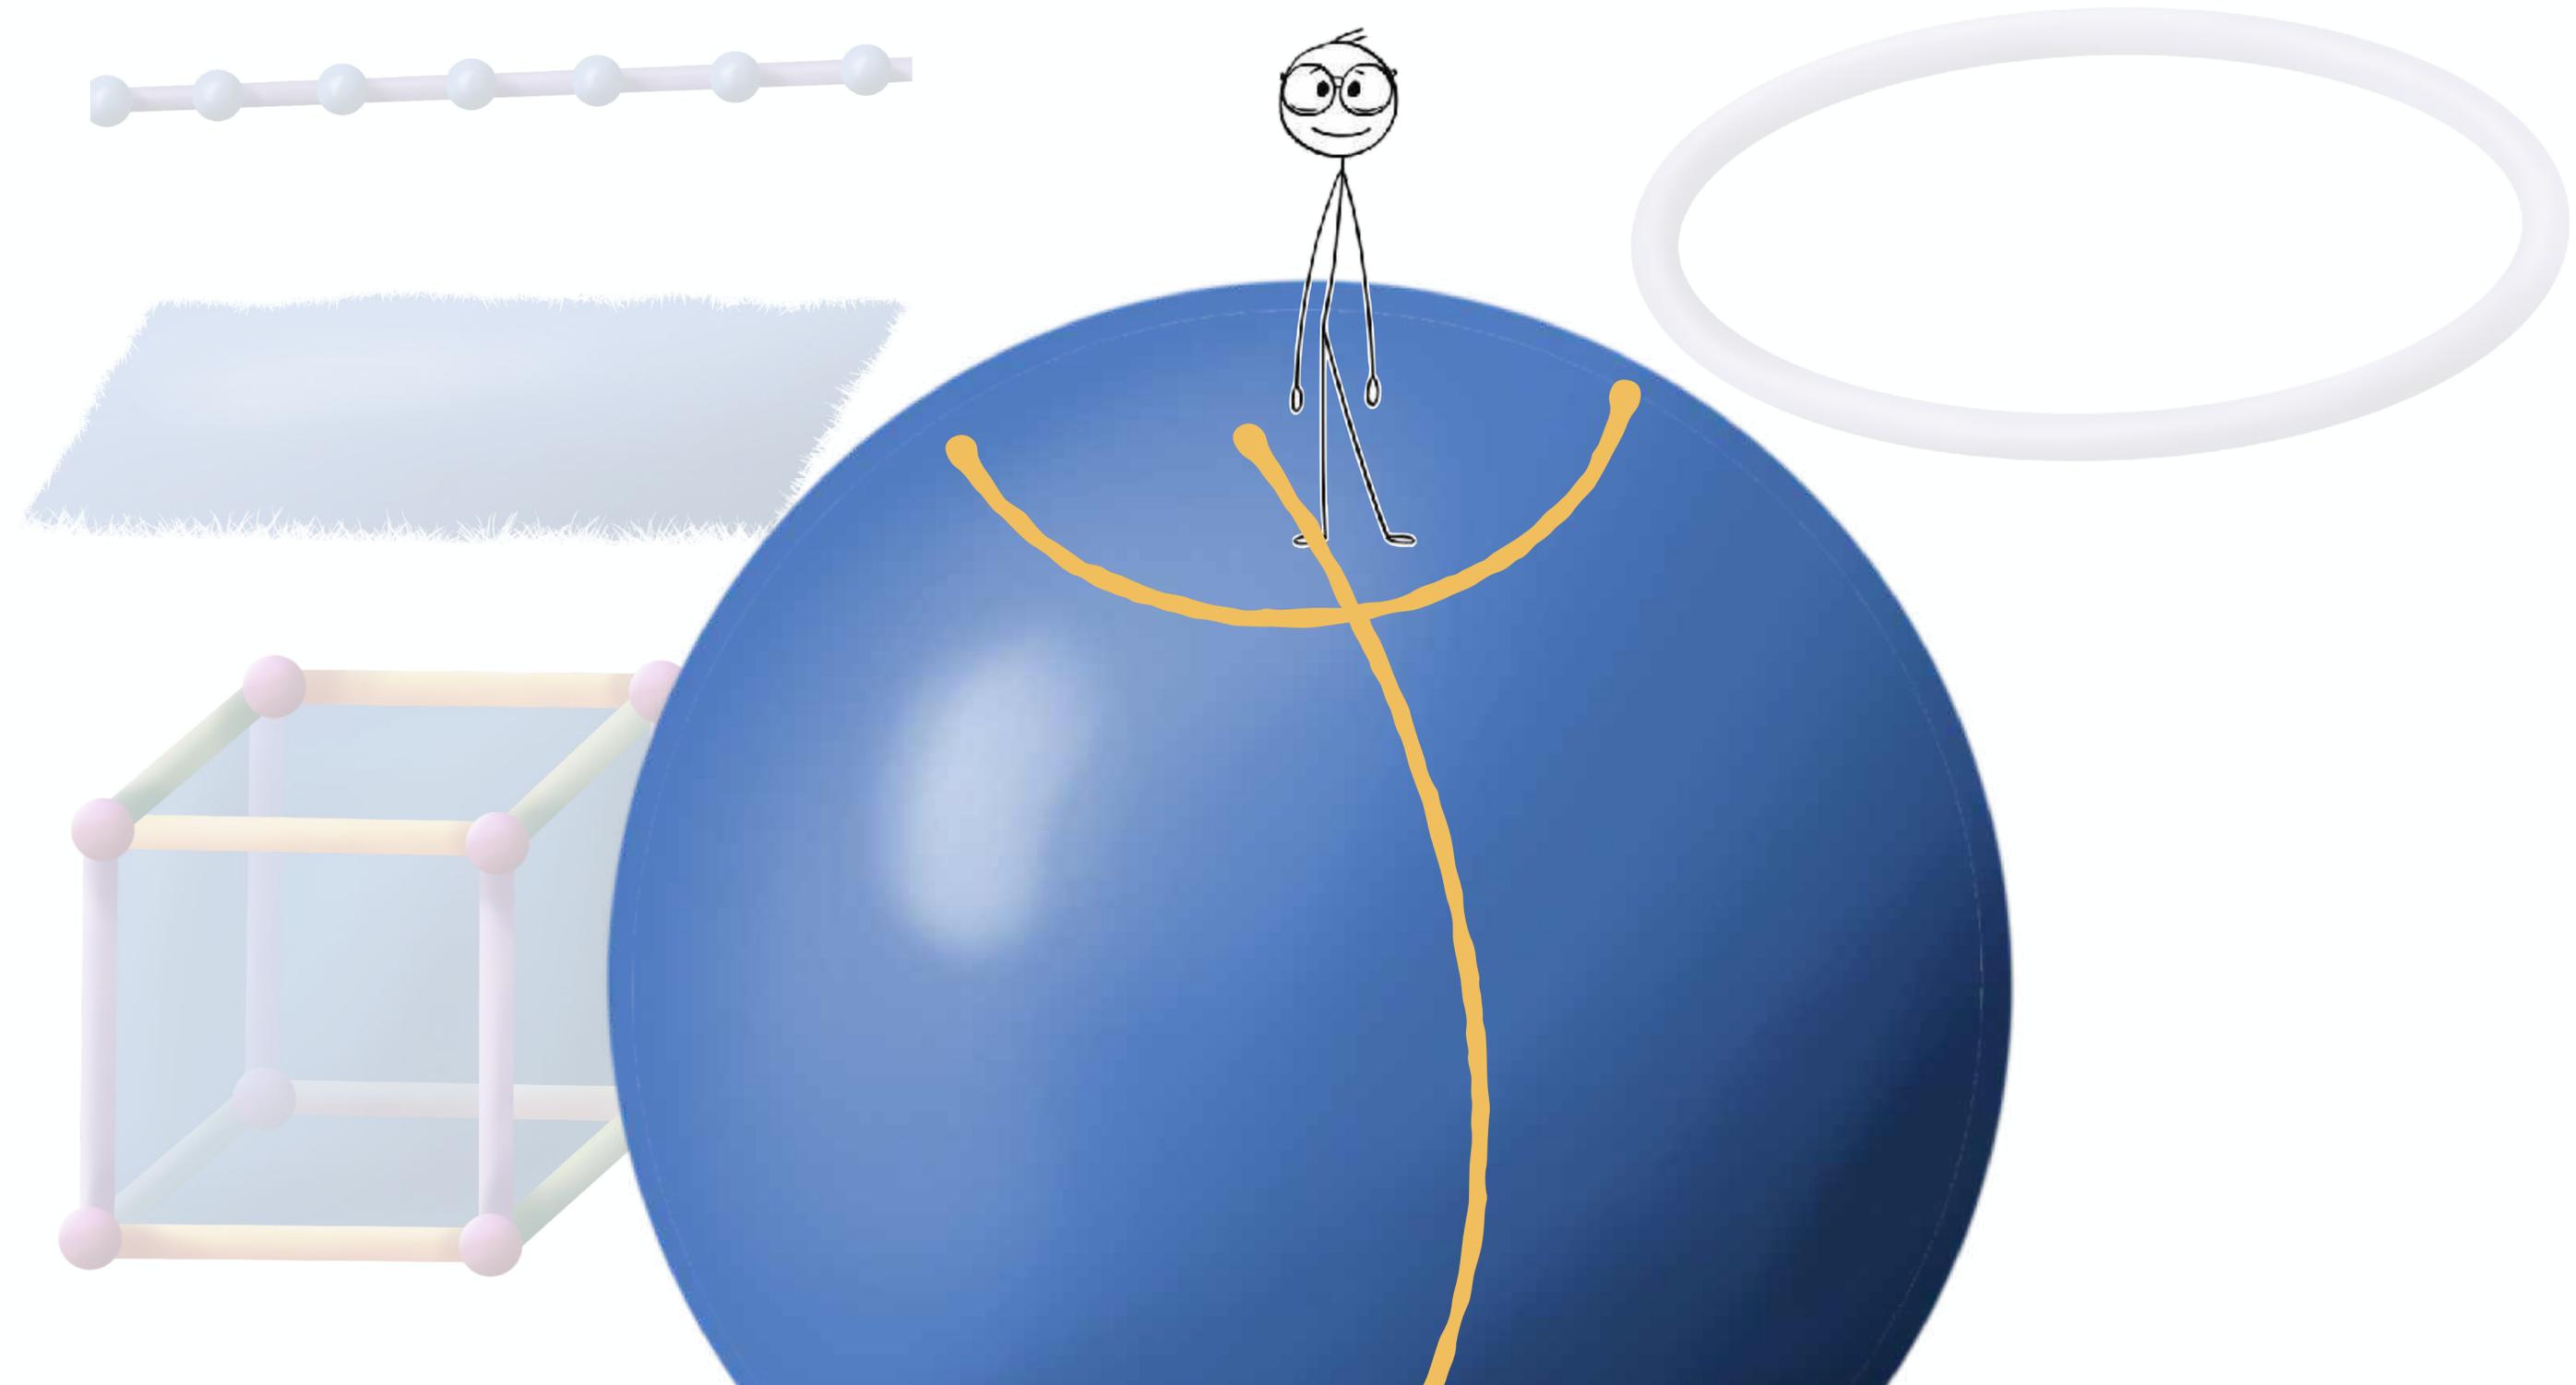
The number of pieces of information you need to completely specify a point.





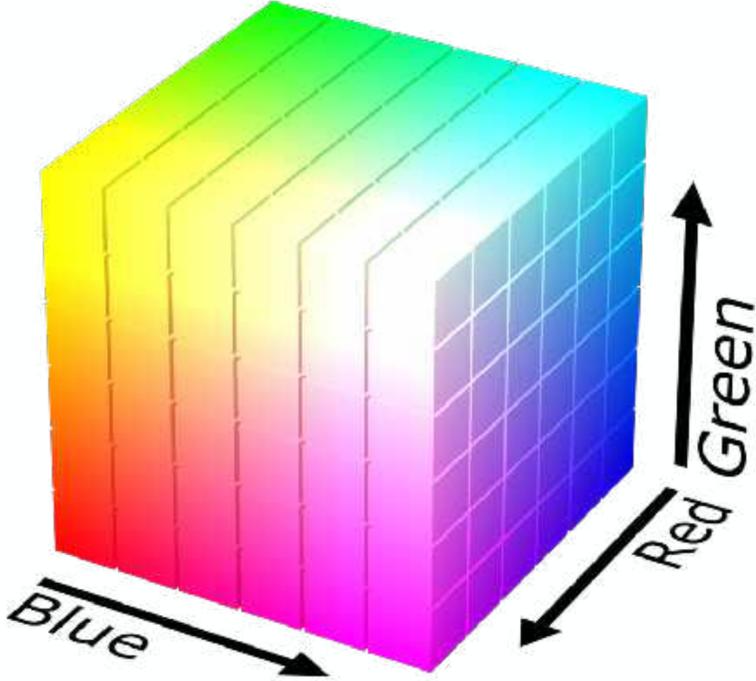
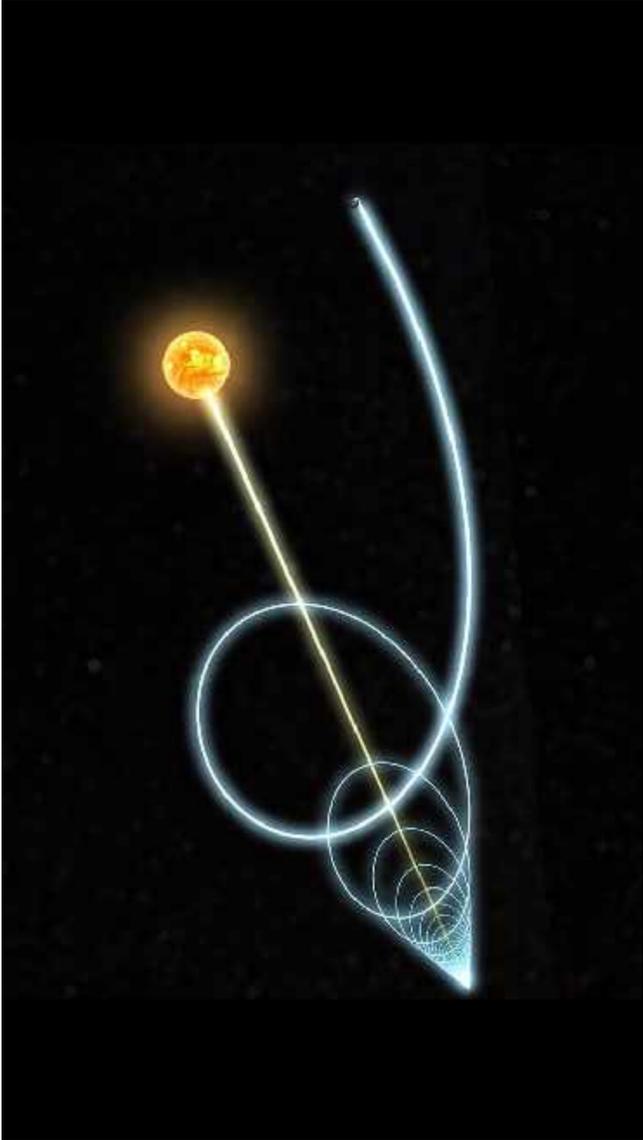






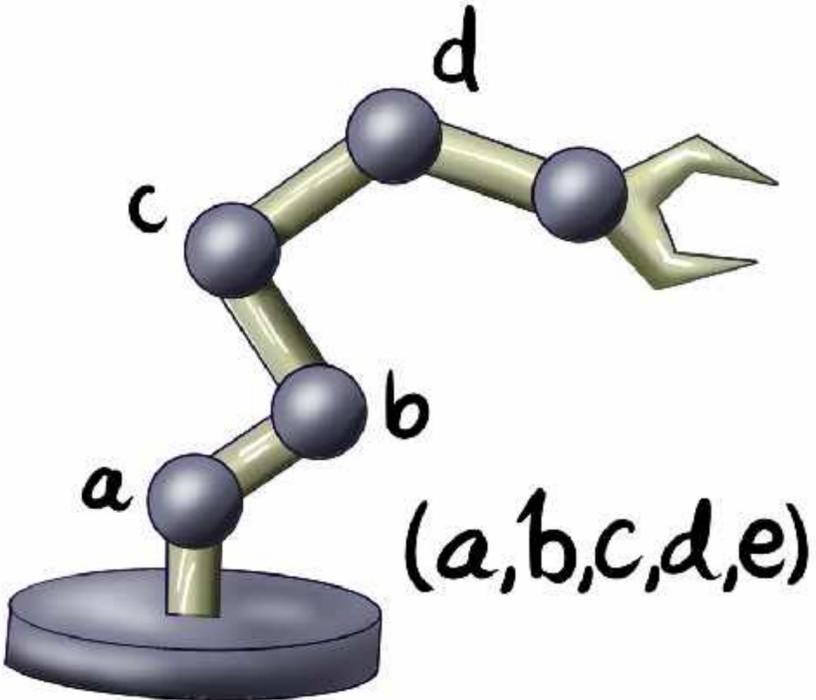
This opened our eyes to a much broader world

Spacetime is 4D

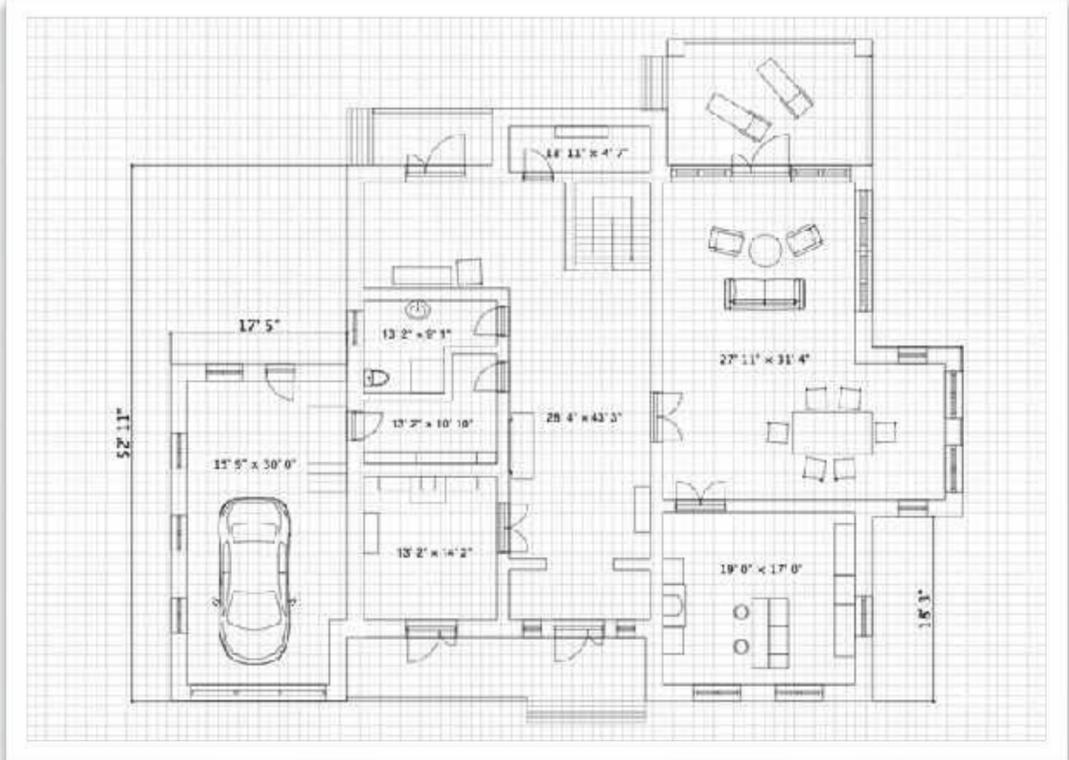


Space of colors is 3D

Configuration spaces of robots, etc



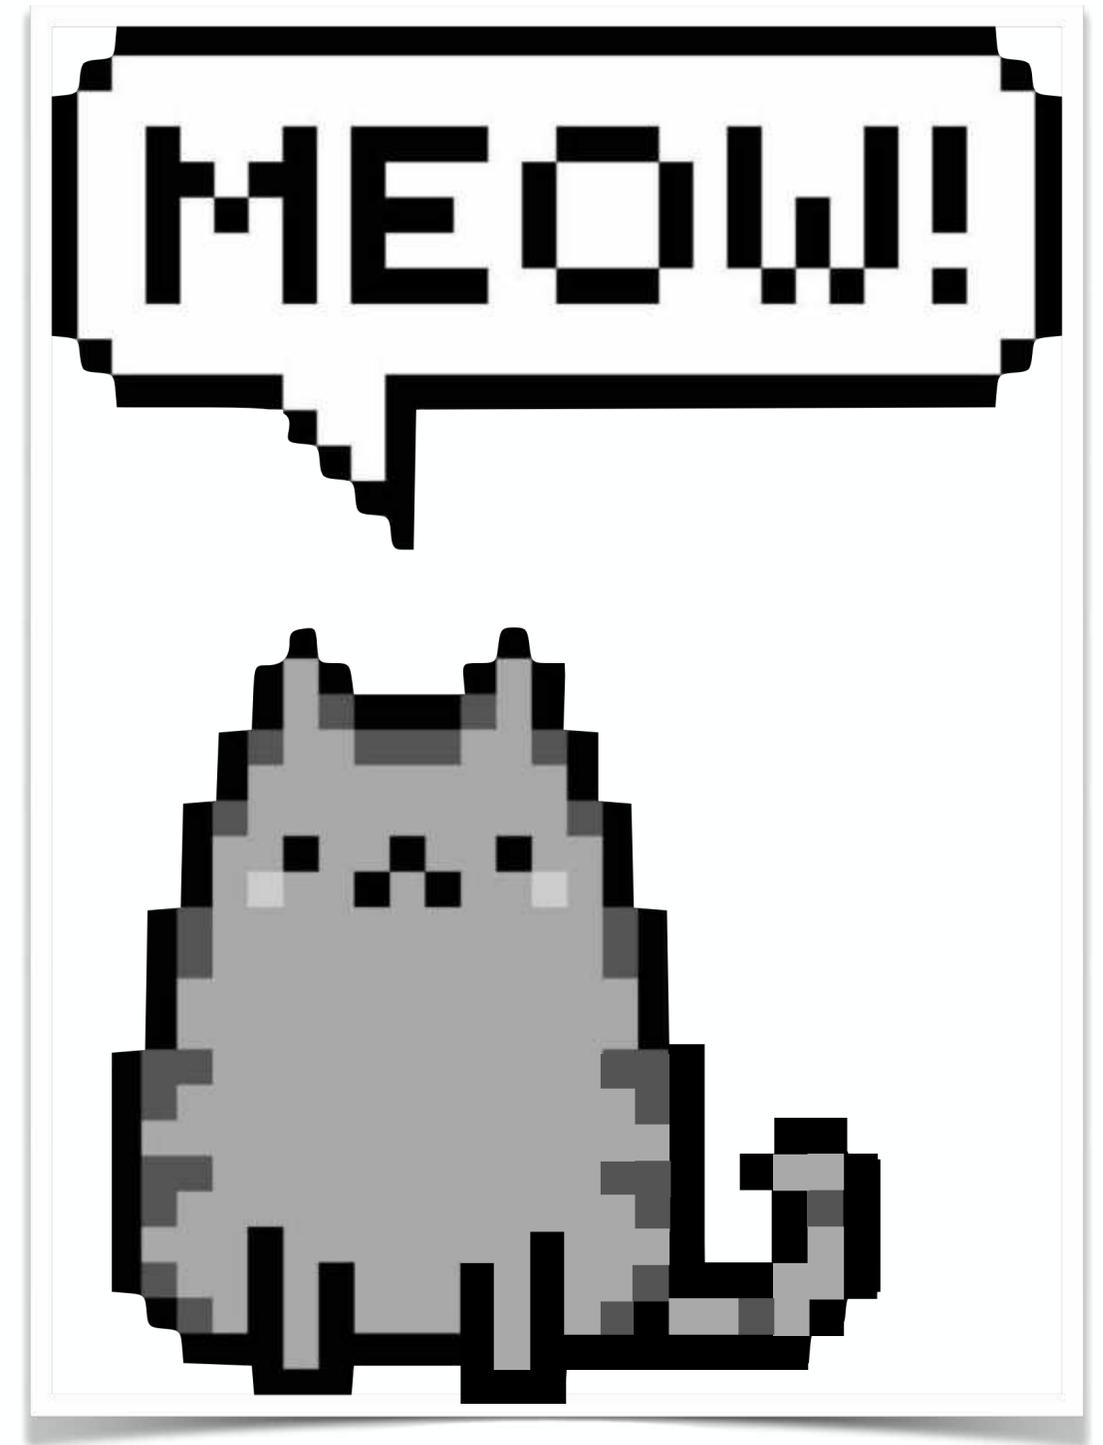
Parameter spaces of design

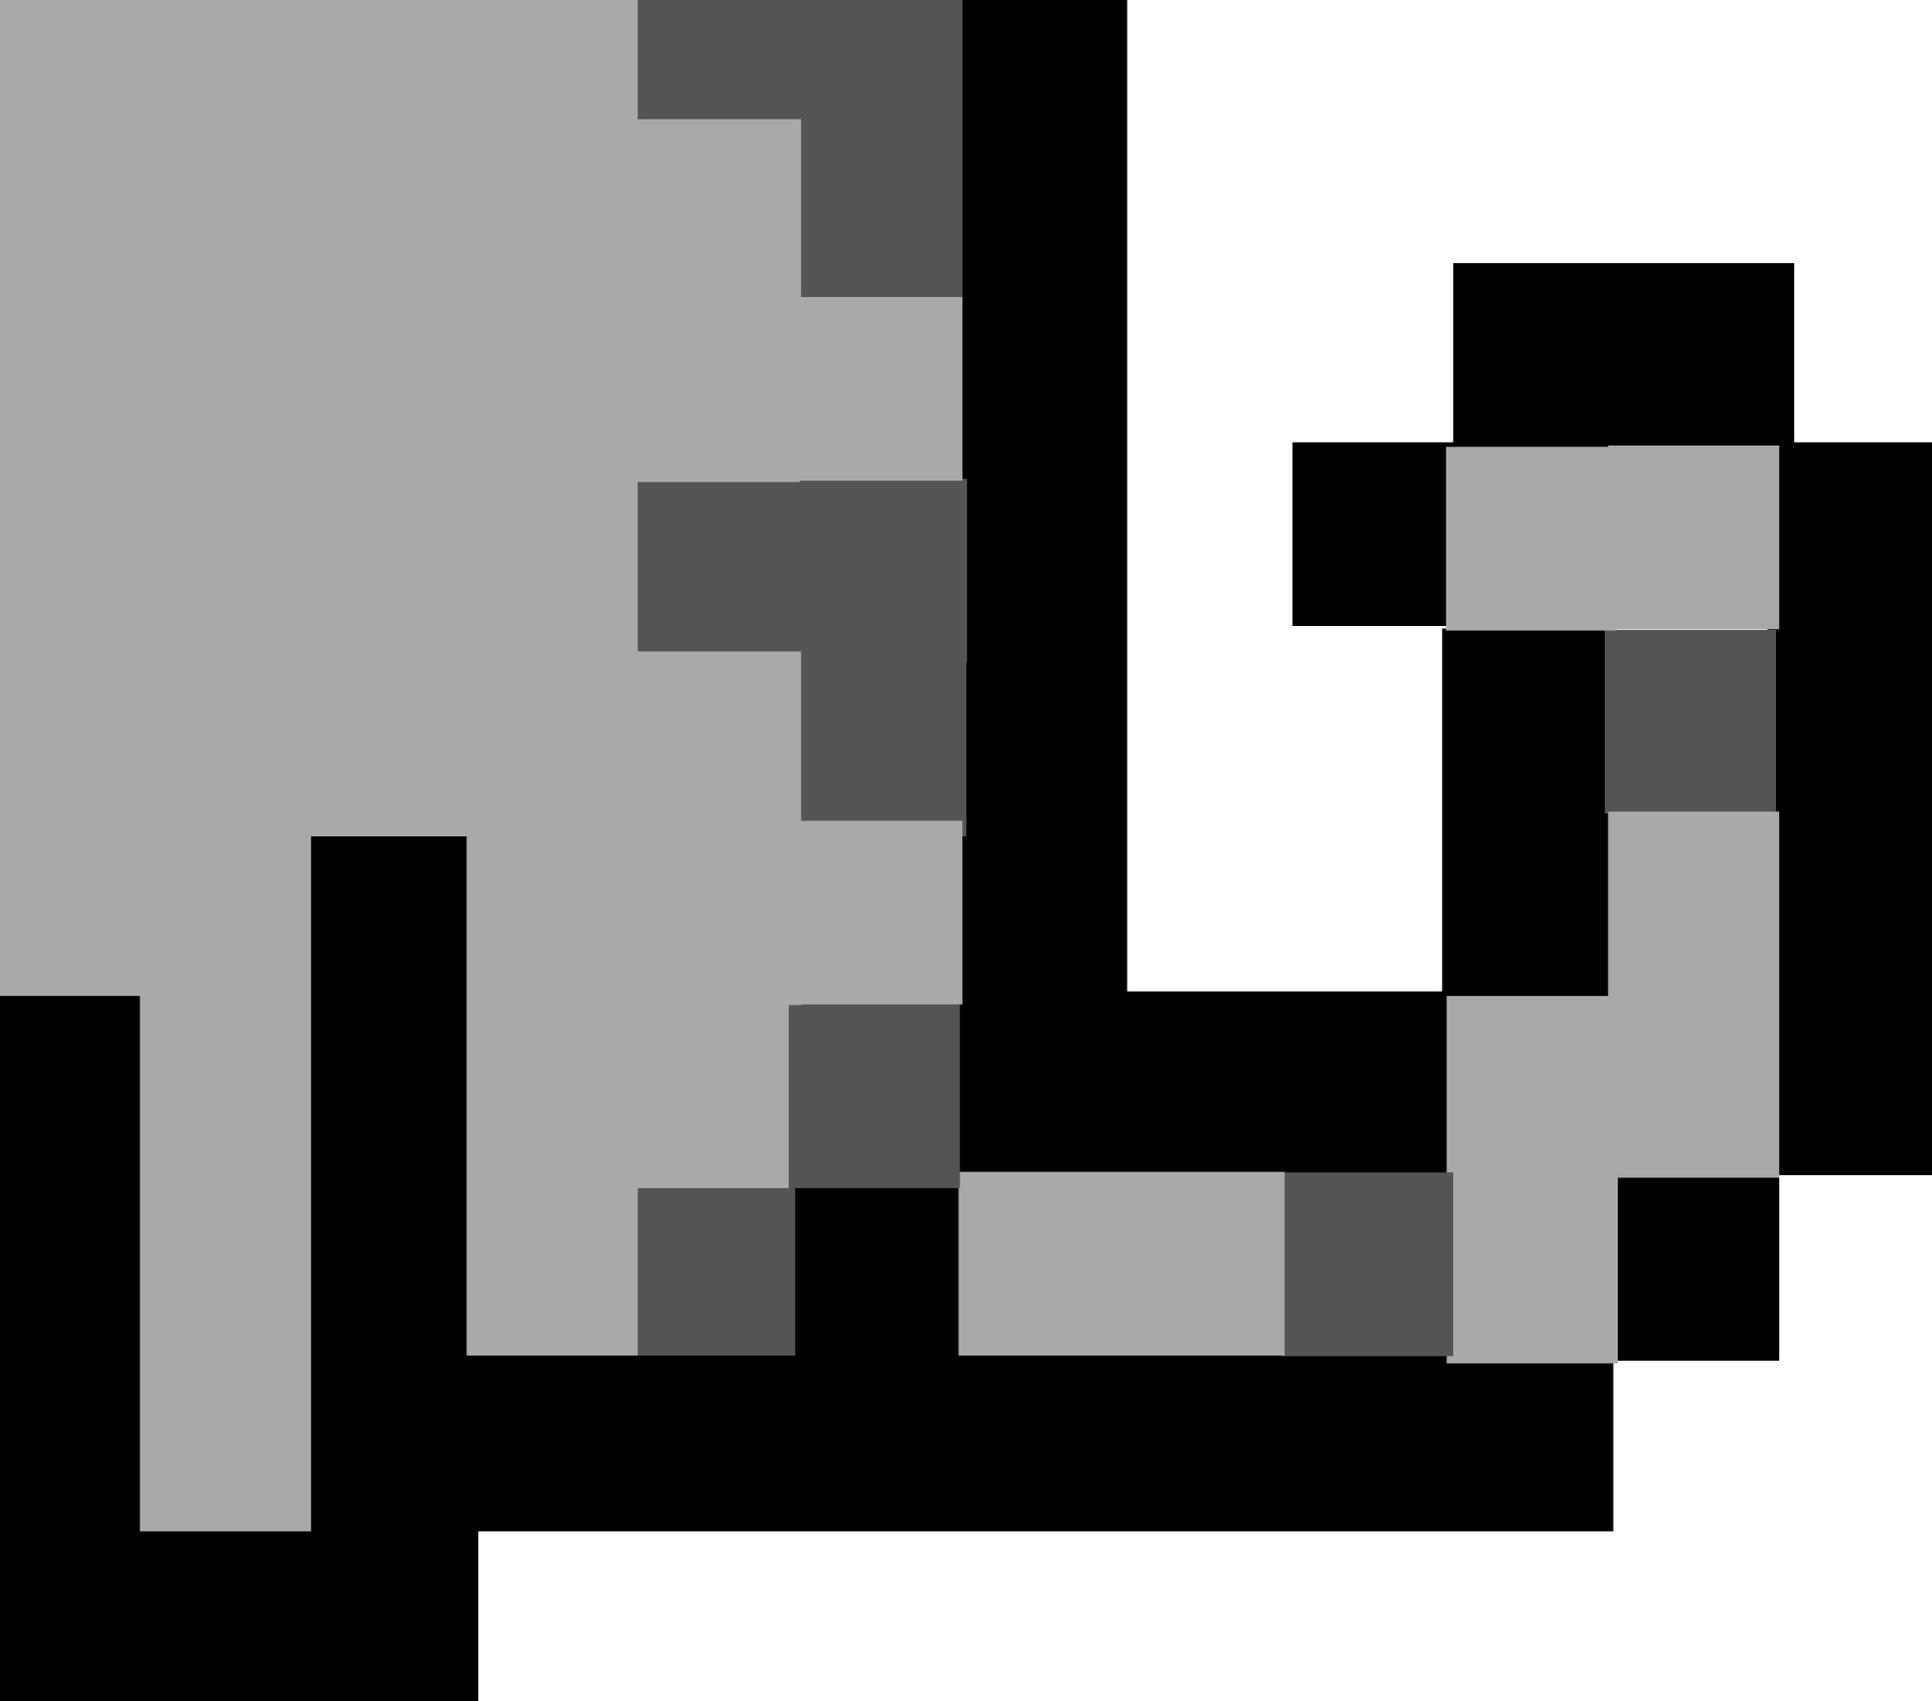


An example: the space of
photographs.



An image is a point in
high dimensional
space





Dimension of the space of images

nm-dimensional

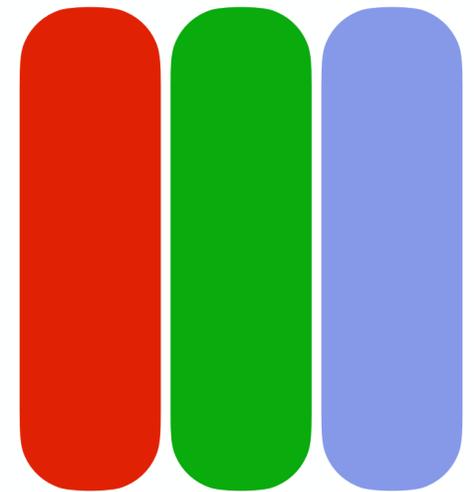
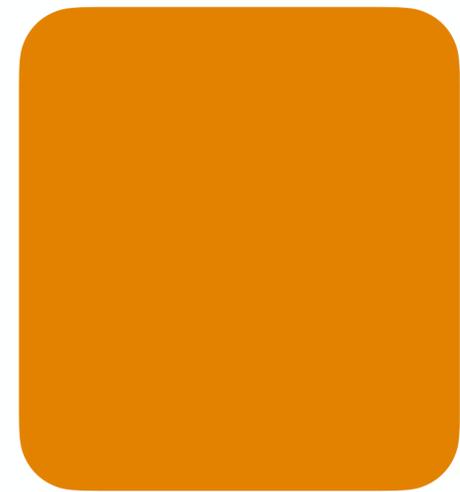


m Pixels



n Pixels

Color Pixels



3 numbers per pixel: red, green, blue.

Space of images = $3nm$ -dimensional

Dimension of the space of images

nm-dimensional

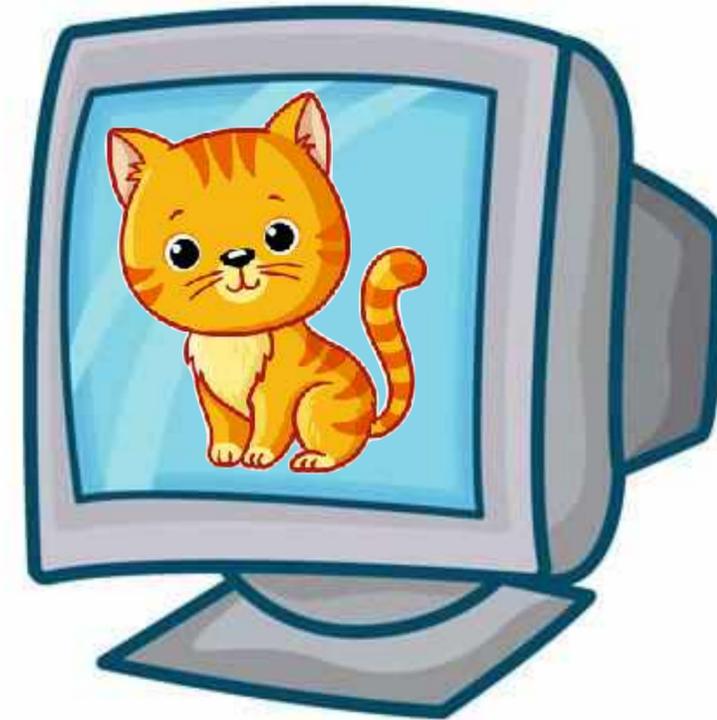


m Pixels

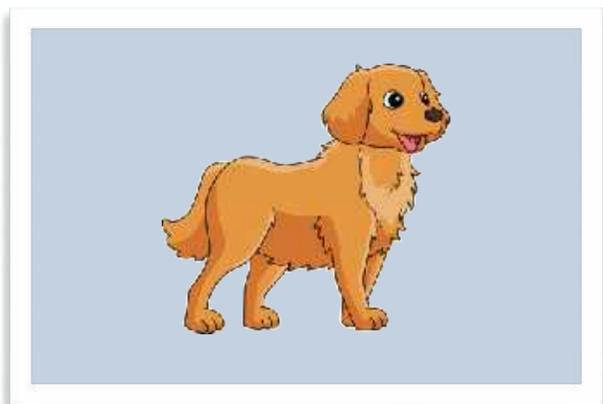


n Pixels

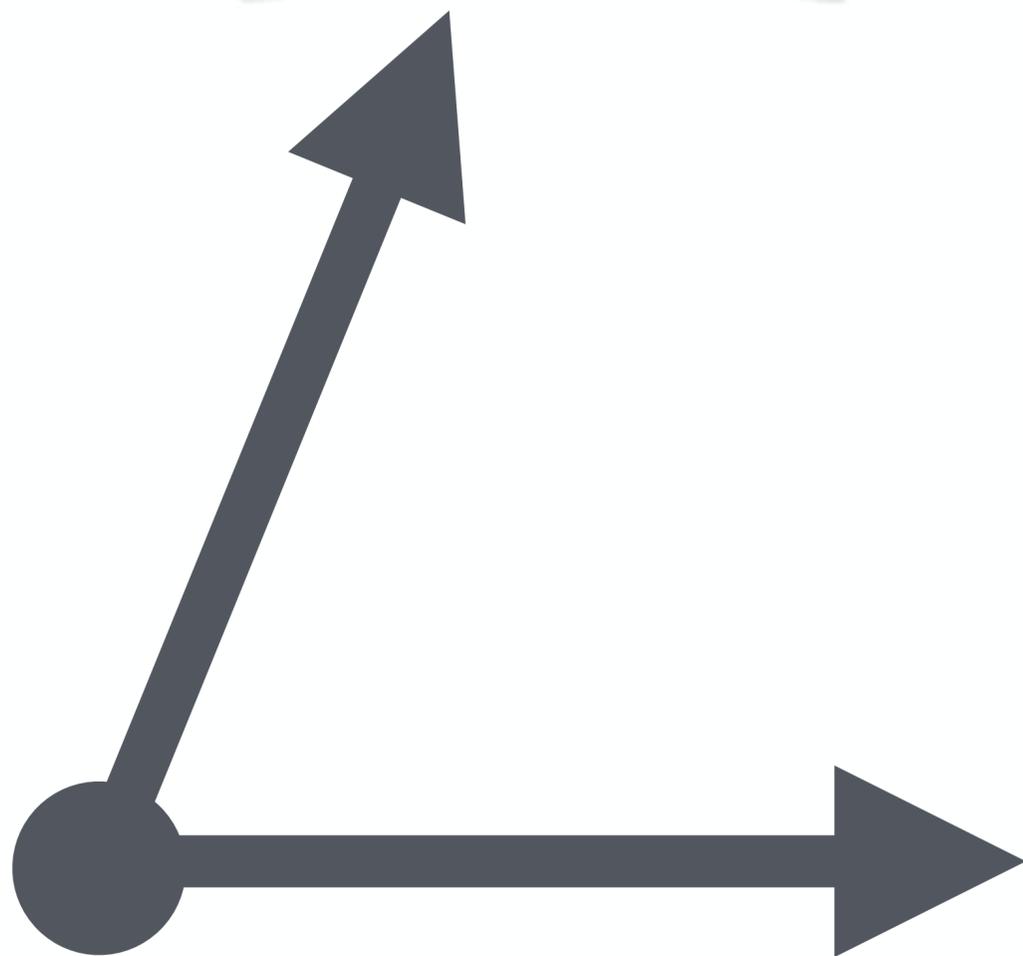
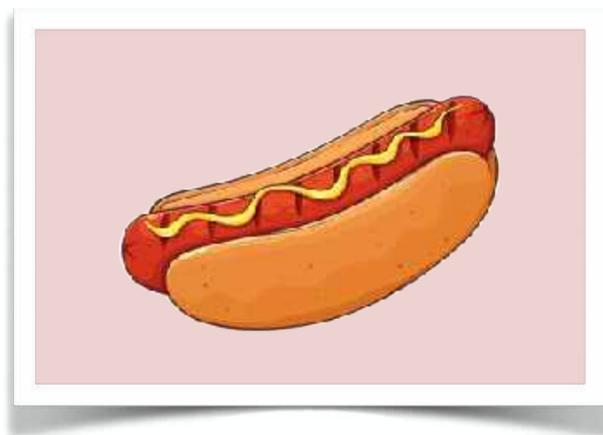
Our Example: 480p Color Images

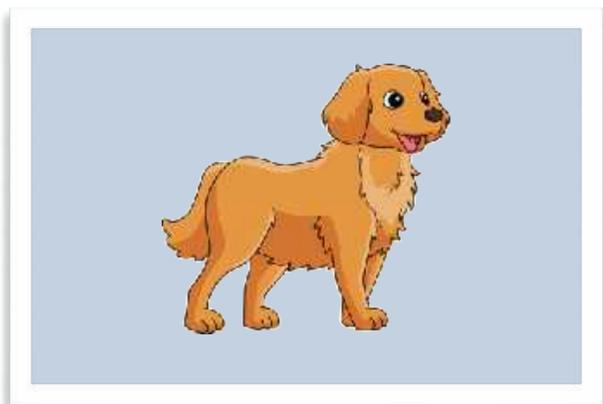


$$3 \times 480 \times 320 = 460,800$$



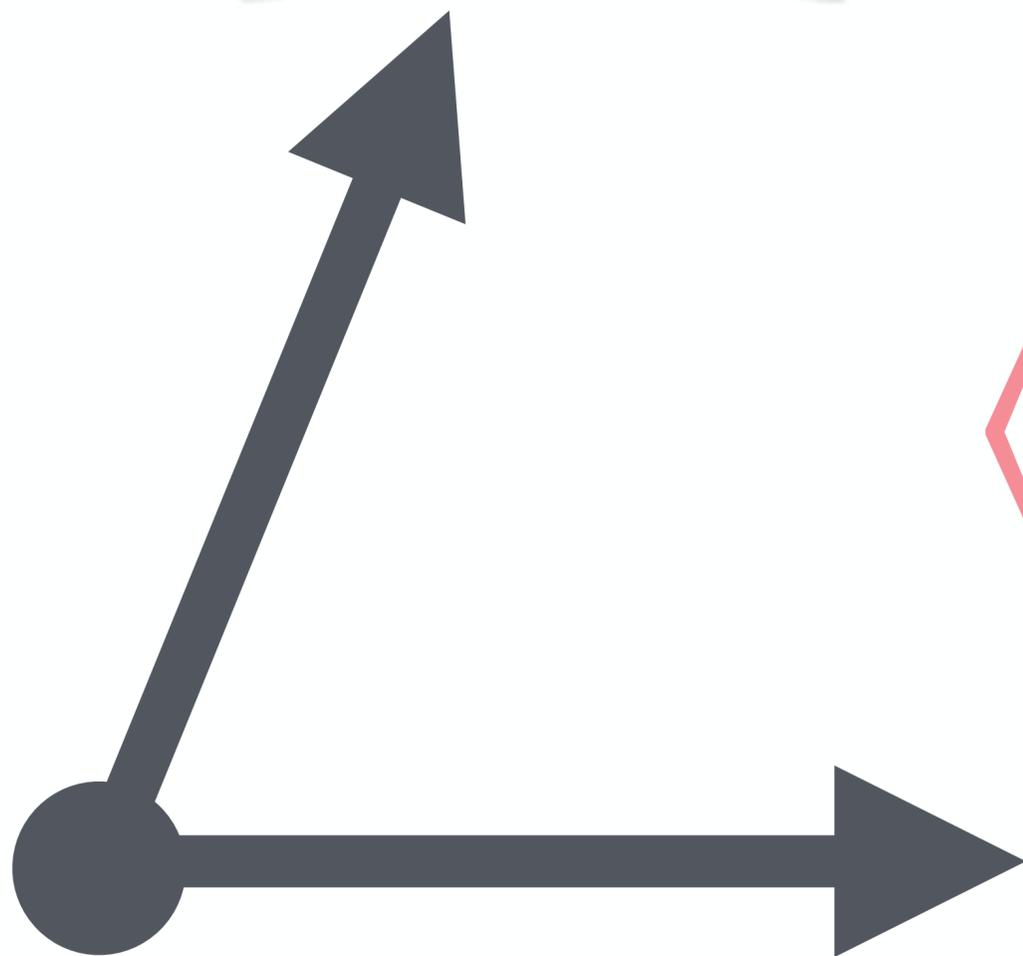
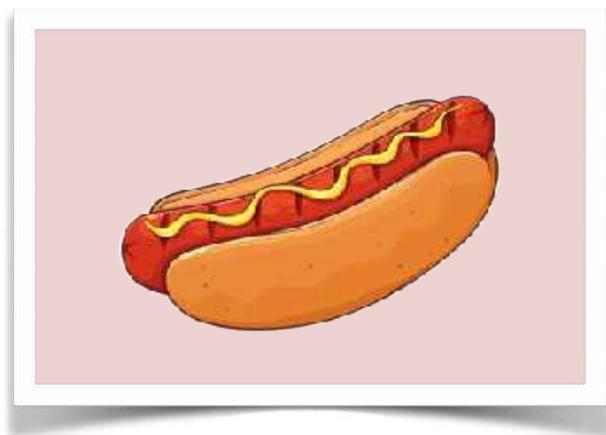
What is the *distance* between a dog and a hot dog?

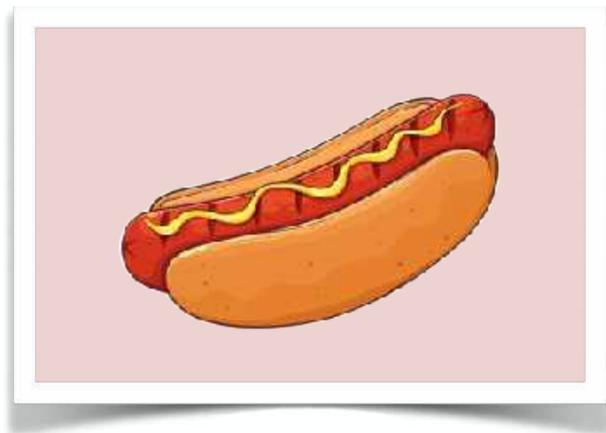
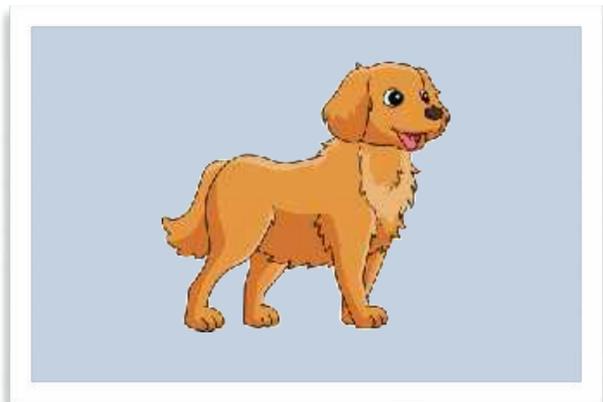




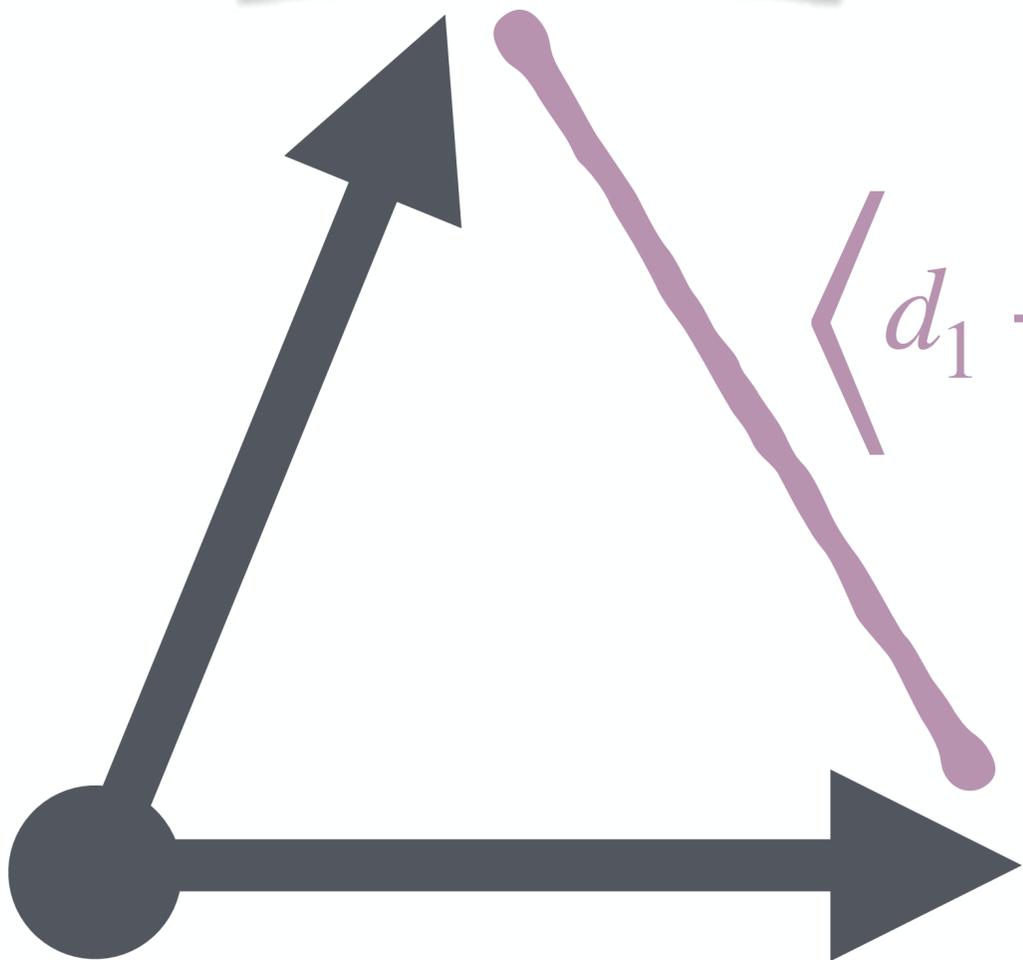
$\langle d_1, d_2, d_3, d_4, d_5, d_6, d_7, d_8 \dots d_{460,800} \rangle$

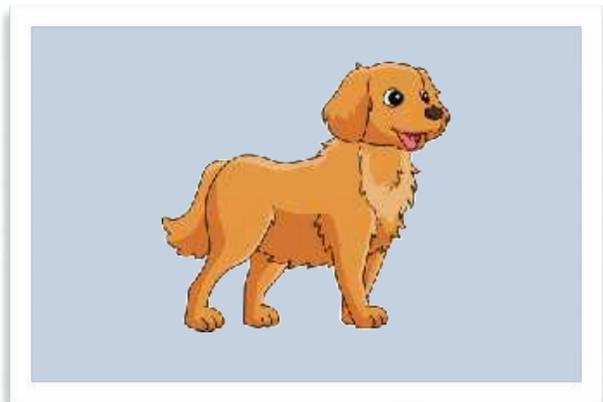
$\langle h_1, h_2, h_3, h_4, h_5, h_6, h_7, h_8 \dots h_{460,800} \rangle$



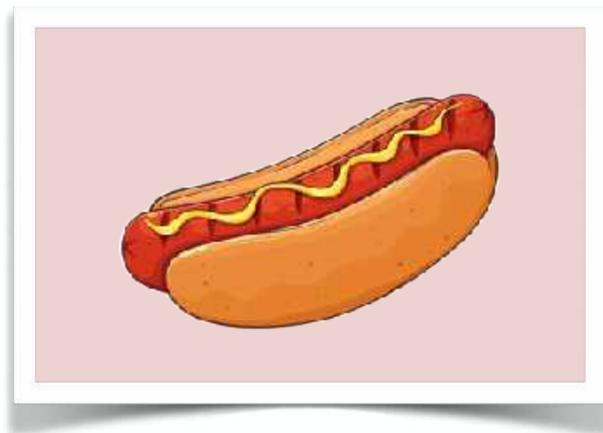


$\langle d_1 - h_1, d_2 - h_2, \dots, d_{460,800} - h_{460,800} \rangle$





$$\text{dist} = \sqrt{\sum_{i=1}^{460,800} (d_i - h_i)^2}$$



Thinking of the set of images as having *geometry* is the essential insight to many modern Machine Learning applications.

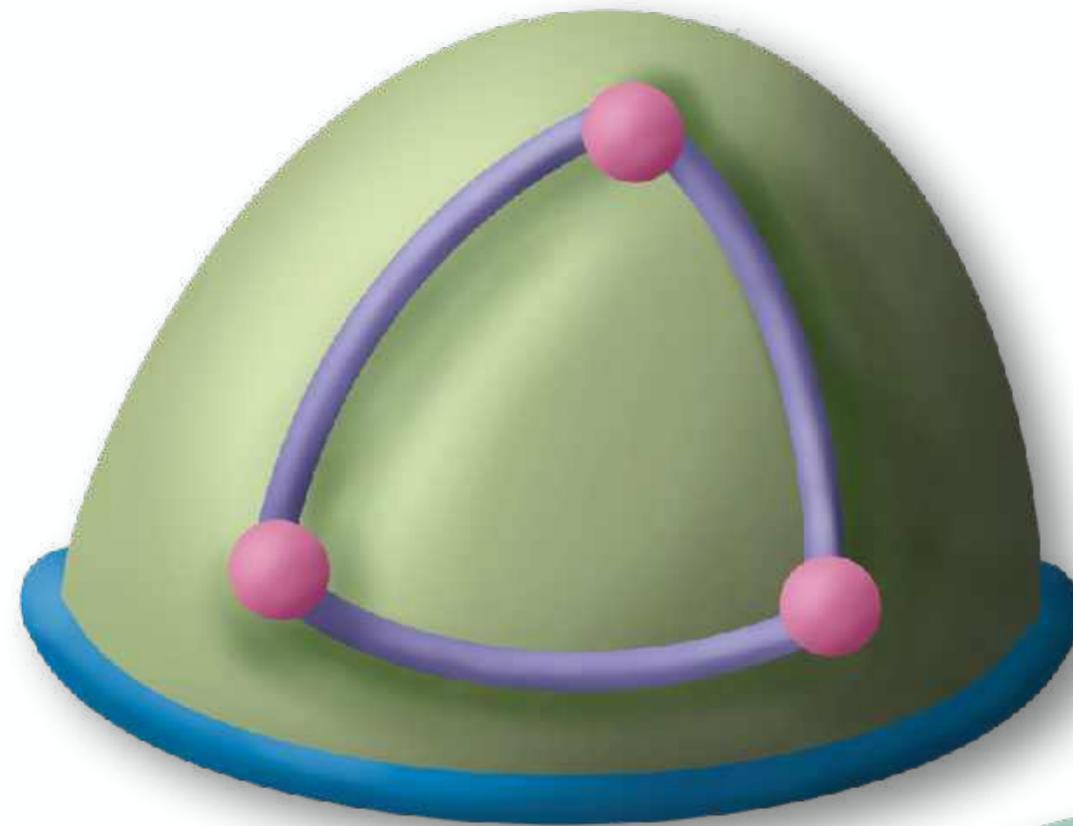


Space can be curved

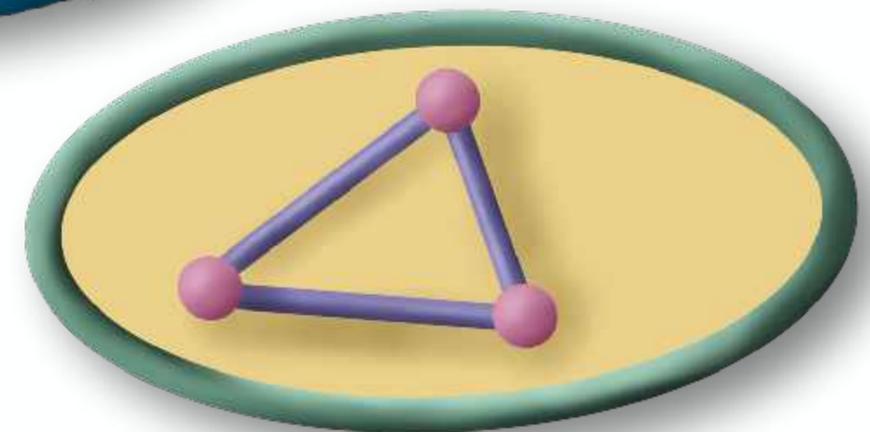


Gauß

As a surveyor, Gauss had to think about 'realistic' geometry.

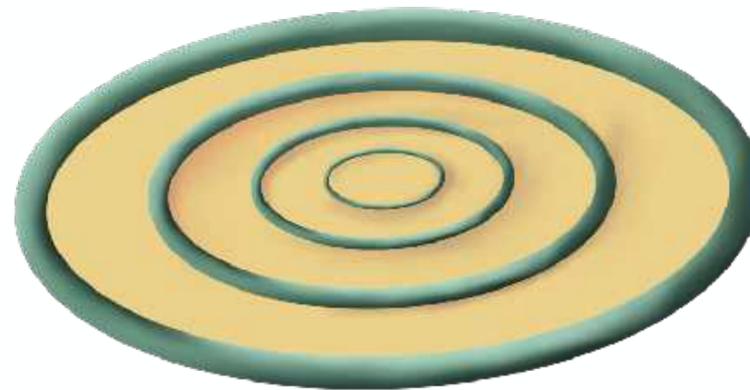
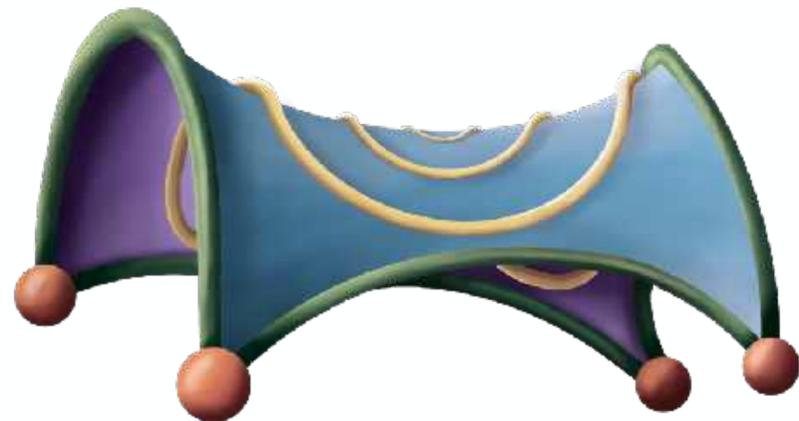
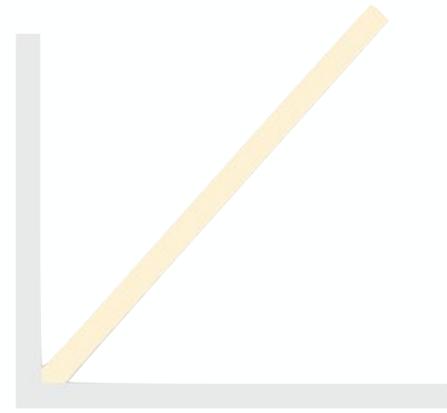
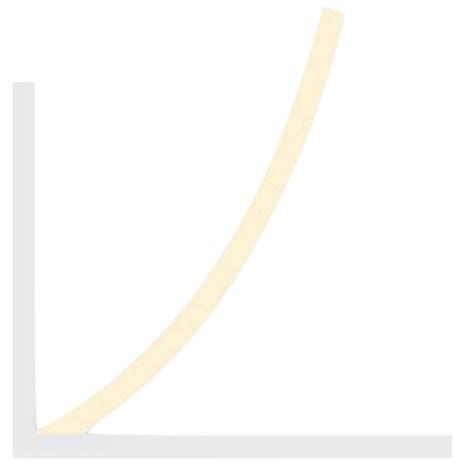


Geometry of a hilltop: triangles don't add to 180!

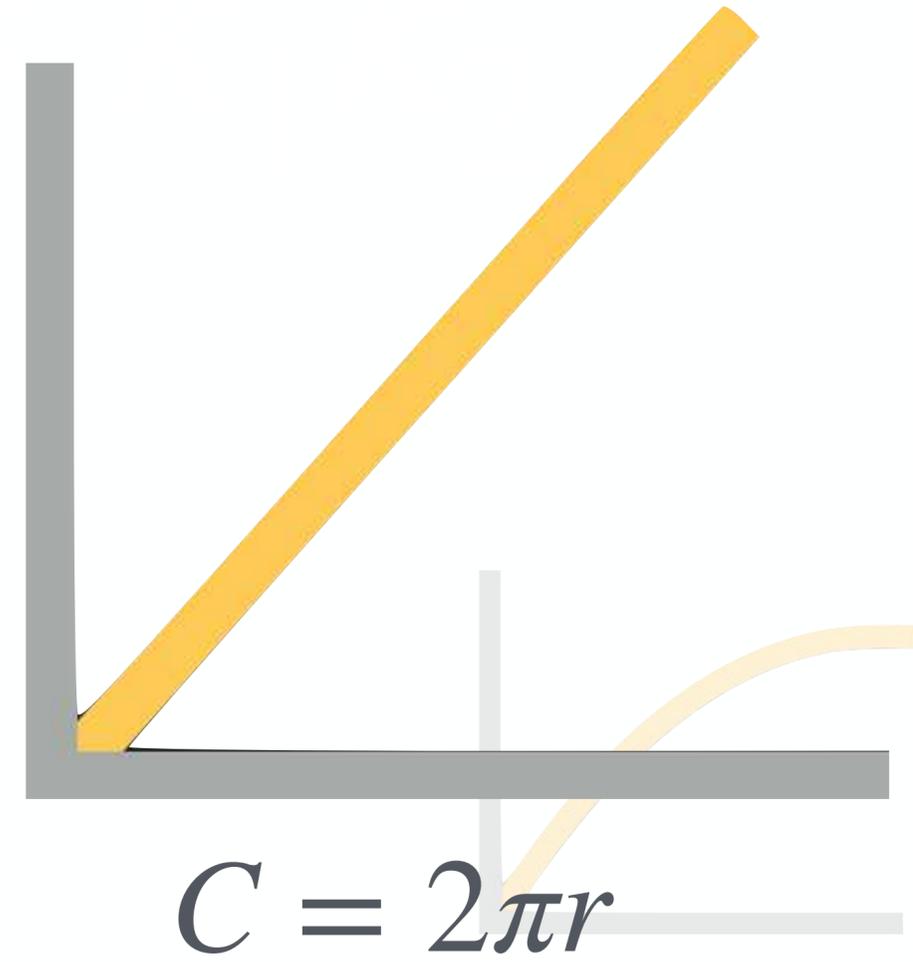
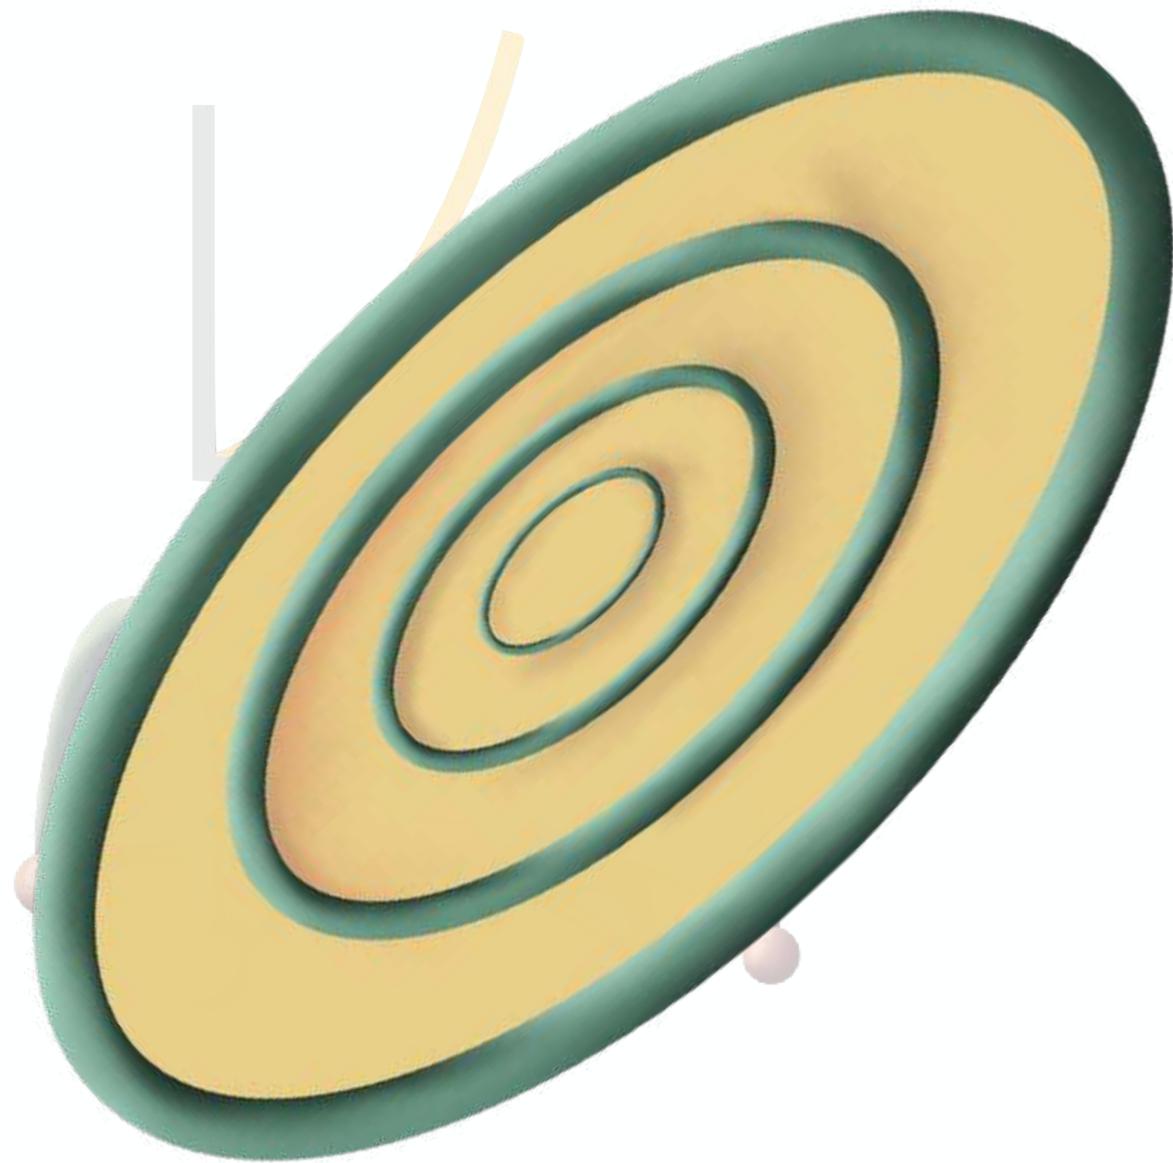


What is Curvature?

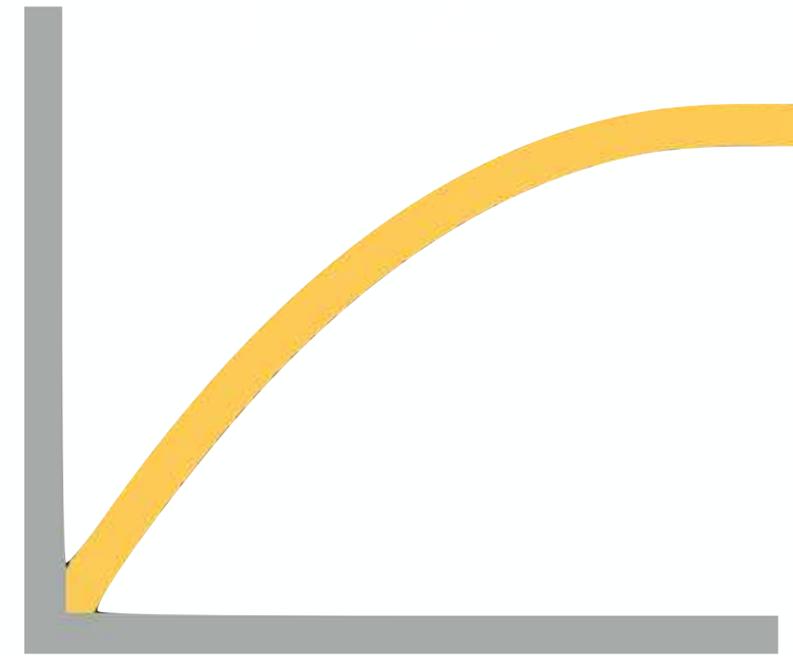
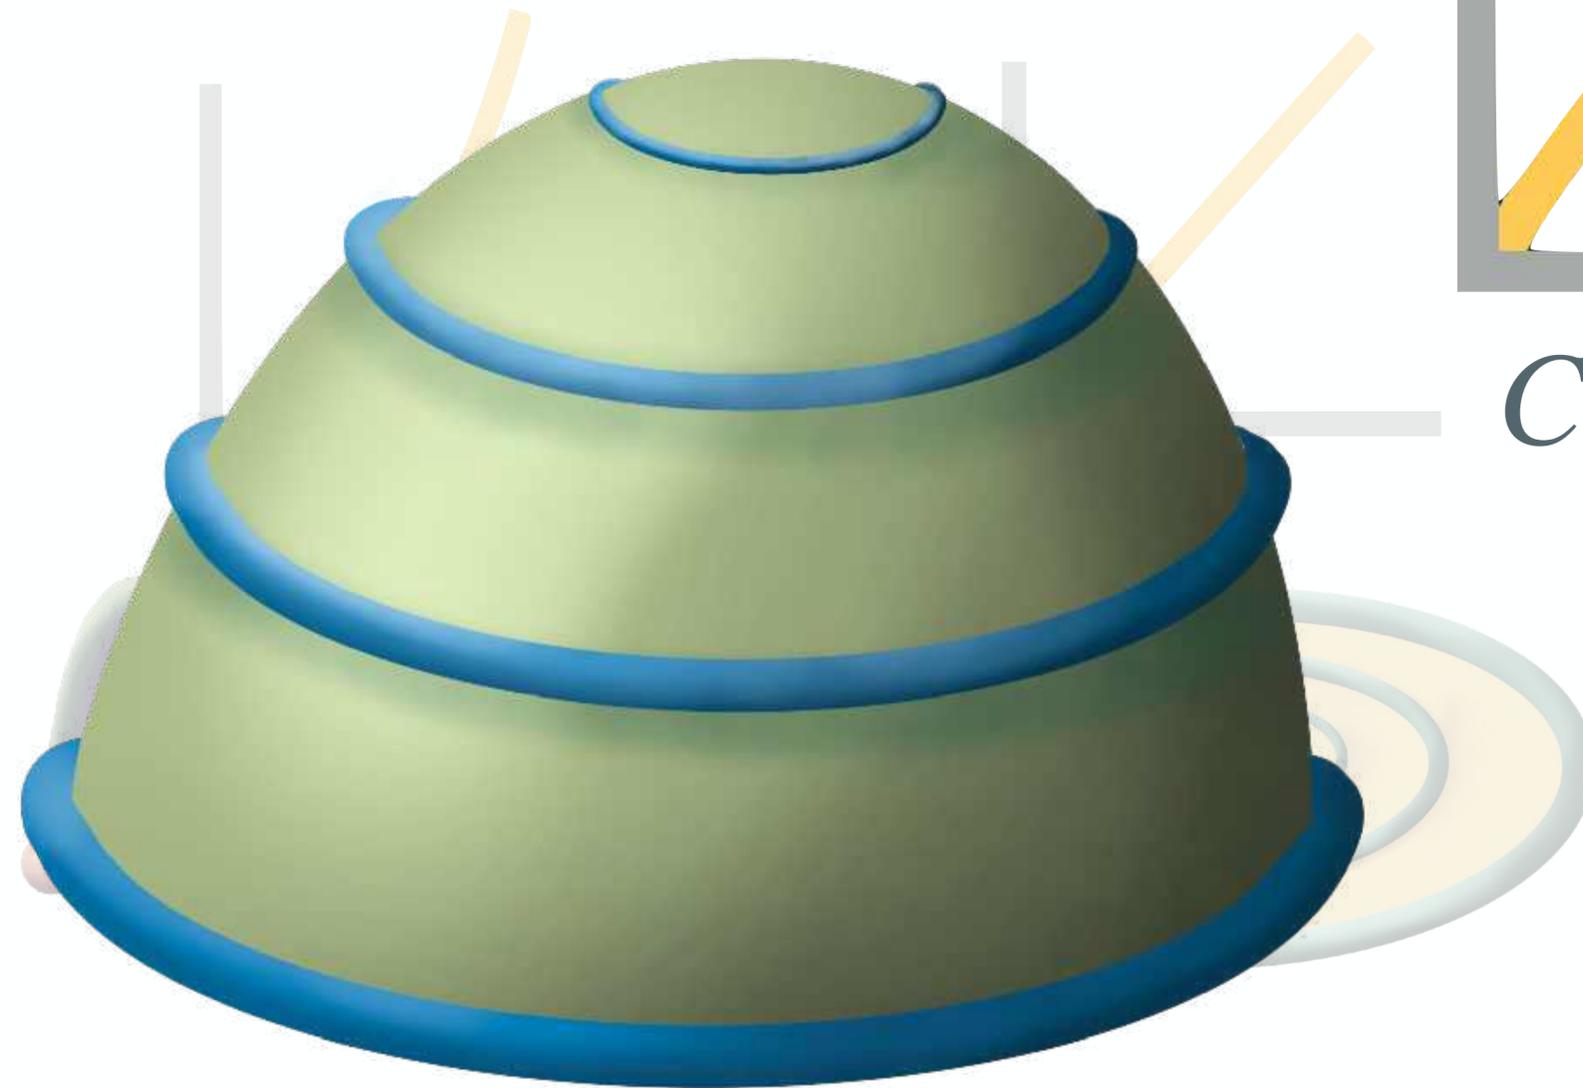
How can we quantify the deviation of a surface from the plane?



What is Curvature?

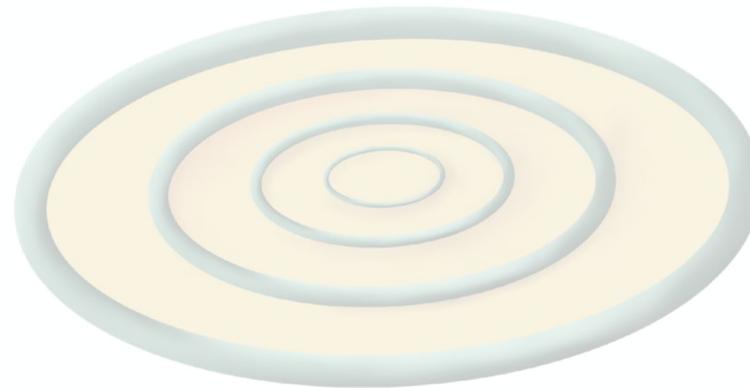
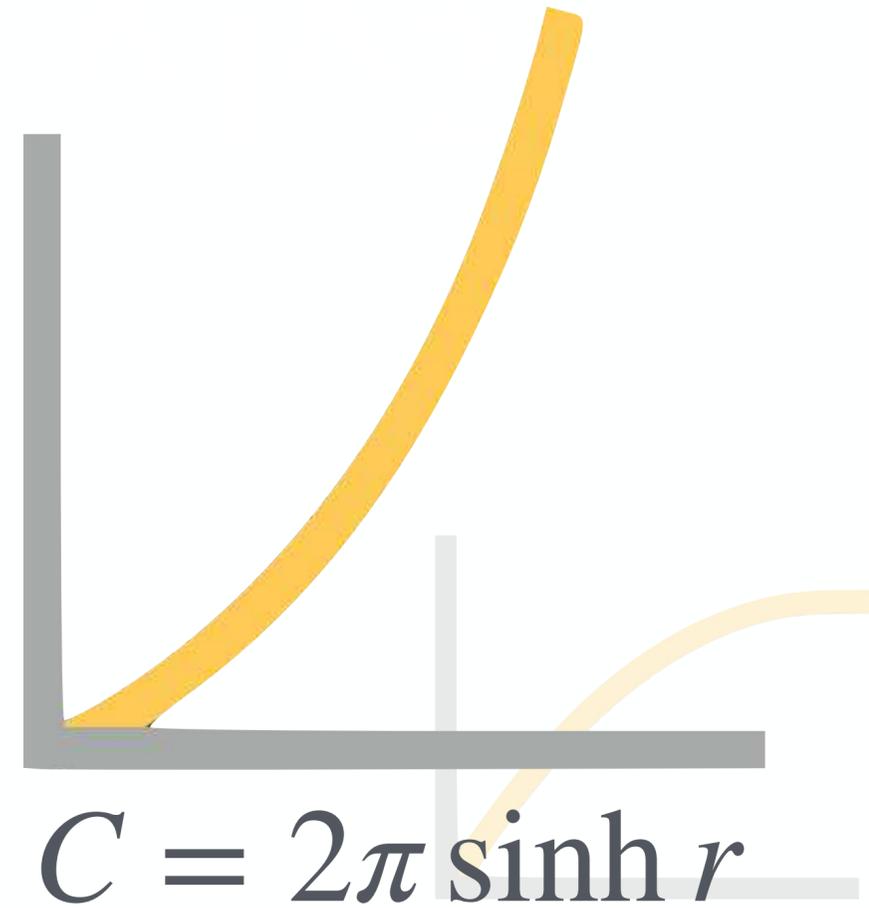
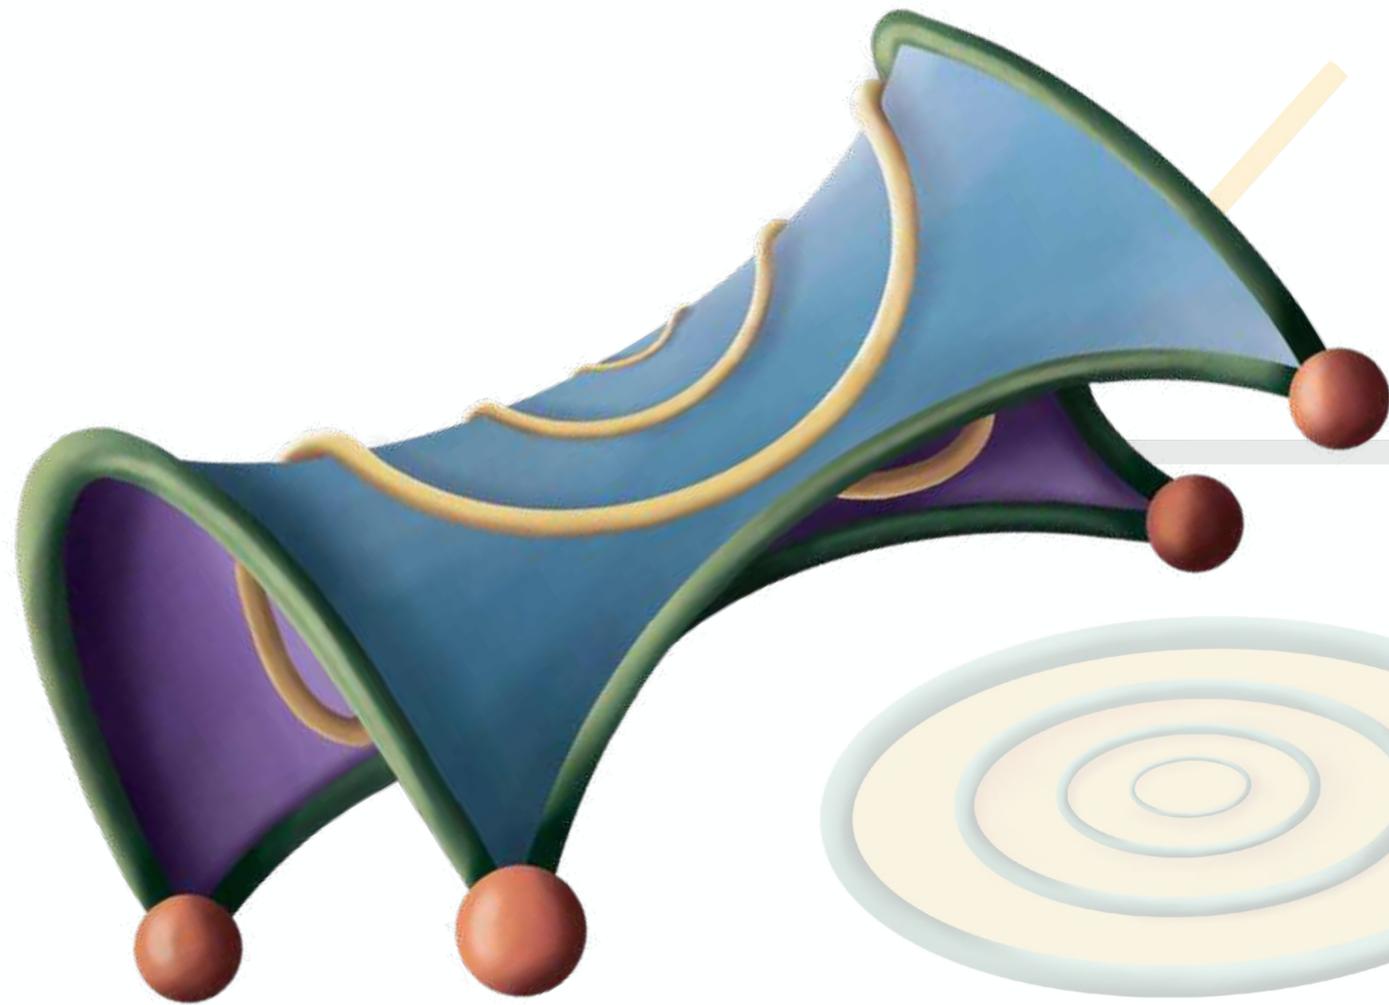


What is Curvature?



$$C = 2\pi \sin r$$

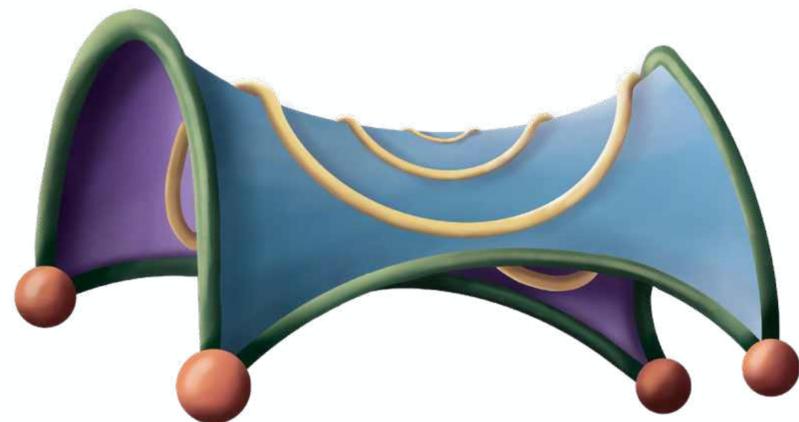
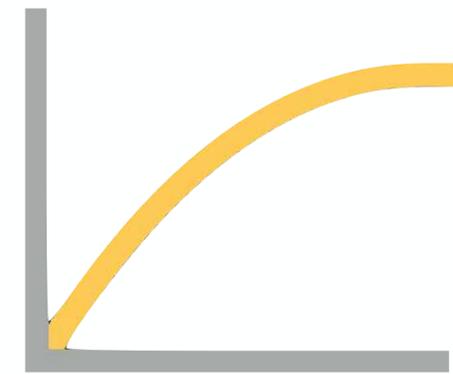
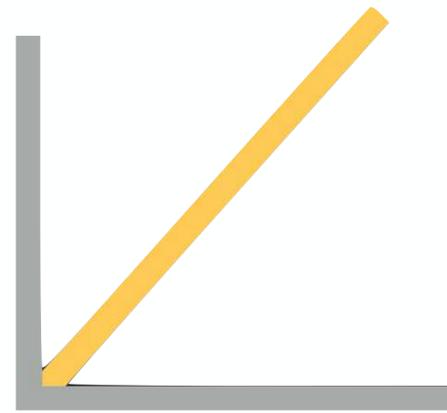
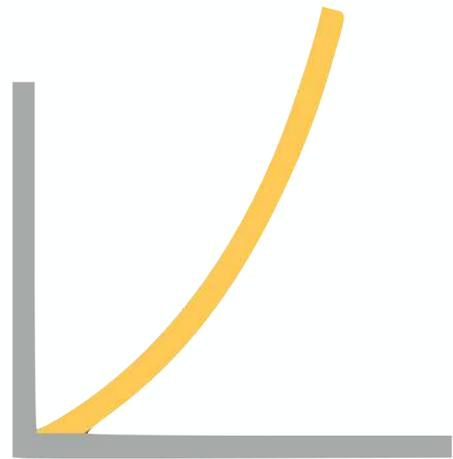
What is Curvature?



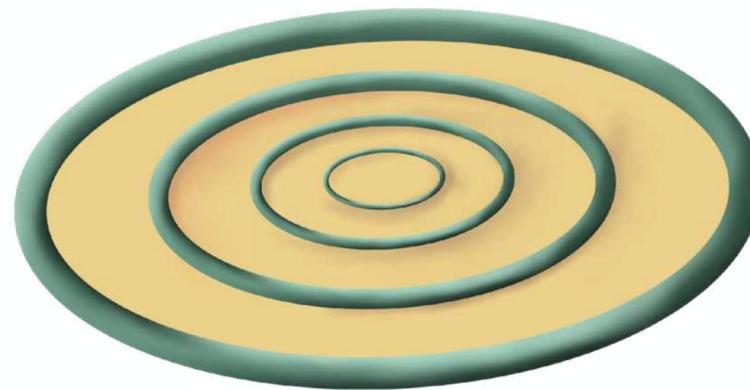
What is Curvature?

Curvature measures this difference.

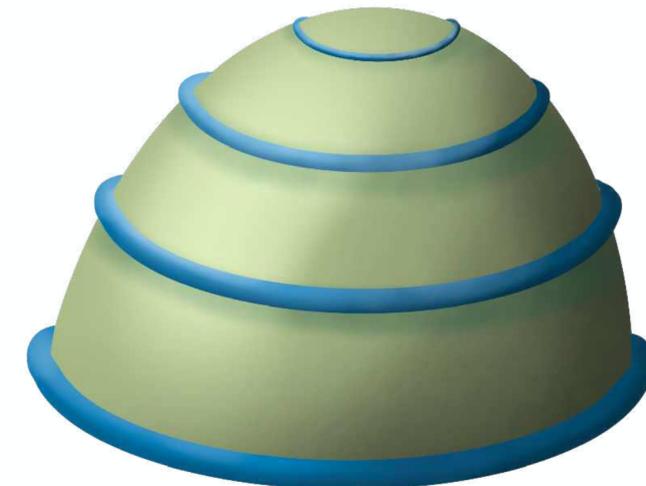
$$\kappa \propto \lim_{r \rightarrow 0} \frac{2\pi r - C(r)}{r^3}$$



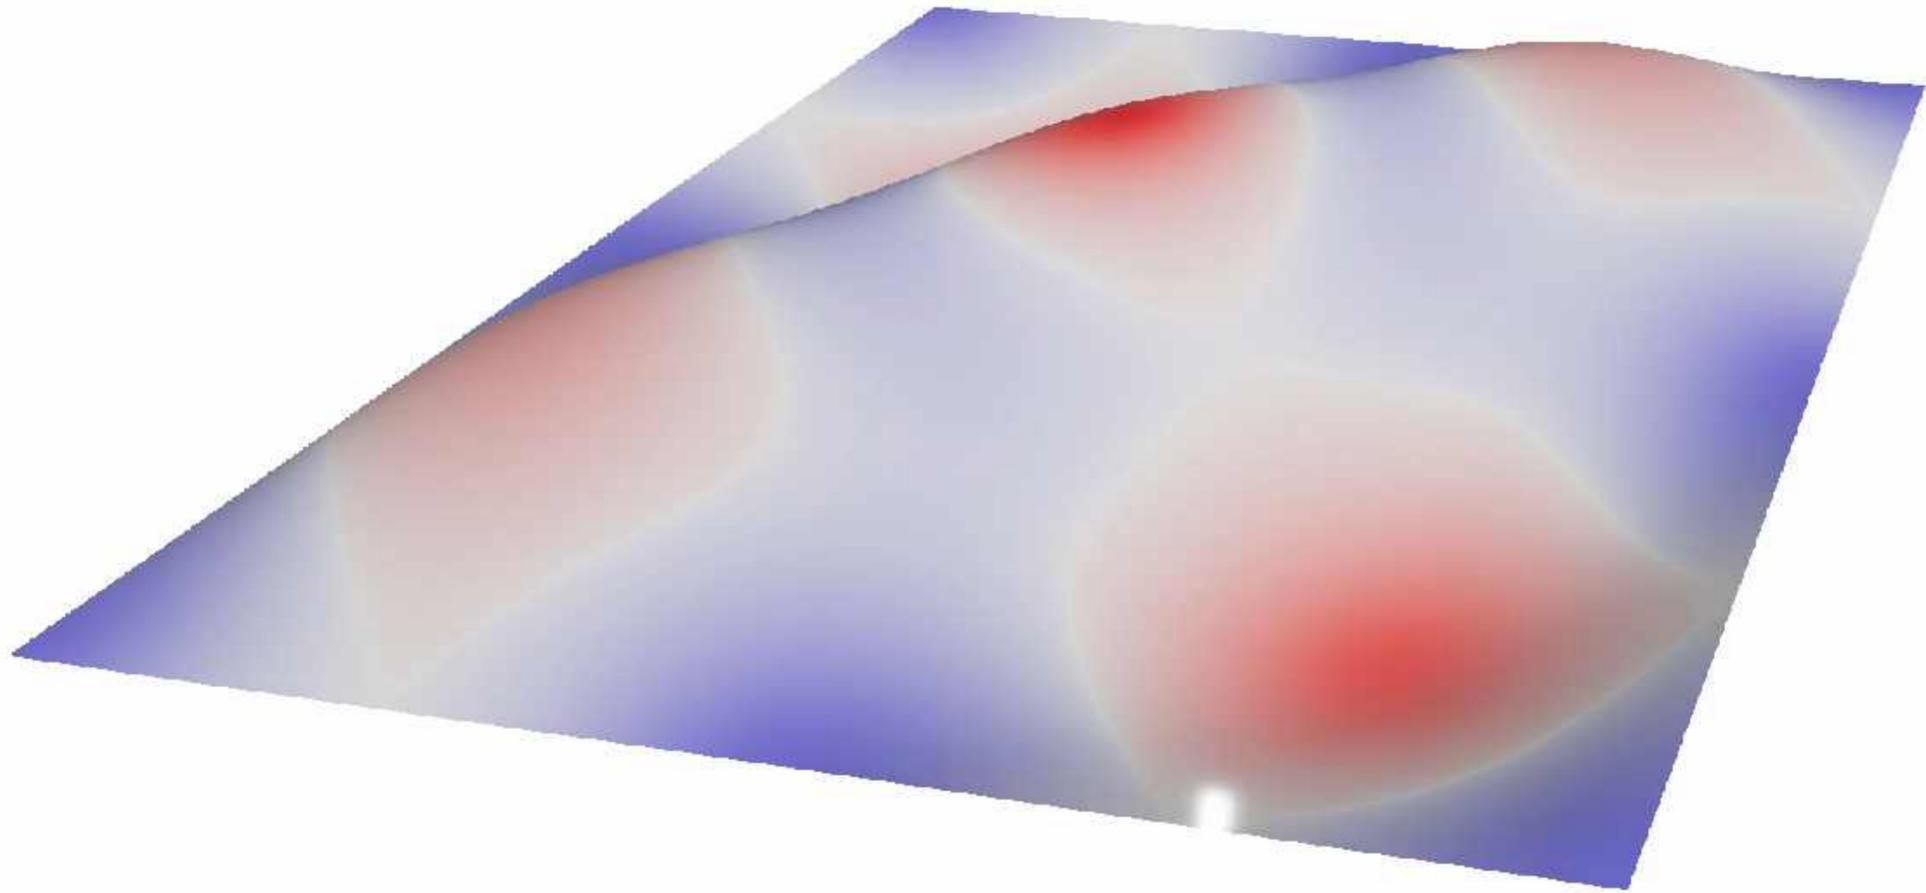
$$\kappa < 0$$



$$\kappa = 0$$



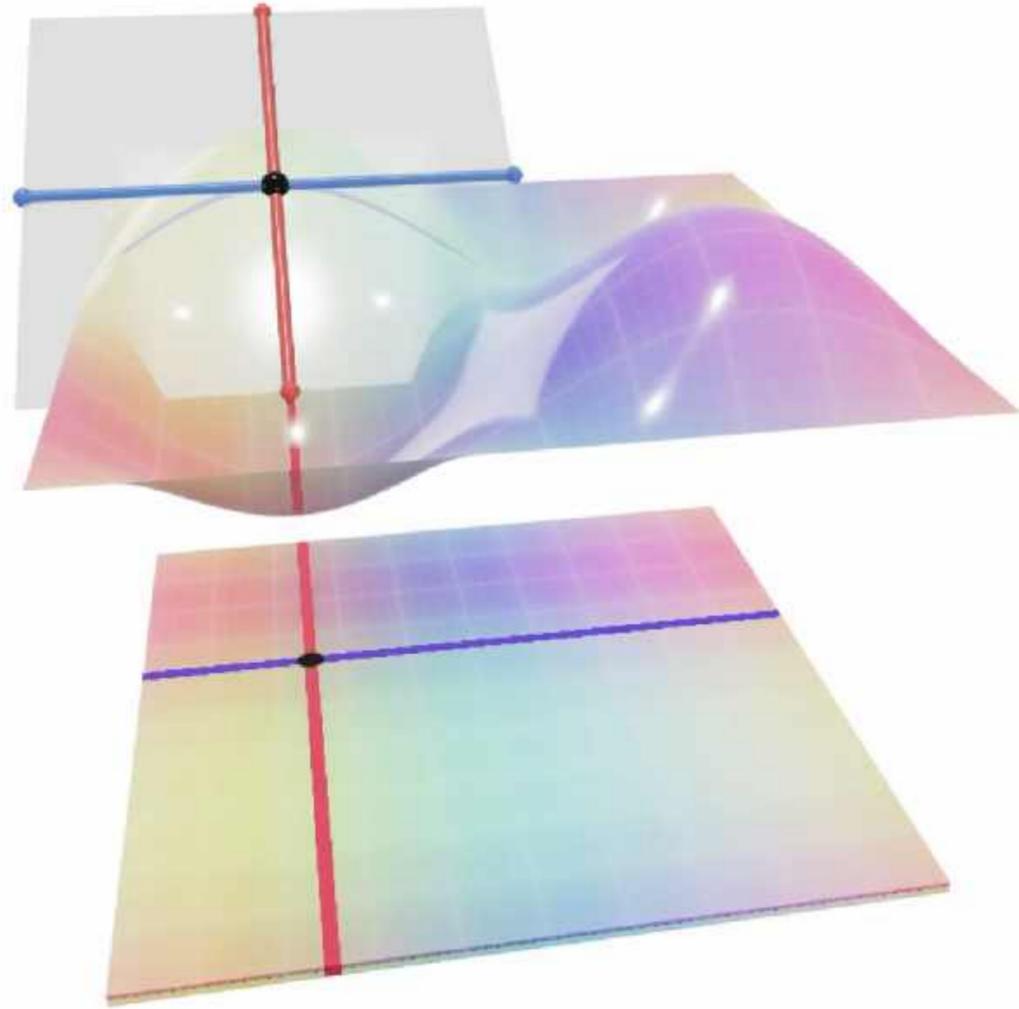
$$\kappa > 0$$



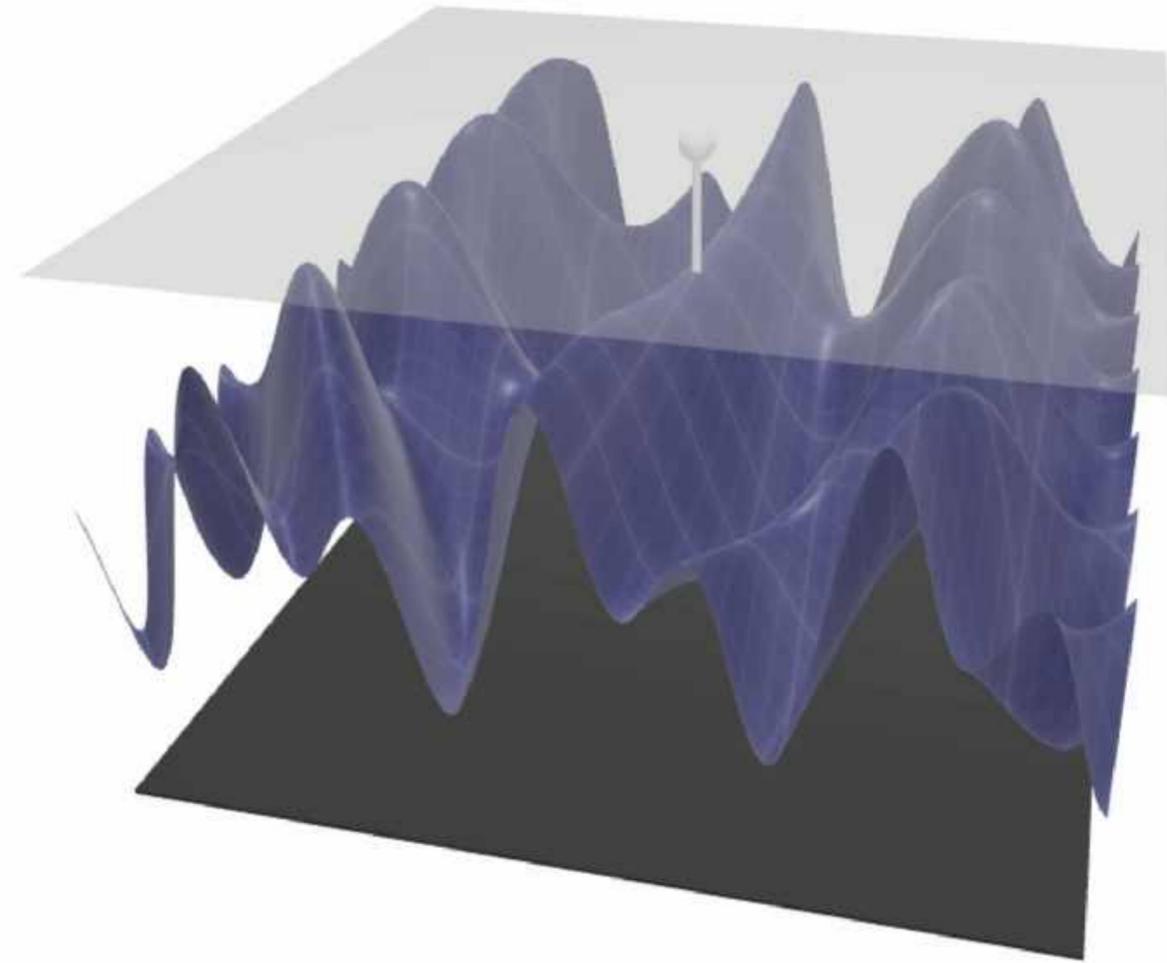
Positive Curvature

Negative Curvature

Applications to Mathematics:



Differential geometry: careful thinking about vector calculus and distances.



Differential topology: understanding spaces / graphs / functions by slicing.

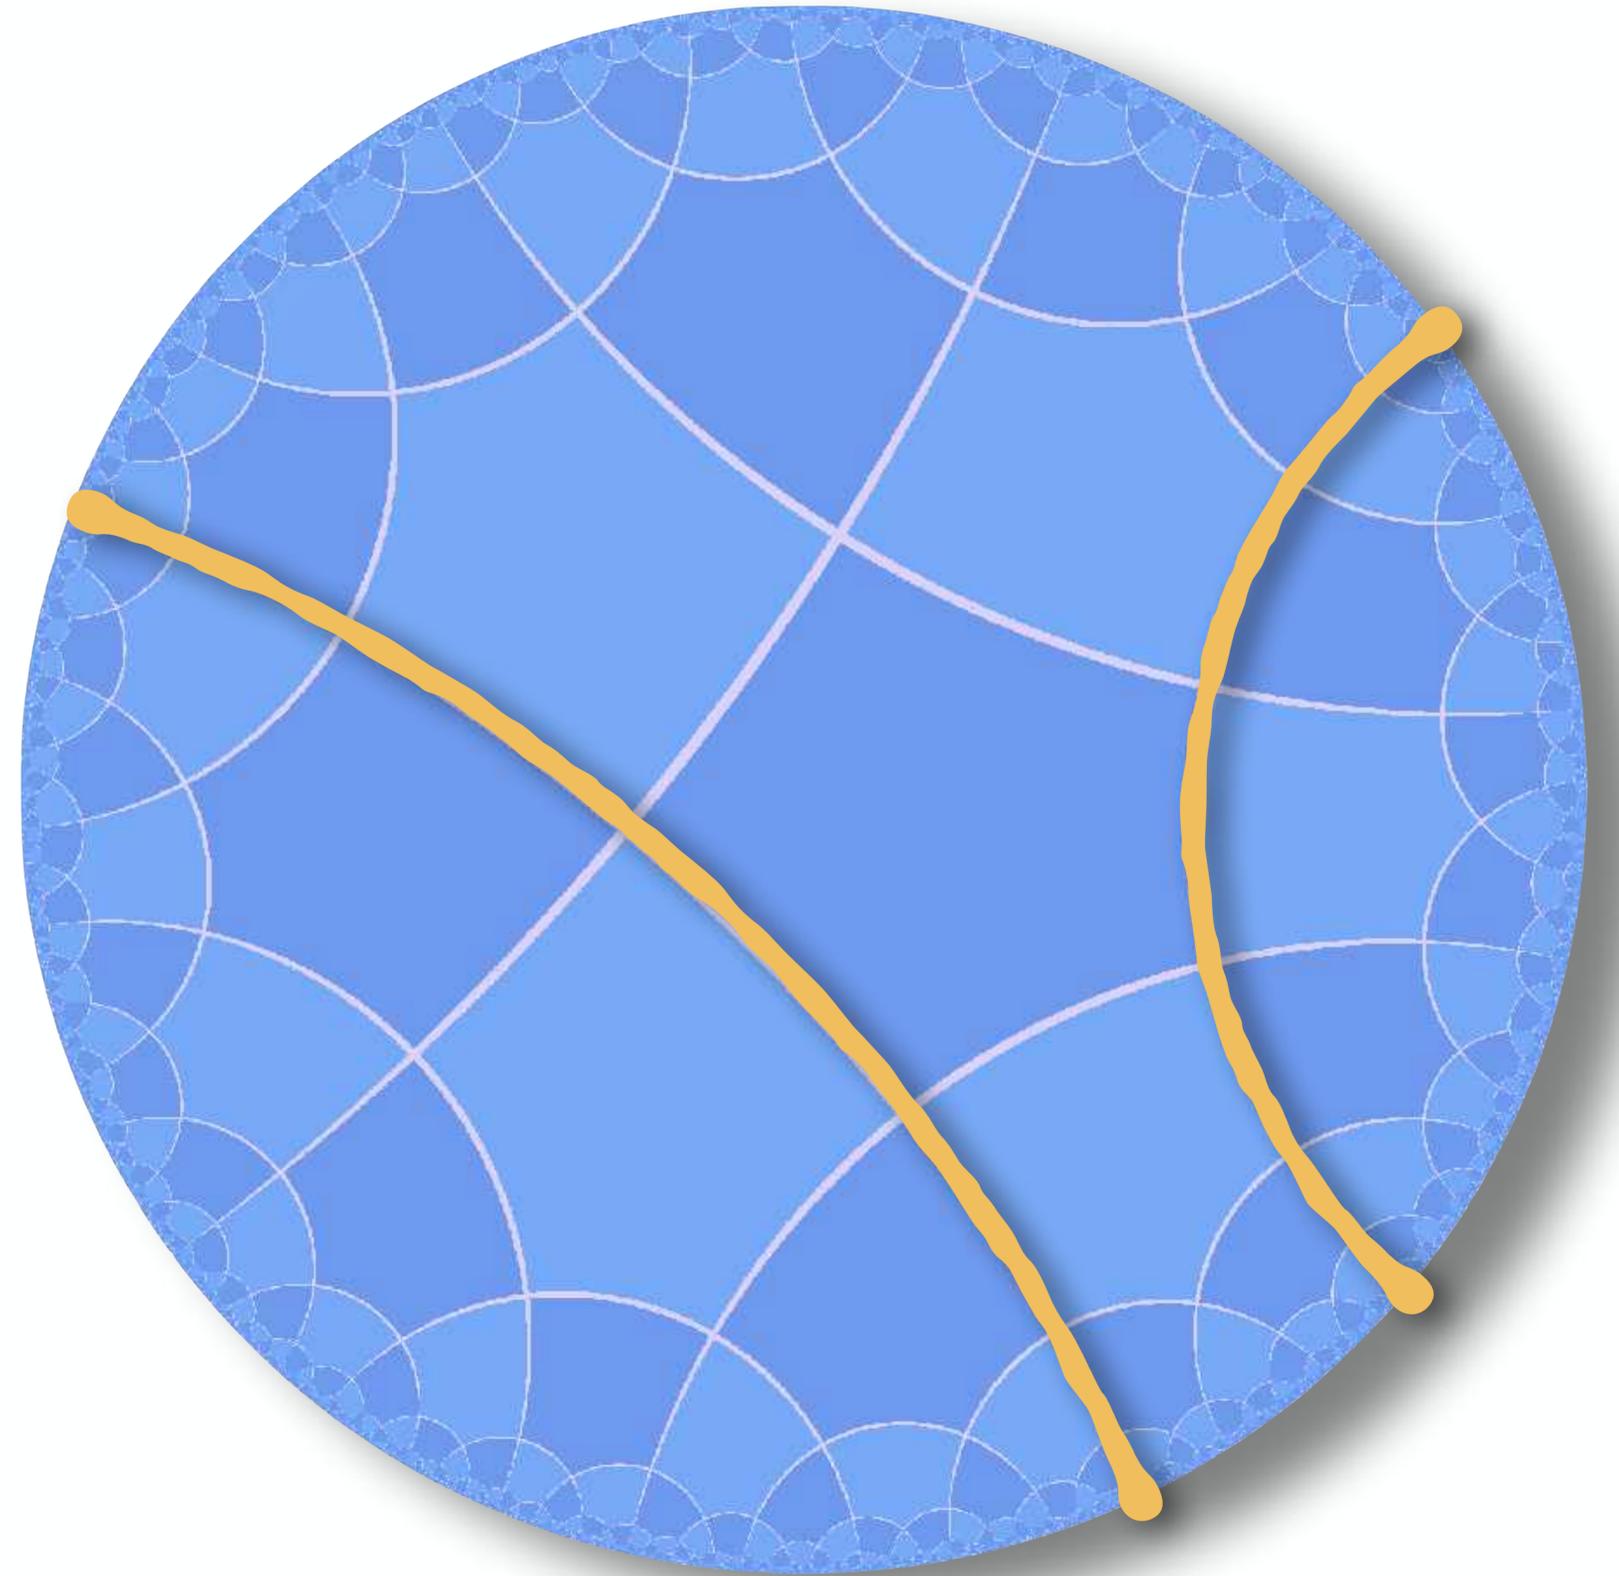
Applications to Mathematics:

Taking curvature seriously led to the discovery of *Hyperbolic Geometry*



Волюаі Лобачевский Гауџ

Proved that Euclid's 5th postulate is independent of the first 4, which was open for nearly 2,000 years.

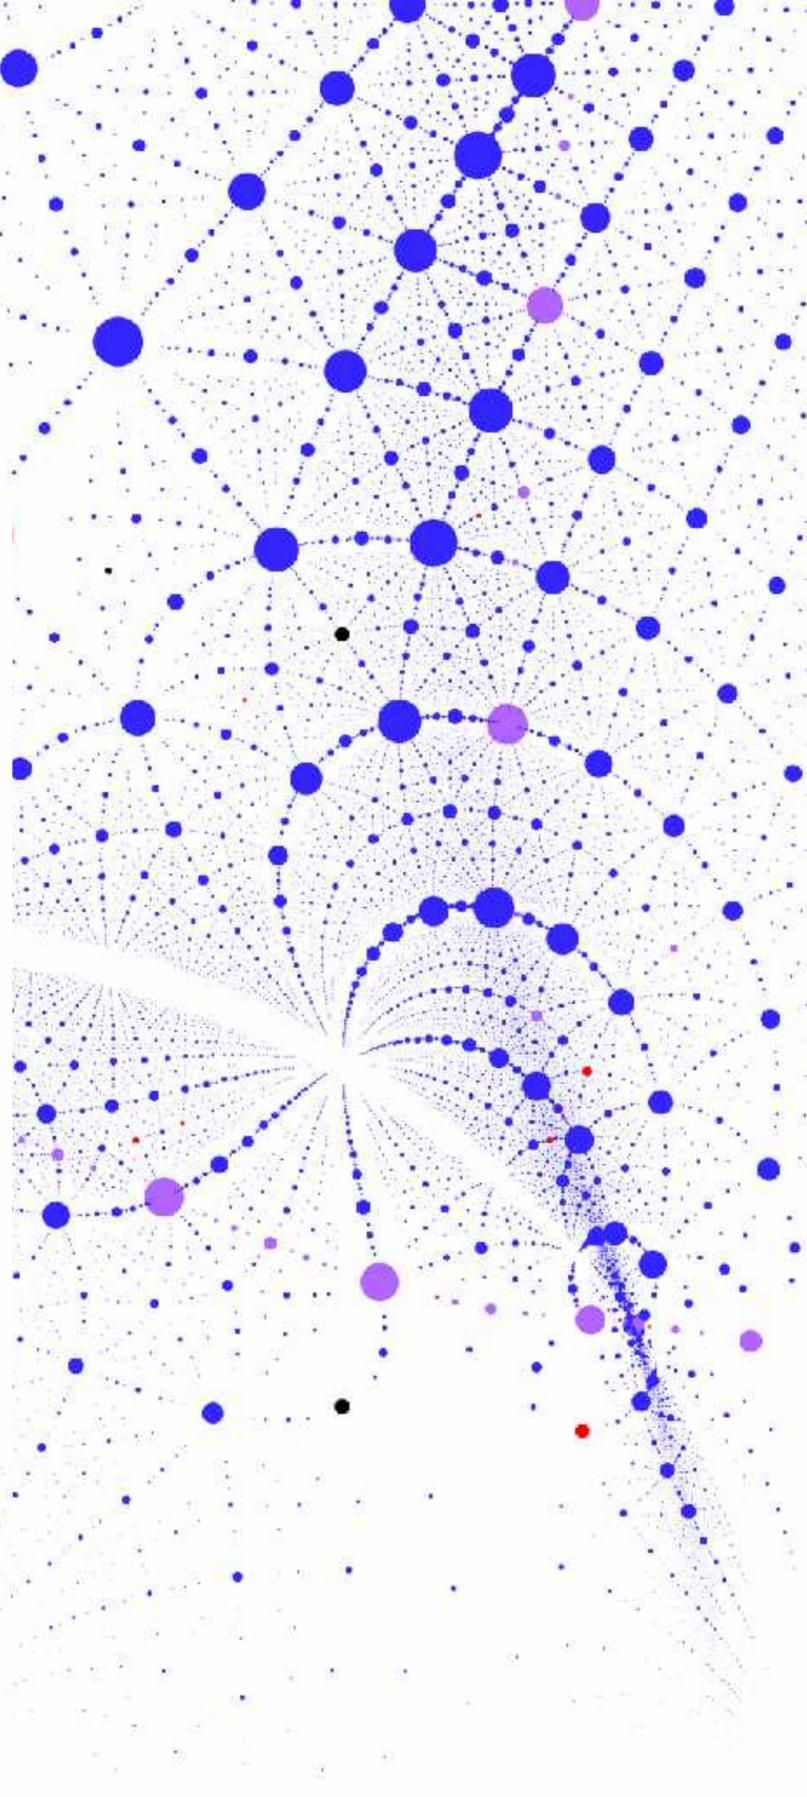
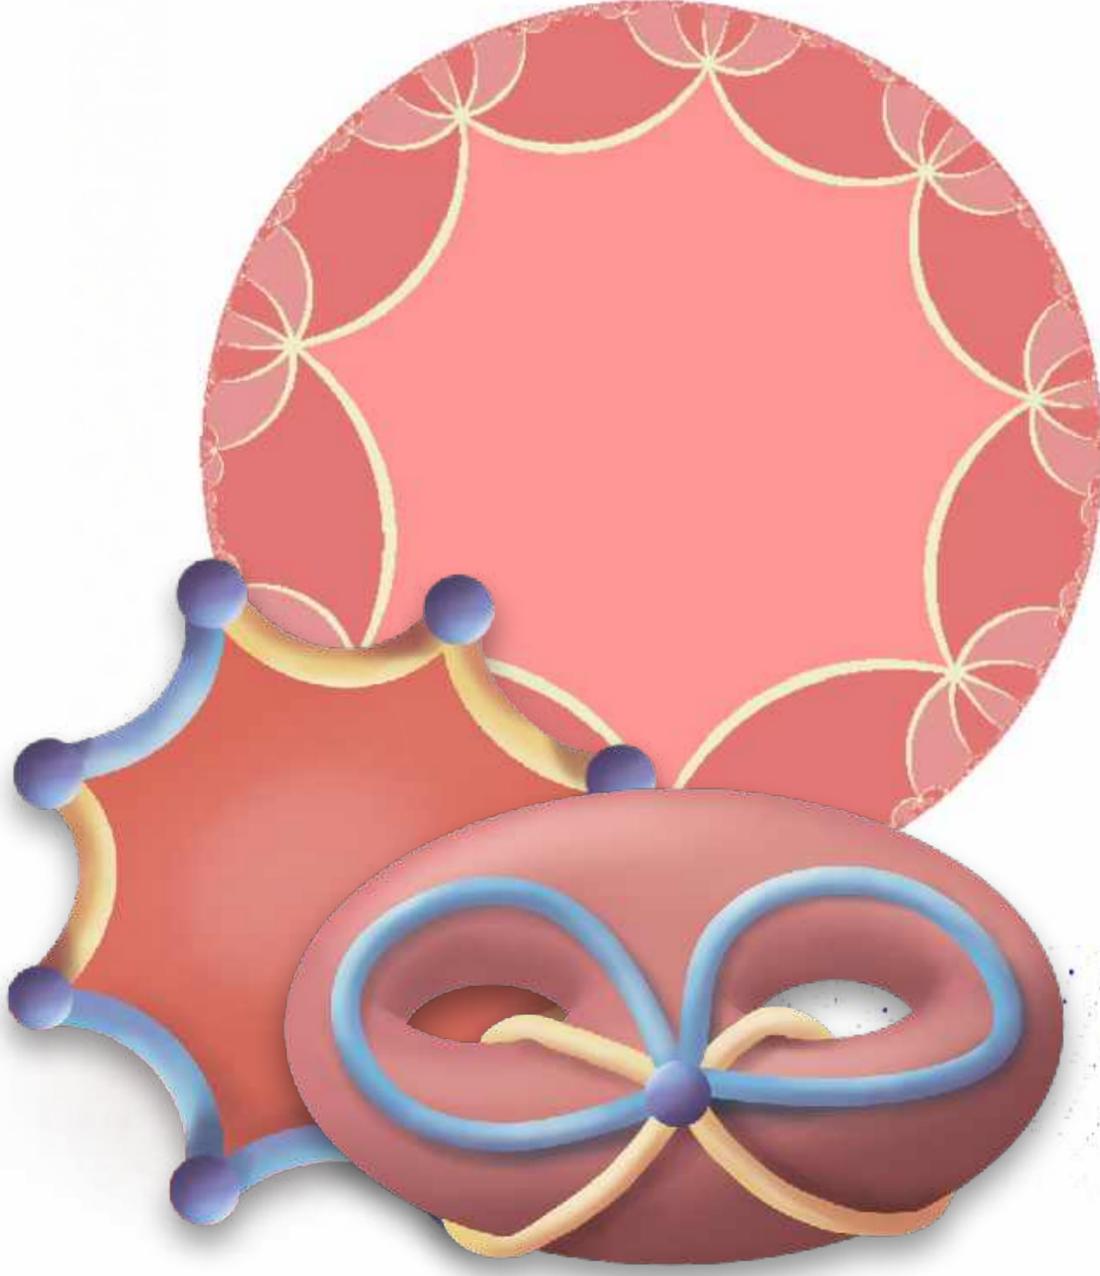
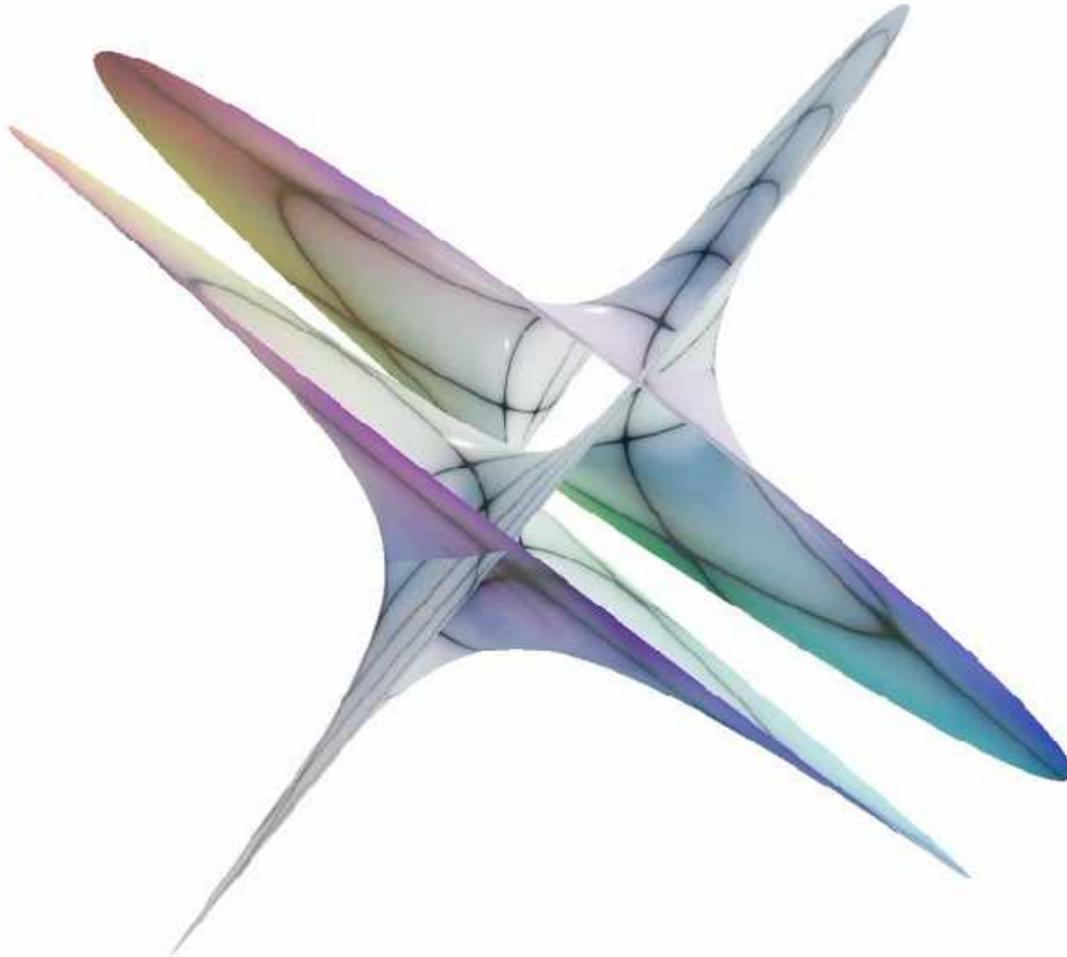
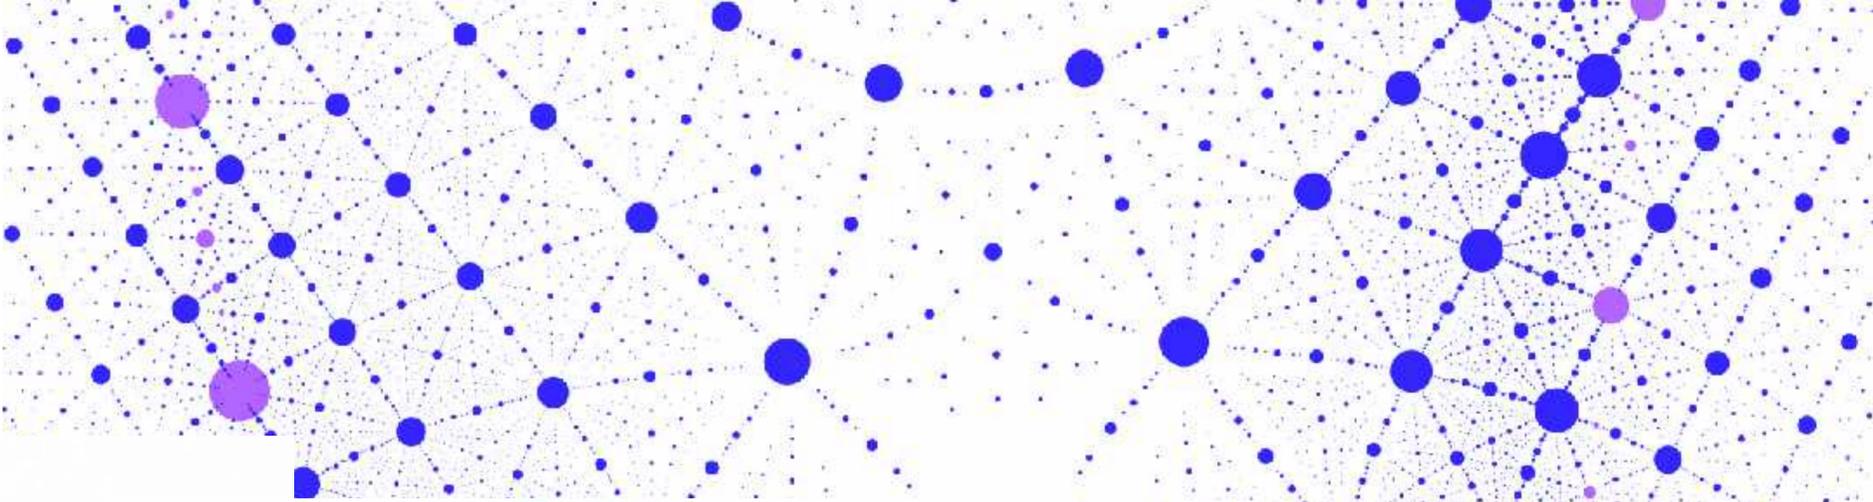


Hyperbolic geometry is everywhere!

Complex Analysis

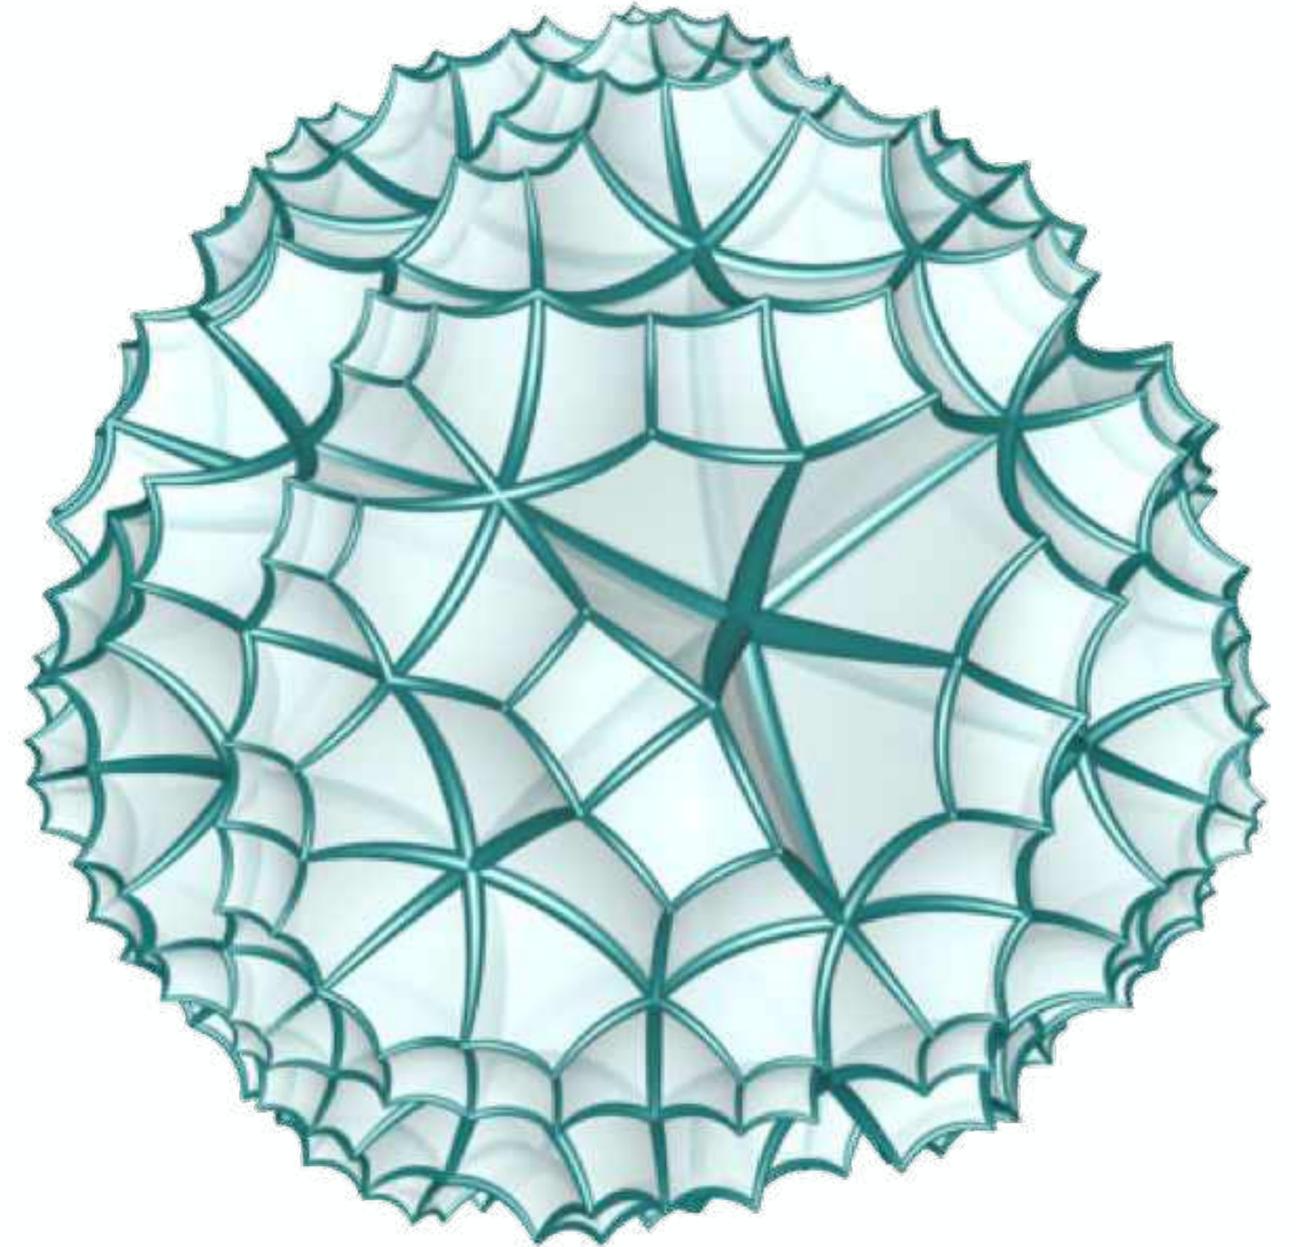
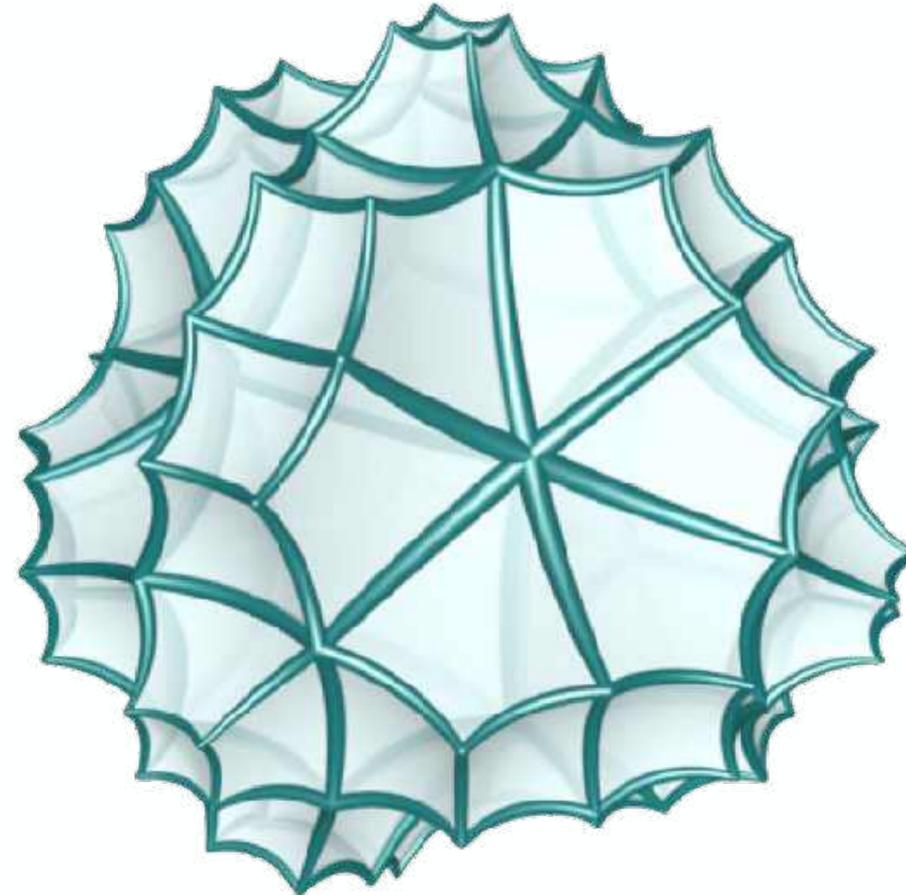
Topology

Number Theory



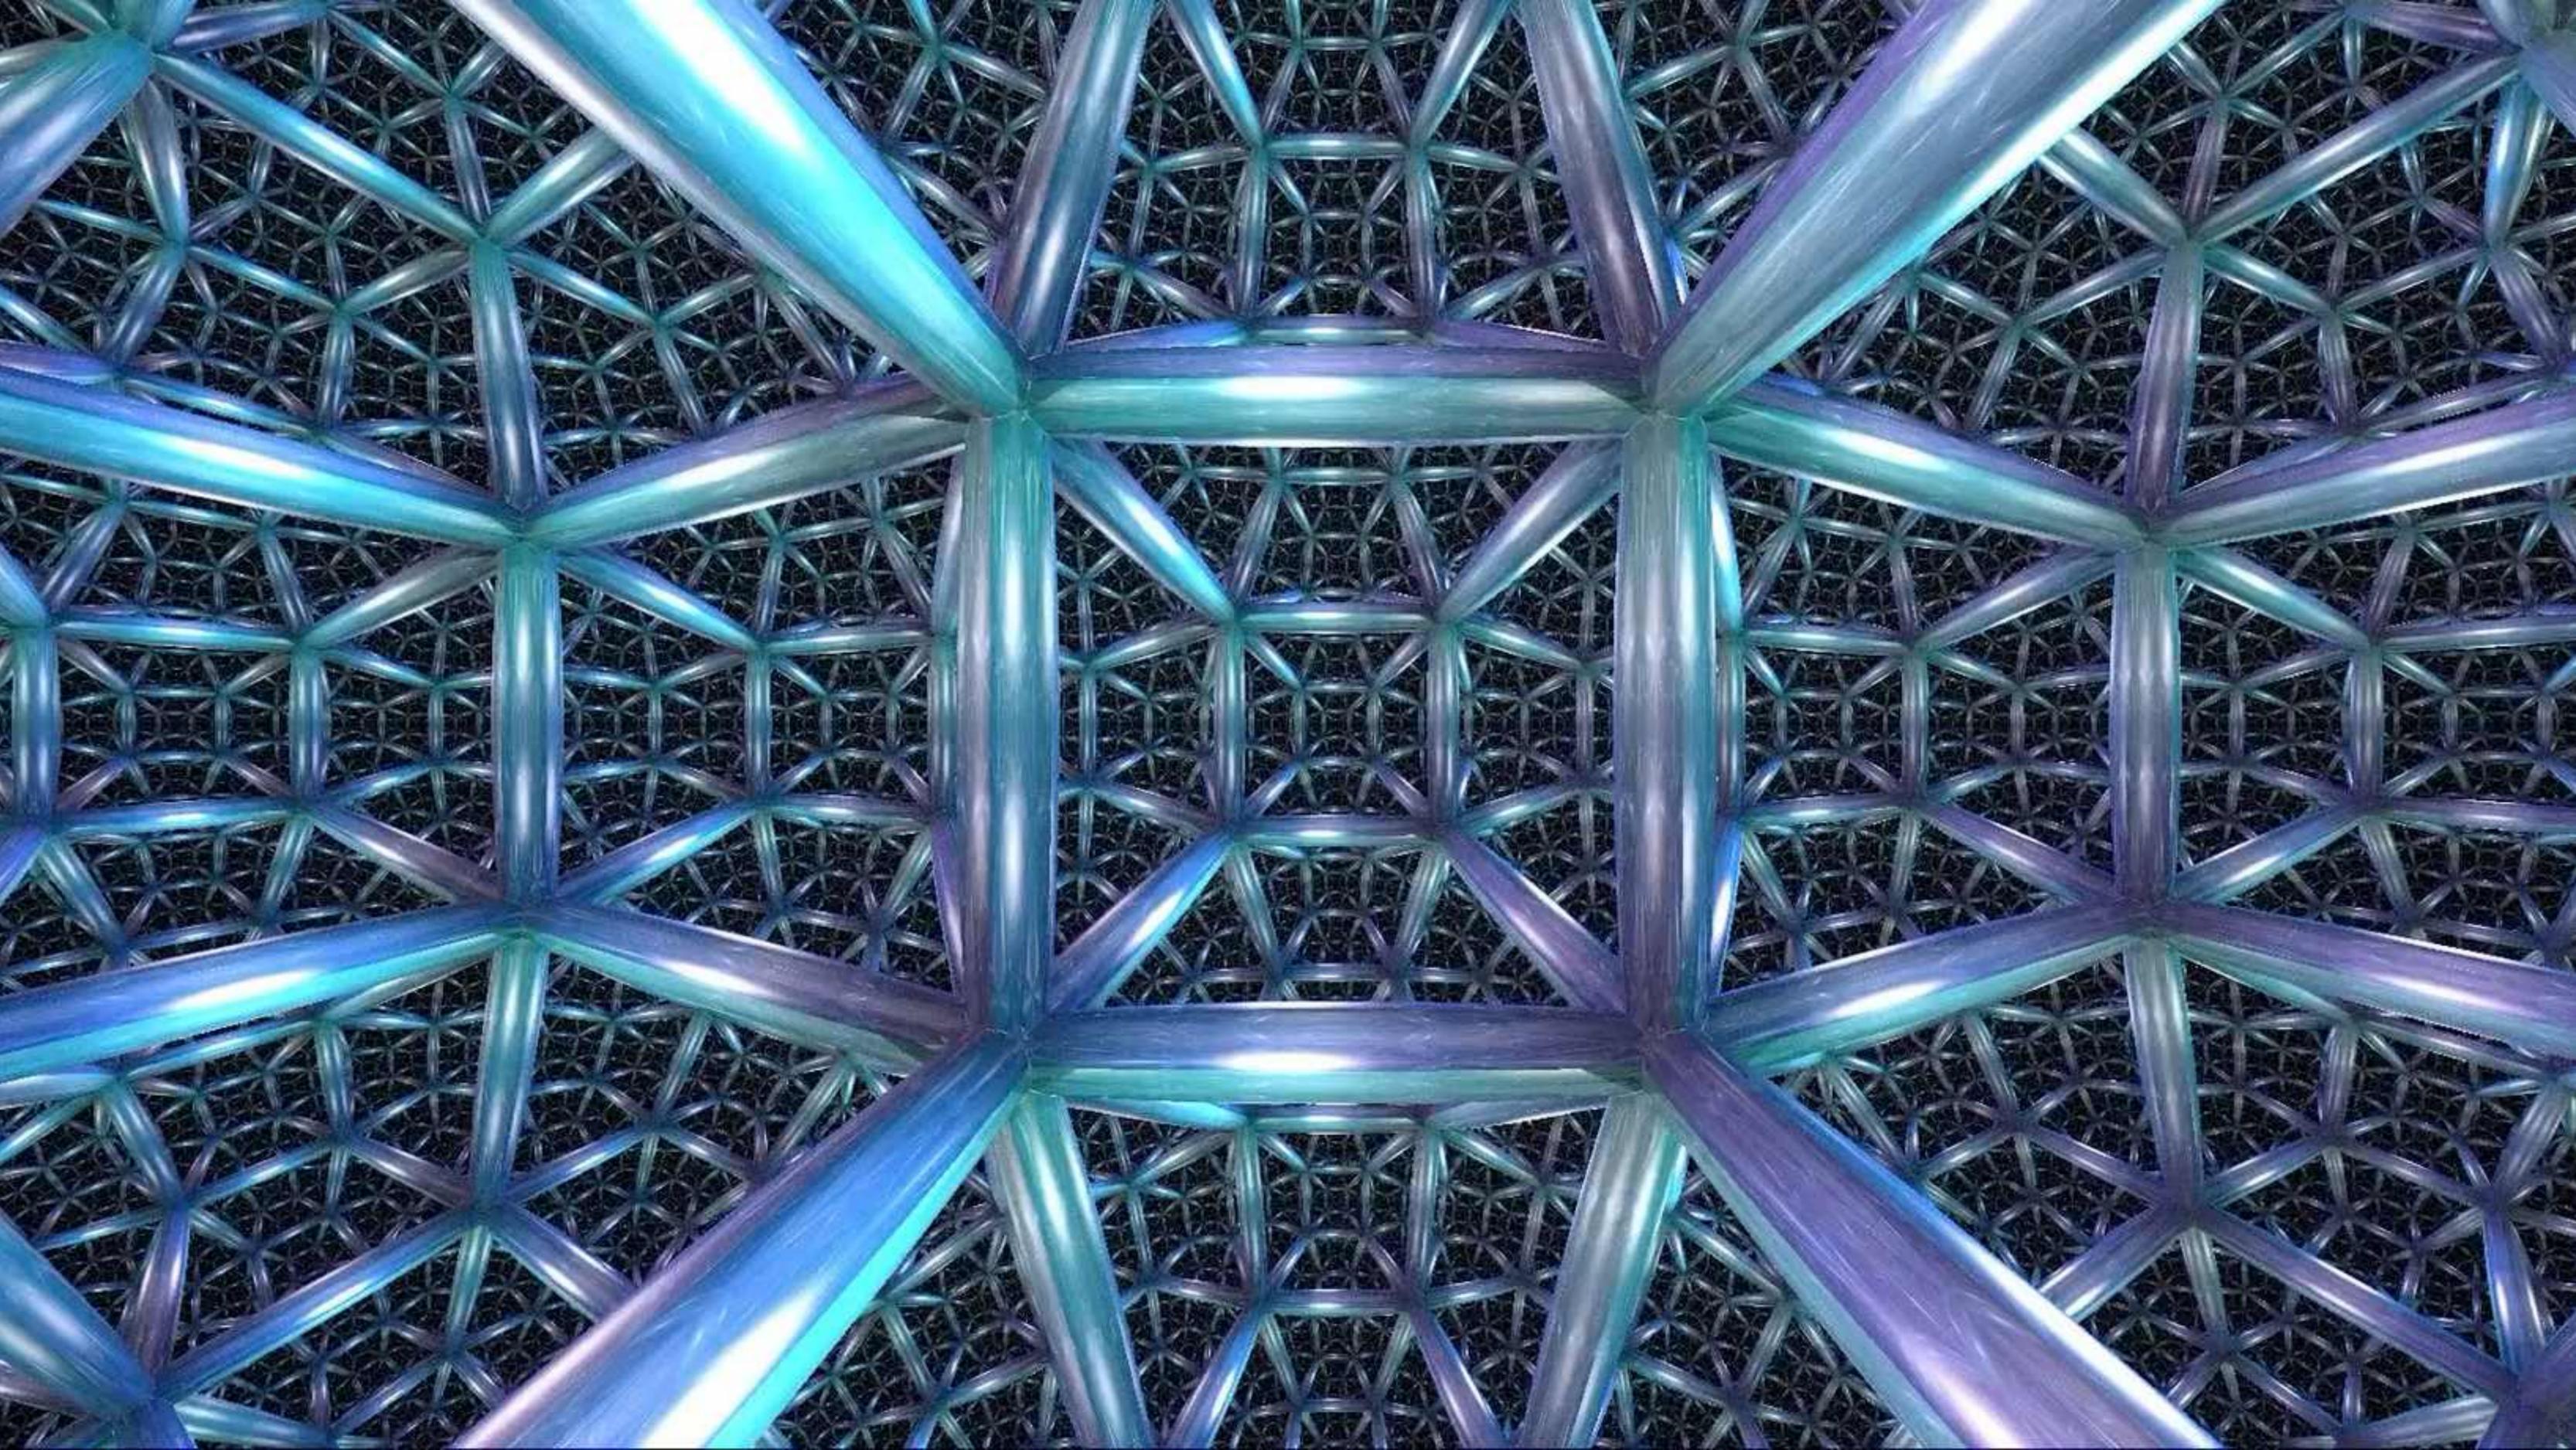


Negatively curved space is HUGE.
Most of the interesting questions
about 3-dimensional space turn out
to be questions about hyperbolic
geometry.



**Exponentially many cubes
stack around a central cube.**

Image by Roice Nelson

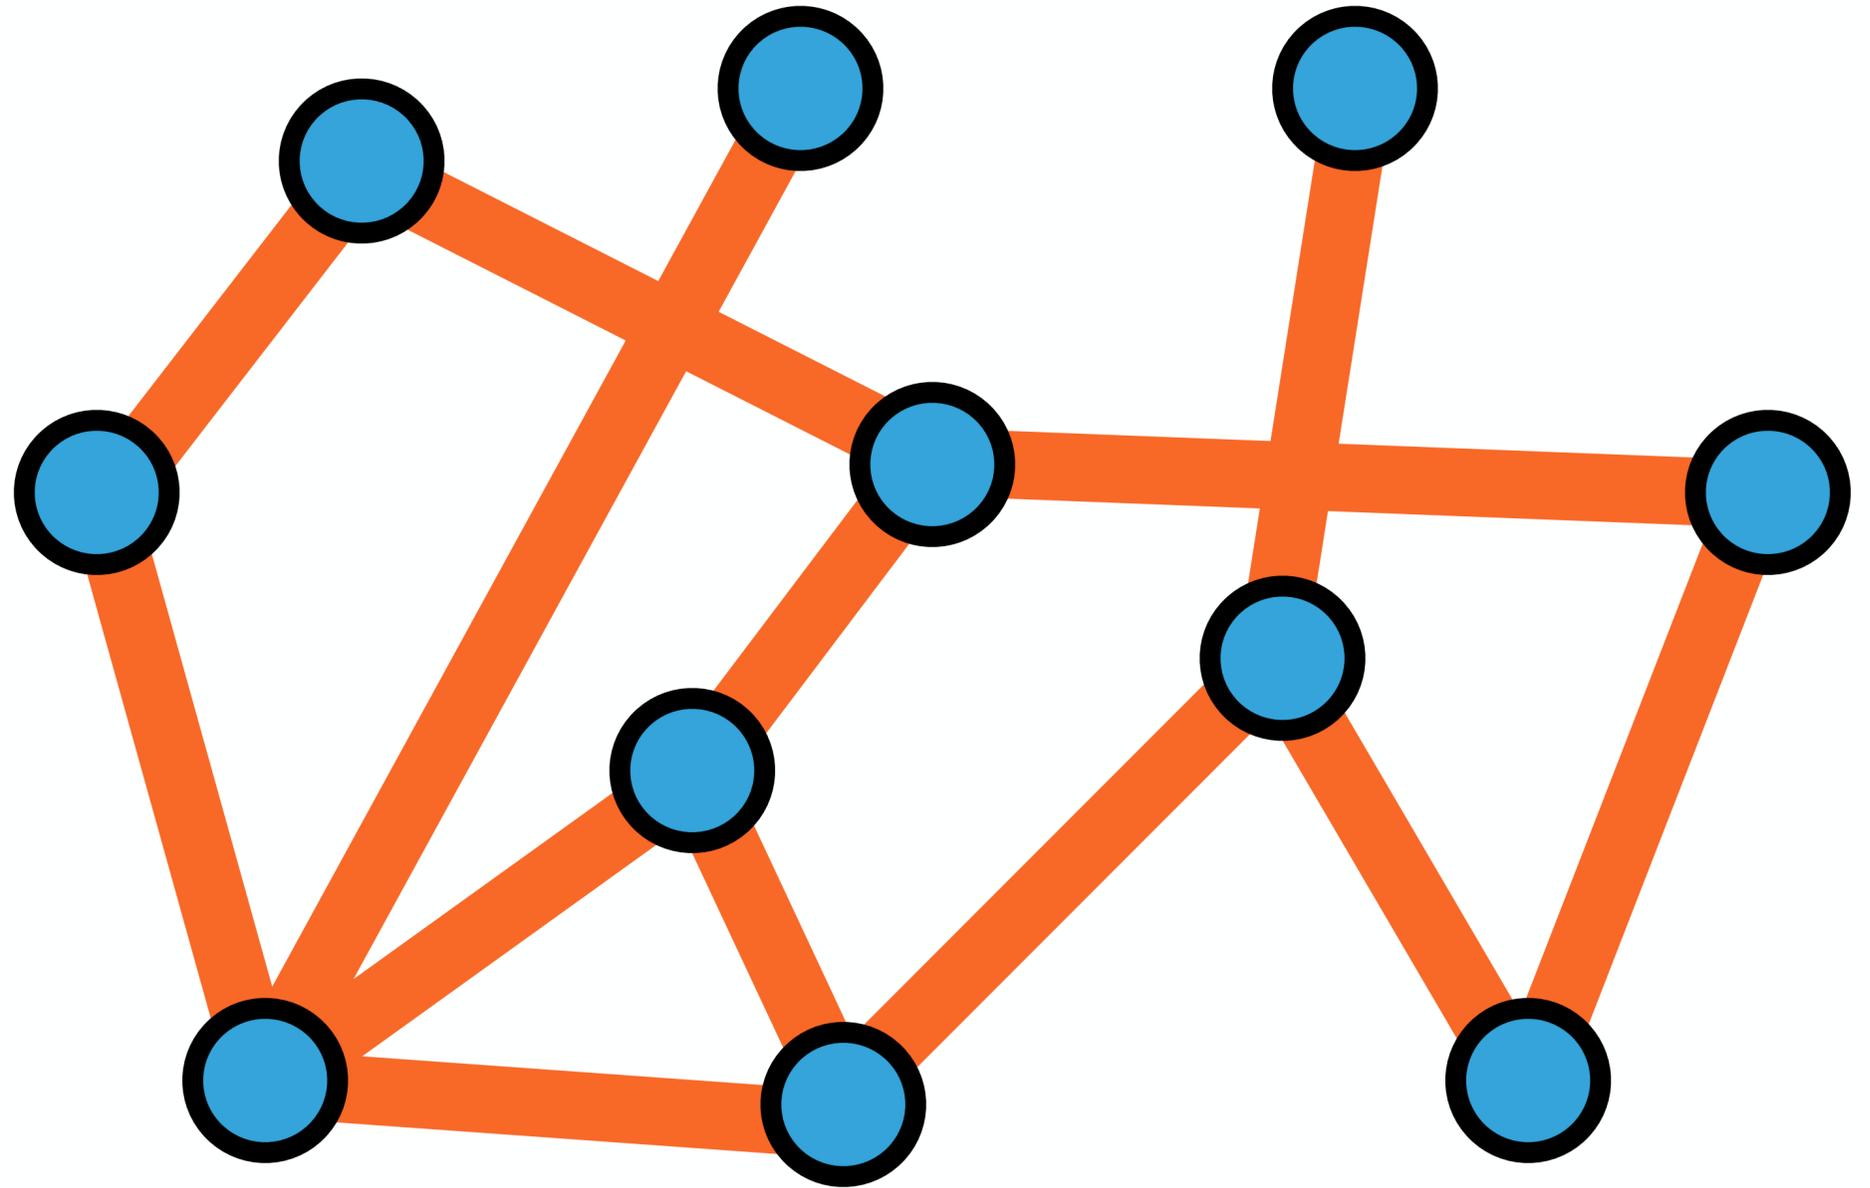


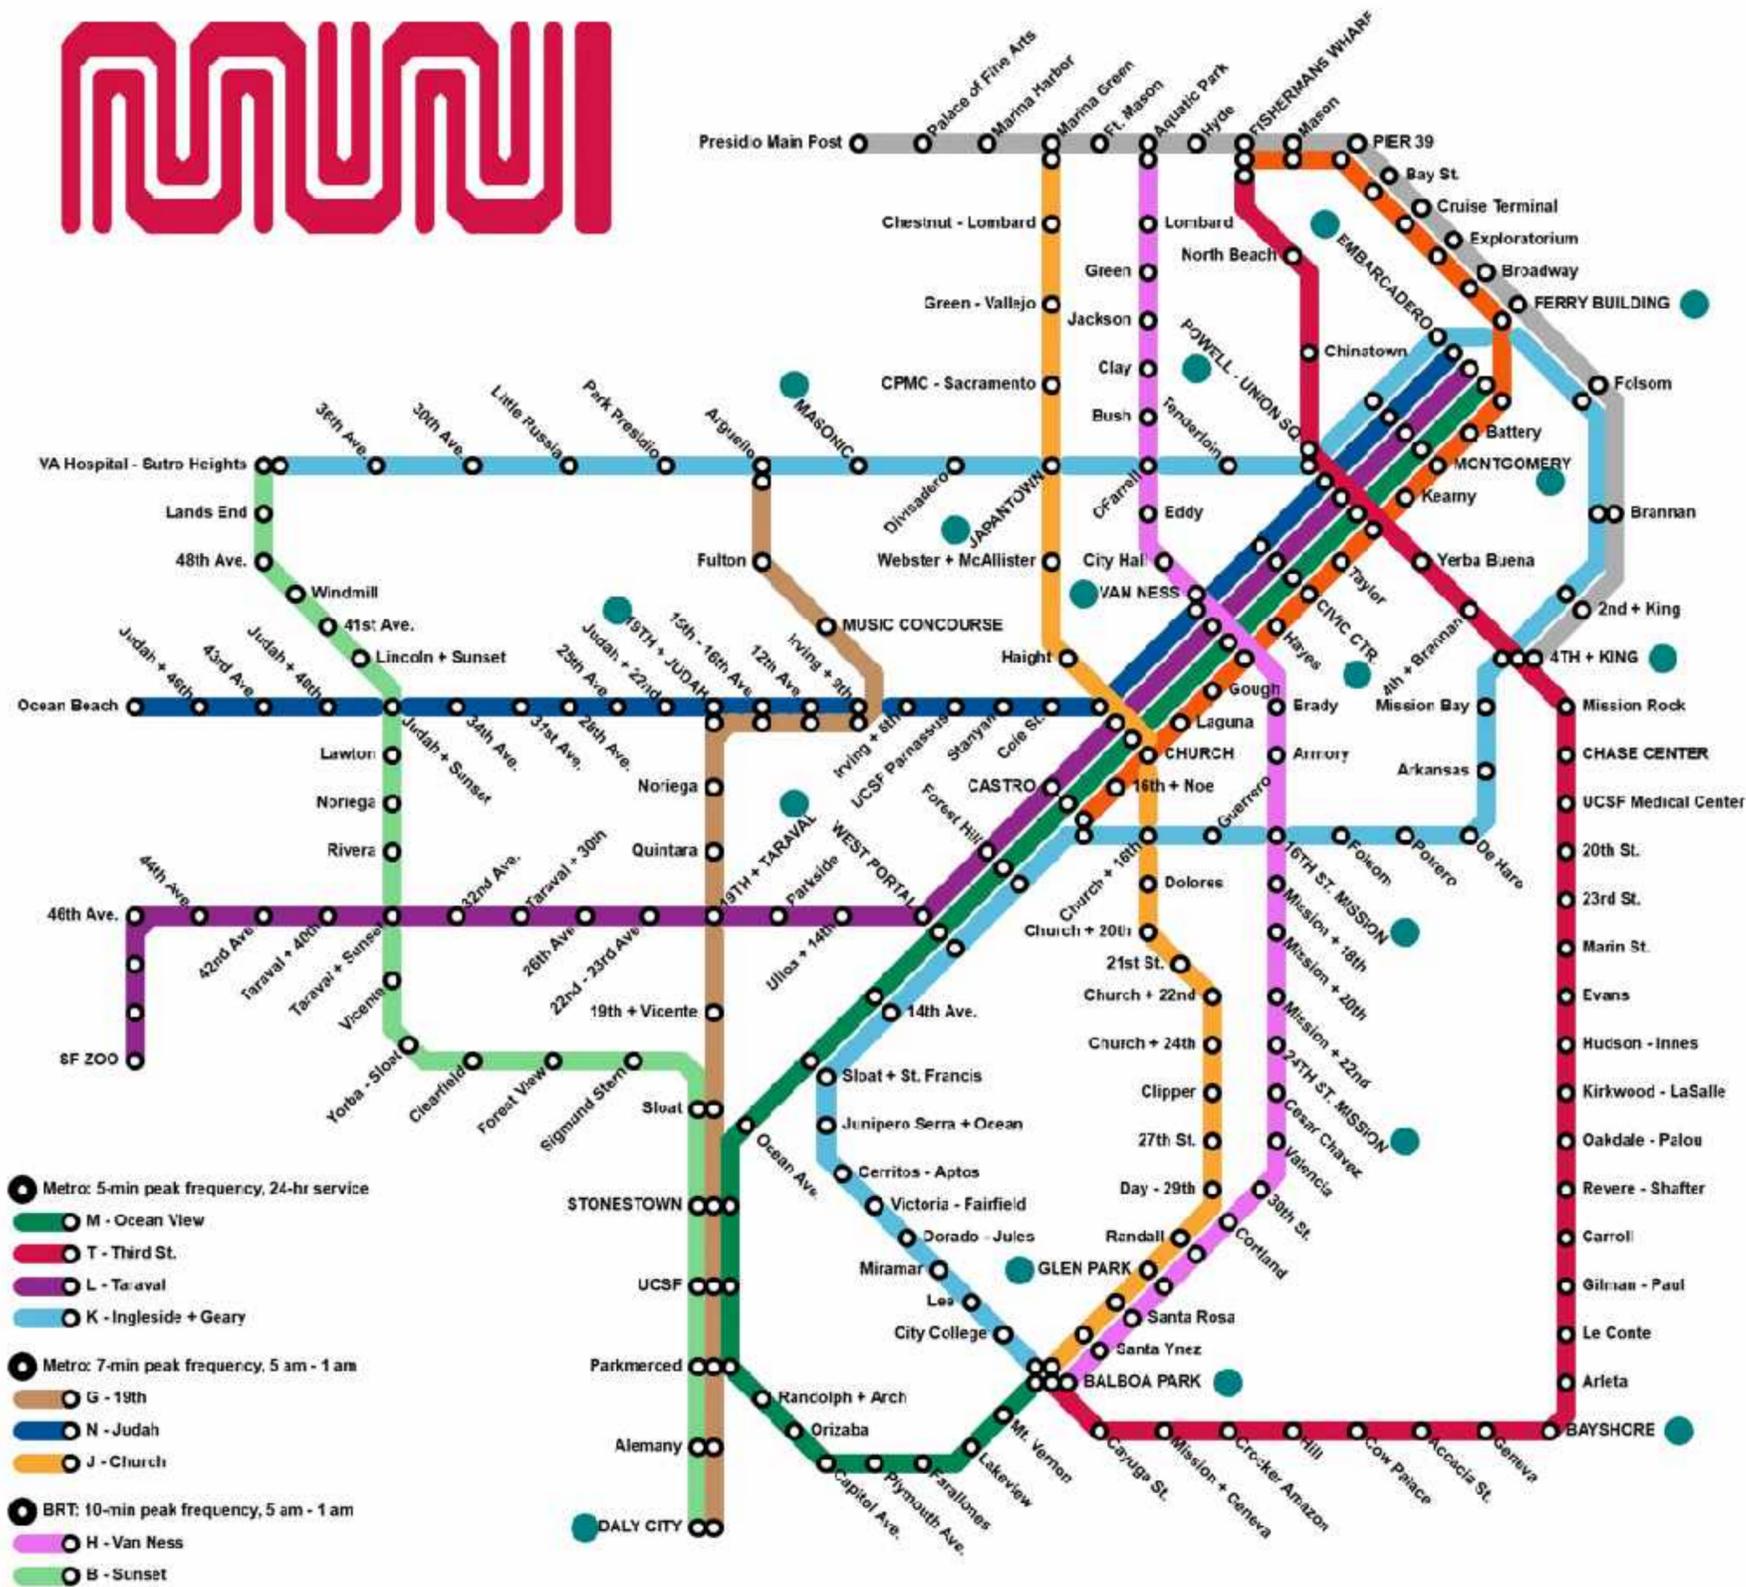
Untangling Webs

The geometry of graph embeddings.

Data and Graphs

Vertices are data
Edges are relations





The Muni Graph

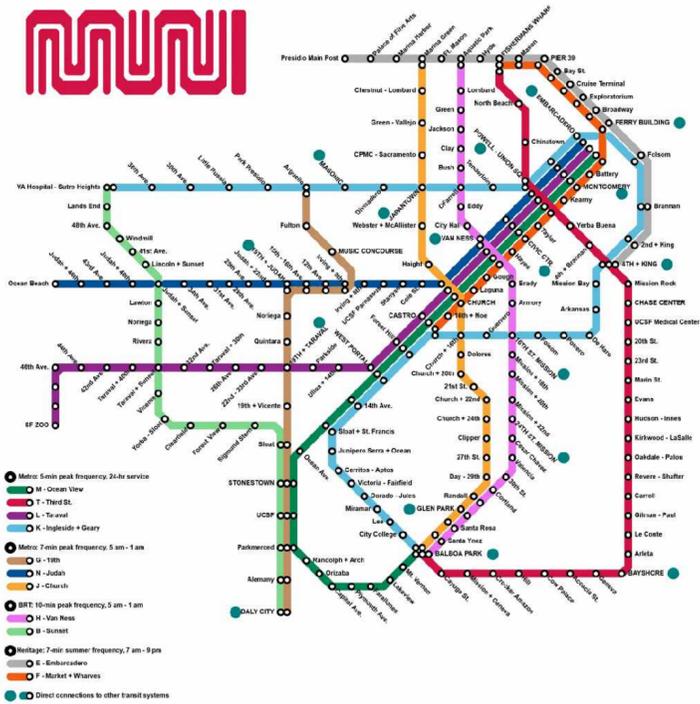
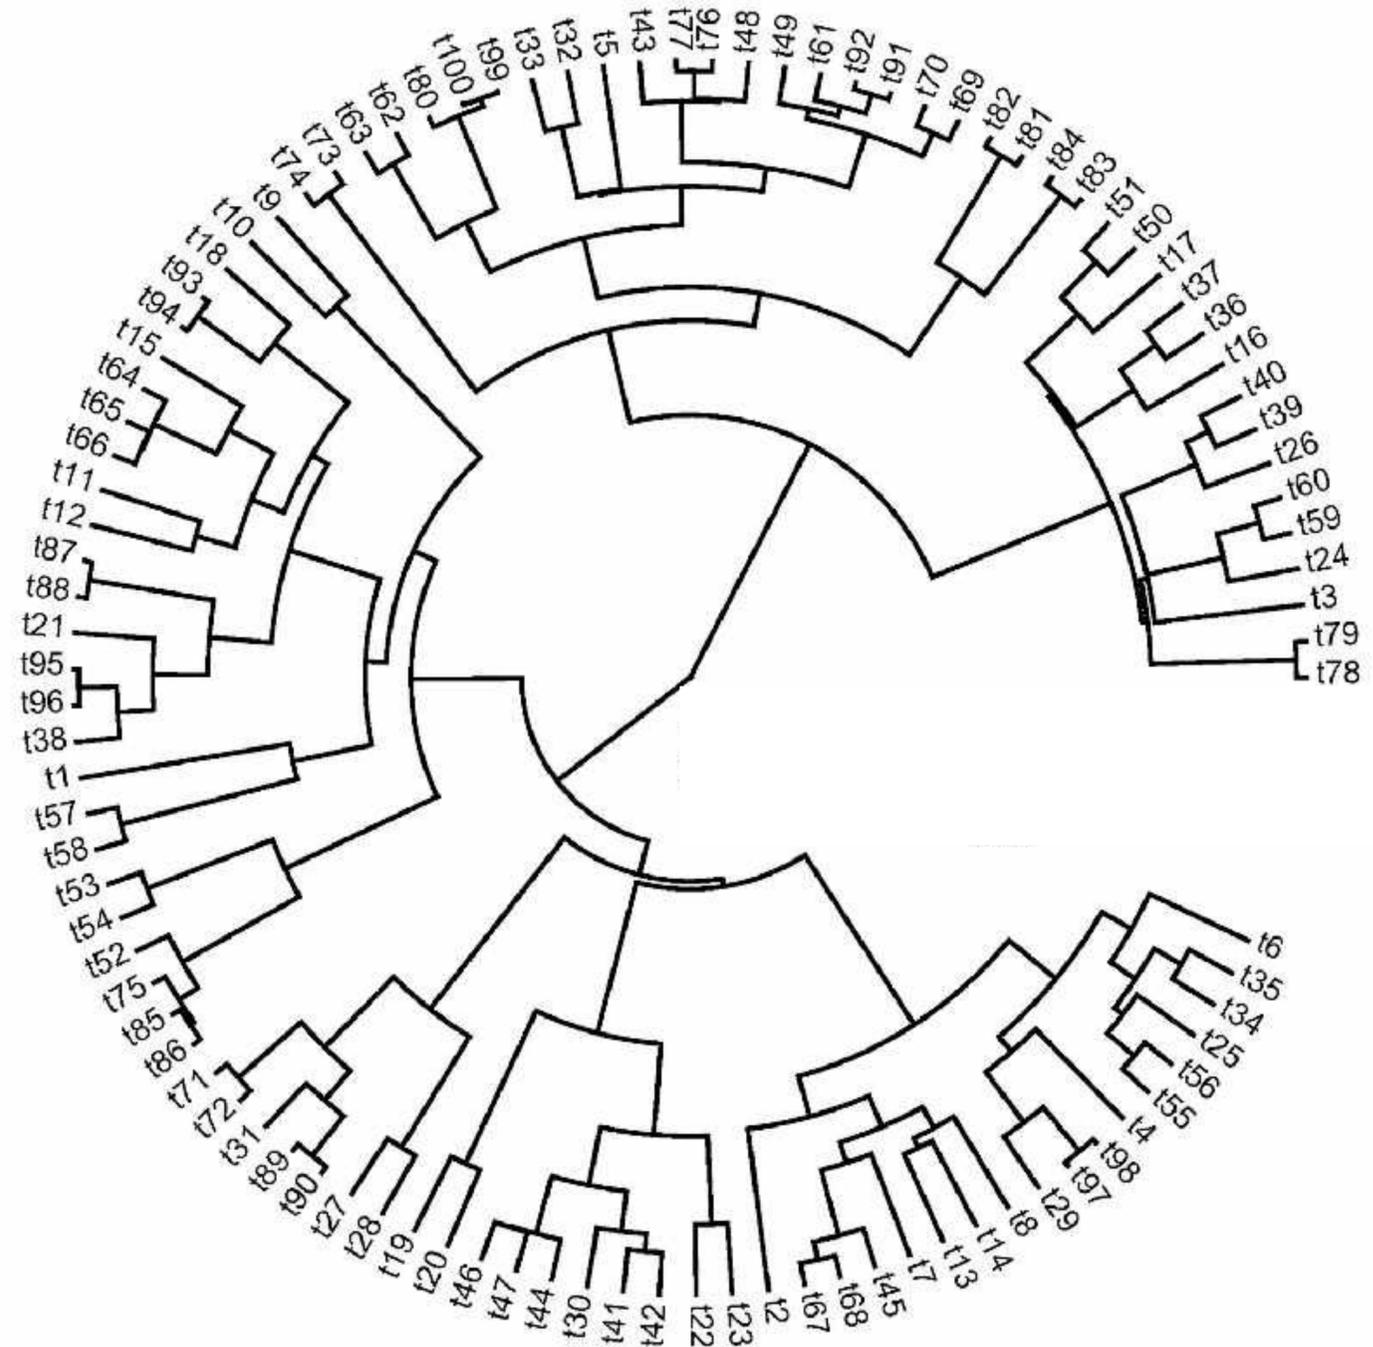
Vertices: Bus/Tram Stops

Edges: Bus/Tram Routes

Phylogenetic Trees

Vertices: Species or Genes

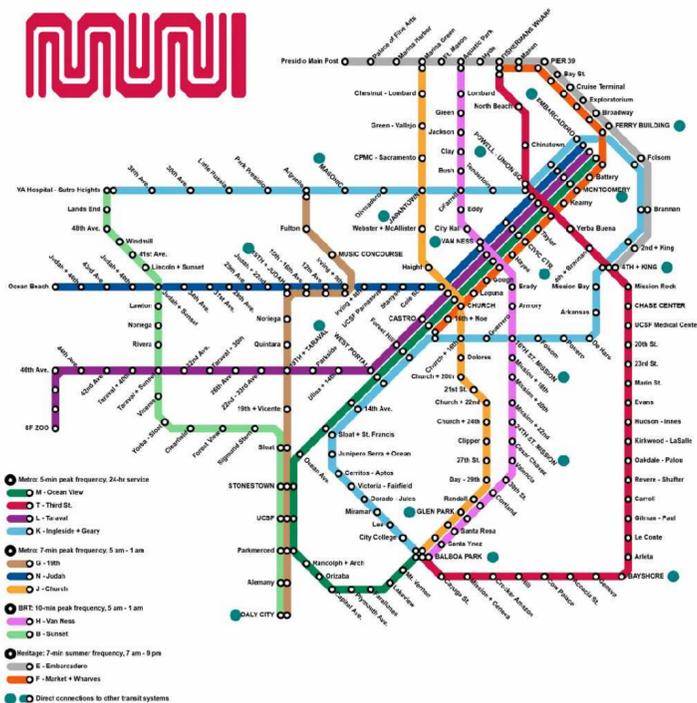
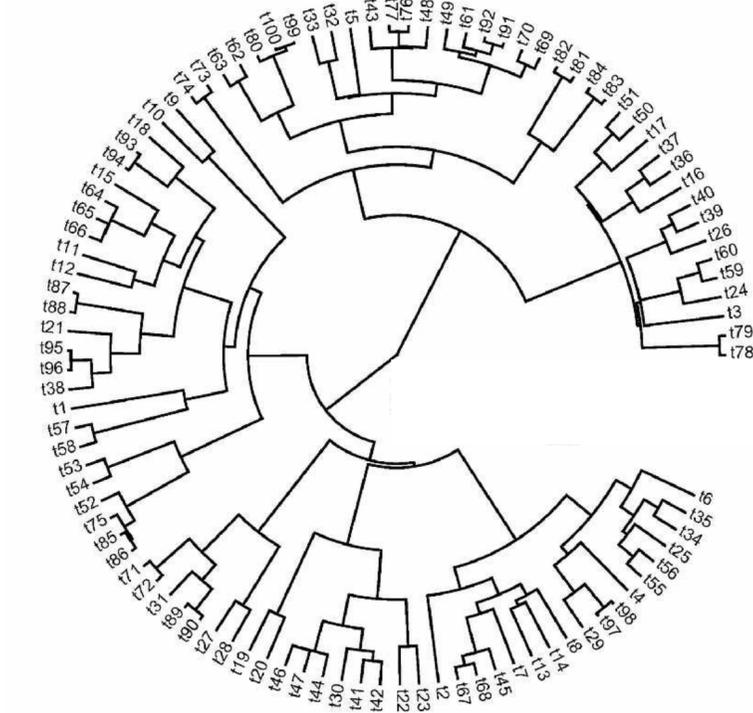
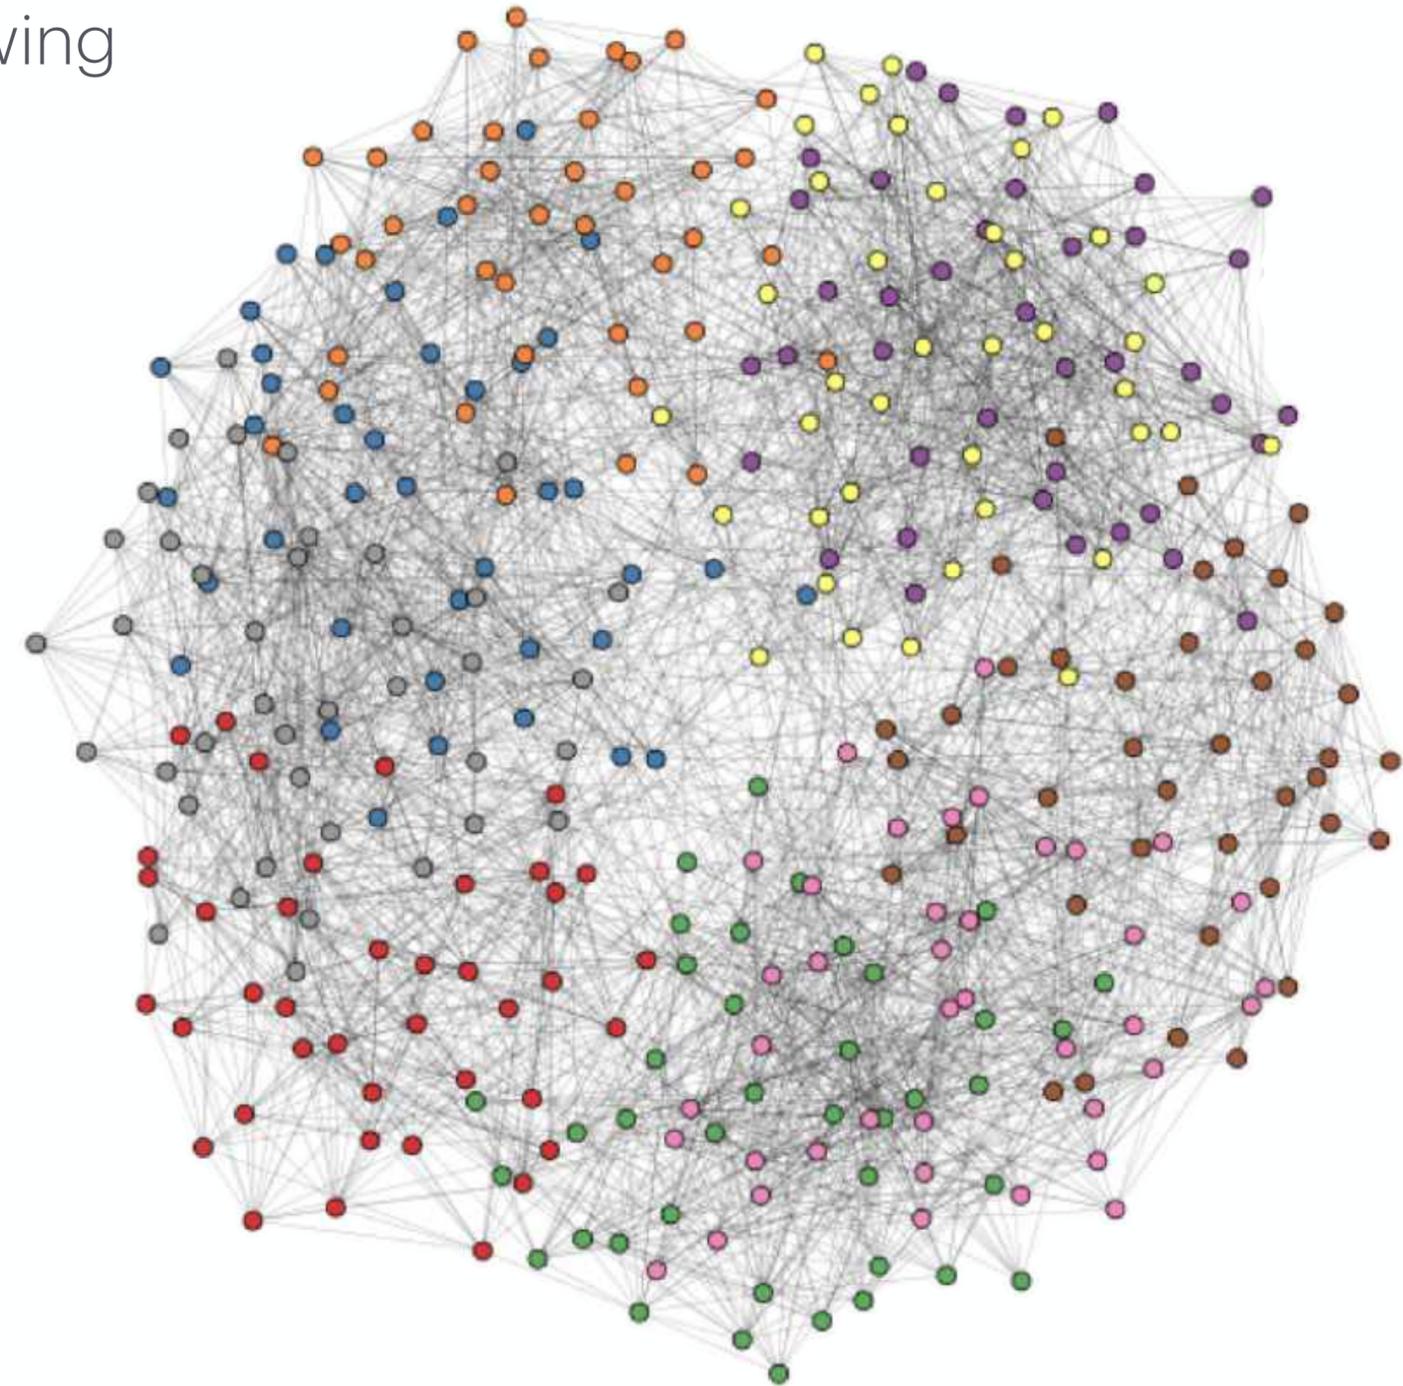
Edges: Ancestry



Social Networks

Vertices: Users

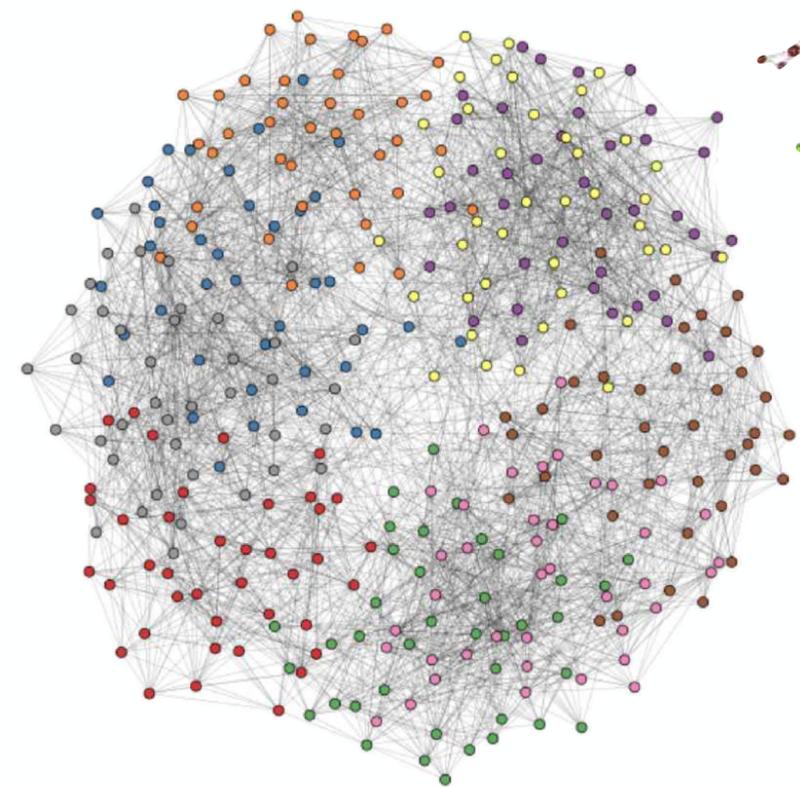
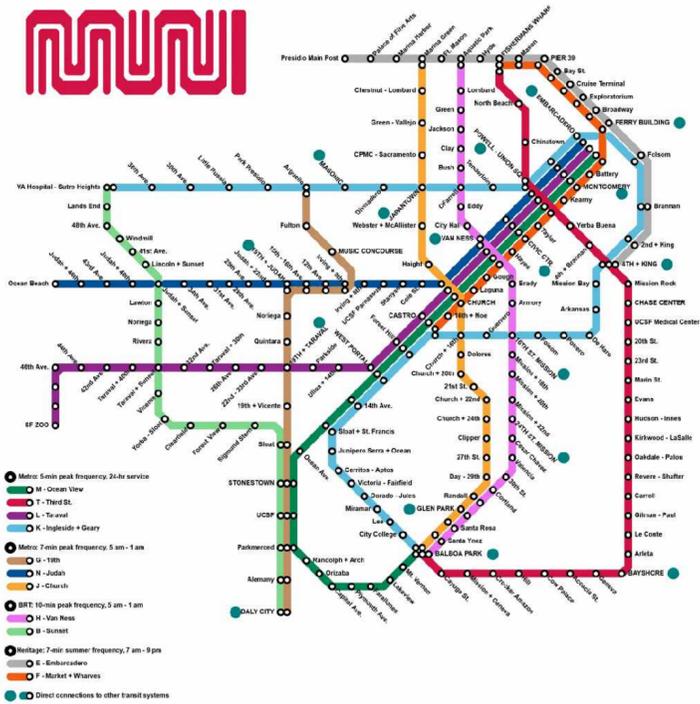
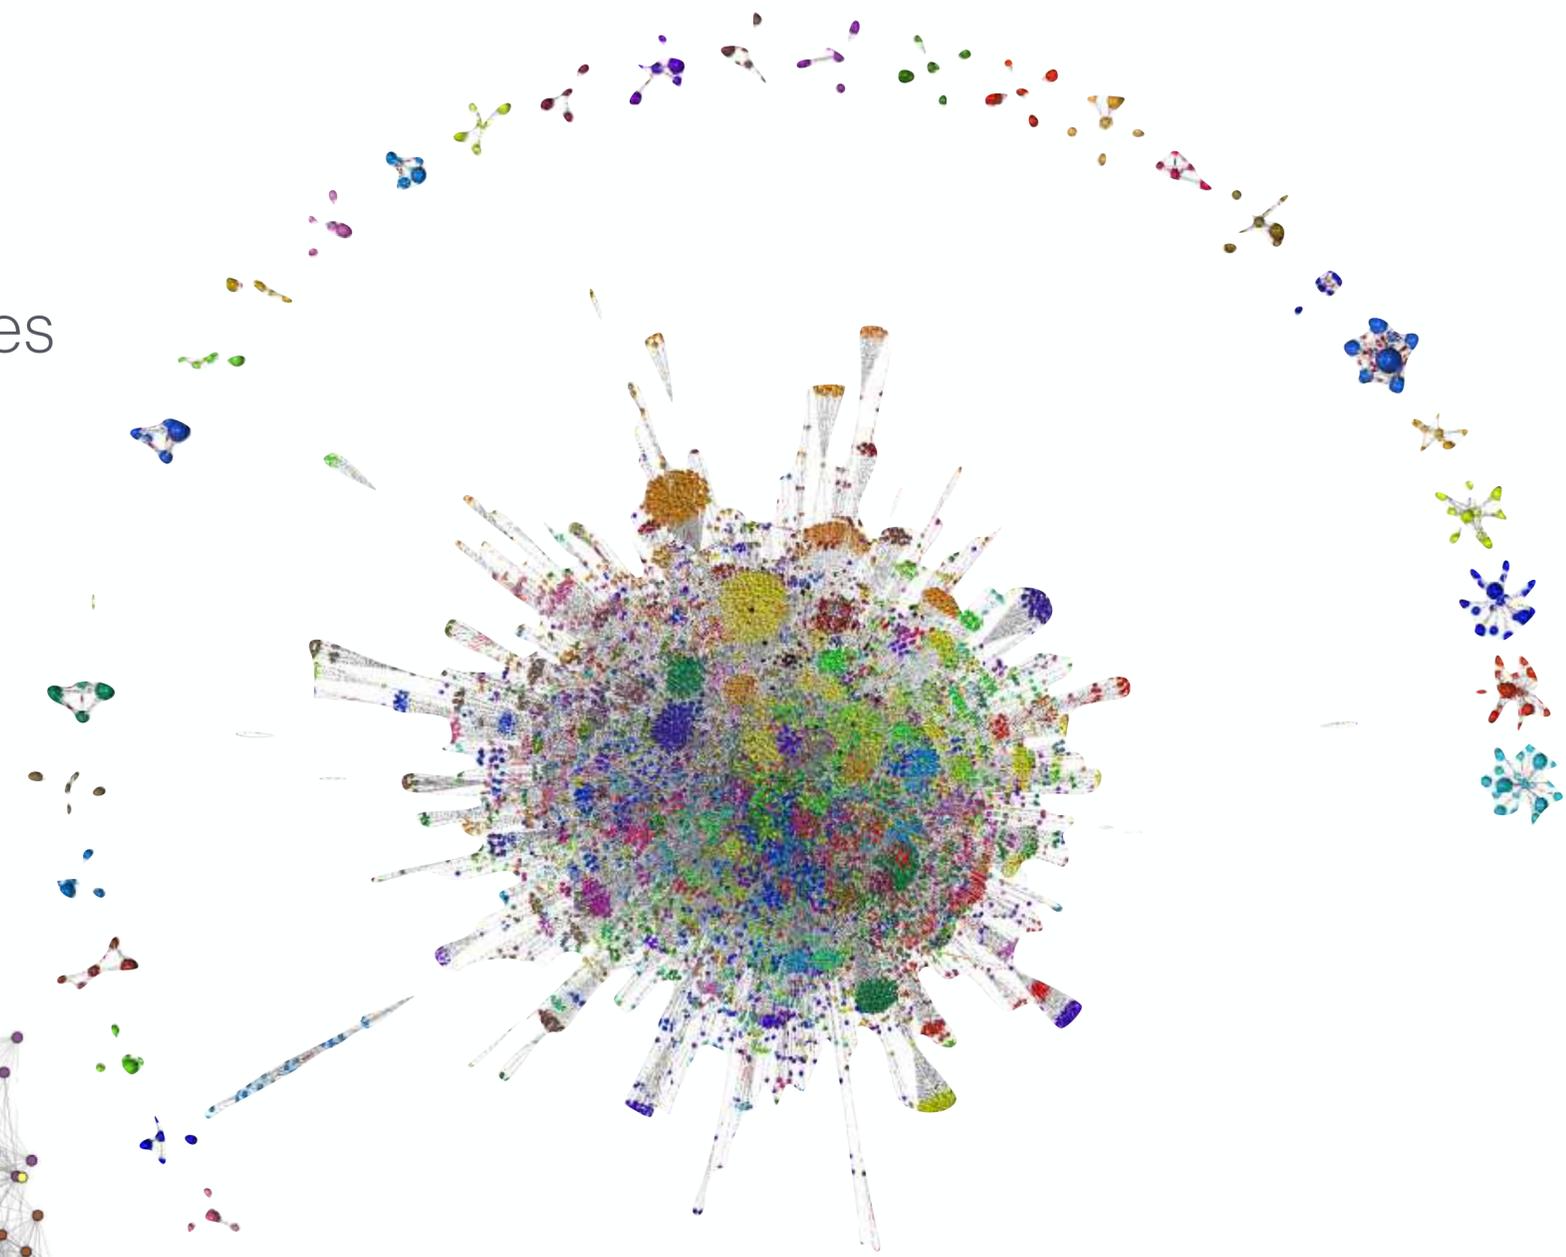
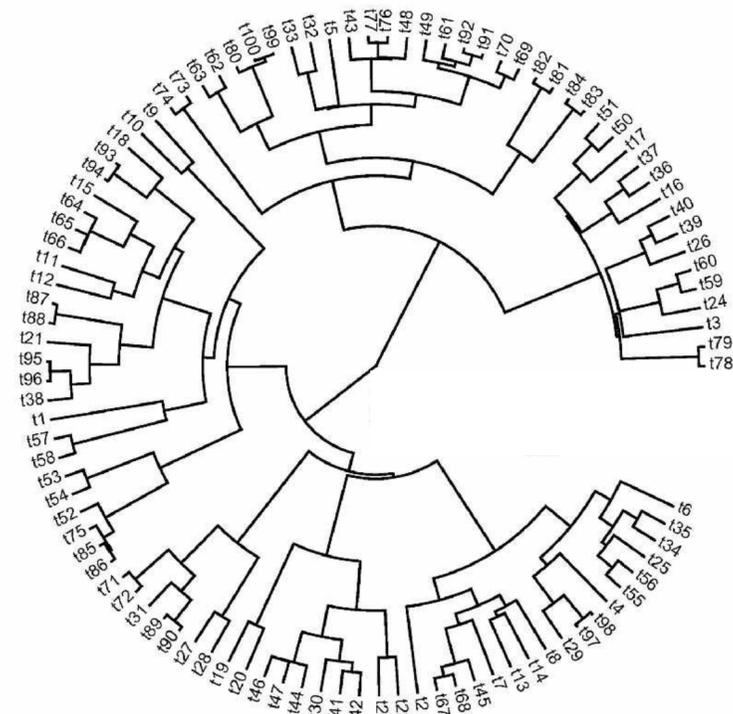
Edges: Following

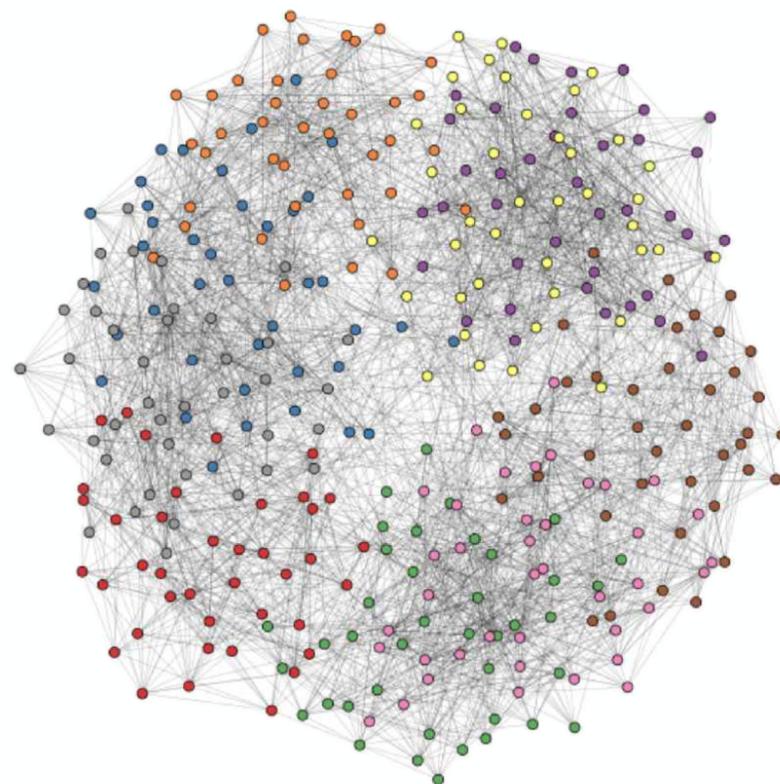
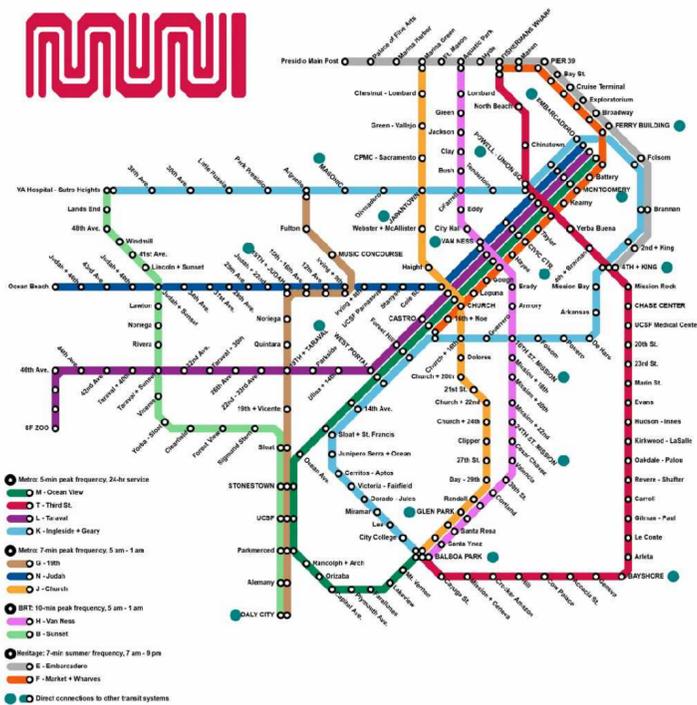
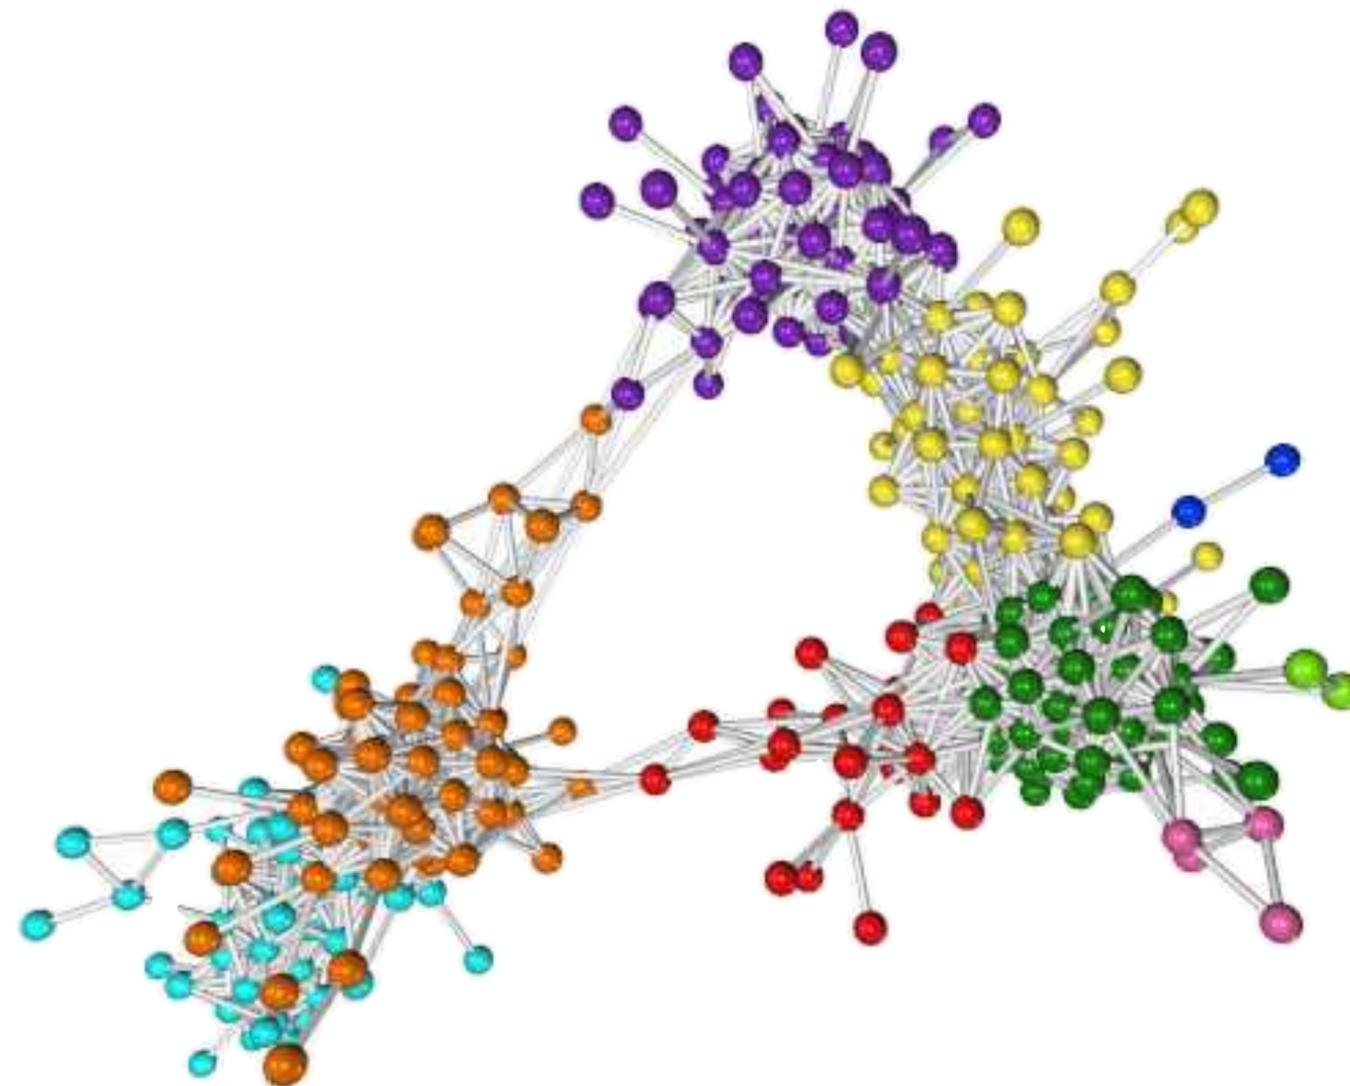
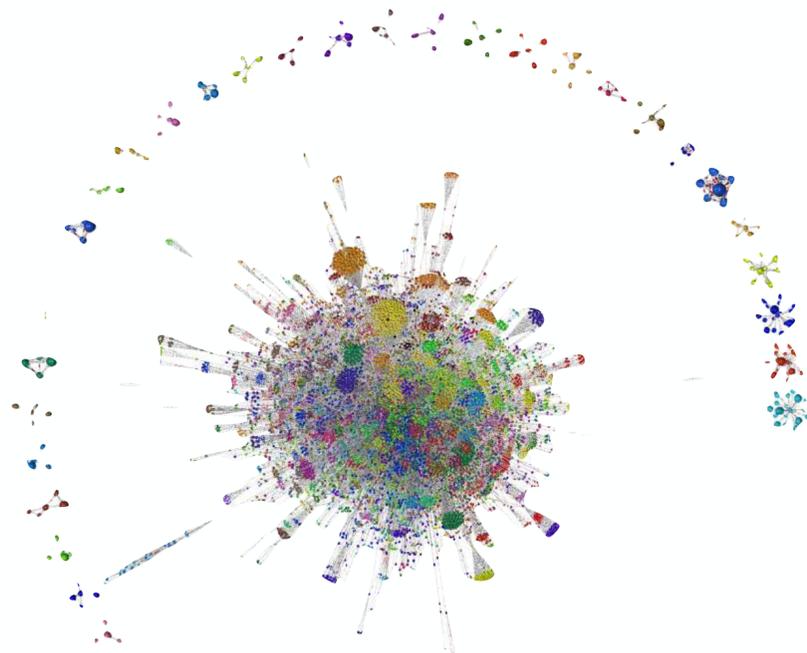
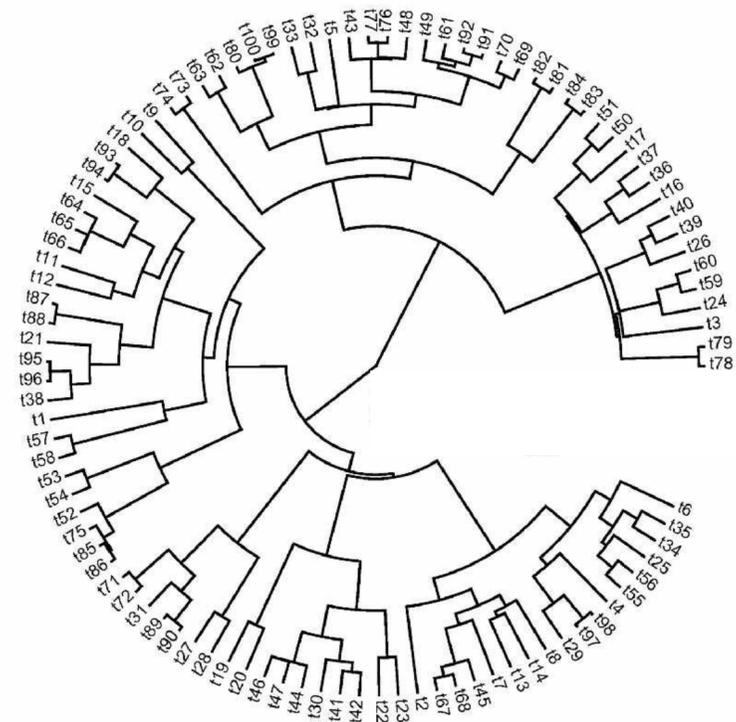


Reddit

Vertices: Subreddits

Edges: Crossreferences

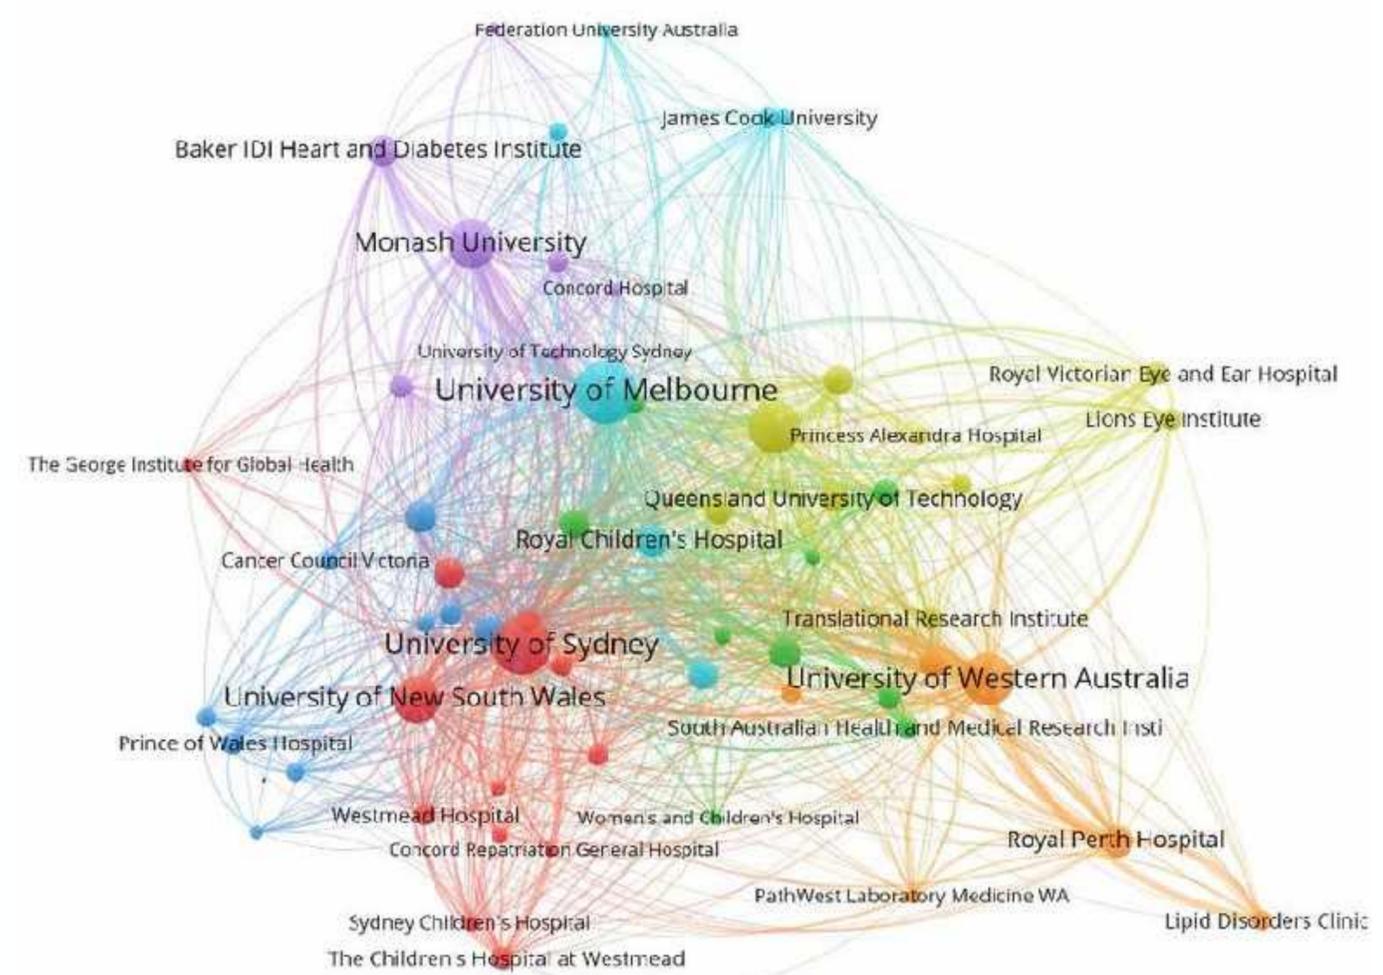
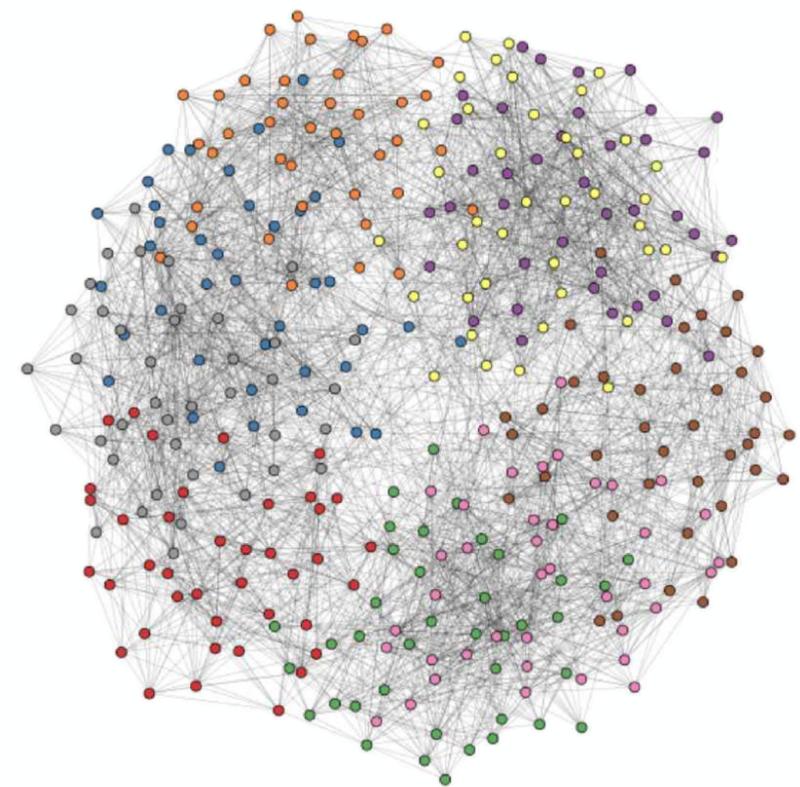
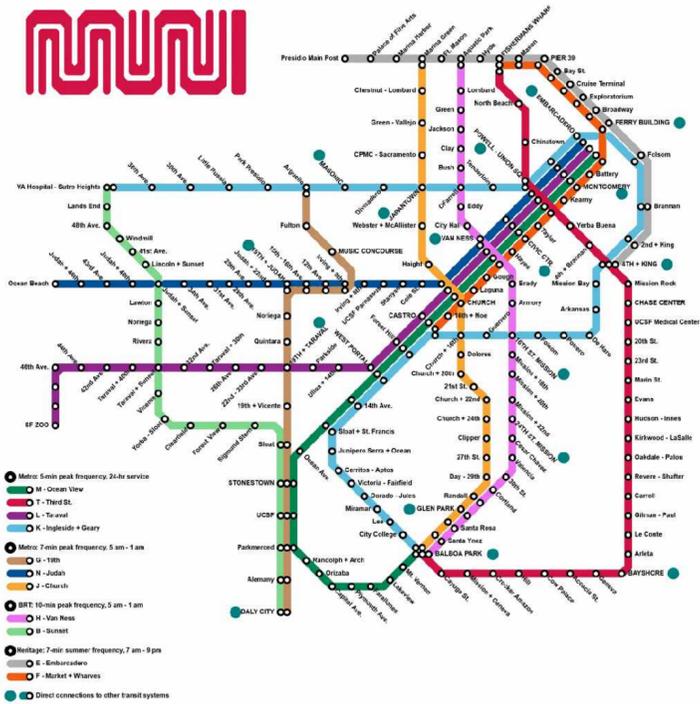
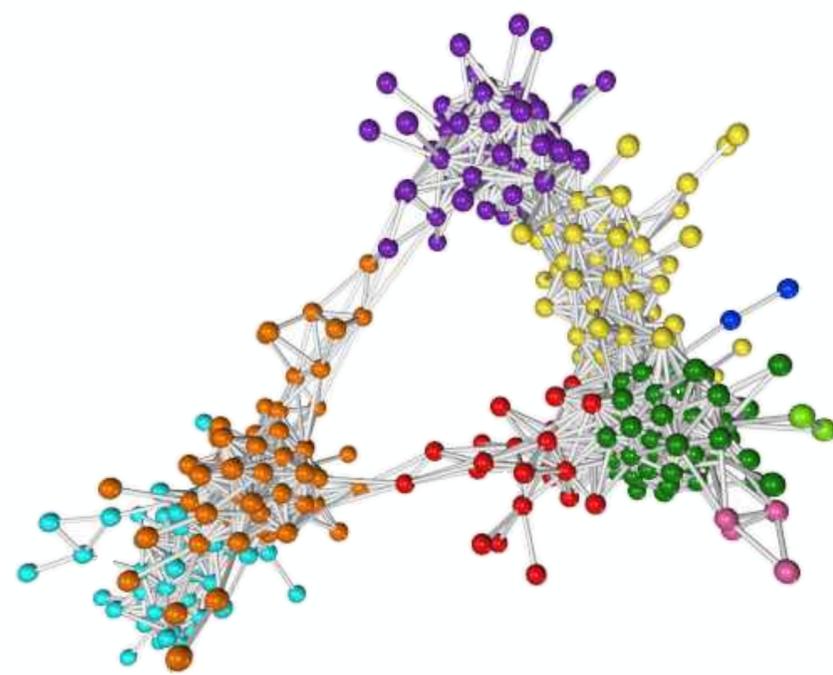
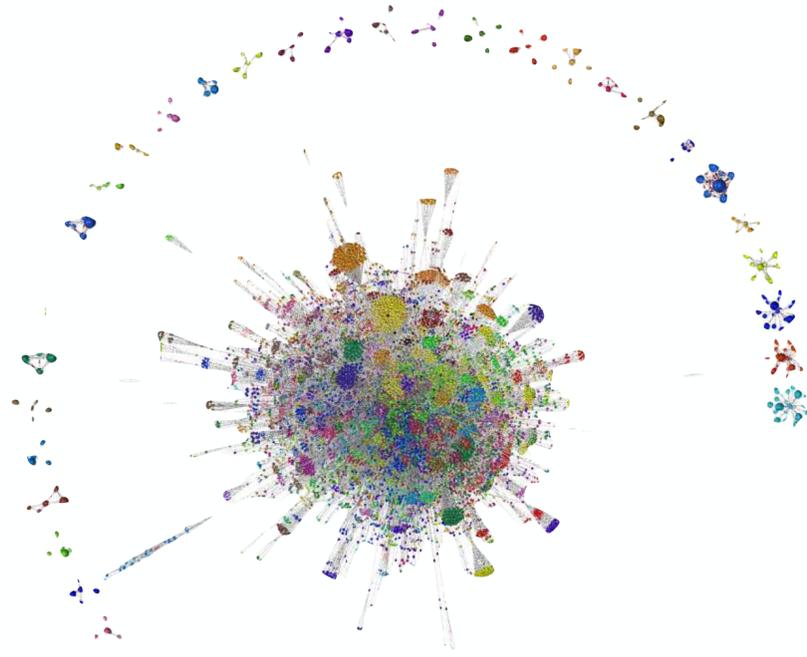
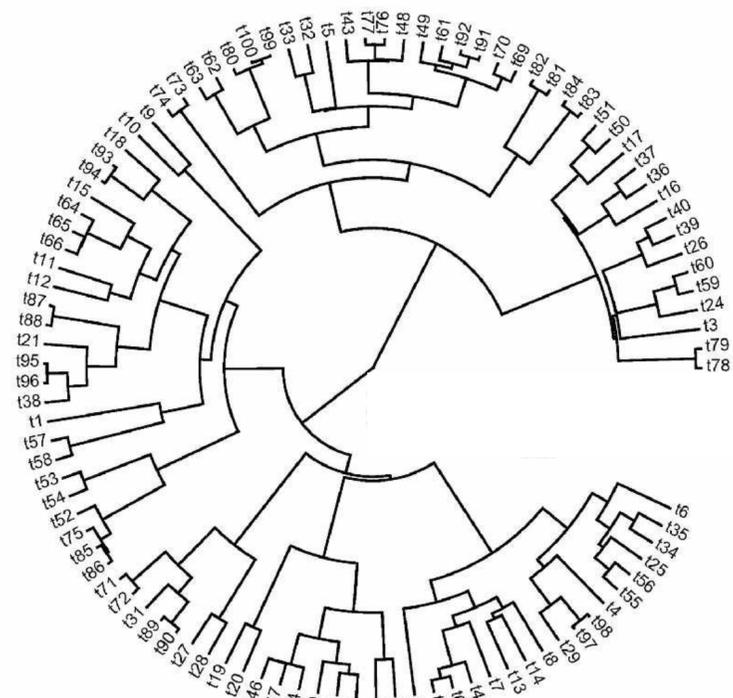


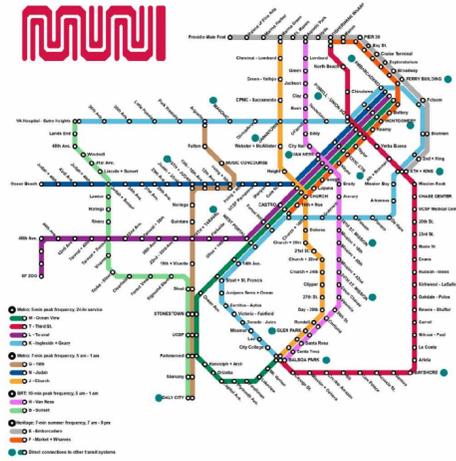


Biochemistry

Vertices: Mouse Genes

Edges: Co-expression

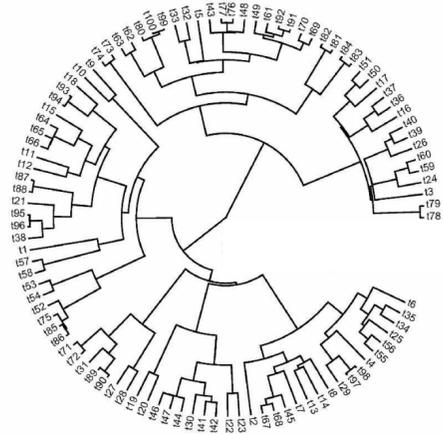




Route "Length" :

Travel Time

Some connections are stronger than others.

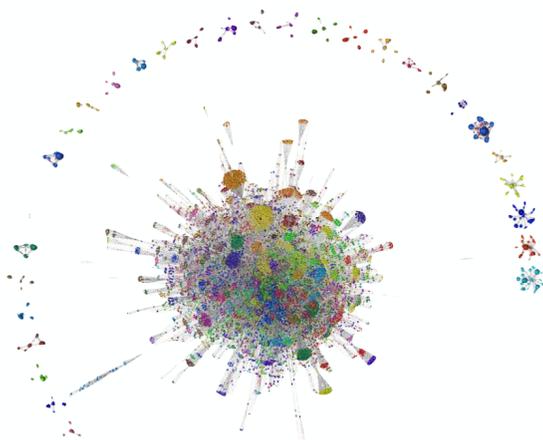


Genetic "Distance" :

Number of base pair changes



If we have a 'weight' on each edge, we can treat it as a length. **This gives the graph geometry!**

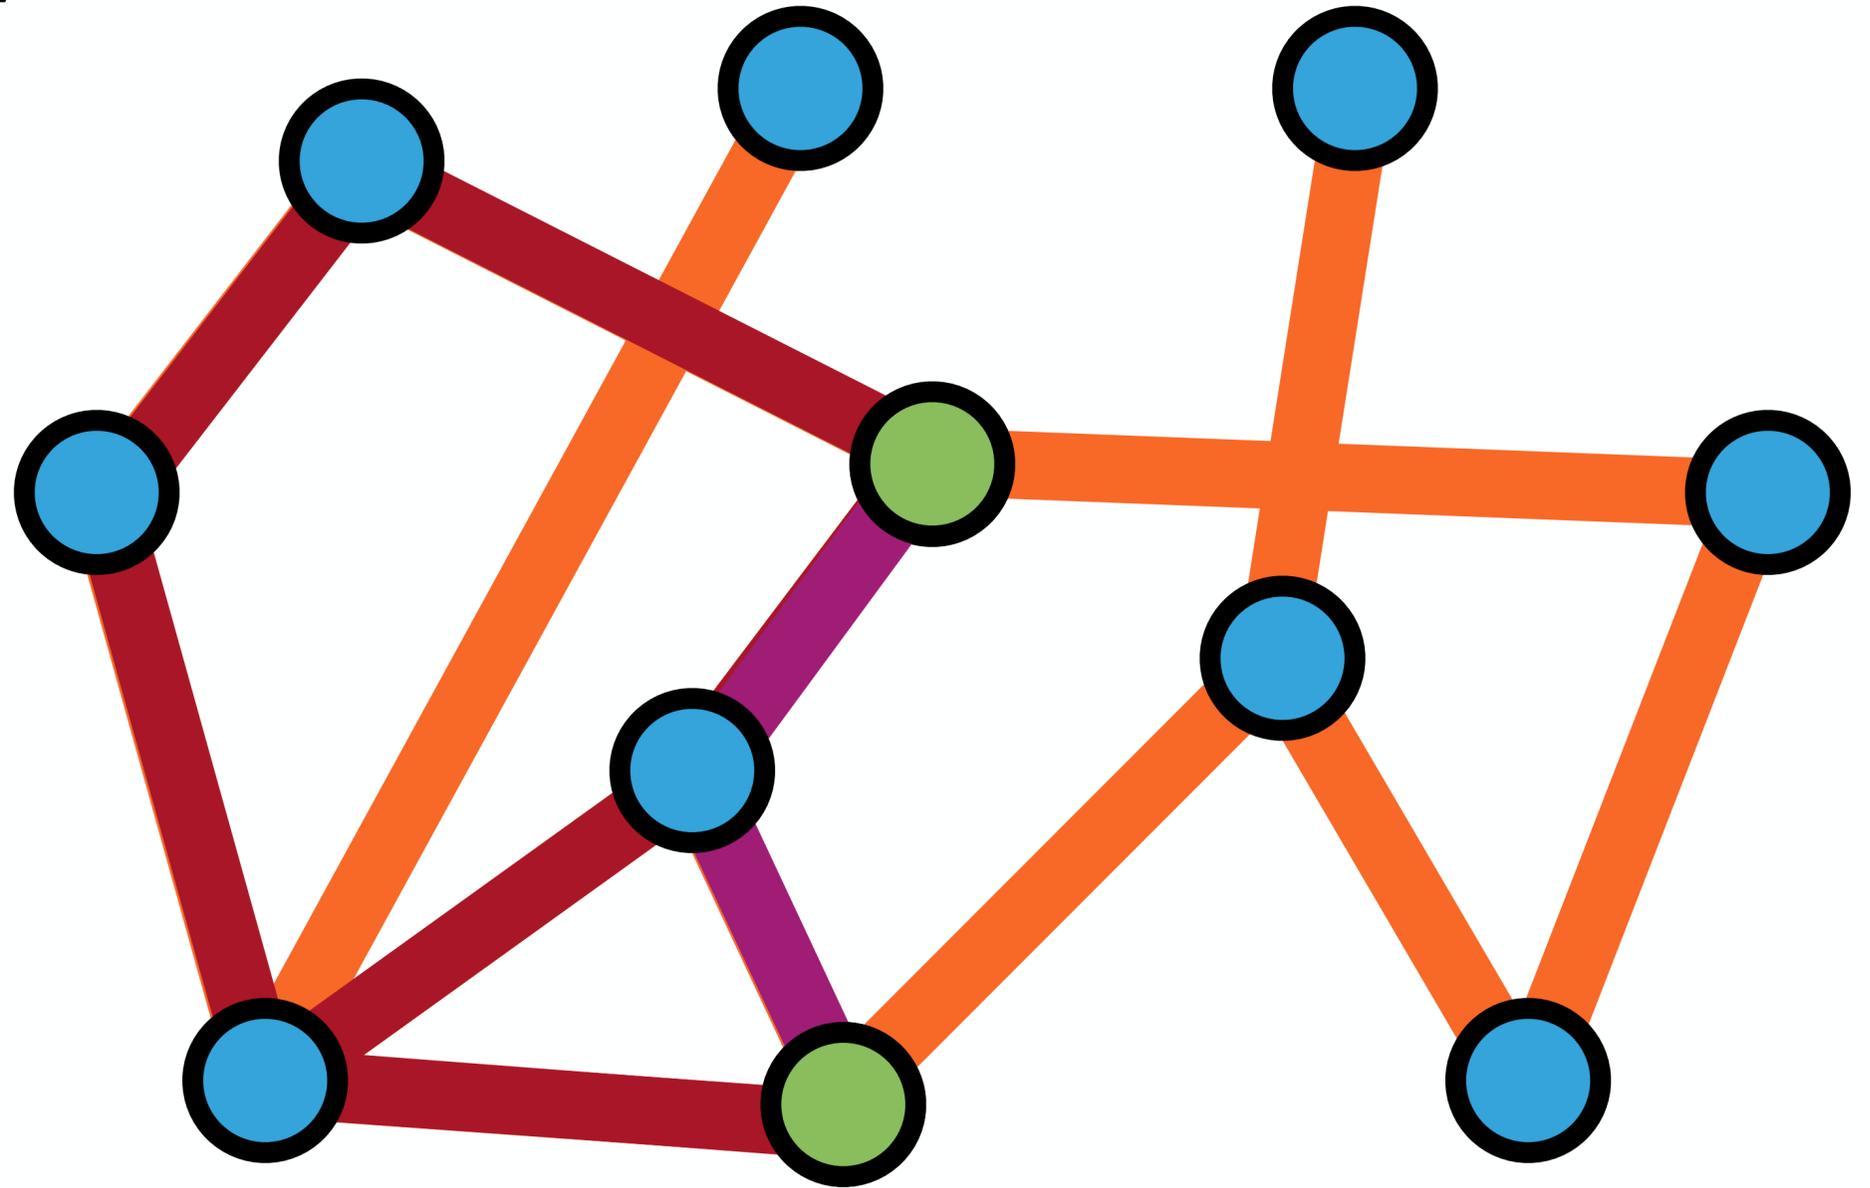


Community "Closeness" :

$1/(\text{Number Crossrefs})$

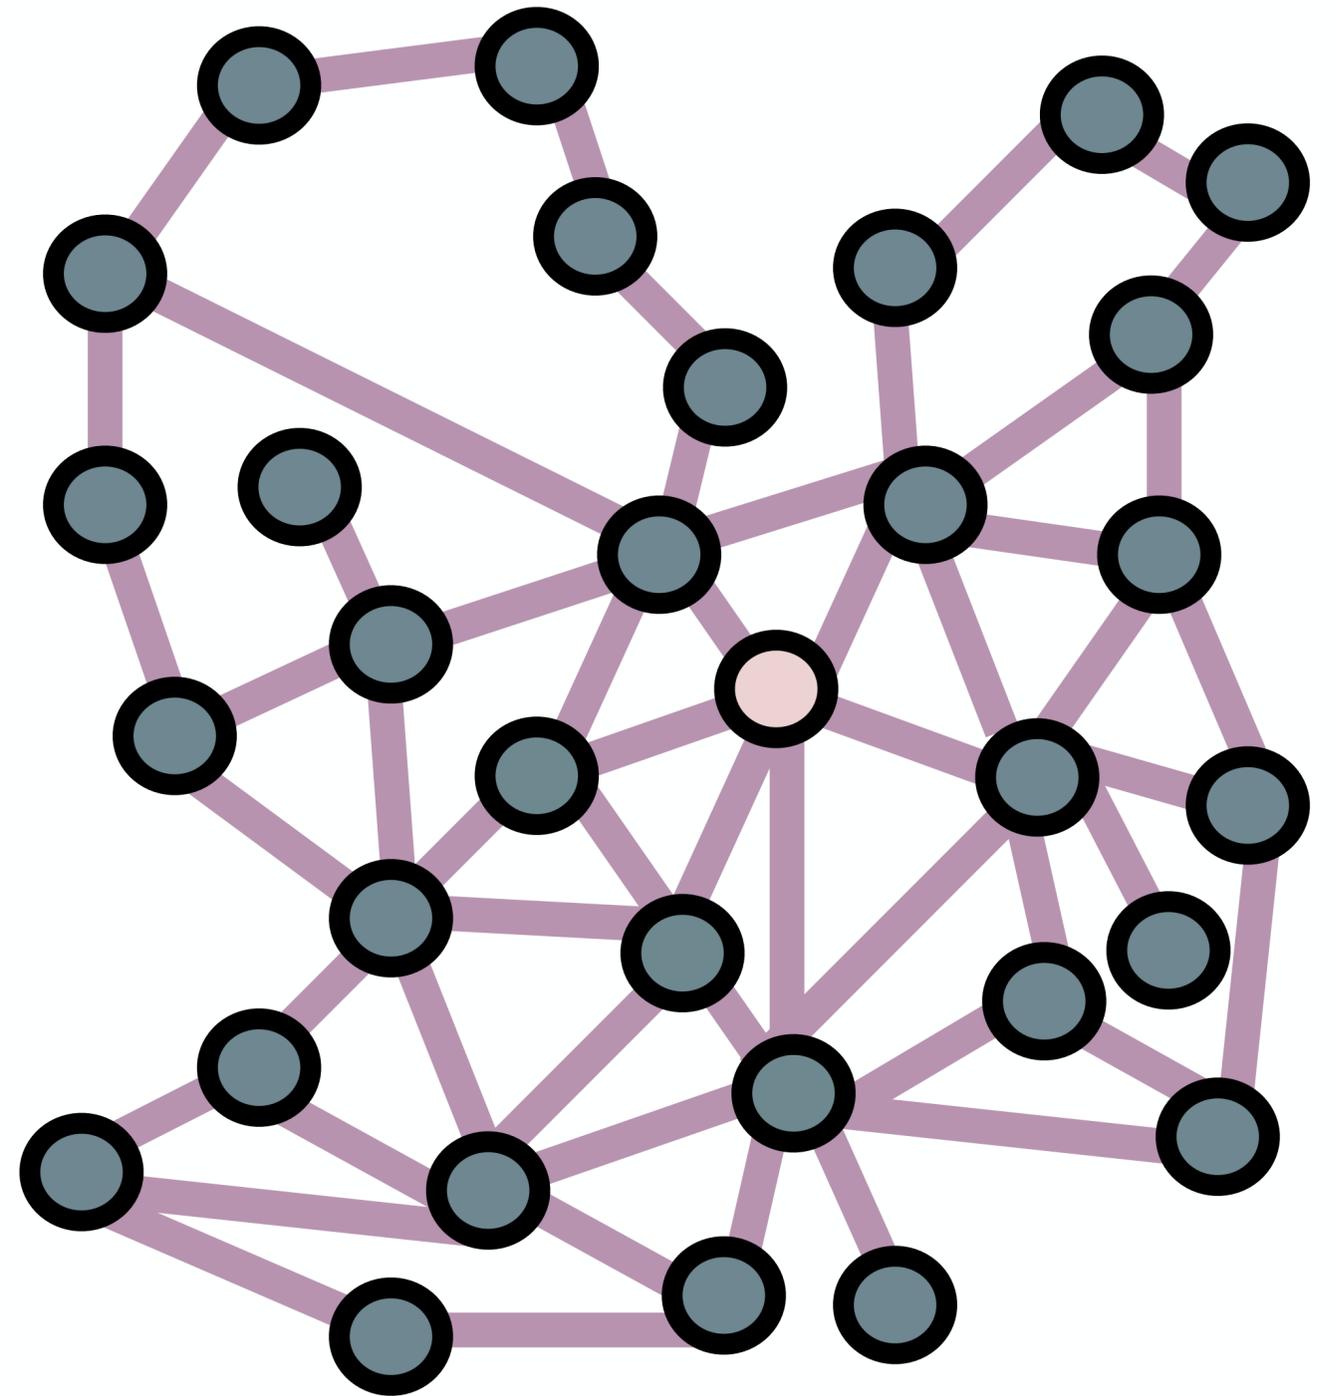
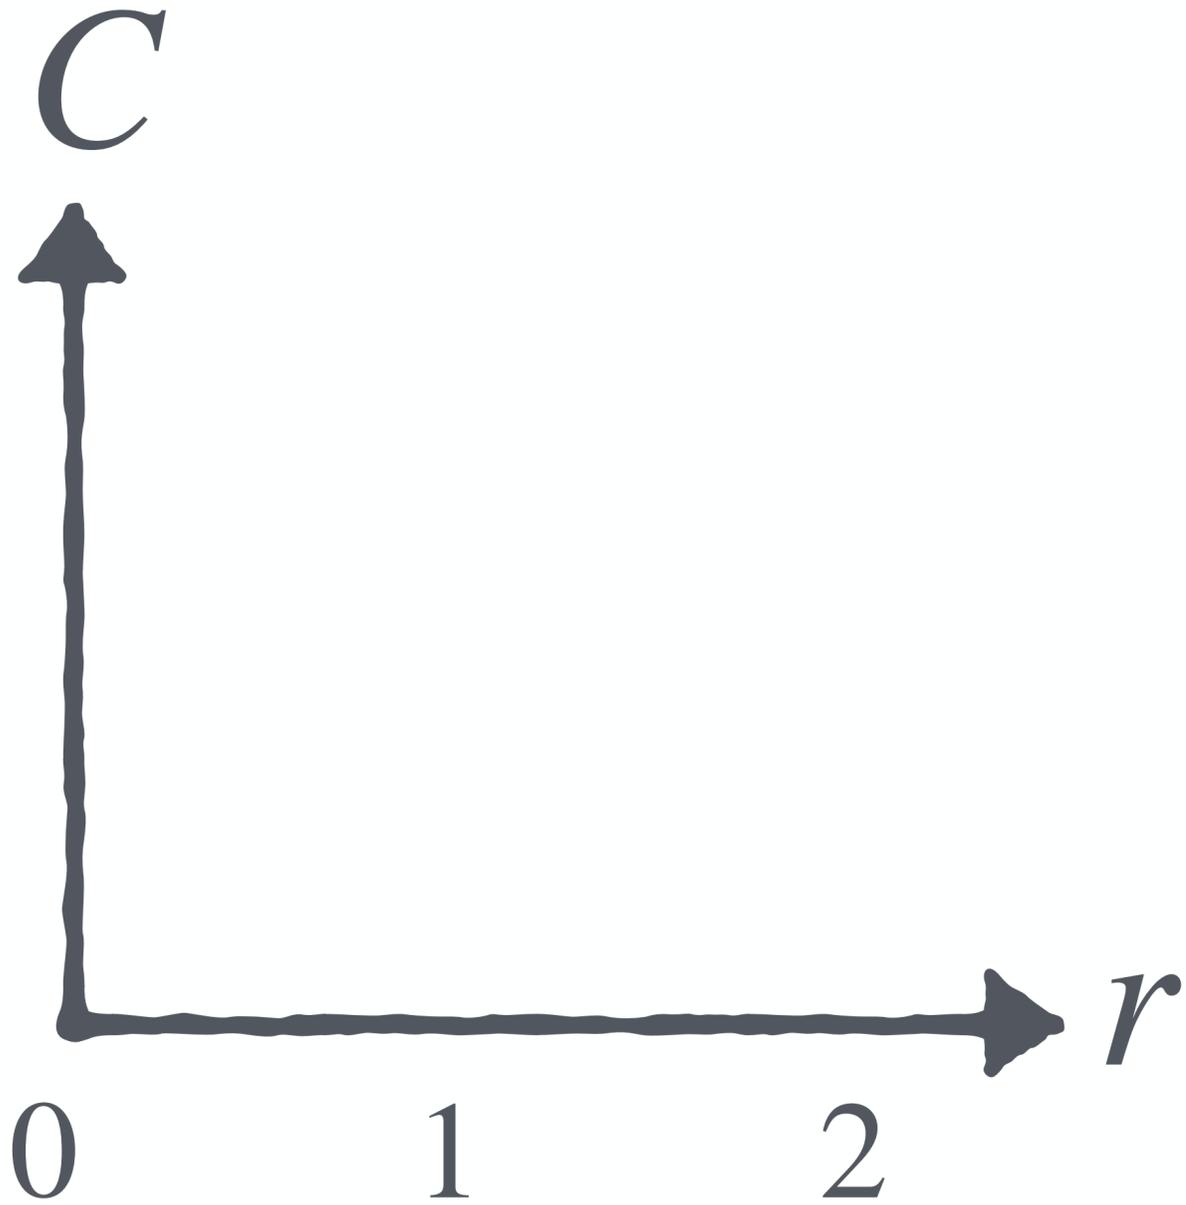
Distance gives Geometry
to a Graph

The **distance**
between two
vertices is the **length**
of the shortest path.

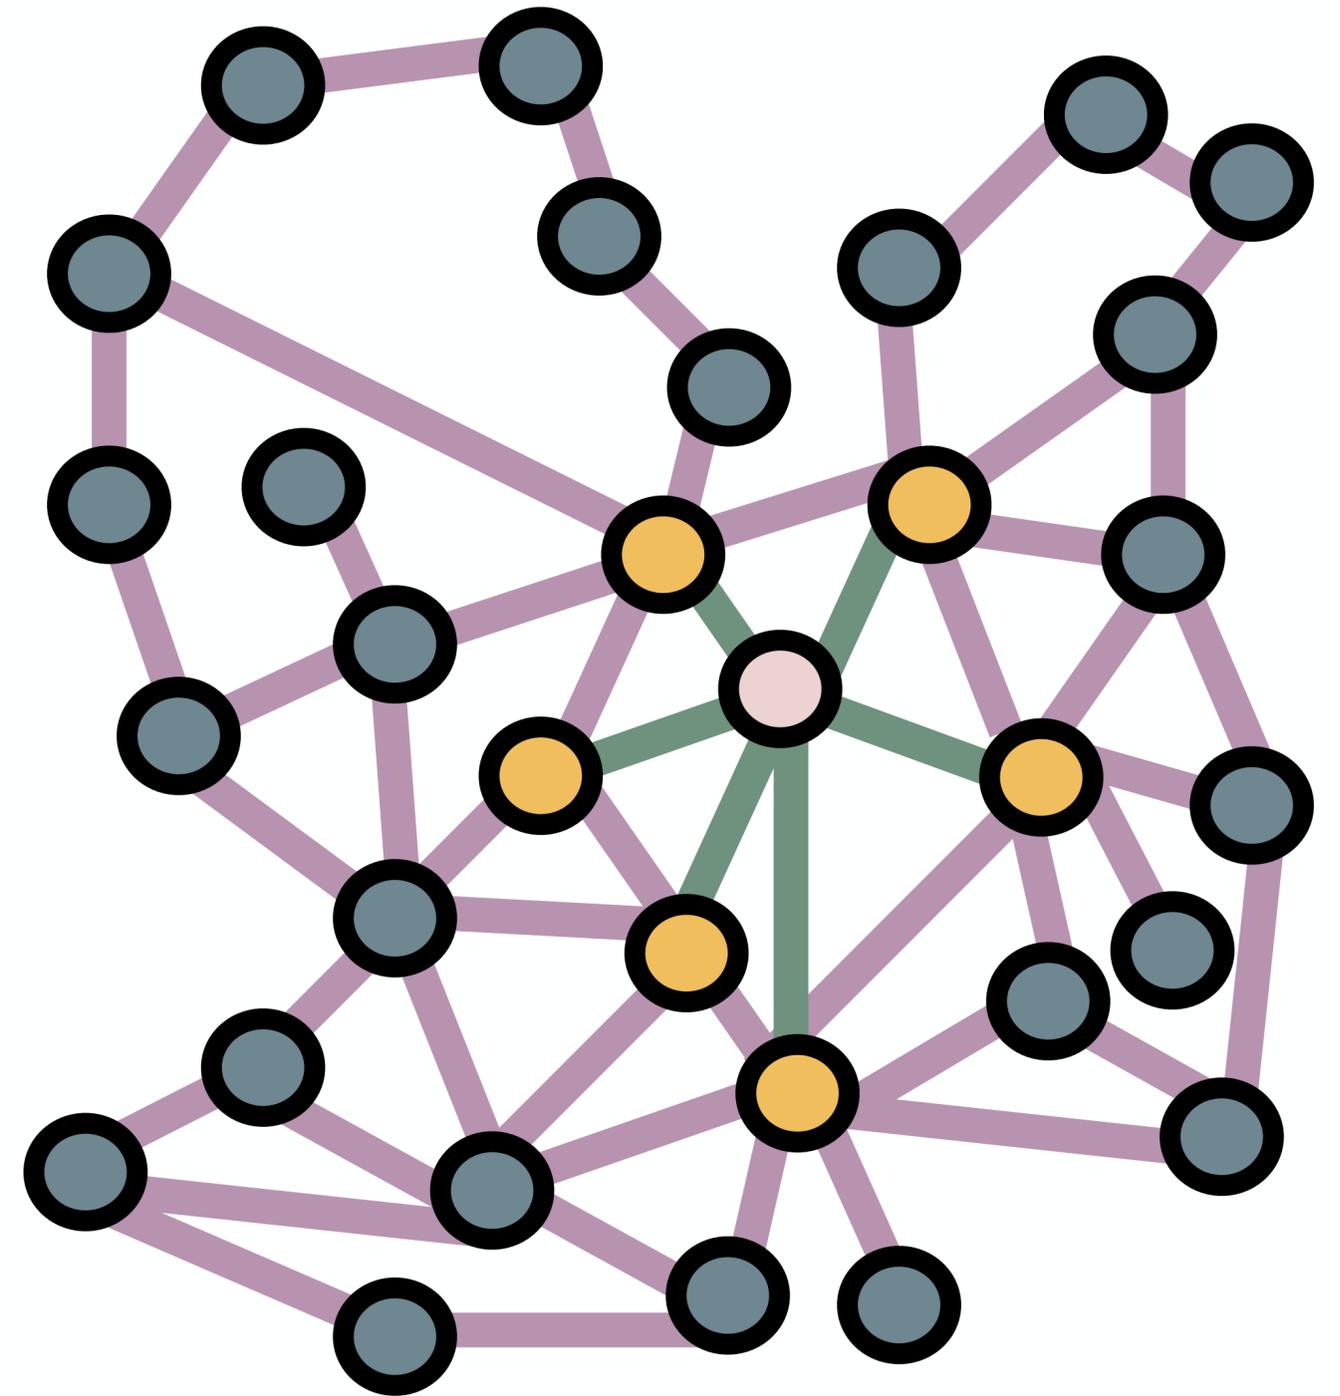
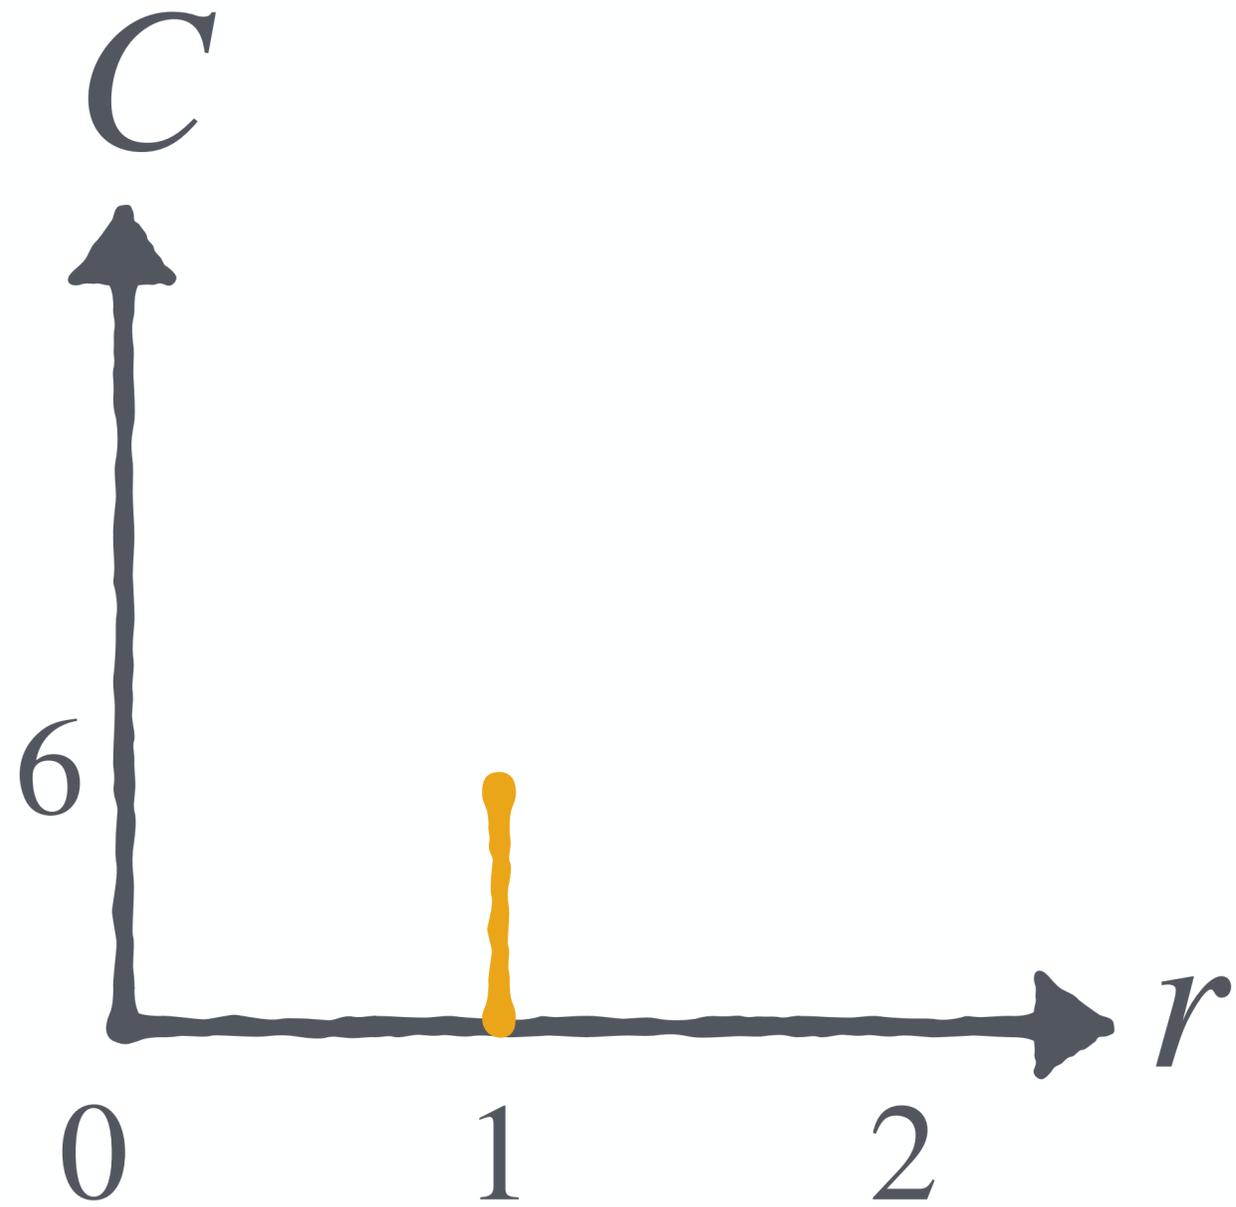


$$\ell = \mathcal{L} = \text{dist}$$

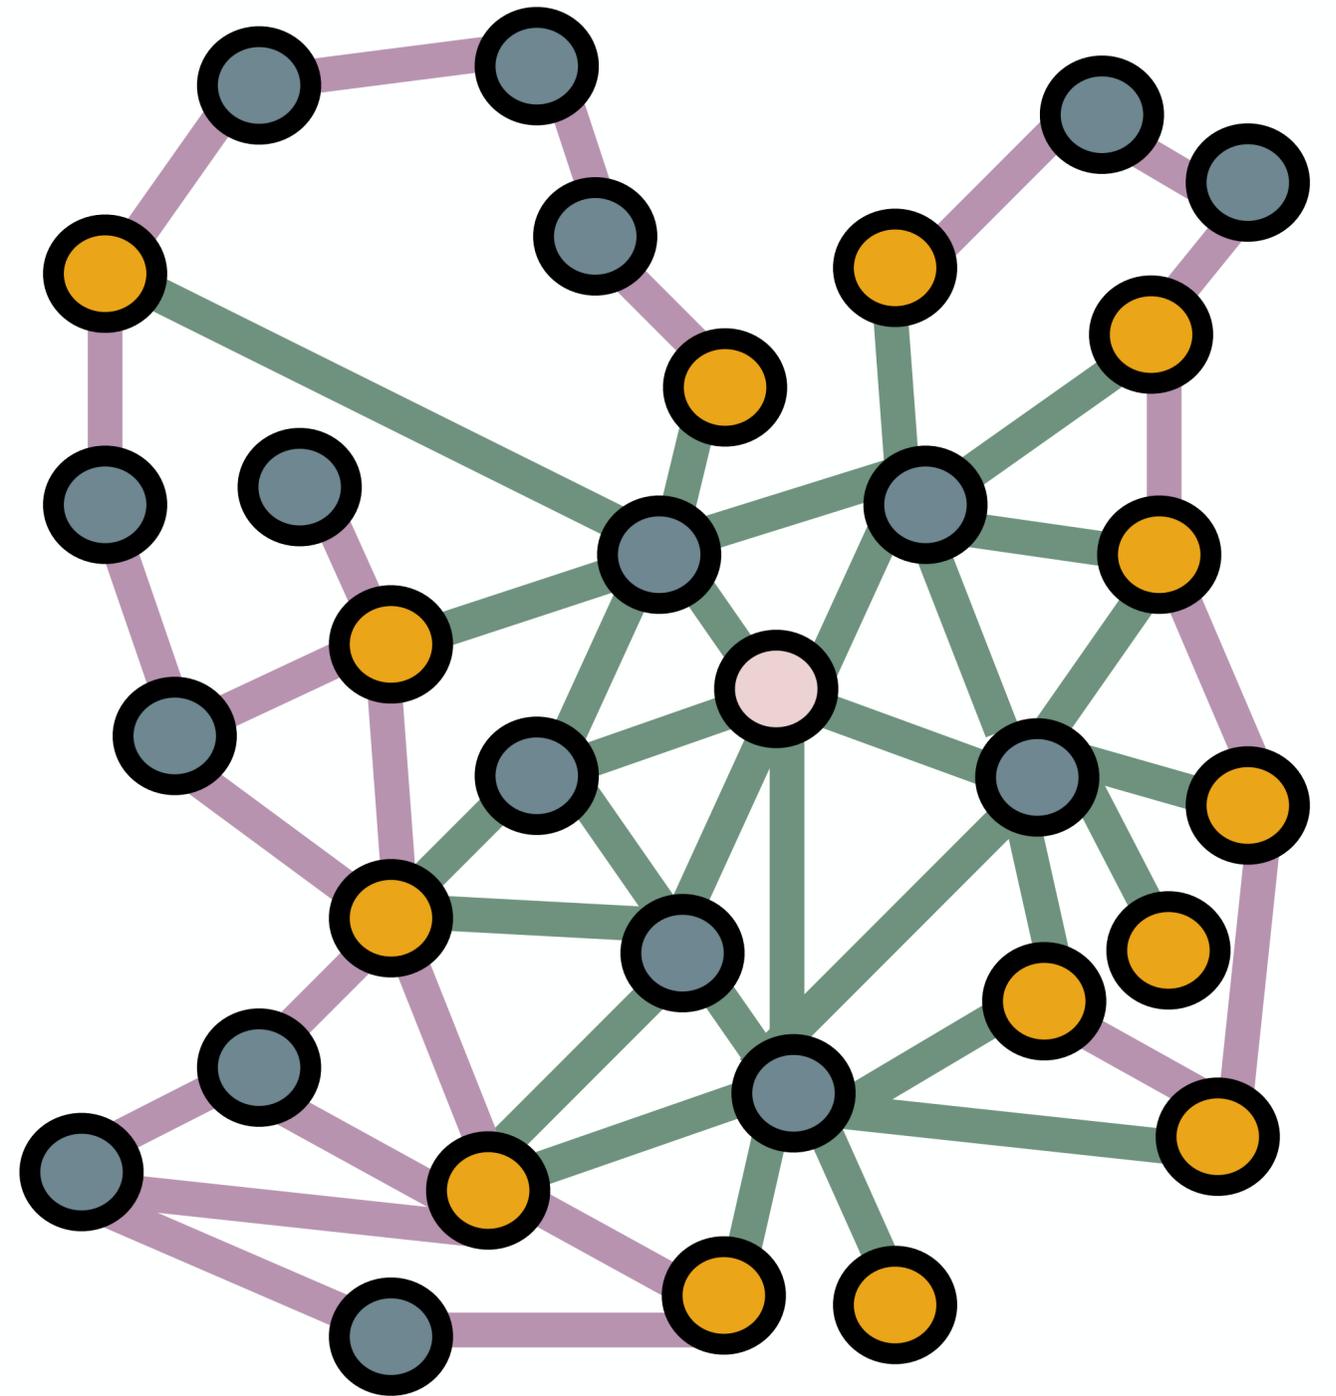
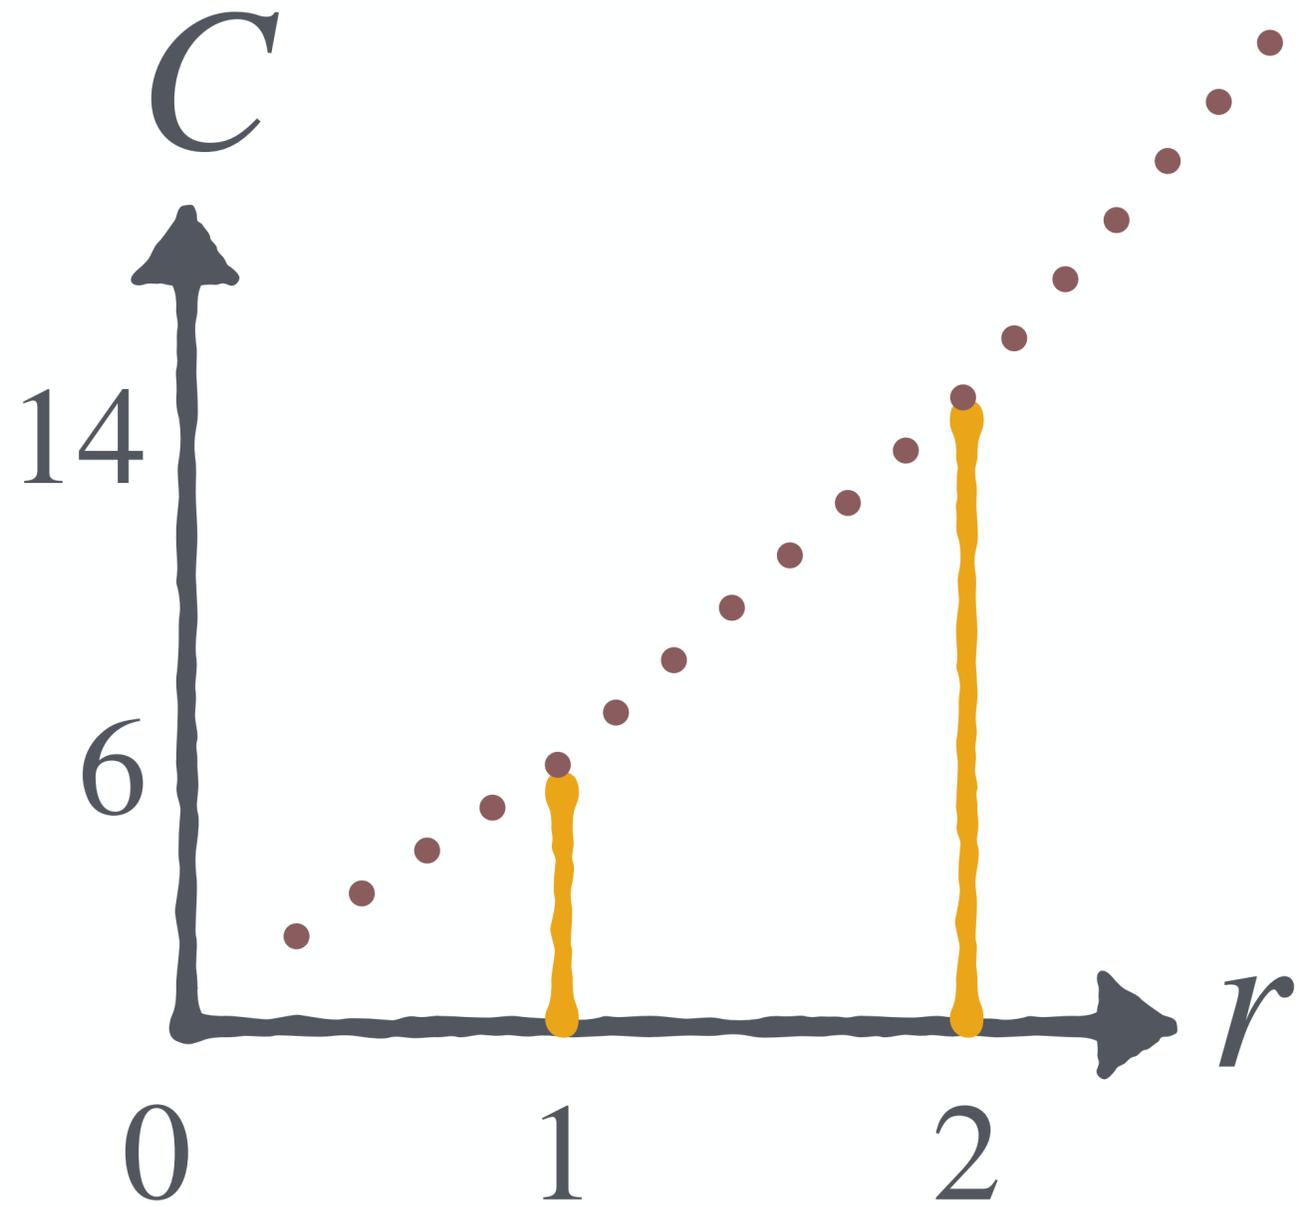
Distance gives Geometry
to a Graph



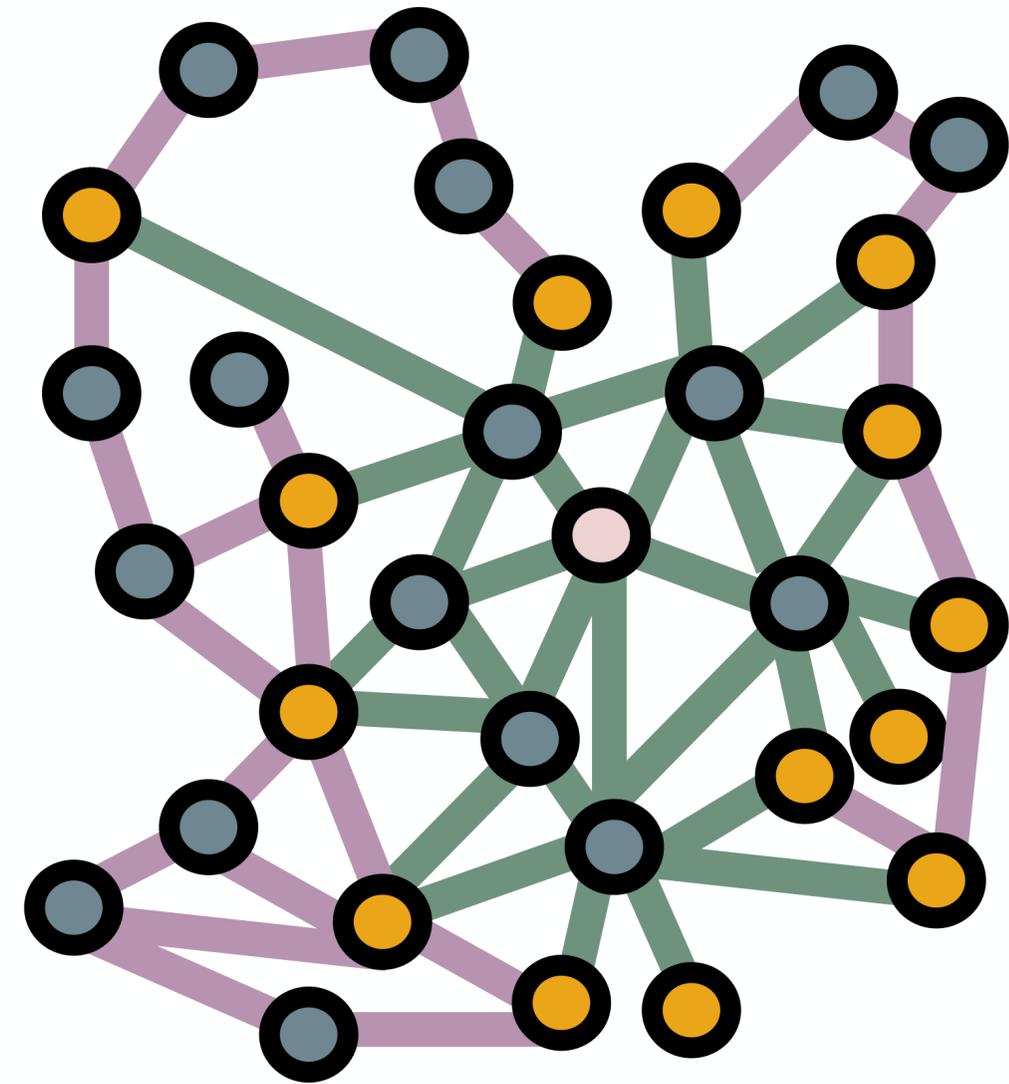
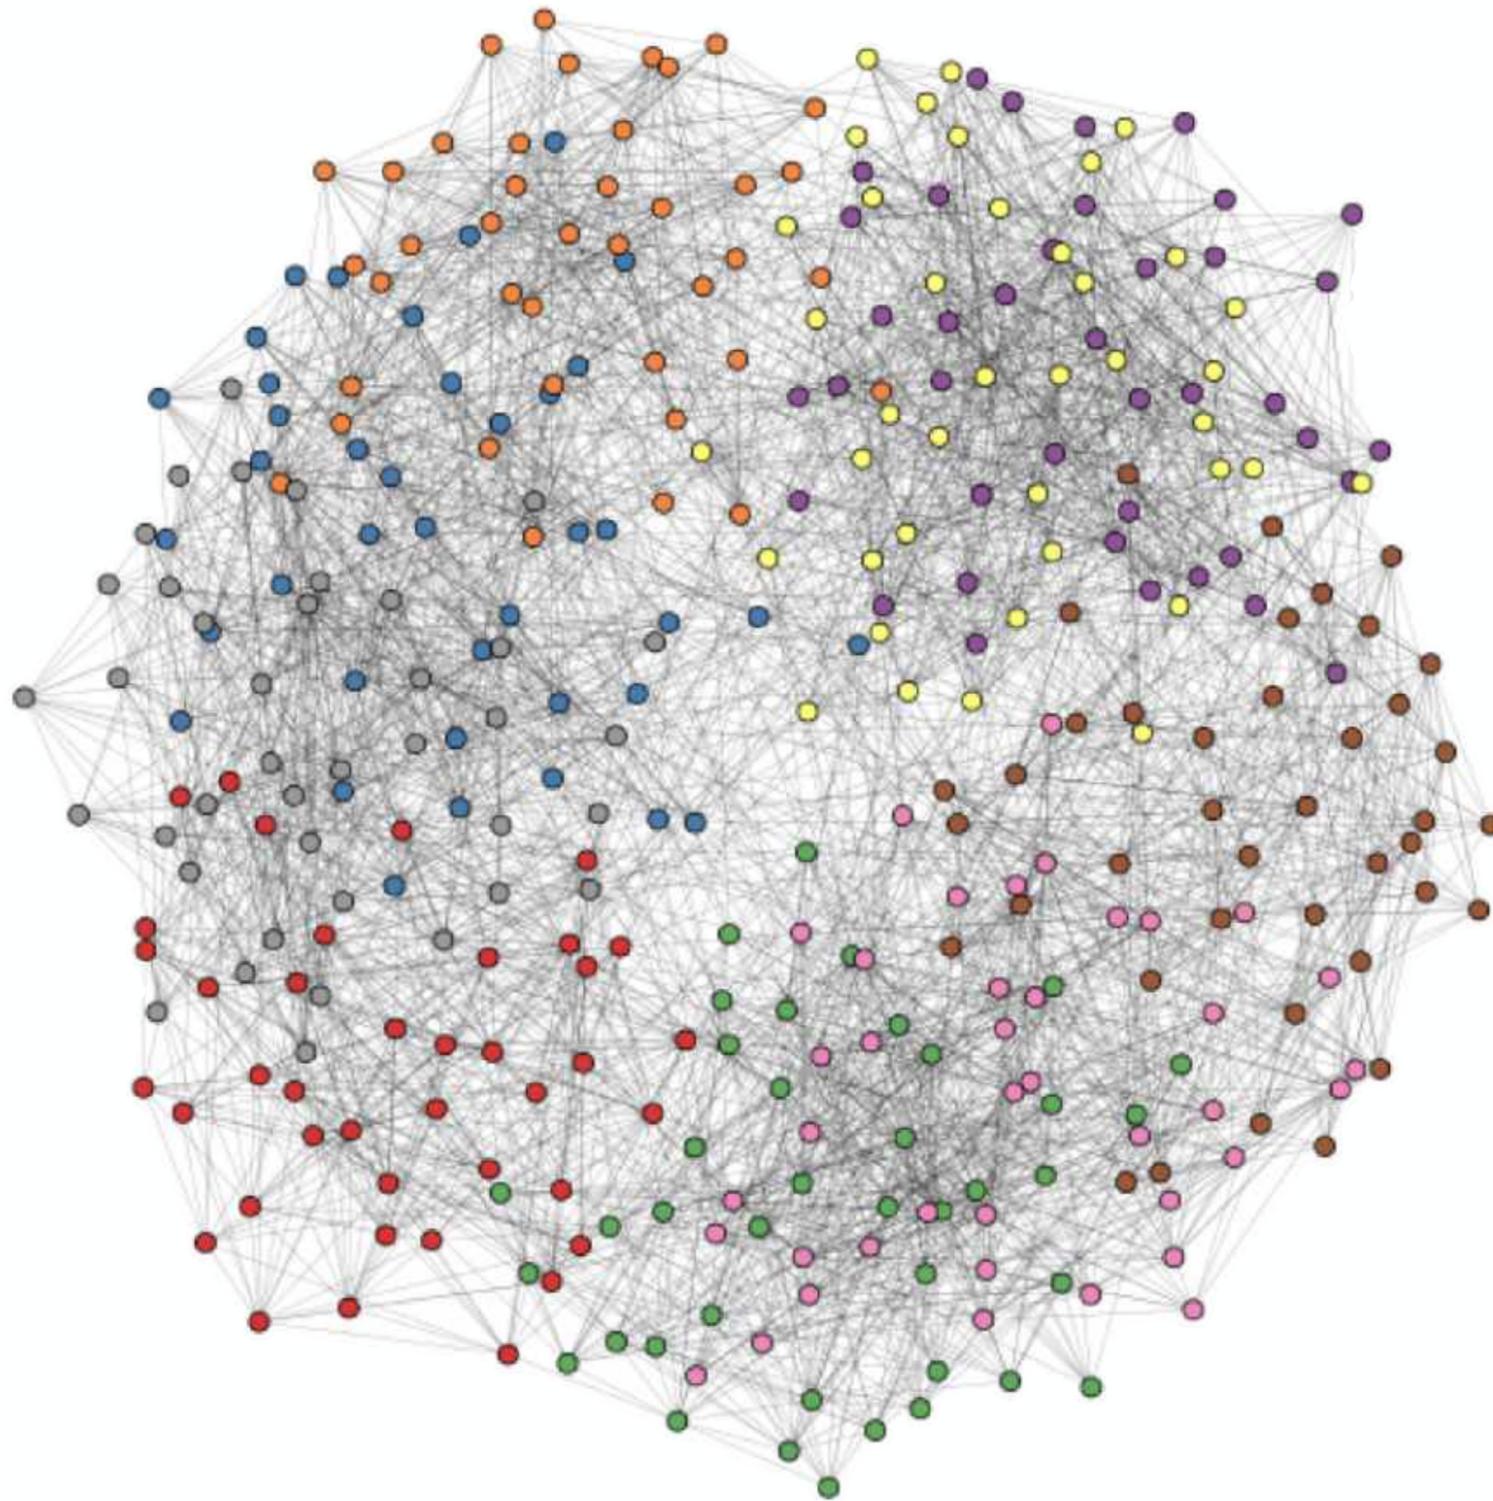
Distance gives Geometry to a Graph



Distance gives Geometry to a Graph



Unfortunately, this can be difficult to compute





But not always! It's
easy on this graph.

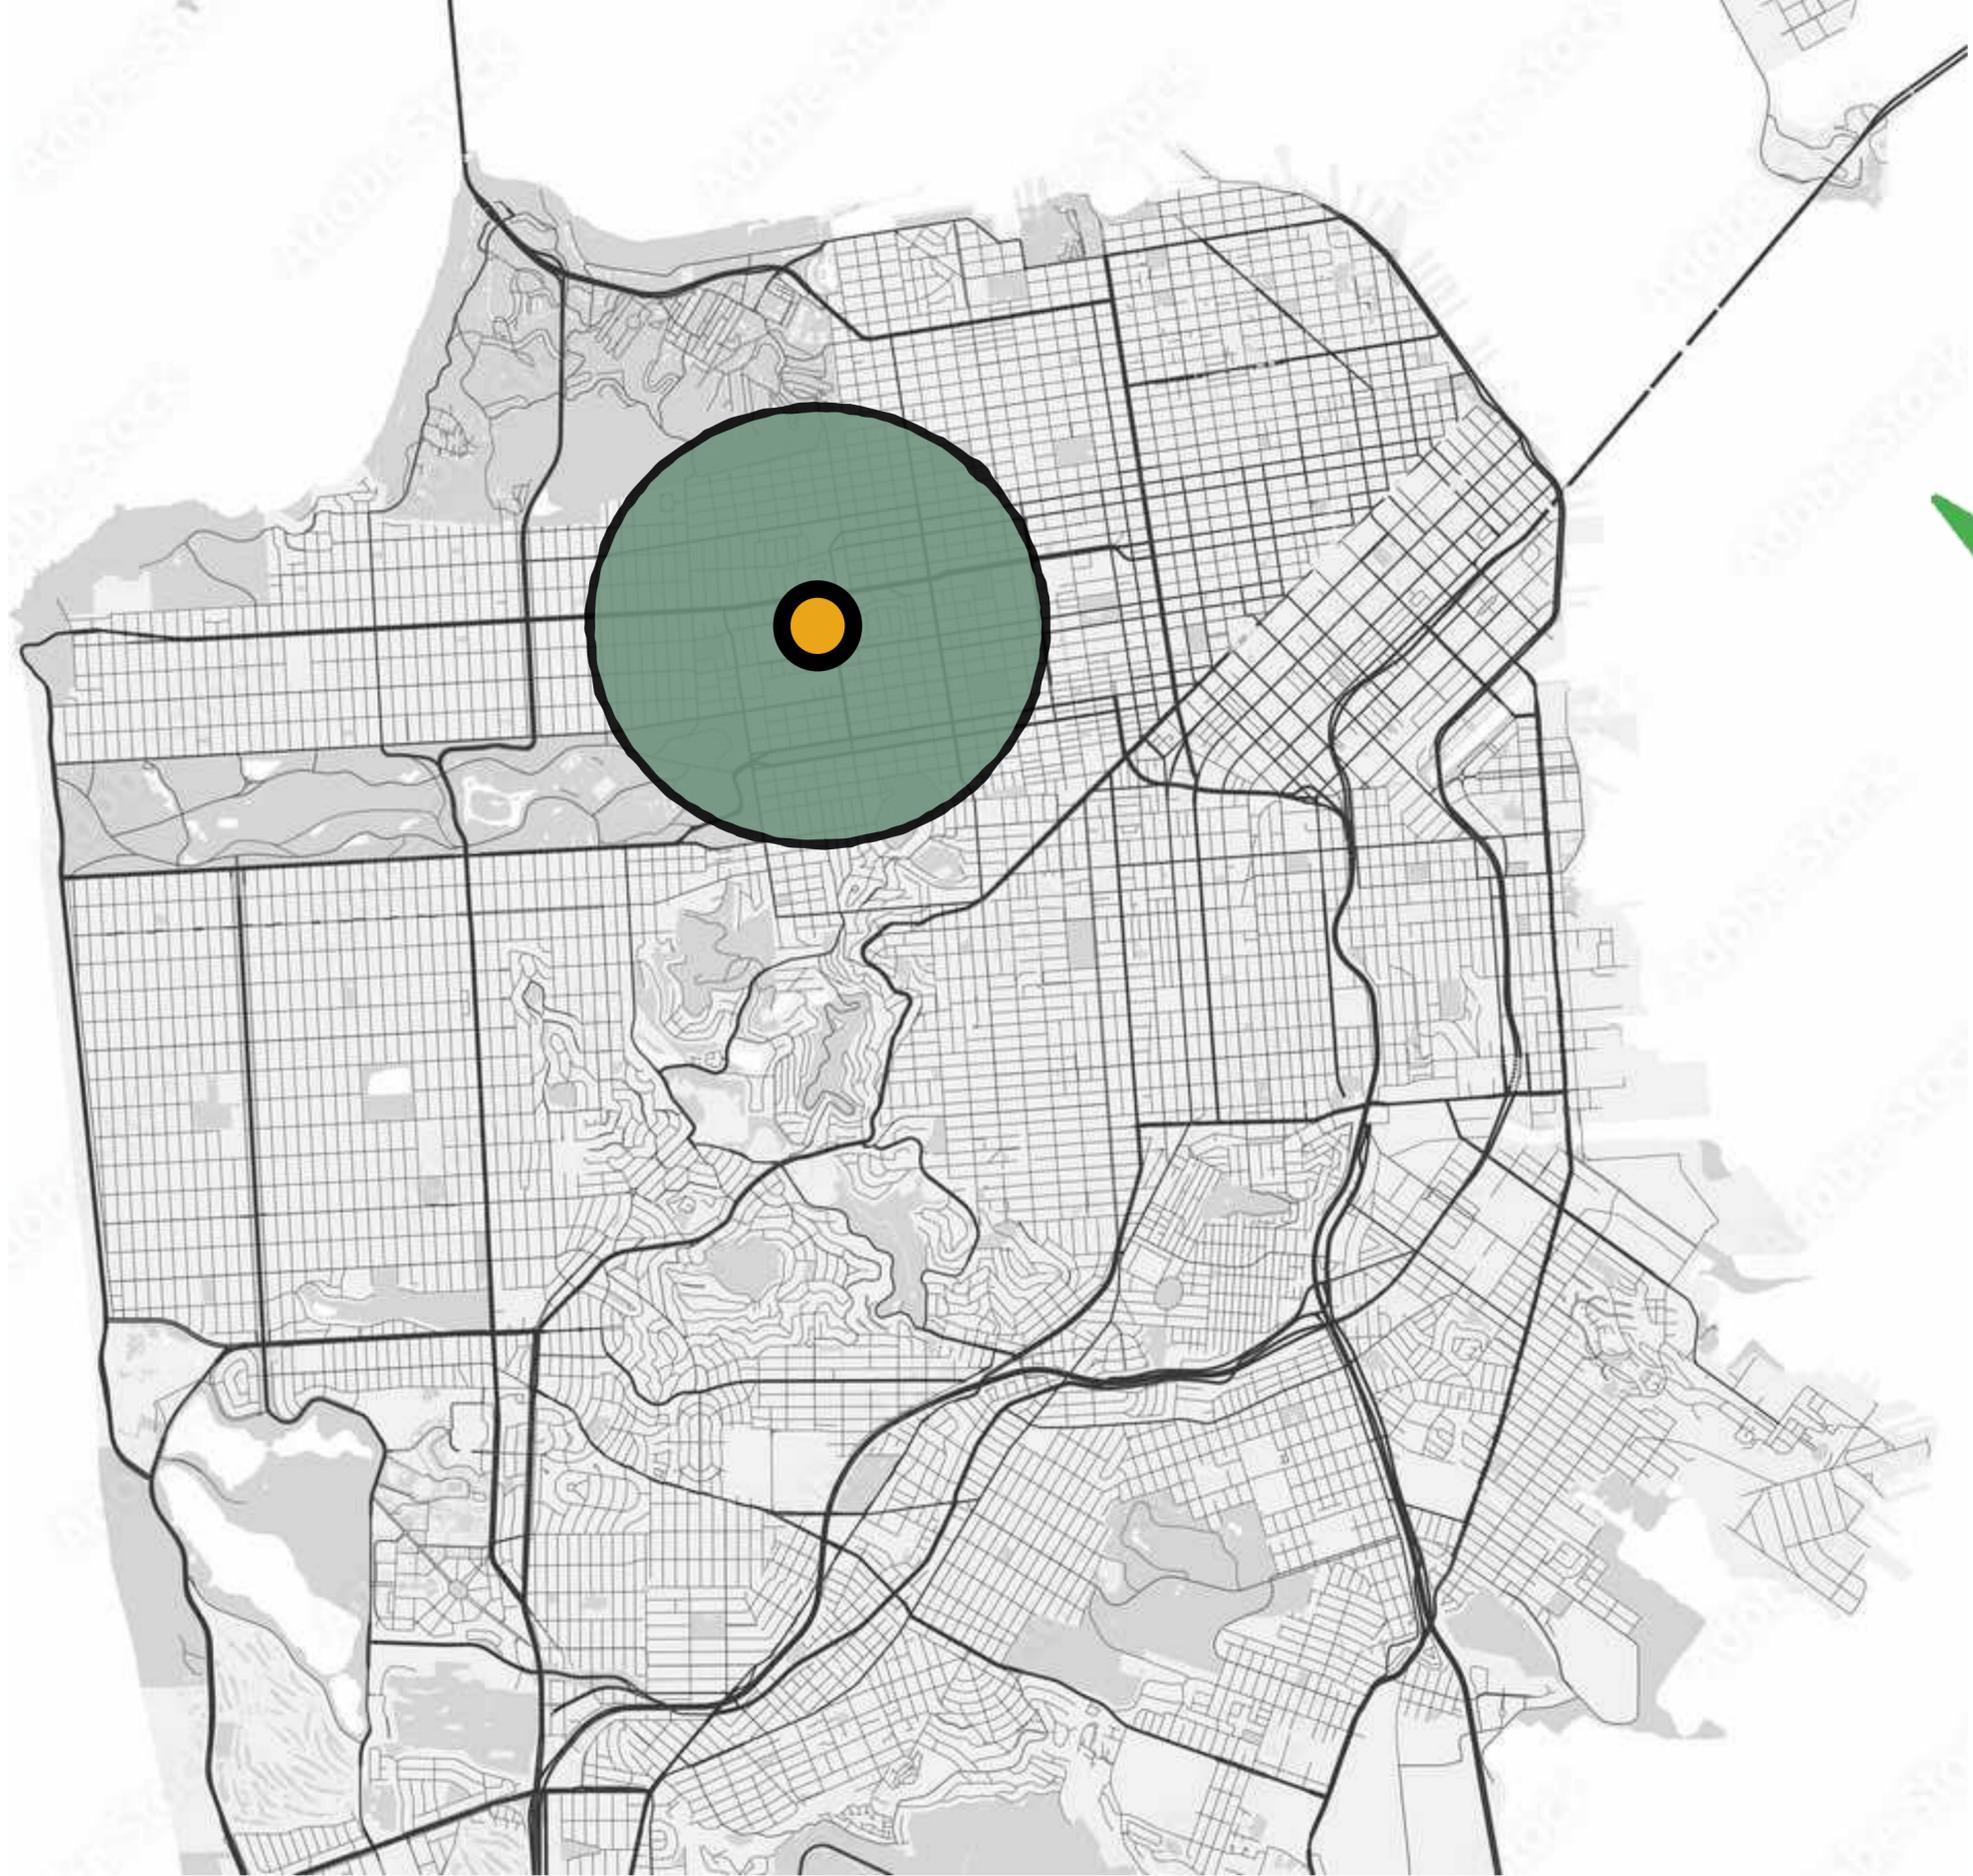
Finding Distances



But not always! It's easy on this graph.



Finding Distances

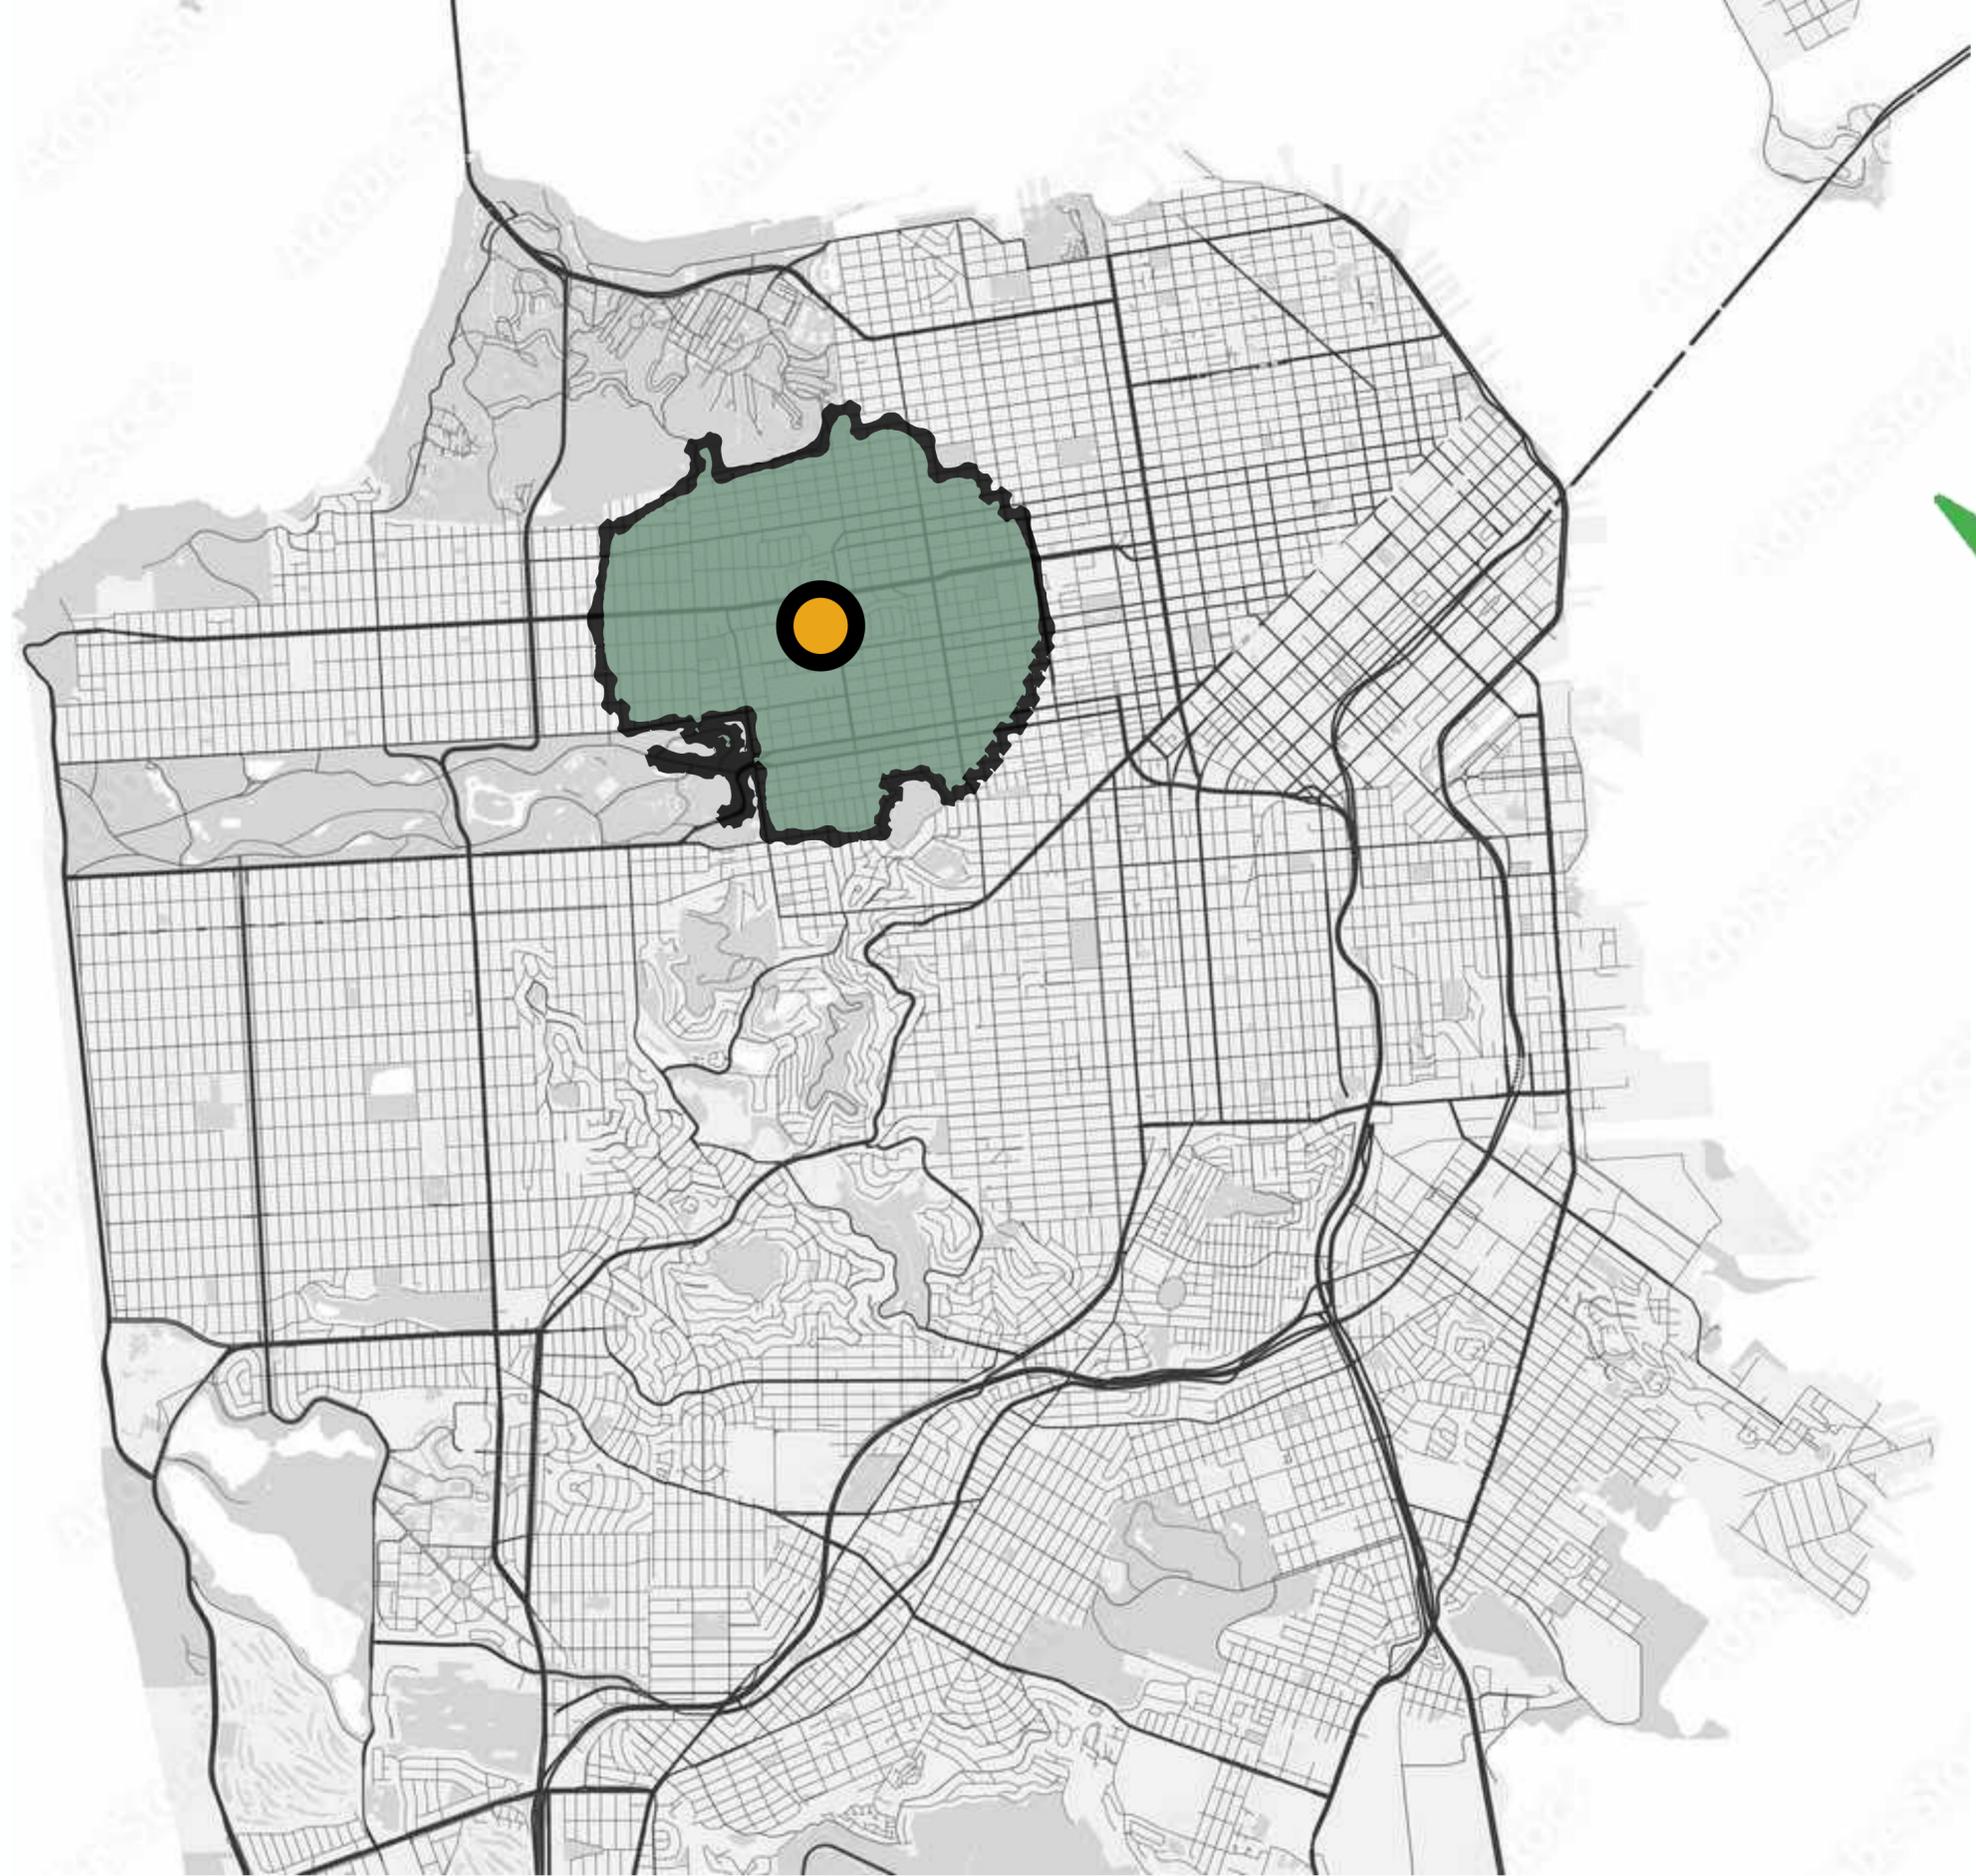


But not always! It's
easy on this graph.



Finding Distances

Finding balls



But not always! It's easy on this graph.

- ✓ Finding Distances
- ✓ Finding balls

Why its easy isn't a feature of the graph, but rather a **property of how its laid out**

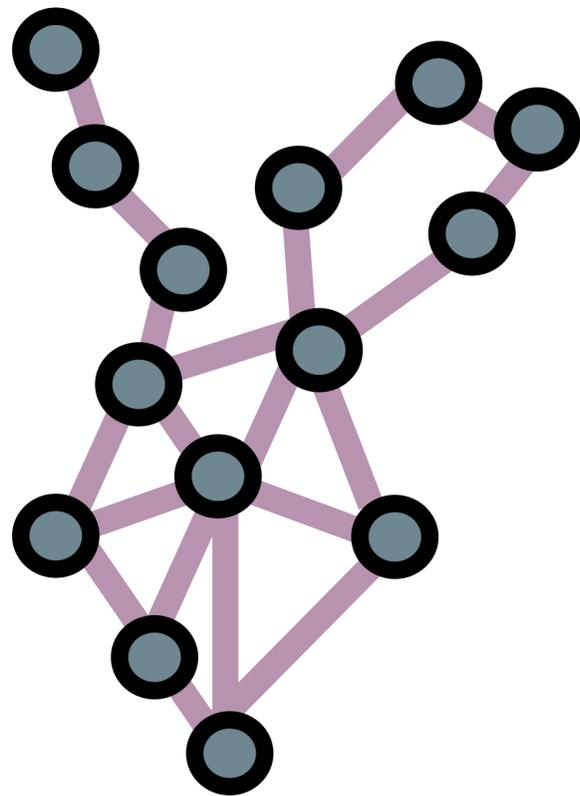
Graph distances match well with distances in the plane.

This lets us use plane geometry to estimate graph quantities.

Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning

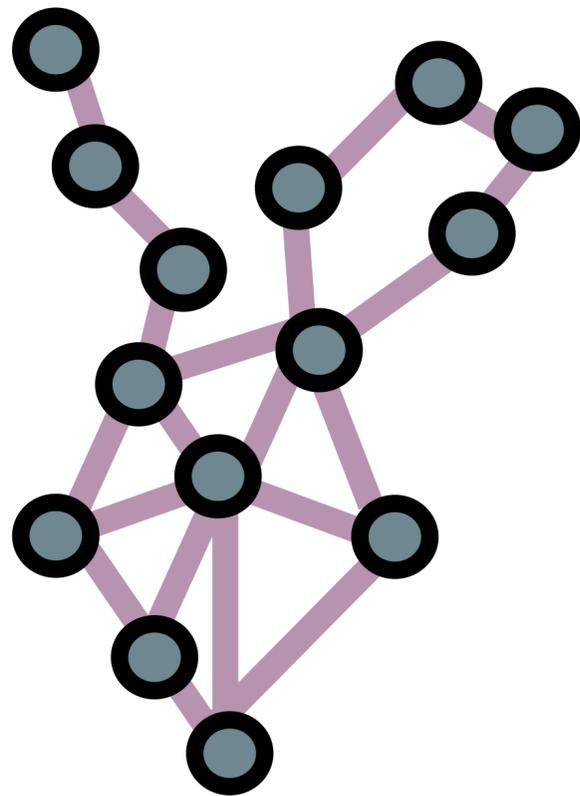
Start with a graph
 G with V vertices



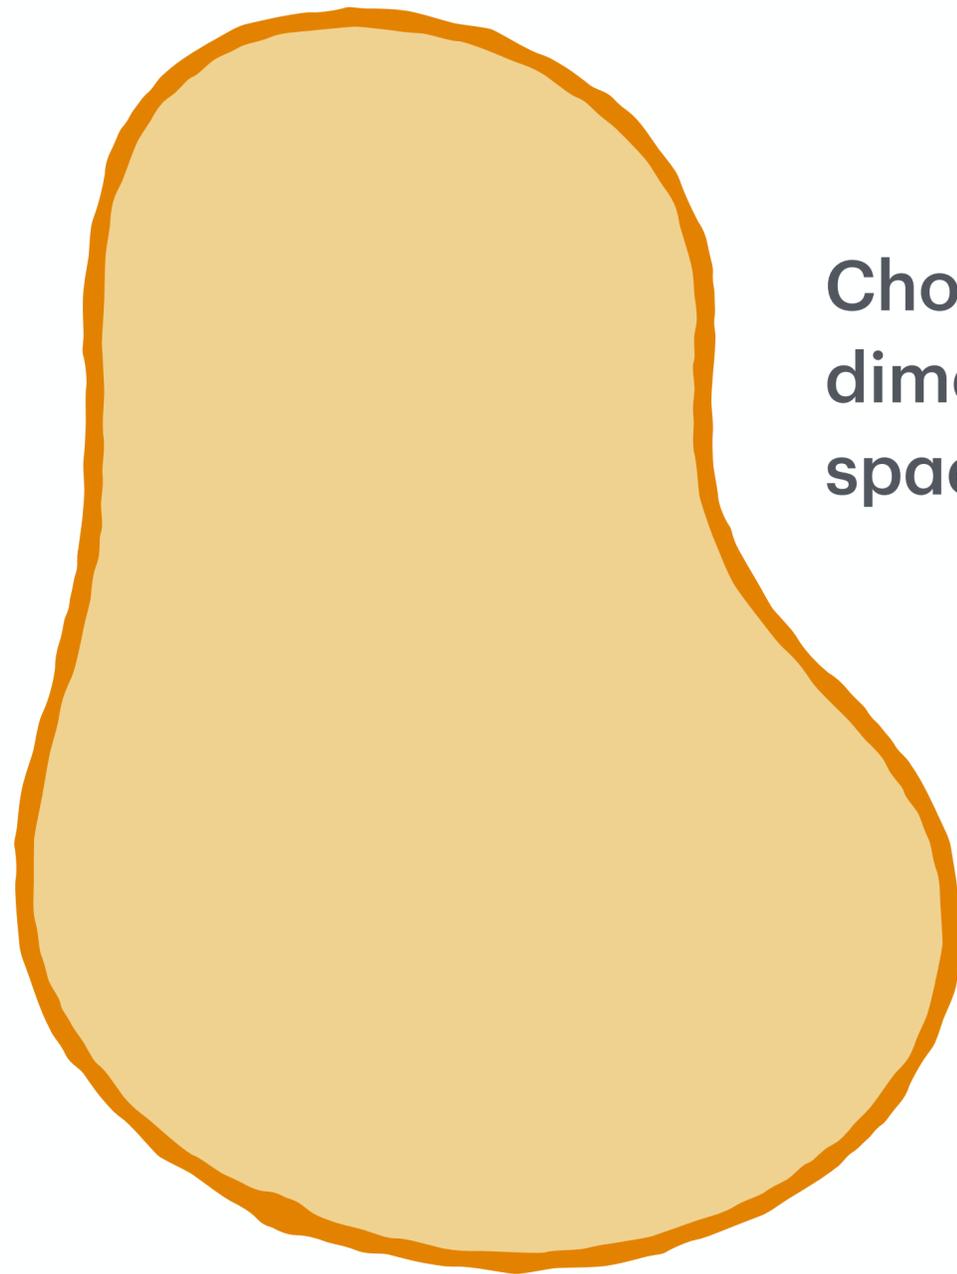
Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning

Start with a graph
 G with V vertices



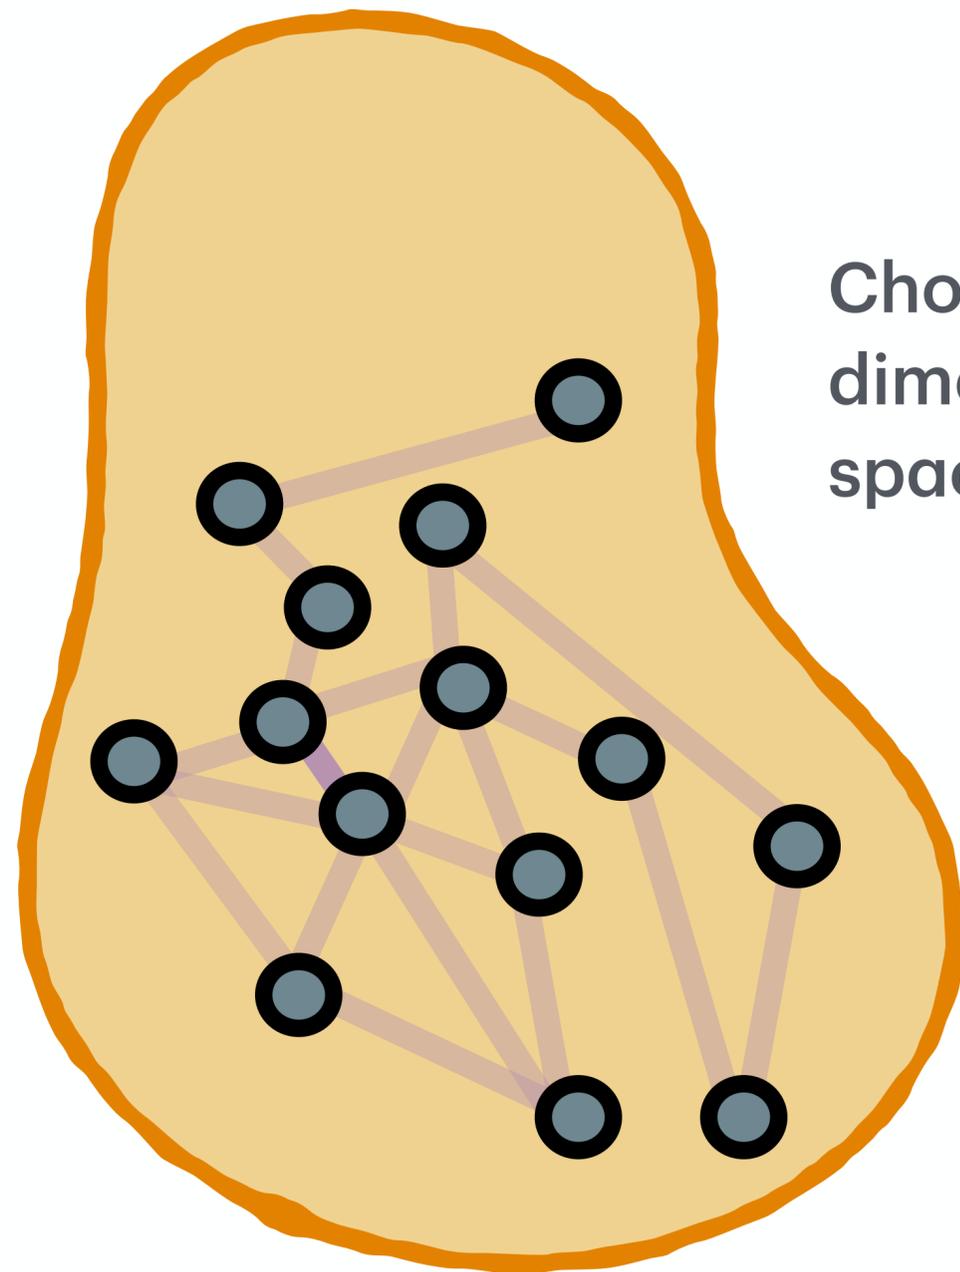
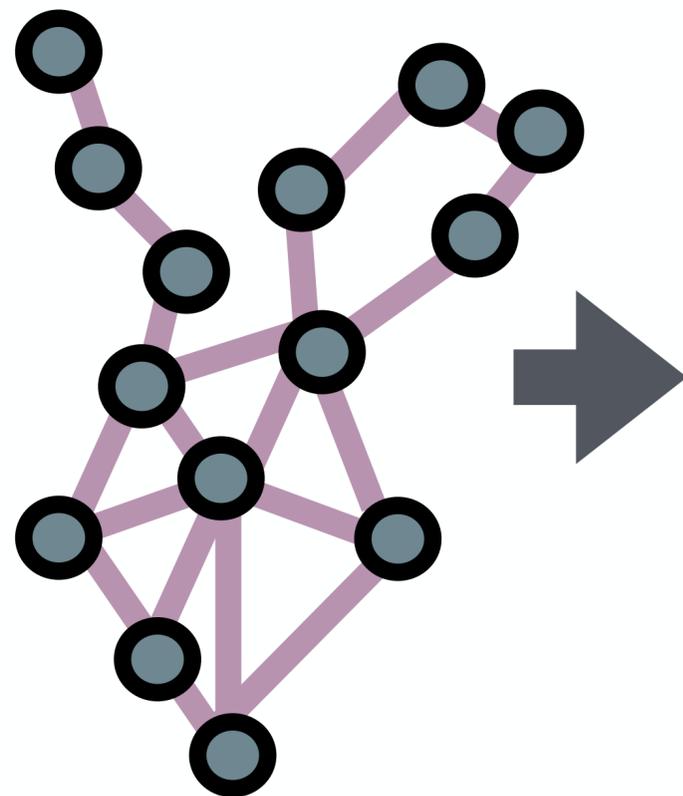
Choose a N -
dimensional
space X



Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning

Start with a graph G with V vertices



Choose a N -dimensional space X

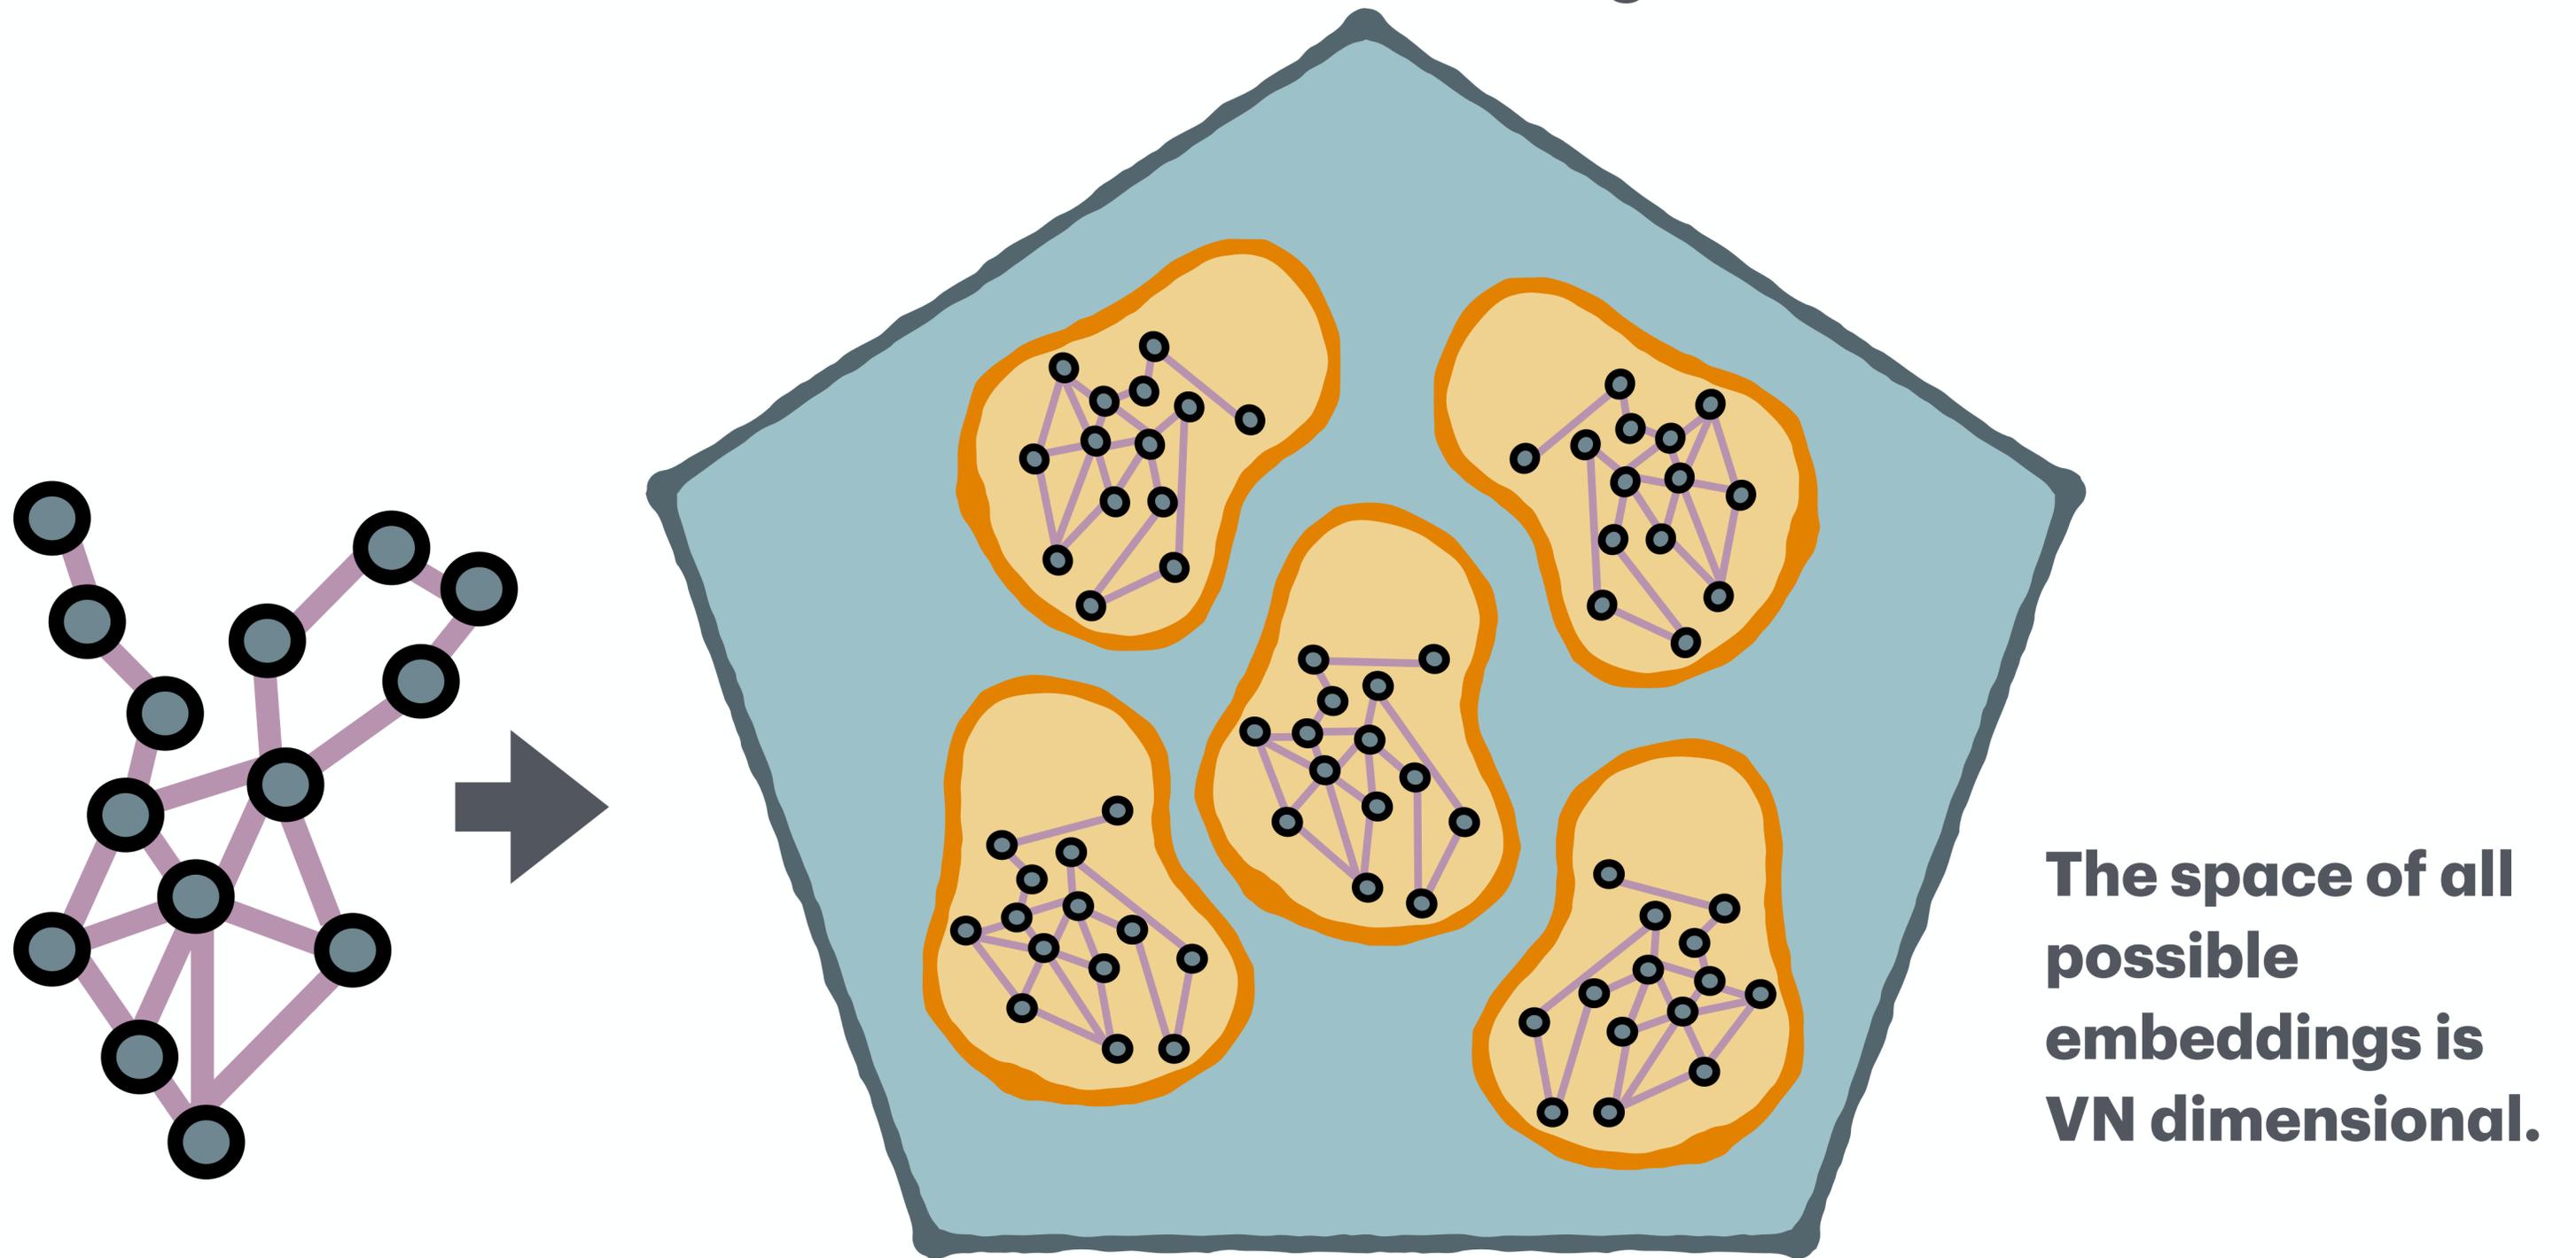
An embedding of the G in X is a choice of V points.

Each point is specified by N numbers.

An embedding of G in X is specified by VN numbers.

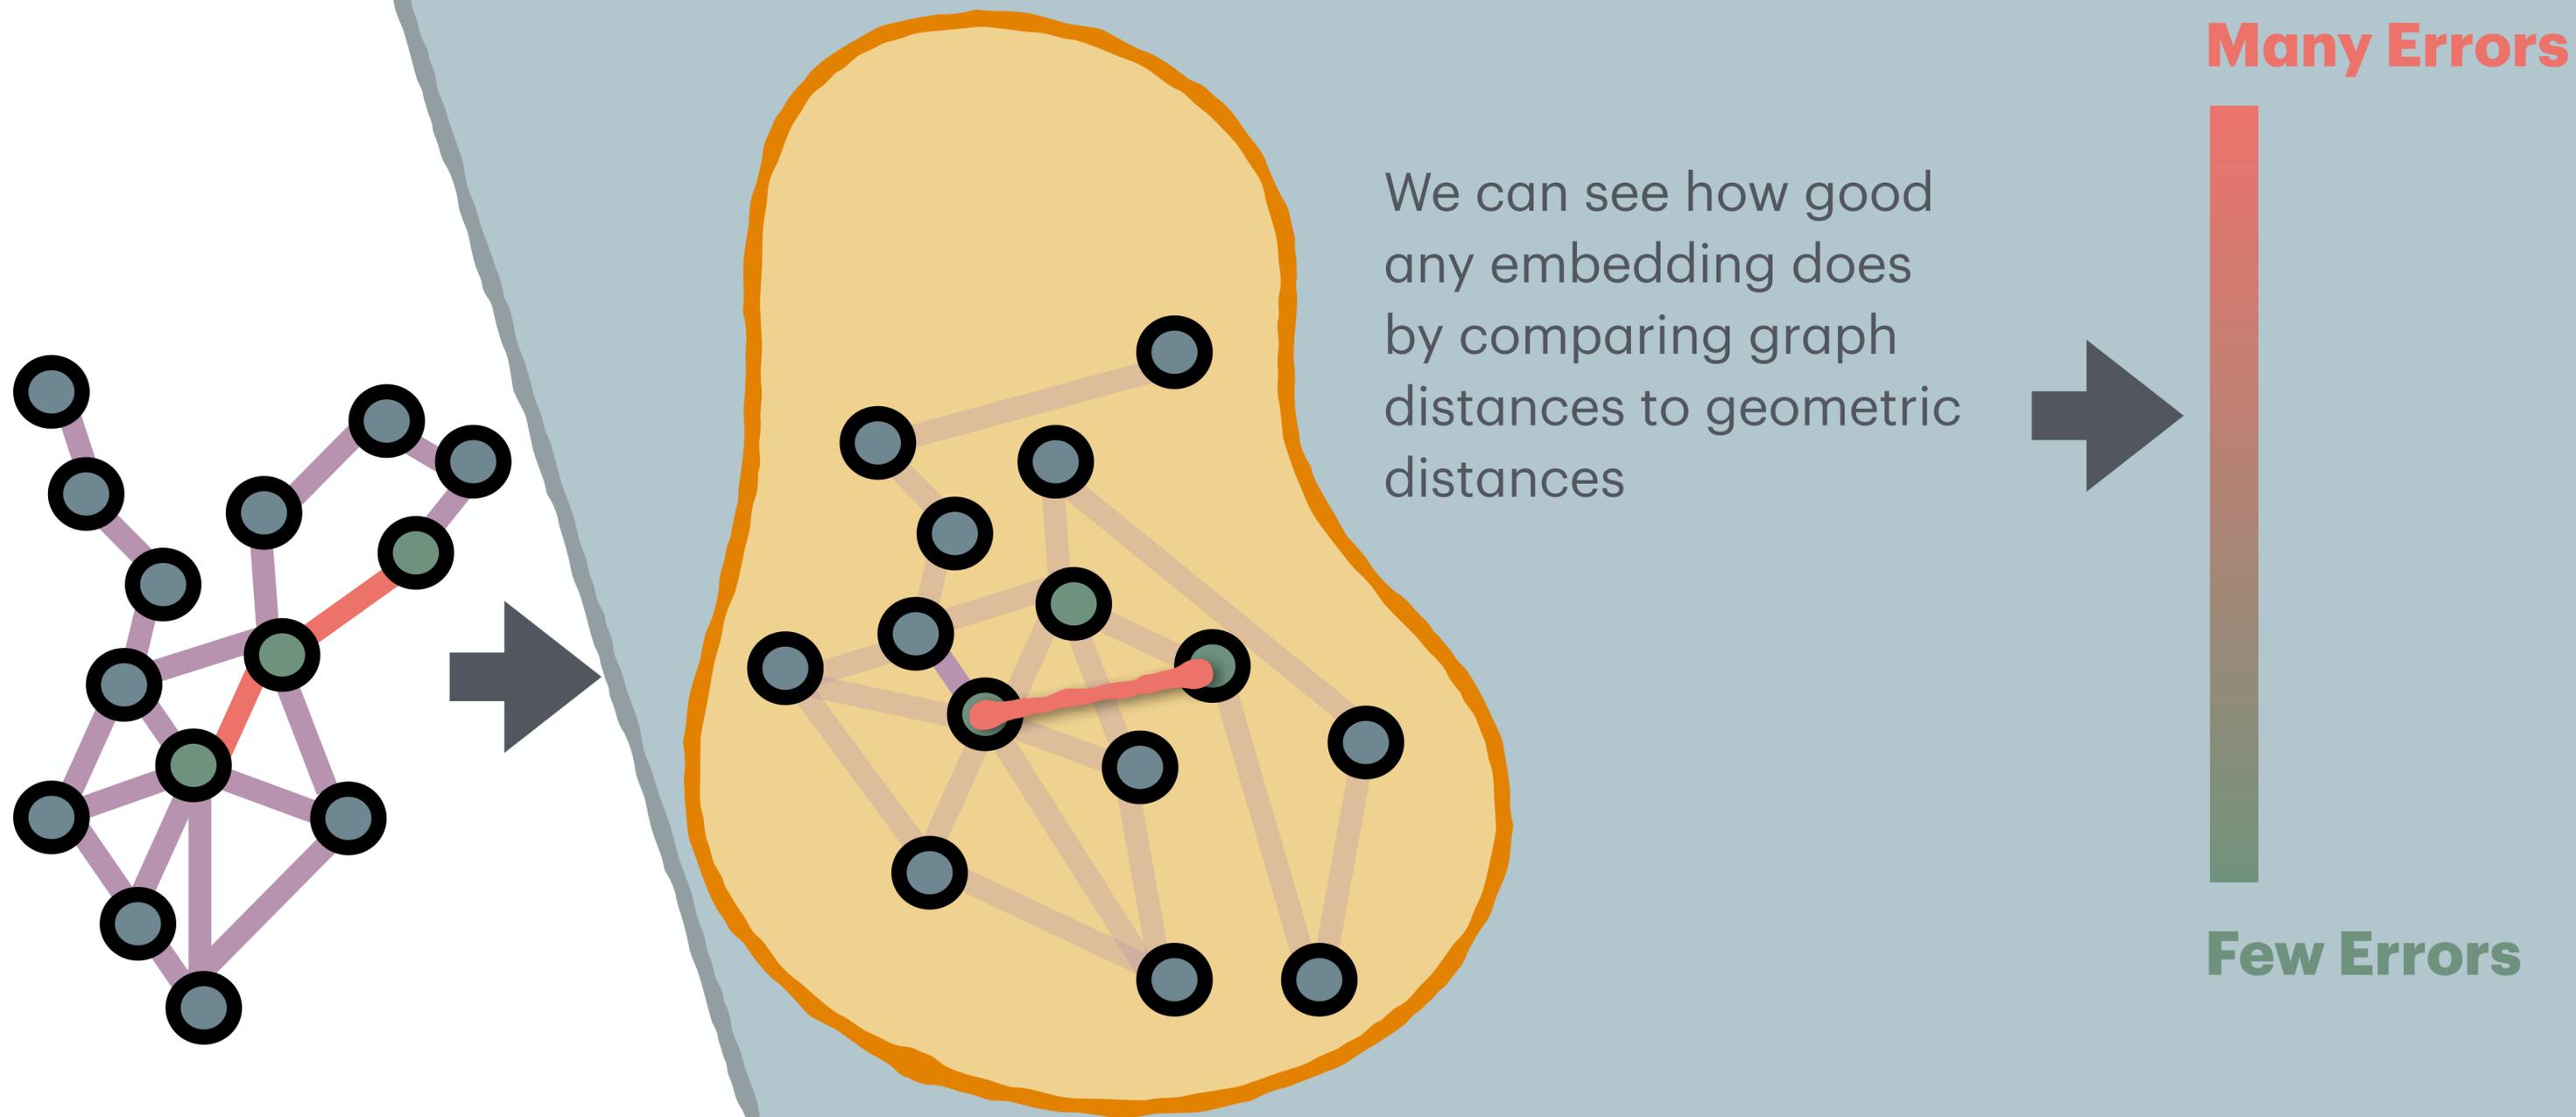
Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning



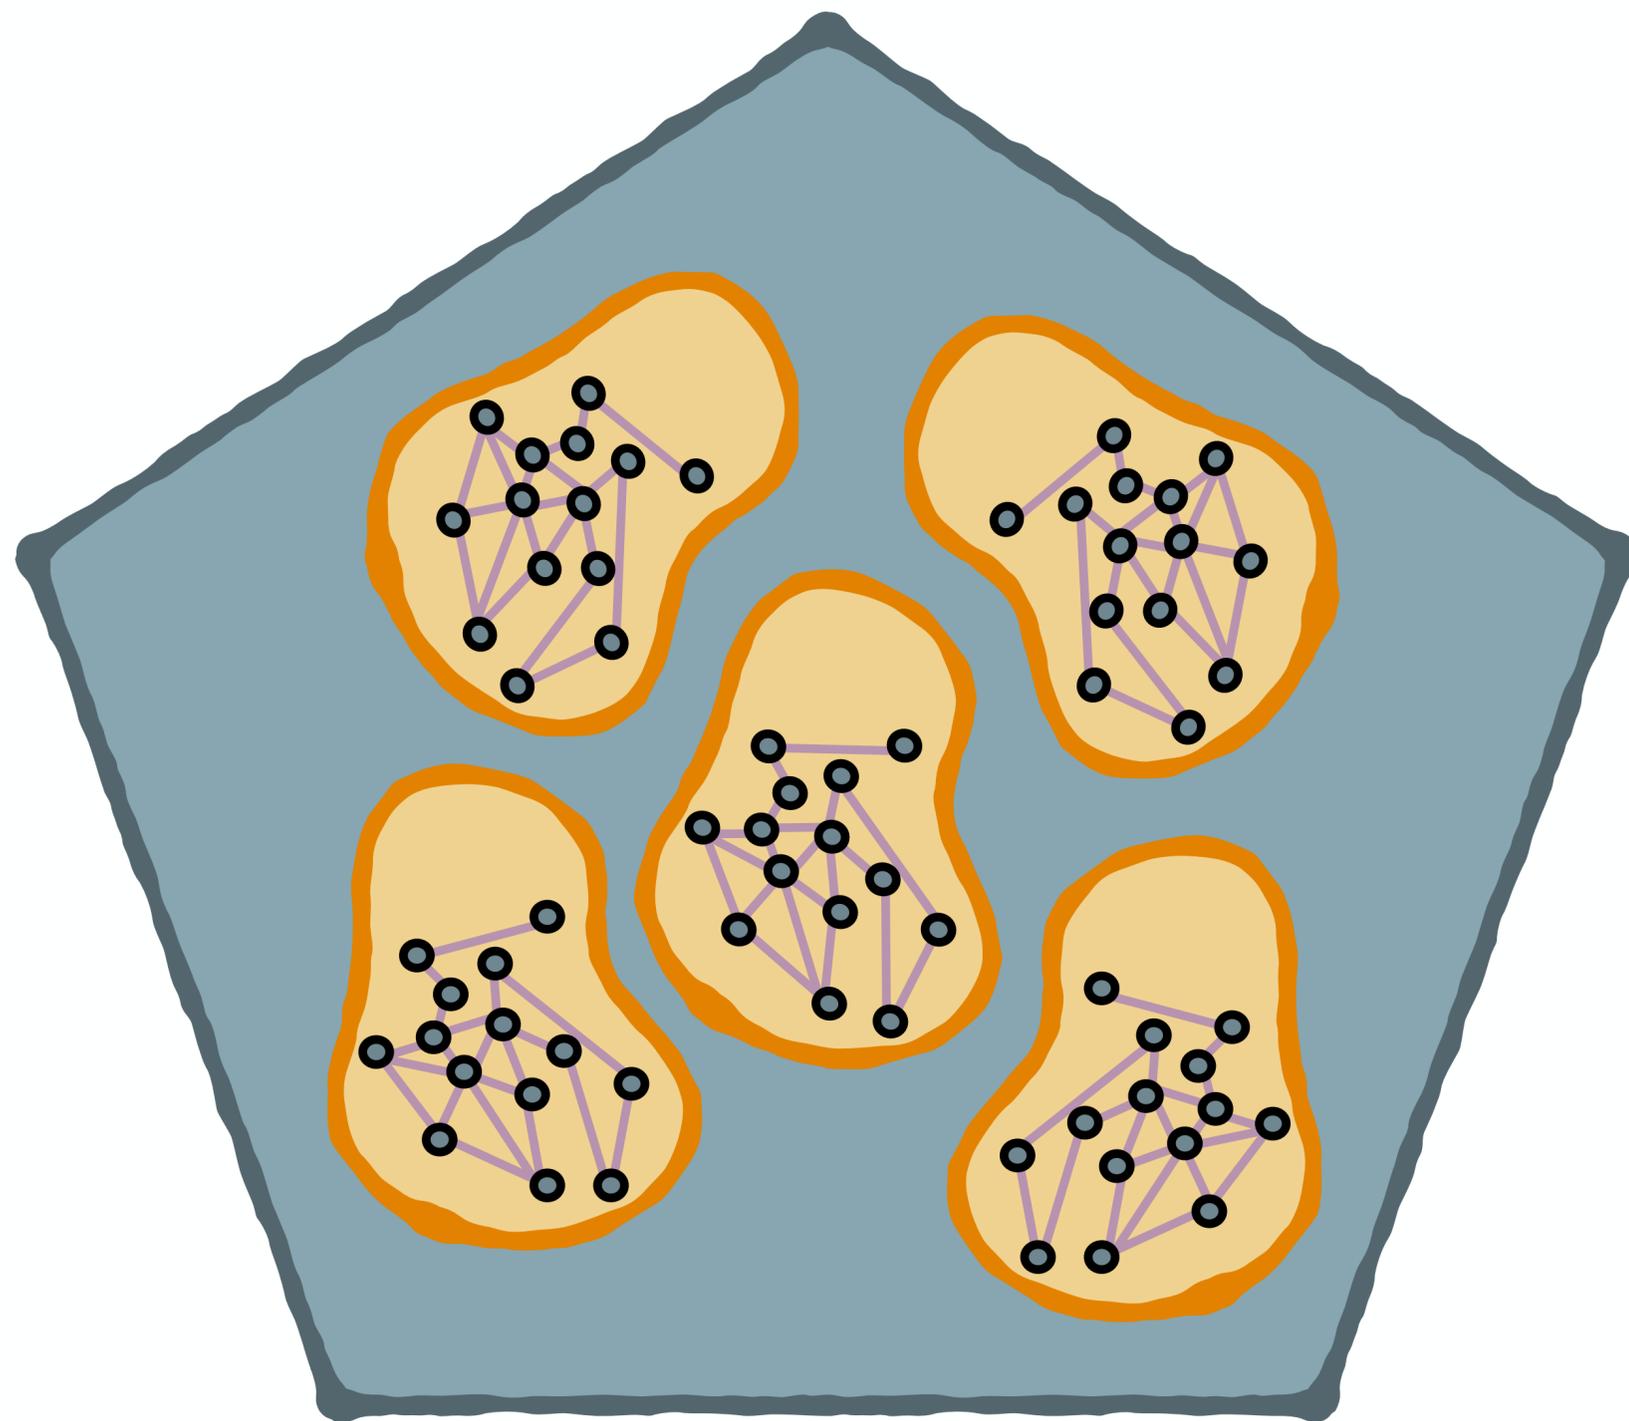
Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning



Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning



This defines a real valued function on the space of embeddings



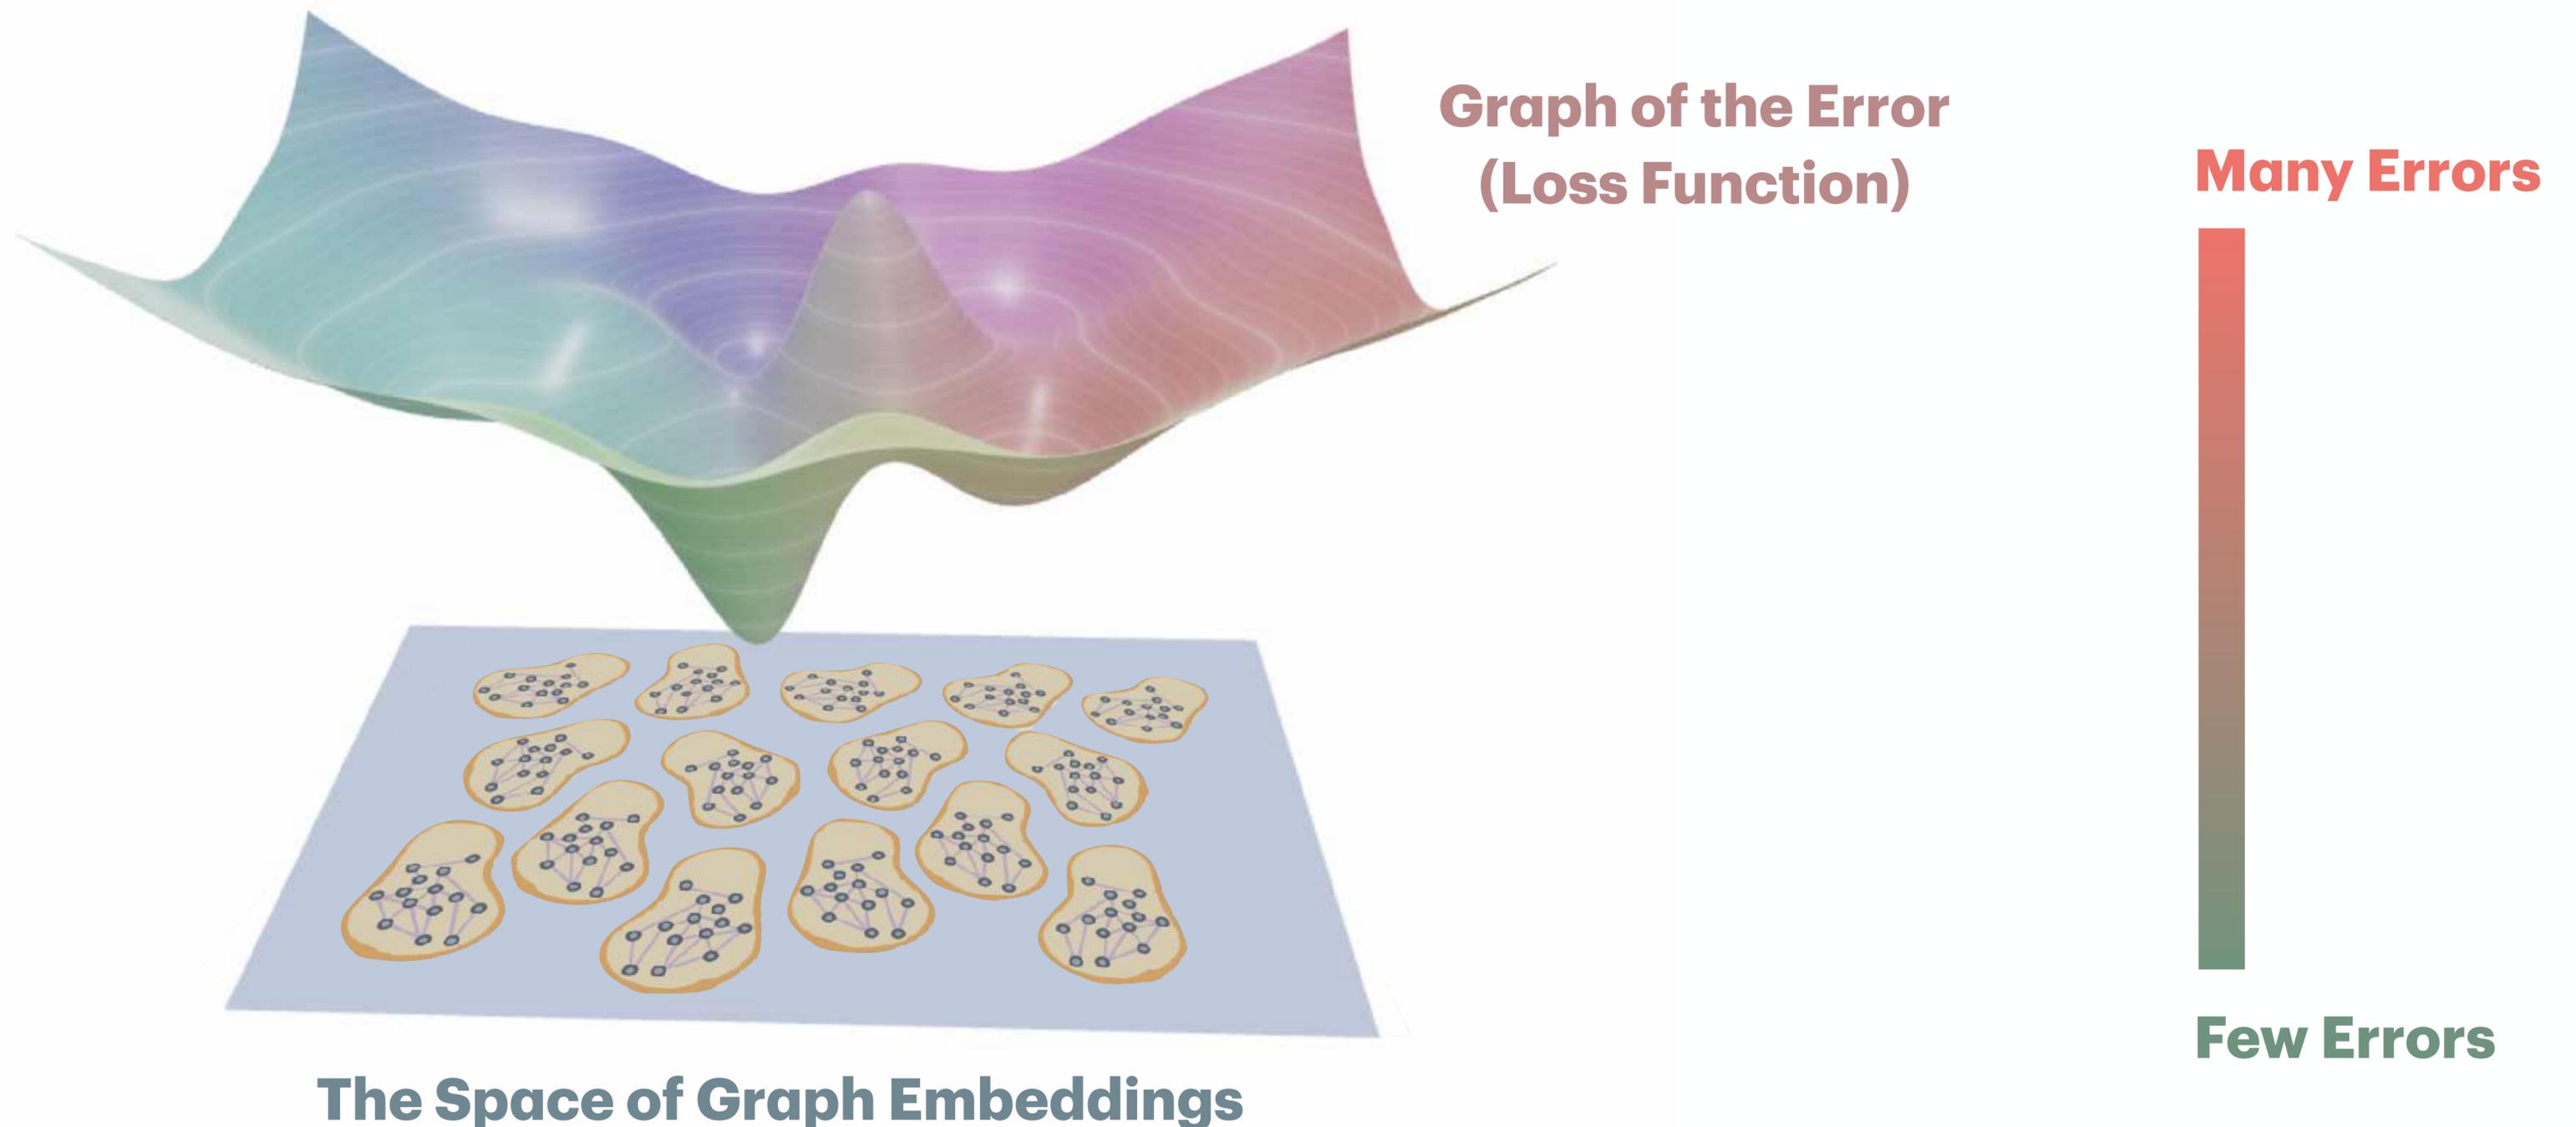
Many Errors



Few Errors

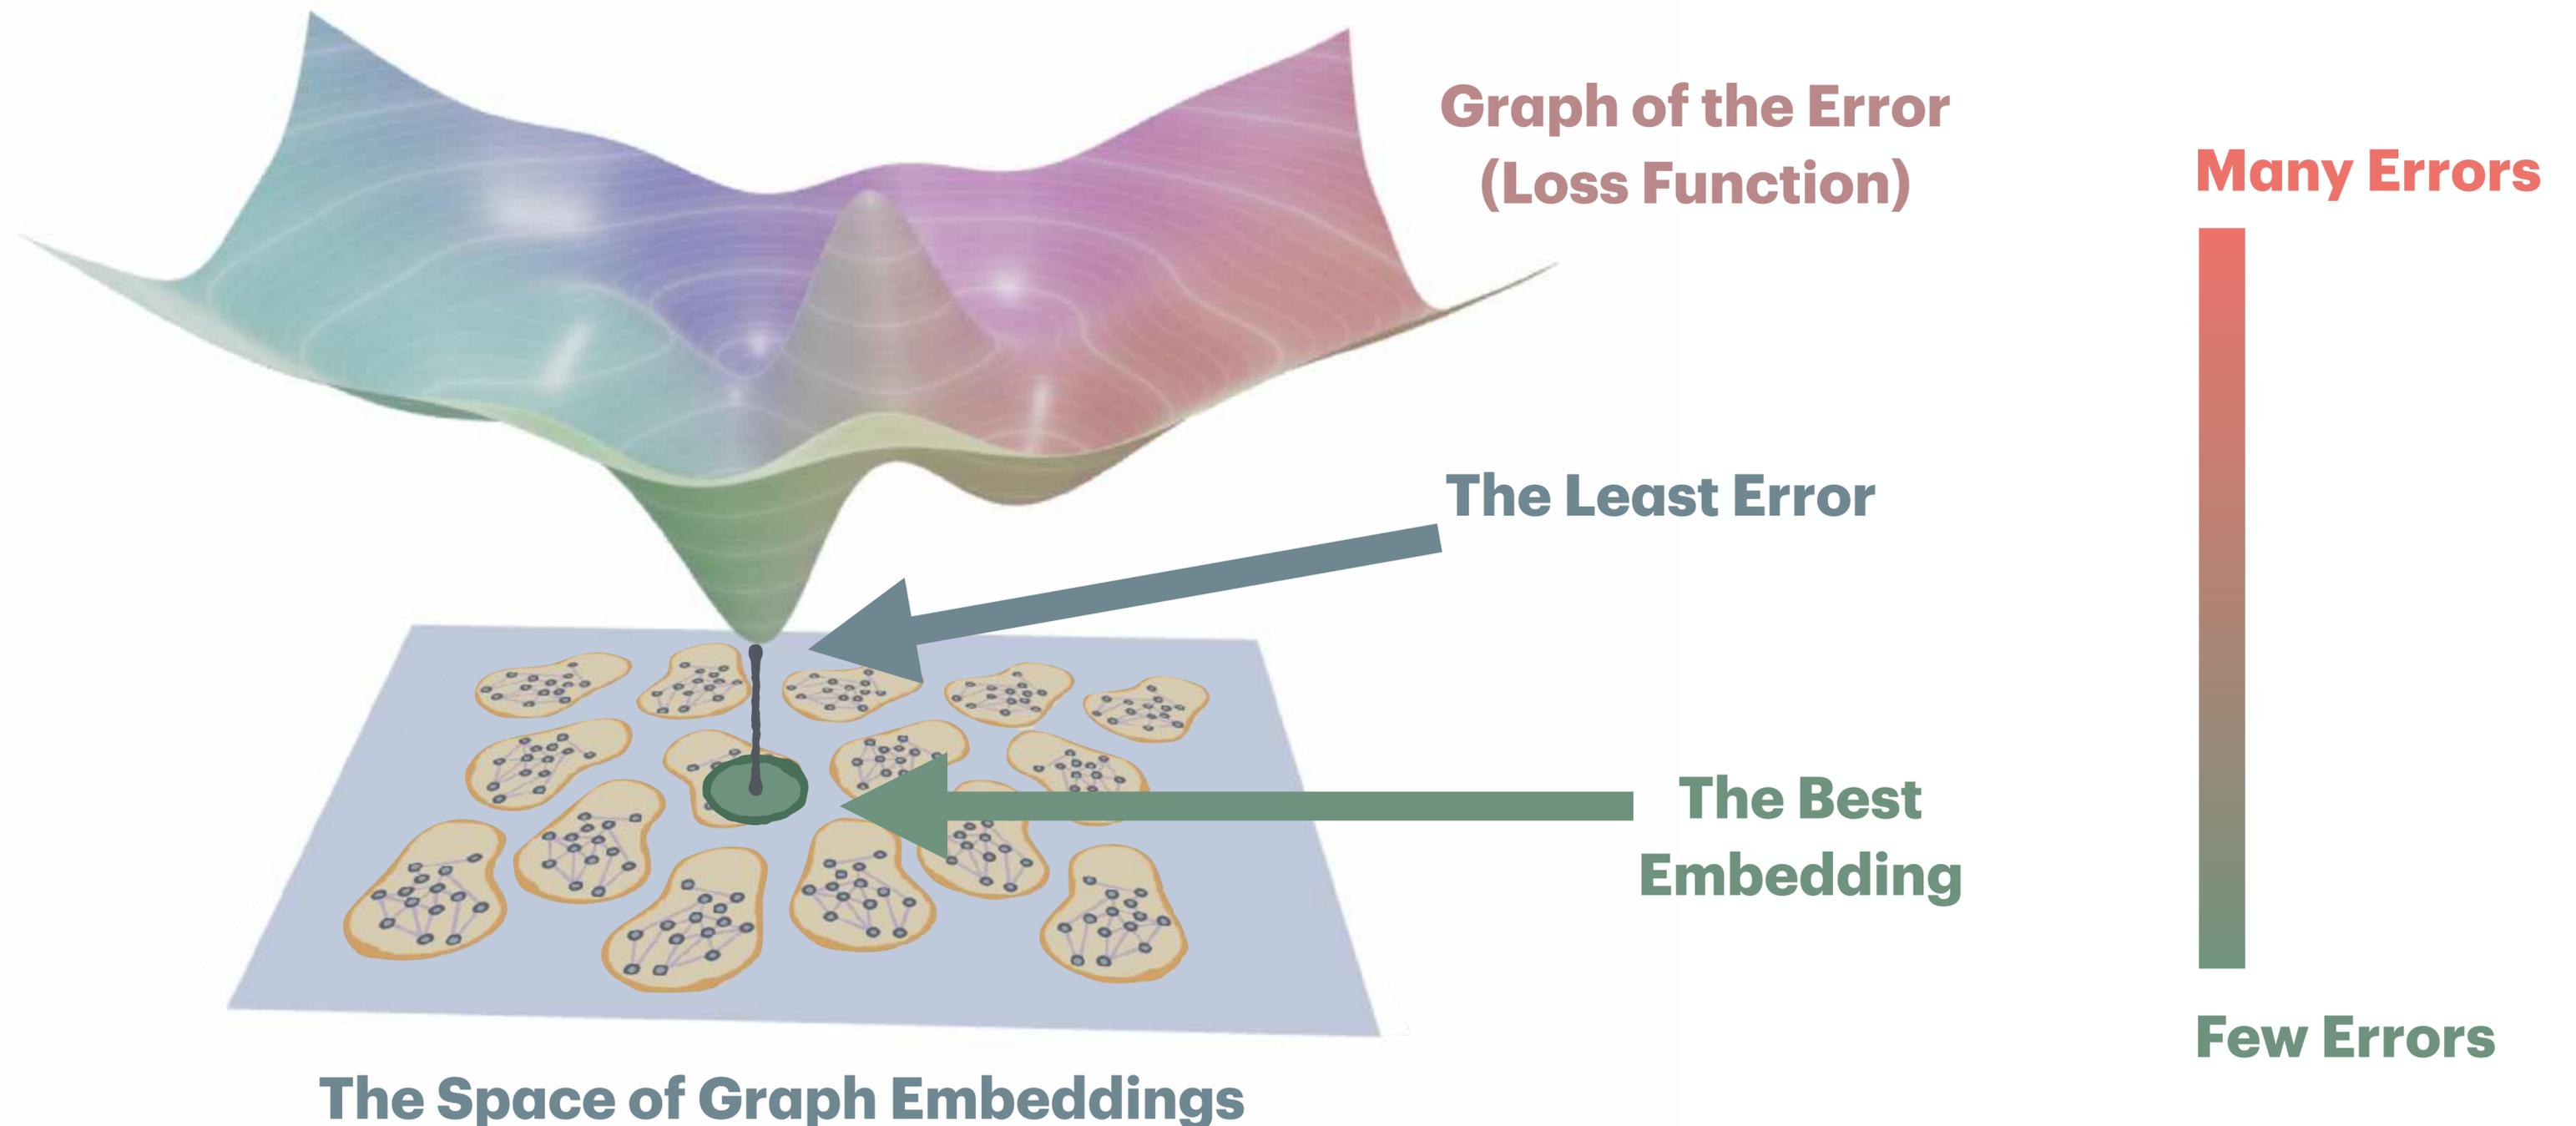
Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning



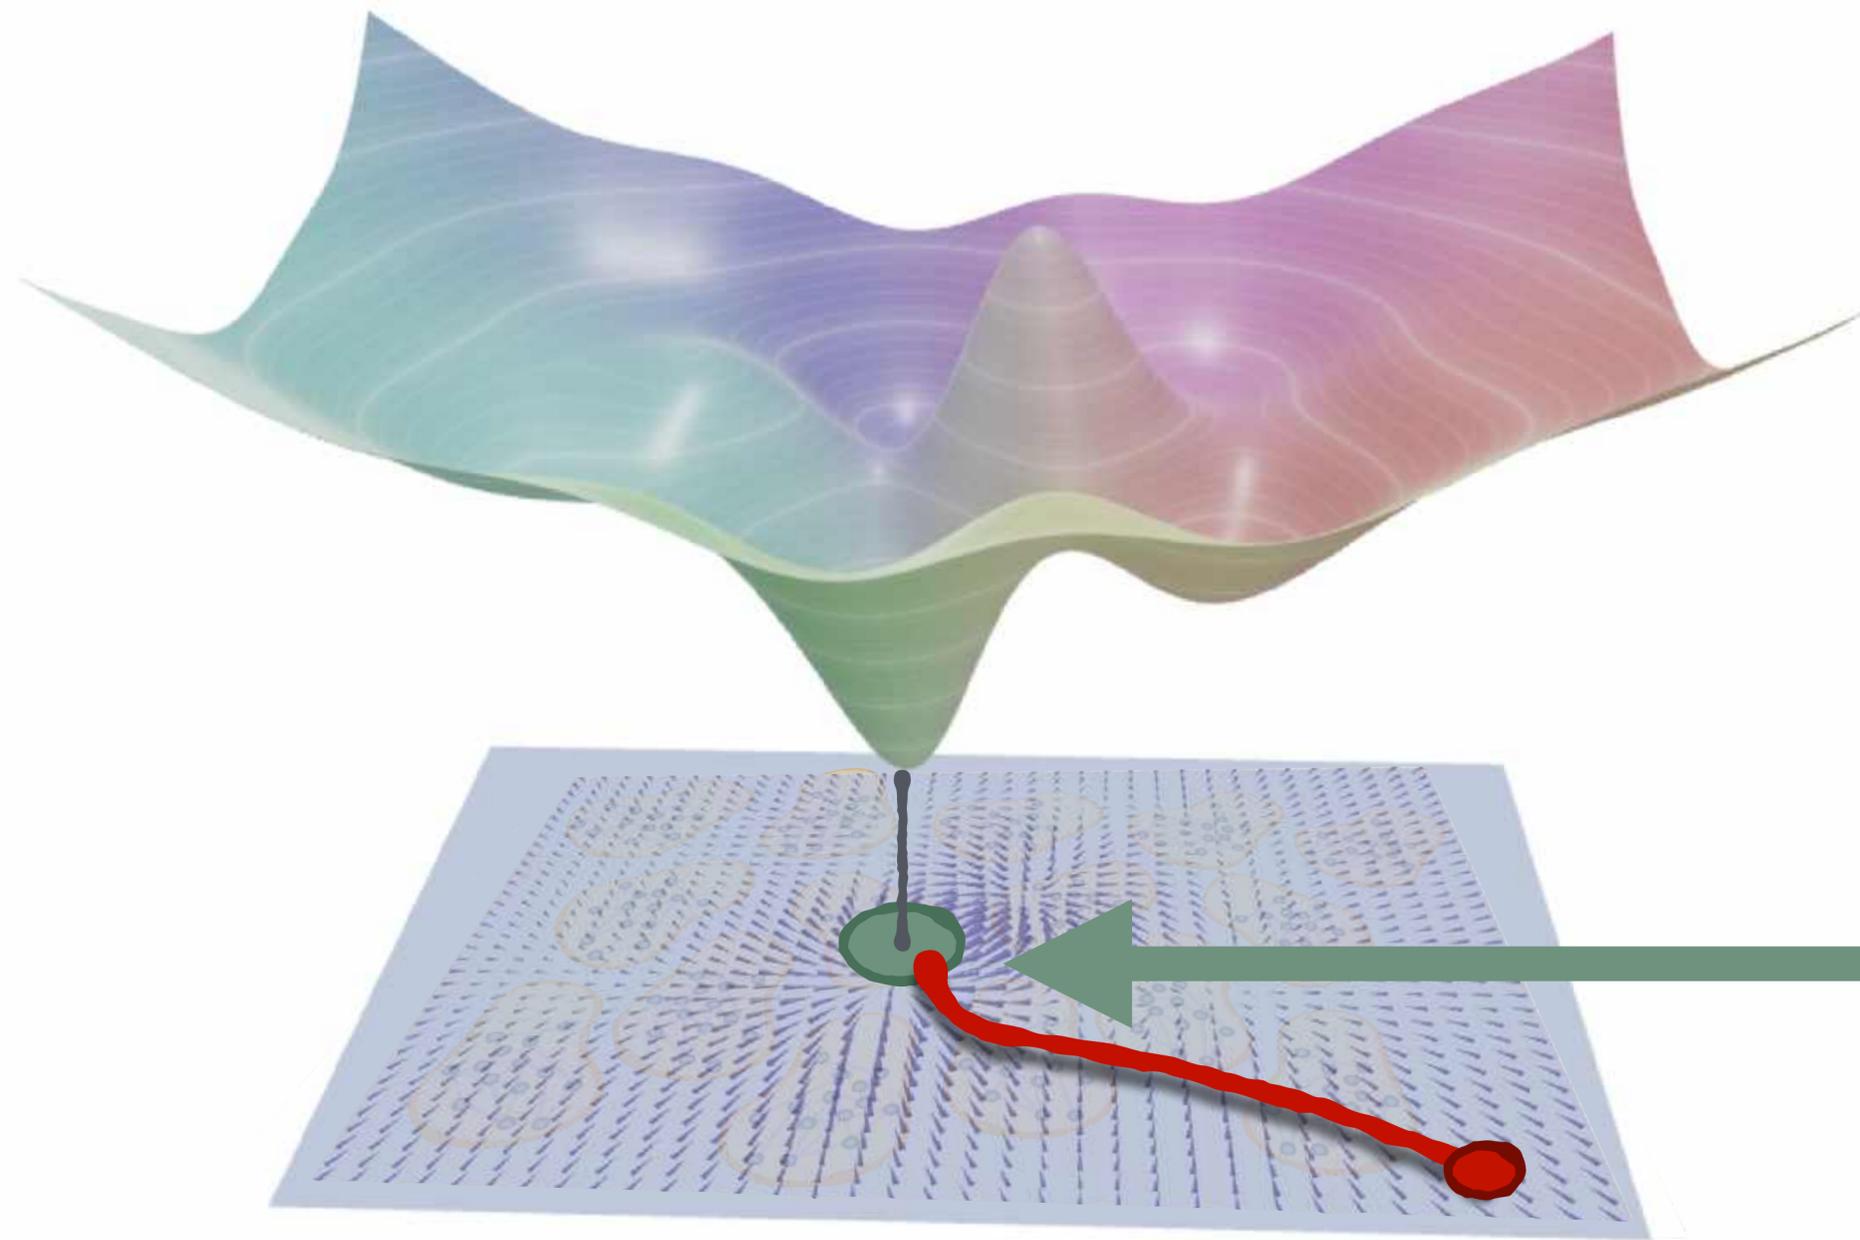
Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning



Embedding graphs in space

Beautiful Ideas at the Heart of Machine Learning



The Space of Graph Embeddings

Gradient Descent

– ∇ Error points in the direction of *steepest decrease*.

Starting at a random point, following the gradient is one tool to help find a minimum

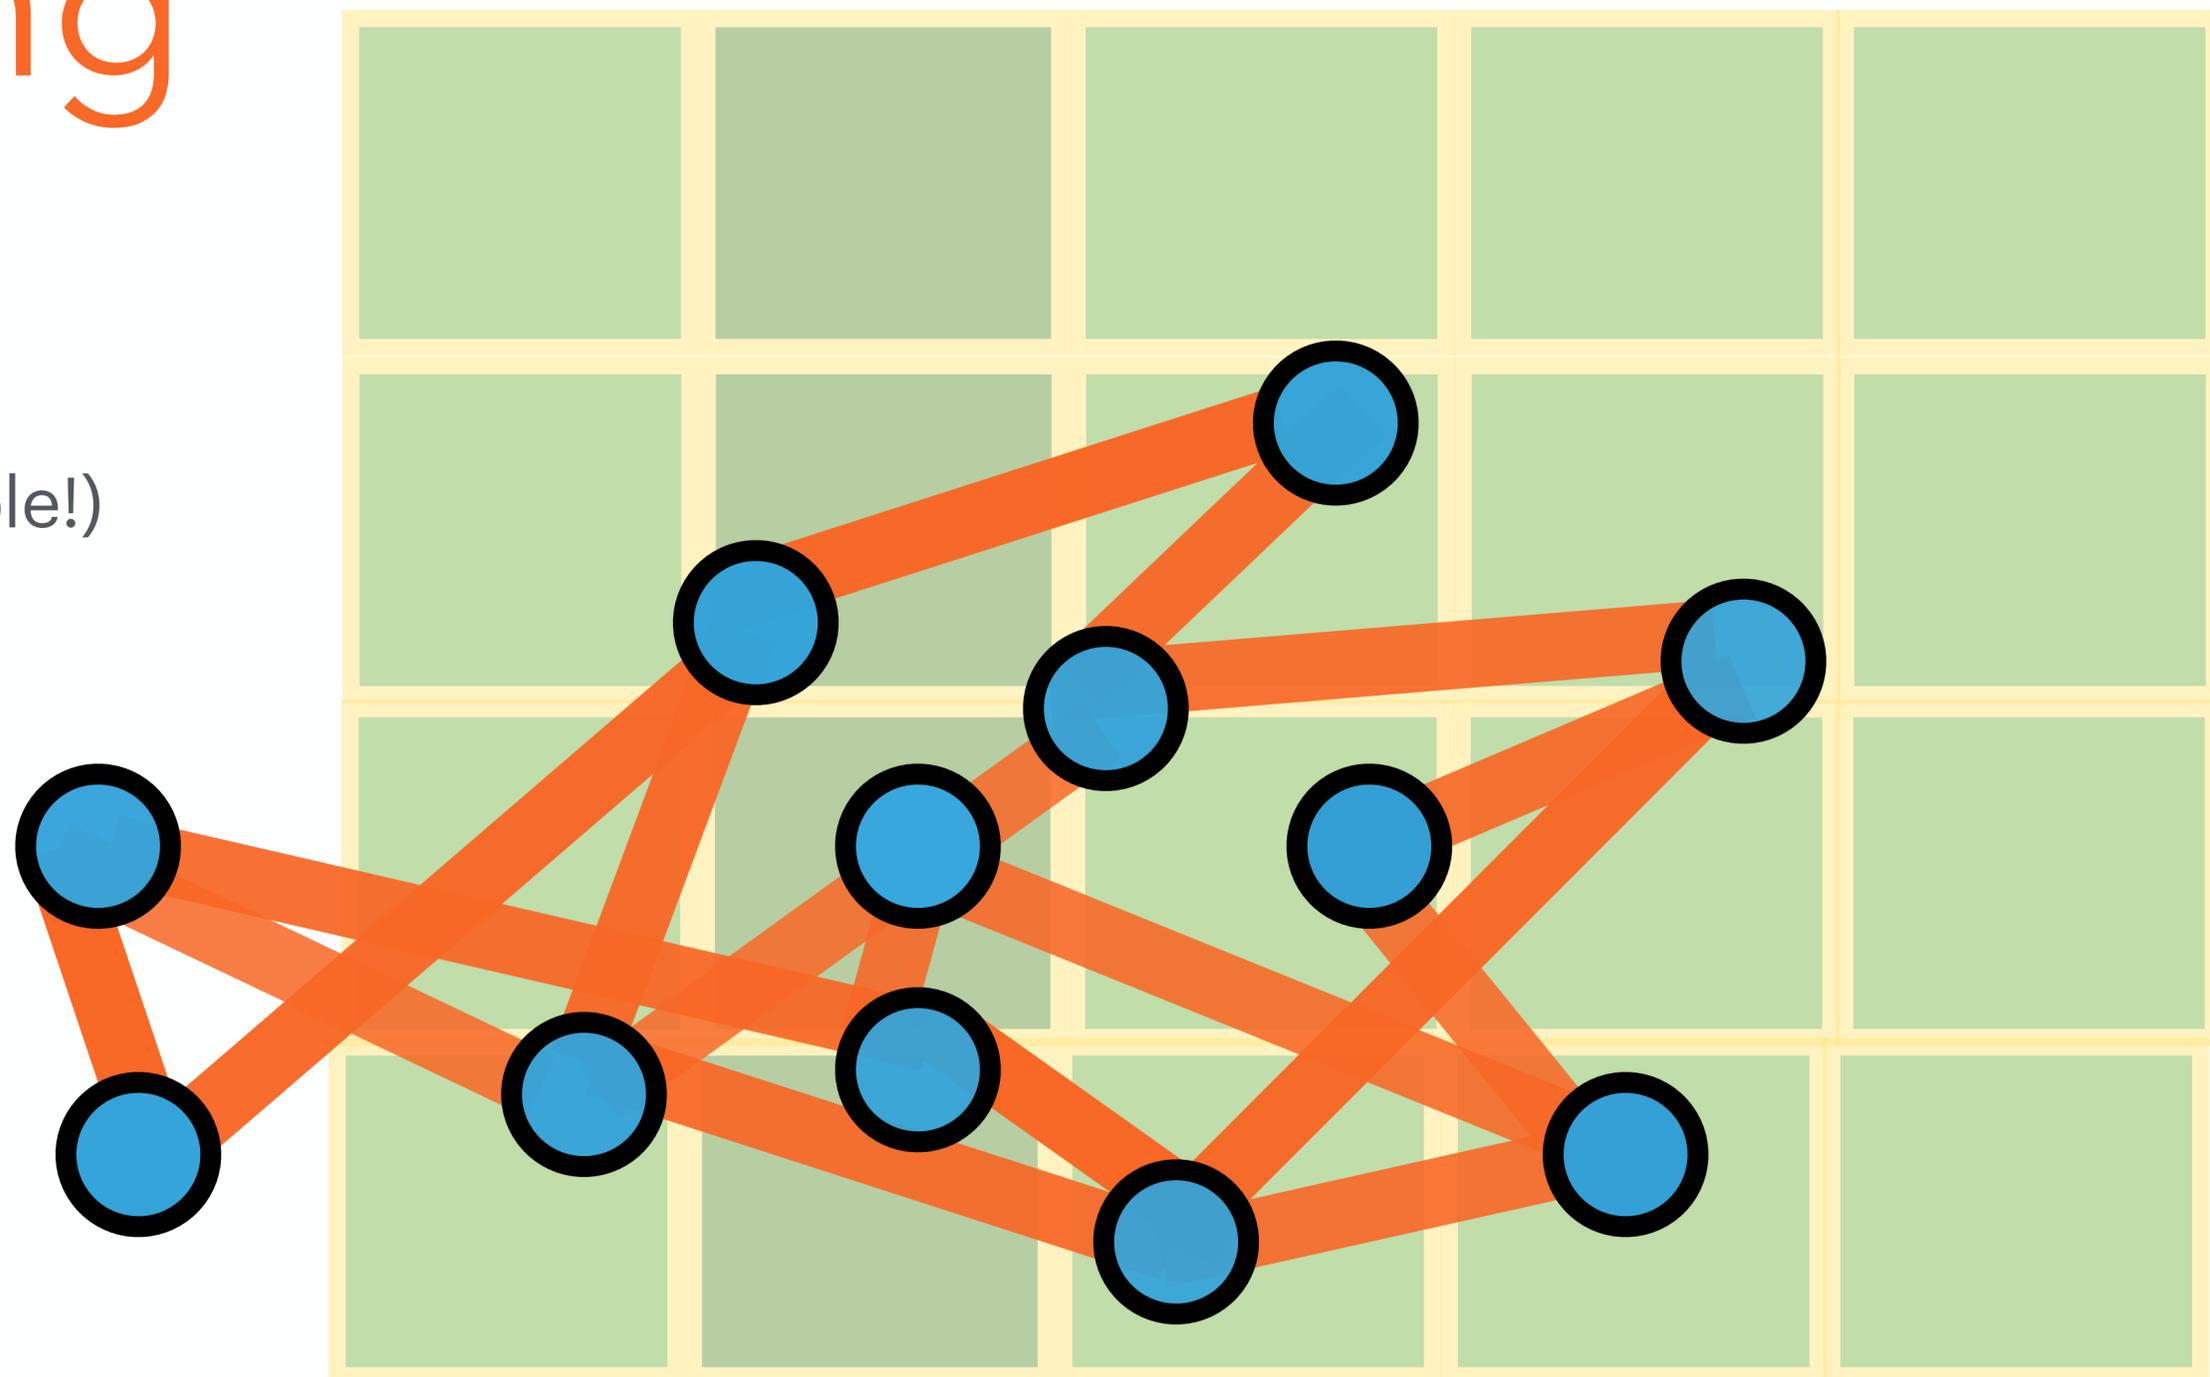
The Best Embedding

Graph Embedding

An example:

Randomly assign a point to each vertex, giving a (terrible!) embedding.

Run Gradient Descent



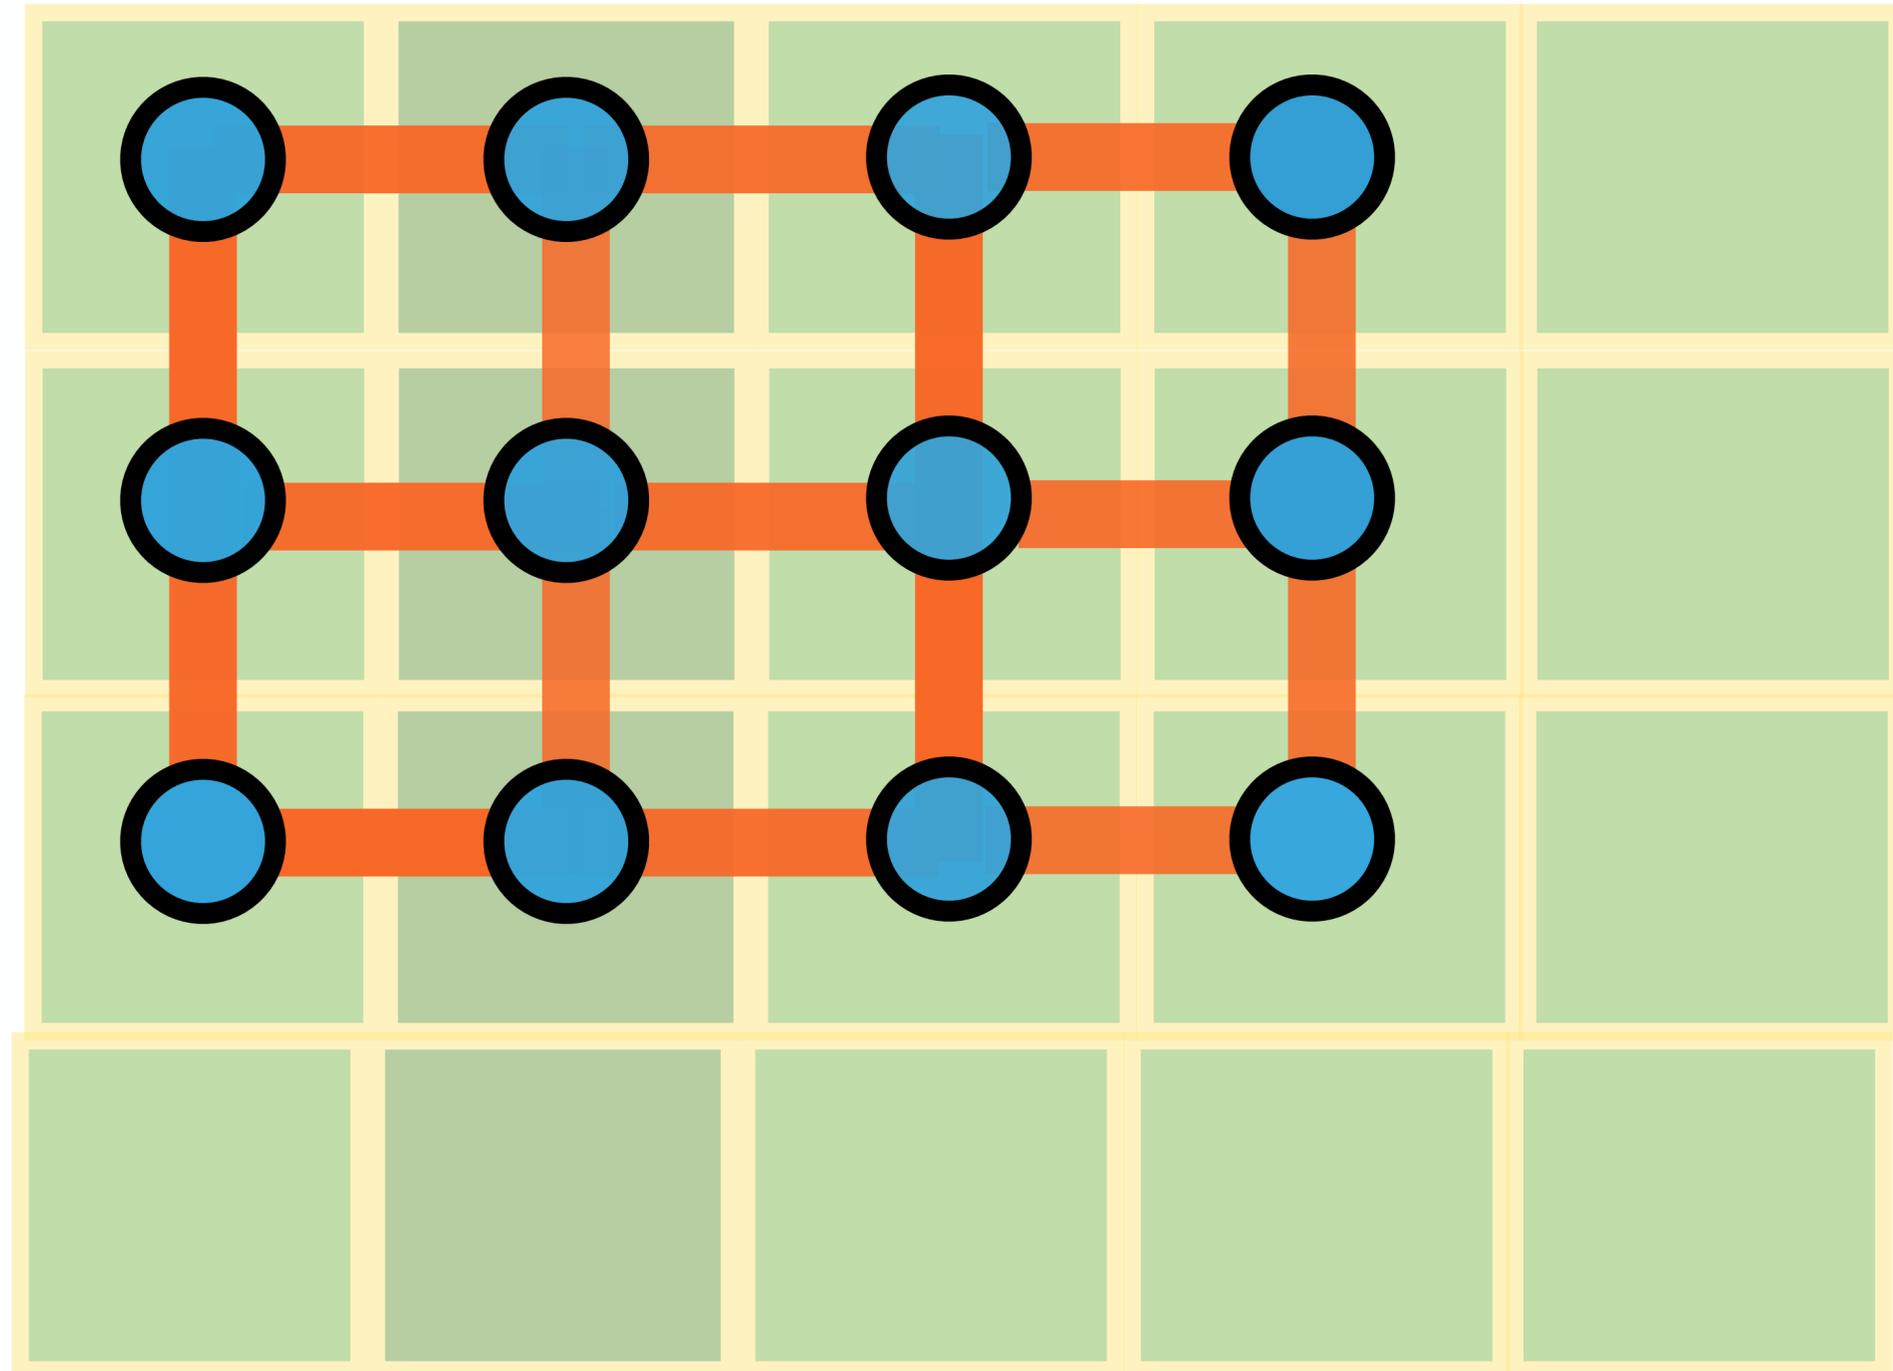
Graph Embedding

An example:

Randomly assign a point to each vertex, giving a (terrible!) embedding.

Run Gradient Descent

Produce an optimal embedding.



Graph Embedding

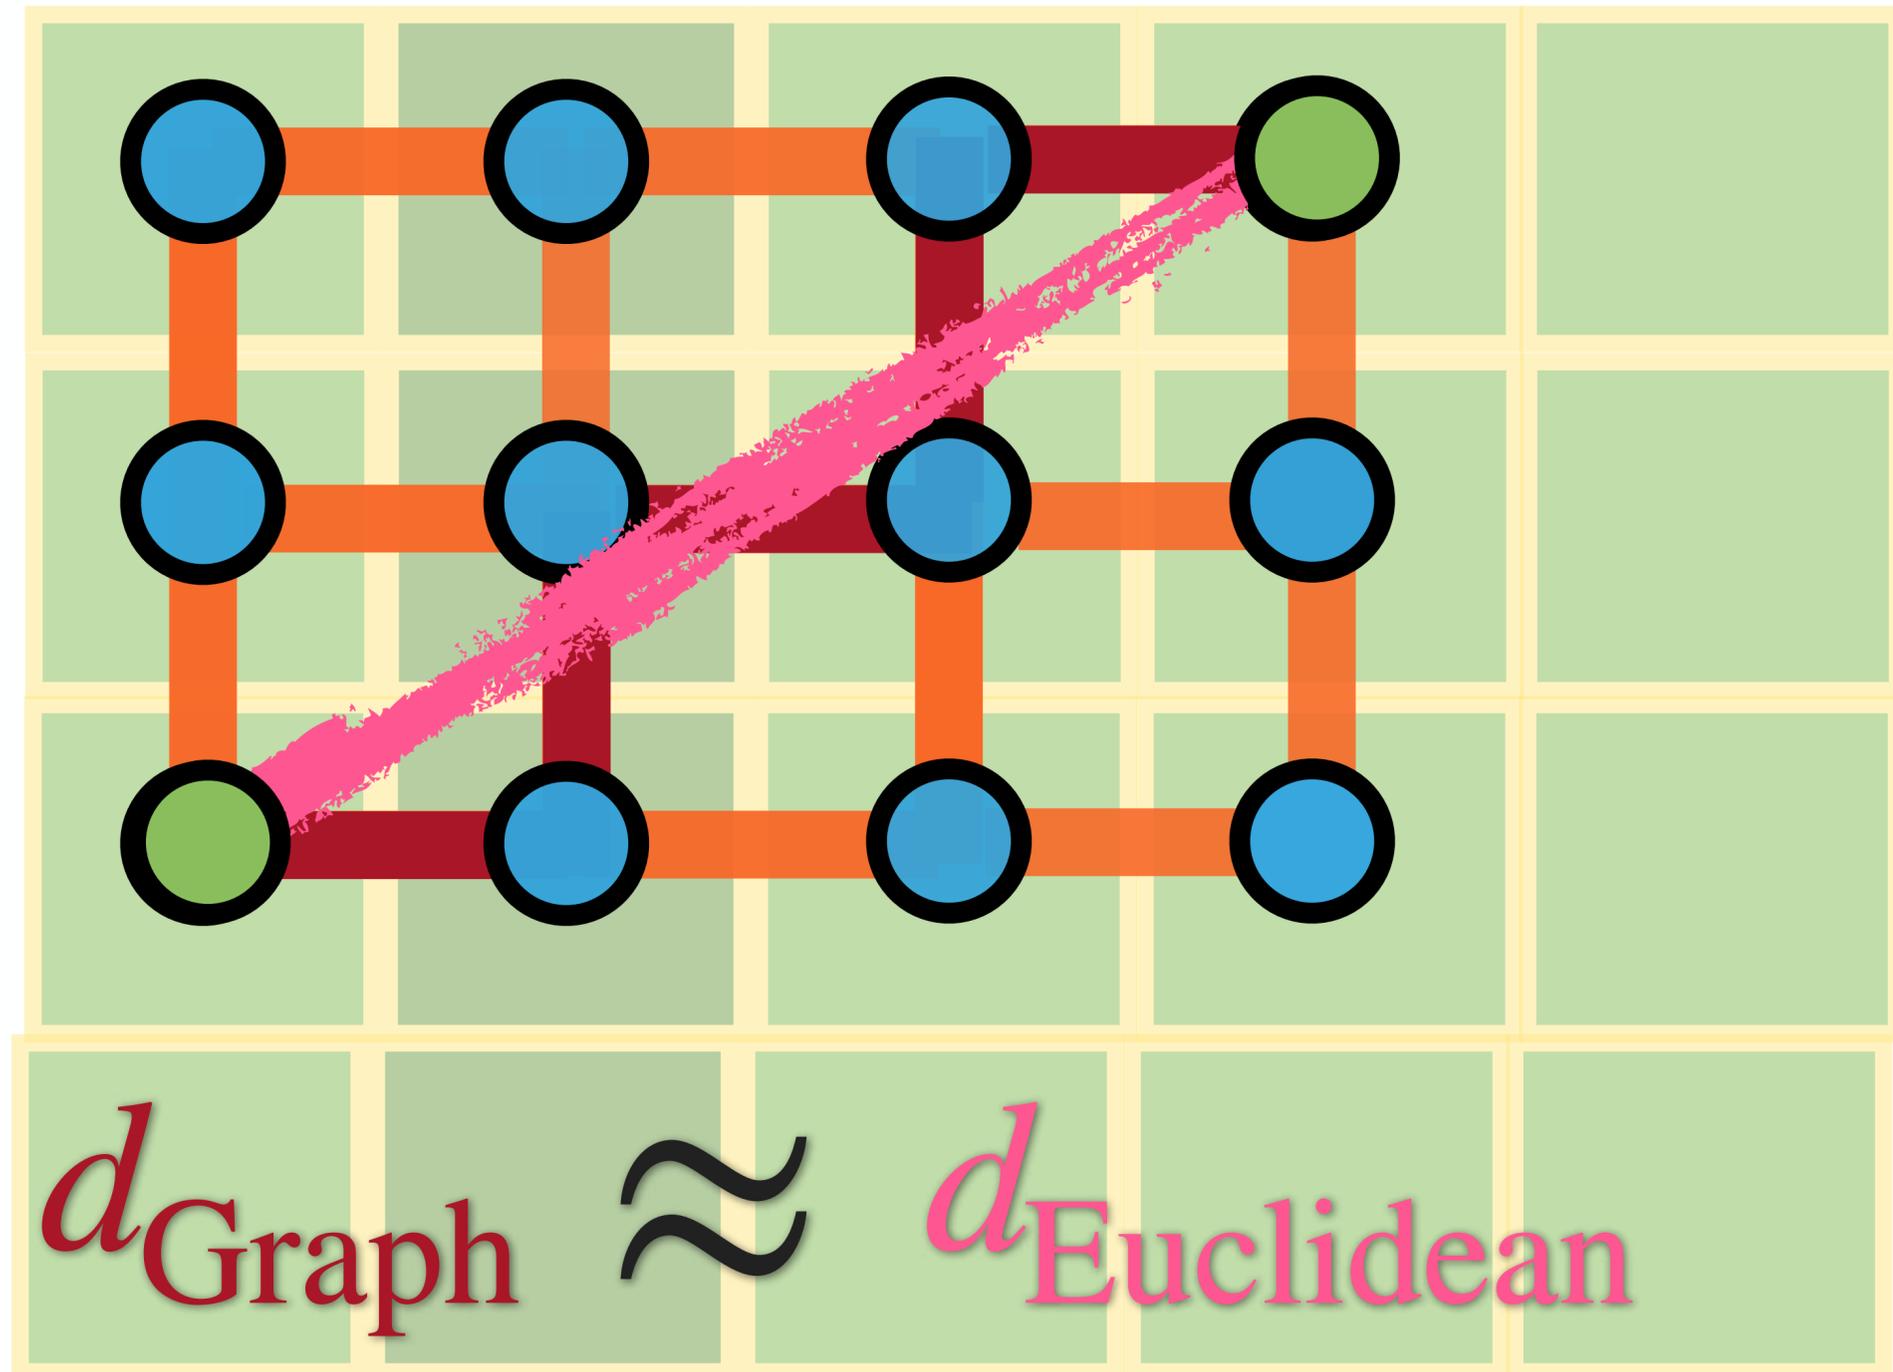
An example:

Randomly assign a point to each vertex, giving a (terrible!) embedding.

Run Gradient Descent

Produce an optimal embedding.

Using this embedding,
(approximate) graph
computations are easy!



Graph Embedding

An example:

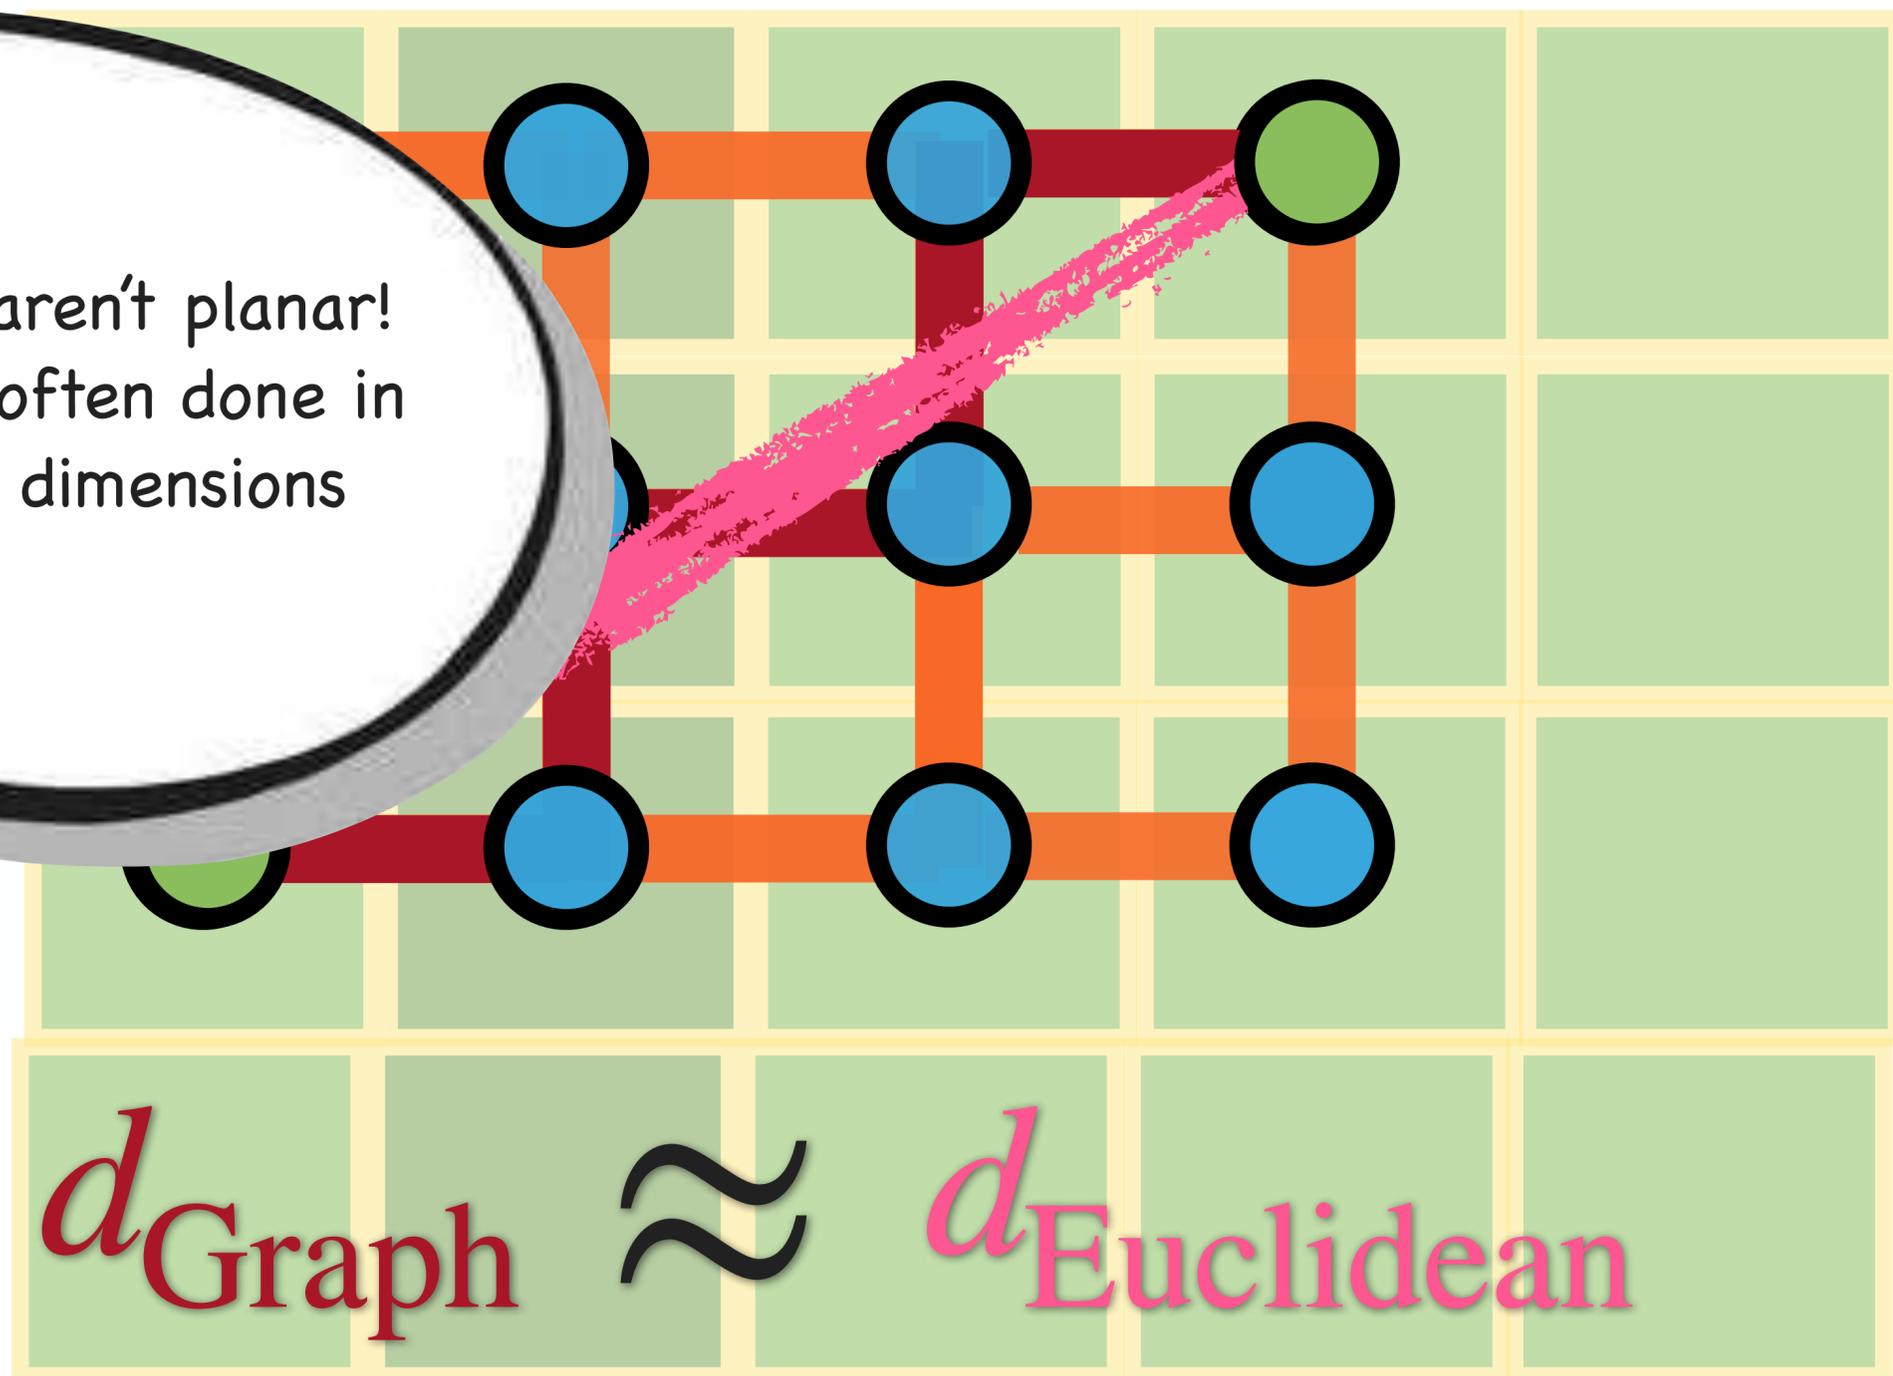
Randomly assign
each vertex a
embedding

Run Gradient Descent

Produce a minimal
embedding

Using this embedding,
(approximate graph
computations are easy!)

Most graphs aren't planar!
Embedding is often done in
hundreds of dimensions



Palash Goyal and Emilio
University of Southern California, Informa
4676 Admiralty Way, Suite 1001, Marina del Rey, CA 90265

A Comprehensive Survey of Graph Embedding Problems, Techniques and Applications

Hongyun Cai, Vincent W. Zheng, and Kevin Chen-Chuan Chang

appears in a wide diversity of real-world scenarios to analyze and understand the data, and thus can benefit a lot of users. However, most graph analytics methods suffer from the curse of dimensionality to solve the graph analytics problem. It often requires that node and graph properties are maximumly preserved in the embedding. We first introduce the formal definition of graph embedding.

A Tutorial on Network Embeddings

Haochen Chen¹, Bryan Perozzi², Rami Al-Rfou², and Steven Skiena¹

¹Stony Brook University

²Google Research

{haocchen, skiena}@cs.stonybrook.edu, bperozzi@acm.org, rmyeid@google.com

August 9, 2018

Abstract

Network embedding methods aim at learning low-dimensional latent representation of nodes

The Link-Prediction Problem for Social Networks

David Liben-Nowell
Department of Computer Science
Carleton College
Northfield, MN 55057 USA
dlibenno@carleton.edu

Jordan Kleinberg
Department of Computer Science
Cornell University
Ithaca, NY 14853
kleinber@cornell.edu

Deep Learning on Graphs

Ziwei Zhang, Peng Cui and Wenwu Zhu,

Abstract—Deep learning has been shown to be successful in a number of domains. Recently, it has been applied to graph data. In this paper, we survey the state-of-the-art in deep learning on graphs.

Network Representation Learning: A Survey

Daokun Zhang, Jie Yin, Xingquan Zhu *Senior Member, IEEE*, Chengqi Zhang *Senior Member, IEEE*

Abstract—With the widespread use of information technologies, information networks are becoming increasingly popular to capture various disciplines, such as social networks, citation networks, telecommunication networks, and transportation networks. These networks shed light on different aspects of social life such as the structure of societies, communication patterns. In reality, however, the large scale of information networks often makes network representation learning expensive or intractable. Network representation learning has been recently proposed as a new learning paradigm that maps nodes into a low-dimensional vector space, by preserving network topology structure, vertex content, and other side information. This survey reviews the current literature on network representation learning in the data mining and machine learning communities to categorize and summarize the state-of-the-art network representation learning methods, underlying learning mechanisms, the network information intended to preserve, as well as the algorithmic complexity. We also summarize evaluation protocols used for validating network representation learning including evaluation methods, and open source algorithms. We also perform empirical studies to compare the performance of different algorithms on common datasets, and analyze their computational complexity. Finally, we suggest some future research directions to facilitate future study.

Keywords: network representation learning, network embedding, graph mining, network representation learning, network embedding.

Abstract

Given a snapshot of a social network, can we infer which new links are likely to occur in the near future? We formalize this question and we develop approaches to link prediction based on measures of nodes in a network. Experiments on large co-authorship networks show that measures for detecting node proximity can outperform more complex methods.

1 Introduction

As part of the recent surge of research on large, complex networks, a considerable amount of attention has been devoted to the computation of network structures whose nodes represent people or other entities embedded in a network. Edges represent interaction, collaboration, or influence between nodes. Networks include the set of all scientists in a particular discipline; the set of all co-authored papers; the set of all employees in a large company; a collection of business leaders, with their interactions; or a collection of people who have met together on a corporate board of directors. The increased availability of such networks has stimulated extensive study of the structure and dynamics of these networks. (See, for example, the survey by Watts [37], Grossman [11], Newman [29], and Adamic and Glance [30].)

Social networks are highly dynamic objects; they grow and change as new nodes are added and edges are formed. Identifying the mechanisms by which they evolve is a central problem in network science, and it forms the motivation for our work.

Representation Learning on Graphs: Methods and Applications

William L. Hamilton
wleif@stanford.edu

Rex Ying
rexying@stanford.edu

Jure Leskovec
jure@cs.stanford.edu

Department of Computer Science
Stanford University
Stanford, CA, 94305

Abstract

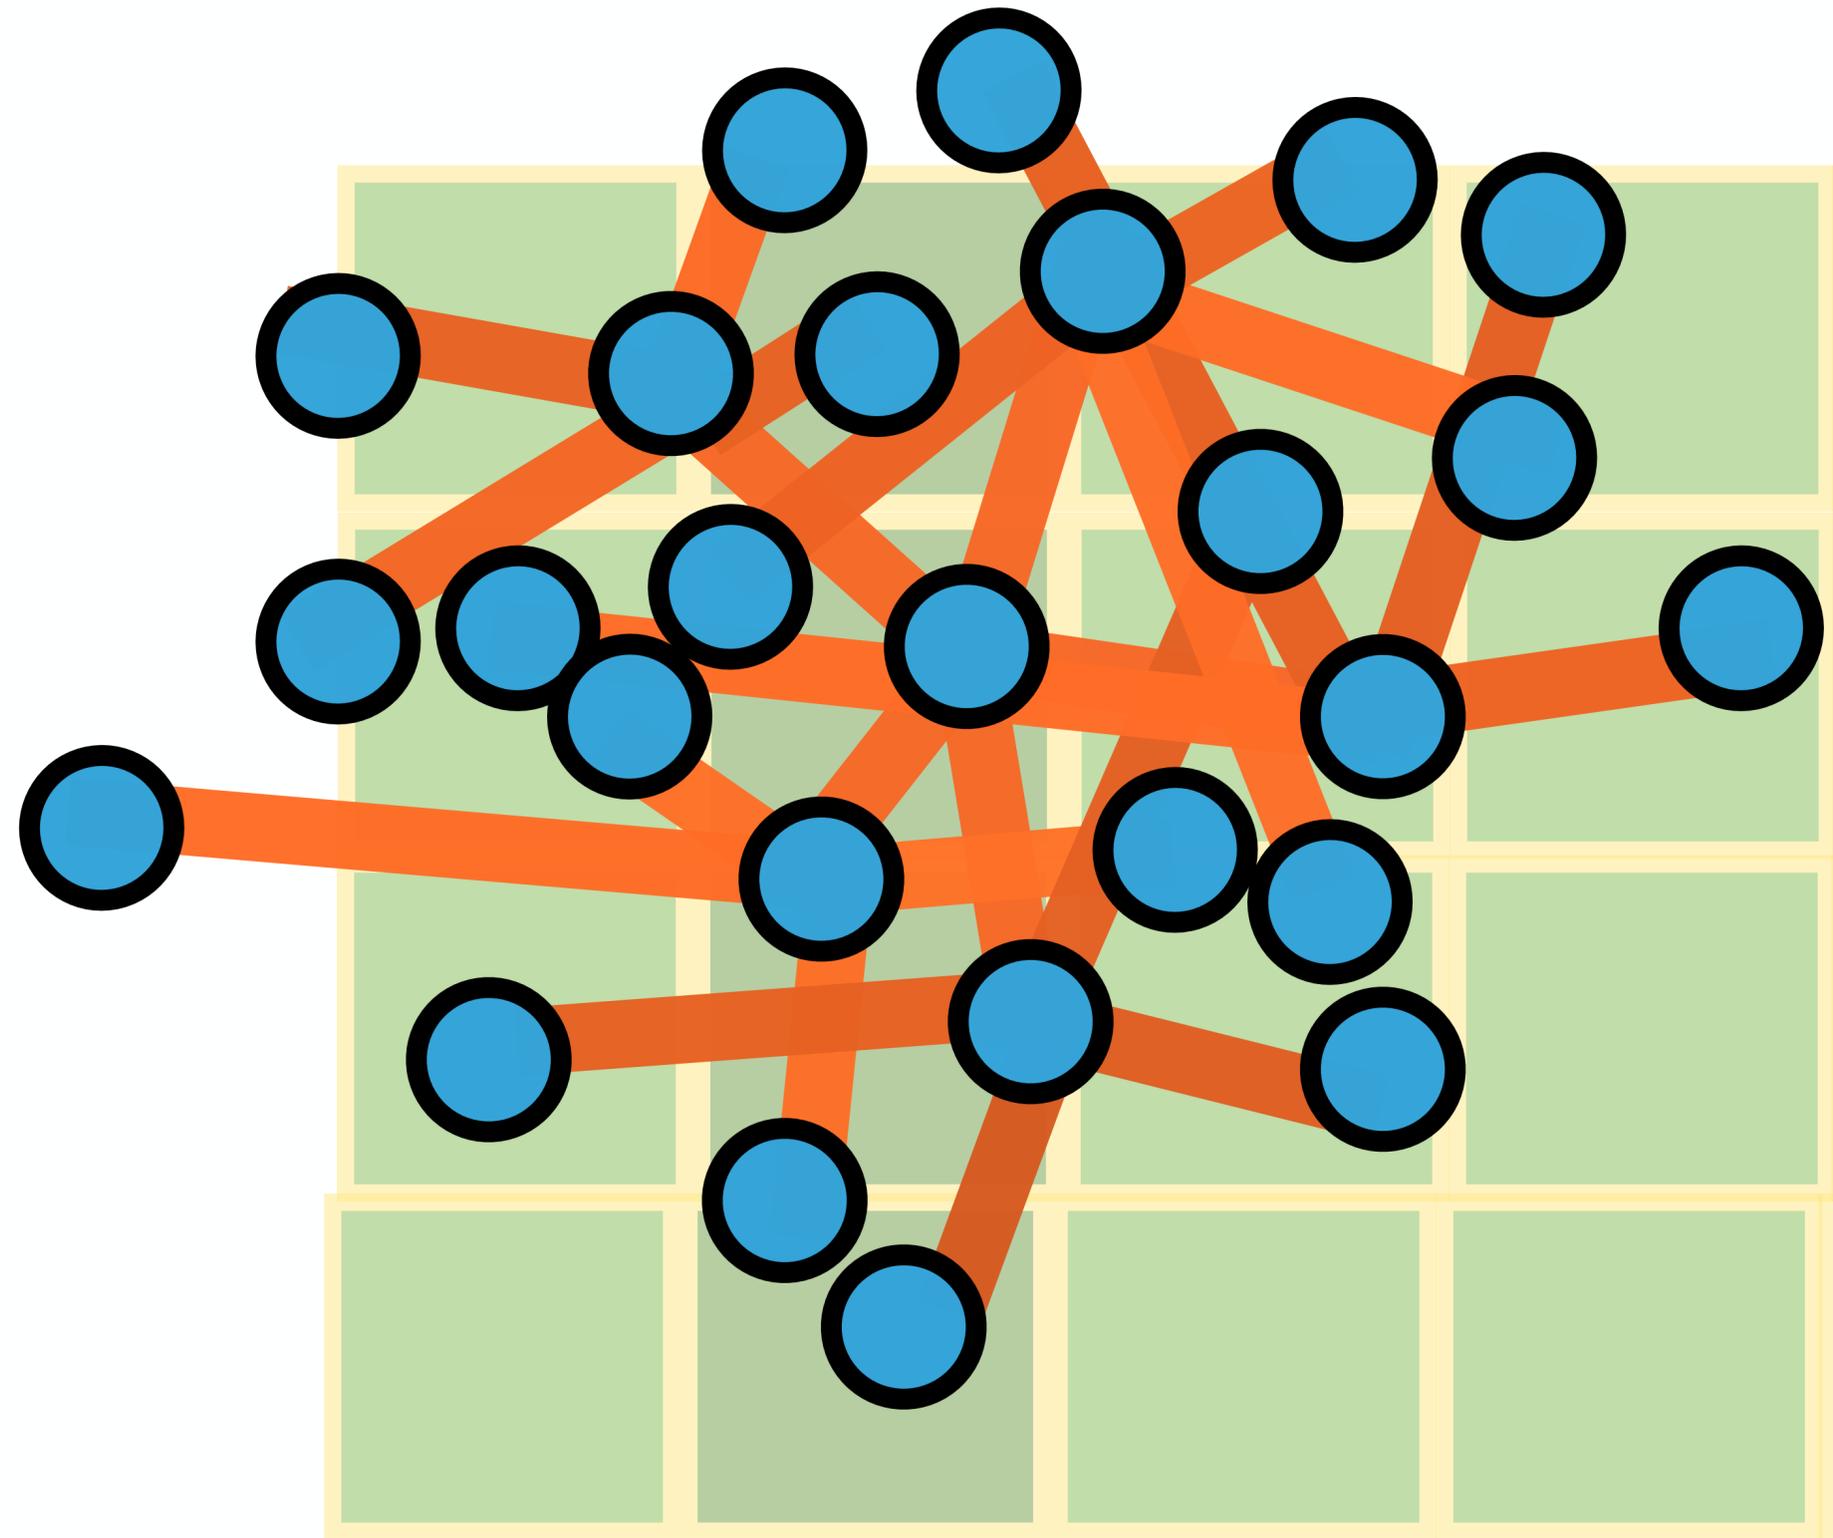
Machine learning on graphs is an important and ubiquitous task with applications ranging from drug design to friendship recommendation in social networks. The primary challenge in this domain is finding a way to represent, or encode, graph structure so that it can be easily exploited by machine learning models. Traditionally, machine learning approaches relied on user-defined heuristics to extract features encoding structural information about a graph (e.g., degree statistics or kernel functions). However, recent years have seen a surge in approaches that automatically learn to encode graph structure into low-dimensional embeddings, using techniques based on deep learning and nonlinear dimensionality reduction. Here we provide a conceptual review of key advancements in this area of representation learning.

[cs.SI] 10 Apr 2018

Nov 10 10:51 AM '18



Problem: this can't work all the time!



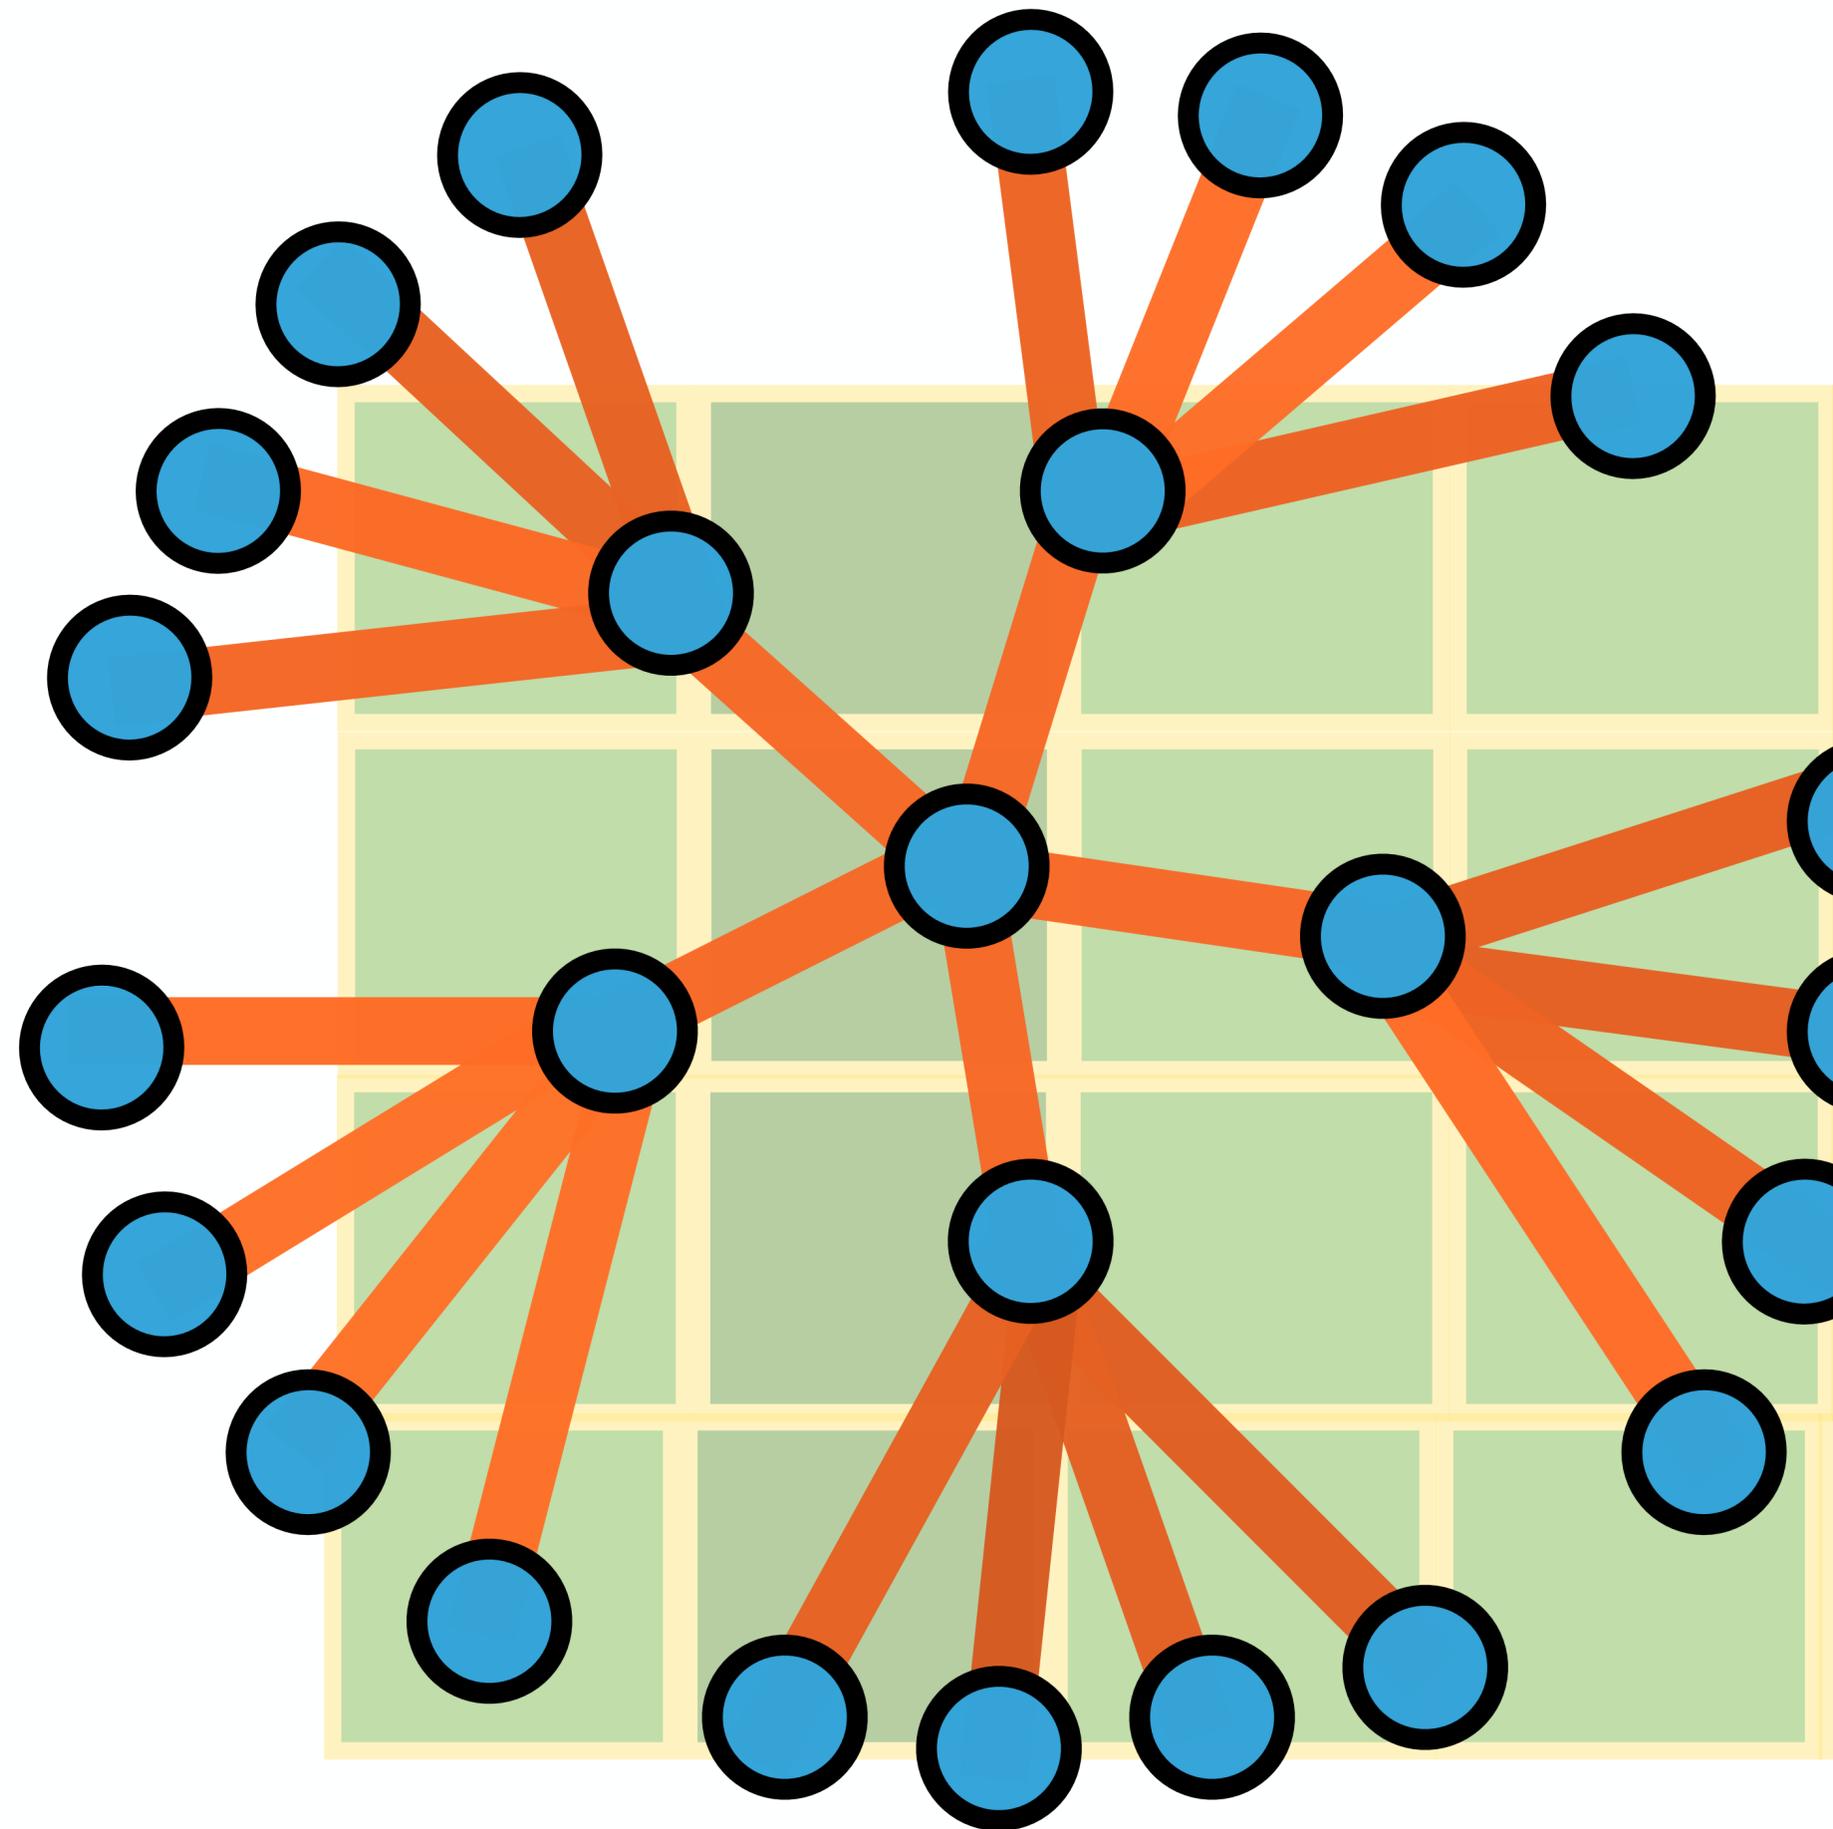


Problem: this can't work all the time!

Volume of balls in n-dimensional space grows like r^n

Balls in a the tree here grow exponentially: like 5^r

Exponentials grow faster than polynomials: big enough trees can't fit in *any* Euclidean space.





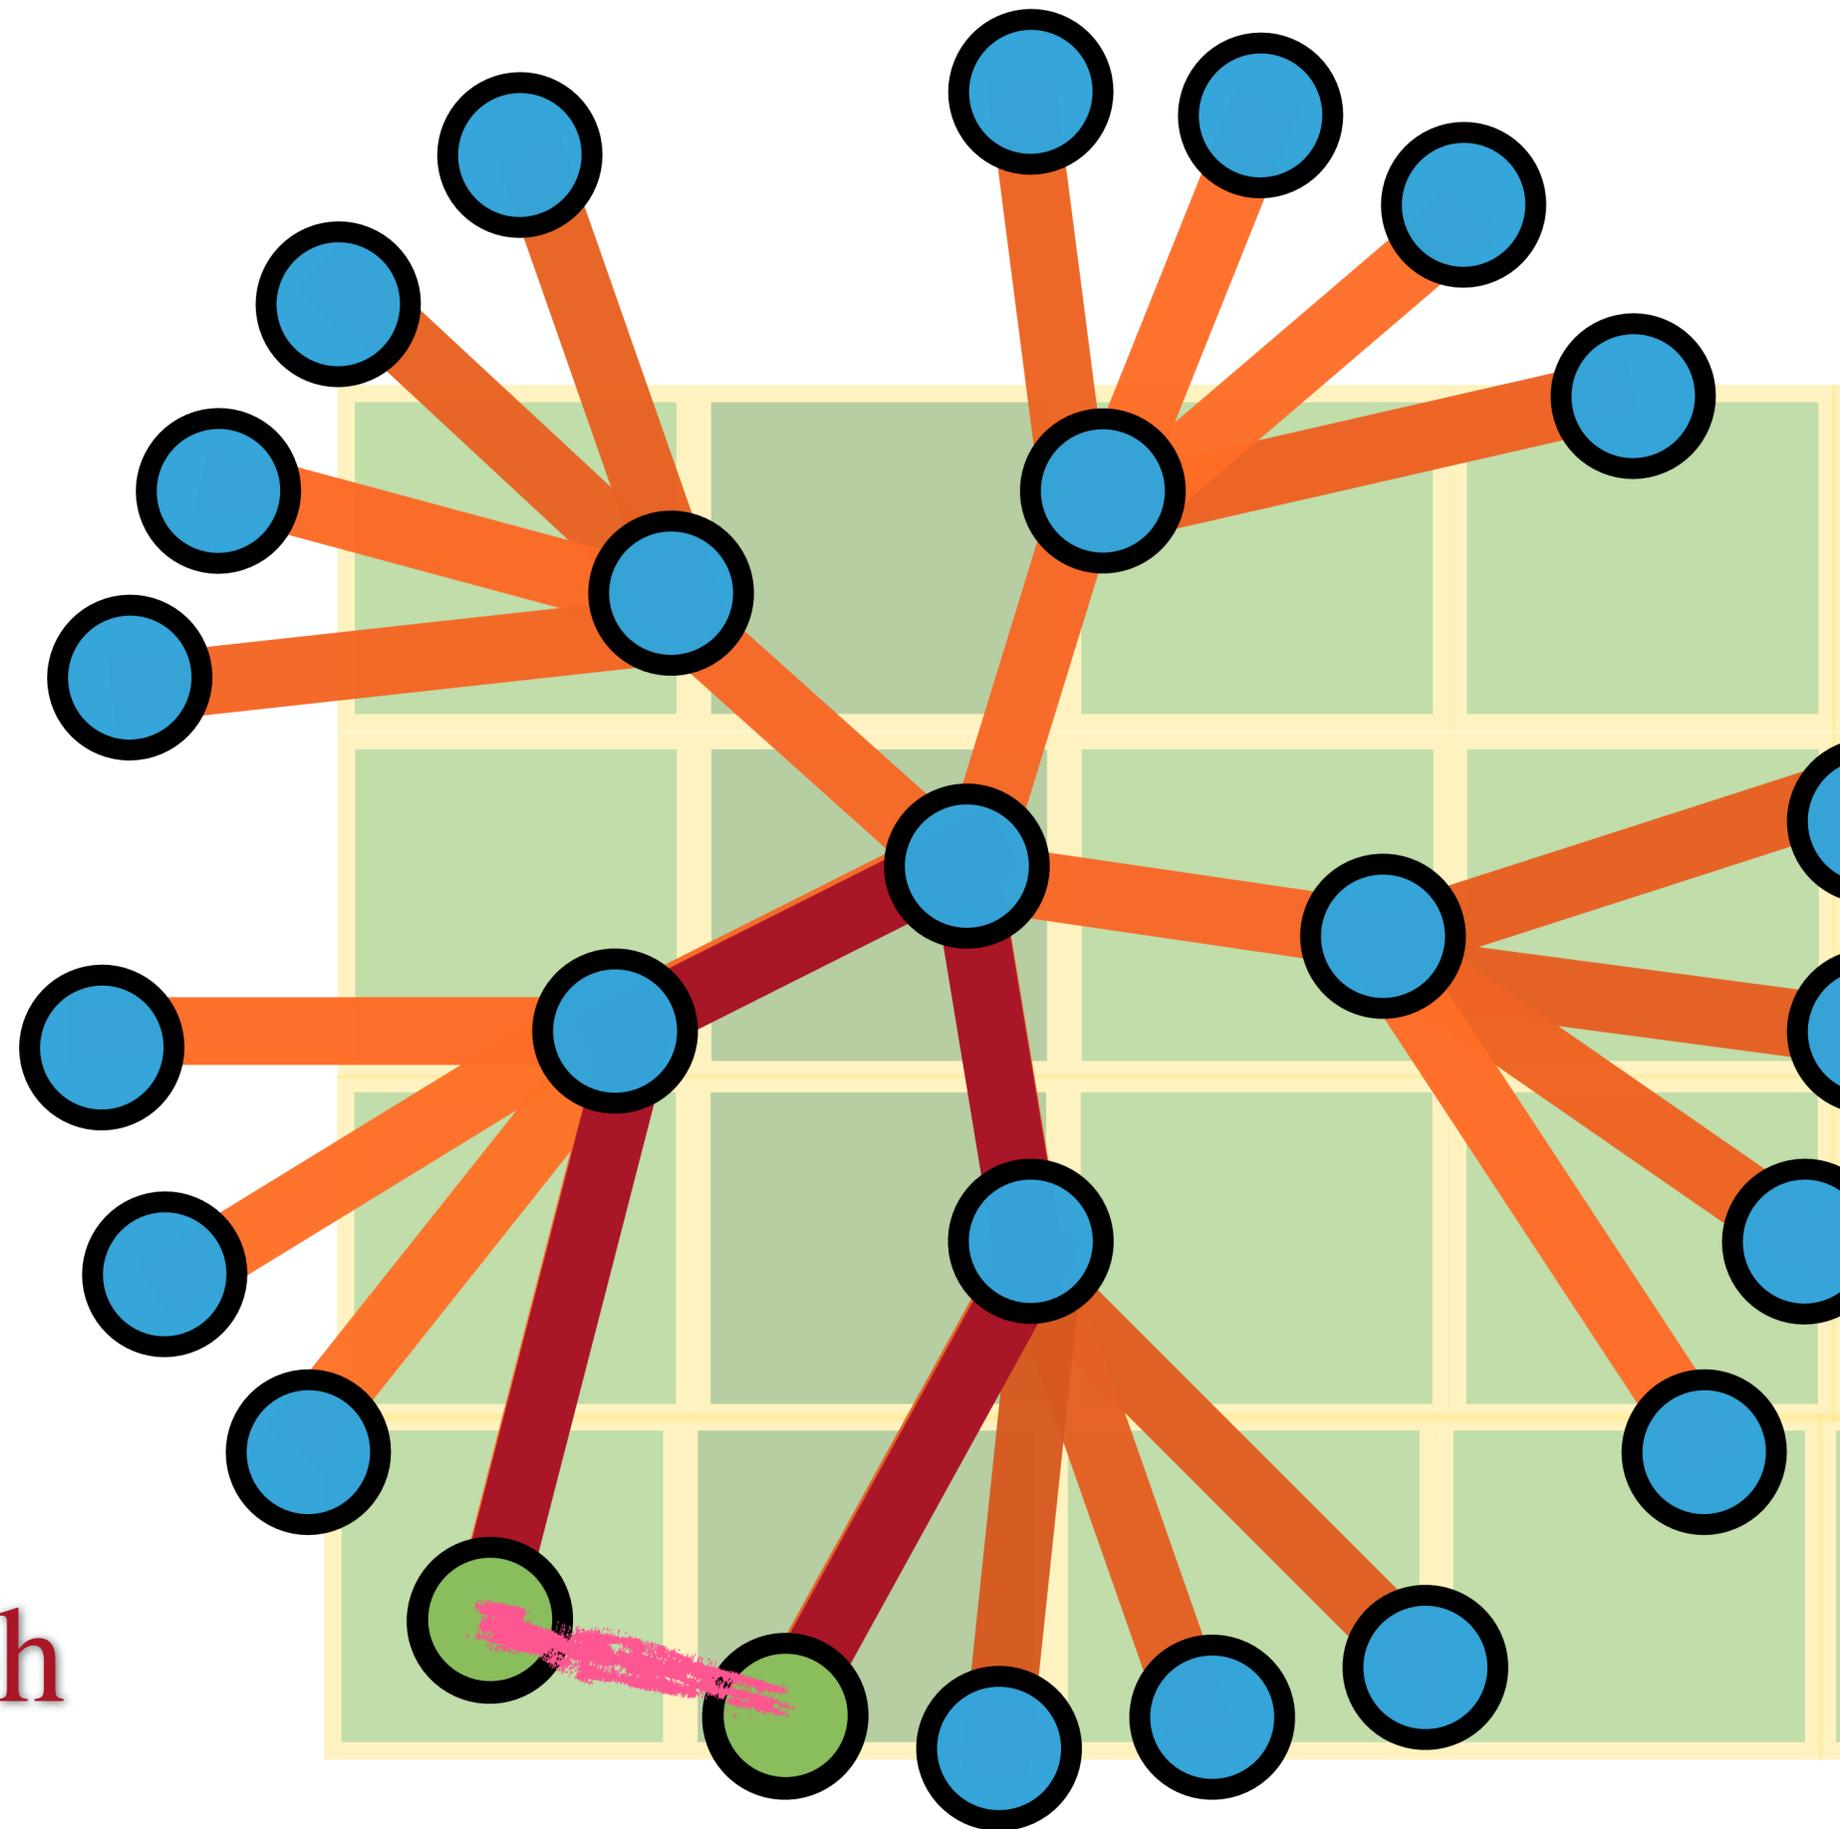
Problem: this can't work all the time!

Volume of balls in n-dimensional space grows like r^n

Balls in a the tree here grow exponentially: like 5^r

Exponentials grow faster than polynomials: big enough trees can't fit in *any* Euclidean space.

$$d_{\text{Euclidean}} < < d_{\text{Graph}}$$





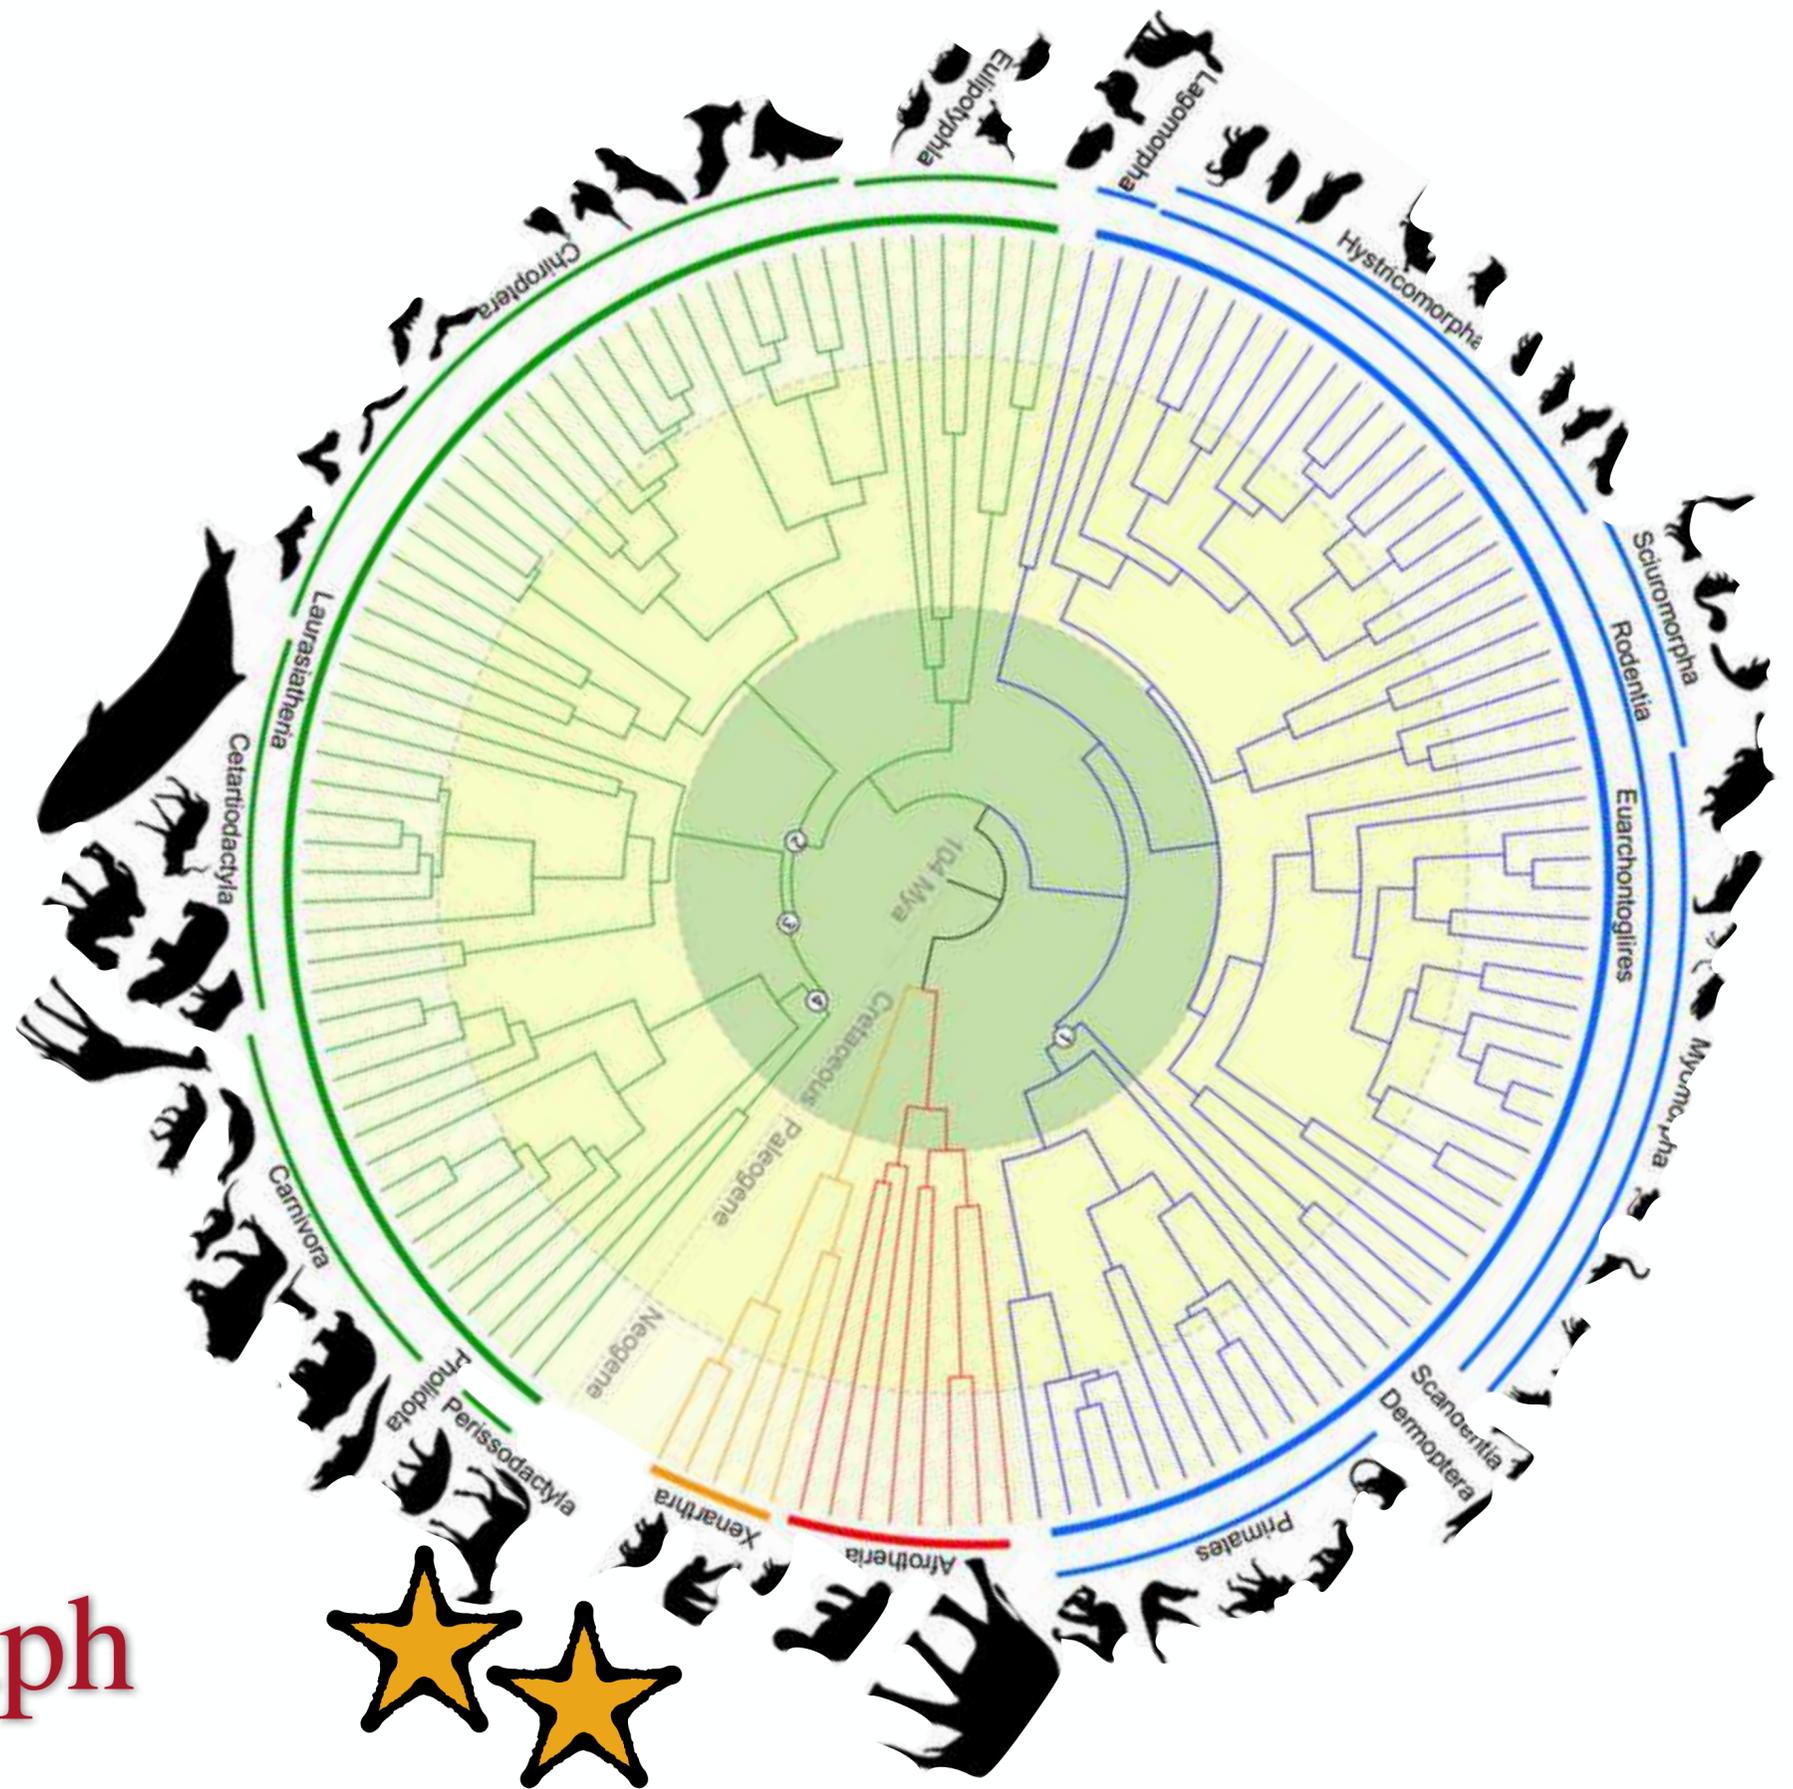
Problem: this can't work all the time!

Volume of balls in n-dimensional space grows like r^n

Balls in a the tree here grow exponentially: like 5^r

Exponentials grow faster than polynomials: big enough trees can't fit in *any* Euclidean space.

$d_{\text{Euclidean}} < < d_{\text{Graph}}$



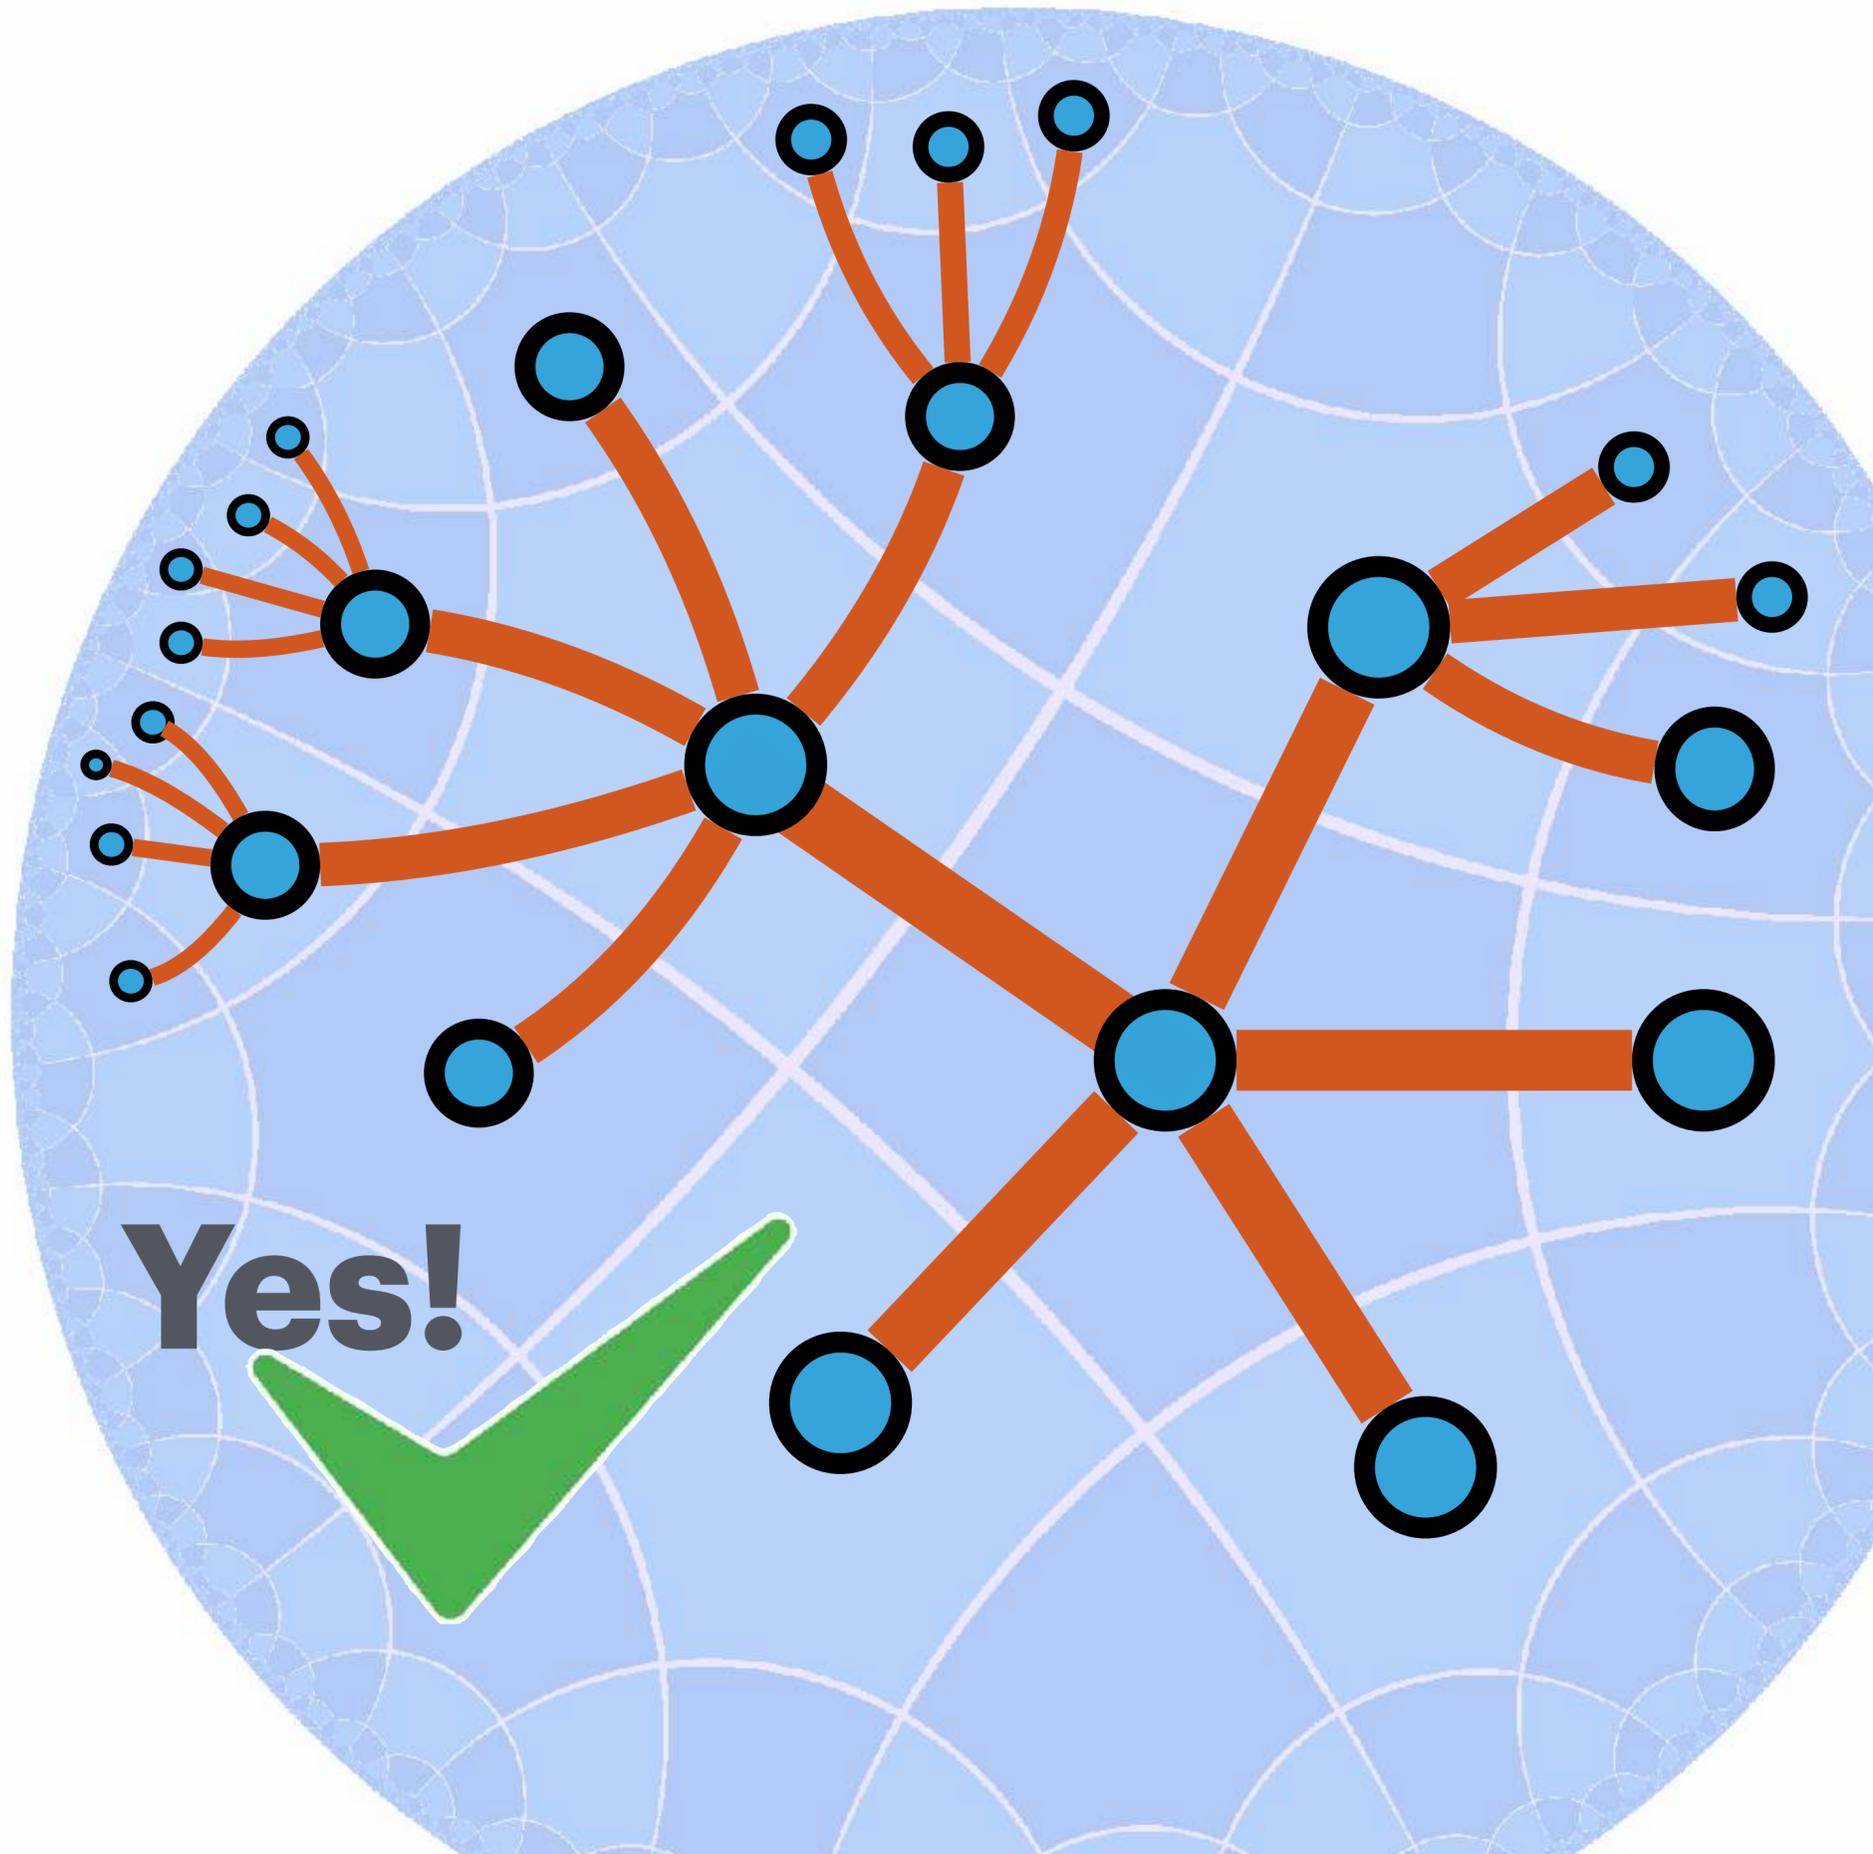
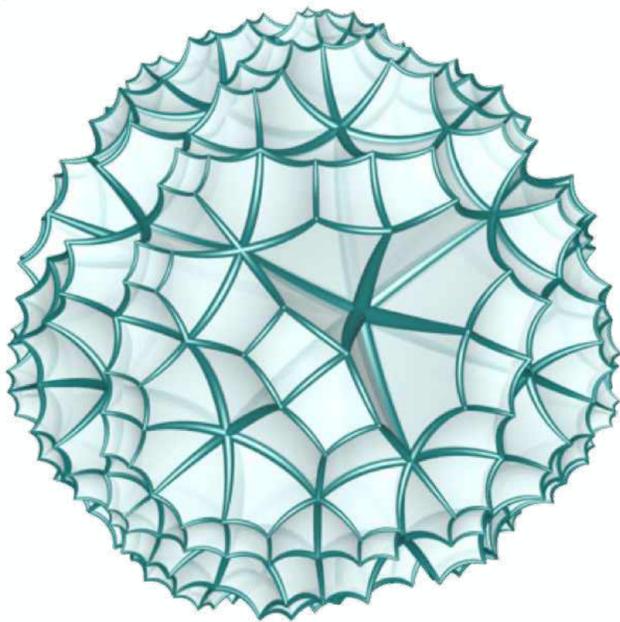


Idea:

The size of balls in negatively curved space grows exponentially!



Can we embed a tree in hyperbolic space?



Hyperbolic Geometry of Complex Networks
Dmitri Krioukov,¹ Fragkiskos Papadopoulos,² Maksim Kitsak,¹ Amin Vahdat,³ and Mari

Hyperbolic Graph Convolutional Neural Networks

Ines Chami[†], Adva Wolf[†], Frederic Sala[†], Christopher Ré[†], and Tommi Leskovec[†]

Neural Embeddings

Benjamin Paul Chamberlain

ABSTRACT
Neural embeddings for language processing that enable high performance on a wide range of applications. Graph-structured data that enable tasks in machine learning and graph analysis. We show that the hyperbolic space provides insights into the performance of graph neural networks.

KEYWORDS
neural embeddings, graph neural networks, hyperbolic geometry

1 INTRODUCTION
Embedding data into a low-dimensional space by first normalizing the data.

Hyperbolic

Learning low-dimensional representations of network data is a challenging problem. Graph neural networks (GNNs) are a powerful tool for analyzing graphs, but they often require high-dimensional representations. In this paper, we propose a novel GNN architecture that leverages hyperbolic geometry to learn low-dimensional representations of graph data. We show that this approach yields improved performance on a variety of graph tasks, including link prediction and node classification.

Low-Dimensional Knowledge Graph Embeddings via Hyperbolic Rotations

Ines Chami, Adva Wolf, Frederic Sala & Christopher Ré
Department of Computer Science
Stanford University

{chami, advaw, fredsalala}@stanford.edu, chrismre@cs.stanford.edu

Abstract

Knowledge graphs (KGs) capture rich relationships between a large number of entities. Embeddings of these structures must preserve these relationships with high fidelity. Recently, hyperbolic embedding methods have achieved state-of-the-art quality in graph representation learning tasks; when embedding certain graphs, they can produce parsimonious embeddings that have higher fidelity while using much fewer dimensions than their Euclidean counterparts. Mirroring work in Euclidean space, we are the first to leverage trainable hyperbolic rotations, a key notion in providing sufficiently rich representations for the complex logical patterns found in KGs. Coupled with trainable curvature, this approach yields improved embeddings in fewer dimensions. We evaluate our method, ROTATIONH, on the WN18RR link prediction task; in the low-dimensional setting, we improve on previous Euclidean-based efforts by 2.2% in mean reciprocal rank (MRR), and in the high-dimensional setting, we achieve a new state-of-the-art MRR of 49.2%, improving on existing methods by 1.1%.

1 Introduction

Knowledge Graphs and Back: Hyperbolic Embeddings and Clustering

Ines Chami^{††}, Christopher Ré[†]

Dimensionality Reduction via Horospherical Projections

Albert Gu^{*1}, Dat Nguyen^{*1}, Christopher Ré¹

Dimensionality reduction in high-dimensional spaces is a fundamental problem in machine learning. In this paper, we propose a novel dimensionality reduction technique called Horospherical Projections (HOROPCA). HOROPCA is a generalization of Principal Component Analysis (PCA) to hyperbolic geometry. Given a core notion of directions, PCA involves the following ingredients:

neighbor search (Krauthgamer & Lee, 2006; Wu & Charikar, 2020), hierarchical clustering (Monath et al., 2019; Chami et al., 2020a), or dimensionality reduction which is the focus of this work.

Euclidean Principal Component Analysis (PCA) is a fundamental dimensionality reduction technique which seeks directions that best explain the original data. PCA is an important primitive in data analysis and has many important uses such as (i) dimensionality reduction (e.g. for memory efficiency), (ii) data whitening and pre-processing for downstream tasks, and (iii) data visualization.

Here, we seek a generalization of PCA to hyperbolic geometry. Given a core notion of directions, PCA involves the following ingredients:

1. A nested sequence of affine subspaces (flags) spanned by a set of directions.
2. A projection method which maps points to these subspaces.



Problem: if you try to run this on real world datasets, it does a bad job

Where we got involved!



Me



Anna Wienhard



Michel Strube



Idea: it does bad when the intrinsic geometry of a graph does not look like hyperbolic space....





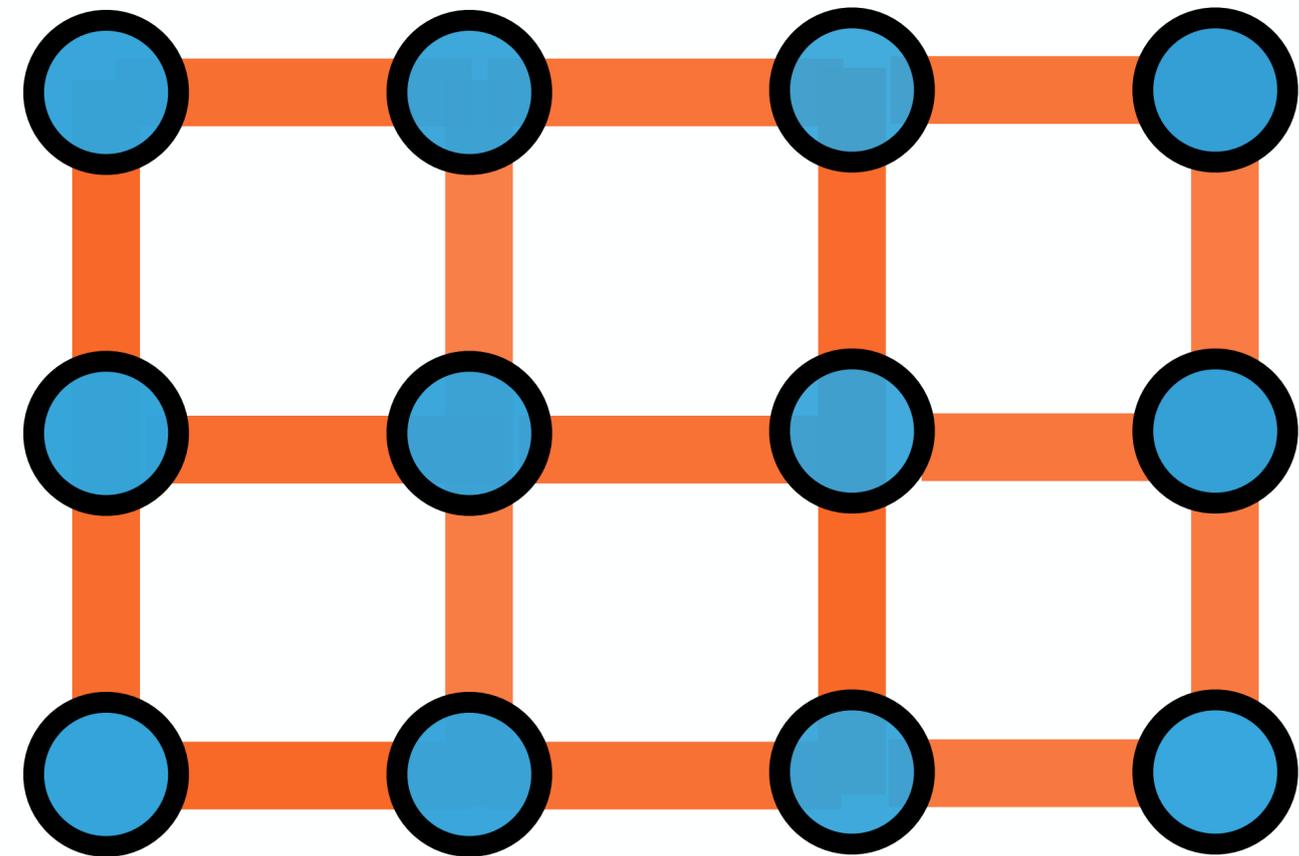
Problem: if you try to run this on real world datasets, it does a bad job

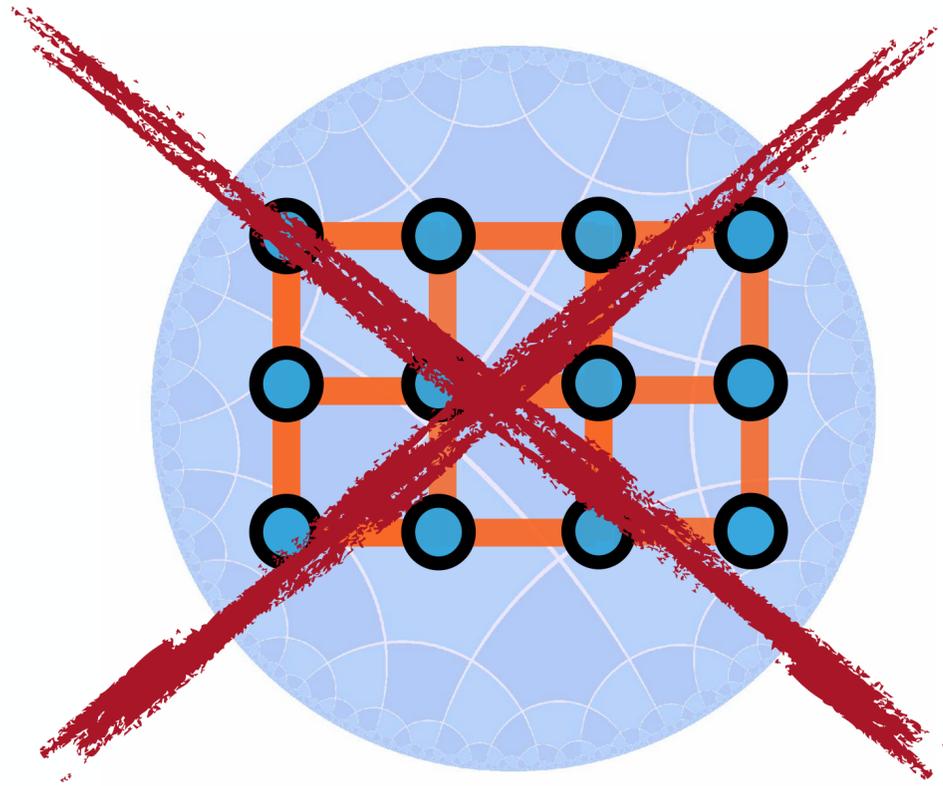
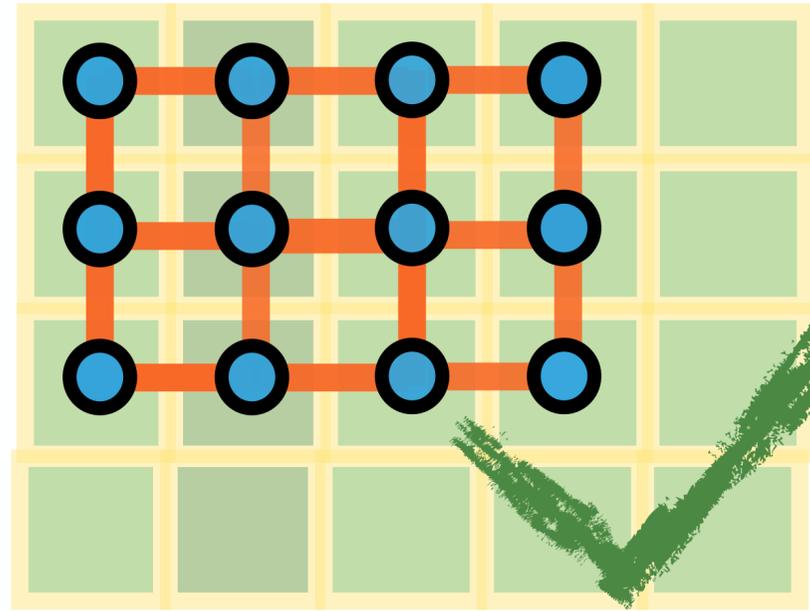
This grid-shaped graph has geometry like the Euclidean plane.

Euclidean space is very different than hyperbolic space.

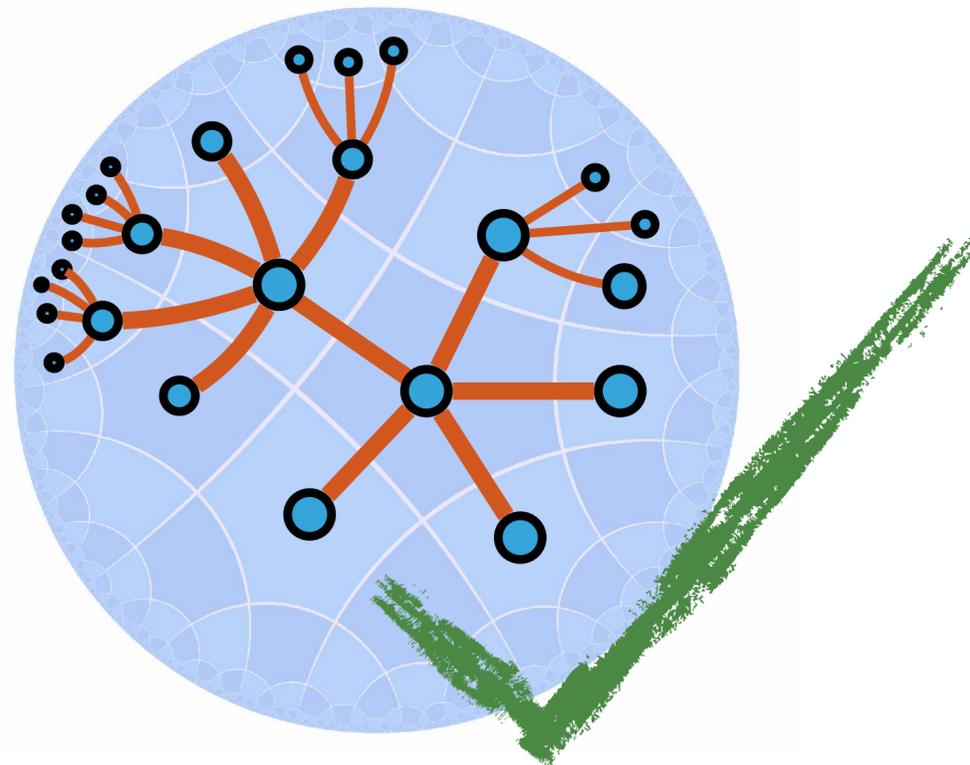
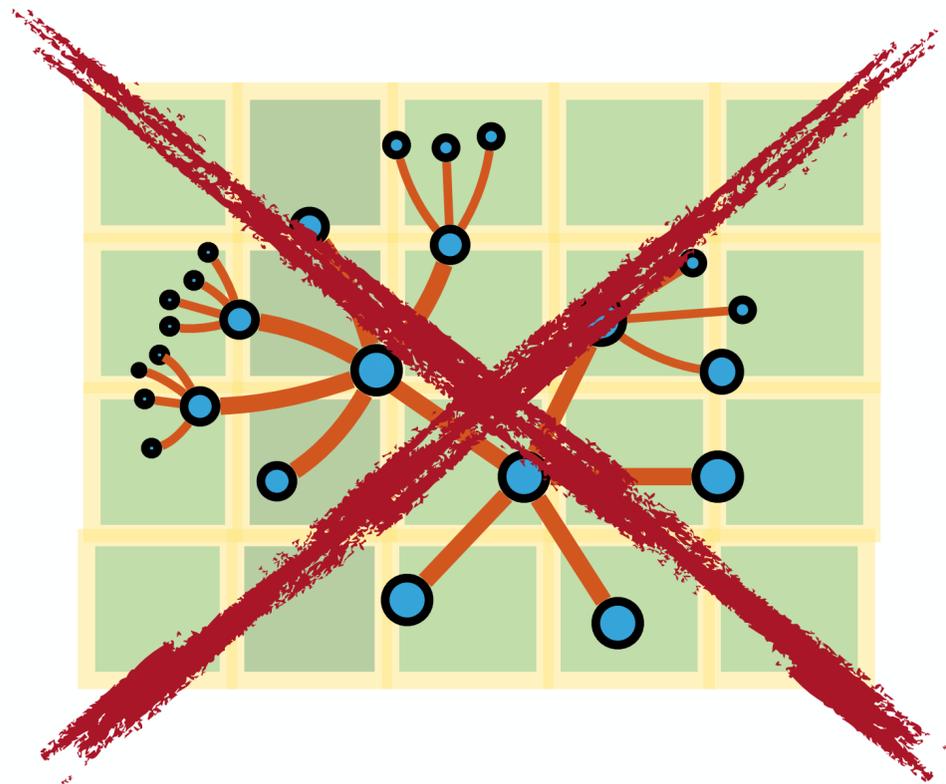
Balls in the graph grow like r^2 . Balls in hyperbolic space grow like e^r

So, this grid won't fit well inside of hyperbolic space.





This is not good. To figure out where to embed our graph we need to know the geometry of our graph.



But we want something for arbitrary, real world data!

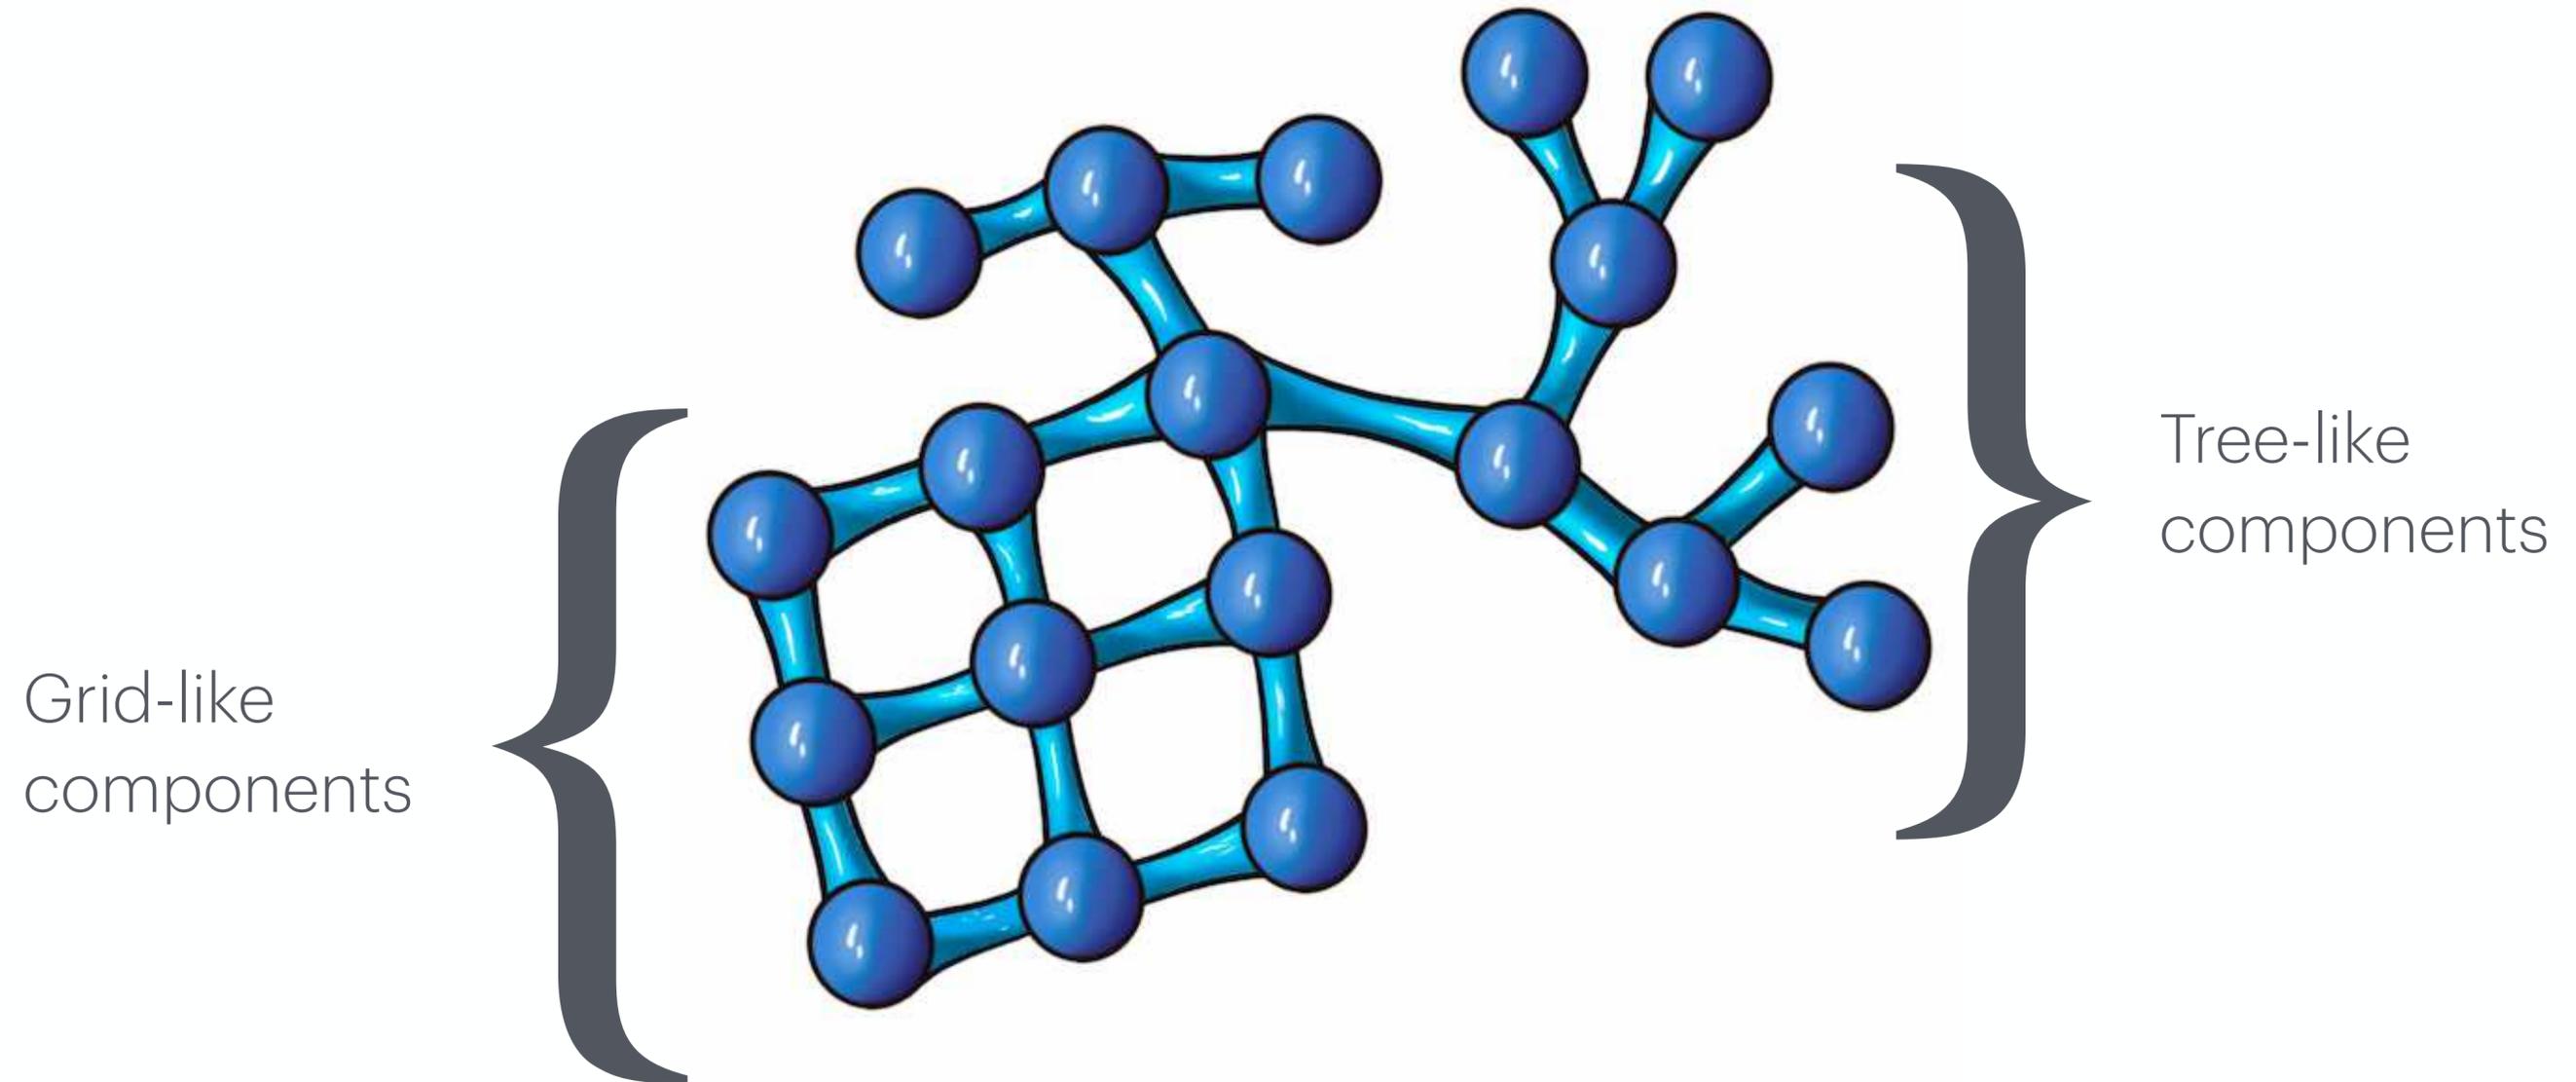


**But it gets
worse...**

Vertices: legal cases in
England and Wales.

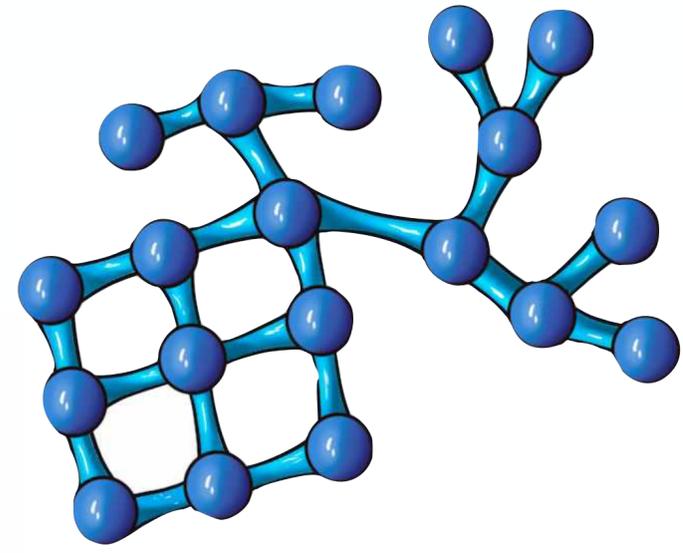
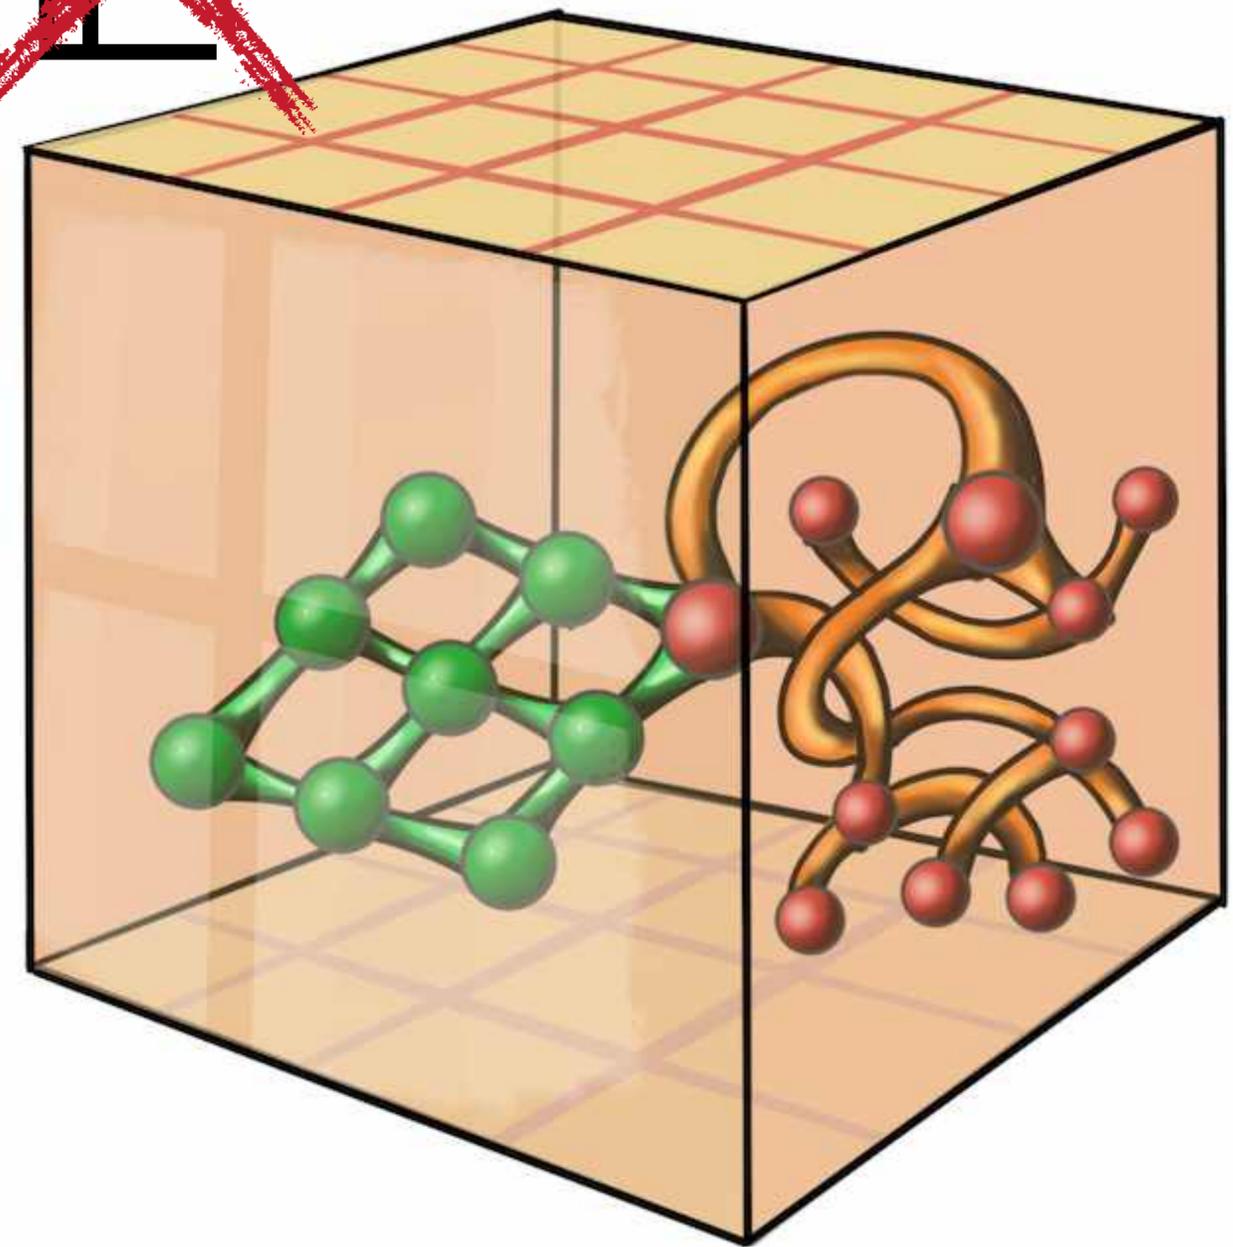
Edges: Citations

A small piece of the legal casework graph



This graph is trouble for **BOTH** spaces

~~\mathbb{R}^n~~



~~\mathbb{R}^n~~



Can we build a “**universal embedder**”: a space that let’s us embed real world graphs accurately, even if we don’t know their geometry?



Me



Anna Wienhard



Michel Strube

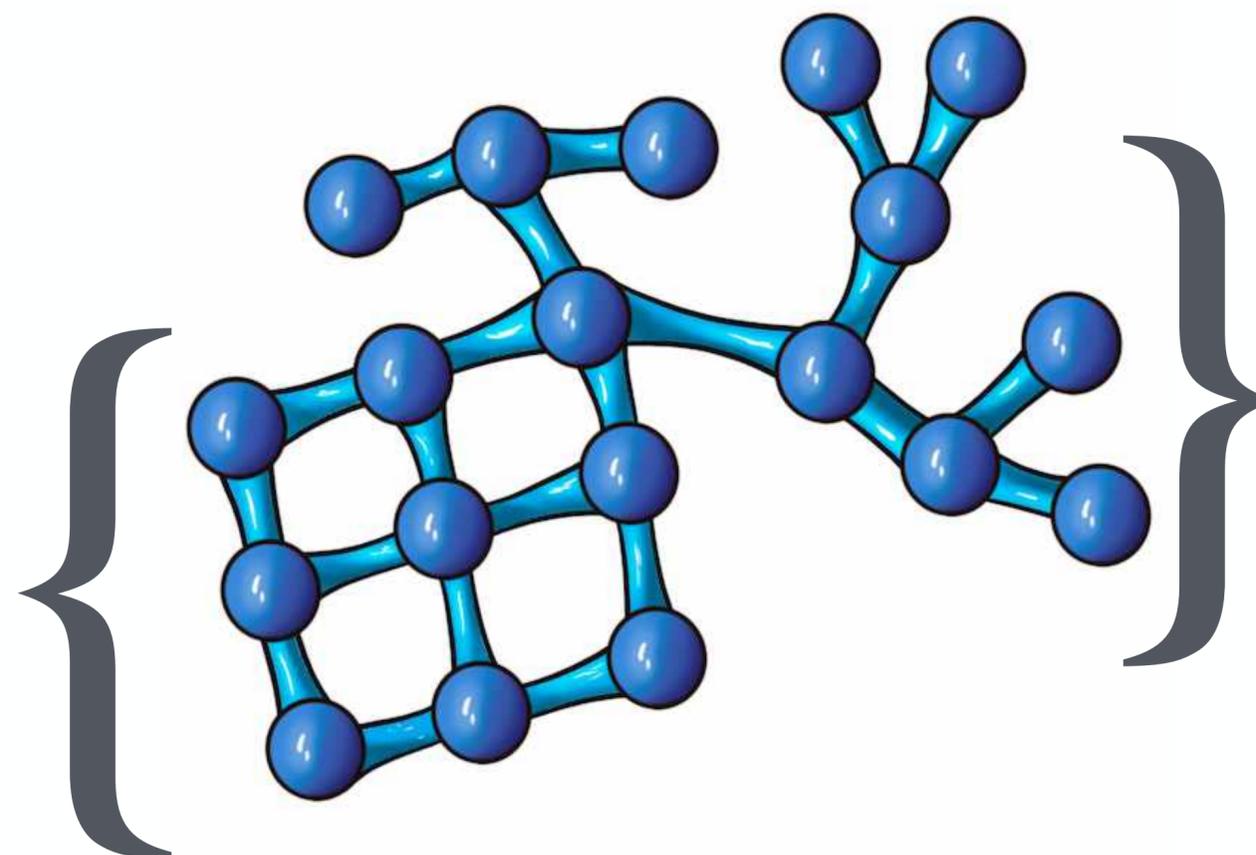


Fede Lopez



Bea Pozzetti

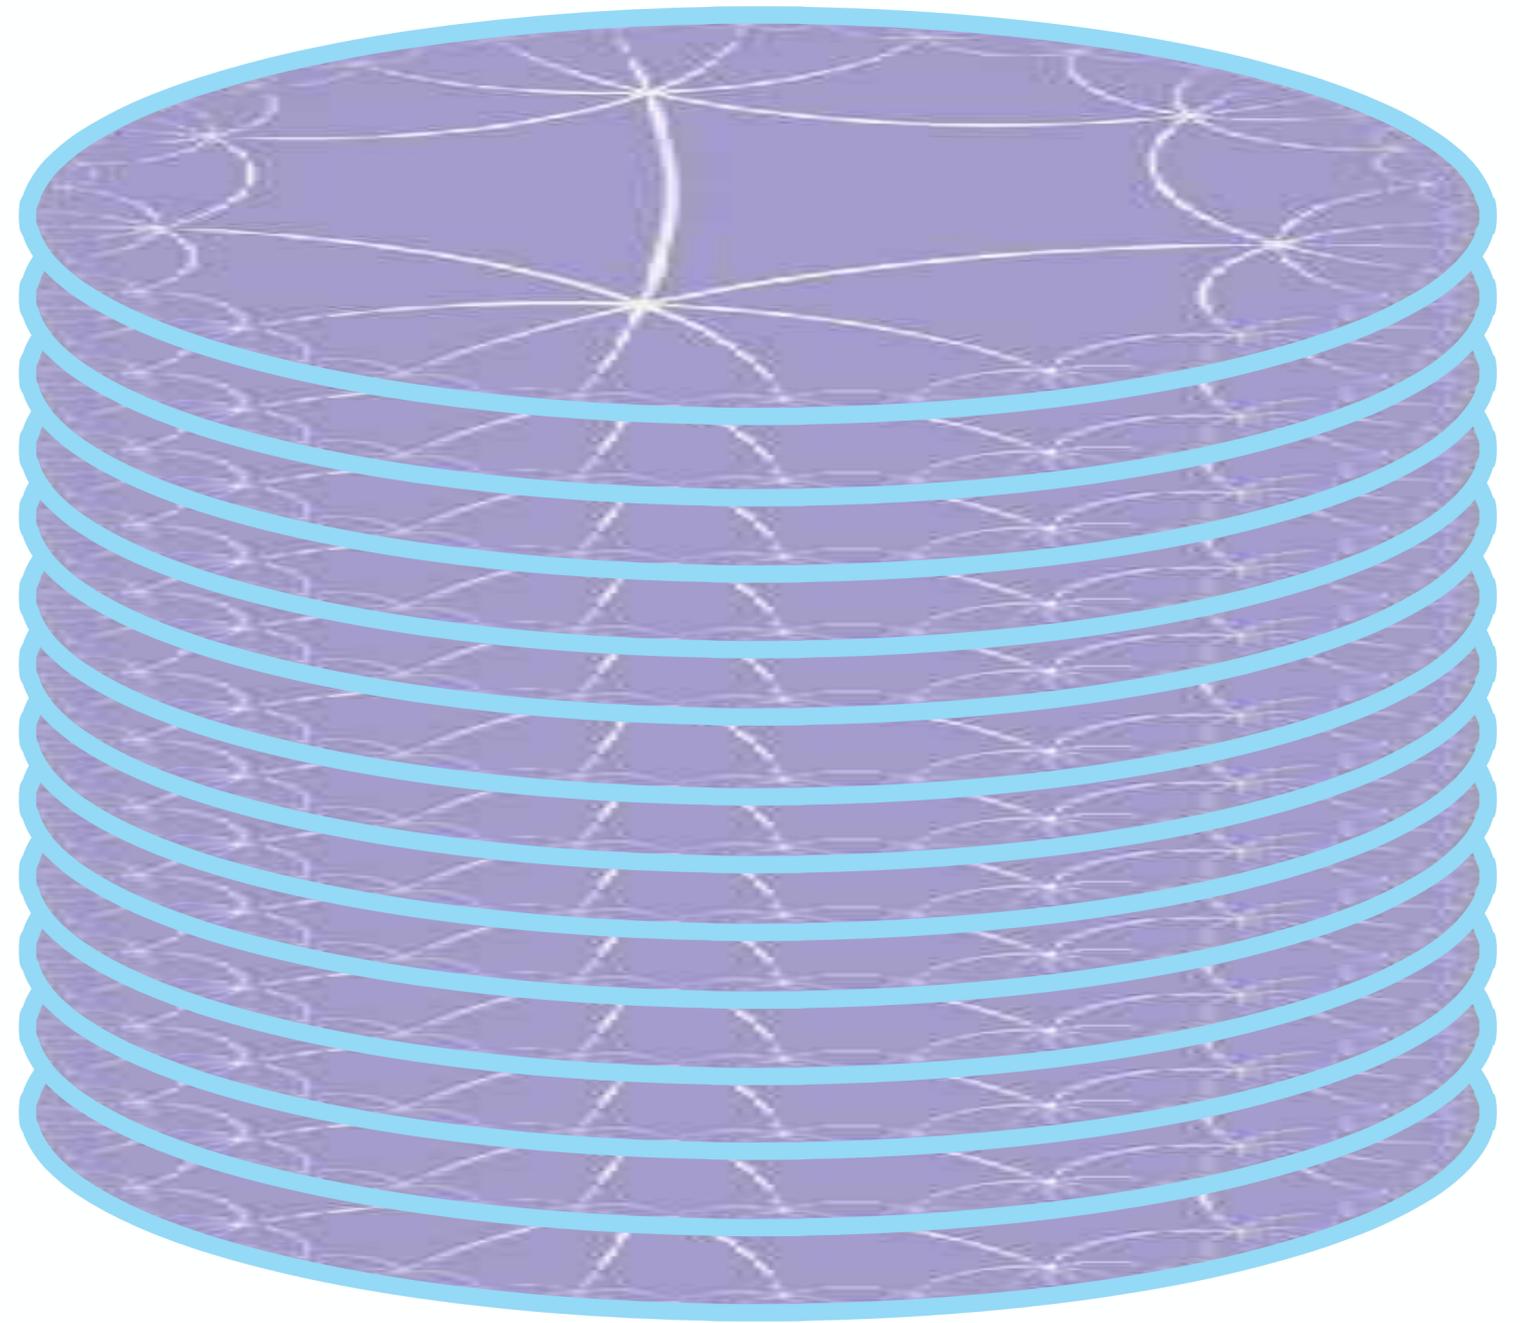
Need
flat
spots



Need
hyperbolic
spots

$$\mathbb{H}^2 \times \mathbb{R}$$

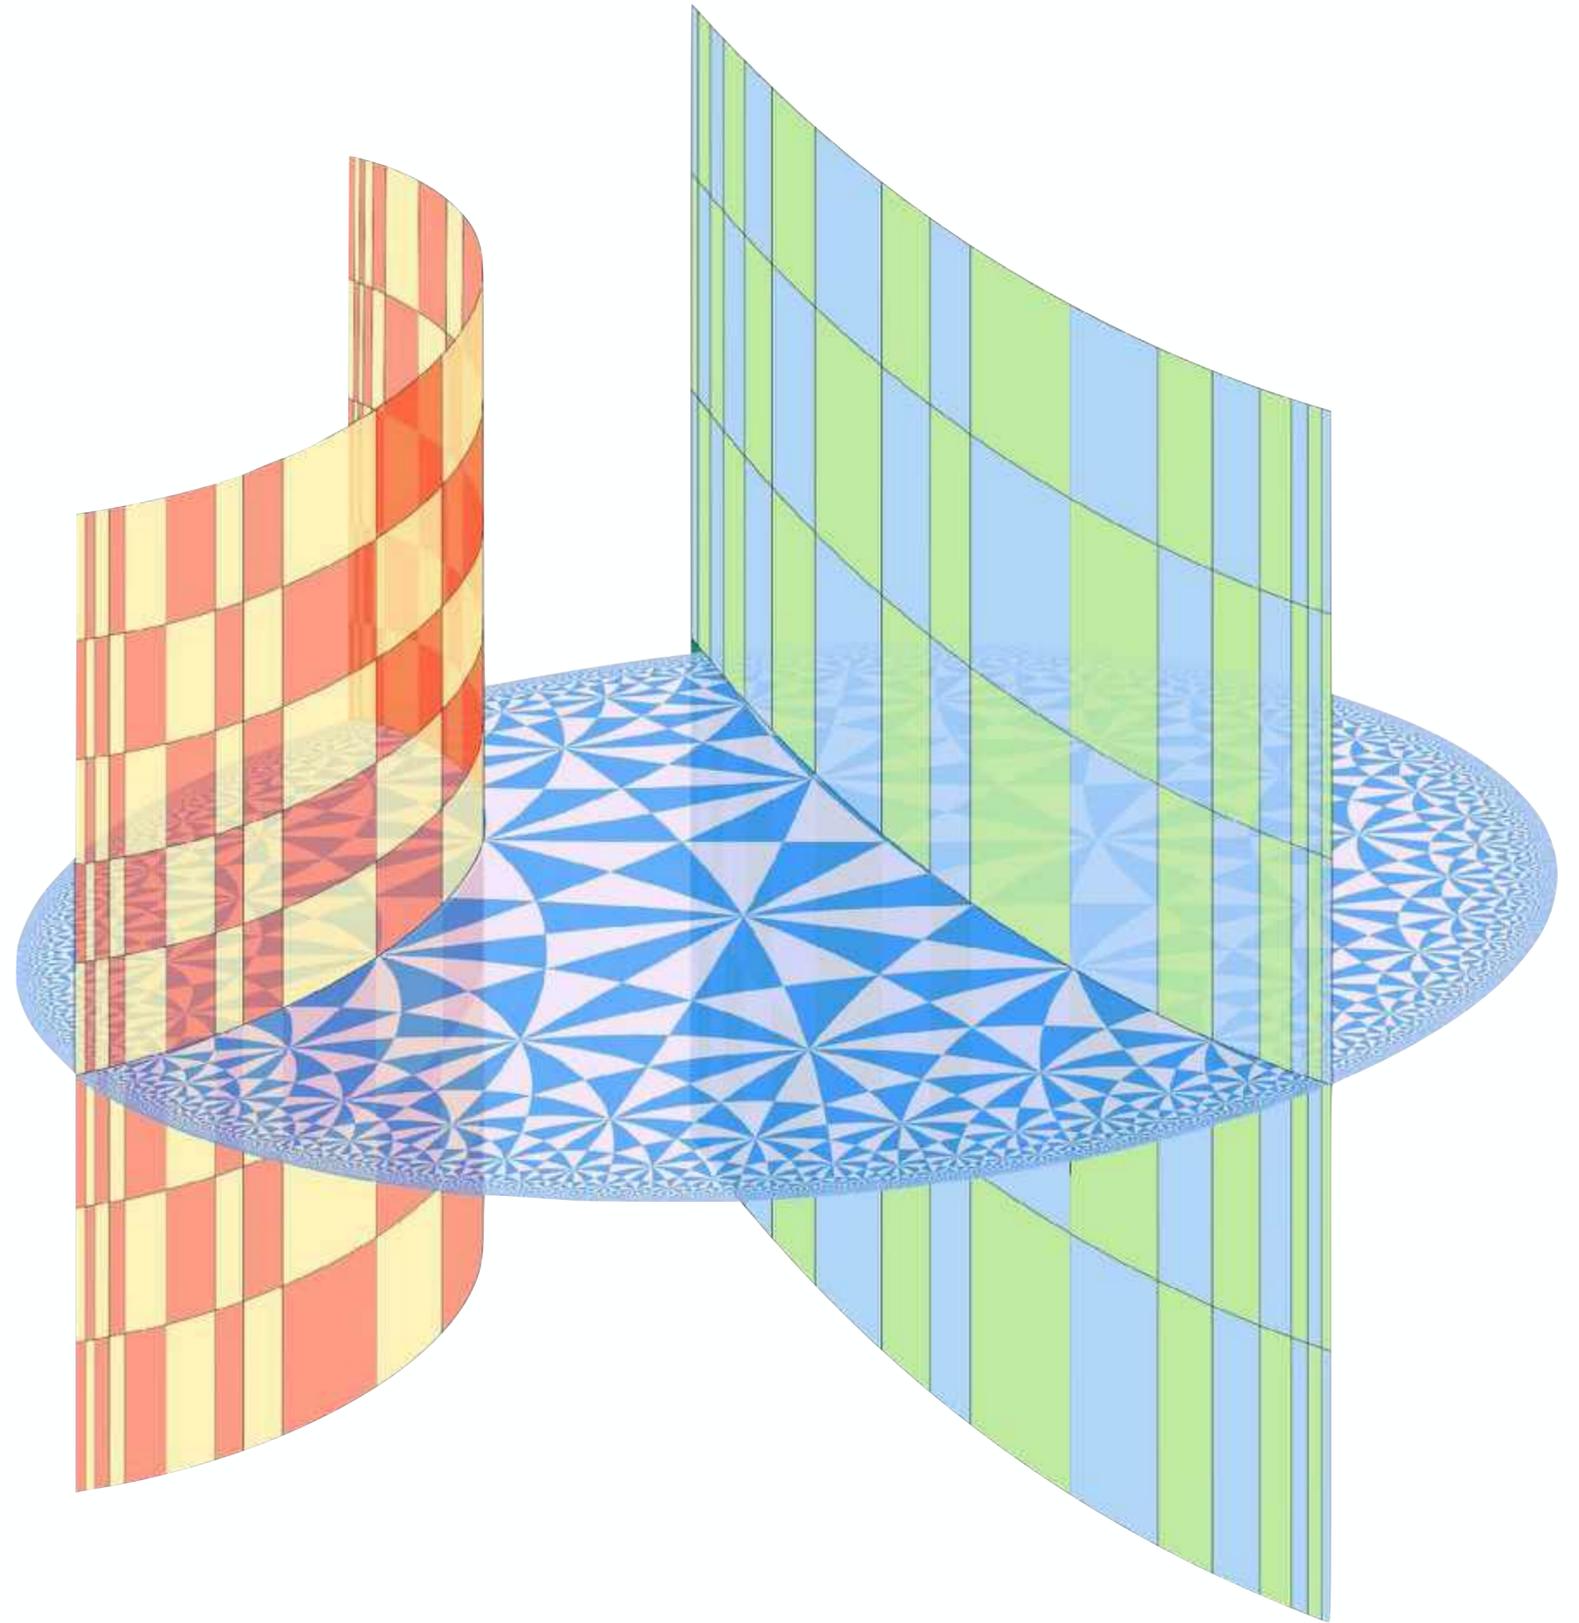
The geometry of
**a stack of
hyperbolic
planes.**

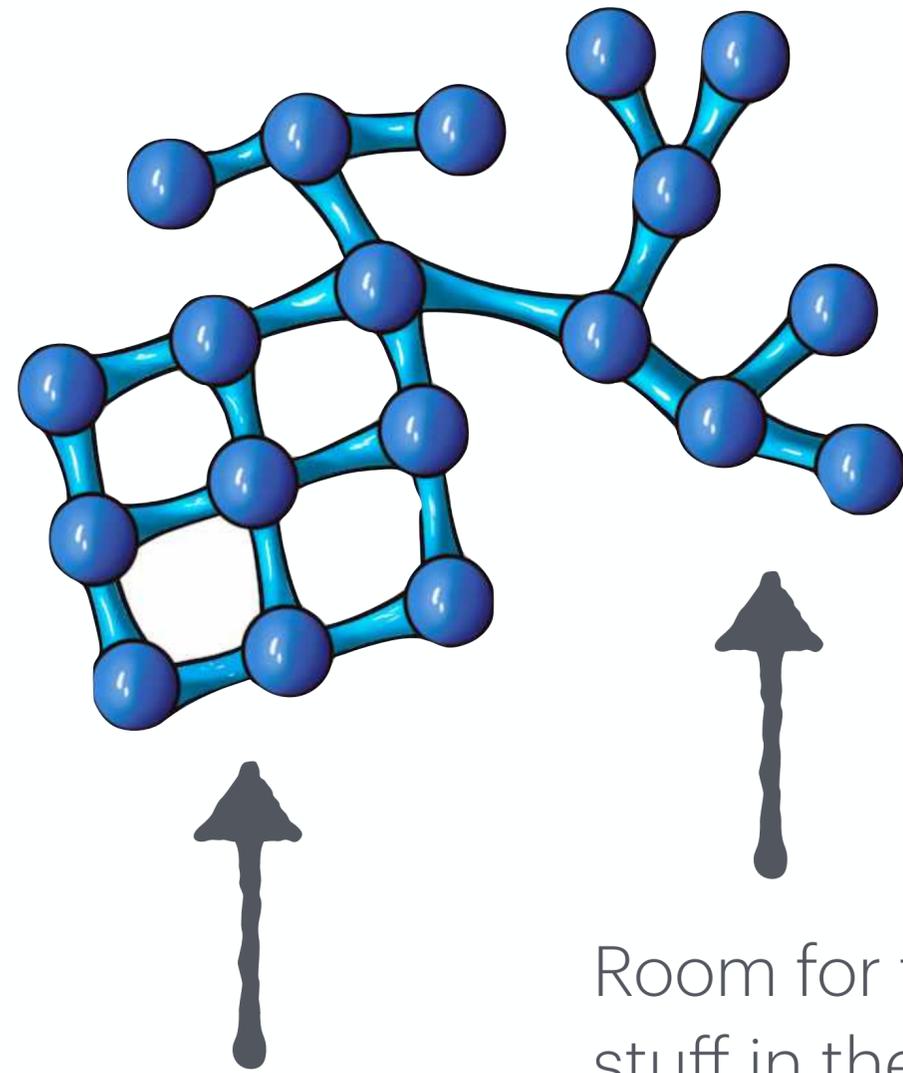


$$\mathbb{H}^2 \times \mathbb{R}$$

Horizontal slices =
Negative Curvature

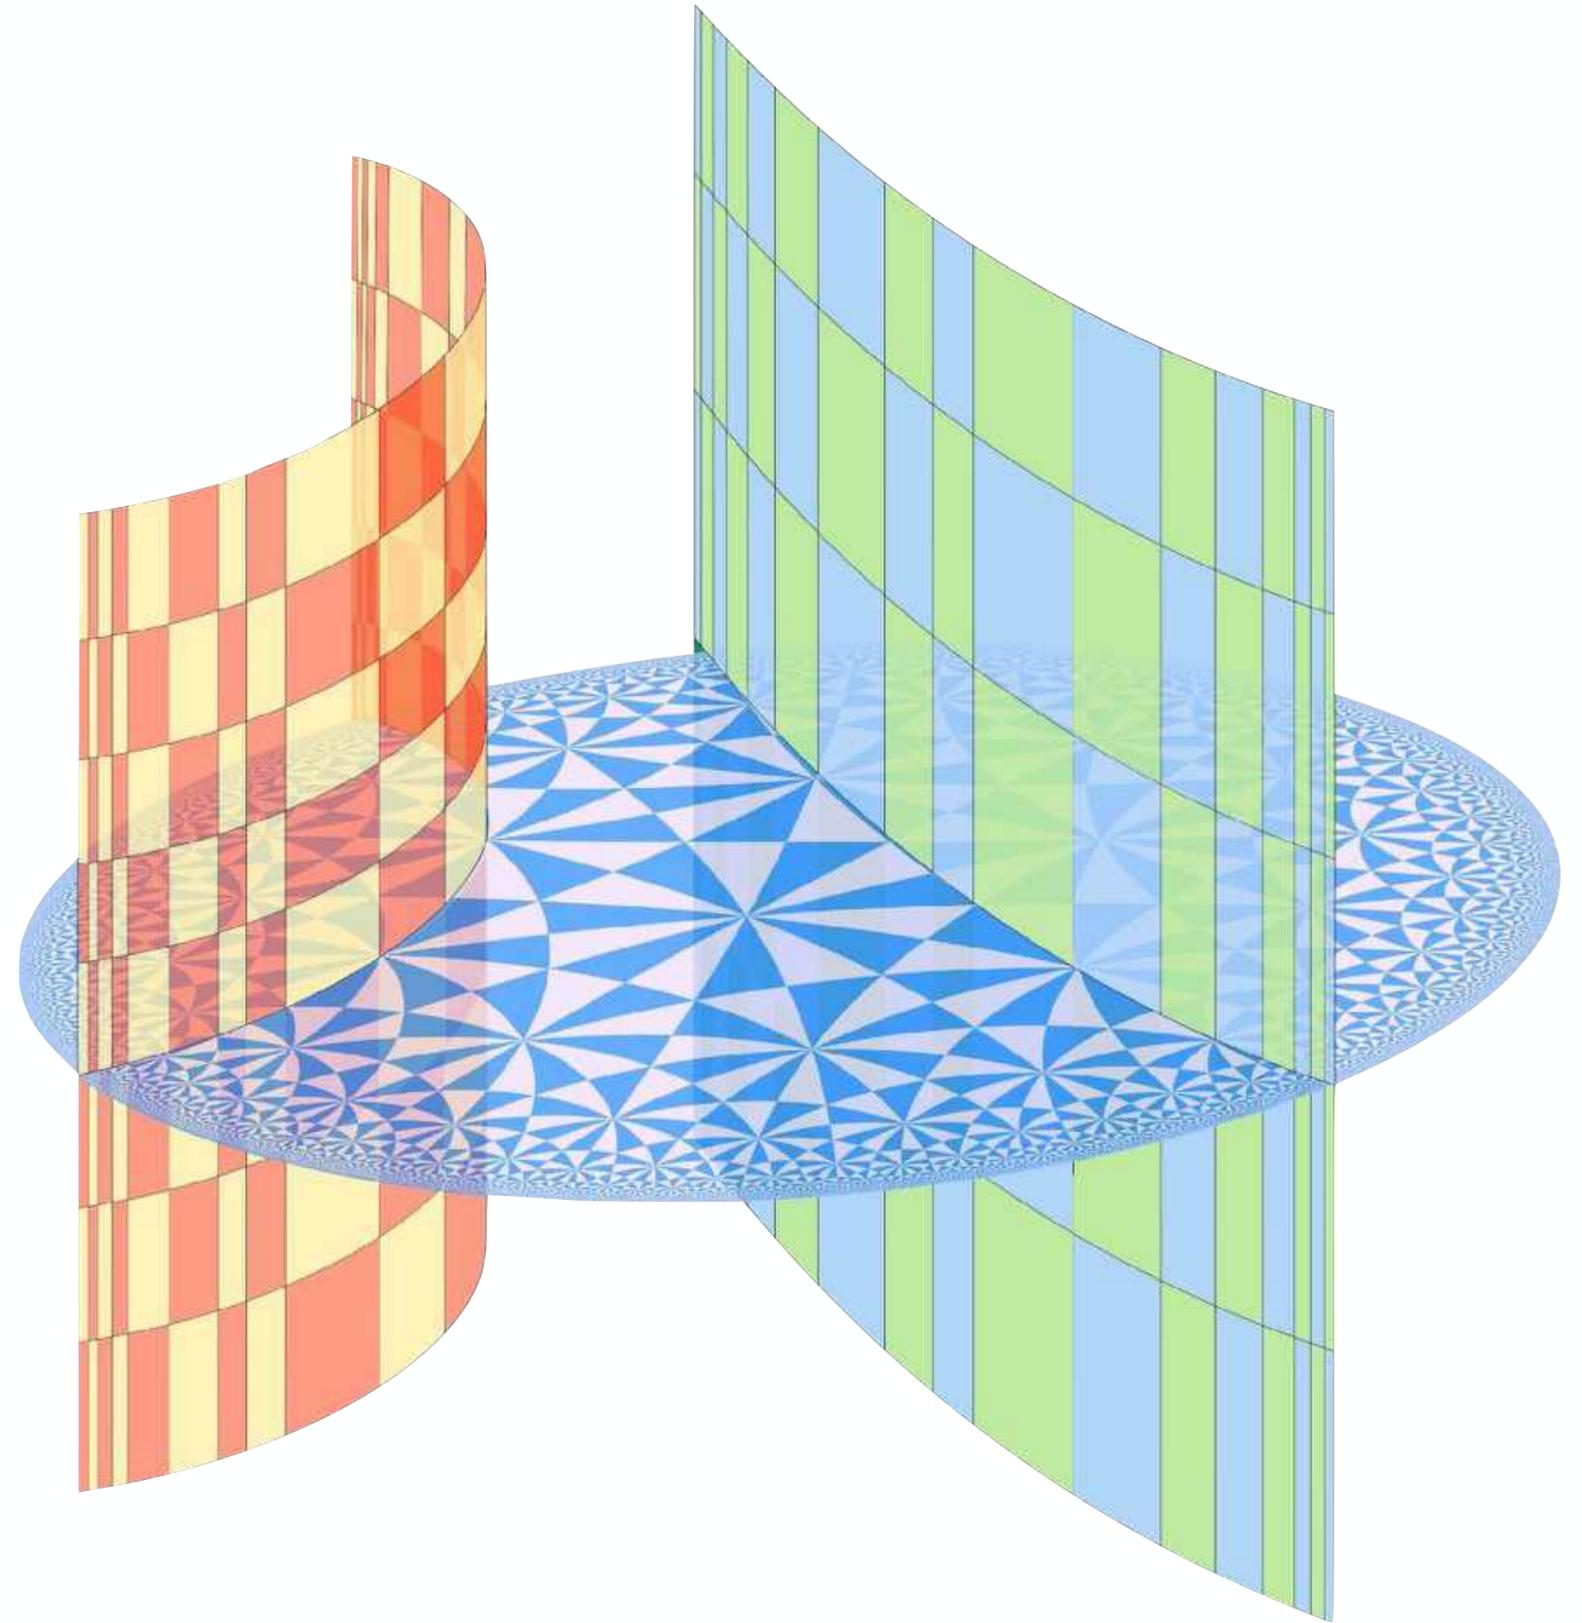
Vertical slices =
Zero Curvature

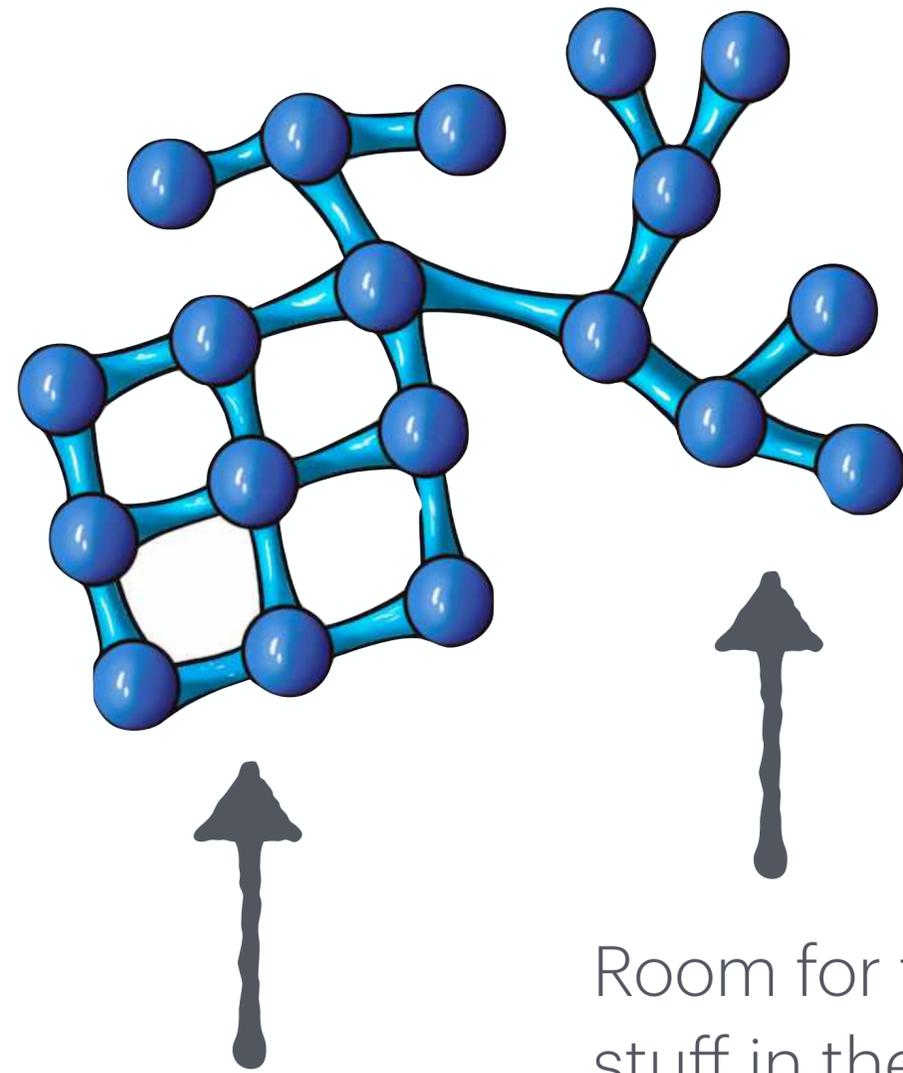




Room for grid-like stuff in the flat parts,

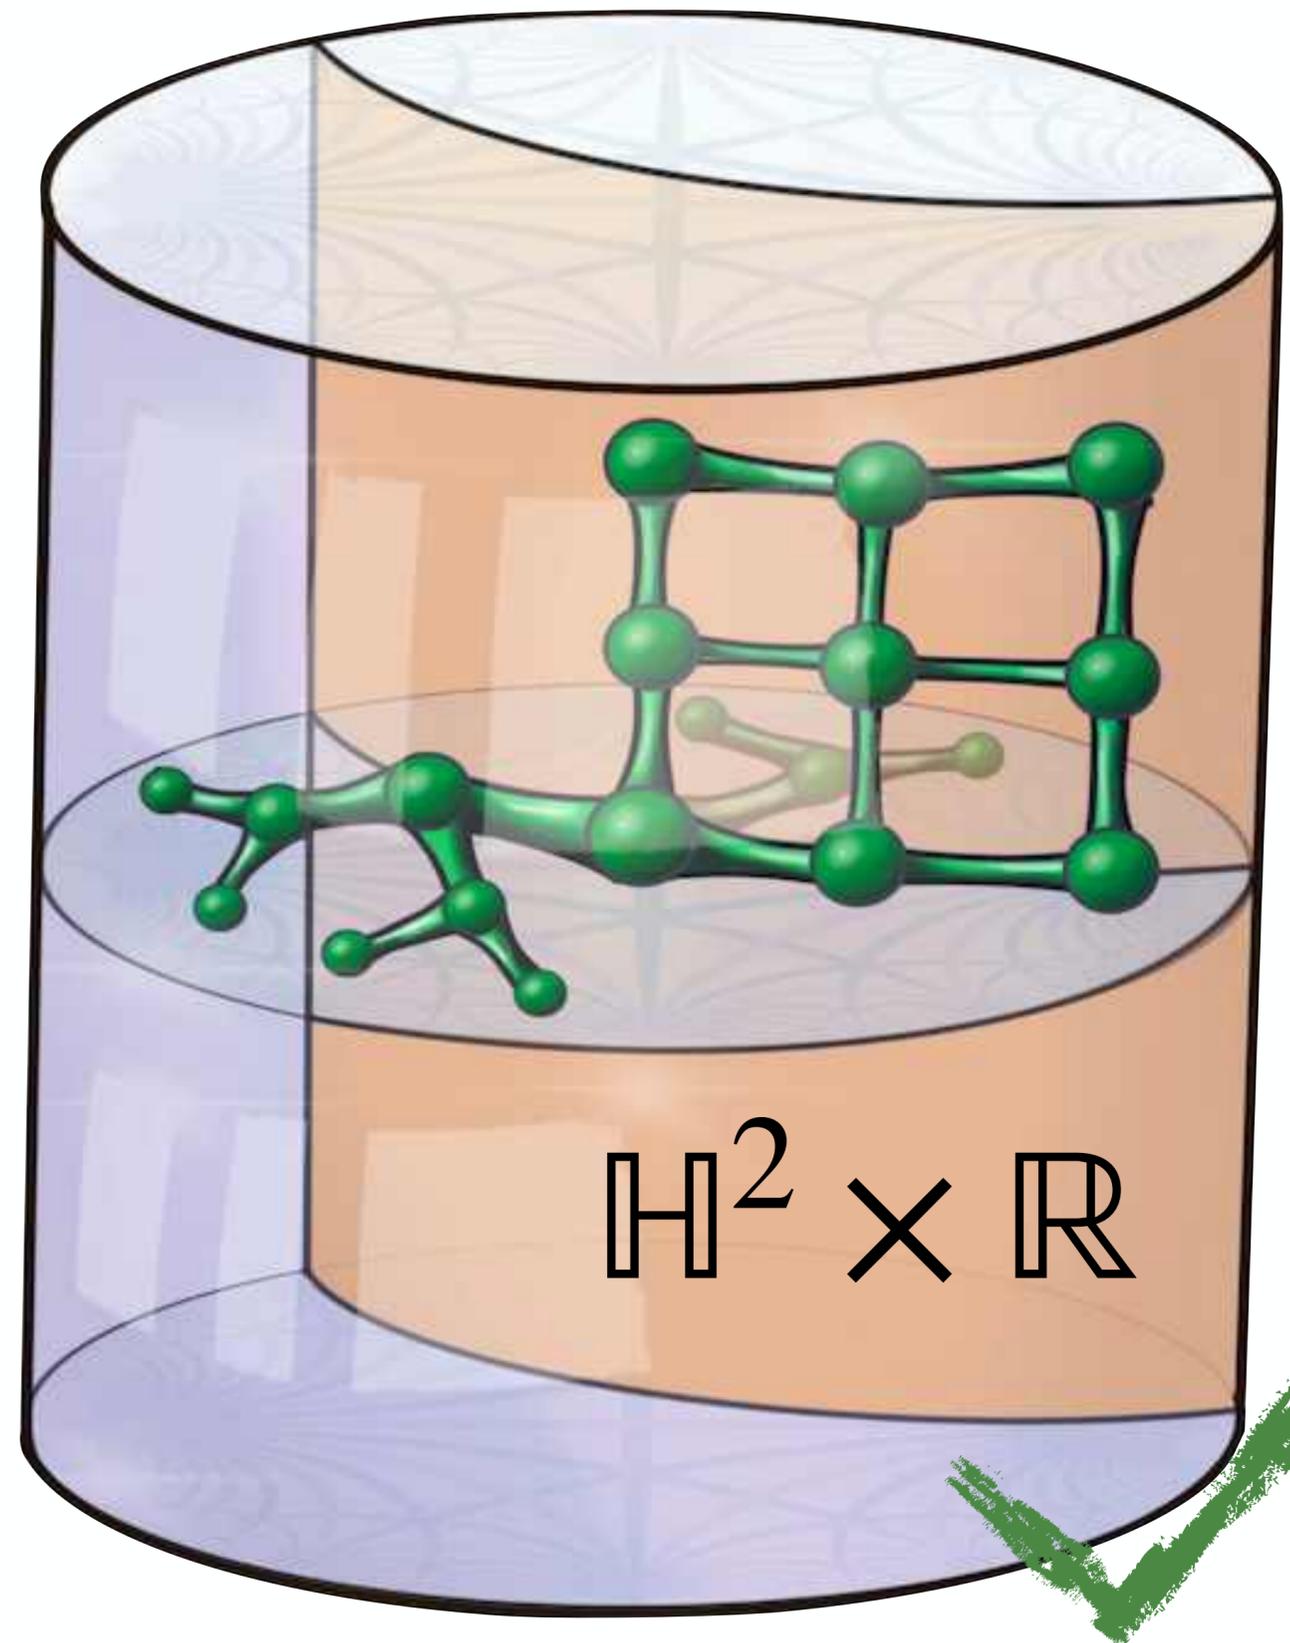
Room for tree-like stuff in the negative curvature



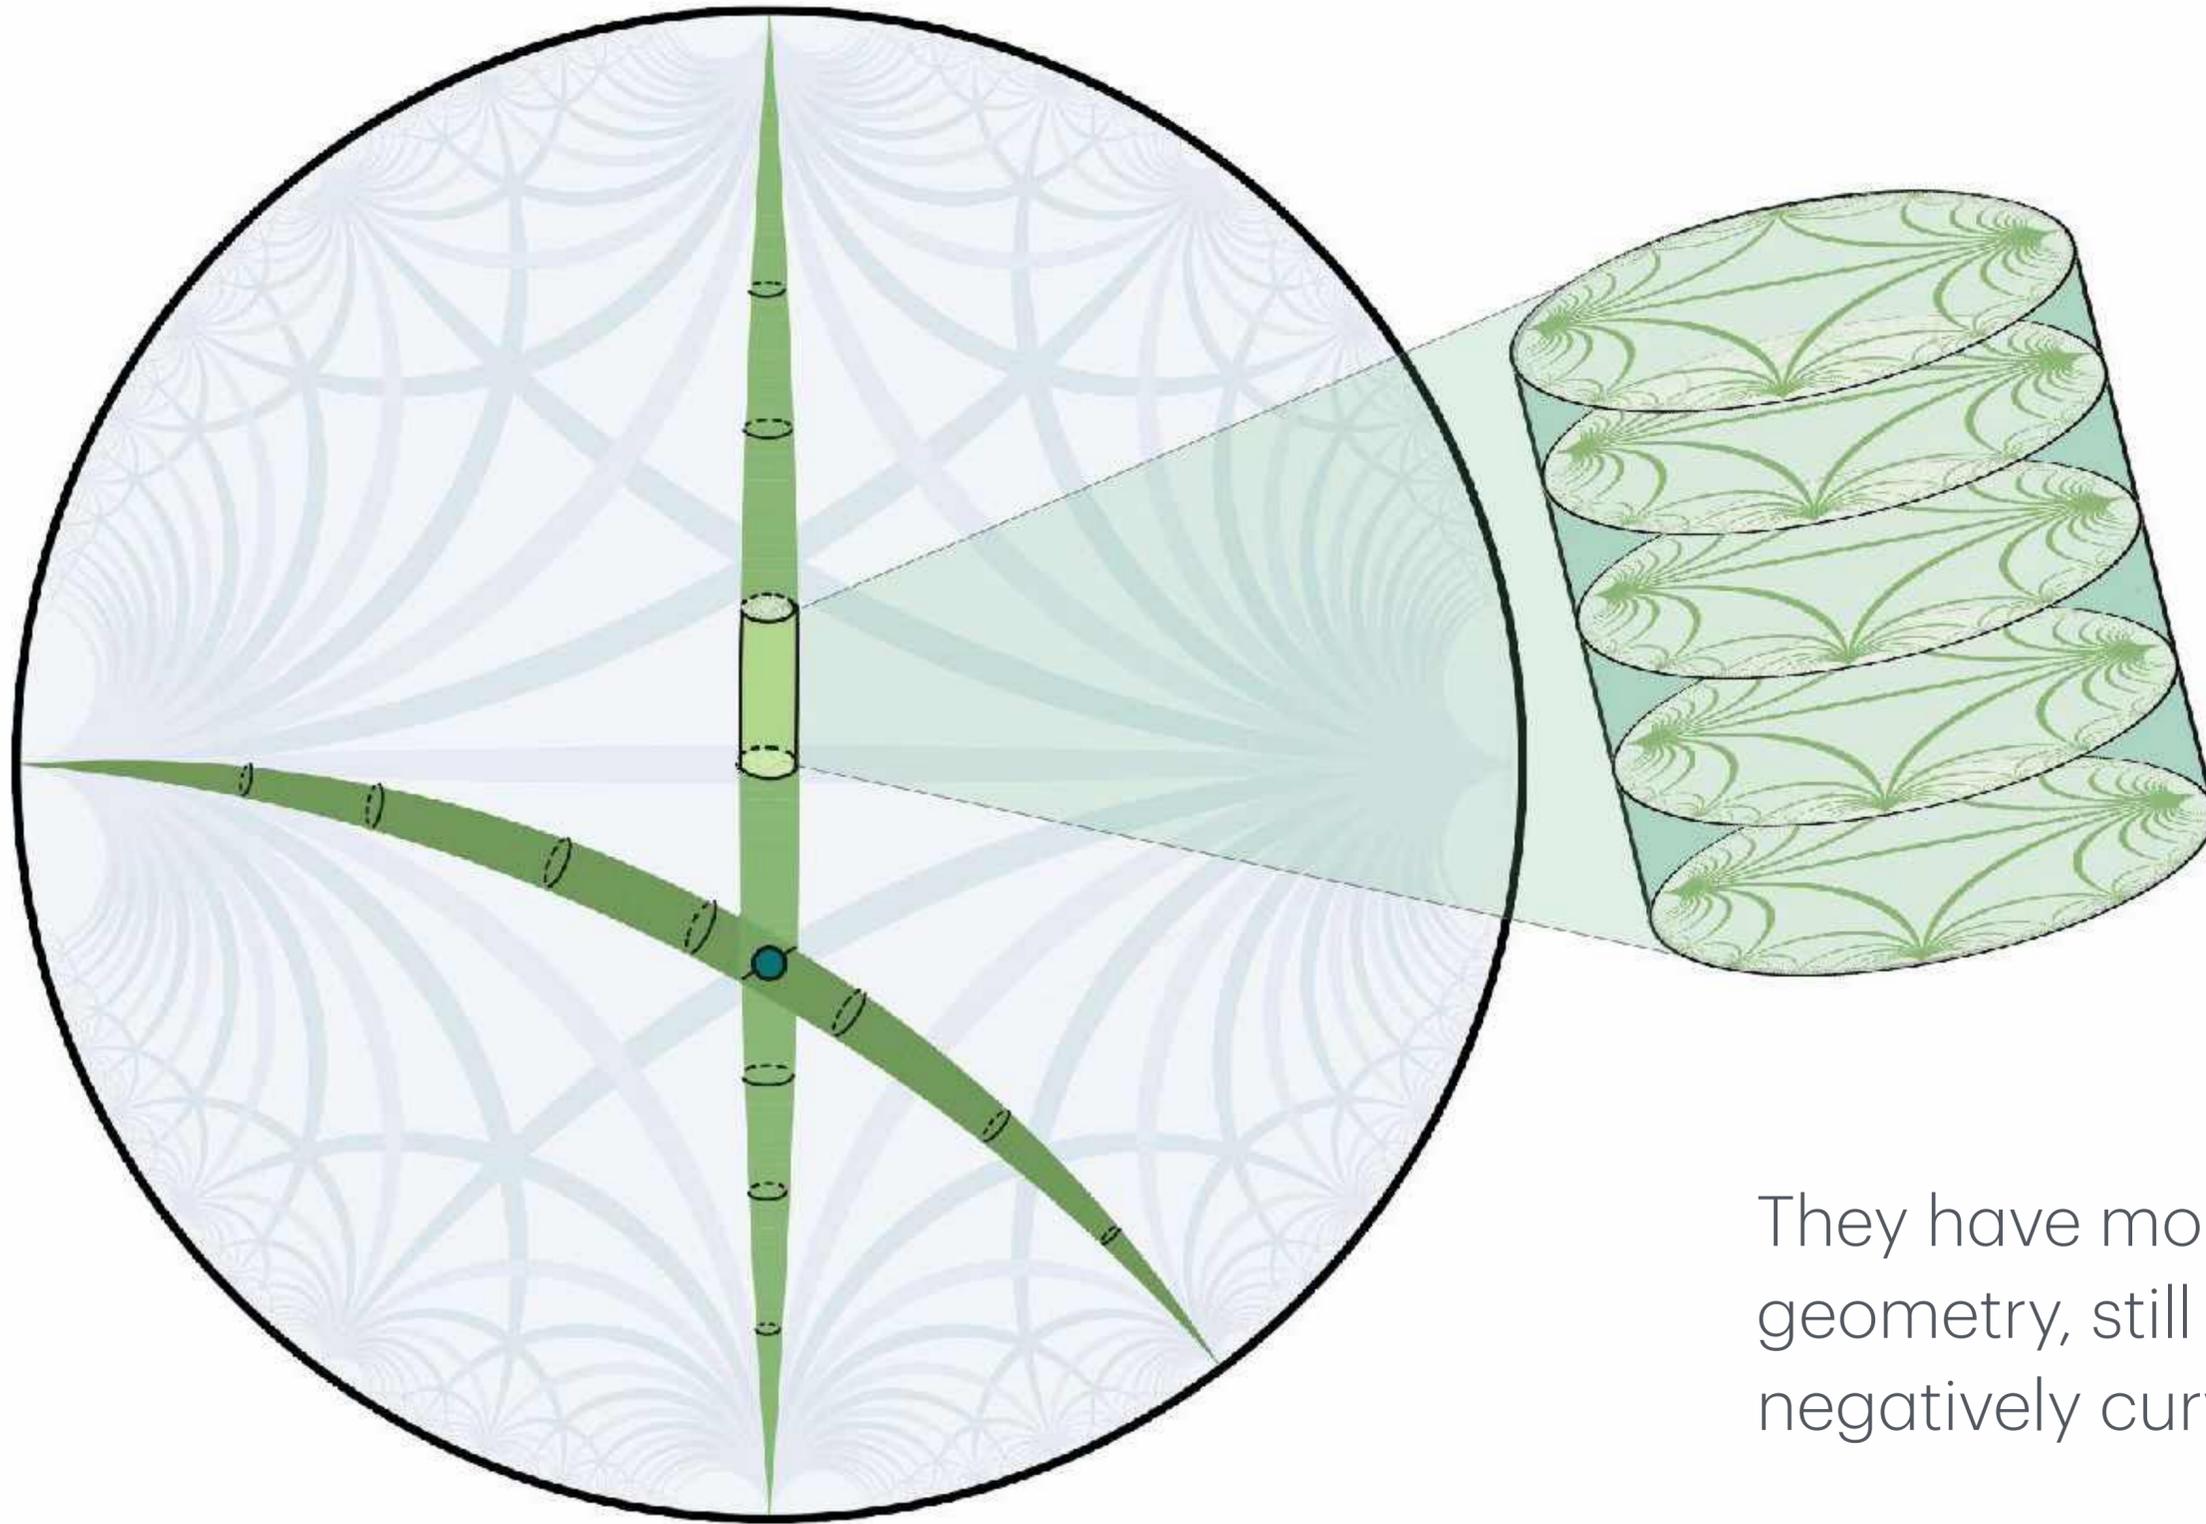


Room for grid-like stuff in the flat parts,

Room for tree-like stuff in the negative curvature

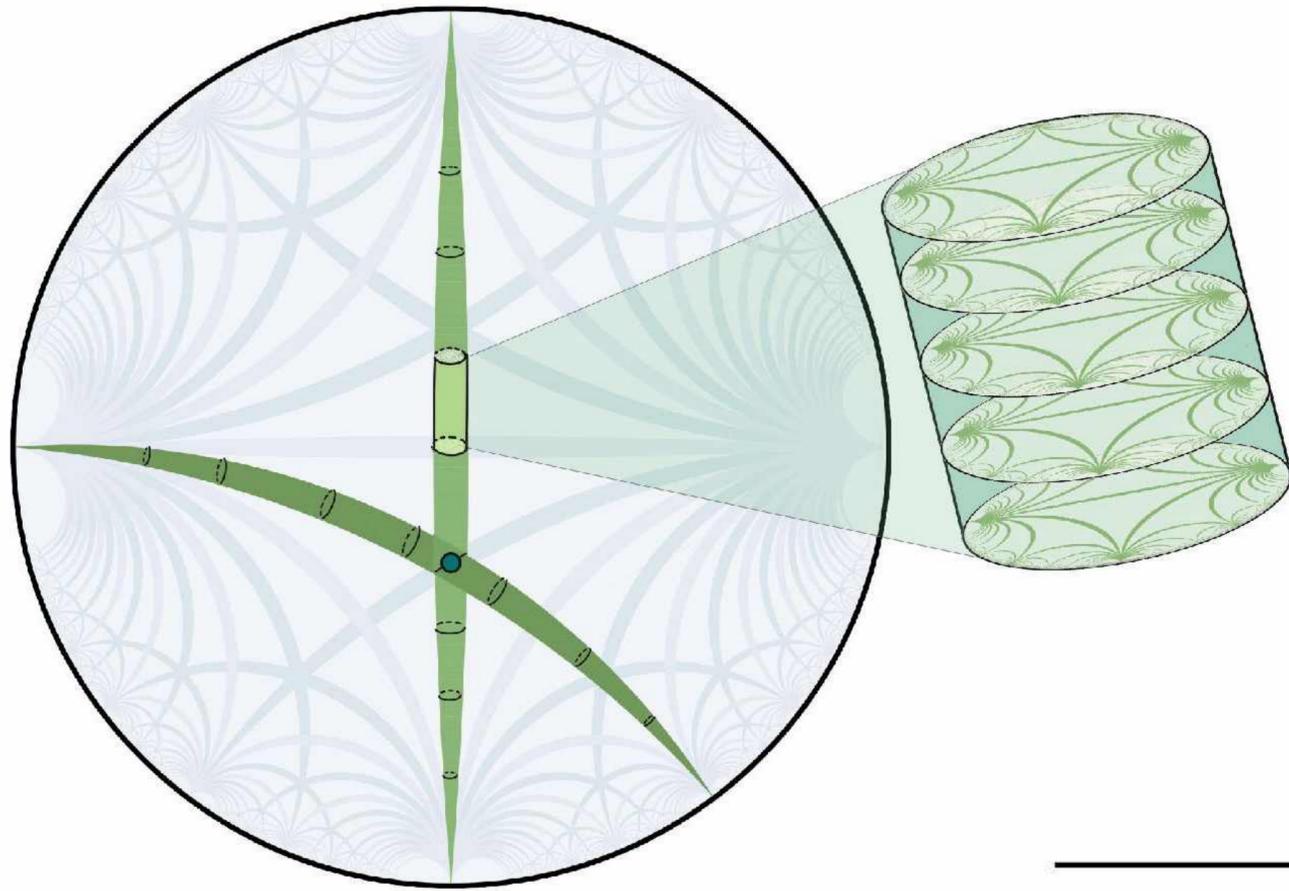


$$\mathbb{H}^2 \times \mathbb{R}$$



Higher dimensional versions of this are called *Symmetric Spaces*:

They have more complicated geometry, still full of both flat and negatively curved regions.



Higher dimensional versions of this are called *Symmetric Spaces*:

All symmetric spaces were classified in the 1920's by Elie Cartan.

Symmetric Spaces for Graph Embeddings

Type	Non-compact	Compact	$\text{rk}_{\mathbb{R}}$	dim
AI	$\text{SL}(n, \mathbb{R})/\text{SO}(n, \mathbb{R})$	$\text{SU}(n)/\text{SO}(n)$	$n - 1$	$\frac{(n-1)(n+2)}{2}$
A	$\text{SL}(n, \mathbb{C})/\text{SU}(2)$	$(\text{SU}(n) \times \text{SU}(n))/\text{SU}(n)$	$n - 1$	$(n + 1)(n - 1)$
BDI	$\text{SO}(p, q)/\text{SO}(p) \times \text{SO}(q)$	$\text{SO}(p + q)/\text{SO}(p) \times \text{SO}(q)$	$\min\{p, q\}$	pq
AIII	$\text{SU}(p, q)/\text{SU}(p) \times \text{SU}(q)$	$\text{SU}(p + q)/\text{SU}(p) \times \text{SU}(q)$	$\min\{p, q\}$	$2pq$
CI	$\text{Sp}(2n, \mathbb{R})/\text{U}(n)$	$\text{Sp}(2n)/\text{U}(n)$	n	$2n(n + 1)$
DIII	$\text{SO}^*(2n)/\text{U}(n)$	$\text{SO}(2n)/\text{U}(n)$	$\lfloor \frac{n}{2} \rfloor$	$n(n - 1)$
CII	$\text{Sp}(p, q)/\text{Sp}(p) \times \text{Sp}(q)$	$\text{Sp}(p + q)/\text{Sp}(p) \times \text{Sp}(q)$	$\min\{p, q\}$	$4pq$
AII	$\text{SL}(n, \mathbb{H})/\text{Sp}(n)$	$\text{SU}(2n)/\text{Sp}(n)$	$n - 1$	$(n - 1)(2n + 1)$
D	$\text{SO}(2n, \mathbb{C})/\text{SO}(2n)$	$(\text{SO}(2n) \times \text{SO}(2n))/\text{SO}(2n)$	n	$n(2n - 1)$
B	$\text{SO}(2n + 1, \mathbb{C})/\text{SO}(2n + 1)$	$(\text{SO}(2n + 1) \times \text{SO}(2n + 1))/\text{SO}(2n + 1)$	n	$n(2n + 1)$
C	$\text{Sp}(n, \mathbb{C})/\text{Sp}(n)$	$(\text{Sp}(n) \times \text{Sp}(n))/\text{Sp}(n)$	n	

Table 6. The classical symmetric spaces. Row CI represents the Siegel spaces and their compact duals.



Proposal: In the machine learning community, we propose to replace Linear Algebra with Riemannian Geometry on the Symmetric Space

Vector-valued Distance and Gyrocalculus on the Space of Symmetric Positive Definite Matrices

Steve Trettel
Stanford University

Federico López*
HITS

Michael Strube
HITS

Symmetric Spaces for Graph Embeddings: A Finsler-Riemannian Approach

Federico López¹ Beatrice Pozzetti² Steve Trettel³ Michael Strube¹ Anna Wienhard²

Abstract

Learning faithful graph representations as sets of vertex embeddings has become a fundamental intermediary step in a wide range of machine learning applications. We propose the systematic use of symmetric spaces in representation learning, a class encompassing many of the previously used embedding targets. This enables us to introduce a new method, the use of Finsler metrics integrated in a Riemannian optimization scheme, that better adapts to dissimilar structures in the graph. We develop a tool to analyze the embeddings and infer structural properties of the data sets. For implementation, we choose Siegel spaces, a versatile family of symmetric spaces. Our approach outperforms competitive baselines for graph reconstruction tasks on various synthetic and real-world datasets. We further demonstrate its applicability on two downstream tasks, recommender systems

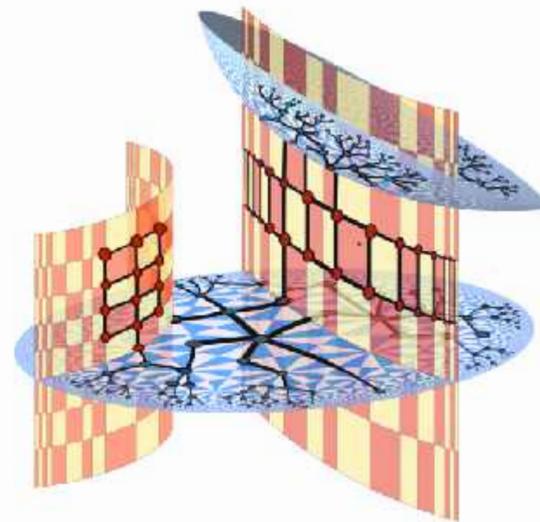


Figure 1. Symmetric spaces have a rich structure of totally geodesic subspaces, including flat subspaces (orange) and hyperbolic planes (blue). This compound, yet computationally tractable

Modeling Graphs Beyond Hyperbolic: Graph Neural Networks in Symmetric Positive Definite Matrices

Wei Zhang

Lopez¹, *J. Maxwell Riestenberg², Michael Strube¹,
Mdin Taha², and Steve Trettel³
¹Institute for Theoretical Studies, Germany
*fede.lastname}@h-its.org
²University of San Francisco, USA
zha}@mathi.uni-heidelberg.de
el@usfca.edu

that alignment between the structure of an embedding space is crucial for the data. The uniformity of the data allows for representing geometric features, such as grids and hierarchical and topological features, in symmetric spaces. In this paper, we develop an innovative approach to implement the neural networks in SPD. Experimental results on SPD and hyperbolic networks in SPD and complex graphs for reconstruction and data

arXiv:2110.13475v1 [cs.LG] 26 Oct 2021

Symmetric Spaces for Graph Embeddings

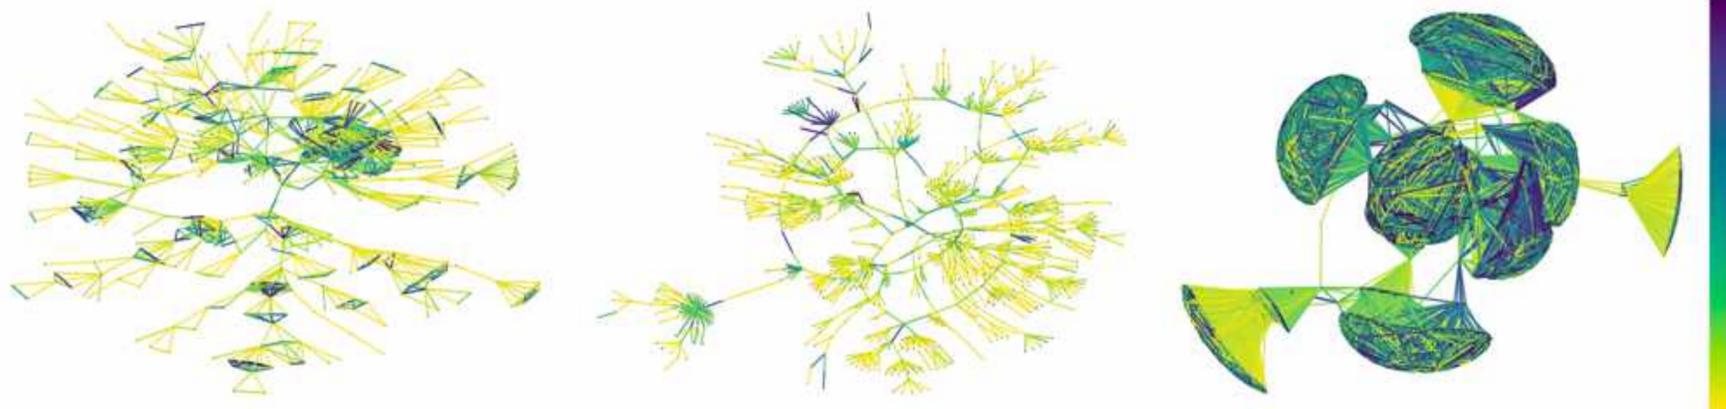


Figure 5. Edge coloring of S_2^{F1} for BIO-DISEASOME (left) and CSPHD (center) and FACEBOOK (right). Edge colors indicate the angle of the vector-valued distance for each edge, on a linear scale from 0 (yellow) to $\pi/4$ (blue).

Table 3: Results for Question Answering.

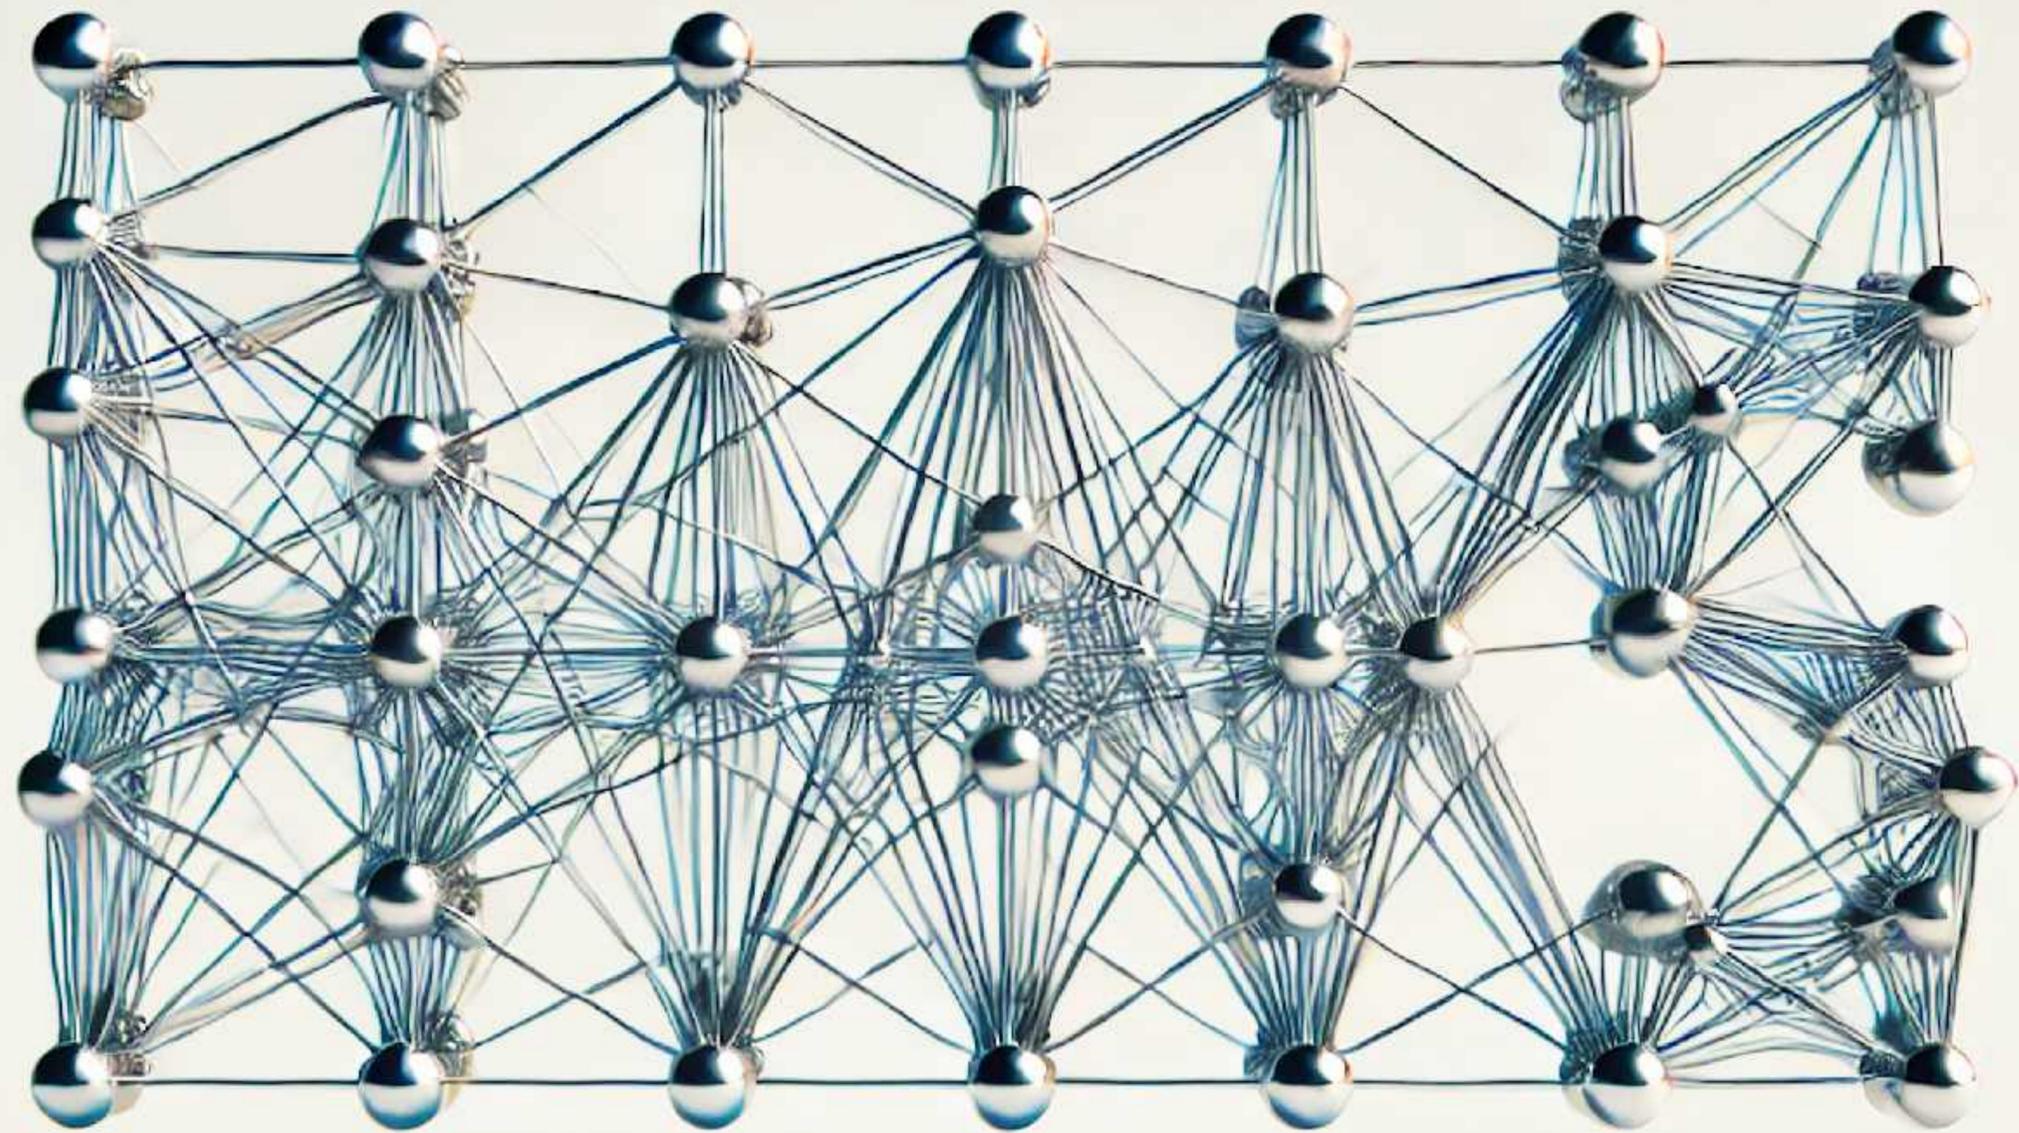
Model	TRECQA		WIKIQA	
	MRR	H@1	MRR	H@1
Euclidean	55.9±2.0	41.0±2.0	43.4±0.3	22.4±1.1
Hyperbolic	58.0±1.3	39.3±2.0	44.0±0.4	22.8±0.6
SPD _{Sca} ^R	55.4±0.1	37.1±0.1	45.5±0.5	24.4±1.1
SPD _{Sca} ^{F1}	57.1±0.7	38.6±0.2	44.8±0.5	24.0±0.6
SPD _{Rot} ^R	58.7±1.5	41.4±2.9	44.6±0.6	23.6±0.6
SPD _{Rot} ^{F1}	58.1±0.5	43.6±1.0	43.7±0.4	23.8±0.8
SPD _{Ref} ^R	57.3±0.3	40.7±1.1	43.9±0.7	23.4±2.0
SPD _{Ref} ^{F1}	59.6±0.5	42.1±1.0	44.7±1.2	25.0±2.5

Table 2: Results for Knowledge graph-based recommender systems.

Model	SOFTWARE		LUXURY		PANTRY		MINDREADER	
	MRR	H@10	MRR	H@10	MRR	H@10	MRR	H@10
TRANSE	28.5±0.1	47.2±0.5	35.6±0.1	52.3±0.1	16.6±0.0	35.3±0.1	19.1±0.4	37.6±0.1
ROTC	28.5±0.3	45.4±1.4	33.0±0.1	49.8±0.2	14.5±0.0	31.3±0.2	25.3±0.3	50.3±0.6
MURE	29.4±0.4	47.1±0.4	35.6±0.7	54.0±0.3	19.4±0.1	39.5±0.2	25.2±0.3	49.9±0.6
MURP	29.6±0.3	47.9±0.3	37.5±0.1	55.2±0.3	19.4±0.1	39.8±0.2	25.3±0.3	49.3±0.2
SPD _{Sca} ^R	29.4±0.4	48.1±0.8	37.5±0.2	55.1±0.2	19.5±0.0	39.6±0.3	25.4±0.1	49.8±0.3
SPD _{Sca} ^{F1}	28.8±0.1	46.9±0.5	37.3±0.3	54.1±0.9	19.0±0.1	38.8±0.2	25.7±0.5	49.5±0.1
SPD _{Rot} ^R	30.3±0.2	48.6±0.9	37.2±0.1	54.8±0.4	20.0±0.1	40.3±0.1	25.3±0.0	50.5±0.3
SPD _{Rot} ^{F1}	30.1±0.1	49.1±0.3	36.9±0.1	54.5±0.6	19.2±0.0	39.3±0.1	25.7±0.0	49.5±0.2
SPD _{Ref} ^R	29.6±0.2	48.0±0.5	37.3±0.2	55.0±0.2	19.3±0.0	39.7±0.3	25.3±0.0	49.1±0.1
SPD _{Ref} ^{F1}	29.3±0.1	47.5±0.6	36.8±0.0	54.8±0.1	18.6±0.2	38.3±0.3	24.8±0.2	47.9±1.8

Table 2. Comparison of different classifiers on Citeseer (top) and Cora (bottom). We bold the best accuracy in each row.

	\mathbb{R}^6		SPD ₃	
	LIN-XE	LIN-XE	SVM-MM	NC-MM
GIN	48.2 ± 6.3	68.0 ± 1.3	67.3 ± 1.2	67.0 ± 0.8
SGC	62.6 ± 3.4	69.4 ± 1.0	69.7 ± 0.8	67.9 ± 1.5
Cheb	63.2 ± 2.1	54.6 ± 10.4	61.4 ± 4.4	64.0 ± 2.3
GAT	55.0 ± 5.2	67.3 ± 1.7	69.2 ± 0.7	68.1 ± 1.1
GCN	64.7 ± 2.3	69.9 ± 0.8	69.2 ± 0.8	68.2 ± 1.0
GIN	77.1 ± 1.0	79.9 ± 0.6	79.5 ± 0.6	78.8 ± 1.0
SGC	75.7 ± 3.6	81.5 ± 0.9	81.8 ± 0.3	81.1 ± 0.6
Cheb	71.9 ± 2.8	75.5 ± 3.9	77.9 ± 2.4	79.2 ± 1.3
GAT	67.9 ± 4.2	79.4 ± 0.9	81.2 ± 1.4	81.2 ± 1.1
GCN	78.1 ± 1.7	79.7 ± 0.9	80.7 ± 0.5	80.2 ± 1.3



THANKS