Comparing Geometric Embeddings of Graphs

Bachelor's Thesis of

Linda Katharina Staerke

At the Department of Informatics
Institute of Theoretical Informatics (ITI)

Reviewer:	T.T.-Prof. Dr. Thomas Bläsius
Second reviewer:	Dr. rer. nat. Torsten Ueckerdt
Advisors:	T.T.-Prof. Dr. Thomas Bläsius
	Jean-Pierre von der Heydt

June 2025 – November 2025

Karlsruher Institut für Technologie
Fakultät für Informatik
Postfach 6980
76128 Karlsruhe

For the proof idea of Lemma 4.4, I have used lytris.

I declare that I have developed and written the enclosed thesis completely by myself. I have not used any other than the aids that I have mentioned. I have marked all parts of the thesis that I have included from referenced literature, either in their original wording or paraphrasing their contents. I have followed the by-laws to implement scientific integrity at KIT.

Karlsruhe, November 2025

Linda Staerke

. . . .

(Linda Katharina Staerke)

Abstract

A *geometric embedding* of a graph is an assignment of positions to vertices such that vertices are positioned “close” to each other if and only if they are connected. They have been an important tool in various fields of studies. Network scientists use hyperbolic and weighted geometry to generate graphs with properties of real-world social networks. In machine learning, one is interested in learning latent space representations of graphs such that the dot product between positions of vertices represents their similarity.

Both applications use social network graphs, thus it is of interest how well we can translate between the underlying geometric embeddings. However, almost no comparisons like this have been done. We make progress on this problem by comparing weighted, dot product and ball embeddings of graphs.

In this thesis, we show that under certain assumptions there is no universal translation from d-dimensional dot product embeddings to d-dimensional weighted embeddings. Then we turn our attention to dot product and ball graphs and compare their random graph models. To do this, we define what it means for one random graph model to be a *generalization up to constant factors* of another random graph model. This definition is inspired by the relation of geometric inhomogeneous random graphs (GIRGs) to hyperbolic random graphs proven by Bringmann et al. [BKL19]. We prove that $(d + 1)$-dimensional random dot product graphs generalize d-dimensional random ball graphs up to constant factors and vice versa.

Zusammenfassung

Eine *geometrische Einbettung* eines Graphen ist eine Zuordnung von Positionen zu Knoten, sodass Knoten genau dann “nahe” beieinander liegen, wenn sie im Graphen verbunden sind. Sie sind ein wichtiges Werkzeug in einer Vielzahl an Fachgebieten. Die Netzwerkwissenschaften benutzen hyperbolische und gewichtete Geometrie um Graphen mit Eigenschaften von echten sozialen Netzwerken zu generieren. Im Bereich des maschinellen Lernens ist man daran interessiert latente Repräsentationen von Graphen zu lernen, sodass das Skalarprodukt zwischen den Positionen von Knoten ihre Ähnlichkeit widerspiegelt.

In beiden Anwendungsfällen beschäftigt man sich mit Graphen von sozialen Netzwerk, deshalb ist es von Interesse, wie gut man die unterliegenden geometrischen Einbettungen ineinander übersetzen kann. Allerdings gab es bisher fast keine Vergleiche zwischen gewichteten, Skalarprodukt- und Kugel-Einbettungen (ball embeddings) von Graphen.

In dieser Abschlussarbeit zeigen wir, dass es unter gewissen Annahmen keine universelle Übersetzung von d-dimensionalen Skalarprodukt-Einbettungen zu d-dimensionalen gewichteten Einbettungen gibt. Anschließend fokussieren wir uns auf Skalarprodukt- und Kugelgraphen und vergleichen deren zugehörige Zufallsgraphmodelle. Dafür definieren wir, wann ein Zufallsgraphmodell ein anderes *bis auf konstante Faktoren verallgemeinert*. Diese Definition ist inspiriert durch den Zusammenhang zwischen geometrischen, inhomogenen Zufallsgraphen (GIRGs) und hyperbolischen Zufallsgraphen, die durch Bringmann et al. bewiesen wurde [BKL19]. Wir zeigen das $(d + 1)$-dimensionale Skalarprodukt-Zufallsgraphen die d-dimensionalen Kugel-Zufallsgraphen bis auf konstante Faktoren verallgemeinern und umgekehrt.

Contents

1 Introduction

Realistic graph models have various applications in different fields such as network science, social science and biology. As real-world instances are often scarce, there is a high interest to recreate their properties, e.g., through machine learning and random graph theory.

A common tool used by both of these disciplines are *geometric embeddings* of graphs. In a geometric embedding of a graph, vertices are assigned positions in a geometric space of some dimension d, e.g., the Euclidean space $\mathbb{R}^d$, such that vertices are connected if and only if their positions satisfy a specified geometrical condition [RRŠ89]. Typically, such a condition checks if the vertices are "close" enough to be connected. Some embeddings also assign further properties to vertices, like a weight, for even more nuanced embeddings. In such cases the condition for connectivity does not only depend on the position of the vertices in question but also their other properties.

In practice, geometric embeddings are used in the following two ways. First, they can be used to represent graphs in lower dimensional, continuous spaces in the following sense. If the dimension of the geometric space is significantly lower than the number of vertices, one can store the positions of the vertices instead of the full discrete adjacency matrix. Connection between vertices can then be checked through the (possibly differentiable) connectivity condition of the chosen geometric embedding. This requires criteria that can easily be checked. One important instance of this principle are latent space representations of graphs used in machine learning [MKNŠ21 | Xu21]. There, the goal is to learn low-dimensional graph embeddings in continuous spaces, such that similar vertices get positioned close together. It is then much faster and more accurate to analyze graphs in this continuous setting instead of using their discrete symbolic representation [Xu21].

Second, one can study random graph models emerging from a specific geometry. In these random graph models positions of vertices are drawn independently from a distribution over the desired geometric space. Usually this is a uniform distribution over a bounded subset of the space. If there exist other properties, they are drawn independently from a chosen distribution over the domain of said properties. Then one uses the connectivity condition from the geometric model to calculate the resulting edges in the resulting random graph.

There is high research interest in random graph models that produce graphs with similar properties than real-world graphs, e.g., as testing instances for graph algorithms. Geometric models have the advantage that they already share certain properties with real-world graphs. One shared property is locality, i.e., the fact, that two vertices with a common neighbor are more likely to be connected. In geometric graphs this results from the fact that vertices which are connected are usually positioned close together. For other properties some models are more suitable than others. For example, many real-world network graphs admit a heterogeneous degree distribution and there is research showing that their inherent geometry is hyperbolic rather than Euclidean [DT19]. Thus, one is interested in geometric embeddings, that have these properties.

The most basic example of geometric embeddings (without any extra properties) are *unit disk graphs* which have been studied extensively [CCJ90]. In this model, vertices are embedded in 2-dimensional Euclidean space and two vertices are connected if and only if their Euclidean distance is at most 1. Their name comes from the fact, that unit disk graphs are precisely the intersection graphs of unit disks in the Euclidean plane. This model easily generalizes into higher dimensions yielding *unit ball graphs*, i.e., intersection graphs of unit d-balls in $\mathbb{R}^d$. In formulas, for two vertices u and v and their positions $p_u, p_v \in \mathbb{R}^d$ it holds that

$$uv \in E \iff \|p_u - p_v\|_2 \leq 2.$$

When embedding graphs, this model has some drawbacks, mostly that certain graphs need a high number of dimensions to be embedded correctly. One example are trees, which need up to logarithmic dimension in the number of the vertices of the tree [RRŠ89]. This can be fixed by introducing a ball radius and looking at their intersections, yielding the model of *ball graphs*. In this model, trees need only two dimensions to be represented. In ball graphs, vertices get assigned not only a position $p_v \in \mathbb{R}^d$ but also a radius r_v. Vertices are connected if and only if the corresponding balls with midpoint p_v and radius r_v intersect. In formulas, for two vertices u and v this reads

$$uv \in E \iff \|p_u - p_v\|_2 \leq r_u + r_v.$$

An embedding that is widely used in machine learning, is using the dot product of the positions as a connectivity criterion [Xu21 | Ahm+13 | Ou+16]. It yields the following model, independently introduced by Reiterman et al. and Fiduccia et al. [RRŠ89 | FSTZ98] A d-dimensional *dot product embedding* of a graph is a geometric embedding of said graph in $\mathbb{R}^d$ where two vertices u and v are connected if and only if the dot product of their corresponding vectors x_u, x_v is greater than 1 [RRŠ89 | FSTZ98]. Contrary to the embeddings above, the inequality in the condition is the other way around, as the dot product gets larger if two vectors point in the same direction. In formulas for vertices u, v it holds that

$$uv \in E \iff \langle x_u, x_v \rangle \geq 1.$$

As an example of finding geometric embeddings whose corresponding random graph models have properties of real-world graphs, we introduce an embedding in hyperbolic space. To do this, the idea that comes to mind is considering unit disk graphs in hyperbolic space. However, unlike in the Euclidean setting, the size of the disk has impacts on which graphs can be represented. In a *hyperbolic unit disk graph*, vertices are assigned positions in the hyperbolic plane. Vertices are connected if and only if their hyperbolic distance is smaller than a fixed threshold. The corresponding random graph model has a lot of properties of real-world network graphs, including the mentioned heterogeneous degree distribution, which have been studied extensively [PKBV10 | GPP12 | BFM13 | BFK18]. One drawback of this model is that generating hyperbolic random graphs is mathematically complex and can be numerically unstable.

Bringmann et al. have suggested a similar model to hyperbolic random graphs, the so-called *geometric inhomogeneous random graphs* (GIRGs), which are easier mathematically more simple [BKL19]. The underlying geometric embedding has a weighted Euclidean geometry. This means that each vertex gets a position p_v in $\mathbb{R}^d$ and a positive weight w_v. Two vertices are connected if and only if their distance is smaller than the product of their weights. In formulas for vertices u and v the connectivity condition reads

$$uv \in E \iff \|p_u - p_v\|_2 \leq w_u w_v.$$

Especially in the past years there has been extensive research in GIRGs and their efficient generation [KL20 | Blä+22 | MS22 | FGKS24].

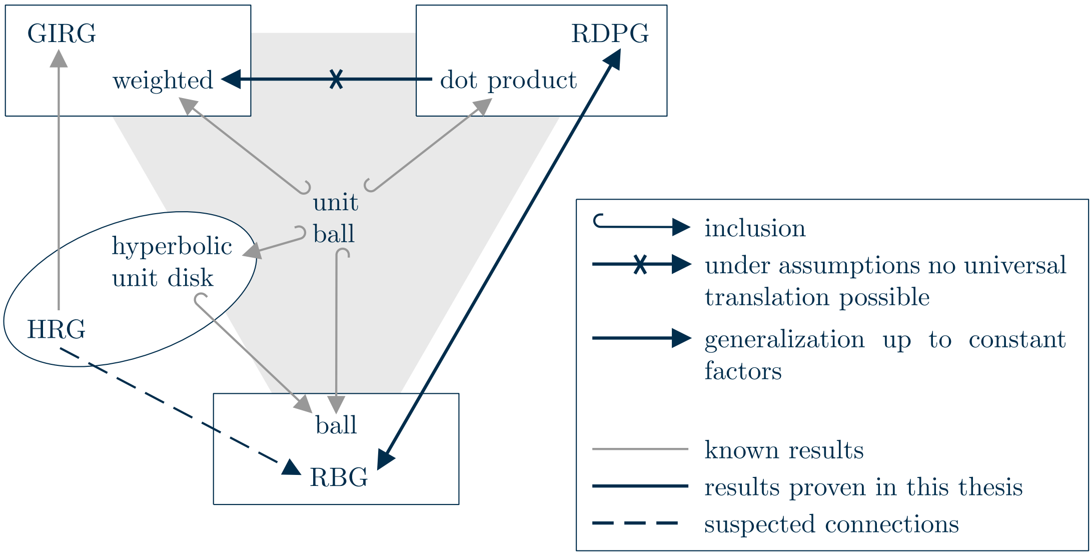

Figure 1.1: Relations of different types of embeddings and their corresponding random graph models. Shown are unit ball graphs, hyperbolic unit disk graphs, ball embeddings, dot product embeddings and weighted embeddings. Their respective random graph models (except for unit ball graphs) are hyperbolic random graphs (HRG), random ball graphs (RBG), random dot product graphs (RDPG) and GIRGs. Note that the inclusion of unit ball graphs into hyperbolic unit disk graphs only works in dimension 2, i.e., for unit disk graphs.

Summarizing, we have introduced three major geometric embeddings with different types of underlying geometry: ball embeddings, which use Euclidean geometry; dot product embeddings, which are related to spherical geometry and weighted embeddings, which are a generalization of hyperbolic geometry. Dot product embeddings are mainly applied for graph embeddings by the machine learning community and weighted embeddings are mainly used in random graph models

Having these different embeddings, there naturally arise questions regarding the connections between these models. For example, one could ask which embeddings are similar to each other or whether one can transform embeddings of one model into an embedding of another model. On a random graph level one could ask whether the different corresponding random graph models can generate the same graphs, or if some models are more expressive than others. However, between these three big embeddings almost no comparisons have been done.

1.1 Contribution

In this thesis we bridge this gap between different graph models. We give an overview about existing relationships and also introduce new ones. All results are summarized in Figure 1.1 with existing comparisons indicated by gray arrows and our contributions indicated by black arrows.

We start by explaining the inclusion arrows in Figure 1.1. It is easy to see that unit ball graphs are a special case of every other model in the sense that every unit ball graph embedding is or can be transformed into an embedding of the other models. For ball graphs and weighted embeddings, d-dimensional unit ball graphs have an inherent d-dimensional ball or weighted embedding by just setting all radii or weights to 1. For dot product graphs one notices that the angular distance on the surface of the d-sphere (calculated through the dot product) is very similar to the d-dimensional Euclidean distance if one restricts itself to a small section. Thus, one can shrink the d-dimensional unit ball embedding and map it onto a sphere to get a $(d + 1)$-dimensional dot product embedding, where all vectors x_v of the vertices have the same lengths. This fact has been proven by Reiterman et al. [RRŠ89]. One can also transform (2-dimensional) unit disk graphs into hyperbolic unit disk graphs, this fact was proven by Bläsius et al. [BFKS23]. Moreover, if one considers the Poincaré disk model of the hyperbolic plane, every hyperbolic disk looks like a Euclidean disk [And05]. Thus, every hyperbolic unit disk graph is a Euclidean disk graph with potentially different-sized disks.

Moving onto the big models, we first show that there is no universal translation from a d-dimensional dot product embedding to d-dimensional weighted embedding of the same graph, that uses the directions of the vectors as positions and turns their lengths into weights. This fact is indicated in Figure 1.1 by the crossed arrow. It suggests that weighted and dot product graphs are hard to compare. We prove however, that one can always transform a 1-dimensional dot product embeddings into 1-dimensional weighted embeddings.

We then show, that there is a somewhat natural way to translate a d-dimensional ball embedding into a $(d + 1)$-dimensional dot product embedding where all vectors have length at least 1. However, this translation is not exact, as it only produces a subgraph of the ball graph. It can be modified to give a supergraph of the ball graph. There one could hope, that the sub- and supergraph are very similar, in the sense that the number of wrongly introduced edges is small. However, there are malicious counterexamples examples, where this is not the case, e.g., the subgraph can be just an independent set. Instead, we consider random models and show that when considering random drawn ball graphs, then the expected degrees of the sub- and supergraph differ only by a constant factor.

To make this statement precise, we define what it means that one random graph model *generalizes* another *up to constant factors*. Loosely this definition will say, that if one draws a random graph from a random graph model $\mathcal{G}$ and then finds a sub- and a supergraph that come from another model $\mathcal{H}$ such that their expected average degrees differ only by a constant factor, then we say that $\mathcal{H}$ generalizes $\mathcal{G}$ up to constant factors. This definition is not arbitrary, as Bringmann et al. have shown that the GIRG model generalizes hyperbolic random graphs up to constant factors [BKL19]. This generalization was already mentioned above, and is indicated as a gray arrow in Figure 1.1.

We consider a random ball graph model where positions are drawn in a small hypercube independently and uniformly at random and the radii are drawn independently from a chosen distribution. Our random dot product graph model draws directions in a part of $\mathbb{R}^{d+1}$ almost uniformly and independently at random and lengths independently from a chosen distribution on $\mathbb{R}_{\geq 1}$. To get the vector for each vertex one scales the direction with the given lengths.

As already indicated above, we will show that that our $(d + 1)$-dimensional random dot product graph model will generalize the d-dimenisional ball graph model up to constant factors. Moreover, we will see that under certain assumptions the translations from ball graphs to dot product graphs also work in the other direction. This will imply that the generalization up to constant factors will work in both directions as indicated by the arrow in Figure 1.1.

We suspect that one could also translate the inclusion of hyperbolic unit disk graphs in ball graphs to the corresponding random models, i.e. that the random ball graph model generalizes the hyperbolic random graph model (HRG) up to constant factors. In Figure 1.1 this is indicated by a dashed arrow.

1.2 Related Work

Apart from GIRGs and hyperbolic graphs, there is also research that uses random dot product graphs as a model for social networks [YS07]. This random graph model differs from the one that we use in this thesis. For each vertex they choose vectors in $[0, 1]^d$ with length at most 1 (rather than at least 1) uniformly and independently at random and then use the dot product between the vectors of two vertices as the edge probability of the corresponding edge. This model is also widely used as latent space representations in machine learning applications [STP14 | SBNH21 | RCTP22].

For different graphs and graph classes the minimum dimension needed for an embedding, the so-called *embedding dimension*, has also been studied. For unit disk graphs this has been done under the term *sphericity* [Fis83 | Mae84 | MQ06]. Especially for the class of dot product graphs, there has been research interest in studying the corresponding embedding dimension. For certain graph classes, the dot product dimension is known e.g. for trees, where it is 3 [RRŠ89 | FSTZ98] or for planar graphs, where it is 4 [KM11a]. There also has been research on the upper bound of the dot product dimension depending on the number of vertices of the graph [FSTZ98 | LC14]. It is known, that deciding whether the dot product dimension of a given graph G is below a threshold $d > 1$ is NP-hard [KM11b]. For $d = 1$ all graphs with this dot product dimension are known by a classification of forbidden induced subgraphs [RRŠ89 | FSTZ98]. There has also been research about which graphs have dot product dimension 2 [JPL21]. As hyperbolic unit disk graphs and weighted embeddings have almost exclusively been studied in the random graph setting, there are very few results about the corresponding embedding dimensions. It is known that trees have weighted dimension [Ber25]. And there are some results about the generalization of hyperbolic unit disk graphs to higher dimension and the respective embedding dimension called *hyperbolic sphericity* [Gro24].

1.3 Outline

The remainder of this thesis is structured as follows. First, in Chapter 2, we introduce the necessary mathematical concepts and fix our notation. In Chapter 3, we compare d-dimensional dot product and weighted embeddings and show that they differ in certain aspects. At the end of the chapter, we show a translation from dot product embeddings to weighted embeddings in dimension 1. In Chapter 4, we compare d-dimensional ball embeddings to $(d + 1)$-dimensional dot product embeddings and their respective random graph models by studying their relation to an intermediate embedding. We study translations from ball graphs to this intermediate embedding that work well with random graph theory in Section 4.1. Having these translations, we show in Section 4.2 that the random graph models corresponding to ball embeddings and dot product embeddings produce similar distributions of graphs. Finally, in Chapter 5 we summarize our findings and provide an outlook to future work.

2 Preliminaries

In this chapter we first fix our notation about graphs. Then, we introduce the different geometric embeddings of graphs, their corresponding random models and a tool to compare them.

2.1 Graph Theory

A *graph* $G = (V, E)$ consists of a set of *vertices* V together with a set of *edges* $E \subseteq \binom{V}{2} = \{e \subseteq V \mid |e| = 2\}$. We write uv as a short notation for an edge $\{u, v\} \in E$. In particular this means that all our graphs are undirected and without loops and multi-edges. We also write $V(G) := V$ and $E(G) := E$ if we want to make clear to which graph the vertex or edge set belongs. For $uv \in E$ we say that u and v are *connected* or u is a *neighbor* of v. The number of neighbors of a vertex v is called its *degree* $\deg(v)$. The *average degree* of G is denoted as $\overline{d} = \frac{1}{|V|} \sum_{v \in V} \deg(v)$.

A graph $H = (V', E')$ is called a *subgraph* of G if $V' \subseteq V$ and $E' \subseteq E$. We then write $H \subseteq G$ and call G a *supergraph* of H.

The *induced subgraph* $G[V']$ of a set of vertices $V' \subseteq V$ of G is the subgraph $H = (V', E')$ of G with edges $E' = \{xy \in E \mid x, y \in V'\}$. In other words E' contains precisely all edges between vertices in V' in G.

For two graphs G_1, G_2 we define their *disjoint union* as the graph G with vertices $V = V(G_1) \sqcup V(G_2)$ and edges $E = E(G_1) \sqcup E(G_2)$, where $\sqcup$ denotes the disjoint union of sets. In particular in G, there are no edges between vertices in $V(G_1)$ and $V(G_2)$.

2.2 Geometric Embeddings of Graphs

In this thesis we study different kinds of geometric embeddings of graphs. Before we define them, we first fix some geometrical notations. For this whole chapter $d \in \mathbb{N}^+$ refers to a dimension.

Euclidean Geometry We use the d-dimensional Euclidean space $\mathbb{R}^d$ with the standard norm $\|\cdot\|_2$. We also use the maximum norm $\|\cdot\|_\infty$. For $x \in \mathbb{R}^d$ it holds that $\|x\|_\infty \leq \|x\|_2 \leq \sqrt{d}\|x\|_\infty$. For $x, y \in \mathbb{R}^d$ we denote their dot product as $\langle x, y \rangle$. For a (measurable) bounded subspace $A \subseteq \mathbb{R}^d$ we denote its volume with $\mathrm{Vol}(A)$.

For a point $p \in \mathbb{R}^d$ and $r \in \mathbb{R}_0^+$ the *d-dimensional ball with center p and radius r* is the subset of $\mathbb{R}^d$ of all points with Euclidean distance at most r from the center p, i.e., $\{x \in \mathbb{R}^d \mid \|p - x\|_2 \leq r\}$.

A d-dimensional *hypercube* is a subset of $\mathbb{R}^d$ of the form $\bigtimes_{i=1}^d [a_i, a_i + s]$ where $a_i \in \mathbb{R}$ for $i \in \{1, \ldots, d\}$ and $s \in \mathbb{R}^+$.

Spherical Geometry The *d-dimensional sphere* S^d is the subset of $\mathbb{R}^{d+1}$ of all points with distance 1 from the origin, i.e., $S^d = \{x \in \mathbb{R}^{d+1} \mid \|x\|_2 = 1\}$. For $d \geq 2$ and $x, y \in S^d$ we define the *angular distance* dist_S of x, y as $\mathrm{dist}_S(x, y) = \arccos(\langle x, y\rangle)$, where arccos is the inverse function of cosine in the interval $[0, \frac{\pi}{2}]$[1]. The angular distance function fulfills the triangle inequality.

For a point $p \in S^d$ and $\theta \in [0, \frac{\pi}{2}]$ the *cap with center p and angular radius* θ is the subset of S^d of all points with angular distance at most θ from the center p, i.e., the subset $\{x \in S^d \mid \mathrm{dist}_S(p, x) = \arccos(\langle p, x\rangle) \leq \theta\}$. This is the spherical analog of a ball in Euclidean geometry. Applying cosine to the condition we get $\{x \in S^d \mid \langle p, x\rangle \geq \cos(\theta)\}$, i.e., a cap is the intersection of a half-space with S^d.

A translation between Euclidean and spherical geometry are polar coordinates, which use Euclidean coordinates as angles for positions on the sphere. We define the *$(d+1)$-dimensional polar coordinates* $\Phi_{d+1} \colon [-\frac{\pi}{2}, \frac{\pi}{2}]^d \to S^d$ inductively as follows.[2] For $d = 1$ we define:

$$\Phi_2(p) = \begin{pmatrix} \cos(p) \\ \sin(p) \end{pmatrix}.$$

For $d \geq 2$ we define

$$\Phi_{d+1}(p) = \begin{pmatrix} \cos(p_d) \cdot \Phi_d(p_{[d-1]}) \\ \sin(p_d) \end{pmatrix},$$

where $p_{[k]} = \begin{pmatrix} p_1 & \dots & p_k \end{pmatrix} \in \mathbb{R}^k$ for $p \in \mathbb{R}^d$ and $k \leq d$. If one unrolls the inductive definition, then the kth entry of $\Phi_{d+1}(p)$ is $\Phi_{d+1}(p)_k = \sin(p_{k-1}) \prod_{i=k}^{d} \cos(p_i)$ where $\sin(p_0) := 1$. We will omit the dimension index and write only Φ for the polar coordinates if the dimension is clear.

Another translation between Euclidean and spherical geometry is the *stereographic projection*. It establishes a bijection $S^d \setminus \{N\} \to \mathbb{R}^d$, where $N = \begin{pmatrix} 0 & \dots & 0 & 1 \end{pmatrix} \in \mathbb{R}^{d+1}$ is the "North Pole" of the sphere, by mapping a point on the sphere to the intersection of the ray from N through that point with the hyperplane spanned by the first d coordinates[3]. In particular, stereographic projection maps caps to balls and vice versa. This is the only property of the stereographic projection that we use in this thesis. However, centers and radii are not preserved.

A geometric embedding of a graph is an assignment of positions to vertices such that vertices are connected if and only if their positions fulfill a specified geometric condition. We work with the following four embeddings in this thesis.

Ball Embedding A d-dimensional *ball embedding* of a graph $G = (V, E)$ is a map of positions $p \colon V \to \mathbb{R}^d, v \mapsto p_v$ and an assignment of radii $r \colon V \to \mathbb{R}_0^+, v \mapsto r_v$ such that two vertices $u, v \in V$ are connected if and only if the balls B_u with center p_u and radius r_u and B_v with center p_v and radius r_v intersect. In formulas, for vertices $u, v \in V$ it holds that

$$uv \in E \iff \|p_u - p_v\|_2 \leq r_u + r_v.$$

We write $(p_\bullet, r_\bullet)$ as a short notation for the embedding and call G a *d-dimensional ball graph*. If there exist an embedding of G such that all vertices have the same radius r we call G a *unit ball graph*.

[1]Not to be confused with $\cos^{-1} = \frac{1}{\cos}$.

[2]Usually the first angle is defined to come from $[-\pi, \pi]$ such that Φ is surjective. We will not need this in this thesis, as we restrict ourselves to small sections of the sphere

[3]This hyperplane defines the "equator" of the sphere.

Cap Embedding A d-dimensional *cap embedding* of a graph $G = (V, E)$ is a map of positions $p\colon V \to S^{d-1}, v \mapsto p_v$ and an assignment of angular radii $\theta\colon V \to [0, \frac{\pi}{2}], v \mapsto \theta_v$ such that two vertices $u, v \in V$ are connected if and only if the caps C_u with center p_u and angular θ_u and C_v with center p_v and angular radius θ_v intersect. In formulas, for vertices $u, v \in V$ it holds that

$$uv \in E \iff \mathrm{dist}_S(p_u - p_v) \leq \theta_u + \theta_v.$$

We write $(p_\bullet, \theta_\bullet)$ as a short notation for the embedding and call G a *d-dimensional cap graph.* This embedding is the analog of ball embeddings in spherical geometry.

Weighted Embedding A d-dimensional *weighted embedding* of a graph $G = (V, E)$ is a map of positions $p\colon V \to \mathbb{R}^d, v \mapsto p_v$ and an assignment of weights $w\colon V \to \mathbb{R}_0^+, v \mapsto w_v$ such that two vertices $u, v \in V$ are connected if and only if the distance of their positions is at most the product of their weights. In formulas, for vertices $u, v \in V$ it holds that

$$uv \in E \iff \|p_u - p_v\|_2 \leq w_u \cdot w_v.$$

We write $(p_\bullet, w_\bullet)$ as a short notation for the weighted embedding. Weighted embeddings are the underlying geometric embeddings of geometric inhomogeneous random graphs (GIRGs).

Dot Product Embedding A d-dimensional *dot product embedding* of a graph $G = (V, E)$ is a map of directions $p\colon V \to S^{d-1}, v \mapsto p_v$ and a map of lengths $\ell\colon V \to \mathbb{R}^+$ such that two vertices $u, v \in V$ are connected if and only if the dot product of their corresponding vectors $x_u = \ell_u p_u$ and $x_v = \ell_v p_v$ is at least 1. In formulas, for vertices $u, v \in V$ it holds that

$$uv \in E \iff \langle \ell_u p_u, \ell_v p_v \rangle \geq 1.$$

We write $(p_\bullet, \ell_\bullet)$ as a short notation for the embedding and call G a *d-dimensional dot product graph.* Note that in this case the inequality goes in the other direction, as the dot product between vectors is large if they point in a similar direction. A dot product embedding is called *nonnegative* if all directions have only nonnegative entries. In particular this implies that all pairwise dot products $\langle \ell_u p_u, \ell_v p_v \rangle$ are nonnegative. A lot of the times we will restrict ourselves to dot product embeddings where every vertex has length at least 1.

2.3 Random Graphs

Before we define the random graph models and the comparison tool, we first introduce some basic probability theory.

Probability Theory We use capital letters for random variables, e.g. $G, H, L, \ldots$ and calligraphic letters for distributions like $\mathcal{G}, \mathcal{H}, \mathcal{L}$. If a random variable G has distribution $\mathcal{G}$ we write $G \sim \mathcal{G}$ as a short notation. For continuous distributions, we assume that they have probability density functions. One well-known continuous distribution that we use is the *uniform distribution* on a hypercube $K \subseteq \mathbb{R}^d$ denoted as $\mathcal{U}(K)$. We use $\mathbb{P}$ for probabilities and $\mathbb{E}$ for the expected value of a random variable in the usual way. For an event A we define its indicator random variable $\mathbb{1}_A$, which is 1, if the event occurred and 0 otherwise. It holds that $\mathbb{E}[\mathbb{1}_A] = \mathbb{P}(A)$.

A *random graph model* is a probability distribution on a set of graphs. This distribution can depend on parameters, which can be constants, functions or other probability distributions. One example for such a parameter is the number of vertices n, i.e., the chosen set of graphs are the graphs on n vertices. All random graph models that we discuss have the number of vertices as a parameter. We use calligraphic letters for random graph models and write the parameters in parentheses, i.e., like $\mathcal{G}(n, \beta)$, where n is the number of vertices and β is a place-holder for other parameters. For a random variable G that has the distribution of a random graph model $\mathcal{G}(n, \beta)$ we write $G \sim \mathcal{G}(n, \beta)$.

In this thesis we want to compare different random graph models to make statements like that the expected average degree of a model $\mathcal{G}$ is lower than the one from another model $\mathcal{H}$. For this to work easily we can show something even stronger. We want to be able to couple random variables $G \sim \mathcal{G}$ and $H \sim \mathcal{H}$ together, such that G has always lower average degree than H. Formally, for this to be well-defined we need that the two random variables $G \sim \mathcal{G}$ and $H \sim \mathcal{H}$ are defined on the same probability space. If this is the case, we call the random variables G and H a *coupling* of the distributions $\mathcal{G}$ and $\mathcal{H}$. We use the notation (G, H) for such a coupling. One can generalize this notion onto an arbitrary number of distributions. For distributions $\mathcal{G}_1, \dots \mathcal{G}_r$, we call random variables $G_1 \sim \mathcal{G}_1, \dots G_r \sim \mathcal{G}_r$ a coupling of the distributions if $G_1, \dots, G_r$ are defined on the same probability space. Like above, we use $(G_1, \dots, G_n)$ as notation for this coupling.

An easy way to construct a coupling between two random graph models $\mathcal{G}$ and $\mathcal{H}$ is the following. Start with a random variable $G \sim \mathcal{G}$ and construct a function f that maps graphs to graphs such that $f(G) \sim \mathcal{H}$. Then G and $f(G)$ are a coupling of $\mathcal{G}$ and $\mathcal{H}$.

Let $\mathcal{G}(n, \beta)$ and $\mathcal{H}(n, \gamma)$ be random graphs models, i.e., probability distributions over the graphs on n vertices depending on parameters β and γ. We say that $\mathcal{H}$ *generalizes* $\mathcal{G}$ *up to constant factors* if for every possible β there exist parameters γ_1 and γ_2 and a coupling (G, H_1, H_2) between $\mathcal{G}(n, \beta)$, $\mathcal{H}(n, \gamma_1)$ and $\mathcal{H}(n, \gamma_2)$ such that $H_1 \subseteq G \subseteq H_2$ and $\mathbb{E}(\bar{d}(H_2)) \in \mathcal{O}(\mathbb{E}(\bar{d}(H_1)))$. The use of $\mathcal{O}$-notation here means that there exists a universal constant c independent of n and β such that $\mathbb{E}(\bar{d}(H_2)) \leq c \cdot \mathbb{E}(\bar{d}(H_1))$. In less formal words, this means that for every random graph G we find a sub- and a supergraph H_1 and H_2 such that the expected average degrees of H_1 and H_2 differ only by a constant factor.

We use the following random graph models for dot product and ball graphs. For the definitions let $d \in \mathbb{N}^+$.

Random Ball Graphs For a distribution of radii $\mathcal{R}$ on $[0, \frac{\pi}{2}]$ and a d-dimensional hypercube $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$, we define the d-dimensional *random ball graph model* $\mathcal{RBG}_d(n, \mathcal{R}, K)$ as the following probability distribution on graphs with $n \in \mathbb{N}^+$ vertices.

Let V be a vertex set with $|V| = n$. First for each vertex $v \in V$ positions P_v are drawn uniformly and independently in K and radii R_v are drawn independently from the distribution $\mathcal{R}$. This means we have independent random variables $\{P_v, R_v\}_{v \in V}$ with $P_v \sim \Phi(\mathcal{U}(K))$ and $R_v \sim \mathcal{R}$. In the resulting graph G an edge uv is present if and only if $\|P_u - P_v\|_2 \leq R_u + R_v$, i.e., G has the ball embedding $(P_\bullet, R_\bullet)$. Note that the restrictions on $\mathcal{R}$ and K are done to make the comparison between random ball graphs and random disk graphs possible. In this model this is however only a minor restriction, that says that radii and positions are bounded by some constant. For any distribution of radii that is bounded and any hypercube K we can scale both with the same factor such that they lie in the desired range.

Random Dot Product Graphs For a distribution of lengths $\mathcal{L}$ on $\mathbb{R}_{\geq 1}$ and a d-dimensional hypercube $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ like above, we define the $(d+1)$-dimensional *random dot product graph model* $\mathcal{RDPG}_{d+1}(n, \mathcal{L}, K)$ as the following probability distribution on graphs with n vertices.

Let V be a vertex set with $|V| = n$ like above. For each vertex $v \in V$ a polar coordinate vector φ_v is drawn from K uniformly and independently at random. Set $P_v = \Phi_{d+1}(\varphi_v)$, where $\Phi_{d+1} \colon \mathbb{R}^d \to S^d$ is the map for the $(d+1)$-dimensional polar coordinates. The lengths L_v are drawn independently from the length distribution $\mathcal{L}$. This means we have independent random variables $\{P_v, L_v\}_{v \in V}$ with $P_v \sim \Phi(\mathcal{U}(K))$ and $L_v \sim \mathcal{L}$. In the resulting graph G an edge uv is present if and only if $\langle L_u P_u, L_v P_v \rangle \geq 1$, i.e., G is the graph with dot product embedding $(P_\bullet, L_\bullet)$. Like above, the restrictions on $\mathcal{L}$ and K are needed to make the comparison between the two random graph models possible.

Choosing positions only from the small section $\Phi(K)$ of S^d instead of the whole sphere does not change the distribution too much, one mostly needs just one more dimension. This follows from a theorem from Fiduccia et al., who showed that for $\varepsilon > 0$, every $(d+1)$-dimensional dot product graph has a dot product embedding in dimension $d+2$, where every vertex is positioned within distance ε of a fixed unit vector [FSTZ98]. When looking into this proof one can see that, if the length of a vertex in the $(d+1)$-dimensional embedding was greater than 1, it also is in the $(d+2)$-dimensional embedding.

3 Comparing Dot Product and Weighted Embeddings

We start with comparing dot product and weighted embeddings because this connection is of most interest.

At first glance it seems that for a vertex v, its weight w_v and the length ℓ_v have similar effects on the degree of the vertex. In both cases a higher weight or length corresponds to a higher degree, because increasing these values decreases the weighted distance and increases the dot-product, making connections more likely. This however is not the case, as the following observation and theorem show.

A first observation one can make is that both embedding systems admit some artifacts in edge cases that do not match up. In a weighted embedding, vertices with the same position are always connected regardless of their weight. In a dot product embedding vertices with the same position, but length smaller than 1 are never connected. Thus, in the following we restrict ourselves to dot product embeddings where all lengths are greater than 1. Another thing that differs is, that in weighted embeddings, if one increases the weights, each pair of vertices can get connected. In dot product embeddings, vertices whose positions point in opposite directions are never connected, regardless of the length. If we restrict ourselves to nonnegative dot product embeddings[1], we reduce this problem, as the only thing that still produces these artifacts are orthogonal vectors.

The following theorem shows that even if we restrict ourselves to embeddings that do not have these artifacts (nonnegative dot product embeddings with length at least 1), length and weight have different effects on the resulting graph. Informally, it says that there is no function that transforms lengths into weights for all graphs without changing the positions. Note that in this case all positions lie in on a sphere.

Theorem 3.1: *For each function $f\colon \mathbb{R}_{\geq 1} \to \mathbb{R}_0^+$ and each dimension $d \geq 1$ there exists a graph G with a nonnegative d-dimensional dot product embedding $(p_\bullet, \ell_\bullet)$ with lengths $\ell_v \geq 1$ such that $(p_v, f(\ell_v))$ is not a d-dimensional weighted embedding of G.*

Proof. Let $f\colon \mathbb{R}_{\geq 1} \to \mathbb{R}_0^+$ be a function and $d \geq 2$. We will now construct a d-dimensional dot product embedding $(p_\bullet, \ell_\bullet)$ of a graph G such that $(p_\bullet, f(\ell_\bullet))$ is not a weighted embedding of G. First consider an arbitrary graph G together with a nonnegative dot product representation $(p_\bullet, \ell_\bullet)$ where $\ell_v \geq 1$. Let $u,v \in V$. Note that using that $\|p_u\| = \|p_v\| = 1$ the following holds, by definition of the norm:

$$\|p_u - p_v\|^2 = \langle p_u - p_v, p_u - p_v\rangle = \|p_u\|^2 - 2\langle p_u, p_v\rangle + \|p_v\|^2 = 2 - 2\langle p_u, p_v\rangle$$

[1]Embeddings where the coordinates of all positions are nonnegative

Using this, we get the following criterion for an edge between u and v, expressed through the distance $\|p_u - p_v\|$ instead of the dot product.

$$\begin{aligned} uv \in E &\Leftrightarrow \langle \ell_u \cdot p_u, \ell_v \cdot p_v \rangle \geq 1 \\ &\Leftrightarrow \ell_u \ell_v \langle p_u, p_v \rangle \geq 1 \\ &\Leftrightarrow \ell_u \ell_v \frac{2 - \|p_u - p_v\|^2}{2} \geq 1 \\ &\Leftrightarrow \|p_u - p_v\|_2^2 \leq 2 - \frac{2}{\ell_u \ell_v} \\ &\Leftrightarrow \|p_u - p_v\|_2 \leq \sqrt{2 - \frac{2}{\ell_u \ell_v}} =: F(\ell_u, \ell_v) \end{aligned}$$

Note that the equivalences (except the first one) also hold for equalities and not only for the inequalities.

In weighted embeddings there is an edge between $u, v \in E$ if and only if $\|p_u - p_v\| \leq w_u w_v$. Now, we have a criterion for each embedding that depends on the distance of the position. Comparing the two criteria, we observe that if our function f fulfills $f(\ell_u) \cdot f(\ell_v) = F(\ell_u, \ell_v)$, then f would be a valid translation from dot product to weighted embeddings.

As we want to construct nonnegative dot product embeddings we note that there exists a constant c_d such that for every $x, y \in [1, c_d)$ there exist nonnegative vectors[2] $p_x, p_y \in S^d$ such that $\|p_x - p_y\|_2 = F(x, y)$. More precisely, we can set c_d to the solution of the equation $F(x, x) = \sqrt{d}$ as for $d \geq 2$ there exist nonnegative $p_x, p_y \in S^d$ with $\|p_x - p_y\|_2 = \sqrt{d}$.

The rest of the proof proceeds as follows: We first show that for every function that does not fulfill $f(x) \cdot f(y) = F(x, y)$ for all $x, y \in [1, c_d)$ there exists a graph G with a d-dimensional dot product embedding $(p_\bullet, \ell_\bullet)$ such that $(p_\bullet, f(\ell_\bullet))$ is not a d-dimensional weighted embedding of G. Then we show that there exists no function, that fulfills the upper equality. Together this proves the theorem.

First assume there are $x, y \in [1, c_d)$ such that $f(x)f(y) < F(x, y)$. Choose nonnegative $p_x, p_y \in S^{d-1}$ such that $f(x)f(y) < \|p_x - p_y\| < F(x, y)$. Then those two points form an edge for the dot product but not for the weighted case.

Second assume there are $x, y \in [1, c_d)$ such that $f(x)f(y) > F(x, y)$. Analogous to the other case choose nonnegative $p_x, p_y \in S^{d-1}$ such that $f(x)f(y) > \|p_x - p_y\| > F(x, y)$. Then those two points form an edge for the weighted but not for the dot product case.

Now assume that f has to fulfill

$$f(x) \cdot f(y) = \sqrt{2 - \frac{2}{x \cdot y}} = \sqrt{2 - \frac{2}{1 \cdot xy}} = f(1) \cdot f(xy)$$

for all $x, y \in [1, c_d)$.

Setting x and y to 1 yields $f(1) \cdot f(1) = \sqrt{2 - 2/(1 \cdot 1)} = 0$, i.e., $f(1) = 0$. Now using the outer part of the observation yields $f(x) \cdot f(x) = f(1) \cdot f(x^2) = 0$ for all $x \in [1, c_d)$, i.e., $f \equiv 0$. This results in a contradiction as for $x > 1$ it holds that $f(x) \cdot f(x) = \sqrt{2 - \frac{2}{x^2}} > \sqrt{2 - \frac{2}{1^2}} = 0$. ∎

However, there is one special case where a translation from dot product graphs to weighted graphs works, namely the case of d being 1.

[2] Vectors where every coordinate is nonnegative.

The following lemma and corollary show that every 1-dimensional dot product graph has a 1-dimensional weighted embedding. The proof heavily uses that the dot product in one dimension is just the product of two numbers, like the weights in a weighted embedding.

Lemma 3.2: *Let G be a 1-dimensional dot product graph with nonnegative embedding $(p_\bullet, \ell_\bullet)$, i.e. $p_v = 1$ for all $v \in V$. Then G has a 1-dimensional weighted embedding with evenly spaced positions $\tilde{p}_v \in [0, 1]$ and weights $w_v = \ell_v^k$ for some integer $k \in \mathbb{N}^+$.*

Proof. Let G and its embedding $(p_\bullet, \ell_\bullet)$ be like in the claim. Let n be the number of vertices of G and identify $V = \{1, \dots, n\}$.

We construct the weighted embedding $(\tilde{p}_\bullet, w_\bullet)$ as follows. For $i \in V$ we set $\tilde{p}_i = \frac{i}{n}$, now comparing distances this yields

$$\frac{1}{n} \leq \|\tilde{p}_i - \tilde{p}_j\|_2 \leq 1 \text{ for all } i, j \in V.$$

Let $i, j \in V$. If we could guarantee

$$w_i w_j \geq 1 \text{ for } ij \in E \text{ and} \tag{3.1}$$

$$w_i w_j < \frac{1}{n} \text{ for } ij \notin E \tag{3.2}$$

then $(\tilde{p}_\bullet, w_\bullet)$ would be a valid weighted embedding. We now construct weights w_i such that the upper two properties are fulfilled, thus showing the claim. Given that $(p_\bullet, \ell_\bullet)$ is a valid dot product embedding with $p_i = 1$ for all i, we already know that

$$\ell_i \ell_j \geq 1 \text{ for } ij \in E \text{ and}$$
$$\ell_i \ell_j < 1 \text{ for } ij \notin E.$$

Taking the kth power for $k \in \mathbb{N}^+$ this yields.

$$\ell_i^k \ell_j^k \geq 1 \text{ for } ij \in E \text{ and}$$
$$\ell_i^k \ell_j^k < 1 \text{ for } ij \notin E.$$

Moreover, if i and j are not connected the sequence $(\ell_i^k \ell_j^k)_k$ is converging to zero. In particular there exists a $k_0 \in \mathbb{N}^+$ such that $\ell_i \ell_j < \frac{1}{n}$ for all $i, j \in V$ with $ij \notin E$. Applying this to the upper inequalities yields

$$\ell_i^{k_0} \ell_j^{k_0} \geq 1 \text{ for } ij \in E \text{ and}$$
$$\ell_i^{k_0} \ell_j^{k_0} < \frac{1}{n} \text{ for } ij \notin E.$$

In particular setting $w_i = \ell_i^{k_0}$ fulfills the Equations (3.1) and (3.2), providing the desired weights for the embedding.

∎

Using this result for nonnegative dot-product embeddings, we can show that every 1-dimensional dot product embedding can be transformed in a 1-dimensional weighted embedding. In a general 1-dimensional dot product embedding, the subgraphs induced by all vertices with position 1 and -1 are not connected. Thus, we can embed them separately using the lemma above and place them far enough away in 1-dimensional weighted space and get a valid weighted embedding.

Corollary 3.3: *Every* 1*-dot product graph has a* 1*-dimensional weighted embedding.*

Proof. Let G be a 1-dimensional dot product graph with embedding $(p_\bullet, r_\bullet)$. We define subsets $V_+, V_- \subseteq V$ with $V_\pm = \{v \in V \mid p_v = \pm 1\}$. Then G is the disjoint union of the induced subgraphs $G[V_+]$ and $G[V_-]$ as the dot product between vertices in different sets is negative.

Both $G[V_+]$ and $G[V_-]$ admit a nonnegative dot product embedding, namely the one inherited from G. In the case of $G[V_-]$ we set the position to 1. By Lemma 3.2 both graphs have a 1-dimensional weighted embedding with positions in $[0, 1]$. Let $(p_\bullet^\pm, w_\bullet^\pm)$ be the weighted embedding of $G[V_\pm]$ as constructed in. We set $W = 1 + \max\{w_u^- w_v^+ \mid u \in V_-, v \in V_+\}$. Then we construct the 1-dimensional weighted embedding of G as follows. For $v \in V_-$ we set $\tilde{p}_v = p_v^-$ and $\tilde{w}_v = w_v^-$. For $v \in V_+$ we set $\tilde{p}_v = p_v^+ + W$ and $\tilde{w}_v = w_v^+$. This is also an embedding of $G[V_+]$ as translating all vertices by a constant does not change the difference in their positions. For vertices $u \in V_-$ and $v \in V_+$ it holds that $\|\tilde{p}_u - \tilde{p}_v\|_2 \geq W > \tilde{w}_u \tilde{w}_v$ by construction of W.

Thus, $(\tilde{p}_\bullet, \tilde{w}_\bullet)$ is a valid 1-dimensional weighted embedding of G. ∎

4 Comparing Dot Product and Ball Graphs

Having seen that weighted embeddings are not a good model to compare dot product graphs to, we now turn our attention to ball graphs. Here, there are already a few indicators that the models are more similar, especially when restricting the lengths in the dot product embedding to $\mathbb{R}_{\geq 1}$. This fact will become clear throughout the chapter.

Let us first consider ball graphs with fixed size balls, i.e. unit ball graphs, and dot product graphs with fixed lengths in $\mathbb{R}_{\geq 1}$. The latter we will call spherical graphs due to the fact that all vertices lie on a sphere of a fixed radius $r \geq 1$. A different definition of spherical graphs is to say that all positions lie on the unit sphere and vertices are connected if and only if their dot product is larger than some threshold $t = \frac{1}{r^2}$ that is not necessarily 1 as before, but can be smaller. This is equivalent to stating that vertices are connected if and only if their angular distance on the sphere is smaller than the threshold $\arccos(t)$. As the surface of a sphere is similar to a plane, especially if one only looks at a small section, one can now translate unit ball graphs to spherical graphs in one dimension higher, by just shrinking the whole embedding and mapping it onto a tiny section of the sphere. In other words, every d-dimensional unit ball graph is a $(d + 1)$-dimensional dot product graph where all vertices have the same lengths. A formal proof of this fact can be found in [RRŠ89].

The hope is that if we now consider different-sized balls, we can choose different lengths such that we can transform every d-dimensional ball graph into a $(d + 1)$-dimensional dot product graph. Sadly, this is not possible, as the following example shows. Every planar graph is a coin graph, i.e., a 2-dimensional ball graph were all intersections between disks are just points [Sac94]. If there were an exact translation from disk graphs to 3-dimensional dot product graphs, then every planar graph would have a 3-dimensional dot product embedding. This however is a wrong statement, Kang and Müller have shown that there exist planar graphs that have no 3-dimensional dot product embedding [KM11a].

Nevertheless, there is still some similarity between the models. We construct translations from d-dimensional ball graphs to $(d + 1)$-dimensional dot product graphs that provide us with a sub- and a supergraph of the original ball graph. These translations will be constructed in a somewhat natural way. The hope is that the resulting sub- and supergraph are not that different, i.e., their average degree differs only by a constant factor that depends on the dimension d but nothing else. We will see, that there are extreme cases where this is not the case. The claim however is true if one considers random ball graphs. We will show that if we draw a ball graph randomly, the expected average degree of the corresponding sub- and supergraph will differ only by a constant factor. Moreover, the translations can be inverted to turn a dot product embedding of a graph into ball embeddings of a sub- and a supergraph. When considering the respective random graph models, the statement about the expected average degree will hold as well.

We first start with constructing the translations from ball graphs to dot product graphs. On our way to this goal we consider an intermediate embedding that is very similar to the ball embedding, but which is easier to compare to dot product embeddings. This intermediate embedding will be the cap graphs, which are the intersection graphs of caps of a sphere. As we have already argued when discussing the similarity between unit ball and spherical graphs,

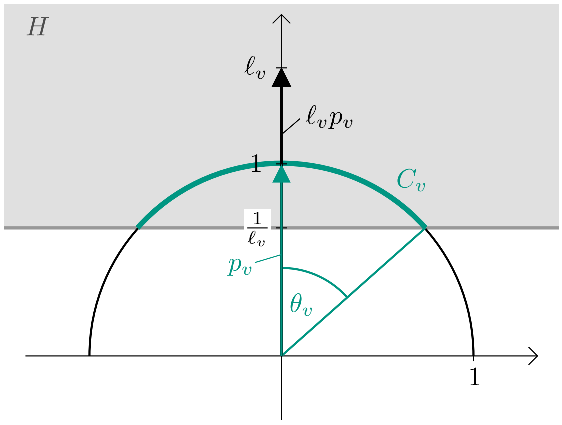

Figure 4.1: Stylized translation between dot product and cap embeddings. For a vertex v and corresponding vector $\ell_v p_v$ the half space H is the subset of all possible positions where vertices would get connected to v. The intersection of H with the unit circle is a cap C_v with center p_v and angular radius $\theta_v = \arccos(\frac{1}{\ell_v})$.

if one considers a tiny section of the sphere S^d, caps in this section look almost like d-balls and a direct translation is possible via stereographic projection. Thus, we can now try to transform cap graphs into dot product graphs.

For a vertex v in a dot product graph with $\ell_v \geq 1$, the intersection of S^d and all possible positions where vertices would be connected to v is a cap with angular radius $\arccos(\frac{1}{\ell_v})$ and center p_v. This fact is illustrated in Figure 4.1. If we look at it the other way around and start with a cap with center p_v and angular radius θ_v, this cap corresponds to the vector $\cos^{-1}(\theta_v)p_v$ in a dot product embedding, where $\cos^{-1}(x) = 1/\cos(x)$[1]. This makes a good candidate for translating caps to positions and lengths. Note that in this construction the lengths will always be at least 1.

However, there is a slight caveat. In a cap embedding, caps that intersect somewhere form an edge. In the translation that we have constructed, we identified the cap as all possible positions on S^d such that other vertices with this position would be connected. In the language of caps, this translates to the fact, that for two connected vertices the center of one cap has to be contained in the other. This is a lot more restrictive than only demanding that the caps intersect. This results in fewer edges in the corresponding dot product graph, in fact the dot product graph is a subgraph of the cap graph. The precise geometrical argument is more nuanced and will be done down below.

A way to ensure that all edges of the cap graph are also present in the dot product graph is to double all angular radii and then translate them, i.e., $\ell_v = \cos^{-1}(2\theta_v)$. This enforces that for intersecting caps, one midpoint is contained in the other cap. Thus, in this case, the dot product graph is a super graph of cap graph.

This construction, that provides us with the desired translations from cap graphs to dot product graphs, is summarized in the following theorem.

Theorem 4.1: *Let $(p_\bullet, \theta_\bullet)$ be a d-dimensional cap embedding of a graph G with $\theta_v \in [0, \frac{\pi}{2}]$ for each $v \in V$. Then the $(d+1)$-dimensional dot product graphs H_1 with dot product embedding $(p_\bullet, \cos^{-1}(\theta_\bullet))$ and H_2 with dot product embedding $(p_\bullet, \cos^{-1}(2\theta_\bullet))$ are a sub- and a supergraph of G.*

[1]The inverse function of cosine is arccos.

Proof. We show the two claims separately, starting with the one for the supergraph

Let $uv \in E$ be an edge in G, i.e., $\arccos(\langle p_u, p_v \rangle) \leq \theta_u + \theta_v$. Applying cosine yields $\langle p_u, p_v \rangle \geq \cos(\theta_u + \theta_v)$ as cosine is monotonically decreasing in $[0, \pi]$. We now have to show, that uv is also an edge in H_2, i.e.

$$\langle p_u, p_v \rangle \geq \frac{1}{\ell_u \ell_v} = \cos(2\theta_u)\cos(2\theta_v).$$

We calculate

$$\langle p_u, p_v \rangle \geq \cos(\theta_u + \theta_v) \geq \min\{\cos(2\theta_u), \cos(2\theta_v)\},$$

where the second inequality follows from $\theta_u + \theta_v \leq 2\max\{\cos(\theta_u), \cos(\theta_v)\}$ using again that cosine is monotonically decreasing in $[0, \pi]$. Then we see that

$$\min\{\cos(2\theta_u), \cos(2\theta_v)\} \geq \cos(2\theta_u)\cos(2\theta_v)$$

as $\cos(\theta) \leq 1$ for all θ, yielding the desired inequality.

Let u,v be disconnected in G, i.e., $\arccos(\langle p_u, p_v \rangle) > \theta_u + \theta_v$. Applying cosine like above yields $\langle p_u, p_v \rangle < \cos(\theta_u + \theta_v)$. We now have to show that they are disconnected in H_1 as well, i.e. $\langle p_u, p_v \rangle < \frac{1}{\ell_u \ell_v} = \cos(\theta_u)\cos(\theta_v)$. We calculate

$$\langle p_u, p_v \rangle < \cos(\theta_u + \theta_v) = \cos(\theta_u)\cos(\theta_v) - \sin(\theta_u)\sin(\theta_v),$$

where the equality is precisely the angle addition formula for cosine. Using that $\sin(\theta) \geq 0$ for $\theta \in [0, \frac{\pi}{2}]$ we get

$$\cos(\theta_u)\cos(\theta_v) - \sin(\theta_u)\sin(\theta_v) \leq \cos(\theta_u)\cos(\theta_v),$$

yielding the desired inequality. ∎

Now that we have established the connection between dot product graphs and cap graphs, we only need a translation from ball graphs to cap graphs. One could use stereographic projection, however this would make the subsequent probabilistic arguments very difficult as the radii and centers of the balls and the caps will differ. Before we do this different kind of translation we first use stereographic projection as a translation to argue why we need a probabilistic approach in the first place.

As already said before, stereographic projection[2] turns ball embeddings into cap embeddings as it sends d-balls to caps in S^d. If we start with a ball embedding where balls intersect only in one point[3], use stereographic projection to get a cap embedding and look at the corresponding dot product subgraph, this subgraph has no edges. This follows from the proof of Theorem 4.1 as for every edge uv in the ball and cap graph it holds that $\arccos(\langle p_u, p_v \rangle) = \theta_u + \theta_v$. Applying cosine yields $\langle p_u, p_v \rangle = \cos(\theta_u + \theta_v) = \cos(\theta_u)\cos(\theta_v) - \sin(\theta_u)\sin(\theta_v)$, which is strictly smaller than $\cos(\theta_u)\cos(\theta_v)$ if $\theta_u, \theta_v \in (0, \frac{\pi}{2})$. This means that u and v are not connected in the corresponding dot product subgraph. This means that the average degree of the subgraph is 0, whereas the average degree of the supergraph is at least as large as the one from the original ball graph, which can be arbitrarily high. Thus, the average degrees differ not by a constant factor.

This is where random graph theory will help us, as these edge cases seldom show up when drawing positions uniformly at random and independently of the radii.

[2] Or rather its inverse.

[3] For $d = 2$ this gives the coin graphs from above.

4.1 Turning Ball Graphs into Cap Graphs

As already said above, stereographic projection is not a good choice for our translation from ball to cap graphs. In this section, we will develop and prove a different kind of translation, that works better with random graph theory.

Like in the case of cap and dot product embeddings we will translate a ball embedding into cap embeddings of a sub- and a supergraph. To achieve this, we use the positions of the balls p_v as polar coordinates for the positions of the caps. As a reminder, the map for $(d+1)$-dimensional polar coordinates Φ_{d+1} takes a coordinate vector $p \in [-\frac{\pi}{2}, \frac{\pi}{2}]^d$ to $\Phi_{d+1}(p) = \begin{pmatrix}\cos(p_d) \cdot \Phi_d(p_{[d-1]}) & \sin(p_d)\end{pmatrix}$ where $\Phi_1 = 1$.

The following theorem states the translations from ball graphs to cap graphs. These translations use the positions in the ball embedding as angles for polar coordinates and scale the radii such that we get cap embeddings of a sub- and a supergraph. For this to work, it uses that the angular distance dist_S on the sphere S^d is similar to the distance in the Euclidean space $\mathbb{R}^d$, especially when looking only at small subspaces. Thus, we restrict the positions of our ball embedding to small hypercubes $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$. The smaller K is, the more similar the cap embeddings of the sub- and supergraph are. In the theorem, the effect of the size of K is measured by a factor $\alpha_K > 0$. This factor approaches 1, when restricting to smaller hypercubes.

However, we note that the respective cap embeddings do not converge to each other, in the sense that the radii of the caps of the subgraph do not converge to the radii of the caps of the supergraph[4]. This is due to the fact that we multiply the radii of the supergraph with $\frac{\alpha_K}{\sqrt{d}}$ to get the radii of the subgraph and not only with α_K. We suspect that it is possible to leave out the factor $\sqrt{d}$[5], such that the cap embeddings converge to each other. However, we were not able to prove this.

Theorem 4.2: *Let $(p_\bullet, r_\bullet)$ be a d-dimensional ball embedding of a graph G in a hypercube $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$. Then the graphs H_1 with cap embedding $\left(\Phi(p_\bullet), \frac{\alpha_K}{\sqrt{d}} r_\bullet\right)$ and H_2 with cap embedding $\left(\Phi(p_\bullet), r_\bullet\right)$ are a sub- and a supergraph of G where $\alpha_K = \min_{p,\tilde{p} \in K} \prod_{i=2}^{d} \cos(p_i)\cos(\tilde{p}_i) > 0$.*

To prove this theorem we need the following lemma that compares the Euclidean distance $\|p - \tilde{p}\|_2$ of two points $p, \tilde{p} \in \mathbb{R}^d$ and the angular distance of the respective polar coordinates $\mathrm{dist}_S(\Phi(p), \Phi(\tilde{p}))$ on S^d. More precisely, it shows that the angular distance $\mathrm{dist}_S(\Phi(p), \Phi(\tilde{p}))$ is bounded by the Euclidean distance from above and bounded below by the maximum norm multiplied by a constant. This constant $\alpha_{p,\tilde{p}} < 1$ depends on the vectors $p, \tilde{p}$ and approaches 1 when moving the vectors closer together. As the maximum norm differs from the Euclidean norm only by a factor $\sqrt{d}$, we thus get a lower bound of $\mathrm{dist}_S(\Phi(p), \Phi(\tilde{p}))$ by the Euclidean distance of the respective vectors multiplied by a constant. This constant approaches $\frac{1}{\sqrt{d}}$, when moving the vectors closer together.

If we look at pictures of the case $d = 2$, it seems that this constant approaches even 1, i.e., for vectors lying close together, the Euclidean lower bound of dist_S converges to the upper bound. However, we were not able to prove this. This is the reason, why the cap embeddings do not converge to each other, as the factor α_K used in Theorem 4.2 is precisely the minimum $\alpha_{p,\tilde{p}}$ for vectors $p, \tilde{p}$ in a hypercube $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$.

[4] The positions however, are the same in the sub- and supergraph.

[5] Potentially with a different definition of α_K.

Lemma 4.3: *Let* $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$. *Then*

$$\alpha_{p,\tilde{p}} \|p - \tilde{p}\|_\infty \leq \mathrm{dist}_S(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) \leq \|p - \tilde{p}\|_2$$

where $\alpha_{p,\tilde{p}} = \prod_{i=2}^d \cos(p_i)\cos(\tilde{p}_i)$.

For the proof of the upper bound we need the following auxiliary lemma, which is a relation of arccos to the Euclidean norm.

Lemma 4.4: *Let* $d \in \mathbb{N}^+$. *For* $x \in \mathbb{R}^d$ *it holds that* $\arccos(\prod_{i=1}^d \cos(x_i)) \leq \|x\|_2$.

We first prove Lemma 4.3 using this result.

Proof of Lemma 4.3. We show both inequalities separately by induction.
We start with showing the upper bound i.e.,

$$\mathrm{dist}_S(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) \leq \|p - \tilde{p}\|_2 \text{ for } p, \tilde{p} \in \Big[-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}\Big]^d.$$

First we prove this claim for p and $\tilde{p}$, where p_i and $\tilde{p}_i$ have the same sign for each $i \in \{1, \ldots, d\}$ (zero matches both signs). We do this by showing inductively, that

$$\langle \Phi(p)_{d+1}, \Phi_{d+1}(\tilde{p}) \rangle \geq \prod_{i=1}^d \cos(p_i - \tilde{p}_i) =: P_d. \tag{4.1}$$

Using that arccos is monotonically decreasing, applying arccos and then Lemma 4.4 this implies the claimed upper bound. Then we generalize the claim of the upper bound to general $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$.

Now onto the inductive proof of Equation (4.1). For $d = 1$ let $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]$ such that they have the same sign[6]. By the angle addition formula for cosine it holds that $\langle \Phi_2(p), \Phi_2(\tilde{p}) \rangle = \cos(p)\cos(\tilde{p}) + \sin(p)\sin(\tilde{p}) = \cos(p - \tilde{p}) = P_1$.

For higher d, let $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ such that p_i and $\tilde{p}_i$ have the same sign for each $i \in \{1, \ldots, d\}$. Using the induction hypothesis (IH) we get

$$\begin{aligned}
\langle \Phi_{d+1}(p), \Phi_{d+1}(\tilde{p}) \rangle &= \left\langle \begin{pmatrix} \cos(p_d) \cdot \Phi_d(p_{[d-1]}) \\ \sin(p_d) \end{pmatrix}, \begin{pmatrix} \cos(\tilde{p}_d) \cdot \Phi_d(\tilde{p}_{[d-1]}) \\ \sin(\tilde{p}_d) \end{pmatrix} \right\rangle \\
&= \cos(p_d)\cos(\tilde{p}_d) \cdot \big\langle \Phi_d(p_{[d-1]}), \Phi_d(\tilde{p}_{[d-1]}) \big\rangle \ + \ \sin(p_d)\sin(\tilde{p}_d) \\
&\overset{\mathrm{IH}}{\geq} \cos(p_d)\cos(\tilde{p}_d) P_{d-1} \ + \ \sin(p_d)\sin(\tilde{p}_d) \\
&= \big(\cos(p_d)\cos(\tilde{p}_d) + \sin(p_d)\sin(\tilde{p}_d)\big) P_{d-1} \ + \ \sin(p_d)\sin(\tilde{p}_d)\,(1 - P_{d-1}) \\
&\geq (\cos(p_d - \tilde{p}_d) P_{d-1} = P_d,
\end{aligned}$$

where the last inequality holds by the angle addition formula for cosine and that the second summand is positive by our assumptions. To be precise $1 - P_{d-1}$ is always positive as $P_{d-1} \leq 1$ and $\sin(p_d)\sin(\tilde{p}_d)$ is positive because of the sign assumption. This concludes the inductive part of the proof.

[6] The sign assumption is not needed in dimension 1.

For general $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ without the sign assumption, we split the straight line from p to $\tilde{p}$ into (at most $d+1$) segments where the sign assumption holds. The splitting points $q^{(1)}, \dots, q^{(r)} \in \mathbb{R}^d$ with $r \leq d$ are the intersection points of the straight line with the coordinate hyperplanes. Set $q^{(0)} := p$, $q^{(r+1)} := \tilde{p}$. Applying the triangle inequality of dist_S to $\mathrm{dist}_S(\Phi(p), \Phi(\tilde{p}))$ yields

$$\mathrm{dist}_S(\Phi(p), \Phi(\tilde{p})) = \mathrm{dist}_S(\Phi(q^{(0)}), \Phi(q^{(r+1)})) \leq \sum_{i=0}^{r} \mathrm{dist}_S(\Phi(q^{(i)}), \Phi(q^{(i+1)})).$$

As the sign assumption holds for endpoints of segments $q^{(i)}, q^{(i+1)}$, the upper bound on dist_S holds by the proof above, and we get $\mathrm{dist}_S(\Phi(q^{(i)}), \Phi(q^{(i+1)})) \leq \|q^{(i)} - q^{(i+1)}\|_2$. Applying this to each summand in the upper equation we get

$$\sum_{i=0}^{r} \mathrm{dist}_S(\Phi(q^{(i)}), \Phi(q^{(i+1)})) \leq \sum_{i=0}^{r} \|q^{(i)}, q^{(i+1)}\|_2 = \|q^{(0)} - q^{(r+1)}\|,$$

where the last equation holds because all $q^{(i)}$ lie on a straight line. Thus, we get the inductive claim $\mathrm{dist}_S(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) \leq \|p - \tilde{p}\|_2$ for all $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$.

Now, we prove the lower bound

$$\alpha_{p,\tilde{p}} \|p - \tilde{p}\|_\infty \leq \mathrm{dist}(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) \text{ for } p, \tilde{p} \in \left[-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}\right]^d$$

by induction on the dimension d.

For $d = 1$ and $p, \tilde{p} \in [-\frac{\pi}{4}, \frac{\pi}{4}]$ we have

$$\begin{aligned}\mathrm{dist}_S(\Phi_2(p), \Phi_2(\tilde{p})) &= \arccos(\langle \Phi_2(p), \Phi_2(\tilde{p})\rangle) = \arccos(\cos(p)\cos(\tilde{p}) + \sin(p)\sin(\tilde{p})) \\ &\overset{(\star)}{=} \arccos(\cos(|p - \tilde{p}|)) = |p - \tilde{p}| = \|p - \tilde{p}\|_\infty,\end{aligned}$$

where $(\star)$ holds as cosine is an even function.

For higher d, let $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$.

We first show that $\mathrm{dist}_S(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) \geq \alpha_{p,\tilde{p}} |p_d - \tilde{p}_d|$. This statement follows from applying arccos to the following inequality

$$\begin{aligned}\langle \Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})\rangle &= \cos(p_d)\cos(\tilde{p}_d) \cdot \langle \Phi_d(p_{[d-1]}), \Phi_d(\tilde{p}_{[d-1]})\rangle + \sin(p_d)\sin(\tilde{p}_d) \\ &\leq \cos(p_d)\cos(\tilde{p}_d) + \sin(p_d)\sin(\tilde{p}_d) = \cos(|p_d - \tilde{p}_d|),\end{aligned}$$

and then using that $\alpha_{p,\tilde{p}} \leq 1$. For the inequality, one uses that cosine is nonnegative in $[-\frac{\pi}{2}, \frac{\pi}{2}]$. The last equality holds because cosine is an even function.

Second we show that

$$\mathrm{dist}(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) \geq \alpha_{p,\tilde{p}} \max_{i \in 1 \dots d-1} |p_i - \tilde{p}_i|. \tag{4.2}$$

This will conclude the proof of the induction step. For this we first note that

$$\langle \Phi_{k+1}(p_{[k]}), \Phi_{k+1}(\tilde{p}_{[k]})\rangle \geq 0 \text{ for } i \in \{1, \dots, d\}. \tag{4.3}$$

This follows from the fact that

$$\arccos(\langle \Phi_{k+1}(p_{[k]}), \Phi_{k+1}(\tilde{p}_{[k]})\rangle) \leq \|p_{[k]} - \tilde{p}_{[k]}\|_2 \leq \sqrt{k} \|p_{[k]} - \tilde{p}_{[k]}\|_\infty \leq \sqrt{k} \frac{\pi}{2\sqrt{d}} \leq \frac{\pi}{2},$$

where the first inequality is just the upper bound for dist_S proven above, and the second is the comparison between the Euclidean norm $\|\cdot\|_2$ and the maximum norm $\|\cdot\|_\infty$. The third inequality holds, because $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ by assumption. This step is where the range assumption in this lemma comes from.

Now onto showing Equation (4.2). For this we make a case distinction on the signs of p_d and $\tilde{p}_d$.

If p_d and $\tilde{p}_d$ have a different sign then $\sin(p_d)\sin(\tilde{p}_d) \leq 0$ and thus, the following chain of inequalities holds:

$$\begin{aligned}\langle \Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})\rangle &= \cos(p_d)\cos(\tilde{p}_d) \cdot \langle \Phi_d(p_{[d-1]}), \Phi_d(\tilde{p}_{[d-1]})\rangle + \sin(p_d)\sin(\tilde{p}_d) \\ &\leq \langle \Phi_d(p_{[d-1]}), \Phi_d(\tilde{p}_{[d-1]})\rangle.\end{aligned}$$

The last inequality holds because of Equation (4.3). Applying arccos and using the induction hypothesis (IH) yields

$$\begin{aligned}\mathrm{dist}_S(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) &\geq \mathrm{dist}_S(\Phi_d(p_{[d-1]}), \Phi_d(\tilde{p}_{[d-1]})) \\ &\overset{IH}{\geq} \alpha_{p_{[d-1]},\tilde{p}_{[d-1]}} \max_{i\in 1\ldots d-1} |p_i - \tilde{p}_i| \geq \alpha_{p,\tilde{p}} \max_{i\in 1\ldots d-1} |p_i - \tilde{p}_i|.\end{aligned}$$

The last inequality follows from the fact that $\alpha_{p,\tilde{p}} \leq \alpha_{p_{[d-1]},\tilde{p}_{[d-1]}}$, this directly follows from the definition of α.

Now let p_d and $\tilde{p}_d$ have the same sign. We first note that arccos is concave in $[0, 1]$, i.e., for $x, y, t \in [0, 1]$ it holds that $\arccos(tx + (1-t)y) \geq t\arccos(x) + (1-t)\arccos(y)$. We want to apply this to

$$\begin{aligned}\mathrm{dist}_S(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) &= \arccos(\Phi_{d+1}(p), \Phi_{d+1}(\tilde{p})) \\ &= \arccos(\cos(p_d)\cos(\tilde{p}_d)\Phi_d(p_{[d-1]}, \Phi_d(\tilde{p}_{[d-1]})) + \sin(p_d)\sin(\tilde{p}_d)).\end{aligned}$$

For this, we start with setting $x = \Phi_d(p_{[d-1]}, \Phi_d(\tilde{p}_{[d-1]}))$ and $t = \cos(p_d)\cos(\tilde{p}_d)$. Both lie in the correct range, as $x \geq 0$ by Equation (4.3) and t because cosine is nonnegative in $[-\frac{\pi}{2}, \frac{\pi}{2}] \supset [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]$. Looking at the condition for concavity we have to set $y = \frac{\sin(p_d)\sin(\tilde{p}_d)}{1-t}$. This lies in the correct range as $\sin(p_d)\sin(\tilde{p}_d)$ is positive by the assumption that p_d and $\tilde{p}_d$ have the same sign and $\sin(p_d)\sin(\tilde{p}_d) + t = \cos(p_d - \tilde{p}_d) \leq 1$. Thus, we can apply the concavity inequality and get

$$\begin{aligned}\arccos(tx + (1-t)y) &\geq t\arccos(x) + (1-t)\arccos(y) \geq t\arccos(x) + 0 \\ &= \cos(p_d)\cos(\tilde{p}_d)\,\mathrm{dist}_S(\Phi_d(p_{[d-1]}), \Phi_d(\tilde{p}_{[d-1]})) \\ &\overset{IH}{\geq} \cos(p_d)\cos(\tilde{p}_d)\alpha_{p_{[d-1]},\tilde{p}_{[d-1]}} \max_{i\in 1\ldots d-1} |p_i - \tilde{p}_i| = \alpha_{p,\tilde{p}} \max_{i\in 1\ldots d-1} |p_i - \tilde{p}_i|,\end{aligned}$$

by the induction hypothesis, which is what we wanted to show.

■

Then we proceed with proving the somewhat technical auxiliary lemma used in the proof above.

Proof of Lemma 4.4. Let $x \in [-\frac{\pi}{2}, \frac{\pi}{2}]^d$.

Define a function $f\colon [0,1] \to \mathbb{R}$ as $f(t) = \arccos(\prod_{i=1}^d \cos(x_i t))$. Note that f is differentiable on $(0,1)$. In particular, as $f(0) = 0$, we can write $f(1) = \int_0^1 f'(t)\mathrm{d}t$. In the rest of the proof we will show that $f'(t) \leq \|x\|_2$ for all $t \in (0,1]$. Using monotony of the integral this yields

$$\arccos\left(\prod_{i=1}^d \cos(x_i)\right) = f(1) = \int_0^1 f'(t)\mathrm{d}t \leq \int_0^1 \|x\|_2 \mathrm{d}t = \|x\|_2,$$

which is the claim we want to show.

Now for proving the desired inequality $f'(t) \leq \|x\|_2$ we first calculate f'. Using that $\arccos'(x) = -\frac{1}{\sqrt{1-x^2}}$, then by the product and chain rule we get

$$f'(t) = \frac{1}{\sqrt{1 - \prod_{i=1}^d \cos^2(x_i t)}} \cdot \sum_{i=1}^d x_i \sin(x_i t) \cdot \prod_{i \neq j} \cos(x_j t).$$

If $x = 0$, then $f'(t) = 0$ because f is constant, showing that in this case the inequality holds. Now assume $x \neq 0$. Applying the Cauchy-Schwarz inequality to the upper formula, this yields

$$|f'(t)|^2 \leq \|x\|^2 \cdot \frac{1}{1 - \prod_{i=1}^d \cos^2(x_i t)} \sum_{i=1}^d \sin^2(x_i t) \prod_{i \neq j} \cos^2(x_j t).$$

The only thing that is left to show is that the second factor is at most 1. We will prove this by induction on d. If $d = 1$ the factor reads $\frac{\sin^2(x_1 t)}{1-\cos^2(x_1 t)}$, which is 1 by the Pythagorean trigonometric identity. Now, let the claim hold for $d \in \mathbb{N}^+$. As $x \neq 0$, we get the following chain of equivalent statements:

$$\begin{aligned}
& & \frac{1}{1 - \prod_{i=1}^{d+1} \cos^2(x_i t)} \sum_{i=1}^{d+1} \sin^2(x_i t) \prod_{i \neq j} \cos^2(x_j t) \leq 1 \\
\Leftrightarrow & & \prod_{i=1}^{d+1} \cos^2(x_i t) + \sum_{i=1}^{d+1} \sin^2(x_i t) \prod_{i \neq j} \cos^2(x_j t) \leq 1 \\
\Leftrightarrow & \underbrace{(\cos^2(x_{d+1} t) + \sin^2(x_{d+1} t))}_{=1} & \prod_{i=1}^{d} \cos^2(x_i t) + \cos^2(x_{d+1} t) \sum_{i=1}^{d} \sin^2(x_i t) \prod_{i \neq j} \cos^2(x_j t) \leq 1 \\
\Leftrightarrow & & \prod_{i=1}^{d} \cos^2(x_i t) + \cos^2(x_{d+1} t) \sum_{i=1}^{d} \sin^2(x_i t) \prod_{i \neq j} \cos^2(x_j t) \leq 1
\end{aligned}$$

The last statement is true because of the induction hypothesis, using that $\cos^2(x_{d+1} t) \leq 1$:

$$\begin{aligned}
&\prod_{i=1}^{d} \cos^2(x_i t) + \cos^2(x_{d+1} t) \sum_{i=1}^{d} \sin^2(x_i t) \prod_{i \neq j} \cos^2(x_j t) \\
&\leq \prod_{i=1}^{d} \cos^2(x_i t) + \sum_{i=1}^{d} \sin^2(x_i t) \prod_{i \neq j} \cos^2(x_j t) \overset{\text{IH}}{\leq} 1
\end{aligned}$$

Thus the claim holds for $d + 1$ completing the induction. ∎

Having done the hard work in comparing the Euclidean distance and the angular distance we can come back to the theorem that we want to prove.

Theorem 4.2: *Let $(p_\bullet, r_\bullet)$ be a d-dimensional ball embedding of a graph G in a hypercube $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$. Then the graphs H_1 with cap embedding $\left(\Phi(p_\bullet), \frac{\alpha_K}{\sqrt{d}} r_\bullet\right)$ and H_2 with cap embedding $\left(\Phi(p_\bullet), r_\bullet\right)$ are a sub- and a supergraph of G where $\alpha_K = \min_{p,\tilde{p}\in K} \prod_{i=2}^{d} \cos(p_i)\cos(\tilde{p}_i) > 0$.*

Proof. The proof is mainly an application of Lemma 4.3, which states that the distance functions in the ball model and the cap model differ only by a constant factor. We first note that for $x \in \mathbb{R}^d$ it holds that $\|x\|_2 \leq \sqrt{d}\|x\|_\infty$. Using this fact and Lemma 4.3, for $p, \tilde{p} \in [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ we get the following chain of inequalities

$$\frac{\alpha_{p,\tilde{p}}}{\sqrt{d}}\|p - \tilde{p}\|_2 \leq \alpha_{p,\tilde{p}}\|p - \tilde{p}\|_\infty \leq \operatorname{dist}_S(\Phi(p), \Phi(\tilde{p})) \leq \|p - \tilde{p}\|_2.$$

We first show the claim for the supergraph. Let $uv \in E$ be an edge in G, i.e., $\|p_u - p_v\|_2 \leq r_u + r_v$. As $\operatorname{dist}_S(\Phi(p_u), \Phi(p_v)) \leq \|p_u - p_v\|_2$ by the inequality above, uv is also an edge in H_2, i.e., $G \subseteq H_2$.

For the claim about the subgraph, let u and v be disconnected in G, i.e., $\|p_u - p_v\|_2 > r_u + r_v$. By the inequality above we get

$$\operatorname{dist}_S(\Phi(p_u), \Phi(p_v)) \geq \frac{\alpha_{p_u,p_v}}{\sqrt{d}}\|p_u - p_v\|_2 > \frac{\alpha_{p_u,p_v}}{\sqrt{d}}(r_u + r_v) \geq \frac{\alpha_K}{\sqrt{d}}(r_u + r_v),$$

as $\alpha_K = \min_{p_u,p_v\in K} \alpha_{p_u,p_v}$ by definition. This shows that u, v are disconnected in H_1 as well, in particular $H_1 \subseteq G$. ∎

4.2 Comparison of the Random Graph Models

Together Theorems 4.1 and 4.2 yield the following translations from a d-dimensional ball graph to $(d+1)$-dimensional dot product graphs, that form a sub- and a supergraph of the ball graph.

Corollary 4.5: *Let $(p_\bullet, r_\bullet)$ be a d-dimensional ball embedding of a graph G in a hypercube $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ with $r_v \in [0, \frac{\pi}{2}]$ for all $v \in V$. Then the graphs H_1 with dot product embedding $\left(\Phi_{d+1}(p_\bullet), \cos^{-1}\left(\frac{\alpha_K}{\sqrt{d}} r_\bullet\right)\right)$ and H_2 with dot product embedding $\left(\Phi_{d+1}(p_\bullet), \cos^{-1}(2r_\bullet)\right)$ are a sub- and a supergraph of G where $\alpha_K = \min_{p,\tilde{p}\in K} \prod_{i=2}^{d} \cos(p_i)\cos(\tilde{p}_i)$.*

Proof. Let G and its ball embedding $(p_\bullet, r_\bullet)$ be like in the claim. Then we can apply Theorem 4.2 as all prerequisites are fulfilled to get graphs $\tilde{H}_1$ with cap embedding $\left(\Phi_{d+1}(p_\bullet), \frac{\alpha_K}{\sqrt{d}} r_\bullet\right)$ and $\tilde{H}_2$ with cap embedding $(\Phi_{d+1}(p_\bullet), r_\bullet)$ who are a sub- and a supergraph of G.

For $\tilde{H}_1$ and $\tilde{H}_2$ the prerequisites of Theorem 4.1 are fulfilled. Applying this theorem to both graphs we get a subgraph H_1 of $\tilde{H}_1$ with dot product embedding $\left(\Phi_{d+1}(p_\bullet), \cos^{-1}\left(\frac{\alpha_K}{\sqrt{d}} r_\bullet\right)\right)$ and a supergraph H_2 of $\tilde{H}_2$ with dot product embedding $\left(\Phi_{d+1}(p_\bullet), \cos^{-1}(2r_\bullet)\right)$.

Thus, H_1 and H_2 are the desired sub- and supergraph of G. ∎

As there is an inverse for polar coordinates, and we change the radii by an invertible function, these translations also work in the other direction. Other direction means that for a dot product graph with lengths at least 1 and positions in a small subset of the sphere, we find ball graphs in one dimension lower that are a sub- and a supergraph of the dot product graph.

Corollary 4.6: *Let $(p_\bullet, \ell_\bullet)$ be a $(d+1)$-dimensional dot product embedding of a graph G in the image of hypercube $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ under stereographic projection $\Phi(K)$ with $\ell_v \geq 1$ for all $\ell_v \in V$. Then the graphs H_1 with ball embedding $\left(\Phi_{d+1}^{-1}(p_\bullet), \frac{1}{2}\arccos\left(\frac{1}{\ell_\bullet}\right)\right)$ and H_2 with ball embedding $\left(\Phi_{d+1}^{-1}(p_\bullet), \frac{\sqrt{d}}{\alpha_K}\arccos\left(\frac{1}{\ell_\bullet}\right)\right)$ are a sub- and a supergraph of G where $\alpha_K = \min_{q,\tilde{q}\in K}\prod_{i=2}^{d}\cos(q_i)\cos(\tilde{q}_i) > 0$.*

Proof. This corollary follows from Corollary 4.5.

Let G and its dot product embedding $(p_\bullet, \ell_\bullet)$ be like in the claim. First consider H_1 as the ball graph with embedding $(q_\bullet, r_\bullet) := \left(\Phi_{d+1}^{-1}(p_\bullet), \frac{1}{2}\arccos\left(\frac{1}{\ell_\bullet}\right)\right)$, thus we have $p_v = \Phi_{d+1}(q_v)$ and $\ell_v = \cos^{-1}(2r_v))$. Thus, G is a supergraph of H_1, by Corollary 4.5. The proof for the claim about H_2 works analogous. ∎

Now that we have established the connection for individual graphs we can study the corresponding random models. We will show that the $(d+1)$-dimensional random dot product graph model generalizes the d-dimenional random ball graph model, but also the other way around. The previous corollary provide us with translations between graphs, which are a good way to construct the necessary couplings. The only other thing, that we need to do, is proving that the expected average degrees differ only by a constant factor, when changing lengths or radii as in the corollaries above. As we will see in the proof, we need this calculation only for one model. This will follow from the fact, that the generalizations work in both directions. This comes in very handy, as the average degree calculation in random dot product graphs is very annoying. In the following lemma we show, that if we scale the radii of a random ball graph by a constant factor $\beta \geq 1$ then the average degree of this new random ball graph will differ by a constant factor c_β that does not depend on the number of vertices n, but only on β and the dimension of the random ball graph model.

Lemma 4.7: *Let $\beta \in \mathbb{R}_{\geq 1}$. For $G \sim \mathcal{RBG}_d(n, \mathcal{R}, K)$ and $G_\beta \sim \mathcal{RBG}_d(n, \beta\cdot\mathcal{R}, K)$ it holds that $\mathbb{E}(\bar{d}(G_\beta)) \leq \beta^d \cdot \mathbb{E}(\bar{d}(G))$, where c_β is a constant independent of n, $\mathcal{R}$ and K.*

Proof. For $\beta \in \mathbb{R}_{\geq 1}$, let $(P_\bullet, \beta\cdot R_\bullet)$ be the dot product embedding of a random graph $G_\beta \sim \mathcal{RBG}_d(n, \beta\cdot\mathcal{R}, K)$. This means that for each $v \in V$ we have random variables $P_v \sim \mathcal{U}(K)$ and $R_v \sim \mathcal{R}$.

We calculate

$$\mathbb{E}[\bar{d}(G_\beta)] = \mathbb{E}[\frac{1}{n}\sum_{u\in V(G_\beta)}\deg(u)] = \frac{1}{n}\sum_{u\in V(G_\beta)}\mathbb{E}[\deg(u)],$$

where for fixed $u \in V$:

$$\begin{aligned}\mathbb{E}[\deg(u)] &= \mathbb{E}[\sum_{v\in V}\mathbb{1}_{\{\|P_u-P_v\|_2\leq\beta(R_u+R_v)\}}]\\ &= \sum_{v\in V}\mathbb{E}[\mathbb{1}_{\{\|P_u-P_v\|_2\leq\beta(R_u+R_v)\}}] = \sum_{v\in V}\mathbb{P}(\{\|P_u-P_v\|_2\leq\beta(R_u+R_v)\}).\end{aligned}$$

This means that it suffices to show that

$$\mathbb{P}(\{\|P_u-P_v\|_2\leq\beta(R_u+R_v)\}) \leq c_\beta\cdot\mathbb{P}(\{\|P_u-P_v\|_2\leq(R_u+R_v)\})$$

for our desired constant c_β. [7]

[7] To make sense of these probabilities we need a coupling (i.e., a joint distribution) between P_u, P_v, R_u and R_v. This joint distribution is given by the fact, that these random variables are assumed to be independent.

Fix $p_u \in K$ and $r_u, r_v \in [0, \frac{\pi}{2}]$. We show that

$$\mathbb{P}(\{\|p_u - P_v\|_2 \leq \beta(r_v + r_w)\}) \leq c_\beta \cdot \mathbb{P}(\{\|p_u - P_v\|_2 \leq (r_u + r_v)\}). \tag{4.4}$$

Then the upper claim follows as an application of the law of total probability, using that P_u, P_v, R_u and R_v are independent.[8]

For $p_v \in K$, it holds that $\|p_u - p_v\|_2 \leq \beta(r_v + r_w)$ if and only if p_v is contained in the d-ball B_β with center p_u and radius $\beta(r_v + r_w)$. As p_v is drawn uniformly from K this yields

$$\mathbb{P}(\{\|p_u - P_v\|_2 \leq \beta(r_v + r_w)\}) = \frac{\operatorname{Vol}(B_\beta \cap K)}{\operatorname{Vol}(K)}. \tag{4.5}$$

Then Equation (4.4) follows from showing $\operatorname{Vol}(B_\beta \cap K) \leq c_\beta \operatorname{Vol}(B_1 \cap K)$. To do this, we note that for a measurable, bounded subspace $A \subseteq \mathbb{R}^d$ it holds that $\operatorname{Vol}(\beta A) = \beta^d \operatorname{Vol}(A)$ where $\beta A := \{\beta \cdot a \mid a \in A\}$. As Vol is invariant under translations, we can translate K, B_β and B_1 such that the center p_u of B_β and B_1 is 0. In the following arguments we use the same letters for the translated subsets.

Then it holds that $B_\beta \cap K \subseteq \beta(B_1 \cap K)$. This follows from the fact that for $x \in B_\beta \cap K$ the straight line from x to the center $p_u \in K$ is contained in K by convexity of the hypercube K. As $\frac{x}{\beta} \in B_1$ is contained in the straight line from x to $0 = p_u$ and thus in K it holds that $x \in \beta(B_1 \cap K)$.

This implies $\operatorname{Vol}(B_\beta \cap K) \leq \operatorname{Vol}(\beta(B_1 \cap K)) \leq \beta^d \operatorname{Vol}(B_1 \cap K)$, which was the last thing we needed to show. In particular $c_\beta = \beta^d$, which is independent of n, $\mathcal{R}$ and K. ∎

Having done this average degree calculation, we can now prove the promised theorems, that$(d + 1)$-dimensional random dot product graphs generalize d-dimensional random ball graphs up to constant factors and vice versa. We start with proving that the random dot product graph model generalizes the random ball graph model up to constant factors.

Theorem 4.8: *Let $d \in \mathbb{N}^+$. The random dot product graph model $\mathcal{RDPG}_{d+1}(n, \mathcal{L}, K)$ generalizes the random ball graph model $\mathcal{RBG}_d(n, \mathcal{R}, K)$ up to constant factors.*

Proof. Let $\mathcal{R}$ be a distribution of radii, $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ be a hypercube. We now have to find parameters $\mathcal{L}_1, K_1$ and $\mathcal{L}_2, K_2$ for the random dot product graph model and a coupling (G, H_1, H_2) of $\mathcal{RBG}_d(n, \mathcal{R}, K)$, $\mathcal{RDPG}_{d+1}(n, \mathcal{L}_1, K_1)$ and $\mathcal{RDPG}_{d+1}(n, \mathcal{L}_2, K_2)$ such that the coupling fulfills the properties specified in the definition of generalization up to constant factors. Namely, it has to hold that $H_1 \subseteq G \subseteq H_2$ and the expected average degrees of H_1 and H_2 differ only by a constant factor independent of $n, \mathcal{R}$ and K. Looking at Corollary 4.5 we choose $\mathcal{L}_1 = \cos^{-1}\left(\frac{\alpha_K}{\sqrt{d}}\mathcal{R}\right)$, $\mathcal{L}_2 = \cos^{-1}(2\mathcal{R})$ and $K_1 = K_2 = K$.

Let $G \sim \mathcal{RBG}_d(n, \mathcal{R}, K)$ with ball embedding (P_v, R_v), i.e., for each $v \in V$ we have independent random variables $P_v \sim \mathcal{U}(K)$ and $R_v \sim \mathcal{R}$.

We define H_1 as the random graph with embedding $\left(\Phi_{d+1}(P_\bullet), \cos^{-1}\left(\frac{\alpha_K}{\sqrt{d}}R_\bullet\right)\right)$ and H_2 as the random graph with embedding $\left(\Phi_{d+1}(P_\bullet), \cos^{-1}(2R_\bullet)\right)$. By Corollary 4.5 it follows that $H_1 \subseteq G \subseteq H_2$ and by construction we have $H_1 \sim \mathcal{RDPG}_{d+1}(n, \mathcal{L}_1, K_1)$ and $H_2 \sim \mathcal{RDPG}_{d+1}(n, \mathcal{L}_2, K_2)$. Moreover, G, H_1 and H_2 are defined on the same probability space, namely the one from G. This means that we have the desired coupling (G, H_1, H_2) for the three random graph models in question.

[8] To apply the law of total probability we need that all included random variables have probability density functions. This was assumed in the preliminaries.

The only thing left to show is the statement about the expected average degrees of H_1 and H_2. As it is hard to calculate the expected average degree of random dot product graphs, we use the translation in the other direction and find random ball graphs $H_1' \subseteq H_1$ and $H_2' \supseteq H_2$ and then compare the expected average degrees of H_1' and H_2'. More precisely we define H_1' as the random graph with embedding $\left(P_\bullet, \frac{1}{2}\frac{\alpha_K}{\sqrt{d}}R_\bullet\right)$ and H_2' as the random graph with embedding $\left(P_\bullet, \frac{\sqrt{d}}{\alpha_K}2R_\bullet\right)$.

By Corollary 4.6 it follows that $H_1' \subseteq H_1$ and $H_2 \subseteq H_2'$ and by construction we have $H_1' \sim \mathcal{RBG}_d(n, \frac{1}{2}\frac{\alpha_K}{\sqrt{d}}\mathcal{R}, K_1)$ and $H_2' \sim \mathcal{RBG}_d(n, \frac{\sqrt{d}}{\alpha_K}2\mathcal{R}, K_2)$. If we set $\mathcal{R}_1 := \frac{1}{2}\frac{\alpha_K}{\sqrt{d}}\mathcal{R}$ and $\mathcal{R}_2 := \frac{\sqrt{d}}{\alpha_K}2\mathcal{R}$, then $\mathcal{R}_2 = 4\left(\frac{\sqrt{d}}{\alpha_K}\right)^2\mathcal{R}_1$. Thus, we can apply Lemma 4.7 to H_1' and H_2' with $\beta_K = 4\left(\frac{\sqrt{d}}{\alpha_K}\right)^2 \geq 4d$ to get that $\mathbb{E}(\bar{d}(H_2)) \leq \mathbb{E}(\bar{d}(H_2')) \leq c_{\beta_K} \cdot \mathbb{E}(\bar{d}(H_1')) \leq c_{\beta_K} \cdot \mathbb{E}(\bar{d}(H_1))$, where the outer inequalities follow from the fact that $H_1' \subseteq H_1$ and $H_2' \supseteq H_2$.

This is almost what we want to show as c_{β_K} is independent of n and $\mathcal{R}$ but not independent of K. However, as $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$, we observe that $\alpha_K \geq \alpha_{[-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d} > 0$. Thus, it follows that $c_{\beta_K} \leq c_{\beta_{[-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d}} =: c$, which implies that the statement about expected degrees also holds with c instead of c_{β_K}. This concludes the proof, as c is the desired constant factor, independent of all parameters of the random ball graph model. ■

As we have not compared the average degrees of the random dot product graphs but constructed another sub- and supergraph, we remark that in reality the constant factor between the average degrees will be smaller than proven in the theorem. Also note that we have proven something even stronger than promised. If we restrict to smaller subsets $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$, the constant factor gets smaller, yielding better approximations of the original random ball graph model.

Now we prove the other direction, i.e., that random ball graphs generalize random dot product graphs up to constant factors. Note that we have implicitly used this statement when doing the degree calculation in the other proof. Thus, the arguments in the following proof should be familiar.

Theorem 4.9: *Let $d \in \mathbb{N}^+$. The random ball graph model $\mathcal{RBG}_d(n, \mathcal{R}, K)$ generalizes the random dot product graph model $\mathcal{RDPG}_{d+1}(n, \mathcal{L}, K)$ up to constant factors.*

Proof. Let $\mathcal{L}$ be a distribution of lengths in $\mathbb{R}_{\geq 1}$ and $K \subseteq [-\frac{\pi}{4\sqrt{d}}, \frac{\pi}{4\sqrt{d}}]^d$ be a hypercube. We now proceed exactly as above and choose parameters $\mathcal{R}_1, K_1$ and $\mathcal{R}_2, K_2$ and a coupling (G, H_1, H_2) of $\mathcal{RDPG}_{d+1}(n, \mathcal{R}_1, K)$, $\mathcal{RBG}_d(n, \mathcal{R}_1, K_1)$ and $\mathcal{RBG}_d(n, \mathcal{R}_2, K_2)$ such that the coupling fulfills the properties specified in the definition of generalization up to constant factors.

Looking at Corollary 4.6, we choose $\mathcal{R}_1 = \frac{1}{2}\arccos\left(1/\mathcal{L}\right)$, $\mathcal{R}_2 = \frac{\sqrt{d}}{\alpha_K}\arccos\left(1/\mathcal{L}\right)$ and $K_1 = K_2 = K$.

Let $G \sim \mathcal{RDPG}_{d+1}(n, \mathcal{L}, K)$ with dot product embedding (P_v, L_v), i.e., for each $v \in V$ we have independent random variables $P_v \sim \Phi(\mathcal{U}(K))$ and $L_v \sim \mathcal{L}$.

We define H_1 as the random graph with embedding $\left(\Phi_{d+1}^{-1}(P_\bullet), \frac{1}{2}\arccos\left(1/L_\bullet\right)\right)$ and H_2 as the random graph with embedding $\left(\Phi_{d+1}^{-1}(P_\bullet), \frac{\sqrt{d}}{\alpha_K}\arccos\left(1/L_\bullet\right)\right)$. By Corollary 4.6 it follows that $H_1 \subseteq G \subseteq H_2$ and by construction we have $H_1 \sim \mathcal{RBG}_d(n, \mathcal{R}_1, K_1)$ and $H_2 \sim \mathcal{RBG}_d(n, \mathcal{R}_2, K_2)$. Moreover, G, H_1 and H_2 are defined on the same probability space, namely the one from G. This means that we have the desired coupling (G, H_1, H_2) for the three random graph models in question.

The only thing left to show is the statement about the expected average degrees of H_1 and H_2. As $\mathcal{R}_2 = \frac{2\sqrt{d}}{\alpha_K}\mathcal{R}_1$ we can apply the degree calculation of Lemma 4.7 with $\beta_K = \frac{2\sqrt{d}}{\alpha_K} \geq 2$ and get that $\mathbb{E}(\bar{d}(H_2)) \leq c_{\beta_K} \cdot \mathbb{E}(\bar{d}(H_1))$. With the same arguments as in the proof of Theorem 4.8 we find a constant c independent of K. This concludes the proof.

■

5 Conclusion

In Chapter 3, we compared d-dimensional dot product and weighted embeddings. We first argued that these embeddings admit different artifacts in certain edge cases. We then showed that, even if we restrict ourselves to embeddings which do not have these artifacts, there is no universal translation from d-dimensional dot product graphs to d-dimensional weighted embeddings that turns lengths to weights but leaves the positions unchanged. This suggests that these models work differently and are hard to compare. We also constructed a translation from 1-dimensional dot product graphs to 1-dimensional weighted embeddings. However, this translation heavily depended on the fact that the 1-dimensional dot product is just a multiplication, like the multiplication of weights. We do not think that this proof generalizes into higher dimensions.

We thus moved onto models, where we suspected that they are more similar and easier to compare. In Chapter 4, we compared $(d + 1)$-dimensional dot product graphs and d-dimensional ball graphs and showed the similarity of the corresponding random graph models. More precisely, we proved that our $(d + 1)$-dimensional random dot product graph model generalizes the d-dimensional random ball graph model up to constant factors and the other way around. To be able to do this we looked at cap graphs which are very similar to ball graphs as proved in Section 4.1. We also found a natural way to turn caps into the positions and lengths of a dot product embedding by looking at the subset of positions that would get connected to a certain vector. Altogether, we established new connections between dot product, weighted and ball embeddings.

5.1 Future Work

The only relation of the three big models that we have not studied in this thesis, is the one between weighted and ball embeddings. Here, we strongly suspect that there is no universal translation from d-dimensional ball embeddings to d-dimensional weighted embeddings, that turns radii into weights, but leaves the positions unchanged, i.e. a similar theorem as Theorem 3.1. We think this statement is true, since in ball embeddings the radii are added but in weighted embeddings the weights are multiplied. Changing a summand has a different effect on the sum as changing a factor in a product.

Both in this case and in the setting of Theorem 3.1 one could try to find translations that also change the positions. For this it might be helpful to determine the embedding dimensions and corresponding embeddings for different graphs.

In the introduction, we also spoke about a fourth model, namely hyperbolic unit disk graphs. We already know that one can translate hyperbolic unit disk graphs into Euclidean disk graphs by using the Poincaré disk model. We suspect that it is possible to translate this connection to random graphs and then use our translations to dot product graphs to show that random dot product graphs are a generalization up to constant factors of hyperbolic random graphs. As GIRGs generalize hyperbolic random graphs up to constant factors, this would provide a new starting point for comparing dot product and weighted embeddings.

However, for this to work correctly, one needs different random graph models for dot product and ball graphs than ours. Most notably, positions would not be drawn uniformly at random, as we would start with a uniform distribution in bounded subset of hyperbolic space and then adjust them, because the center of a hyperbolic disk is not the same as the center of the corresponding Euclidean disk. Moreover, the radius of the disks will depend upon the position. As we scale the disks for the translation from disk graphs to dot product graphs, one also has to decide if one scales the radii of the hyperbolic or the corresponding Euclidean disk as these two notions differ. This could all be adapted to work out, the only critical thing one needs to show, is that the expected average degrees of the dot product graphs will differ only by a constant factor. Here one cannot use our degree calculation as it heavily used that positions were drawn uniformly at random and independent of the corresponding radii.

A last open question would be to relate the random dot product graph model studied in literature to either the models introduced in this thesis or to other existing models like GIRGs. The second one might be more fruitful as for both models it has been shown that they produce graphs with properties of real-world social network graphs [YS07|BKL19].

Bibliography

[Ahm+13] Amr Ahmed, Nino Shervashidze, Shravan Narayanamurthy, Vanja Josifovski, and Alexander J. Smola. “Distributed large-scale natural graph factorization”. In: *Proceedings of the 22nd International Conference on World Wide Web*. Rio de Janeiro, Brazil: Association for Computing Machinery, 2013, pp. 37–48. ISBN: 9781450320351. DOI: *10.1145/2488388.2488393*.

[And05] James W. Anderson. “Planar Models of the Hyperbolic Plane”. In: *Hyperbolic Geometry*. London: Springer London, 2005, pp. 117–144. ISBN: 978-1-84628-220-1. DOI: *10.1007/1-84628-220-9_4*.

[Ber25] Karl Bernhard. “Degree Independent Weights for Weighted Embeddings”. Karlsruhe Institute of Technology, 2025.

[BFK18] Thomas Bläsius, Tobias Friedrich, and Anton Krohmer. “Cliques in Hyperbolic Random Graphs”. In: *Algorithmica* Volume 80 (2018), pp. 2324–2344. ISSN: 1432-0541. DOI: *10.1007/s00453-017-0323-3*.

[BFKS23] Thomas Bläsius, Tobias Friedrich, Maximilian Katzmann, and Daniel Stephan. “Strongly Hyperbolic Unit Disk Graphs”. In: *40th International Symposium on Theoretical Aspects of Computer Science (STACS 2023)*. Edited by Petra Berenbrink, Patricia Bouyer, Anuj Dawar, and Mamadou Moustapha Kanté. Vol. 254. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 13:1–13:17. ISBN: 978-3-95977-266-2. DOI: *10.4230/LIPIcs.STACS.2023.13*.

[BFM13] Michel Bode, Nikolaos Fountoulakis, and Tobias Müller. “On the giant component of random hyperbolic graphs”. In: *The Seventh European Conference on Combinatorics, Graph Theory and Applications*. Edited by Jaroslav Nešetřil and Marco Pellegrini. Pisa: Scuola Normale Superiore, 2013, pp. 425–429. ISBN: 978-88-7642-475-5.

[BKL19] Karl Bringmann, Ralph Keusch, and Johannes Lengler. “Geometric inhomogeneous random graphs”. In: *Theoretical Computer Science* Volume 760 (2019), pp. 35–54. ISSN: 0304-3975. DOI: *https://doi.org/10.1016/j.tcs.2018.08.014*.

[Blä+22] Thomas Bläsius, Tobias Friedrich, Maximilian Katzmann, Ulrich Meyer, Manuel Penschuck, and Christopher Weyand. “Efficiently generating geometric inhomogeneous and hyperbolic random graphs”. In: *Network Science* Volume 10 (2022), pp. 361–380. DOI: *10.1017/nws.2022.32*.

[CCJ90] Brent N. Clark, Charles J. Colbourn, and David S. Johnson. “Unit disk graphs”. In: *Discrete Mathematics* Volume 86 (1990), pp. 165–177. ISSN: 0012-365X. DOI: *https://doi.org/10.1016/0012-365X(90)90358-O*.

[DT19] Mikhail Drobyshevskiy and Denis Turdakov. “Random Graph Modeling: A Survey of the Concepts”. In: *ACM Comput. Surv.* Volume 52 (Dec. 2019). ISSN: 0360-0300. DOI: *10.1145/3369782*.

[FGKS24] Tobias Friedrich, Andreas Göbel, Maximilian Katzmann, and Leon Schiller. “Cliques in High-Dimensional Geometric Inhomogeneous Random Graphs”. In: *SIAM Journal on Discrete Mathematics* Volume 38 (2024), pp. 1943–2000. eprint: *https://doi.org/10.1137/23M157394X*.

[Fis83] Peter C Fishburn. “On the sphericity and cubicity of graphs”. In: *Journal of Combinatorial Theory, Series B* Volume 35 (1983), pp. 309–318. ISSN: 0095-8956. DOI: *https://doi.org/10.1016/0095-8956(83)90057-6*.

[FSTZ98] Charles M. Fiduccia, Edward R. Scheinerman, Ann Trenk, and Jennifer S. Zito. “Dot product representations of graphs”. In: *Discrete Mathematics* Volume 181 (1998), pp. 113–138. ISSN: 0012-365X. DOI: *10.1016/S0012-365X(97)00049-6*.

[GPP12] Luca Gugelmann, Konstantinos Panagiotou, and Ueli Peter. “Random Hyperbolic Graphs: Degree Sequence and Clustering”. In: *Automata, Languages, and Programming*. Edited by Artur Czumaj, Kurt Mehlhorn, Andrew Pitts, and Roger Wattenhofer. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 573–585. ISBN: 978-3-642-31585-5.

[Gro24] Lennart Großkreutz. “Hyperbolic Sphericity”. Karlsruhe Institute of Technology, 2024.

[JPL21] Matthew Johnson, Daniël Paulusma, and Erik Jan van Leeuwen. “What Graphs are 2-Dot Product Graphs?” In: *International Journal of Computational Geometry & Applications* Volume 31 (2021), pp. 1–16. eprint: *https://doi.org/10.1142/S0218195921500011*.

[KL20] Júlia Komjáthy and Bas Lodewijks. “Explosion in weighted hyperbolic random graphs and geometric inhomogeneous random graphs”. In: *Stochastic Processes and their Applications* Volume 130 (2020), pp. 1309–1367. ISSN: 0304-4149. DOI: *https://doi.org/10.1016/j.spa.2019.04.014*.

[KM11a] Ross J. Kang and Tobias Müller. “Dot Product Representations of Planar Graphs”. In: *Graph Drawing*. Edited by Ulrik Brandes and Sabine Cornelsen. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 287–292. ISBN: 978-3-642-18469-7.

[KM11b] Ross J. Kang and Tobias Müller. “Sphere and dot product representations of graphs”. In: *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry*. Paris, France: Association for Computing Machinery, 2011, pp. 308–314. ISBN: 9781450306829. DOI: *10.1145/1998196.1998249*.

[LC14] Bo-Jr Li and Gerard Jennhwa Chang. “Dot product dimensions of graphs”. In: *Discrete Applied Mathematics* Volume 166 (2014), pp. 159–163. ISSN: 0166-218X. DOI: *https://doi.org/10.1016/j.dam.2013.10.014*.

[Mae84] Hiroshi Maehara. “Space graphs and sphericity”. In: *Discrete Applied Mathematics* Volume 7 (1984), pp. 55–64. ISSN: 0166-218X. DOI: *https://doi.org/10.1016/0166-218X(84)90113-6*.

[MKNŠ21] Ilya Makarov, Dmitrii Kiselev, Nikita Nikitinsky, and Lovro Šubelj. “Survey on graph embeddings and their applications to machine learning problems on graphs”. In: *PeerJ Computer Science* Volume 7 (2021), e357–e357. ISSN: 2376-5992. DOI: *10.7717/peerj-cs.357*.

[MQ06] T.S. Michael and Thomas Quint. "Sphericity, cubicity, and edge clique covers of graphs". In: *Discrete Applied Mathematics* Volume 154 (2006), pp. 1309–1313. ISSN: 0166-218X. DOI: *https://doi.org/10.1016/j.dam.2006.01.004*.

[MS22] Riccardo Michielan and Clara Stegehuis. "Cliques in geometric inhomogeneous random graphs". In: *Journal of Complex Networks* Volume 10 (Feb. 2022), cnac002. ISSN: 2051-1329. eprint: *https://academic.oup.com/comnet/article-pdf/10/1/cnac002/42466016/cnac002.pdf*.

[Ou+16] Mingdong Ou, Peng Cui, Jian Pei, Ziwei Zhang, and Wenwu Zhu. "Asymmetric Transitivity Preserving Graph Embedding". In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. San Francisco, California, USA: Association for Computing Machinery, 2016, pp. 1105–1114. ISBN: 9781450342322. DOI: *10.1145/2939672.2939751*.

[PKBV10] Fragkiskos Papadopoulos, Dmitri Krioukov, Marián Boguñá, and Amin Vahdat. "Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces". In: *Proceedings of the 29th Conference on Information Communications*. San Diego, California, USA: IEEE Press, 2010, pp. 2973–2981. ISBN: 9781424458363.

[RCTP22] Patrick Rubin-Delanchy, Joshua Cape, Minh Tang, and Carey E. Priebe. "A Statistical Interpretation of Spectral Embedding: The Generalised Random Dot Product Graph". In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* Volume 84 (June 2022), pp. 1446–1473. ISSN: 1369-7412. eprint: *https://academic.oup.com/jrsssb/article-pdf/84/4/1446/49463141/rssb12509-sup-0001-supinfo.pdf*.

[RRŠ89] J. Reiterman, V. Rödl, and E. Šiňajová. "Geometrical embeddings of graphs". In: *Discrete Mathematics* Volume 74 (1989), pp. 291–319. ISSN: 0012-365X. DOI: *https://doi.org/10.1016/0012-365X(89)90142-8*.

[Sac94] Horst Sachs. "Coin graphs, Polyhedra, and conformal mapping". In: *Discrete Mathematics* Volume 134 (1994), pp. 133–138. ISSN: 0012-365X. DOI: *https://doi.org/10.1016/0012-365X(93)E0068-F*.

[SBNH21] Francesco Sanna Passino, Anna S. Bertiger, Joshua C. Neil, and Nicholas A. Heard. "Link prediction in dynamic networks using random dot product graphs". In: *Data Mining and Knowledge Discovery* Volume 35 (Sept. 2021), pp. 2168–2199. ISSN: 1573-756X. DOI: *10.1007/s10618-021-00784-2*.

[STP14] Daniel L. Sussman, Minh Tang, and Carey E. Priebe. "Consistent Latent Position Estimation and Vertex Classification for Random Dot Product Graphs". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* Volume 36 (2014), pp. 48–57. DOI: *10.1109/TPAMI.2013.135*.

[Xu21] Mengjia Xu. "Understanding Graph Embedding Methods and Their Applications". In: *SIAM Review* Volume 63 (2021), pp. 825–853. eprint: *https://doi.org/10.1137/20M1386062*.

[YS07] Stephen J. Young and Edward R. Scheinerman. "Random Dot Product Graph Models for Social Networks". In: *Algorithms and Models for the Web-Graph*. Edited by Anthony Bonato and Fan R. K. Chung. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 138–149. ISBN: 978-3-540-77004-6.