Chapter 5

Orthogonal representations and their dimension

An *orthogonal representation* of a simple graph G in $\mathbb{R}^d$ assigns to each $i \in V$ a vector $\mathbf{u}_i \in \mathbb{R}^d$ such that $\mathbf{u}_i^\mathsf{T} \mathbf{u}_j = 0$ whenever $ij \in \overline{E}$. An *orthonormal representation* is an orthogonal representation in which all the representing vectors have unit length. Clearly we can always scale the nonzero vectors in an orthogonal representation this way, and usually this does not change any substantial feature of the problem.

Note that we did not insist that adjacent nodes are mapped onto non-orthogonal vectors. If this condition also holds, then we call the orthogonal representation *faithful*.

Example 5.0.1 For $d = 1$, the vector labels are just real numbers u_i, and the constraints $u_i u_j = 0$ $(ij \in \overline{E})$ mean that no two nodes labeled by nonzero numbers are adjacent; in other words, $\mathrm{supp}(\mathbf{u})$ is a stable set of nodes. Orthonormal representations are obtained as the incidence vectors of stable sets, with arbitrary signs.

Since very simple problems about stable sets are NP-hard (for example, their maximum size), this example should warn us that orthogonal representations can be very complex. ♦

Example 5.0.2 Every graph has a trivial orthonormal representation in $\mathbb{R}^V$, in which node i is represented by the standard basis vector $\mathbf{e}_i$. (This representation is not faithful unless the graph has no edges. However, it is easy to perturb this representation to make it faithful.) Of course, we are interested in "nontrivial" orthogonal representations, which are more "economical" than the trivial one. ♦

Example 5.0.3 Every graph has a faithful orthogonal representation in $\mathbb{R}^E$, in which we label a node i by the incidence vector ∇_i of the set of edges incident with it. It is perhaps natural to expect that this simple representation will be rather "uneconomical" for most purposes. ♦

Example 5.0.4 Figure 5.1 below shows that for the graph obtained by adding a diagonal to the pentagon a simple orthogonal representation in 2 dimensions can be constructed. ♦

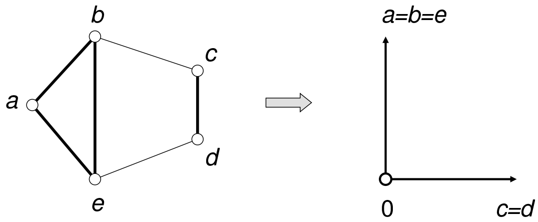

Figure 5.1: An (almost) trivial orthogonal representation

Example 5.0.5 The previous example can be generalized. In an orthogonal representation, we can label a set of nodes with the same nonzero vector if and only if these nodes form a clique. Let $k = \overline{\chi}(G)$, and let $\{B_1, \dots, B_m\}$ be a family of disjoint complete subgraphs covering all the nodes. Let $\{\mathbf{e}_1, \dots, \mathbf{e}_m\}$ be the standard basis of $\mathbb{R}^m$. Then mapping every node of B_i to $\mathbf{e}_i$ is an orthonormal representation. ♦

Example 5.0.6 In the introduction we have seen an orthogonal representation with a more interesting geometric content. The previous example gives an orthogonal representation of C_5 in 3-space (Figure 5.2, left). The "umbrella" representation defined in the introduction gives another orthogonal representation of C_5 in 3-space. ♦

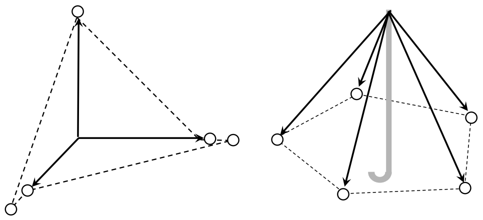

Figure 5.2: Two orthogonal representations of C_5.

Perhaps the most natural way to be "economic" in constructing an orthogonal representation is to minimize the dimension. We can say only a little about the minimum dimension of all orthogonal representations, but we get interesting results if we impose some "non-degeneracy" conditions.

5.1 Minimum dimension with no restrictions

We start with some easy facts about the minimum dimension in which G has an orthonormal representation. It is trivial that $d \geq \alpha(G)$, since there are $\alpha(G)$ mutually orthogonal vector labels.

The minimum dimension of an orthonormal representation is (loosely) tied to the chromatic number of the complement. Let us consider the (infinite) graph H_d, whose nodes are all vectors in $\mathbb{R}^d$, two of them being adjacent if and only if the vectors are orthogonal. The chromatic number of this graph is not known precisely, but (for large d) rather good bounds are known: there is a constant $c > 1$ for which

$$c^d \le \chi(H_d) \le 4^d. \tag{5.1}$$

Using this, it will be easy to prove the following.

Proposition 5.1.1 *The minimum dimension d in which a graph G has an orthonormal representation satisfies*

$$\frac{1}{2}\log\chi(\overline{G}) \le d \le \chi(\overline{G}).$$

Proof. The orthogonal representation of G in $\mathbb{R}^d$ gives a homomorphism $\overline{G} \to H_d$, and hence

$$\chi(\overline{G}) \le \chi(H_d) \le 4^d,$$

proving the lower bound in the proposition. On the other hand, the trivial orthonormal representation in Example 5.0.5 shows that G has an orthonormal representation in dimension $\chi(\overline{G})$. □

The bounds above are essentially tight. The Erdős–de Bruijn Theorem implies that H_d has a finite subgraph H with $\chi(H) = \chi(H_d) \ge c^d$ (where $c > 1$ is the constant in (5.1)), showing that the lower bound is tight up to a constant factor. The graph on d nodes and no edges shows the tightness of the upper bound.

5.2 General position and connectivity

The first non-degeneracy condition we study is *general position*: we assume that any d of the representing vectors in $\mathbb{R}^d$ are linearly independent. (This implies that neither one of them is the zero vector.) Perhaps surprisingly, there is an exact condition for this type of geometric representability.

Theorem 5.2.1 *A graph with n nodes has a general position orthogonal representation in $\mathbb{R}^d$ if and only if it is $(n-d)$-connected.*

It is useful to remark that a graph is $(n-d)$-connected if and only if its complement does not contain a complete bipartite graph with more than d nodes (where for a complete bipartite graph we always assume that its bipartition classes are nonempty). So we could formulate this theorem as follows: *If a graph contains no complete bipartite subgraph with more than $d+1$ nodes, then it has a dual orthogonal representation in $\mathbb{R}^d$.*

The condition that the given set of representing vectors is in general position is not easy to check (it is NP-hard). A weaker, but very useful condition will be that the vectors representing the nodes nonadjacent to any node v are linearly independent. We say that such a representation is in *locally general position*. This condition implies that every node is represented by a nonzero vector unless it is connected to all the other nodes. In this case there is no condition on the representing vector, so we may assume that all vectors are nonzero (equivalently, unit vectors).

Theorem 5.2.1 will be proved in the following slightly more general form:

Theorem 5.2.2 *If G is a graph with n nodes, then the following are equivalent:*

(i) *G has a general position orthogonal representation in $\mathbb{R}^d$;*

(ii) *G has an orthogonal representation in $\mathbb{R}^d$ in locally general position;*

(iii) *G is $(n-d)$-connected.*

We describe the simple proofs of two implications. The third one will take most of this chapter.

(i)⇒(ii). Note that (i) implies that all degrees are at least $n-d$. Indeed, if $\deg(v) < n-d$ then consider any set of d non-neighbors $S \subseteq \overline{N}(v)$, $|S| = d$. In an orthogonal representation in general position, the positions of the nodes in S are linearly independent, so they span the space, and so v cannot be represented by a vector orthogonal to all of them.

This implies that every node has at most $d-1$ non-neighbors, which are therefore represented by linearly independent vectors, so the representation is in locally general position.

(ii)⇒(iii). Let V_0 be a cutset of nodes of G, then $V = V_0 \cup V_1 \cup V_2$, where $V_1, V_2 \neq \emptyset$, and no edge connects V_1 and V_2. This implies that the vectors representing V_1 are non-neighbors of any node in V_2, and hence they are represented by linearly independent vectors in an orthogonal representation in locally general position. Similarly, the vectors representing V_2 are linearly independent. Since the vectors representing V_1 and V_2 are mutually orthogonal, all vectors representing $V_1 \cup V_2$ are linearly independent. Hence $d \geq |V_1 \cup V_2| = n - |V_0|$, and so $|V_0| \geq n-d$ (Figure 5.3).

The difficult part of the proof will be (iii)⇒(i): this requires the construction of a general position orthogonal representation for $(n-d)$-connected graphs. We describe and analyze the algorithm constructing the representation first. As a matter of fact, to describe the construction is quite easy, the difficulty lies in the proof of its validity.

5.2.1 G-orthogonalization

The following procedure can be viewed as an extension of the Gram-Schmidt orthogonalization algorithm. Let G be any simple graph, where $V = [n]$, and let $\mathbf{u}: V \to \mathbb{R}^d$ be any vector

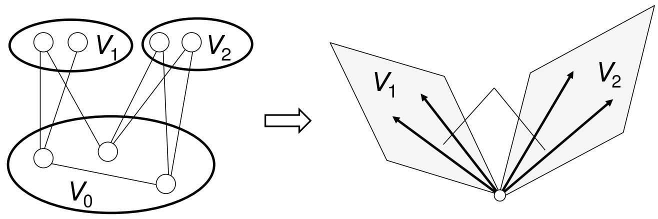

Figure 5.3: General position orthogonal representation in low dimension implies high connectivity.

labeling. Let us choose vectors $\mathbf{f}_1, \mathbf{f}_2, \ldots$ consecutively as follows. Let $\mathbf{f}_1 = \mathbf{u}_1$. Supposing that the vectors $\mathbf{f}_i$ $(i < j)$ are already chosen, we define $\mathbf{f}_j$ as the orthogonal projection of $\mathbf{u}_j$ onto the subspace $L_j = \text{lin}\{\mathbf{f}_i : \ i < j, \, ij \notin E\}^\perp$. The sequence $(\mathbf{f}_1, \ldots, \mathbf{f}_n)$ is a vector labeling of G, which we call the *G-orthogonalization* of the vector labeling $\mathbf{u}$. If $E = \emptyset$, then $(\mathbf{f}_1, \ldots, \mathbf{f}_n)$ is just the Gram–Schmidt orthogonalization of $(\mathbf{u}_1, \ldots, \mathbf{u}_n)$. If G is complete, then $(\mathbf{f}_1, \ldots, \mathbf{f}_n) = (\mathbf{u}_1, \ldots, \mathbf{u}_n)$.

It is important to note that the G-orthogonalization of a vector labeling depends on the ordering of V. In this section we will have to compare G-orthogonalizations obtained through different orderings of a fixed graph G.

It is trivial that $\mathbf{f}$ is an orthogonal representation of G, and it follows by straightforward induction on j that for every $j \in [n]$,

$$\text{lin}\{\mathbf{u}_1, \ldots, \mathbf{u}_j\} = \text{lin}\{\mathbf{f}_1, \ldots, \mathbf{f}_j\}.$$

From now on, we start with a *random vector labeling* $\mathbf{u}$: this means that the vectors $\mathbf{u}_i$ are random, chosen independently and uniformly from S^{d-1}. Its G-orthogonalization with respect to an ordering σ of V can be considered as a mapping $\mathbf{f}^\sigma : V \to \mathbb{R}^d$, which can also be considered as a single point in the dn dimensional space $\mathbb{R}^d \times V$. It is important to remember that $\mathbf{f}_i^\sigma$ denotes the vector assigned to the node i, and not to the i-th node in the ordering σ. It will be convenient to define $N^\sigma(i)$ and $\overline{N}^\sigma(i)$ as the set of preceding i in the ordering σ that are adjacent and nonadjacent to i, respectively.

If every node of an ordered graph G has degree at least $n - d$, then the vectors in the G-orthogonalization of a random vector labeling $\mathbf{u}$ are almost surely nonzero. Indeed, for any node i, $|\overline{N}^\sigma(i)| \le d - 1$, hence the vectors $\mathbf{f}_j^\sigma$, $j \in \overline{N}^\sigma(i)$ don't span $\mathbb{R}^d$, and hence almost surely $\mathbf{u}_i$ is linearly independent of them, and then its projection onto $\text{lin}(\overline{N}^\sigma(i))^\perp$ is nonzero.

The main fact we want to prove is the following. (This will complete the proof of Theorem 5.2.2.)

Lemma 5.2.3 *If an ordered graph G is $(n-d)$-connected, and $\mathbf{u}$ is a random vector labeling of $V(G)$, then the vectors in its G-orthogonalization are almost surely in general position.*

Let us consider a simple example.

Example 5.2.4 Let $V = \{a, b, c, d\}$ and $E = \{ac, bd\}$. Consider a random vector labeling $\mathbf{u}$ in $\mathbb{R}^3$ and compute its G-orthogonalization $\mathbf{f}$, associated with the given ordering. Since every node has degree 1, the vectors $\mathbf{f}_i$ are nonzero almost surely. We have $\mathbf{f}_a = \mathbf{u}_a$, $\mathbf{f}_b \in \mathbf{f}_a^\perp$, and $\mathbf{f}_c \in \mathbf{f}_b^\perp$; almost surely, $\mathbf{f}_c$ will not be parallel to $\mathbf{f}_a$, so together they span the plane $\mathbf{f}_b^\perp$. This means that $\mathbf{f}_d$, which is orthogonal to both $\mathbf{f}_a$ and $\mathbf{f}_c$, must be parallel to $\mathbf{f}_b$, no matter what $\mathbf{u}_d$ is.

Now suppose that we reverse the order of the nodes. Let $\mathbf{f}'_d, \mathbf{f}'_c, \mathbf{f}'_b, \mathbf{f}'_a$ be the vectors obtained by the G-orthogonalization in this order. Almost surely $\mathbf{f}'_d$ will not be parallel to $\mathbf{f}'_b$, but $\mathbf{f}'_c$ will be parallel with $\mathbf{f}'_a$. So not only are the two distributions different, but a particular event, namely $\mathbf{f}_b \| \mathbf{f}_d$, occurs with probability 0 in one and probability 1 in the other.

Let us modify this example by connecting b and c by an edge. Processing the vectors in the order (a, b, c, d) again, the vectors $\mathbf{f}_a = \mathbf{u}_a$ and $\mathbf{f}_b$ will be orthogonal again. No condition on $\mathbf{f}_c$, so almost surely $\mathbf{f}_c = \mathbf{u}_c$ will be linearly independent of $\mathbf{f}_a$ and $\mathbf{f}_b$, but not orthogonal to either one of them. The direction of $\mathbf{f}_d$ is still determined, but now it will not be parallel to $\mathbf{f}_b$; in fact, depending on $\mathbf{f}_c$, it can have any direction in $\mathbf{f}_a^\perp$. Doing this in the reverse order, we get vectors $\mathbf{f}'_d, \mathbf{f}'_c, \mathbf{f}'_b, \mathbf{f}'_a$ that have similar properties. The distributions of $(\mathbf{f}_d, \mathbf{f}_c, \mathbf{f}_b, \mathbf{f}_a)$ and $(\mathbf{f}'_d, \mathbf{f}'_c, \mathbf{f}'_b, \mathbf{f}'_a)$ are still different, but any event that occurs with probability 0 in one will also occur with probability 0 in the other (this is not quite obvious). ⧫

This example motivates the following considerations. As noted, the distribution of the G-orthogonalization may depend on the ordering of the nodes; the key to the proof will be that this dependence is not too strong. To define what this means, consider two random variables X and Y with values in $\mathbb{R}^M$. We say that these are *mutually absolutely continuous*, if for every Borel set B,

$$\mathsf{P}(X \in B) = 0 \quad \Leftrightarrow \quad \mathsf{P}(Y \in B) = 0.$$

Informally, if a "reasonable" property of X almost surely holds (for example, its coordinates satisfy an algebraic inequality), then this is also true for Y, and vice versa.

We need the following property of being mutually absolutely continuous. (The property is quite natural, and not hard to prove, but we don't prove it, since we have not developed the necessary measure theoretic foundations of probability theory.) Consider two random variables, each of which consist of two components: $X = (A, B)$ and $X' = (A', B')$, where $A, A' \in \mathbb{R}^M$ and $B, B' \in \mathbb{R}^N$. Assume that A and A' are mutually absolutely continuous. Furthermore, assume that both B and B' can be conditioned on every $a \in \mathbb{R}^M$. (Informally, this means that we generate X by generating a value $a \in \mathbb{R}^M$ of A, and then depending on this, we generate a value b of the random variable $B|a$.) We also assume that for every a, $B|a$ and $B'|a$ are mutually absolutely continuous. Then X and Y are mutually absolutely continuous. This is not hard to prove, using measure theory.

Observe that the probability that the first d vectors in a G-orthogonalized random representation are linearly dependent is 0. We are going to show that this event has probability 0 if the G-orthogonalization is carried out in any other order. Since we can start the ordering with any d-tuple of nodes, this will imply that the representation is almost surely in general position, proving Lemma 5.2.3.

So it suffices to prove the the following.

Lemma 5.2.5 (Main Lemma) *If G is $(n-d)$-connected, then for any two orderings σ and τ of V, the random variables $\mathbf{f}^\sigma$ and $\mathbf{f}^\tau$ are mutually absolute continuous.*

Proof. It suffices to prove this in the case when τ is the ordering obtained from σ by swapping the nodes in positions j and $j+1$ $(1 \le j \le n-1)$. Let us label the nodes so that $\sigma = (1, 2, \dots, n)$.

Clearly $\mathbf{f}_i^\sigma = \mathbf{f}_i^\tau$ for $i < j$. The main part of the proof is to show that the distributions of $(\mathbf{f}_1^\sigma, \dots, \mathbf{f}_{j+1}^\sigma)$ and $(\mathbf{f}_1^\tau, \dots, \mathbf{f}_{j+1}^\tau)$ are mutually absolute continuous. We prove this by induction on j, distinguishing several cases.

Case 1. j and $j+1$ are adjacent in G. In this case the vector $\mathbf{f}_{j+1}$ does not depend on $\mathbf{u}_j$ and vice versa, so $\mathbf{f}_j^\sigma = \mathbf{f}_j^\tau$ and $\mathbf{f}_{j+1}^\sigma = \mathbf{f}_{j+1}^\tau$.

Case 2. j and $j+1$ are not adjacent, but they are connected by a path that lies entirely in $\{1, \dots, j, j+1\}$. Let P be a shortest such path and let t be its length (number of edges), so $2 \le t \le j$. We argue by induction on t (and on j). Let i be any internal node of P. We swap j and $j+1$ by the following steps (Figure 5.4):

(1) Interchange i and j, by several swaps of consecutive elements among the first j.

(2) Swap i and $j+1$.

(3) Interchange $j+1$ and j, by several swaps of consecutive elements among the first j.

(4) Swap j and i.

(5) Interchange $j+1$ and i, by several swaps of consecutive elements among the first j.

In each step, the new and the previous distributions of the G-orthogonalized vectors are mutually absolute continuous: In steps (1), (3) and (5) this is so because the swaps take place among the first j nodes, and we can invoke induction on j; in steps (2) and (4), because the nodes swapped are at a smaller distance than t in the graph distance, and we can invoke induction on t.

Case 3. There is no path connecting j to $j+1$ within $\{1, \dots, j+1\}$. This case is tedious but not particularly deep. It suffices to show that the distributions of $(\mathbf{f}_j^\sigma, \mathbf{f}_{j+1}^\sigma)$ and $(\mathbf{f}_j^\tau, \mathbf{f}_{j+1}^\tau)$, conditioned on the previous vectors $\mathbf{v}_1, \dots, \mathbf{v}_{j-1}$, are mutually absolute continuous.

For $S \subseteq V$, let $\mathrm{lin}(S)$ denote the linear subspace of $\mathbb{R}^d$ generated by the vectors $\{\mathbf{f}_i : i \in S\}$. Clearly $V \setminus \{1, \dots, j+1\}$ is a cutset, whence by our hypothesis on the connectivity of G

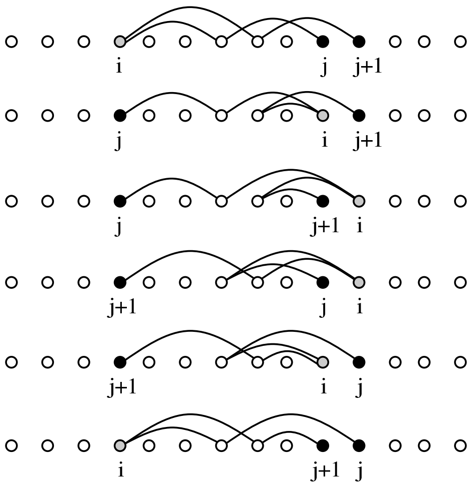

Figure 5.4: Interchanging j and $j+1$.

it follows that $n-j-1 \geq n-d$ and so $j \leq d-1$. Let $B = \mathrm{lin}\{\mathbf{v}_1, \ldots, \mathbf{v}_{j-1}\}$ and $C = B^\perp$, then $t = \dim(C) = d - \dim(B) \geq d - (j-1) \geq 2$.

The random vector $\mathbf{v}_j$ can be decomposed as $\mathbf{v}_j = \sin(\theta_j)\mathbf{b}_j + \cos(\theta_j)\mathbf{c}_j$, where $\mathbf{b}_j \in B$, $\mathbf{c}_j \in C$, $|\mathbf{b}_j| = |\mathbf{c}_j| = 1$, and $0 \leq \theta_j \leq \pi/2$. The unit vectors $\mathbf{b}_j$ and $\mathbf{c}_j$ are uniformly distributed on the unit sphere of B and C, respectively, and θ_j has some distribution which depends on the dimensions, but what is important for us is that it has a positive and continuous density function on $(0, \pi/2)$.

The subspace which is orthogonal to all previous non-neighbors of j (in the ordering σ) is $L_j = (\overline{N}^\sigma(j))^\perp$. Thus $C \subseteq L_j$. Let $\mathbf{b}'_j$ be the orthogonal projection of $\mathbf{b}_j$ onto L_j. Since $\mathbf{b}_j$ is orthogonal to C, so is $\mathbf{b}'_j$, thus $\mathbf{b}'_j \in B$.

We define L_{j+1} similarly for the node $j+1$ and ordering τ, along with the decomposition $\mathbf{v}_{j+1} = \sin(\theta_{j+1})\mathbf{b}_{j+1} + \cos(\theta_{j+1})\mathbf{c}_{j+1}$, and the orthogonal projection $\mathbf{b}'_{j+1}$ of $\mathbf{b}_{j+1}$ onto L.

The following observation will be important:

$$\mathbf{b}'_j \perp \mathbf{b}'_{j+1}. \tag{5.2}$$

Here we have to use that there is no path in J connecting j and $j+1$. This implies that the set $\{1, \ldots, j-1\}$ has a partition $W' \cup W''$ so that $N^\sigma(j) \subseteq W'$, $N^\tau(j+1) \subseteq W''$, and there is no edge between W' and W''. Let $B' = \mathrm{lin}(W')$ and $B'' = \mathrm{lin}(W'')$, then $B' \perp B''$, and $L_j \perp B''$, which implies that $\mathbf{b}'_j \in B'$. Similarly, $\mathbf{b}'_{j+1} \in B''$, which implies (5.2).

To get $\mathbf{f}^\sigma_j$, we have to project $\mathbf{v}_j$ onto L_j; using that $C \subseteq L_j$, the component of $\mathbf{v}_j$ in C remains unchanged, and this projection can be expressed as

$$\mathbf{f}^\sigma_j = \sin(\theta_j)\mathbf{b}'_j + \cos(\theta_j)\mathbf{c}_j \tag{5.3}$$

Similarly, we have

$$\mathbf{f}^\tau_{j+1} = \sin(\theta_{j+1})\mathbf{b}'_{j+1} + \cos(\theta_{j+1})\mathbf{c}_{j+1}. \tag{5.4}$$

To describe the other two orthogonalized vectors, notice that the only difference between $\mathbf{f}^\tau_{j+1}$ and $\mathbf{f}^\sigma_{j+1}$ is that we have to make $\mathbf{f}^\sigma_{j+1}$ orthogonal to $\mathbf{f}^\sigma_j$ as well. Since the vectors $\mathbf{b}'_j$ and $\mathbf{b}'_{j+1}$ are already orthogonal to each other, as well as to all vectors in C, this simply means that we have to modify (5.4) by replacing $\mathbf{c}_{j+1}$ by its projection onto the orthogonal complement of $\mathbf{c}_j$. So we can write

$$\mathbf{f}^\sigma_{j+1} = \sin(\theta_{j+1})\mathbf{b}'_{j+1} + \cos(\theta_{j+1})\cos(\phi)\mathbf{d}_{j+1}, \tag{5.5}$$

where $\mathbf{d}_{j+1} \in C$ is the unit vector pointing in the direction of the component of $\mathbf{c}_{j+1}$ orthogonal to $\mathbf{c}_j$, and $\phi = \measuredangle(\mathbf{c}_j, \mathbf{c}_{j+1})$. Similarly

$$\mathbf{f}^\tau_j = \sin(\theta_j)\mathbf{b}'_j + \cos(\theta_j)\cos(\phi)\mathbf{d}_j. \tag{5.6}$$

We can generate the vectors $\mathbf{f}^\sigma_j$, $\mathbf{f}^\sigma_{j+1}$, $\mathbf{f}^\tau_{j+1}$ and $\mathbf{f}^\tau_{j+1}$ as follows. First we generate $\mathbf{v}_1, \dots, \mathbf{v}_{j-1}$ and the unit vectors $\mathbf{b}_j, \mathbf{b}_{j+1} \in B$. We select two orthogonal random unit vectors $\mathbf{c}_j$ and $\mathbf{d}_{j+1}$ in C. We choose ϕ from the appropriate distribution, and obtain $\mathbf{c}_{j+1}$ by rotating $\mathbf{c}_j$ in the plane of $\mathbf{c}_j$ and $\mathbf{d}_{j+1}$ by an angle of ϕ in the direction of $\mathbf{d}_{j+1}$.

From this, we can construct the vectors $\mathbf{b}'_j$, $\mathbf{b}'_{j+1}$ and $\mathbf{d}_{j+1}$. Note that the distribution of the pair $(\mathbf{c}_j, \mathbf{d}_{j+1})$ is the same as that of $(\mathbf{d}_j, \mathbf{c}_{j+1})$, independently of the angle ϕ: uniform on all orthogonal pairs of unit vectors in C.

Now we come to the scalars: the pair of coefficients $(\cos(\theta_j), \cos(\theta_{j+1})\cos(\phi))$ has a density function that is positive on the unit square, and so does the pair $(\cos(\theta_j)\cos(\phi), \cos(\theta_{j+1}))$. It follows that these pairs are mutually absolutely continuous, which implies that so are $(\mathbf{f}^\sigma_j, \mathbf{f}^\sigma_{j+1}$ and $(\mathbf{f}^\tau_j, \mathbf{f}^\tau_{j+1})$.

Thus we know that the distributions of $(\mathbf{f}^\sigma_1, \dots, \mathbf{f}^\sigma_{j+1})$ and $(\mathbf{f}^\tau_1, \dots, \mathbf{f}^\tau_{j+1})$ are mutually absolute continuous. The remaining vectors $\mathbf{f}^\sigma_k$ and $\mathbf{f}^\tau_k$ are generated in the same way from these two sequences, and hence the distributions of $\mathbf{f}^\sigma$ and $\mathbf{f}^\tau$ are mutually absolutely continuous.

This completes the proof of Lemma 5.2.5 and of Theorem 5.2.2. □

5.3 Faithfulness

An orthogonal representation is *faithful* if adjacent nodes are represented by non-orthogonal vectors. It is probably difficult to decide whether a given graph has a faithful orthogonal representation in a given dimension. In other words, we do not know how to determine the minimum dimension of a faithful orthogonal representation. What we can do is to give a not-quite-trivial lower bound, and—as an application of the results in the previous section—a nontrivial upper bound.

Let us start with some examples.

Example 5.3.1 We have constructed orthogonal representations from cliques covering the nodes of a graph (Example 5.0.5). These representations are far from being faithful in general. However, using a family of cliques covering all *edges*, we get a faithful representation. Let $\{B_1, \dots, B_k\}$ be any family of complete subgraphs covering all edges of a graph G. We can construct the vector labeling $\mathbf{u}$ in $\mathbb{R}^k$, defined by $(\mathbf{u}_i)_j = \mathbb{1}(i \in B_j)$. Then $\mathbf{u}$ is a faithful orthogonal representation of G: if i and j are adjacent nodes, then they are contained in one of the cliques B_i, and hence $\mathbf{u}_i^\top \mathbf{u}_j > 0$; if they are nonadjacent, then their supports are disjoint, and so $\mathbf{u}_i^\top \mathbf{u}_j = 0$.

As a special case, we get a faithful orthogonal representation in dimension $n-1$ of the path P_n with n nodes. The path P_n has no faithful orthogonal representation in lower dimension. Indeed, in any faithful orthogonal representation $(\mathbf{v}_i : \ i \in V(P_n)$, for every $1 \le k \le n-1$, the vector $\mathbf{v}_{k+1}$ is orthogonal to $\mathbf{v}_1, \dots, \mathbf{v}_{k-1}$, but not orthogonal to $\mathbf{v}_k$, which implies that $\mathbf{v}_k$ is linearly independent of $\mathbf{v}_1, \dots, \mathbf{v}_{k-1}$. Hence the vectors $\mathbf{v}_1, \dots, \mathbf{v}_{n-1}$ are linearly independent, and so the ambient space must have dimension at least $n-1$. ♦

This example motivates the following lower bound on the dimension of a faithful orthogonal representation. It is trivial that $\alpha(G)$ is a lower bound on the dimension of any orthogonal representation by nonzero vectors (faithful or not). To strengthen this bound in the case of faithful representations, we say that a subset $S \subseteq V(G)$ is *almost stable*, if for every connected subgraph H of $G[S]$ with at least one edge there is a node $j \in V(G) \setminus V(H)$ that is adjacent to exactly one node in $V(H)$. (The node j may or may not belong to S.) It is trivial that every stable set is almost stable.

Proposition 5.3.2 *The dimension of any faithful orthogonal representation of a graph G is at least as large as its largest almost stable set.*

Proof. Let $\mathbf{v}$ be a faithful orthogonal representation of G in dimension d, and let S be an almost stable set. We prove by induction on $|S|$ that the vectors in $\mathbf{v}(S)$ are linearly independent. If S is a stable subset of $V(G)$, then this is trivial. Else, let H be a connected component of $G[S]$ with at least one edge, and let $j \notin V(H)$ be a node adjacent to a unique node $i \in V(H)$. The set $S' = S \setminus \{i\}$ is almost stable, and hence the vectors in $\mathbf{v}(S')$ are linearly independent by the induction hypothesis.

Next, observe that the vectors $\mathbf{v}(V(H) \setminus \{i\})$ cannot span $\mathbf{v}_i$; this follows from the fact that $\mathbf{v}_j$ is orthogonal to each of these vectors, but not to $\mathbf{v}_i$. Since the vectors $\mathbf{v}_k$ with $k \in S \setminus V(H)$ are all orthogonal to every vector $\mathbf{v}_k$ with $k \in V(H)$, it follows that S' cannot span $\mathbf{v}_i$. This proves that $\mathbf{v}(S) = \mathbf{v}(S') \cup \{\mathbf{v}_i\}$ is a linearly independent set of vectors. □

From the results in Section 5.2.1, it is easy to derive an upper bound on the minimum dimension of faithful orthogonal representations.

Proposition 5.3.3 *Every $(n-d)$-connected graph on n nodes has a faithful orthogonal representation in $\mathbb{R}^d$.*

Proof. It suffices to show that in a G-orthogonalized random representation, the probability that two adjacent nodes are represented by orthogonal vectors is zero. By the Lemma 5.2.5, it suffices to prove this for the G-orthogonalization based on an ordering starting with these two nodes. But then the assertion is obvious. $\square$

Using that a graph with minimum degree D cannot contain a complete bipartite graph with more than $2D$ nodes, Proposition 5.3.3 implies the following bound.

Corollary 5.3.4 *If the maximum degree of a graph G is D, then G has a faithful dual orthogonal representation in $2D$ dimensions.*

It is conjectured that the bound on the dimension can be improved to $D+1$. Proposition 5.3.3 shows that the conjecture is true if we strengthen its assumption by requiring that G is $(n-D)$-connected.

We conclude this section with two examples illustrating the use (and insufficiency) of the upper and lower bounds above.

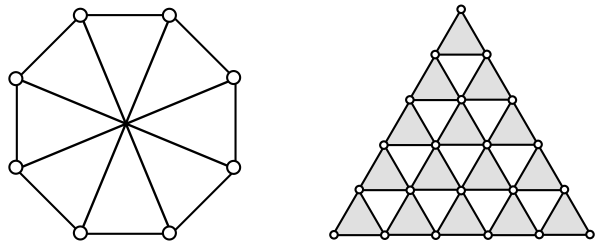

(a) (b)

Figure 5.5: The graph V_8 and the triangular grid Δ_6.

Example 5.3.5 Let V_8 denote the graph obtained from the cycle C_8 by adding its longest diagonals (Figure 5.5(a)). This graph is 3-connected, and hence by Proposition 5.3.3, it has a faithful orthogonal representation in $\mathbb{R}^5$. It does not have a faithful orthogonal representation in $\mathbb{R}^4$. On the other hand, it does not have an almost stable set of size 5. (The last two facts need some case analysis, not reproduced here.) So the lower bound in Proposition 5.3.2 does not always give the right value. ♦

Example 5.3.6 (Triangular grid I) Consider the triangular grid Δ_k with $k > 2$ nodes along the bottom line Figure 5.5(b)). This graph has $n = \binom{k+1}{2}$ nodes. The dark triangles form a family of cliques covering all edges, so the construction in Example 5.3.1 yields a faithful orthogonal representation in dimension is $n-k$, the number of dark triangles. This is much better than the upper bound from Proposition 5.3.3, which is $n-2$.

This dimension is in fact the smallest possible; this follows by Proposition 5.3.2, since it is easy to see that the set of nodes above the bottom line is almost stable, and hence the minimum dimension of a faithful representation is $n - k$. ♦