Northeastern University

Research Experience for Undergraduates

Summer 2024

A Bi-Lipschitz Coorbit Embedding of Point Clouds

Authors

Ang Barrett, Dan Kiem, and Mitchell Wadas

Faculty Advisor: Prof. Harm Derksen

Graduate Mentor: Shengnan Huang

Abstract

In this paper, we study clouds of n points in $\mathbb{R}^d$, which can be identified with matrices in $\mathbb{R}^{d\times n}$ modulo the action of S_n by column permutation. We construct a bi-Lipschitz embedding of this orbit space into $\mathbb{R}^m$ where $m = O(dn^2)$. We show that our construction is a coorbit embedding (Definition 3.7) and provide sufficient conditions on the parameters (templates) to ensure that this map is injective and bi-Lipschitz. Finally, we perform numerical experiments to identify a set of parameters which yield small distortion in practice. The constructive nature, modest embedding dimension, and small distortion of our map make it suitable for applications.

Contents

1 Introduction and Background

1.1 Motivation

Given a Euclidean space and a group of orthogonal transformations, we wish to embed the orbit space into another Euclidean space by a bi-Lipschitz map. We are motivated by the analysis of data that resides in such an orbit space. For example:

> **Graphs**: A graph with n vertices can be represented by its adjacency matrix, an element of $\mathbb{R}^{n\times n}$, but this representation is unique only up to the relabelling of nodes, corresponding to conjugation by a permutation matrix. We may therefore identify each graph not with an adjacency matrix, but rather an orbit under the conjugation action of the symmetric group S_n.
>
> **Point Clouds**: A point cloud is a finite, unordered collection of vectors. A cloud of n points in $\mathbb{R}^d$ may be represented by an element of $\mathbb{R}^{d\times n}$ with uniqueness up to a permutation of columns. Hence, point clouds are better described as orbits under the permutation action of S_n.

In each of the examples above, data resides in a vector space V equipped with an inner product, and as a consequence, a metric, which provides a notion of distance between vectors. However, the natural metric is not suitable for studying data points which we naturally identify with elements lying in the same orbit, as vectors which represent the same object may have a non-zero distance. To correctly measure distance, we must treat the data as residing in the orbit space corresponding to the group action. However, most existing data analysis algorithms assume that data lies in a Euclidean space. To employ such algorithms, we wish to embed the orbit space into Euclidean space.

We will study this question for finite groups acting on Euclidean vector spaces by orthogonal transformations. To begin, we will define the metric quotient V/G whose points are the G-orbits of vectors in V. We will then discuss some invariant theory based approaches and their limitations. In Section 3, we will discuss a new class of maps, called coorbit embeddings, which enable the construction of bi-Lipschitz embeddings in many cases. In Section 4 we construct an explicit embedding of the orbit space of point clouds modulo column permutation into Euclidean space which is bi-Lipschitz in the quotient metric. Finally, in Section 5, we discuss the complexity and distortion of this embedding. The small distortion and modest embedding dimension of this map make it suitable for a wide variety of applications with point cloud data.

1.2 Metric Quotient

When V is an inner product space, the automorphisms of V are orthogonal transformations, and we wish to consider linear representations to the orthogonal group $\mathrm{O}(V)$. We say that a group G acts **orthogonally** on V if the action is given by a representation $\rho : G \to \mathrm{O}(V)$. The inner product on V induces a norm $||v|| = \sqrt{\langle v, v\rangle}$, and a metric $d_V(u, v) = ||v - u||$ on the space. Orthogonal transformations are isometries with respect to this metric:

$$||Tv - Tu||^2 = \langle Tv - Tu, Tv - Tu\rangle = \langle v - u, v - u\rangle = ||v - u||^2.$$

Hence, when G acts orthogonally, it acts by isometries. In this situation, we define the **orbit space** V/G to be the set of orbits $V/G = \{G \cdot v : v \in V\}$. We equip V/G with the metric

$$d_{V/G} : V/G \times V/G \to \mathbb{R}, \quad d_{V/G}(G \cdot v, G \cdot u) = \min\{d_V(g \cdot v, h \cdot u) : g, h \in G\}.$$

Since G acts by isometries, we have $d_V(g\cdot v, h\cdot u) = d_V(v, g^{-1}h\cdot u) = d_V(h^{-1}g\cdot v, u)$ and consequently,

$$d_{V/G}(G \cdot v, G \cdot u) = \min_{g\in G} d_V(v, g \cdot u) = \min_{h\in G} d_V(v, h \cdot u).$$

Lemma 1.1. *When $G \leq \mathrm{O}(d)$ is a finite subgroup acting orthogonally on $V = \mathbb{R}^d$, the orbit space V/G is a metric space with respect to $d_{V/G}$.*

Proof. Let $u, v, w \in V$. It is clear that $d_{V/G}(G \cdot v, G \cdot v) = 0$, and that $d_{V/G}(G \cdot v, G \cdot u) > 0$ when $G \cdot v \neq G \cdot u$. We have already verified that this map is symmetric, so all that remains is the triangle inequality. Let g, h be the elements of G such that $d_{V/G}(G \cdot u, G \cdot v) = d_V(u, g \cdot v)$ and $d_{V/G}(G \cdot v, G \cdot w) = d_V(v, h \cdot w)$. We know $d_{V/G}(G \cdot u, G \cdot w) \leq d_V(u, gh \cdot w)$, and we have

$$d_V(u, gh \cdot w) \leq d_V(u, g \cdot v) + d_V(g \cdot v, gh \cdot w) = d_V(u, g \cdot v) + d_V(v, h \cdot w).$$

This right hand side is simply $d_{V/G}(G \cdot u, G \cdot v) + d_{V/G}(G \cdot v, G \cdot w)$ so we are done. □

1.3 Bi-Lipschitz Embeddings

We are interested in embedding the orbit space into some other metric space by a map which preserves the structure given by the quotient metric, namely, the distance between orbits. We formalize this notion with the next definition.

A map $f : X \to Y$ between metric spaces (X, d_X) and (Y, d_Y) is **bi-Lipschitz** if it is both **Lipschitz** and **lower Lipschitz**. More concretely, f is bi-Lipschitz if there exist positive constants $\alpha,\ \beta \in \mathbb{R}$ such that

$$\alpha \cdot d_X(x_1, x_2) \leq d_Y(f(x_1), f(x_2)) \leq \beta \cdot d_X(x_1, x_2) \tag{1}$$

for all $x_1,\ x_2 \in X$. We call α the lower Lipschitz constant, and β the (upper) Lipschitz constant. Notice that the existence of a positive lower Lipschitz constant implies that f is injective: if there is a non-zero α satisfying the above then $f(x_1) = f(x_2)$ implies $\alpha \cdot d_X(x_1, x_2) = 0$, and since $\alpha \neq 0$, we must have $x_1 = x_2$. If f is bi-Lipschitz, we define the **distortion** of f by

$$\text{dist}(f) = \inf \left\{ \frac{\beta}{\alpha} : \alpha,\ \beta \text{ satisfy (1)} \right\}.$$

The distortion serves as a means of quantifying the quality of the embedding. As such, we wish to minimize it to the optimal case where $\alpha = \beta$ and consequently, $\text{dist}(f) = 1$. When the distortion is 1, embedding f preserves the structure of the quotient metric.

Fix a natural number m. If $f : V \to \mathbb{R}^m$ is constant on G-orbits in V, then f factors through $f^{\downarrow} : V/G \to \mathbb{R}^m$ defined by $G \cdot v \mapsto f(v)$, meaning $f = f^{\downarrow} \circ \pi$ where $\pi : V \to V/G$ is the projection $v \mapsto G \cdot v$. We follow [4] and reserve this notation to denote the factor map on the metric quotient induced by a G-invariant map. We summarize the situation with the commutative diagram below.

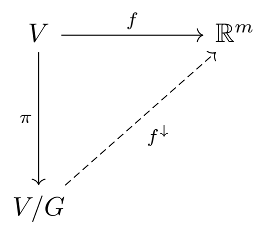

If $V = \mathbb{R}^d$ and $G \leq \text{O}(d)$ is a finite subgroup, then results from [5] and [1] guarantee the existence of a bi-Lipschitz embedding of the orbit space V/G into $\mathbb{R}^m$. However, relatively little is known about how to construct these embeddings. This invites the following question:

Question 1.1. Given a finite group G acting orthogonally on a Euclidean vector space V, can we construct a G-invariant map $f : V \to \mathbb{R}^m$ so that $f^{\downarrow} : V/G \to \mathbb{R}^m$ is bi-Lipschitz?

Question 1.1 will be the central focus of this paper, and we will answer it explicitly in Section 4 for a specific group action by constructing a bi-Lipschitz embedding of the orbit space of point clouds (matrices modulo column permutation) into $\mathbb{R}^m$.

2 Invariant Theory

Given a group acting on a real vector space V by linear representation $\rho : G \to \text{GL}(V)$, we define the **dual representation** $\rho^* : G \to \text{GL}(V^*)$ by $\rho^*(g)l = l \circ \rho(g^{-1})$ for a linear function l in the dual space. This definition respects the natural pairing $\langle \cdot, \cdot \rangle : V^* \times V \to K$ given by $\langle l, v \rangle = l(v)$ so that $\langle \rho^*(g)l, \rho(g)v \rangle = \langle l, v \rangle$ for all $l \in V^*$ and $v \in V$. The action of G on V^* induced by the dual representation is $g \cdot l = l \circ \rho(g^{-1})$, which has the effect $(g \cdot l)(v) = l(g^{-1} \cdot v)$ for $v \in V$.

We may regard $\mathbb{R}[x_1, \cdots, x_n]$, the ring of polynomials in n indeterminates with real coefficients, as a ring of functions on $V = \mathbb{R}^n$ by substituting the indeterminates for a basis of V^*. We use the notation $\mathbb{R}[V]$ to denote this ring of polynomial functions. The action of G on V^* induces an action on $\mathbb{R}[V]$ by $g \cdot f = f \circ \rho(g^{-1})$ for $f \in \mathbb{R}[V]$. In essence we have, $(g \cdot f)(v) = f(g^{-1} \cdot v)$ for $v \in V$.

Invariant theory is the study of polynomials $f \in \mathbb{R}[V]$ which remain invariant under this group action, i.e. those $f \in \mathbb{R}[V]$ for which $g \cdot f = f$ for all $g \in G$. The set of invariant polynomials (also referred to as invariants) forms a subring of $\mathbb{R}[V]$ which is denoted $\mathbb{R}[V]^G$. When G is finite, a result due to Emmy Noether states that the the invariant ring is finitely generated as an algebra over $\mathbb{R}$

(see Proposition 3.0.6 in [6]). Moreover, for finite groups, there exist algorithms for computing a generating set (see for example [6, Chapter 3]).

2.1 Separating Sets

Definition 2.1. Let G be a finite group acting linearly on $V = \mathbb{R}^d$. A subset $S \subseteq \mathbb{R}[V]^G$ is called **separating** if for any two vectors $u, v \in V$, we have $G \cdot v \neq G \cdot u$ implies that there exists $f \in S$ such that $f(v) \neq f(u)$.

Separating sets (for finite group actions) may be equivalently characterized by saying that $S \subseteq \mathbb{R}[V]^G$ is separating if, for all $u, v \in V$,

$$G \cdot v = G \cdot u \iff f(v) = f(u) \quad \forall f \in S.$$

The left implication follows immediately from the definition, and the converse follows from the fact that $f \in S$ is G-invariant. This property makes our interest in separating sets clear: if $S \subseteq \mathbb{R}[V]^G$ is a separating set, then defining

$$\phi : V \to \mathbb{R}^{|S|}, \quad \phi(v) = (f(v) : f \in S)$$

induces a map $\phi^{\downarrow} : V/G \to \mathbb{R}^{|S|}$ which is injective. We call a map which is injective on the orbit space **separating**, and we say that it separates G-orbits in V. In summary, every separating set defines an embedding V/G to Euclidean space which is given component-wise by the members of the set. We will use this property in Algorithm 2.1 to construct a separating embedding of the orbit space.

Proposition 2.1. *If $S \subseteq \mathbb{R}[V]^G$ is a finite generating set, then S is also a separating set.*

Proof. Take $u, v \in V$ with distinct orbits, and let v_i denote the ith component of v. Let $\mathbf{e}_1, \cdots, \mathbf{e}_d$ be the standard basis for V, and $x_1, \cdots, x_d$ the dual basis. Define the polynomial

$$p_v = \prod_{\sigma \in G} \left(\sum_{i=1}^{d} (v_i - \sigma \cdot x_i)^2 \right) \in \mathbb{R}[V]^G.$$

Evaluating p_v at some $w \in V$, we have $p_v(w) = \prod_{\sigma \in G} ||v - \sigma^{-1} \cdot w||^2$, and consequently $p_v(w) = 0$ exactly when v and w lie in the same orbit. Certainly $p_v(v)$ is zero, and p_v is G-invariant, so if $S = \{f_1, \cdots, f_m\}$, then we may write p_v as the finite sum

$$p_v = \sum_{\alpha} c_\alpha f_1^{\alpha_1} \cdots f_k^{\alpha_m} \quad \text{where } \alpha \in \mathbb{N}^k.$$

If $f_1, \cdots, f_d$ all agree at u and v, then we must have $p_v(v) = p_v(u) = 0$, implying that $G \cdot u = G \cdot v$ and contradicting our assumption. Therefore, we conclude there must be some $f \in S$ for which $f(u) \neq f(v)$, making S a separating set. □

This discussion leads to the following algorithm for constructing embeddings of V/G to $\mathbb{R}^m$. Calling this procedure an algorithm is justified by the existence of algorithms for computing a generating set in the cases we consider.

Algorithm 2.1. Given a finite group $G \leq \mathrm{O}(d)$ acting orthogonally on $V = \mathbb{R}^d$, compute a generating set $S = \{f_1, \cdots, f_m\}$ and take $\phi : V \to \mathbb{R}^m$ to be map $\phi(v) = (f_1(v), \cdots, f_m(v))$. Then $\phi^{\downarrow} : V/G \to \mathbb{R}^m$ is injective.

2.2 Differentiability

In the last subsection, we described an algorithm which constructs an embedding $V/G \to \mathbb{R}^n$ when $V = \mathbb{R}^d$ and $G \leq \mathrm{O}(d)$ is finite. Unfortunately, the maps obtained in this way will rarely be bi-Lipschitz. In fact, in many cases the bi-Lipschitz condition requires that a factor map $f^\downarrow$ is induced by a non-differentiable invariant $f : V \to \mathbb{R}^m$.

Let G be a finite group acting orthogonally on $V = \mathbb{R}^d$, and suppose $f : V \to \mathbb{R}^m$ is G-invariant. We say a map $f : V \to W$ between Euclidean spaces is differentiable at a point $x \in V$ if the Jacobian (matrix) of f exists at x. In essence, f is differentiable at x if all its first order partial derivatives exist at x.

Theorem 2.1 (Theorem 21 in [4])**.** *If $x \in V$ is fixed by some non-identity element of G and f is differentiable at x, then $f^\downarrow : V/G \to \mathbb{R}^m$ is not lower Lipschitz.*

As a result, maps given component wise by elements of $\mathbb{R}[V]^G$ cannot yield a lower Lipschitz embedding unless the action of G on V is free. The requirement that G act freely on V is quite restrictive, so we desire invariants which are not differentiable everywhere, and therefore might be used to induce a bi-Lipschitz embedding of the orbit space. This is a notable departure from classical invariant theory, which focuses primarily on the algebraic structure of polynomial invariants.

3 Coorbit Embedding

3.1 Max Filtering

In the previous section, we concluded that polynomial invariants fail to yield bi-Lipschitz embeddings of the orbit space since they are differentiable everywhere, and in particular, at fixed points of the group action. Therefore, we desire an alternative family of invariants which are not differentiable at any points fixed by any element in the group. One such family of maps, called max filters, was introduced in [5] in which it was shown that there exist separating invariants constructed from max filters for all finite subgroups of the orthogonal group. In this subsection we will describe the max filter bank (Definition 3.3) and discuss its efficacy as a theoretical and practical tool for constructing bi-Lipschitz embeddings of orbit spaces.

Definition 3.1 (Definition 1 in [5])**.** Given a real inner product space V and a group of linear isometries G, the **max filtering map** $\langle\langle \cdot, \cdot \rangle\rangle : V/G \times V/G \to \mathbb{R}$ is defined by

$$\langle\langle G \cdot x, G \cdot y \rangle\rangle := \sup_{p \in G\cdot x, q \in G \cdot y} \langle p, q \rangle.$$

We are interested in finite groups, so we may replace the supremum with a maximum. Additionally, we assume G acts orthogonally on V, so this definition can be simplified to

$$\langle\langle G \cdot x, G \cdot y \rangle\rangle = \max_{\sigma \in G} \langle \sigma \cdot x, y \rangle = \max_{\sigma \in G} \langle x, \sigma \cdot y \rangle.$$

Definition 3.2 (Definition 3 in [5])**.** Given a **template** $z \in V$, we refer to $\langle\langle G \cdot z, \cdot \rangle\rangle : V/G \to \mathbb{R}$ as the corresponding **max filter**.

The max filter is a map $V/G \to \mathbb{R}$ which is parameterized by a vector $z \in V$; we follow [5] in referring to these parameters as "templates".

Definition 3.3 (Definition 3 in [5]). Given a (possibly infinite) sequence $\{z_i\}_{i\in I}$ of templates in V, the corresponding **max filter bank** is $\Phi : V/G \to \mathbb{R}^I$ defined by $\Phi(G\cdot x) = \{\langle\langle G\cdot z_i, G\cdot x\rangle\rangle\}_{i\in I}$.

We would like to construct a injective max filter bank $\Phi : V/G \to \mathbb{R}^m$ which embeds the orbit space into a low dimensional space (i.e. we desire m to be small) and may be evaluated efficiently. It turns out that for $V = \mathbb{R}^d$ and finite $G \leq \mathrm{O}(d)$, there exists a finite collection of templates which yield an injective max filter bank.

Definition 3.4. We define the semi-algebraic sets of $\mathbb{R}^d$ as the smallest family of sets in $\mathbb{R}^d$ which contain the algebraic sets $\{x \in \mathbb{R}^d : p(x) = 0\}$ and sets of the form $\{x \in \mathbb{R}^d : q(x) > 0\}$ for $p, q \in \mathbb{R}[x_1, \cdots, x_d]$, and is closed under complementation, finite unions, and finite intersections. Cf. [3, Section 2.3].

Definition 3.5. A semi-algebraic homeomorphism $f : S \to T$ between semi-algebraic sets S and T is a homeomorphism whose graph $\{(x, f(x)) : x \in S\} \subseteq S \times T$ is a semi-algebraic set.

Every semi-algebraic set $S \subseteq \mathbb{R}^d$ may be decomposed into the disjoint union of finitely many semi-algebraic sets, each of which is semi-algebraically homeomorphic to the hyper-cube $(0,1)^i \subseteq \mathbb{R}^d$ for some $i \leq d$ (with the convention that $(0,1)^0$ is a point) [3, Theorem 5.19]. The dimension of semi-algebraic set S is the largest i which appears in this decomposition [3, Propositions 5.27, 5.28].

Theorem 3.1 (Corollary 13 in [5]). *Consider any finite subgroup $G \leq \mathrm{O}(d)$. For generic templates $z_1, \cdots, z_m \in \mathbb{R}^d$, the max filter bank $x \mapsto \{\langle\langle G\cdot z_i, G\cdot x\rangle\rangle\}_{i=1}^m$ separates G-orbits in $\mathbb{R}^d$ provided $m \geq 2d$.*

Theorem 3.1 guarantees the existence of a separating max filter bank, but it does not suggest a way to select templates so that the associated max filter bank is separating. The proof of this theorem demonstrates that the subset $\mathcal{Z} \subseteq (\mathbb{R}^d)^n$ where a collection of templates fails to separate G-orbits is semi-algebraic of dimension at most $dn-1$, but in practice, this set is difficult to describe.

Let $S^{d-1} = \{v \in \mathbb{R}^d : ||v|| = 1\}$ be the unit $(d-1)$-sphere in $\mathbb{R}^d$, and let $\mathrm{Unif}(S^{d-1})$ be the uniform distribution on the unit $(d-1)$-sphere.

Theorem 3.2 (Theorem 18 in [5]). *Fix a finite group $G \leq \mathrm{O}(d)$ of order N and select m at least $12N^2 d \log\left(\frac{2}{\delta}+1\right)$ where*

$$\delta := \left(\frac{\pi}{128N^4}\cdot\frac{1}{2d+3\log(4N^2)}\right)^{1/2}.$$

Draw independent random vectors $z_1, \cdots, z_m \sim \mathrm{Unif}(S^{d-1})$. With probability at least $1-e^{-m/(12N^2)}$, the max filter bank $\Phi : \mathbb{R}^d/G \to \mathbb{R}^m$ with templates $\{z_i\}_{i=1}^m$ has lower Lipschitz bound δ and upper Lipschitz bound $m^{1/2}$.

We may summarize this result by stating that sufficiently many templates, drawn at random from the unit sphere, yield a bi-Lipschitz embedding of the orbit space with positive probability. This provides a method to construct the bi-Lipschitz maps of interest, but this approach is limited by the dimension of the co-domain as well as the computational complexity, both of which grow with the order of the group.

If we fix d and regard the number of templates, m, required to apply the theorem above as a function of the group order, we can show that m grows faster than N^2 since the logarithmic term is increasing in N. Moreover, $\log(2/\delta + 1) \geq 1$ for d and N at least 1 (which is always the case), so it follows that $m \geq N^2$. This makes the above result somewhat impractical for groups of large order. For example, consider the orbit space of point clouds (that is $\mathbb{R}^{d\times n}$ modulo the action of S_n by column permutation). In this case, the embedding dimension must be larger than $(n!)^2$, which is untenable particularly for applications involving point cloud data.

3.2 Coorbit Embedding

The authors of [1] generalize the max filter map by taking subsets of the sorted list of inner products rather than just the largest, as in the max filter map. Importantly, these generalized maps, called **coorbit embeddings**, are necessarily bi-Lipschitz when they are separating. For a natural number n, let $[n]$ denote the set $\{1, 2, 3, \cdots, n\}$.

Definition 3.6. Let $\texttt{sort} : \mathbb{R}^d \to \mathbb{R}^d$ be the operator that sorts a vector's components in non-increasing order. In other words, $\texttt{sort}$ returns a vector with the same components as its input so that the output, call it $v = (v_1, \cdots, v_d)$, satisfies $v_1 \geq \cdots \geq v_d$.

Let G be a finite group of order N acting orthogonally on $V = \mathbb{R}^d$. Fix a natural number k, and let $\boldsymbol{w} = (w_1, \cdots, w_k)$ belong to V^k. For $i \in [k]$ and $j \in [N]$, let $\xi_{ij} : V \to \mathbb{R}$ be the map which sends $x \in V$ to the jth coordinate of the vector

$$\texttt{sort}(\langle x, \sigma \cdot w_i\rangle : \sigma \in G).$$

Fix an index set $S \subseteq [N] \times [k]$. For $i \in [k]$, let S_i be the set $S_i = \{j \in [N] : (j, i) \in S\}$, and define

$$\xi_i : V \to \mathbb{R}^{|S_i|}, \quad v \mapsto (\xi_{ij}(v) : j \in S_i).$$

Definition 3.7. The **coorbit embedding** associated to $\boldsymbol{w} \in V^k$ and $S \subseteq [N] \times [k]$ is the map

$$\xi_{\boldsymbol{w},S} : V \to \mathbb{R}^{|S|}, \quad v \mapsto (\xi_1(v), \cdots, \xi_k(v)).$$

Theorem 3.3 (Theorem 2.1 in [1])**.** *Let G be a finite subgroup of $\mathrm{O}(d)$ with $|G| = N$. Fix $\boldsymbol{w} \in V^k$ and $S \subseteq [N] \times [k]$. If the coorbit embedding $\xi^{\downarrow}_{\boldsymbol{w},S} : V/G \to \mathbb{R}^{|S|}$ is injective, then it is bi-Lipschitz.*

As a consequence, we see that the max filter bank, which can be obtained as a coorbit embedding by taking $S = \{1\} \times [k]$, is necessarily bi-Lipschitz when it is injective. Hence, we can revisit Theorem 3.1, which states that $m = 2d$ generic templates separate orbits, and conclude that for finite groups $G \leq \mathrm{O}(d)$ acting orthogonally on $\mathbb{R}^d$ there exists a collection of $m = 2d$ templates in $\mathbb{R}^d$ for which the associated max filter bank is a bi-Lipschitz embedding of the orbit space $\mathbb{R}^d/G$ into $\mathbb{R}^m$.

4 Quotients by Column Permutation

4.1 A Point Cloud Embedding

The results from Section 3 demonstrate that bi-Lipschitz coorbit embeddings exist for all finite groups acting orthogonally on $\mathbb{R}^d$, but they do not suggest a practical way to construct these maps. In this section, we will investigate the action of S_n on $V = \mathbb{R}^{d\times n}$ by column permutation and construct an explicit bi-Lipschitz embedding of the orbit space V/S_n into a Euclidean space of dimension $O(dn^2)$.

Let $\rho : S_n \to \mathrm{O}(n)$ be the defining representation of S_n, and define $\sigma \cdot M = M\rho(\sigma)^T$. It is easy to verify that this action is an orthogonal transformation of $V \cong \mathbb{R}^{dn}$, so there exists a bi-Lipschitz embedding of the orbit space into $\mathbb{R}^{2dn}$ by Theorem 3.1 and Theorem 3.3. For matrices $M = [m_1, \cdots, m_n] \in \mathbb{R}^{d\times n}$, the action of S_n has the effect of permuting columns according to

$$\sigma \cdot [m_1, \cdots, m_n] = \left[m_{\sigma^{-1}(1)}, \cdots, m_{\sigma^{-1}(n)}\right].$$

We would like to construct a bi-Lipschitz embedding of V/S_n to a Euclidean space of small dimension. To do so, we will construct a coorbit embedding and find explicit templates which make this map injective.

Definition 4.1. Given $z \in \mathbb{R}^d$, define $\phi_z : V \to \mathbb{R}^n$ by $\phi_z(M) = \mathtt{sort}(M^T z)$. For a finite collection $\boldsymbol{z} = \{z_1, \cdots, z_k\} \subseteq \mathbb{R}^d$, define $\Phi_{\boldsymbol{z}} : V \to \mathbb{R}^{kn}$ by

$$\Phi_{\boldsymbol{z}}(M) = (\phi_{z_1}(M), \cdots, \phi_{z_k}(M)).$$

Each component of $\Phi_{\boldsymbol{z}}$ is constant on S_n-orbits, and therefore, $\Phi_{\boldsymbol{z}}$ induces a well defined map $\Phi_{\boldsymbol{z}}^{\downarrow} : V/S_n \to \mathbb{R}^{kn}$. We will now show that $\Phi_{\boldsymbol{z}}$ is an instance of coorbit embedding. To begin, for each $z \in \boldsymbol{z}$ let $z' = [z, 0, \cdots, 0] \in V$ be the matrix which has the vector z as its first and only non-zero column. Notice that, for $M = [m_1, \cdots, m_n] \in V$, we have $\langle M, z'\rangle = \mathrm{tr}(M^T z) = \langle m_1, z\rangle$.

Proposition 4.1. *Take $\boldsymbol{w} = (z_1', \cdots, z_k') \in V^k$, and let $S = \{(n-1)!, 2(n-1)!, \cdots, n!\} \times [k]$. The coorbit embedding $\xi_{w,S}$ is equal to $\Phi_{\boldsymbol{z}}$.*

Proof. Both maps are given by the concatenation of k smaller functions; for $\Phi_{\boldsymbol{z}}$ these are the functions ϕ_{z_i} for $z_i \in \boldsymbol{z}$, and for $\xi_{w,S}$ the component functions are ξ_i for $1 \le i \le k$. Hence, it suffices to show that $\phi_{z_i} = \xi_i$. Let $M = [m_1, \cdots, m_n]$ belong to V, and observe that

$$\xi_i(M) = (\xi_{i,(n-1)!}(M), \xi_{i,2(n-1)!}(M), \cdots, \xi_{i,n!}(M)).$$

Recall $\xi_{i,j}(M)$ is the jth component of the sorted vector $\mathtt{sort}(\langle M, \sigma \cdot z_i'\rangle : \sigma \in S_n)$ which is equivalent to $\mathtt{sort}(\langle m_{\sigma(1)}, z_i\rangle : \sigma \in S_n)$. Since the letter 1 has a stabilizer isomorphic to S_{n-1}, each value $\langle m_j, z_i\rangle$ for $1 \le j \le n$ will appear $(n-1)!$ times in the sorted vector. Therefore, by taking $S_i = \{(n-1)!, 2(n-1)!, \cdots, n!\}$, we find

$$\xi_i(M) = \mathtt{sort}(\langle m_{\sigma(1)}, z_i\rangle : \sigma H \in S_n/H) = \mathtt{sort}(\langle m_j, z_i\rangle : 1 \le j \le n).$$

where $H \cong S_{n-1}$ is the stabilizer of 1 in S_n. The right-most vector is simply $\phi_{z_i}(V) = \mathtt{sort}(M^T z_i)$, so we conclude $\xi_{w,S} = \Phi_{\boldsymbol{z}}$. $\square$

Corollary 4.1. *If $\boldsymbol{z} \subseteq \mathbb{R}^d$ is a finite collection of templates such that $\Phi_{\boldsymbol{z}}^{\downarrow}$ is injective on the orbit space, then $\Phi_{\boldsymbol{z}}^{\downarrow}$ is bi-Lipschitz.*

This invites the question, how can we choose $\boldsymbol{z} \subseteq \mathbb{R}^d$ so that $\Phi_{\boldsymbol{z}}^{\downarrow}$ is injective? To investigate, let M and W belong to V, and let m_j denote the jth column of M. Let $\mathbf{e}_1, \cdots, \mathbf{e}_d$ be the standard basis for $\mathbb{R}^d$, and $x_1, \cdots, x_d$ the dual basis. Associate to each column m_j the linear form

$$m_j^* := \sum_{i=1}^{d} m_{ij} x_i \in \mathbb{R}[x_1, \cdots, x_d].$$

Notice that, for $v \in \mathbb{R}^d$, we have $m_j^*(v) = \langle m_j, v\rangle$. Now, to each matrix $M \in V$, associate the polynomial

$$P_M := \prod_{j=1}^{n} \left(t - m_j^*\right) \in \mathbb{R}[x_1, \cdots, x_d][t] \tag{2}$$

Lemma 4.1. *Two matrices in V have the same associated polynomial if and only if they lie in the same S_n-orbit.*

Proof. Let $M, W \in V$. The polynomial ring $\mathbb{R}[x_1, \cdots, x_d][t]$ is a unique factorization domain, and each term in the product is irreducible so $P_M = P_W$ exactly when there is permutation $\sigma \in S_n$ so that $m_j^* = w_{\sigma(j)}^*$ for all $1 \leq j \leq n$. This equality implies $m_j = w_{\sigma(j)}$, so we have $M = \sigma^{-1} \cdot W$. The converse follows immediately from the fact that multiplication is commutative in the polynomial ring $\mathbb{R}[x_1, \cdots, x_d][t]$. □

Let e_j denote the jth elementary symmetric polynomial in n indeterminates, and observe that the polynomial P_M expands as

$$P_M = t^n - e_1\left(m_1^*, \cdots, m_n^*\right) t^{n-1} + \cdots + (-1)^n e_n\left(m_1^*, \cdots, m_n^*\right). \tag{3}$$

Hence, two matrices M and W in V lie in the same S_n-orbit if and only if

$$e_j\left(m_1^*, \cdots, m_n^*\right) = e_j\left(w_1^*, \cdots, w_n^*\right) \quad \text{for all } 1 \leq j \leq n. \tag{4}$$

The polynomials appearing in (4) are homogeneous of degree j, so they belong to the homogeneous piece $\mathbb{R}[x_1, \cdots, x_d]_{(j)}$, which is a real vector space of dimension $\binom{j+d-1}{d-1}$.

Definition 4.2. A set of points $X \subseteq \mathbb{R}^d$ is **unisolvent** for a vector space $F \subseteq \mathbb{R}[x_1, \cdots, x_d]$ if the zero polynomial is the only polynomial which vanishes at every point in X.

If a set of points is unisolvent for $F \subseteq \mathbb{R}[x_1, \cdots, x_d]$, then any two polynomials in F which agree at all of these points must be identical (because their difference is the zero polynomial). Hence, solutions to the Lagrange interpolation problem are unique in F (see Chapter 2 in [10] for more information). This explains the etymology of "unisolvent".

Lemma 4.2. *If $\boldsymbol{z} \subseteq \mathbb{R}^d$ is unisolvent for $F_j = \mathbb{R}[x_1, \cdots, x_d]_{(j)}$ for all $j \in [n]$, then $\Phi_{\boldsymbol{z}}^{\downarrow}$ is injective.*

Proof. It suffices to show that $\Phi_{\boldsymbol{z}}(M) = \Phi_{\boldsymbol{z}}(W)$ implies $S_n \cdot M = S_n \cdot W$. If $\Phi_{\boldsymbol{z}}(M) = \Phi_{\boldsymbol{z}}(W)$ then $\phi_{z_i}(M) = \phi_{z_i}(W)$ for all $1 \leq i \leq k$. Since $m_j^*(z_i) = \langle m_j, z_i\rangle$, it follows that

$$\prod_{j=1}^{n}(t - m_j^*(z_i)) = \prod_{j=1}^{n}(t - w_j^*(z_i)) \in \mathbb{R}[t]$$

which implies

$$e_j(m_1^*, \cdots, m_n^*)(z_i) = e_j(w_1^*, \cdots, w_n^*)(z_i) \quad \text{for all } 1 \leq j \leq n.$$

The polynomials $e_j(m_1^*, \cdots, m_n^*)$ and $e_j(w_1^*, \cdots, w_n^*)$ belong to F_j, so by the assumption that $\boldsymbol{z}$ is unisolvent, the above implies

$$e_j(m_1^*, \cdots, m_n^*) = e_j(w_1^*, \cdots, w_n^*) \quad \text{for all } 1 \leq j \leq n.$$

Consequently, $P_M = P_W$ and we conclude that M and W lie in the same orbit. □

Given a basis $\{u_1, \cdots, u_N\}$ for a vector space $F \subseteq \mathbb{R}[x_1, \cdots, x_d]$, we can determine whether a set of points $z_1, \cdots, z_k \in \mathbb{R}^d$ is unisolvent by looking at the kernel of the matrix

$$\begin{bmatrix} u_1(z_1) & \cdots & u_N(z_1) \\ u_1(z_2) & \cdots & u_N(z_2) \\ \vdots & \ddots & \vdots \\ u_1(z_k) & \cdots & u_N(z_k) \end{bmatrix}$$

The set is unisolvent exactly when this matrix has a trivial kernel. From this perspective, it is clear that we must choose a set points in $\mathbb{R}^d$ so that the matrix has rank N, and therefore, we must choose $k \geq N$. Of course, this is not a sufficient condition, but it places a lower bound on k which makes our results in the next two subsections optimal.

4.2 The 2-Dimensional Case

To determine when $\boldsymbol{z}$ is unisolvent, we start by focusing on the case $d = 2$. A homogeneous polynomial $f \in \mathbb{R}[x_1, x_2]_{(n)}$ of degree n satisfies $f(\lambda v) = \lambda^n f(v)$, so we may regard the zeros of f as points on the real projective line

$$\mathbb{P}^1 = (\mathbb{R}^2 \setminus \{0\})/\sim \quad \text{, where } x \sim y \iff x = \lambda y \text{ for some } \lambda \in \mathbb{R} \setminus \{0\}.$$

We use the notation $[v_1 : v_2] \in \mathbb{P}^1$ to denote the equivalence class $[v]$ for a vector $v = (v_1, v_2) \in \mathbb{R}^2$. Every point in $\mathbb{P}^1$, spare the horizontal subspace, has a canonical representative of the form $[\beta : 1]$. The horizontal subspace has representative $[1 : 0]$, and we call this point ∞. We call a projective point $[v_1 : v_2] \in \mathbb{P}^1$ a root of f if $f(v_1, v_2) = 0$.

Lemma 4.3. *A non-zero homogeneous polynomial $f \in \mathbb{R}[x_1, x_2]_{(n)}$ of degree n can have at most n distinct projective roots.*

Proof. A typical $f \in \mathbb{R}[x, y]$ can be written $f = a_n x^n + a_{n-1} x^{n-1} y + \cdots + a_1 x y^{n-1} + a_0 y^n$ where some of the coefficients may be zero. If f has a projective root at $[\beta : 1] \in \mathbb{P}^1$, then

$$f(\beta, 1) = a_n \beta^n + a_{n-1} \beta^{n-1} + \cdots + a_1 \beta^1 + a_0 = 0$$

meaning β is a root of the univariate polynomial $f(x, 1)$. Hence, f may have at most n finite projective roots. If f has a root at ∞, then $f(1, 0) = a_n = 0$. Therefore, $f(x, 1)$ is univariate of degree at most $n - 1$, and it follows that f may have at most $n - 1$ additional roots in $\mathbb{P}^1$. □

Corollary 4.2. *Any collection of $n + 1$ vectors in $\mathbb{R}^2$ corresponding to distinct projective points is unisolvent for the homogeneous piece $\mathbb{R}[x_1, x_2]_{(j)}$ for all $1 \leq j \leq n$.*

The top elementary symmetric function is homogeneous of degree n, so we make take $\boldsymbol{z}$ to be any collection of $n + 1$ points in $\mathbb{R}^2$, none of which lie on the same line through the origin, and conclude that $\Phi_{\boldsymbol{z}}^{\downarrow}$ is injective by Lemma 4.2. In summary, we have proven the following theorem:

Theorem 4.1. *Let $V = \mathbb{R}^{2 \times n}$. If $\boldsymbol{z} \subseteq \mathbb{R}^2$ is a collection of $n + 1$ vectors, none of which lie on the same line through the origin, then $\Phi_{\boldsymbol{z}}^{\downarrow} : V/S_n \to \mathbb{R}^{n(n+1)}$ is bi-Lipschitz.*

4.3 The d-Dimensional Case

To generalize this result beyond the case $d = 2$, we will transform the coefficients appearing in (3) to homogeneous bivariate polynomials of larger degree, and show that equality of these transformed coefficients implies that matrices M and W lie in the same S_n orbit. To begin, define

$$\omega_k := \binom{d-1}{k}^{1/2} x_1^{d-k-1} x_2^k \in \mathbb{R}[x_1, x_2]_{(d-1)}.$$

We obtain a ring homomorphism $\eta : \mathbb{R}[x_1, \cdots, x_d] \to \mathbb{R}[x_1, x_2]$ by substituting $x_i = \omega_{i-1}$.

Lemma 4.4. *Let $v_1, \cdots, v_n \in \mathbb{R}^d$. Any set of $n(d-1)+1$ vectors in $\mathbb{R}^2$, none of which lie on the same line through the origin, is unisolvent for the polynomials $\eta(e_j(v_1^*, \cdots, v_n^*))$ for $1 \le j \le n$.*

Proof. Notice that $\eta(e_j(v_1^*, \cdots, v_n^*))$ is bivariate and homogeneous of degree $j(d-1)$ for all $j \in \mathbb{N}$. Hence, the result follows from Corollary 4.2. $\square$

Let $M = [m_1, \cdots, m_n]$ and $W = [w_1, \cdots, w_n]$ belong to V. If we take $z_1, \cdots, z_k$ to be a collection of $k = n(d-1)+1$ vectors in $\mathbb{R}^2$ satisfying the hypothesis of the above lemma, then for all $1 \le j \le n$ we have

$$\eta(e_j(m_1^*, \cdots, m_n^*))(z_i) = \eta(e_j(w_1^*, \cdots, w_n^*))(z_i) \quad \text{for all } 1 \le i \le k$$
$$\implies \eta(e_j(m_1^*, \cdots, m_n^*)) = \eta(e_j(w_1^*, \cdots, w_n^*)).$$

Since η is a ring homomorphism, $\eta(e_j(m_1^*, \cdots, m_n^*)) = e_j(\eta(m_1^*), \cdots, \eta(m_n^*))$, so invoking the identity (3) again, we have

$$\prod_{j=1}^{n} \left(t - \eta(m_j^*)\right) = \prod_{j=1}^{n} \left(t - \eta(w_j^*)\right) \in \mathbb{R}[x_1, x_2][t].$$

Since $\mathbb{R}[x_1, x_2][t]$ is a unique factorization domain, this implies that there is some $\sigma \in S_n$ such that

$$\eta(m_j^*) = \sum_{i=1}^{d} m_{ij}\eta(x_i) = \sum_{i=1}^{d} w_{ij}\eta(x_i) = \eta\left(w_{\sigma(j)}^*\right).$$

The image of $x_1, \cdots, x_d$ under η is linearly independent, so it follows that $m_j = w_{\sigma(j)}$, and consequently, we have $M = \sigma^{-1} \cdot W$. In summary, taking $z_1, \cdots, z_k$ to be a collection of $k = n(d-1)+1$ vectors in $\mathbb{R}^2$ which define distinct projective points,

$$\eta(e_j(m_1^*, \cdots, m_n^*))(z_i) = \eta(e_j(w_1^*, \cdots, w_n^*))(z_i) \quad \text{for all } 1 \le j \le n, \text{ and } 1 \le i \le k$$

implies that M and W lie in the same S_n-orbit. Evaluating $\eta(e_j(m_1^*, \cdots, m_n^*))$ at z_i is equivalent to evaluating $e_j(m_1^*, \cdots, m_n^*)$ at $\psi(z_i)$ where $\psi : \mathbb{R}^2 \to \mathbb{R}^d$ is given by

$$(v_1, v_2) \mapsto (\omega_0, \cdots, \omega_{d-1}), \quad \omega_k := \binom{d-1}{k}^{1/2} v_1^{d-k-1} v_2^k.$$

Hence, we may take $\boldsymbol{z} = \{\psi(z_1), \cdots, \psi(z_k)\}$ and find that $\Phi_{\boldsymbol{z}} : V/S_n \to \mathbb{R}^{kn}$ is bi-Lipschitz. In conclusion, we have proven our central result:

Theorem 4.2. *Let $z_1, \cdots, z_k$ be a collection of $k = n(d-1)+1$ vectors in $\mathbb{R}^2$, no two of which lie on the same line through the origin. Take $\boldsymbol{z} = (\psi(z_i))_{1\leq i\leq k}$, then $\Phi_{\boldsymbol{z}}^{\downarrow} : V/S_n \to \mathbb{R}^{kn}$ is bi-Lipschitz.*

Remark 4.1. The coefficients on ω_k were chosen so that ψ sends vectors on the unit circle in $\mathbb{R}^2$ to vectors on the unit $(d-1)$-sphere in $\mathbb{R}^d$. Observe that, for $v \in \mathbb{R}^2$, applying the binomial theorem yields

$$||\psi(v)||^2 = \sum_{k=0}^{d-1} \binom{d-1}{k} v_1^{2(d-k-1)} v_2^{2k} = (v_1^2 + v_2^2)^{d-1} = ||v||^{2(d-1)}.$$

In many cases, the Lipschitz constants depend on the norm of the templates selected; to make analysis of these constants simpler, we would like to choose templates on the unit $(d-1)$-sphere. Additionally, in order to make general statements about the distortion of our map, we desire that the norm of our templates does not change as d does.

5 Discussion

In the previous section, we constructed an explicit embedding of the orbit space of point clouds into Euclidean space, and gave sufficient conditions on the templates to ensure that our map is injective and bi-Lipschitz. In this section we will analyze our construction, beginning with the embedding dimension and computational complexity, and concluding with a discussion of distortion. The results of our analysis suggest that $\Phi_{\boldsymbol{z}}$ is well suited for applications with point cloud data.

5.1 Complexity

For applications involving point cloud data, we would like the embedding dimension (i.e. the dimension of the co-domain as an $\mathbb{R}$-vector space) to be small. In addition, we must be able evaluate our map efficiently, ideally in polynomial time or better.

Recall that our map requires $k = n(d-1)+1$ templates, and embeds the orbit space into $\mathbb{R}^{kn}$. Therefore, the embedding dimension is $m = n^2(d-1)+n = O(dn^2)$. Compare this to the max filter bank obtained by selecting random templates (see Theorem 3.2) where the embedding dimension must be at least $(n!)^2$.

We will now analyze the time complexity of $\Phi_{\boldsymbol{z}}$ for constant d. For $M \in \mathbb{R}^{d\times n}$, the product $M^T z$ may be computed in linear time. The complexity of the component map ϕ_z for some template z is therefore dominated by the sorting step, which can be performed in $O(n \log n)$ time. It follows that the time complexity of $\Phi_{\boldsymbol{z}}$ is $kO(n \log n) = O(n^2 \log n)$ when d is constant.

5.2 Distortion

We now turn our attention to investigating the distortion of our embedding. In particular, we would like to answer the following question:

Question 5.1. How can we choose a collection $\boldsymbol{z} \subseteq \mathbb{R}^d$ of $k = n(d-1)+1$ templates in order to achieve small distortion?

To investigate, we performed some numerical experiments. Previous discussion in Section 4 suggests that we may select $\boldsymbol{z}$ to be any collection of k vectors in $\mathbb{R}^2$ such that no two lie on the same line through the origin. This sufficient condition leaves a great deal of choice, so we would like to compare different methods of selecting templates.

Method 5.1. The first method we consider is choosing k equally spaced vectors from the upper half of the unit circle. More precisely, let $R \in \mathbb{R}^{2\times 2}$ be the matrix which rotates vectors counterclockwise through π/k, and let $z_1 = \mathbf{e}_1$ be the first standard basis vector in $\mathbb{R}^2$. Select $\boldsymbol{z} = \{R^0 z_1, R^1 z_1, \cdots, R^{k-1} z_1\}$.

Intuition suggests that this choice will yield small distortion, but to evaluate $\boldsymbol{z}$, we perform the following experiment: Let $\mathcal{N}_t(0, I)$ be the standard t-variate normal distribution, and let $x \sim \mathcal{N}_{dn}(0, I)$ be a random variable representing a matrix in $\mathbb{R}^{d\times n}$. Draw $r = 100$ samples from x, say $x_1, \cdots, x_r$, and compute the $\binom{100}{2} = 4950$ distinct pairwise distances in both the orbit space and the target space. With this data, compute the **realized distortion**, which we define by

$$\mathrm{dist}^*(\Phi_{\boldsymbol{z}}) = \frac{\beta^*}{\alpha^*}, \text{ with } \beta^* := \max_{i\neq j} \frac{||\Phi_{\boldsymbol{z}}(x_i) - \Phi_{\boldsymbol{z}}(x_j)||}{d_{V/G}(G\cdot x_i, G\cdot x_j)}, \text{ and } \alpha^* := \min_{i\neq j} \frac{||\Phi_{\boldsymbol{z}}(x_i) - \Phi_{\boldsymbol{z}}(x_j)||}{d_{V/G}(G\cdot x_i, G\cdot x_j)}.$$

We start with $V = \mathbb{R}^{2\times 3}$ and perform the experiment with $\boldsymbol{z}$ selected according to Method 5.1. In this case, we attain a realized distortion $\mathrm{dist}^*(\Phi_{\boldsymbol{z}}) \approx 1.66$. We visualize the results of this experiment below in Figure 1.

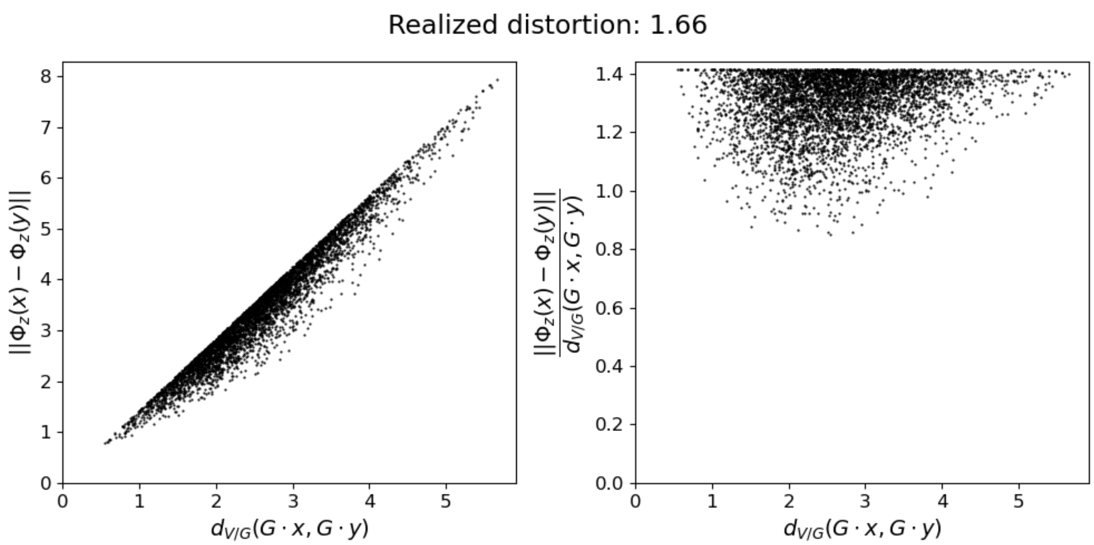

Figure 1: Distortion attained with templates selected by Method 5.1.

For many applications, we may be satisfied with distortion less than 2, so this first method seems promising. However, we desire a benchmark to which we can compare this result.

Method 5.2. We select templates at random from $\mathcal{N}_d(0, I)$ while verifying that no two of these templates lie on the same line through the origin.

Performing the same experiment as described above (again with $V = \mathbb{R}^{2\times 3}$), Method 5.2 yields a realized distortion of $\mathrm{dist}^*(\Phi_{\boldsymbol{z}}) \approx 6.83$. We visualize the results of this method below in Figure 2.

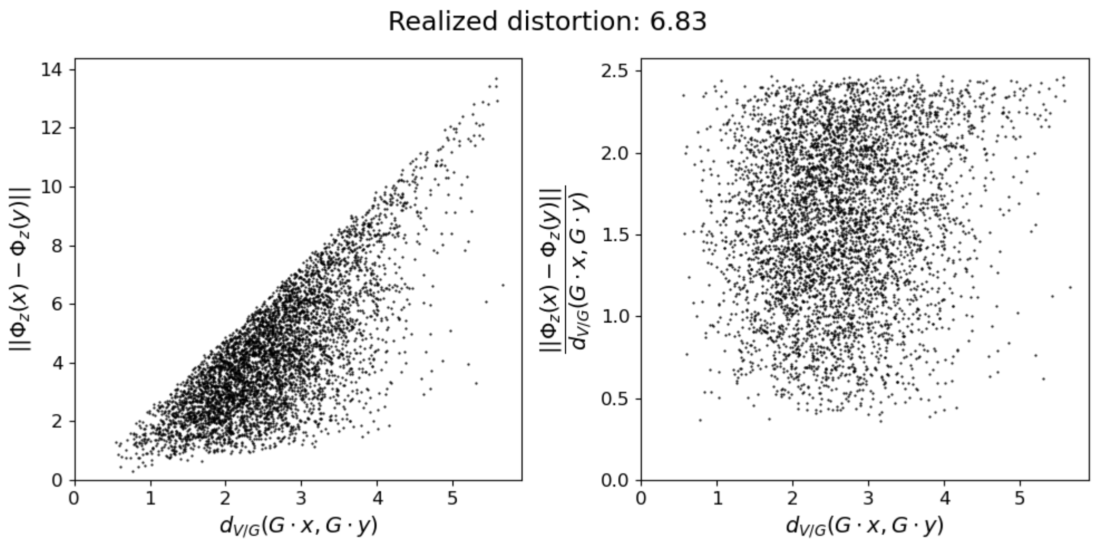

Figure 2: Distortion attained with templates selected by Method 5.2.

6 Conclusion

In Section 4, we constructed an explicit coorbit embedding $\mathbb{R}^{d\times n}/S_n \to \mathbb{R}^m$ where $m = O(dn^2)$, and provided sufficient conditions on the templates to ensure that this map is injective and bi-Lipschitz. In Section 5, we described a method of choosing templates which yields small distortion in practice; when coupled with the small embedding dimension, and modest computational complexity of our map, this makes our construction suitable for applications with point cloud data.

To continue this investigation, we would like to find upper and lower Lipschitz constants for the map $\Phi_z^{\downarrow}$, which would provide a theoretical upper bound on the distortion. Recent developments in the theory of coorbit embedding, namely [8], may assist in this analysis. In addition, it would be useful to study the behavior of the distortion (realized or otherwise) as n and d vary. We are currently studying this question through numerical experiments like those described in Section 5.

Finally, we would like to generalize this construction to other group actions, and we are in the process of investigating this possibility. Progress will likely factor through an associated polynomial (as in Equation 2) whose coefficients separate orbits for the action of interest. We may be able to employ [6, Theorem 3.9.13] to construct such polynomials for arbitrary finite group actions.

Acknowledgements

We would like to thank our advisor Prof. Harm Derksen, and our mentor Shengnan Huang, for their guidance and support throughout this project. We also want to thank the organizers of the Northeastern University Summer Math Research Program; they are Dr. Hunter Dinkins, Dr. Josh Wen, Prof. Harm Derksen, Prof. Iva Halacheva, Prof. Valerio Toledano Laredo, and Prof. Xuwen Zhu. The authors were supported by the following sources of funding:

- NSF-RTG grant "Algebraic Geometry and Representation Theory at Northeastern University" (DMS-1645877),
- NSF grant "Invariant Theory and Complexity Theory for Quiver Representations and Tensors" (DMS-2147769),
- The Northeastern University Department of Mathematics,
- The Northeastern University College of Science.

References

[1] Radu Balan and Efstratios Tsoukanis. *G-Invariant Representations using Coorbits: Bi-Lipschitz Properties*. 2023. arXiv: 2308.11784 [math.RT].

[2] Radu Balan and Efstratios Tsoukanis. *G-Invariant Representations using Coorbits: Injectivity Properties*. 2023. arXiv: 2310.16365 [math.RT].

[3] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*. Second. Vol. 10. Algorithms and Computation in Mathematics. Springer-Verlag, Berlin, 2006, pp. x+662. ISBN: 978-3-540-33098-1.

[4] Jameson Cahill, Joseph W. Iverson, and Dustin G. Mixon. *Towards a Bilipschitz Invariant Theory*. 2024. arXiv: 2305.17241 [math.FA].

[5] Jameson Cahill, Joseph W. Iverson, Dustin G. Mixon, and Daniel Packer. *Group-invariant Max Filtering*. 2022. arXiv: 2205.14039 [cs.IT].

[6] Harm Derksen and Gregor Kemper. *Computational invariant theory*. Vol. 130. Encyclopedia of Mathematical Sciences. Springer, Heidelberg, 2015, pp. xxii+366. ISBN: 978-3-642-07796-8.

[7] Dustin G. Mixon and Yousef Qaddura. *Injectivity, Stability, and Positive Definiteness of Max Filtering*. 2022. arXiv: 2212.11156 [math.FA].

[8] Dustin G. Mixon and Yousef Qaddura. *Stable Coorbit Embeddings of Orbifold Quotients*. 2024. arXiv: 2403.14042 [math.FA].

[9] Mara D. Neusel. *Invariant Theory*. Vol. 36. Student Mathematical Library. American Mathematical Society, Providence, RI, 2015, pp. viii+314. ISBN: 978-0-8218-4132-7.

[10] Holger Wendland. *Scattered data approximation*. Vol. 17. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, Cambridge, 2005, pp. x+336. ISBN: 978-0521-84335-5.