Petros T. Boufounos, Shantanu Rane, Hassan Mansour

**Abstract** Low-dimensional embeddings have emerged as a key component in modern signal processing theory and practice. In particular, embeddings transform signals in a way that preserves their geometric relationship but makes processing more convenient. The literature has, for the most part, focused on lowering the dimensionality of the signal space while preserving distances between signals. However, there has also been work exploring the effects of quantization, as well as on transforming geometric quantities, such as distances and inner products, to metrics easier to compute on modern computers, such as the Hamming distance.

Embeddings are particularly suited for modern signal processing applications, in which the fidelity of information represented by the signals is of interest, instead of the fidelity of the signal itself. Most typically, this information is encoded in the relationship of the signal to other signals and templates, as encapsulated in the geometry of the signal space. Thus, embeddings are very good tools to capture the geometry, while reducing the processing burden.

In this chapter, we provide a concise overview of the area, including foundational results and recent developments. Our goal is to expose the field to a wider community, to provide, as much as possible, a unifying view of the literature, and to demonstrate the usefulness and applicability of the results.

Shantanu Rane

Hassan Mansour

Petros T. Boufounos

Mitsubishi Electric Research Laboratories, 201 Broadway, Cambridge, MA 02139, USA, e-mail: petrosb@merl.com

Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304, USA e-mail: srane@parc.com

Mitsubishi Electric Research Laboratories, 201 Broadway, Cambridge, MA 02139, USA, e-mail: mansour@merl.com

# **1** Introduction

Signal representation theory and practice has primarily focused on how to best represent or approximate signals, while incurring the smallest possible distortion. Advances such as frames, compressive sensing and sparse approximations have all been applied in improving the representation accuracy or sampling complexity using a fidelity metric as the principal figure of merit. On the other hand, as computation becomes more prevalent, signal representations are increasingly important in inference and estimation applications. Such applications typically exploit the geometry of the signal space, usually captured mathematically by norms and inner products. In these cases, the representation should faithfully preserve the geometry of the signal space, but not necessarily the signals themselves.

This chapter explores embeddings as a signal representation mechanism that preserves the geometry of the signal space. Embeddings are transformations from one signal space to another—the embedding space—which exactly or approximately preserve signal geometry. The use of an embedding is beneficial if the transformation provides some convenience in its use. For example, the embedding space might have significantly lower dimensionality than the signal space, might allow for easier computation of certain quantities, or might enable efficient transmission by quantizing in the embedding space.

In this chapter, we explore several aspects of embedding design. We start with the foundational work by Johnson and Lindenstrauss [40], and continue with more recent developments. We describe embeddings that preserve distances, inner products, and angles between signals, while reducing the dimension and the bit-rate. We also describe embedding design strategies, both data-agnostic and universal, as well as learning-based and data-driven. Our discussion also explores the effect of quantization, which becomes necessary when the embeddings are used to reduce the bit-rate of the representation.

Our goal is to expose the field to a wide community and show that embeddings are essential data processing tools. In our exposition, we attempt to provide, as much as possible, a unifying view of the literature. However, we remark that recent advances have reinvigorated research in this area, often making such unification elusive.

## 1.1 Notation

In the remainder of the chapter, we use regular typeface, e.g., x and y, to denote scalar quantities. Lowercase boldface such as **x** denotes vectors and uppercase boldface such as **A** denotes matrices. The m<sup>th</sup> element of vector **x** is denoted using  $x_m$ . Functions are denoted using regular lowercase typefaces, e.g.,  $g(\cdot)$ . Unless explicitly noted, all functions are scalar functions of one variable. In abuse of notation, a vector input to such functions, e.g.,  $g(\mathbf{x})$  means that the function is applied element-

2



Fig. 1 Distance-preserving embeddings approximately preserve a function  $g(\cdot)$  of the distance, allowing distances to be computed in a space that (typically) has fewer dimensions or has other desirable properties.

wise to all the elements of **x**. Sets and vector spaces are denoted using calligraphic fonts, e.g.,  $\mathcal{W}$ ,  $\mathcal{S}$ .

# 1.2 Outline

The next section describes distance-preserving embeddings. Starting with general definitions and foundational results, the section explores embedding design strategies—both data-agnostic and data-driven—and discusses the nature of distancepreserving guarantees. Section 3 examines embeddings that preserve angles and inner products, including kernel inner products. Quantization strategies and the effects of quantization on the embedding guarantees are discussed in Sec. 4. Section 5 provides a higher-level discussion and concludes the chapter.

# 2 Preserving Distances

The best-known embeddings preserve the geometry of the space by preserving the distance between signals. In this section, we examine distance-preserving embeddings, and explore some ways to design their distance-preserving properties.

## 2.1 Randomized Linear Embeddings

An embedding is a transformation of a set of signals in a high-dimensional space to a (typically) lower-dimensional one such that some aspects of the geometry of the set are preserved, as depicted in Fig. 1. Since the set geometry is preserved, distance computations can be performed directly on the low-dimensional—and often low bit-rate—embeddings, rather than the underlying signals. For the purposes of this chapter, we define an embedding as follows. **Definition 1.** A function  $f : \mathscr{S} \to \mathscr{W}$  is a  $(g, \delta, \varepsilon)$  embedding of  $\mathscr{S}$  into  $\mathscr{W}$  if, for all  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$ , it satisfies

$$(1-\delta)g\left(d_{\mathscr{S}}(\mathbf{x},\mathbf{x}')\right) - \varepsilon \le d_{\mathscr{W}}\left(f(\mathbf{x}),f(\mathbf{x}')\right) \le (1+\delta)g\left(d_{\mathscr{S}}(\mathbf{x},\mathbf{x}')\right) + \varepsilon.$$
(1)

In this definition,  $g: \mathbb{R} \to \mathbb{R}$  is an invertible function mapping distances in  $\mathscr{S}$  to distances in  $\mathscr{W}$  and  $\delta$  and  $\varepsilon$  quantify, respectively, the multiplicative and the additive ambiguity of the mapping. We will often refer to  $g(\cdot)$  as the distance map and to  $f(\cdot)$  as the embedding map. In most known embeddings, such as the ones discussed in this section, the distance map is the identity g(d) = d or a simple scaling. The similarity metrics  $d_{\mathscr{S}}(\cdot, \cdot)$  and  $d_{\mathscr{W}}(\cdot, \cdot)$  are typically distances, but could also be correlations, divergences, or other functions capturing signal geometry and similarity<sup>1</sup>.

The best known embeddings are the Johnson-Lindenstrauss (JL) embeddings [40]. These are functions  $f : \mathscr{S} \to \mathbb{R}^M$  from a finite set of signals  $\mathscr{S} \subset \mathbb{R}^N$  to an *M*-dimensional vector space such that, given two signals  $\mathbf{x}$  and  $\mathbf{x}'$  in  $\mathscr{S}$ , their images satisfy:

$$(1-\delta) \|\mathbf{x} - \mathbf{x}'\|_2^2 \le \|f(\mathbf{x}) - f(\mathbf{x}')\|_2^2 \le (1+\delta) \|\mathbf{x} - \mathbf{x}'\|_2^2.$$
(2)

In other words, these embeddings preserve Euclidean, i.e.,  $\ell_2$ , distances of point clouds within a small factor, measured by  $\delta$ , and using the identity as a distance map.

In the context of Def. 1, a JL embedding is a  $(g_I, \delta, 0)$  embedding of squared Euclidean distances— $d_{\mathscr{S}}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2^2$  and  $d_{\mathscr{W}}(f(\mathbf{x}), f(\mathbf{x}')) = \|f(\mathbf{x}) - f(\mathbf{x}')\|_2^2$ —with an identity distance map  $g_I(d) = d$ . In this context, the JL theorem can be stated as:

**Theorem 1.** Given  $\delta \in (0,1)$  and a set  $\mathscr{S} \subset \mathbb{R}^N$  of  $\#(\mathscr{S}) = L$  points and  $M = O(\delta^{-2} \ln L)$ , there exists a Lipschitz map  $f : \mathbb{R}^N \to \mathbb{R}^M$  that is a  $(g_I, \delta, 0)$  embedding of  $\mathscr{S}$ , with  $g_I(d) = d$ ,  $d_{\mathscr{S}}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2^2$  and  $d_{\mathscr{W}}(f(\mathbf{x}), f(\mathbf{x}')) = \|f(\mathbf{x}) - f(\mathbf{x}')\|_2^2$ .

Johnson and Lindenstrauss demonstrated that a distance-preserving embedding, as described above, exists in a space of dimension  $M = O(\delta^{-2} \log L)$ , where *L* is the number of signals in  $\mathscr{S}$  (its cardinality) and  $\delta$  the desired tolerance in the embedding. Remarkably, *M* is independent of *N*, the dimensionality of the signal set  $\mathscr{S}$ . Subsequent work showed that it is straightforward to compute such embeddings using a linear mapping. In particular, the function  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ , where **A** is an  $M \times N$  matrix whose entries are drawn randomly from specific distributions, satisfies (2) for all  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$  with probability  $1 - c_1 e^{\log L - c_2 \delta^2 M}$ , for some universal constants  $c_1, c_2$ , where the probability is with respect to the measure of **A**. Commonly used distributions for the entries of **A** are i.i.d. Gaussian, i.i.d. Rademacher, or i.i.d.

4

<sup>&</sup>lt;sup>1</sup> Technically, we could incorporate  $g(\cdot)$  into  $d_{\mathscr{S}}(\cdot, \cdot)$  and remove it from this definition. However, we choose to make it explicit here and consider it a distortion to be explicitly analyzed. In an abuse of nomenclature, we generally refer to  $d(\cdot, \cdot)$  as distance, even if in some cases it is not strictly a distance metric but might be an inner product, or another geometric quantity of interest.

uniform [1, 25]. More recent work has shown that the embedding dimensionality  $M = O(\delta^{-2} \log L)$  is also necessary, making these constructions tight [38].

Most proofs involve constructing a randomized map such that (1) holds with very high probability on a pair of points  $\mathbf{x}, \mathbf{x}' \in \mathcal{S}$ . Using a concentration of measure argument, such as Hoeffding's inequality or a Chernoff bound, it can typically be shown that the guarantee fails with probability that decays exponentially with the number of measurements, i.e., with the dimensionality of the embedding space M =dim( $\mathcal{W}$ ). In other words, the embedding fails on a pair of points with probability bounded by  $\Omega(e^{-Mw(\delta,\varepsilon)})$ , where  $w(\delta,\varepsilon)$  is an increasing function of  $\varepsilon$  and  $\delta$  that quantifies the concentration of measure exhibited by the randomized construction.

Once the embedding guarantee is established for a pair of signals, a union bound or chaining argument can be used to extend it to a finite set of signals. If the set  $\mathscr{S}$ is finite, containing *L* points, then the probability that the embedding fails is upper bounded by  $\Omega(L^2 e^{-M_W(\delta,\varepsilon)}) = \Omega(e^{2\log L - M_W(\delta,\varepsilon)})$ , which decreases exponentially with *M*, as long as  $M = O(\log L)$ .

More recently, in the context of compressive sensing, such linear embeddings have been shown to embed infinite sets of signals. For example, the restricted isometry property (RIP) is an embedding of *K*-sparse signals and has been shown to be achievable with  $M = O(K \log \frac{N}{K})$  [10, 23, 50]. A near equivalence of RIP with the JL lemma has also been established: an RIP matrix with its columns randomly multiplied with  $\pm 1$  will satisfy the JL lemma [41]. Similar properties have been shown for other signal set models, such as more general unions of subspaces and manifolds [9, 11, 12, 21, 28, 29, 50].

Typically, these generalizations are established by first proving that the embedding holds in a sufficiently dense point cloud on the signal set and exploiting linearity and smoothness to extend it to all the points of the set. The resulting guarantee uses the covering number of the set, i.e., its Kolmogorov complexity—instead of the number of points L—to measure the complexity of the set and determine the dimensionality required of the projection. A fairly general exposition of this approach, as well as generalizations for non-smooth embedding maps can be found in [20].

An alternative characterization of the complexity of  $\mathscr{S}$  is its Gaussian width.

**Definition 2.** Given a set  $\mathscr{S} \subseteq \mathbb{R}^N$ , the quantity

$$W(\mathscr{S}) = \mathbb{E}\left\{\sup_{\mathbf{x}\in\mathscr{S}}\mathbf{g}^{T}\mathbf{x}\right\},\tag{3}$$

where the expectation is taken over  $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})$  is called *the Gaussian width* of  $\mathcal{S}$ .

The Gaussian width of a set can sometimes be easier to characterize than its Kolmogorov complexity, although the latter can be bounded by the former [28].

Beyond the discussion above, in the remainder of this chapter, we defer on the rigorous development required to extend embedding guarantees to hold for infinite signal sets. Nevertheless, in many cases we will mention if such generalizations are possible or exist in the literature.

#### 2.2 Embedding Map Design

One of the key elements in the embedding definition (1) is the embedding map  $g(\cdot)$ . The JL guarantee in (2) implies an embedding map g(d) = d, that does not distort the distance measure. However, it is often desirable to introduce such distortions and understand their effect. For example, if the interest is in preserving only local distances, the distance map can be used to describe and characterize the distance preserving properties of the embedding [17, 19, 20].

A general approach to embedding design would use  $g(\cdot)$  to derive an embedding function  $f(\cdot)$ , possibly randomized, that achieves (1) given sufficient dimensionality of the embedding space  $\mathcal{W}$ . Unfortunately, such a design is still an open problem. Furthermore, an arbitrary  $g(\cdot)$  is not always possible. For example, any realizable  $g(\cdot)$  satisfies a generalized subadditivity property [20].

Instead, [20] demonstrates a general probabilistic approach to designing the embedding function  $f(\cdot)$  and deriving the embedding map. The mapping function takes the form  $\mathbf{y} = f(\mathbf{x}) = h(\mathbf{A}\mathbf{x} + \mathbf{w})$ , where the elements of  $\mathbf{A}$  are randomly chosen from an i.i.d. distribution and the elements of the dither  $\mathbf{w}$  are chosen from an i.i.d. distribution uniform in [0, 1). The embedding is designed through h(t), a bounded periodic scalar function with period 1, applied element-wise to its argument. The Fourier series coefficients of  $h(\cdot)$  are denoted using  $H_k$  and  $\bar{h} = \sup_t h(t) - \inf_t h(t)$ .

**Theorem 2** ([20], Thm. 4.1). Consider a set  $\mathscr{S}$  of Q points in  $\mathbb{R}^N$ , measured using  $\mathbf{y} = h(\mathbf{A}\mathbf{x} + \mathbf{w})$ , with  $\mathbf{A}$ ,  $\mathbf{w}$ , and h(t) as above. With probability greater than  $1 - e^{2\log Q - 2M\frac{\varepsilon^2}{h^4}}$  the following holds

$$g(d) - \varepsilon \le \frac{1}{M} \left\| \mathbf{y} - \mathbf{y}' \right\|_2^2 \le g(d) + \varepsilon$$
(4)

for all pairs  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$  and their corresponding measurements  $\mathbf{y}, \mathbf{y}'$ , where

$$g(d) = 2\sum_{k} |H_{k}|^{2} (1 - \phi_{l}(2\pi k|d)).$$
(5)

defines the distance map of the embedding.

In the theorem above,  $\phi_l(l|d)$  is a characteristic function depending on the density of **A**. For example, if the elements of **A** are drawn from an i.i.d. Normal distribution, then the characteristic function is  $\phi_l(\xi|d) = \phi_{\mathcal{N}(0,\sigma^2d^2)}(\xi) = e^{-\frac{1}{2}(\sigma d\xi)^2}$  and the distance map becomes

$$g(d) = 2\sum_{k} |H_{k}|^{2} \left(1 - e^{-2(\pi\sigma dk)^{2}}\right),$$
(6)

with *d* measuring the  $\ell_2$  distance.

If, instead, elements of **A** are drawn from an i.i.d. Cauchy distribution with zero location parameter and scale parameter  $\gamma$ , then the characteristic function is  $\phi_l(\xi|d) = e^{-\gamma d|\xi|}$  and the corresponding distance map is

$$g(d) = 2\sum_{k} |H_{k}|^{2} (1 - e^{-2\pi\gamma dk}),$$
(7)

with d in this case measuring the  $\ell_1$  distance.

The guarantee in Thm. 2 is about embedding the  $\ell_1$  or  $\ell_2$  distance into  $\ell_2^2$ . By taking the square root, the guarantee can be provided for embedding into  $\ell_2$  instead.

**Corollary 1** ( [20], Cor. 4.1). Consider the signal set  $\mathscr{S}$ , defined and measured as in Thm. 2. With probability greater than  $1 - e^{2\log Q - 2M\left(\frac{\varepsilon}{h}\right)^4}$  the following holds

$$\widetilde{g}(d) - \varepsilon \le \frac{1}{\sqrt{M}} \left\| \mathbf{y} - \mathbf{y}' \right\|_2 \le \widetilde{g}(d) + \varepsilon$$
(8)

for all pairs  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$  and their corresponding measurements  $\mathbf{y}, \mathbf{y}'$ , where  $\widetilde{g}(d) = \sqrt{g(d)}$ .

# 2.3 Distance-preserving properties of the map

Typically, when designing a distance map, it is desirable to understand how accurate the embedding is in representing distances. In particular, embedding guarantees, as stated above and in the literature, bound how much the distance in the embedding space might deviate from the true distance between the signals.

However, in practice, embeddings are used as a proxy for the true distance of the signals. Given two signals, **x** and **x'**, and their embedding distance,  $d_{\mathcal{W}}(f(\mathbf{x}), f(\mathbf{x}'))$ , a natural estimate of the true signal distance is [19,20]

$$\widetilde{d}_{\mathscr{S}} = g^{-1} \left( d_{\mathscr{W}} \left( f(\mathbf{x}), f(\mathbf{x}') \right) \right), \tag{9}$$

assuming  $g(\cdot)$  is differentiable. Thus, the approximation guarantee is often more useful when stated with respect to the estimate,  $\tilde{d}_{\mathcal{S}}$ .

$$\left| \widetilde{d}_{\mathscr{S}} - d_{\mathscr{S}}(\mathbf{x}, \mathbf{x}') \right| \lesssim \frac{\varepsilon + \delta d_{\mathscr{W}}(f(\mathbf{x}), f(\mathbf{x}'))}{g'\left(\widetilde{d}_{\mathscr{S}}\right)}.$$
(10)

An important component of this guarantee is its dependence on the gradient of the embedding map  $g'(\cdot)$  around the distance of the signals. In regions where the embedding map is flatter, the ambiguity is higher. In hindsight, this is expected: estimates of a variable observed through a non-linear map and observation ambiguity are less accurate at regions of the map that are flatter.

Figure 2 demonstrates this effect using a  $(g,0,\varepsilon)$  embedding as an example. The solid line in the left figure depicts the distance map  $g(\cdot)$ . The two dashed lines depict the upper and lower bounds of the guarantee, separated by  $\varepsilon$  above and below the distance map. In other words, the vertical ambiguity is constant across the range of  $d_{\mathscr{S}}$ . The figure also shows two example points on which the embedding distance is

Petros T. Boufounos, Shantanu Rane, Hassan Mansour



**Fig. 2** Effect of the gradient of the distance map on the distance ambiguity of the embedding. (left) Even though the vertical ambiguity is constant across the distance map, the corresponding horizontal ambiguity varies significantly, depending on the slope of the map. (right) Example embedding exhibiting similar behavior as described by the map on the left.

computed,  $d_{\mathscr{W}}$  and  $d'_{\mathscr{W}}$ . The corresponding estimates of the true signal distance are  $g^{-1}(d_{\mathscr{W}})$  and  $g^{-1}(d'_{\mathscr{W}})$ , respectively. However, the ambiguity of these estimates is significantly higher for  $d_{\mathscr{W}}$  than for  $d'_{\mathscr{W}}$ , because of the difference in slope of  $g(\cdot)$  at the corresponding points. Simulations using an actual embedding design exhibiting the same behavior are shown on the right hand side.

Embedding maps designed using the approach in Sec. 2.2 eventually saturate and become flat beyond a certain signal distance. Thus, the ambiguity becomes infinite; the embedding does not preserve distances beyond a range. Given an embedding map  $h(\cdot)$ , this range can be controlled by the scaling parameters of the distribution of **A**, such as  $\sigma$  and  $\gamma$  in (6) and (7), respectively. The same parameters also scale the gradient of the embedding, thus controlling the ambiguity, as described in (10). In other words, varying the scale parameters is equivalent to navigating a trade-off between smaller ambiguity while representing a smaller range of distances, and greater ambiguity while representing a larger range of distances. In fact, similar trade-offs are possible with any embedding function, simply by scaling the argument and replacing  $f(\mathbf{x})$  with  $f(a\mathbf{x})$  for any a > 0.

The distance preserving ambiguity described above characterizes distance preservation through  $g(\cdot)$  along a full range of distances. However, it is often sufficient to only guarantee the locality of the embedding, i.e., that small distances remain small and larger distances do not become too small. Recent work has attempted to define locality in the context of binary embeddings [46, Def. 2.3], as well as, implicitly, in the context of learning an embedding for classification [31, Eq. (6)]. In the same spirit, guarantees on using JL embeddings for classification have been recently established, assuming specific signal models. In particular, in [7] it is shown that separated convex ellipsoids remain separated when randomly projected to a space with sufficiently high dimensionality. However, an appropriate and useful characterization of locality is still a pending question.

One important property of the embeddings described so far is their universality. Their randomized construction does not take the data into account. The guarantees hold with very high probability on any set of points  $\mathscr{S}$  to be embedded, as long as the set complexity is known. Thus, there is no adversarial selection of the data for which the embedding will fail, assuming the data set is generated independently of the embedding. The next section explores embeddings designed while taking sample data into account, their advantages, as well as their disadvantages.

# 2.4 Learning the Embedding Map

A key advantage of the embeddings described above is their universality and the simplicity in computing them. However, it is often advantageous to tune the embedding to an application using available training data. The main assumption is that the training data is representative of the data to be observed by the application; tuning the embedding to the data should provide an embedding that performs well on all future data on which the embedding will be used.

Inspired by the JL lemma, recent work [31,54] demonstrates that given a set of *L* points  $\mathscr{S} = {\mathbf{x}_i \in \mathbb{R}^N, i = 1, ..., L}$  as training data, it is possible to formulate a convex optimization problem and determine a linear embedding map,  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$ , that preserves the squared Euclidean distance. The resulting map provides a  $(g_I, \delta, 0)$  embedding. The problem can be formulated to either minimize the dimensionality of the embedding space under a fixed multiplicative distortion  $\delta$  or minimize the distortion given a fixed embedding dimensionality.

In formulating the problem, the objects of interest are not the signals  $\mathbf{x}_i$  but their differences  $\mathbf{x}_i - \mathbf{x}_j$ . Thanks to the linearity of the map, to guarantee a  $1 \pm \delta$  multiplicative ambiguity it is sufficient to guarantee a  $\delta$  distortion of the normalized difference  $\frac{\mathbf{x}_i - \mathbf{x}_j}{\|\mathbf{x}_i - \mathbf{x}_j\|_2}$ . Thus, the formulation starts with the secant set

$$\mathscr{X} = \left\{ \mathbf{v}_{ij} = \frac{\mathbf{x}_i - \mathbf{x}_j}{\|\mathbf{x}_i - \mathbf{x}_j\|_2}, \, \mathbf{x}_i, \mathbf{x}_j \in \mathscr{S}, \, i \neq j \right\}$$
(11)

The map  $f(\mathbf{x}) = \mathbf{A}\mathbf{x}$  satisfies the guarantee for all  $\mathbf{v}_{ij} \in \mathscr{X}$  if

$$\left| \|\mathbf{A}\mathbf{v}_{ij}\|_{2}^{2} - \|\mathbf{v}_{ij}\|_{2}^{2} \right| \le \delta,$$

$$(12)$$

where, by construction,  $\|\mathbf{v}_{ij}\|_2^2 = 1$  for all i, j.

The squared norm can be expressed as a quadratic form  $\|\mathbf{A}\mathbf{v}_{ij}\|_{2}^{2} = \mathbf{v}_{ij}^{T}\mathbf{A}^{T}\mathbf{A}\mathbf{v}_{ij}$ which is linear in  $\mathbf{P} = \mathbf{A}^{T}\mathbf{A}$ . Furthermore, if  $\mathbf{A} \in \mathbb{R}^{M \times N}$ , then  $\mathbf{P}$ , which is positive semidefinite, has rank( $\mathbf{P}$ ) = M. Thus, the  $\mathbf{P}$  corresponding to the embedding that satisfies (12) for all pairs  $i \neq j$  with the minimum number of measurements can be found using the following optimization [31]:

$$\widehat{\mathbf{P}} = \arg\min_{\mathbf{P}^T = \mathbf{P} \succeq 0} \operatorname{rank}(\mathbf{P})$$
subject to  $|\mathbf{v}_{ij}^T \mathbf{P} \mathbf{v}_{ij} - 1| \le \delta$  for all  $i \ne j$ .
$$(13)$$

This is a non-convex and combinatorially complex program. To solve it, [31] proposes the relaxation of the rank using the nuclear norm, which results in the following polynomial-time semidefinite program:

$$\widehat{\mathbf{P}} = \arg\min_{\mathbf{P}^T = \mathbf{P} \succeq 0} \|\mathbf{P}\|_*$$
subject to  $|\mathbf{v}_{ij}^T \mathbf{P} \mathbf{v}_{ij} - 1| \le \delta$  for all  $i \ne j$ .
$$(14)$$

Alternatively, [54] modifies the formulation to determine the optimal  $\delta$  using a fixed number of measurements M, also adding an energy constraint on the coefficients of the matrix **A**. The resulting problem constrains both the rank and the trace norm of **P**.

$$\widehat{\mathbf{P}} = \arg\min_{\mathbf{P}^T = \mathbf{P} \succ \mathbf{0}} \max_{i \neq i} |\mathbf{v}_{ij}^T \mathbf{P} \mathbf{v}_{ij} - 1|$$
(15)

subject to rank(
$$\mathbf{P}$$
)  $\leq M$  and  $\|\mathbf{P}\|_* \leq b$ , (16)

where b is the energy constraint. Using a game-theoretic formulation, [54] also derives an algorithm to solve (16) with performance guarantees. It is also shown that the performance of the embedding can be guaranteed on new data, similar to the training set, using a continuity argument similar to the one in [10].

As mentioned in Sec. 2.3, a notion of semantic locality is also introduced in [31], in the context of classification. In particular, for elements i and j from the training data that belong in the same class, the embedding should guarantee that their distances do not increase significantly but does not need to limit how much they may shrink. On the other hand if elements i and j belong to different classes, the embedding should guarantee that their distances do not shrink significantly but may allow them to grow unconstrained. Under those conditions, the embedding guarantees that each cluster stays together, even though two different clusters may separate from each other. Thus, classification is still possible in the embedded data. The resulting optimization is less constrained than (15).

$$\widehat{\mathbf{P}} = \arg \min_{\mathbf{P}^{T} = \mathbf{P} \succeq 0} \|\mathbf{P}\|_{*}$$
subject to  $\mathbf{v}_{ij}^{T} \mathbf{P} \mathbf{v}_{ij} \ge 1 - \delta$  for all  $i \neq j$  in different classes.
$$\mathbf{v}_{ij}^{T} \mathbf{P} \mathbf{v}_{ij} \le 1 + \delta \text{ for all } i \neq j \text{ in the same class.}$$
(17)

In all the formulations above, **A** can be determined from  $\widehat{\mathbf{P}}$  using a simple factorization. For example, the economy-sized singular value decomposition is  $\widehat{\mathbf{P}} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{U}^T$ , where  $\mathbf{U} \in \mathbb{R}^{N \times M}$  has orthonormal columns and  $\boldsymbol{\Sigma} \in \mathbb{R}^{M \times M}$  is diagonal. The embedding can be computed using  $\widehat{\mathbf{A}} = \boldsymbol{\Sigma}^{1/2} \mathbf{U}^T$ .

#### **3** Preserving Inner Products, Angles, and Correlations

The embeddings discussed in the previous section are designed to preserve distances between signals in the embedding space. However, in a number of problems, inner products and correlations should be preserved instead. In this section we consider how distance embeddings can be used to preserve regular inner products and kernel inner products, as well as how binary and phase embeddings can be used to preserve normalized correlations, i.e., angles, without preserving distances.

10

#### 3.1 Inner Product Embeddings

When the signal and the embedding spaces are inner product spaces, then the inner product can be determined using the signal distances that are preserved. The inner product of the measurements  $\langle \mathbf{y}, \mathbf{y}' \rangle$  can be derived from the  $\ell_2^2$  difference of the measurements,  $\|\mathbf{y} - \mathbf{y}'\|_2^2$ . Specifically,

$$\|\mathbf{y} - \mathbf{y}'\|_{2}^{2} = \|\mathbf{y}\|_{2}^{2} + \|\mathbf{y}'\|_{2}^{2} - 2\langle \mathbf{y}, \mathbf{y}' \rangle \Longrightarrow \langle \mathbf{y}, \mathbf{y}' \rangle = \frac{\|\mathbf{y}\|_{2}^{2} + \|\mathbf{y}'\|_{2}^{2} - \|\mathbf{y} - \mathbf{y}'\|_{2}^{2}}{2}.$$
 (18)

When all these norms are preserved by the embedding, it is straightforward to show that JL-type random projections satisfy [2]

$$\left| \langle \mathbf{y}, \mathbf{y}' \rangle - \langle \mathbf{x}, \mathbf{x}' \rangle \right| \le \delta(\|\mathbf{x}\|_2^2 + \|\mathbf{x}'\|_2^2) \tag{19}$$

With a little more care, exploiting the linearity of the embedding, a tighter bound can be derived [27]

$$\left| \langle \mathbf{y}, \mathbf{y}' \rangle - \langle \mathbf{y}, \mathbf{y}' \rangle \right| \le \delta \|\mathbf{x}\|_2 \|\mathbf{x}'\|_2 \tag{20}$$

In addition to standard inner products, appropriately designed embeddings can also be used to approximate kernel inner products. Kernel inner product embeddings were first introduced in [51] and significantly generalized in [17,20]. Common kernels include the Gaussian  $K(\mathbf{x}, \mathbf{x}') = e^{\|\mathbf{x}-\mathbf{x}'\|_2^2/\sigma^2}$  and the Laplacian  $K(\mathbf{x}, \mathbf{x}') = e^{\gamma \|\mathbf{x}-\mathbf{x}'\|_1^2}$ . Since computing those kernels relies on computing distances, the development in Sec. 1 could be used to directly estimate the distance and compute the kernel. However, the resulting ambiguity would manifest itself in the exponent, making it difficult to characterize and control.

Instead, guarantees based on computing the inner product in the embedding domain can be derived, exploiting the design approach in Sec. 2.2. Similarly to standard inner products, establishing the guarantees relies on (18). However, the difficulty lies in bounding  $\|\mathbf{y}\|_2^2$  which is necessary, in addition to the distance between  $\mathbf{y}$  and  $\mathbf{y}'$ . When using the embedding design in Thm. 2, it is straightforward to show that, in the embedding space,

$$\sum_{k} |H_k|^2 - \varepsilon \le \frac{1}{M} \|\mathbf{y}\|_2^2 \le \sum_{k} |H_k|^2 + \varepsilon,$$
(21)

with probability greater than  $1 - 2e^{\log Q - 2M\frac{\varepsilon^2}{h}}$ . Thus, if  $d_{\mathcal{W}}(\mathbf{y}, \mathbf{y}') = \|\mathbf{y} - \mathbf{y}'\|_2^2$  in Def. 1, and substituting (4) and (21) in (18), we can show that the embedding can be designed to approximate a kernel.

**Theorem 3 (Thm. 4.4 in [20]).** Consider a set  $\mathscr{S}$  of Q points in  $\mathbb{R}^N$ , measured using  $\mathbf{y} = h(\mathbf{Ax} + \mathbf{w})$ , with  $\mathbf{A}$ ,  $\mathbf{w}$ , and h(t) as in Thm. 2. With probability greater than  $1 - e^{2\log Q - \frac{8}{9}M\frac{e^2}{\hbar^4}}$  the following holds

Petros T. Boufounos, Shantanu Rane, Hassan Mansour

$$K(d) - \varepsilon \le \frac{1}{M} \langle \mathbf{y}, \mathbf{y}' \rangle \le K(d) + \varepsilon$$
(22)

for all pairs  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$  and their corresponding measurements  $\mathbf{y}, \mathbf{y}'$ , where

$$K(d) = \sum_{k} |H_k|^2 \phi_l(k|d).$$
<sup>(23)</sup>

defines the kernel of the embedding.

Thus, to embed a Gaussian kernel and linear combinations of it, it suffices to draw the elements of **A** from an i.i.d. Gaussian distribution. Alternatively, to embed a Laplacian kernel and linear combinations of it, the elements of **A** should be drawn from a Cauchy distribution. The resulting kernels will be described by plugging the corresponding  $\phi(\cdot|d)$  in (23), in a similar manner as in (6) and (7):

$$K(d) = \sum_{k} |H_k|^2 e^{-2(\pi\sigma dk)^2}$$
(24)

$$K(d) = \sum_{k} |H_k|^2 e^{-2\pi\gamma dk}$$
<sup>(25)</sup>

where **A** is generated using an i.i.d. zero-mean Gaussian distribution with variance  $\sigma^2$  or an i.i.d. Cauchy distribution with scale parameter  $\gamma$  and *d* is the  $\ell_2$  or the  $\ell_1$  distance between signals, respectively.

## 3.2 Angle Embeddings

Another geometric quantity of interest in a number of applications is the angle between signals.

$$d_{\perp}(\mathbf{x}, \mathbf{x}') = \frac{1}{\pi} \arccos \frac{\langle \mathbf{x}, \mathbf{x}' \rangle}{\|\mathbf{x}\|_2 \|\mathbf{x}'\|_2}$$
(26)

The cosine of the angle is the correlation coefficient of the signals, i.e., their inner product normalized by their respective norms.

Since JL-type embeddings preserve distances and inner products, it is expected that they should preserve angles as well. A tighter bound than a naive application of the definition and the bounds of the previous section was shown in [30] in the context of sparse signals and the RIP. Specifically,

**Theorem 4 (Adapted from Thm. 1 and Remark 1 in [30]).** Consider an embedding satisfying the RIP for K-sparse vectors with RIP constant  $\delta \leq 1/3$ . For any K-sparse **x** and **x'** with the same support, such that  $d_{\geq}(\mathbf{x}, \mathbf{x}') \leq 1/2$  then the angle between the embedded vectors **y**, and **y'** satisfies

$$-\sqrt{3\delta} \le d_{\angle}(\mathbf{x}, \mathbf{x}') - d_{\angle}(\mathbf{y}, \mathbf{y}') \le 3\delta$$
(27)

13

$$\Longrightarrow \left| d_{\angle}(\mathbf{x}, \mathbf{x}') - d_{\angle}(\mathbf{y}, \mathbf{y}') \right| \le \sqrt{3\delta}.$$
(28)

This result can be used to derive a generalized notion of the RIP, linking the inner product of the embeddings with the geometry of the signals in the signal space [30, Cor. 1].

More recently, an embedding was derived in the context of 1-bit CS, explicitly preserving only angles of signals, not their inner products or magnitudes [37]. In particular, the embedding map

$$\mathbf{y} = f(\mathbf{x}) = \operatorname{sign}(\mathbf{A}\mathbf{x}),\tag{29}$$

where **A** has i.i.d. Normally distributed, entries maps the signals to an *M*-dimensional binary space, denoted  $\mathscr{B}^M$ , in which the normalized Hamming distance, defined as  $d_H(\mathbf{y}, \mathbf{y}') = (\sum_m y_m \oplus y'_m)/M$ , is the natural metric.

In [37] it is shown that (29) preserves the angle between signals in the normalized Hamming distance between the measurements, making it a Binary  $\varepsilon$ -stable embedding

**Definition 3.** Let  $\varepsilon \in [0,1)$  a mapping  $f : \mathscr{S} \to \mathscr{B}^M$  is a Binary  $\varepsilon$ -stable embedding (B $\varepsilon$ SE) of  $\mathscr{S}$  if for all  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$ ,

$$d_{\angle}(\mathbf{x}, \mathbf{x}') - \boldsymbol{\varepsilon} \le d_H(\mathbf{y}, \mathbf{y}') \le d_{\angle}(\mathbf{x}, \mathbf{x}') + \boldsymbol{\varepsilon}.$$
(30)

In other words, a B $\varepsilon$ SE is a  $(g_I, 0, \varepsilon)$  embedding according to Def. 1, with  $d_{\mathscr{S}} = d_{\angle}$  and  $d_{\mathscr{W}} = d_H$ . While the result has been developed for *K*-sparse vectors, it is straightforward to show that it holds for finite sets of *L* points using  $M = O(\varepsilon^{-2} \log L)$ .

**Theorem 5** (Adapted from Thm. 3 in [37]). Let  $\mathbf{A} \in \mathbb{R}^{M \times N}$  be a matrix generated from an i.i.d. Normal distribution and  $\mathscr{S}$  be a set of L points. The map (29) is a  $B\varepsilon SE$  of  $\mathscr{S}$  with probability greater than  $1 - 2e^{2(\log P - \varepsilon^2 M)}$ .

Subsequent work [3,46–49] demonstrated variations of this result for infinite signal sets, as a function of their mean width, with varying dependence on  $\varepsilon$ . Furthermore, with some constraints on the signals, it can also be shown for more general matrix ensembles, with elements drawn from subgaussian distributions.

The generalization of the sign function to complex numbers is the phase. As expected in hindsight, similar to sign measurements, phase measurements of the form

$$\mathbf{y} = \angle (\mathbf{A}\mathbf{x}) \tag{31}$$

can also provide stable angle embeddings [14–16]. In particular, if two signals  $\mathbf{x}, \mathbf{x}'$  in a finite set  $\mathcal{W}$  of size *L* are measured with a complex random Gaussian matrix, the expected value of the *m*<sup>th</sup> element of the measured phase difference is equal to

Petros T. Boufounos, Shantanu Rane, Hassan Mansour

$$E\left\{\left|\angle\left(e^{i(y_m-y'_m)}\right)\right|\right\} = \pi d_{\angle}(\mathbf{x},\mathbf{x}'),\tag{32}$$

Note that this way of calculating the phase difference naturally takes phase wrapping into account.

Similarly to the concentration of measure proofs so far, Hoeffding's inequality bounds the probability that the average of *M* random variables  $|\angle (e^{i(y_m - y'_m)})|$  deviates from (32). A natural distance metric in the embedding space is

$$d_{\text{phase}}(\mathbf{y}, \mathbf{y}') = \frac{1}{M} \sum_{m} \left| \frac{1}{\pi} \angle \left( e^{i(y_m - y'_m)} \right) \right|$$
(33)

Using the union bound on  $L^2$  point pairs, a stable embedding guarantee follows

**Theorem 6** ( [16]). Consider a finite set  $\mathscr{S}$  of L points measured using (31), with  $\mathbf{A} \in \mathbb{C}^{M \times N}$  consisting of *i.i.d* elements drawn from the standard complex normal distribution. With probability greater than  $1 - 2e^{2\log L - 2\varepsilon^2 M}$  the following holds for all  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$  and corresponding measurements  $\mathbf{y}, \mathbf{y}' \in \mathbb{R}^M$ .

$$\left| d_{\text{phase}}(\mathbf{y}, \mathbf{y}') - d_{\angle}(\mathbf{x}, \mathbf{x}') \right| \le \varepsilon \tag{34}$$

A complex-valued measurement matrix **A** is necessary here. If **A** only contains real elements, the information in **y** is essentially the sign of the measurement—0 and  $\pi$  for positive and negative measurements, respectively. In that case, the embedding becomes a B $\varepsilon$ SE. Furthermore, even though the embedding has an additive ambiguity—i.e., is a  $(g_I, 0, \varepsilon)$  embedding—it is conjectured that a multiplicative ambiguity guarantee should be possible to derive—i.e., that it is, in fact, a  $(g_I, \delta, 0)$ embedding [16].

Figure 3 compares the performance of this embedding with the B $\varepsilon$ SE, and demonstrates that, as expected, it exhibits lower ambiguity for the same number of measurements *M*. Furthermore, it shows that the becomes tighter as signals become similar, supporting the conjecture that a multiplicative-only ambiguity exists.



Fig. 3 Comparison of  $B\varepsilon SE$  (red) with continuous angle embedding (blue) for the same number of measurements. The continuous embedding becomes tighter as signals become more similar. As expected, the binary embedding has higher ambiguity for the same number of measurements.

14



Fig. 4 Illustration of the bits vs. measurements trade-off for quantized JL embeddings. (a) A sketch of the tradeoff between bits per coefficient and embedding dimension given a fixed bit-rate for quantized JL embeddings. The error due to the JL ambiguity  $\delta$  also depends on the norm of the signals being compared, thus affecting the true optimum in practice. Constants were arbitrarily selected for illustration puproses; the true optimum also depends on the true value of the constants. (b) Three different simulation examples using the same M = 256, quantized at 2, 4, and 32 bits per dimension, consuming R = 512, 1024, and 8192 bits, respectively. As expected, the 32 bit embedding performs best, but at a significant rate penalty. (c) Three simulation examples using rate R = 256, quantized at 2, 4, and 32 bits per dimension, requiring M = 128, 64, and 8 dimensions, respectively. As evident, quantizing at 32 bits per coefficient is now suboptimal; the JL-type error due to  $\delta$  dominates. In this example, 4 bits per coefficient quantization seems to provide the best tradeoff overall.

## 4 Quantized Embeddings

Quite frequently, the embedding is performed not simply as a dimensionaliy reduction, but as a compression method. In those cases, the quantity of interest is not the embedding dimensionality, but the number of bits it uses. Therefore, it is necessary to understand how quantization affects the embedding performance, and what the quantizer design trade-offs are.

# 4.1 Quantization of Continuous Embeddings

Although quantization of some embeddings can be analyzed using the periodic embedding framework we describe above, it is often more convenient, especially in the case of high-rate quantization, to consider it separately, as an additional step after the projection. The following development closely follows [20] and references within.

In particular, consider a  $(g, \delta, \varepsilon)$  embedding which is subsequently quantized using an *M*-dimensional vector quantizer  $Q(\cdot)$ . We assume the quantization error is bounded, i.e.,  $d(Q(\mathbf{x}), \mathbf{x}) \leq E_Q$ . The triangle inequality,  $|d_{\mathscr{W}}(f(\mathbf{x}), f(\mathbf{w})) - d_{\mathscr{W}}(Q(f(\mathbf{x})), Q(f(\mathbf{w})))| \leq 2E_Q$ , implies that the quantized embedding guarantee becomes a  $(g, \delta, \varepsilon + 2E_Q)$  embedding, with guarantee

$$(1 - \delta)g(d_{\mathscr{S}}(\mathbf{x}, \mathbf{y})) - \varepsilon - 2E_{Q} \leq d_{\mathscr{W}}(Q(f(\mathbf{x})), Q(f(\mathbf{y}))) \leq (1 + \delta)g(d_{\mathscr{S}}(\mathbf{x}, \mathbf{y})) + \varepsilon + 2E_{Q}.$$
 (35)

**Theorem 7 (Thm. 3.3 in [20]).** Consider a  $(g, \delta, \varepsilon)$  embedding  $f(\cdot)$  and a quantizer  $Q(\cdot)$  with worst case quantization error  $E_Q$ , then the quantized embedding,  $Q(f(\cdot))$ , is a  $(g, \delta, \varepsilon + 2E_Q)$  embedding.

In the specific case of a uniform scalar quantizer with quantization interval  $\Delta$ , the *M*-dimensional quantization  $\ell_2$  error is bounded by  $E_Q \leq \sqrt{M}\Delta/2$ , assuming the quantizer is designed such that it does not saturate or such that the saturation error is negligible. The interval of the quantizer is a function of the number of bits *B* used per coefficient  $\Delta = 2^{-B+1}S$ , where *S* is the saturation level of the quantizer. Given a fixed rate to be used by the embedding, R = MB, the guarantee becomes

$$(1-\delta)g(d_{\mathscr{S}}(\mathbf{x},\mathbf{y})) - \varepsilon - 2^{-\frac{R}{M}+1}\sqrt{M}S$$

$$\leq \|\mathcal{Q}(f(\mathbf{x})) - \mathcal{Q}(f(\mathbf{y}))\|_{2} \leq (1+\delta)g(d_{\mathscr{S}}(\mathbf{x},\mathbf{y})) + \varepsilon + 2^{-\frac{R}{M}+1}\sqrt{M}S. \quad (36)$$

Note that the  $\sqrt{M}$  factor can often be removed, depending on the normalization of the embedding.

Of course,  $\ell_2$  is not always the appropriate fidelity metric. If the  $d_{\mathscr{S}}(\cdot, \cdot)$  corresponds to the  $\ell_1$  distance, the quantization error is bounded by  $E_Q \leq M\Delta/2$ . Again, with care in the normalization, the *M* factor can be removed. If, instead, the  $\ell_{\infty}$  norm is desired, the quantization error is bounded by  $E_Q \leq \Delta/2$ .

One of the issues to consider in designing quantized embeddings using a uniform scalar quantizer is the trade-off between the number of bits per dimension and the total number of dimensions used. Since R = MB, increasing the number of bits per dimension *B* under a fixed bit budget *R*, requires decreasing the number of dimensions *M*. While the former reduces the error due to quantization, the latter will typically increase the uncertainty in the embedding by increasing  $\delta$  and  $\varepsilon$ .

In the case of randomized embeddings, this trade-off can be quantified through the function  $w(\varepsilon, \delta)$ . Given a fixed probability lower bound to guarantee the embedding, then  $M = \Omega(1/w(\varepsilon, \delta))$ . Since  $w(\cdot, \cdot)$  is an increasing function of  $\varepsilon$  and  $\delta$ , which quantify the ambiguity of the embedding, reducing *M* increases this ambiguity. On the other hand, the quantization ambiguity, given by  $2^{-\frac{R}{M}+2}S\sqrt{M}$  decreases with *M*.

This trade-off is explored, for example, in the context of quantized JL embeddings in [43,53]. In particular, randomly generated JL embeddings exhibit ambiguity  $\delta \sim 1/\sqrt{M}$ . On the other hand, the quantization error scales as  $E_Q \sim 2^{-B} \sim 2^{-1/M}$ . A illustrative example is shown in Fig. 4(a): as more bits are used per measurement the ambiguity due to quantization decreases; since fewer measurements are used, the ambiguity due to the embedding's  $\delta$  increases. Figures 4(b) and (c) further demonstrate this using a simulation experiment. In practice, the optimum depends on assumptions on the signal distance and assumptions about the constants of proportionality. The same issue exists for non-uniform quantizers and for vector quantizers, manifested with different constants but with the same order of magnitude effects (e.g., see [36]), as well as other embeddings, such as phase embeddings [14]. Unfortunately, other than experimentation with sample data, there is no known principled way to determine the optimal point in the trade-off.

In addition to the generic guarantees above, it is often possible to provide more explicit guarantees under certain conditions. For example, the 1-bit embedding guarantees in Sec. 3.2 were explicitly established from the embedding map. More recently, [34] draws similarities with the Buffon's needle problem to provide a tighter bound on the  $\ell_1$  embedding distance of quantized dithered JL-type embeddings

**Theorem 8 (Adapted from Prop. 2 in [34]).** Let  $\mathscr{S} \subset \mathbb{R}^N$  be a set of L points. Consider the map

$$\mathbf{y} = f(\mathbf{x}) = Q_{\varepsilon}(\mathbf{A}\mathbf{x} + \mathbf{w}), \tag{37}$$

where  $\mathbf{A} \in \mathbb{R}^{M \times N}$  has elements drawn from an i.i.d., standard Normal distribution, the dither  $\mathbf{w} \in \mathbb{R}^M$  has elements drawn from an i.i.d. distibution, uniform in  $[0, \varepsilon]$ , and  $Q_{\varepsilon}(\cdot)$  is an infinite uniform quantizer with interval  $\varepsilon$ .

Given  $0 < \delta < 1$ ,  $\varepsilon > 0$ , and  $M > C\delta^{-2}L$ , then, with probability greater than  $1 - e^{c''\varepsilon^2 M}$ , the map (37) satisfies

$$(1-\delta)\|\mathbf{x}-\mathbf{x}'\|_2 - c\varepsilon\delta \le \frac{c'}{M}\|\mathbf{y}-\mathbf{y}'\|_1 \le (1+\delta)\|\mathbf{x}-\mathbf{x}'\|_2 + c\varepsilon\delta$$
(38)

for all pairs of points  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$ .

A key insight in this result is the switch to the  $\ell_1$  norm in the embedding space, instead of the  $\ell_2$  norm used by the JL lemma and earlier results.

## 4.2 Universal Quantization and Embeddings

In constrast to conventional quantization analysis, universal scalar quantization, first introduced in [13], fundamentally revisits scalar quantization and redesigns the quantizer to have non-contiguous quantization regions. Unfortunately the discontinuous quantization regions render some of the tools introduced in Sec. 4 impractical. Fortunately, analysis based on the design described in Sec. 2.2 can be used instead.

A universal embedding also relies on a JL-style projection, followed by scaling, dithering and scalar quantization:

$$\mathbf{y} = f(\mathbf{x}) = Q(\Delta^{-1}(\mathbf{A}\mathbf{x} + \mathbf{w})), \tag{39}$$

where **A** is a  $M \times N$  random matrix with  $\mathcal{N}(0, \sigma^2)$ -distributed, i.i.d. elements,  $\Delta^{-1}$  a scaling factor, **w** a length-*M* dither vector with i.i.d. elements, uniformly distributed in  $[0, 2^B \Delta]$ , and  $Q(\cdot)$  a *B*-bit scalar quantizer operating element-wise on its input.

The key component is a modified *B*-bit scalar quantizer. Fitting the analysis of Thm. 2, the quantizer is designed to be a periodic function with non-contiguous quantization intervals, as shown in Fig. 5(a) for B = 1 and 2. The quantizer can be

thought of as a regular uniform quantizer, computing a multi-bit representation of a signal and preserving only the least significant bits (LSB) of the representation. For example, for a 1-bit quantizer, scalar values in [2l, 2l + 1) quantize to 1 and scalar values in [2l+1, 2(l+1)), for any integer *l*, quantize to 0. If  $Q(\cdot)$  is a 1-bit quantizer, this method encodes using as many bits as the rows of **A**, i.e., *M* bits.

This form of quantization, first proposed in [13] in the context of frame expansions and first used in an embedding in [18] is extensively analyzed in [20].

**Theorem 9 (Adapted from Thm. 3.2 in [18]).** Consider a set  $\mathscr{S} \subset \mathbb{R}^N$  with L points embedded using (39), as described above. For all  $\mathbf{x}, \mathbf{x}' \in \mathscr{S}$ , the embedding satisfies

$$g\left(\|\mathbf{x} - \mathbf{y}\|_{2}\right) - \varepsilon \leq d_{H}\left(\mathbf{y}, \mathbf{y}'\right) \leq g\left(\|\mathbf{x} - \mathbf{y}\|_{2}\right) + \varepsilon, \tag{40}$$

with probability  $1 - 2e^{2\log L - 2\varepsilon^2 M}$  with respect to the measure of **A** and **w**. In (40),  $d_H(\cdot, \cdot)$  is the Hamming distance of the embedded signals, the function  $f(\cdot)$  is as specified in (39), and g(d) is the map

$$g(d) = \frac{1}{2} - \sum_{i=0}^{+\infty} \frac{e^{-\left(\frac{\pi(2i+1)\sigma d}{\sqrt{2\Delta}}\right)^2}}{\left(\pi(i+1/2)\right)^2},\tag{41}$$

Furthermore, the distance map g(d) can be bounded using,

$$g(d) \ge \frac{1}{2} - \frac{1}{2}e^{-\left(\frac{\pi\sigma d}{\sqrt{2}\Delta}\right)^2},\tag{42}$$

$$g(d) \le \frac{1}{2} - \frac{4}{\pi^2} e^{-\left(\frac{\pi\sigma d}{\sqrt{2}\Delta}\right)^2},\tag{43}$$

$$g(d) \le \min\left(\sqrt{\frac{2}{\pi}} \frac{\sigma d}{\Delta}, \frac{1}{2}\right),$$
(44)

as shown in Fig. 5(b).

The upper bound (44) also provides a very good approximation of the embedding, as also evident in the figure. The map is approximately linear for small d and becomes constant, equal to 1/2, exponentially fast as d exceeds a threshold  $D_0$ . The slope of the linear section is determined by the parameter ratio  $\sigma/\Delta$ , thus specifying the distance threshold  $D_0 \approx \Delta \sqrt{\pi}/2\sqrt{2}\sigma$ . In other words, the embedding ensures that the Hamming distance of the embedded signals is approximately proportional to the  $\ell_2$  distance between the original signals, as long as that  $\ell_2$  distance was smaller than  $D_0$ . Distances greater than  $D_0$  are shrunk to Hamming distance  $\approx 1/2$ . In other words the embedding can only reveal that the distance is greater than approximately  $D_0$  but not how much greater.

This embedding enables a trade-off between the threshold  $D_0$  and the slope of the linear part, which determines its ambiguity through (10). Assuming the linear approximation in (44), it is straightforward to show that the ratio of the range of distances preserved, as measured through  $D_0$ , to the ambiguity in preserving distances



**Fig. 5** (a) This non-monotonic quantization function  $Q(\cdot)$  allows for universal rate-efficient scalar quantization. This function is equivalent to using a classical multibit scalar quantizer, and preserving only the least significant bits while discarding all other bits. 1-bit shown on top, multi-bit shown on bottom (b) The embedding map g(d) and its bounds produced by the 1-bit quantization function in (a). (c) Experimental verification of the embedding for small and large  $\Delta$  at high (left) and low (right) bit-rates.

in the linear part, as measured through (44) remains constant as the embedding parameters  $\Delta$  and  $\sigma$  change keeping a fixed embedding dimension *M*, and, therefore, a fixed rate R = M. In contrast to the tradeoff depicted in Fig. 4, both  $D_0$  and the slope of the linear part are straightforward to compute and do not depend on difficult-to-characterize constants.

Figure 5(c) illustrates how the embedding behaves in simulations for smaller (red) and larger (blue)  $\Delta$  and for higher (left) and lower (right) bit-rates. The figure plots the embedding (Hamming) distance as a function of the signal distance for randomly generated pairs of signals. The thickness of the curve is quantified by  $\varepsilon$ , whereas the slope of the upward sloping part is quantified by  $\Delta$ .

In addition to 1-bit universal embedding for finite signal sets, [20] generalizes the guarantees to infinite sets and to multi-bit embeddings. Of course, multibit embeddings re-introduce a similar trade-off as in Fig. 4, which has not been explored in the literature.

In addition to the embedding properties, information-theoretic arguments can be used to guarantee that universal embeddings can preserve the query privacy [18,39]. This can be a very useful property in implementing secure protocols for signal-based querying and retrieval in privacy-sensitive applications [52].

## **5** Discussion

As evident from the discussion above, embeddings can play a significant role in modern data processing systems. In this chapter, we have only presented a selective overview of the area, some important results, and pointers for further reading. However, increasing demand for efficient data processing has reinvigorated the field, leading to a flurry of new results in a number of interesting directions.

While we have only discussed  $\ell_p$  distance and angle embeddings in their various forms, there exist embeddings for more exotic distance metrics, such as the edit distance [5, 8, 42, 44]. Furthermore, while JL embeddings and the RIP preserve  $\ell_2$  distances, there is a large body of work in preserving other similarity measurements, such as  $\ell_p$  distances for various *p*'s [32,35,36,45,46]. It should be noted that in some cases, such as embedding the  $\ell_1$  distance into a smaller  $\ell_1$  space, such embeddings have been proven impossible [22]. Still, even in such cases, embeddings have been developed that hold with high, but not exponentially decreasing, probability [32].

A principal motivation for dimensionality reduction is often a reduction in computational complexity. However, the cost of storing and using a dense, fully randomized, embedding matrix can often be prohibitive. Fast transforms have been developed in a number of cases [4, 24, 57], enabling efficient computation of the transform, often without explicitly storing the matrix but using an algorithm, such as the fast Fourier transform (FFT). Still, even when the computation is efficient and the cost of storing the matrix is mitigated, the complexity of using the embedding for very large data retrieval can sometimes be daunting. While the dimensionality reduction definitely helps, the amount of the data, i.e., the number of data points, can be such that search is impossible even if the complexity is linear in the amount of data.

In such cases, locality-sensitive hashing (LSH) methods-which significantly reduce the computational complexity of near-neighbor computation-can be very helpful [6, 26, 33]. These methods are intimately connected to randomized embeddings. The LSH literature shares a lot of the tools, especially with quantized embeddings, such as randomized projections, dithering and quantization. The goal, however, is different. Given a query point, LSH will return its near neighbors very efficiently, using O(1) computation. This efficiency comes at a cost: no attempt is made to represent the distances between neighbors. When LSH is used to compare signals it only provides a binary decision, namely whether the distance between the signals is smaller than a radius or not. There is no guarantee that further information will be preserved. Thus, LSH may not be suitable for applications that require more accurate distance information. Still, the similarity of the methods suggests that some of the quantized embedding designs can be used as LSH functions. While there are examples of such use, [39], this is still an underdeveloped connection, especially for recent embedding designs. Techniques that learn a hash, such as spectral hashing [56] and LDAHash [55], also have strong similarities with embeddings and embedding learning.

Of course, this is a rich topic and it is impossible to exhaustively cover in this chapter. Our hope is that our development exposes the basic principles, some of the

foundational work, and some interesting recent developments. Our goal is to expose embeddings to a wider community, establishing them as an important tool, essential in the belt of any data scientist.

## References

- Achlioptas, D.: Database-friendly Random Projections: Johnson-lindenstrauss With Binary Coins. Journal of Computer and System Sciences 66, 671–687 (2003)
- Achlioptas, D., Mcsherry, F., Schölkopf, B.: Sampling techniques for kernel methods. In: Advances in Neural Information Processing Systems, pp. 335–342 (2002)
- Ai, A., Lapanowski, A., Plan, Y., Vershynin, R.: One-bit compressed sensing with nongaussian measurements. Linear Algebra and its Applications 441, 222–239 (2014)
- Ailon, N., Chazelle, B.: Approximate nearest neighbors and the fast johnson-lindenstrauss transform. In: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing, pp. 557–563 (2006)
- Andoni, A., Deza, M., Gupta, A., Indyk, P., Raskhodnikova, S.: Lower bounds for embedding edit distance into normed spaces. In: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms, pp. 523–526 (2003)
- Andoni, A., Indyk, P.: Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. Commun. ACM 51(1), 117–122 (2008). DOI 10.1145/1327452.1327494
- Bandeira, A.S., Mixon, D.G., Recht, B.: Compressive classification and the rare eclipse problem. arXiv preprint arXiv:1404.3203 (2014)
- Bar-Yossef, Z., Jayram, T., Krauthgamer, R., Kumar, R.: Approximating edit distance efficiently. In: Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on, pp. 550–559. IEEE (2004)
- Baraniuk, R., Cevher, V., Duarte, M., Hegde, C.: Model-based compressive sensing. IEEE Trans. Info. Theory 56(4), 1982–2001 (2010)
- Baraniuk, R., Davenport, M., DeVore, R., Wakin, M.: A simple proof of the restricted isometry property for random matrices. Const. Approx. 28(3), 253–263 (2008)
- Baraniuk, R., Wakin, M.: Random projections of smooth manifolds. Foundations of Computational Mathematics 9(1), 51–77 (2009)
- Blumensath, T., Davies, M.: Sampling theorems for signals from the union of finitedimensional linear subspaces. IEEE Trans. Info. Theory 55(4), 1872–1882 (2009)
- Boufounos, P.T.: Universal rate-efficient scalar quantization. IEEE Trans. Info. Theory 58(3), 1861–1872 (2012). DOI 10.1109/TIT.2011.2173899
- Boufounos, P.T.: Angle-preserving quantized phase embeddings. In: Proc. SPIE Wavelets and Sparsity XV. San Diego, CA (2013)
- 15. Boufounos, P.T.: On embedding the angles between signals. In: Proc. Signal Processing with Adaptive Sparse Structured Representations (SPARS). Lausanne, Switzerland (2013)
- Boufounos, P.T.: Sparse signal reconstruction from phase-only measurements. In: Proc. Int. Conf. Sampling Theory and Applications (SampTA). Bremen, Germany (2013)
- Boufounos, P.T., Mansour, H.: Universal embeddings for kernel machine classification. In: Proc. Sampling Theory and Applications. Washington, DC (2015)
- Boufounos, P.T., Rane, S.: Secure binary embeddings for privacy preserving nearest neighbors. In: Proc. Workshop on Information Forensics and Security (WIFS). Foz do Iguau, Brazil (2011). DOI 10.1109/WIFS.2011.6123149
- Boufounos, P.T., Rane, S.: Efficient coding of signal distances using universal quantized embeddings. In: Proc. Data Compression Conference (DCC). Snowbird, UT (2013)
- Boufounos, P.T., Rane, S., Mansour, H.: Representation and coding of signal geometry. arXiv preprint arXiv:1512.07636 (2015)

- Bourgain, J., Dirksen, S., Nelson, J.: Toward a unified theory of sparse dimensionality reduction in euclidean space. Geometric and Functional Analysis 25(4), 1009–1088 (2015)
- Brinkman, B., Charikar, M.: On the impossibility of dimension reduction in l 1. Journal of the ACM (JACM) 52(5), 766–788 (2005)
- Candès, E.: The restricted isometry property and its implications for compressed sensing. Comptes rendus de l'Académie des Sciences, Série I 346(9-10), 589–592 (2008)
- Candes, E.J., Romberg, J.K., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. Communications on pure and applied mathematics 59(8), 1207–1223 (2006)
- Dasgupta, S., Gupta, A.: An elementary proof of a theorem of Johnson and Lindenstrauss. Random Structures & Algorithms 22(1), 60–65 (2003)
- Datar, M., Immorlica, N., Indyk, P., Mirrokni, V.S.: Locality-sensitive hashing scheme based on p-stable distributions. In: Proceedings of the twentieth annual symposium on Computational geometry, SCG '04, pp. 253–262. ACM, New York, NY, USA (2004)
- Davenport, M.A., Boufounos, P.T., Wakin, M.B., Baraniuk, R.G.: Signal processing with compressive measurements. IEEE Journal of Selected Topics in Signal Processing 4(2), 445–460 (2010). DOI 10.1109/JSTSP.2009.2039178. URL http://dx.doi.org/10.1109/JSTSP.2009.2039178
- Dirksen, S.: Dimensionality reduction with subgaussian matrices: a unified theory. Foundations of Computational Mathematics pp. 1–30 (2014)
- Eldar, Y., Mishali, M.: Robust recovery of signals from a structured union of subspaces. IEEE Trans. Info. Theory 55(11), 5302–5316 (2009)
- Haupt, J., Nowak, R.: A generalized restricted isometry property. Tech. rep., University of Wisconsin-Madison (2007)
- Hegde, C., Sankaranarayanan, A., Yin, W., Baraniuk, R.: NuMax: A convex approach for learning near-isometric linear embeddings. IEEE Trans. Signal Processing 63(22), 6109–6121 (2015). DOI 10.1109/TSP.2015.2452228
- Indyk, P.: Stable distributions, pseudorandom generators, embeddings, and data stream computation. Journal of the ACM (JACM) 53(3), 307–323 (2006)
- Indyk, P., Motwani, R.: Approximate nearest neighbors: towards removing the curse of dimensionality. In: ACM Symposium on Theory of computing, pp. 604–613 (1998)
- Jacques, L.: A quantized johnson-lindenstrauss lemma: The finding of buffon's needle. IEEE Trans. Info. Theory 61(9), 5012–5027 (2015). DOI 10.1109/TIT.2015.2453355
- Jacques, L., Hammond, D.K., Fadili, J.M.: Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine. Information Theory, IEEE Transactions on 57(1), 559–571 (2011)
- Jacques, L., Hammond, D.K., Fadili, J.M.: Stabilizing nonuniformly quantized compressed sensing with scalar companders. IEEE Trans. Info. Theory 59(12), 7969–7984 (2013)
- Jacques, L., Laska, J.N., Boufounos, P.T., Baraniuk, R.G.: Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. IEEE Trans. Info. Theory 59(4) (2013). DOI 10.1109/TIT.2012.2234823. URL http://dx.doi.org/10.1109/TIT.2012.2234823
- Jayram, T., Woodruff, D.P.: Optimal bounds for johnson-lindenstrauss transforms and streaming problems with subconstant error. ACM Transactions on Algorithms (TALG) 9(3), 26 (2013)
- Jimenez, A., Raj, B., Portelo, J., Trancoso, I.: Secure modular hashing. In: IEEE International Workshop on Information Forensics and Security (WIFS), pp. 1–6. IEEE (2015)
- Johnson, W., Lindenstrauss, J.: Extensions of Lipschitz mappings into a Hilbert space. Contemporary Mathematics 26, 189 –206 (1984)
- Krahmer, F., Ward, R.: New and improved johnson-lindenstrauss embeddings via the restricted isometry property. SIAM Journal on Mathematical Analysis 43(3), 1269–1281 (2011)
- Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions, and reversals. Soviet Physics Doklady 10(8), 707–710 (1966)
- 43. Li, M., Rane, S., Boufounos, P.T.: Quantized embeddings of scale-invariant image features for mobile augmented reality. In: Proc. IEEE International Workshop on Multimedia Signal Processing (MMSP). Banff, Canada (2012)

- Ostrovsky, R., Rabani, Y.: Low distortion embeddings for edit distance. Journal of the ACM (JACM) 54(5), 23 (2007)
- Otero, D., Arce, G.R.: Generalized restricted isometry property for alpha-stable random projections. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 3676–3679 (2011)
- Oymak, S., Recht, B.: Near-optimal bounds for binary embeddings of arbitrary sets. arXiv preprint arXiv:1512.04433 (2015)
- Plan, Y., Vershynin, R.: One-bit compressed sensing by linear programming. Communications on Pure and Applied Mathematics 66(8), 1275–1297 (2013)
- Plan, Y., Vershynin, R.: Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. Information Theory, IEEE Transactions on 59(1), 482–494 (2013)
- Plan, Y., Vershynin, R.: Dimension reduction by random hyperplane tessellations. Discrete & Computational Geometry 51(2), 438–461 (2014)
- Puy, G., Davies, M., Gribonval, R.: Recipes for stable linear embeddings from hilbert spaces to R<sup>m</sup>. arXiv preprint arXiv:1509.06947 (2015)
- 51. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. In: Advances in neural information processing systems, pp. 1177–1184 (2007)
- Rane, S., Boufounos, P.T.: Privacy-preserving nearest neighbor methods: Comparing signals without revealing them. IEEE Signal Processing Magazine (2013). DOI 10.1109/MSP.2012.2230221. URL http://dx.doi.org/10.1109/MSP.2012.2230221
- Rane, S., Boufounos, P.T., Vetro, A.: Quantized embeddings: An efficient and universal nearest neighbor method for cloud-based image retrieval. In: Proc. SPIE Applications of Digital Image Processing XXXVI. San Diego, CA (2013)
- Sadeghian, A., Bah, B., Cevher, V.: Energy-aware adaptive bi-Lipschitz embeddings. In: Proc. Int. Conf. Sampling Theory and Applications (SampTA). Bremen, Germany (2013)
- Strecha, C., Bronstein, A., Bronstein, M., Fua, P.: LDAHash: Improved matching with smaller descriptors. IEEE Trans. Pattern Analysis and Machine Intelligence 34(1), 66 –78 (2012). DOI 10.1109/TPAMI.2011.103
- Weiss, Y., Torralba, A., Fergus, R.: Spectral hashing. In: Advances in Neural Information Processing Systems 21, pp. 1753–1760 (2009)
- Yi, X., Caramanis, C., Price, E.: Binary embedding: Fundamental limits and fast algorithm. In: Proc. 32nd International Conference on Machine Learning, vol. 37, pp. 2162–2170. Lille, France (2015)