

Embedding-based Representation of Signal Distances

Petros T. Boufounos and Shantanu Rane
Mitsubishi Electric Research Laboratories
Cambridge, MA 02139, USA.
{petrosb, rane}@merl.com

Abstract—Traditional signal representation and coding theory is focused on how to most efficiently represent and encode a signal with the goal of preserving it as best as possible. However, very often, the processing only concerns specific information in the signal and does not require conserving the signal itself. In this work we examine the problem of representing signals such that some function of their distance is preserved. For that goal, we consider randomized embeddings as a representation mechanism and provide a framework to design them and analyze their performance. This work generalizes previously developed universal embeddings, already proven quite successful in practice.

Index Terms—distance representation, randomized embeddings

I. INTRODUCTION

SIGNAL representation theory and practice has primarily focused on how to best represent or code a signal incurring the smallest possible distortion. For example, image or video representations typically aim to minimize the distortion in the signal so that the visual quality of the signal is maintained when displayed to a user. Quite often, however, the user of a signal is not a human observer, but an algorithm extracting information. In this case, the goal is different: the representation should not destroy the information that the algorithm requires, even if the signal itself cannot be completely recovered.

The representations examined in this paper target applications in which only the information about the distance of one signal to another, or a function of that distance is required. To represent distances, we exploit the distance-preserving properties of randomized embeddings. Our approach generalizes well-known embeddings in a manner that enables the design and control of the distance distortion.

An *embedding* is a transformation of a set of signals in a high-dimensional space to a (typically) lower-dimensional space such that some aspects of the geometry of the set are preserved. Thus, computations requiring distance evaluations can be performed directly using the embedding, rather than the signals themselves. The best known embeddings are due to Johnson and Lindenstrauss, which preserve ℓ_2 distances using a randomized linear map [1].

II. EMBEDDING DESIGN FOR REPRESENTATION OF DISTANCES

In this work we consider a fairly general approach to design embeddings that preserve distances of signals subject to a distance map $g(\cdot)$. Such embeddings allow the design of representations and coding schemes with control on the distortion in the distances, rather than the distortion in the signals.

Specifically we examine embeddings of the form $\mathbf{y} = h(\mathbf{A}\mathbf{x} + \mathbf{w})$, where $h(\cdot)$ is a scalar function, applied element-wise to each coefficient of its argument. We assume it is periodic with period 1, i.e., $h(t) = h(t + 1)$, and has Fourier series denoted by H_k . It is also bounded, with $\bar{h} = \sup_t h(t) - \inf_t h(t)$. We assume the rows, \mathbf{a}_i , of \mathbf{A} are randomly chosen from a rotationally invariant i.i.d. vector distribution and the elements of \mathbf{w} are chosen from an i.i.d. distribution uniform in $[0, 1)$. Thus, any projection $l_i = \langle \mathbf{a}_i, \mathbf{x} \rangle$ is a random variable which we assume only depends on the length $d = \|\mathbf{x}\|_2$. We use $\phi(\xi; d)$ to denote the characteristic function of this random variable. For example, if elements of \mathbf{A} are i.i.d. normally distributed with variance σ^2 , then $\phi_{N(\sigma)}(\xi; d) = e^{-\frac{1}{2}(\sigma d \xi)^2}$.

Under those assumptions, we define a distance map

$$g(d) = 2 \sum_k |H_k|^2 (1 - \phi(k; d)). \quad (1)$$

which, in the special case of normally distributed \mathbf{a} , becomes

$$g(d) = 2 \sum_k |H_k|^2 \left(1 - e^{-\frac{1}{2}(\sigma d k)^2} \right). \quad (2)$$

The embedding represents distances subject to this distance map, with accuracy characterized by the following theorem.

Theorem 2.1: Consider a set \mathcal{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 failure probability $P_F \leq 2Q^2 e^{-2M \frac{\epsilon^2}{\bar{h}^4}}$ the following holds

$$g(\|\mathbf{x} - \mathbf{x}'\|_2) - \epsilon \leq \frac{1}{M} \|\mathbf{y} - \mathbf{y}'\|_2^2 \leq g(\|\mathbf{x} - \mathbf{x}'\|_2) + \epsilon \quad (3)$$

for all pairs $\mathbf{x}, \mathbf{x}' \in \mathcal{S}$ and corresponding measurements \mathbf{y}, \mathbf{y}' .

If $h(t)$ is Lipschitz continuous, Thm. 2.1 can be generalized to arbitrary compact sets \mathcal{S} , such as bounded sparse signals. With more care we can allow for piecewise continuous $h(t)$. This approach generalizes the results in [2], [3] and allows their extension from binary embeddings of point clouds to multi-bit embeddings of compact sets. Details are reserved for an extended version of this paper.

Although (3) characterizes the embedding ambiguity, this form is not useful when representing or coding distances [4]. In such applications, the goal is to estimate the true distance between signals, given their representation, and characterize the estimate's ambiguity.

Specifically, given the embedding distance $d_e = \|\mathbf{y} - \mathbf{y}'\|_2^2 / M$, the signal distance $\|\mathbf{x} - \mathbf{x}'\|_2$ can be estimated using $\hat{d}_s = g^{-1}(d_e)$. The uncertainty of this estimate is approximately equal to $\epsilon / g'(d_s)$, i.e., inversely proportional to $g'(\cdot)$, the slope of $g(\cdot)$. Thus, the steepest parts of the distance map preserve distances best, while flatter parts introduce more ambiguity. Note that scaling $h(\cdot)$ also affects the ratio ϵ^2 / \bar{h}^4 in the probability P_F with which the embedding fails. To keep this probability constant when scaling $h(\cdot)$, ϵ will scale accordingly and, therefore, the ambiguity will remain constant.

The ability to control the error and the ambiguity is integral to the function of a representation. Scalar quantizers, for example, are often non-uniform in order to better accommodate the signal statistics or the distortion that users tolerate at different signal values. Similarly, by manipulating the slope of $g(\cdot)$ it is possible to adjust the distortion of the distance representation and make it distance-dependent.

REFERENCES

- [1] W. Johnson and J. Lindenstrauss, "Extensions of Lipschitz mappings into a Hilbert space," *Contemporary Mathematics*, vol. 26, pp. 189–206, 1984.
- [2] P. T. Boufounos, "Universal rate-efficient scalar quantization," *IEEE Trans. Info. Theory*, vol. 58, no. 3, pp. 1861–1872, March 2012.
- [3] P. T. Boufounos and S. Rane, "Secure binary embeddings for privacy preserving nearest neighbors," in *Proc. Workshop on Information Forensics and Security (WIFS)*, Foz do Iguazu, Brazil, November 29–December 2 2011.
- [4] —, "Efficient coding of signal distances using universal quantized embeddings," in *Proc. Data Compression Conference (DCC)*, Snowbird, UT, March 20–22 2013.