

DISTANCE-BASED DISCRETIZATION OF PARAMETRIC SIGNAL MANIFOLDS

Elif Vural and Pascal Frossard

Ecole Polytechnique Fédérale de Lausanne (EPFL)
Signal Processing Laboratory - LTS4
CH-1015 Lausanne, Switzerland

ABSTRACT

The characterization of signals and images in manifolds often lead to efficient dimensionality reduction algorithms based on manifold distance computation for analysis or classification tasks. We propose in this paper a method for the discretization of signal manifolds given in a parametric form. We present an iterative algorithm for the selection of samples on the manifold that permits to minimize the average error in the manifold distance computation. Experimental results with image appearance manifolds demonstrate that the proposed discretization algorithm outperforms baseline solutions based on random or regular sampling, both in terms of projection accuracy and image registration.

Index Terms— Manifold discretization, image appearance manifolds, manifold distance, pattern transformations

1. INTRODUCTION

Despite the high dimensionality of images, there are many settings where the visual content has an underlying low-dimensional manifold structure. For example, an image appearance manifold (IAM) [1] represents a family of images that can be described by a small set of parameters such as the pose of an object, or the transformation applied to a visual pattern. Low-dimensional manifolds are typically used for transformation-invariant analysis or classification tasks. Transformation invariance is then achieved via the computation of manifold distance or the utilization of alternative distance measures such as the tangent distance [2],[3],[4].

While such distance metrics have proven performance in classification and alignment problems, they are quite expensive in terms of computational complexity. It may therefore be preferable to approximate the manifold by a discrete set of manifold samples for faster computation, though at the price of a lower distance accuracy. However, a regular discretization of a manifold in its parameter space in general does not lead to good performance. The manifold should rather be sampled depending on the manifold geometry and the targeted applications.

In this work, we propose a method for discretizing manifolds of known parameterization, which favors accurate manifold distance estimation. We describe an algorithm that minimizes the manifold distance estimation error caused by representing the manifold with only a finite number of samples. We adopt an iterative approach that alternates between the computation of partition cells and representative samples. We demonstrate the performance of our method with experimentation on two different types of analytically expressible

image appearance manifolds. Our algorithm provides smaller registration error and better projection accuracy than baseline sampling solutions based on random and respectively regular sampling of the manifold.

2. MANIFOLD DISCRETIZATION

Let $\mathcal{M} \subset \mathbb{R}^n$ be a set of signals conforming to a manifold model defined over the parameter domain $\Lambda \subset \mathbb{R}^m$. Then \mathcal{M} can be written as $\mathcal{M} = \{U(\eta), \eta \in \Lambda\}$, where U denotes the mapping from the m -dimensional parameter domain Λ to the examined set of n -dimensional signals.

Given a point $x \in \mathbb{R}^n$, the distance of x to the manifold \mathcal{M} is $d(x, \mathcal{M}) = \min\{d(x, U(\eta)), U(\eta) \in \mathcal{M}\}$, where we take the distance function $d(x, y)$ as the ℓ_2 -distance $\|x - y\|$. The discretization of the manifold consists in selecting a predetermined number N of manifold points, i.e. a sample set $\{S_i\}_{i=1}^N = \{\bigcup_{i=1}^N U(\eta_i), \eta_i \in \Lambda\} \subset \mathcal{M}$.

We would like to select the samples in order to minimize the total manifold distance estimation error E over R , where $R = \{x \in \mathbb{R}^n : a_i \leq x_i \leq b_i\}$ is the hyperrectangular region of \mathbb{R}^n corresponding to n -dimensional signals¹. We define E by

$$E = \int_R (d^2(x, \{S_i\}_{i=1}^N) - d^2(x, \mathcal{M})) dx, \quad (1)$$

where $d(x, \{S_i\}_{i=1}^N) = \min\{d(x, S_i) : i = 1, 2, \dots, N\}$ denotes the distance between x and the sample set $\{S_i\}_{i=1}^N$.

For a given sample set, one can partition R into N cells as $R = \bigcup_{i=1}^N R_i$, where each cell R_i is a polytope consisting of points with smaller ℓ_2 -distance to S_i than to any other S_j , i.e., $R_i = \{x \in R : d(x, S_i) < d(x, S_j), \forall j \neq i\}$. Hence, the total manifold distance estimation error becomes

$$E = \sum_{i=1}^N E_i = \sum_{i=1}^N \int_{R_i} (d^2(x, S_i) - d^2(x, \mathcal{M})) dx. \quad (2)$$

In order to minimize the error E , we follow an iterative optimization procedure. In each iteration of the algorithm, we first compute the partition cells corresponding to the samples, and then optimize each sample individually such that the error E_i in the regarding cell is minimized. Once partition cells are computed, minimization of the manifold distance estimation error E_i within a specific cell R_i is achieved as follows: E_i can be rearranged as

$$E_i = \int_{R_i} d^2(x, S_i) dx - \int_{R_i} d^2(x, \mathcal{M}) dx,$$

This work has been partly funded by the Swiss National Science Foundation under grant number 200021_120060.

¹The parameters a_i and b_i have typically the values 0 and 255 for images.

where the second integration depends only on R_i , and is constant with respect to S_i . Therefore, E_i is given by

$$\begin{aligned} E_i &= \int_{R_i} d^2(x, S_i) dx + c_i \\ &= \int_{R_i} x^T x dx - 2S_i^T \int_{R_i} x dx + V_i S_i^T S_i + c_i, \end{aligned}$$

where $V_i = \int_{R_i} dx$ is the volume of the cell R_i . Denoting the centroid of the cell R_i by $G_i = (\int_{R_i} x dx) / (\int_{R_i} dx)$, we get

$$\begin{aligned} E_i &= \int_{R_i} x^T x dx + V_i(-2S_i^T G_i + S_i^T S_i) + c_i \\ &= V_i(S_i^T S_i - 2S_i^T G_i) + c_i', \end{aligned}$$

where we express the sum of the terms independent of S_i by c_i' . As E_i differs from $\|S_i - G_i\|^2$ only up to a positive multiplicative factor and an additive term constant with respect to S_i , one can equivalently minimize

$$E_i' = \|S_i - G_i\|^2 \quad (3)$$

at each iteration of the algorithm, which actually means that S_i should be selected as the manifold point closest to the centroid of the cell R_i .

To sum up these discussions, the following is a summary of the procedure we apply for obtaining a manifold discretization that minimizes the total manifold distance estimation error: Given the available domain of parameters and the mapping defining the manifold, we begin with an initial sample set $\{S_i^{(0)}\}_{i=1}^N$ on the manifold, which is possibly randomly selected. We optimize the sample set iteratively. In each k -th iteration of the algorithm, we first compute the cells $\{R_i^{(k)}\}_{i=1}^N$ that partition the hyperrectangular region of interest of \mathbb{R}^n with respect to the manifold samples $\{S_i^{(k)}\}_{i=1}^N$; and then we perturb each sample $S_i^{(k)}$ individually to obtain the new sample $S_i^{(k+1)}$ such that the manifold distance estimation error given by (3) is minimized in the corresponding cell. The new sample $S_i^{(k+1)}$ is the projection of the centroid $G_i^{(k)}$ onto the manifold. Iterations are repeated until improvements become negligible. We call this algorithm Registration-Efficient Manifold Discretization (REMD). An iteration of the algorithm is illustrated in Fig. 1, and the pseudocode is given in Alg. 1.

Assuming that the feasible domain Λ of parameter vectors is compact and the mapping U is bounded, for a given number of samples N , there exists a solution $\{S_i^*\}_{i=1}^N$ that globally minimizes the total error E in (1). At each iteration of the discussed method, first the partition cells are updated and then the samples are readjusted, both of which are modifications that either reduce E or retain it. Since the error E is non-increasing throughout the iterations and is also lower bounded, the algorithm converges. However, in general the cost function is a non-convex, complicated function of the parameter vectors; therefore, the algorithm is not guaranteed to converge to the globally optimal solution. Finally, we note that the discretization algorithm proposed above is similar in several aspects to the LBG vector quantization algorithm [5], which however targets the minimization of the error in the signal approximation and not in the manifold distance.




Fig. 1. Illustration of a single iteration of the algorithm

Algorithm 1 Registration-Efficient Manifold Discretization

1: **Input:**

- Λ : Feasible domain of parameter vectors
- U : Mapping from parameter domain Λ to manifold \mathcal{M}
- N : Number of manifold samples

2: **Output:**

- $\{S_i\}_{i=1}^N$: A set of manifold samples

3: **Initialization:**

- 4: Choose an initial set of manifold samples $\{S_i^{(0)}\}_{i=1}^N$.
- 5: $k = 0$.

- 6: **while** The difference between $\{S_i^{(k)}\}_{i=1}^N$ and $\{S_i^{(k-1)}\}_{i=1}^N$ is above some threshold (handling $k = 0$ properly) **do**

- 7: Determine the partition cells $\{R_i^{(k)}\}_{i=1}^N$.
- 8: Compute the centroids $\{G_i^{(k)}\}_{i=1}^N$ of partition cells.
- 9: Update the sample set $\{S_i^{(k+1)}\}_{i=1}^N$ by optimizing each sample such that the manifold distance estimation error $E_i^{(k)}$ in (3) is minimized individually in each cell.

- 10: $k = k + 1$.

11: **end while**

- 12: $\{S_i\}_{i=1}^N = \{S_i^{(k)}\}_{i=1}^N$.
-

3. EXPERIMENTAL RESULTS

The performance of the proposed method is evaluated on two different kinds of image appearance manifolds, namely a 2D pattern transformation manifold, and the manifold generated by the observations of a synthetic 3D object model under varying illumination and viewpoint. In both experiments, the algorithm starts with a randomly selected sample set, and the output discretization is compared to the initial random discretization as well as the one obtained on a regular grid in the parameter domain.

We achieve the computation of the centroids of partition cells experimentally. The centroid of a given region of \mathbb{R}^n can be estimated simply by taking random training points in space, checking if they are in the inquired region, and then computing the arithmetic average of inliers when a sufficient number of them are accumulated [6]. In our setup, we do not select training points as arbitrary space points, instead, we use image databases for training. Actually, in this case the partition cells R_i can be conceived as special regions of interest inside n -dimensional polytopes rather than the polytopes themselves as previously discussed. After the centroids are computed, we estimate their projections onto the manifold with the aid of a dense grid on the manifold. We first locate the projection coarsely by finding the grid point that has the smallest distance to the centroid, and then refine the location of the projection by minimizing its distance to the centroid using gradient descent tools.

3.1. Results on pattern transformation manifolds

In this experiment, we study the discretization of the pattern transformation manifold generated by the 2D rotation and translation of a visual pattern. The transformation manifold of a visual pattern p is

$$\mathcal{M} = \{U(\eta)p : \eta = (\theta, t_x, t_y) \in \Lambda\}, \quad (4)$$

where θ is the rotation parameter, t_x and t_y are the horizontal and vertical translation parameters, and Λ is the domain of transformation parameter vectors.

For experimentation, we use a database of top-view images of 5 different objects, where each object has 500 different images captured under different orientation and positions. An example image for each object is displayed in Fig. 2. Note that due to the above positioning of the camera and the limitations on object positions, the 2D pattern transformation model in (4) constitutes an approximate model for the observations.

We consider the pattern transformation manifold of each object separately. For each of the objects we build the transformation manifold of a fixed representative pattern which is picked randomly among the database images. The image set of each object is grouped randomly into 300 training and 200 test images. For computational convenience, all images are converted to greyscale, downsampled to a resolution of 50×60 pixels, and background pixels are set to the luminance value of 0 by simple thresholding. Manifold points are generated by rotating and translating the representative pattern (cropped previously near the boundary) over a 50×60 pixel zero background within the parameter range $\theta \in [-\pi, \pi]$; $t_x \in [-7, 7]$; $t_y \in [-12, 12]$. All images and generated manifold points are normalized to have unit norm.

We compare the sample set obtained by the REMD algorithm to the initial random sample set and to the sample set given by a regular grid over the parameter domain, where the performance evaluation criterion is the accountability of the discretization for accurate manifold distance estimation. For each discretization, the distances of the test points to the sample set are computed and the average registration error is calculated. The registration error is taken as the ℓ_2 -distance between the exact projection of the test point onto the manifold and the manifold sample with smallest ℓ_2 -distance to the test point. For each object the experiment is repeated 5 times with different random initializations, where the categorization into training and test sets is changed randomly at each run. The results presented in Fig. 3 are averaged over all realizations and all objects. In Fig. 3(a), average registration errors are plotted for various numbers of manifold samples. In Fig. 3(b), we illustrate the distribution of test points in percentage with respect to which set the manifold sample nearest to them belongs to. Here, we compute the distance of each test point to the manifold samples coming from all sets, and report the percentage of the test points that have their closest manifold sample within the REMD output, random and regular sample sets. These experiments intend to measure the capability of the discretizations to provide an accurate approximation of the projection onto the manifold. As demonstrated in the figures, the discretization obtained by the REMD method yields the least registration error when compared to the random discretization, and the regular discretization in the parameter domain. In addition, for the majority of the test points the most accurate approximation of the projection lies within the REMD algorithm output sample set.

3.2. Results on synthetical object observation manifolds

Now we evaluate the performance of our method on the observation manifolds of synthetical 3D objects. For a 3D object model m , we






Fig. 2. Example images from database



(a) Average registration errors



(b) Distribution of test points with respect to closest projection approximations

Fig. 3. Discretization results on pattern transformation manifolds



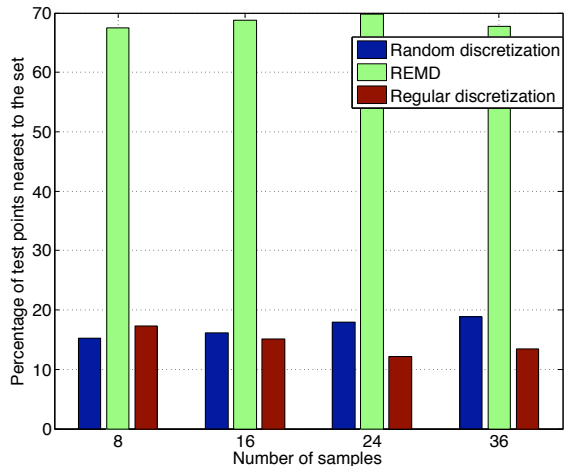


Fig. 4. Example objects from the airplane class



(a) Average registration errors



(b) Distribution of test points with respect to closest projection approximations

Fig. 5. Discretization results on synthetical object observation manifolds

consider the observation manifold of the object defined by

$$\mathcal{M} = \{U(\eta)m : \eta = (I, \theta_x, \theta_y, \theta_z) \in \Lambda\}, \quad (5)$$

where $U(\eta)m$ is the image of the object rendered under the illumination intensity I and the orientation specified by the three rotation angles $\theta_x, \theta_y, \theta_z$.

We use the Princeton Shape Benchmark² database of 3D models, where we conduct our experiments on 8 different classes of objects (car, airplane, ship, tank, human, animal, table, bottle) with several (4-30) objects belonging to each class. Some example objects belonging to the airplane class are shown in Fig 4. For each class we choose a representative object, and sample the observation manifold of the representative object in the parameter range $I \in [0.5, 1]$; $\theta_x \in [0, \pi/2]$; $\theta_y, \theta_z \in [-\pi/4, \pi/4]$. All rendered images are converted to greyscale and downsampled to the resolution 50×50 pixels. The training and test sets for each manifold consist of random observations of the objects belonging to the same class within the

²<http://shape.cs.princeton.edu/benchmark>

same parameter range. The experimentation setup is the same as in Sec. 3.1. Experiments are repeated 4 times for each class with different random initializations, and the representative object of each class is changed randomly at each run. The results, which are presented in Fig. 5, are averaged over all realizations and all objects. In the computation of the registration errors plotted in Fig. 5(a), the registration error for each test image is normalized by the image norm. Fig. 5(b) shows the percentage of test images with most accurate projection approximations within the REMD output sample set, random sample set and the sample set regular in the parameter domain. The results confirm the findings of Sec. 3.1.

4. CONCLUSIONS

We have proposed a method for the discretization of signal manifolds of known parameterization. We optimize manifold samples in an iterative procedure in order to improve the manifold distance estimation error resulting from the approximate representation of the manifold with the finite set of selected points. Experimentation on different type and intrinsic dimension of image appearance manifolds suggests that the proposed method is capable of determining samples that yield a fair estimation of the projection onto the manifold. On the other hand, state-of-the-art methods accomplishing manifold distance computation are considerably demanding in terms of computation cost. For instance, the algorithm proposed in [2] involves a complexity of $O(K \cdot n_1 \cdot n_2)$, where K is the number of atoms used in decomposition and $n_1 \times n_2$ is the image resolution, while a similar algorithm complexity is reported in [3]. Considering that the burden of manifold distance estimation is reduced significantly when the manifold is represented by discrete samples, our approach seems to offer an acceptable compromise between accuracy and computational effort.

5. ACKNOWLEDGMENTS

We would like to thank Ozan Şener for his contribution in the conduction of the experiments of Sec. 3.2.

6. REFERENCES

- [1] M. B. Wakin, D. L. Donoho, H. Choi, and R. G. Baraniuk, "The multiscale structure of non-differentiable image manifolds," 2005, vol. 5914, SPIE.
- [2] E. Kokkopoulos and P. Frossard, "Minimum distance between pattern transformation manifolds: Algorithm and applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 7, pp. 1225–1238, 2009.
- [3] N. Vasconcelos and A. Lippman, "A multiresolution manifold distance for invariant image similarity," *IEEE Transactions on Multimedia*, vol. 7, no. 1, pp. 127–142, 2005.
- [4] P. Simard, Y. LeCun, J. S. Denker, and B. Victorri, "Transformation invariance in pattern recognition-tangent distance and tangent propagation," in *Neural Networks: Tricks of the Trade*. 1998, pp. 239–27, Springer-Verlag.
- [5] Y. Linde, A. Buzo, and R. Gray, "An algorithm for vector quantizer design," *IEEE Transactions on Communications*, vol. 28, no. 1, pp. 84–95, Jan 1980.
- [6] A. Gersho and R. M. Gray, *Vector quantization and signal compression*, Kluwer Academic Publishers, Norwell, MA, USA, 1991.