Optimal Intrinsic Descriptors for Non-Rigid Shape Analysis

Thomas Windheuser
Matthias Vestner
Emanuele Rodolà
Rudolph Triebel
Daniel Cremers

Computer Vision Group,
Department of Computer Science,
Technische Universität München

We propose novel point descriptors for 3D shapes with the potential to match two shapes representing the same object undergoing natural deformations. These deformations are more general than the often assumed isometries, and we use labeled training data to learn optimal descriptors for such cases. Furthermore, instead of explicitly defining the descriptor, we introduce new Mercer kernels, for which we formally show that their corresponding feature space mapping is a generalization of either the Heat Kernel Signature (HKS) [3] or the Wave Kernel Signature (WKS) [1]. I.e. the proposed descriptors are guaranteed to be at least as precise as any Heat Kernel Signature or Wave Kernel Signature of any parameterisation.

A point descriptor $\phi : \mathcal{P} \to \mathbb{R}^T$ takes points from a set of shapes $\mathcal{P} := \cup_i \mathcal{M}_i$ and maps them to a space $\mathbb{R}^T$. Ideally, the descriptors of points that are at corresponding locations on the shapes should have a small distance in the descriptor space. Points at distinct locations on the shapes should be mapped to distinct locations in the descriptor space (see Figure 1).

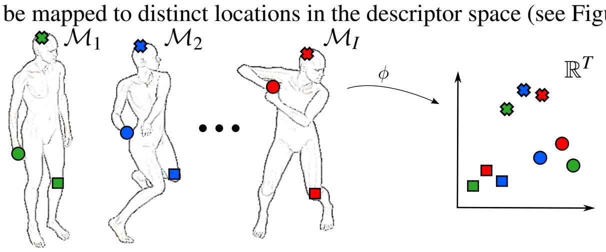

Figure 1: **Comparing points with a point descriptor:** Ideally the point descriptor $\phi : \mathcal{P} \to \mathbb{R}^T$ should map corresponding points to nearby locations and non-corresponding points to distinct locations.

In general one cannot assume that a given descriptor ϕ groups similar points as well as depicted in Fig. 1. The proposed method optimizes for the positive semi-definite matrix $M = L^T L$ inducing a pseudo distance in the descriptor space $\mathbb{R}^T$ via $d_M^2(x,y) = \langle x-y, x-y \rangle_M$ such that the point descriptors are grouped as good as possible. Optimizing for M is equivalent to looking for the best linear transformation L of the descriptor space with respect to the Euclidean distance, since $d_M(x,y) = \|L(x-y)\|$. In Figure 2 we see that L projects the images of ϕ onto the dotted line resulting in the much better descriptor $L \circ \phi$. As an optimization criterion for L we use LMNN [4] (see Figure 3).

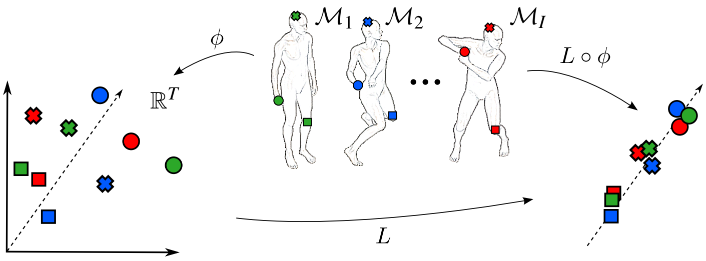

Figure 2: **Optimal distance in the descriptor space:** We are optimizing for a linear transformation L of the descriptor space $\mathbb{R}^T$. This is equivalent to a new distance $d_M^{\phi}(x,y) := \langle x-y, x-y \rangle_M$, where $M = L^\top L$, in the original descriptor space.

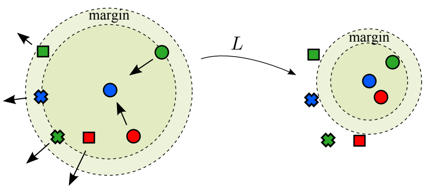

Figure 3: **Optimal LMNN distance [4]:** The neighbourhood of an input sample (blue circle) changes as a result of the training process. In this example, the learned distance is such that the nearest intra-class neighbours lie within a smaller radius after application of the linear mapping L. Similarly, the extra-class neighbours are left outside this optimized neighbourhood by a fixed margin.

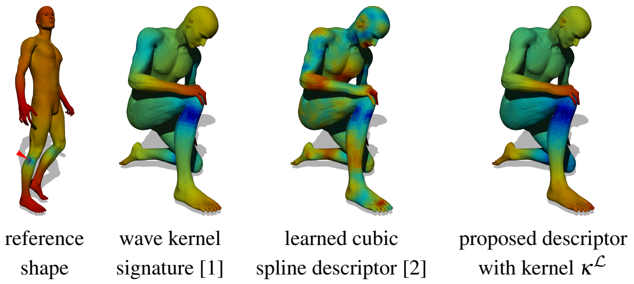

Figure 4: **Qualitative comparison of descriptors:** Distance maps between the descriptor at a reference point (indicated by a red arrow) and the descriptors computed on the shape after deformation. Colours range from blue (small distance) to red (large distance). Qualitatively, WKS and the proposed method do very well at indicating the right location while the cubic spline descriptor exhibits several local minima across the shape. Both test shapes are from the class *michael* (TOSCA), whereas the proposed descriptor and the spline descriptor were trained on the class *david*. The distances on the reference shape a generated by the proposed method.

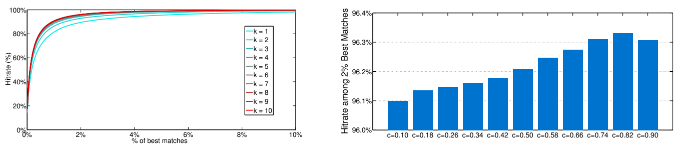

Figure 5: **Parameter sensitivity analysis** **Left:** Precision of the learned descriptor on the test set *michael* with kernel $\kappa^{\mathcal{L}}$, fixed $c = 0.5$ depending on different values for learning parameter k. Precision directly increases with higher values of k. **Right:** Same experimental setup as on the left, with fixed $k = 7$ and different values of c. The plot shows the hitrate when looking at the 2% best matches. The difference in precision among different values of c is only visible in this close-up.

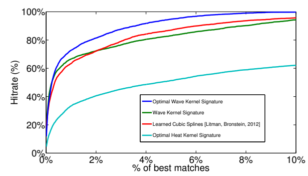

Figure 6: **Descriptor comparison** CMC curves of different descriptors on the TOSCA dataset; the method proposed in this paper with kernel $\kappa^{\mathcal{L}}$ is plotted as a blue curve.

The contributions of our paper can be summarized as follows:

- The method eliminates the need of tuning descriptor parameters. Neither does it have time parameters such as the HKS, nor do we need to choose the dimensionality of the descriptor as in [2]. In contrast, the adjustment of the descriptor is completely driven by the data, i. e. the shapes' deformations fed to the training process. The only two parameters of the objective function are directly related to the descriptor precision. Experiments suggest they can be fixed to constant values across applications, making the framework virtually parameter free.
- The method is a true generalization of the WKS and HKS and can potentially generalize other descriptors as well. Most importantly, we formally show that the proposed descriptors are guaranteed to be at least as accurate as WKS and HKS under any parameterisation with respect to the given shapes. Applications using WKS or HKS can avoid the parameter tuning problem by plugging in the proposed descriptor and are guaranteed to get optimal precision.

[1] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Computer Vision Workshops (ICCV Workshops)*, 2011.

[2] Roee Litman and Alexander M Bronstein. Learning spectral descriptors for deformable shape correspondence. *Transactions on Pattern Analysis and Machine Intelligence*, 1(1), 2013.

[3] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer Graphics Forum*, volume 28, pages 1383–1392, 2009.

[4] Kilian Q Weinberger and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. *The Journal of Machine Learning Research*, 10:207–244, 2009.