Non-rigid 3D Shape Retrieval via Large Margin Nearest Neighbor Embedding

Ioannis Chiotellis [1], Rudolph Triebel [1,2], Thomas Windheuser [1], and Daniel Cremers [1]

[1] Technical University of Munich [2] German Aerospace Center (DLR)

Computer Vision Group

Overview

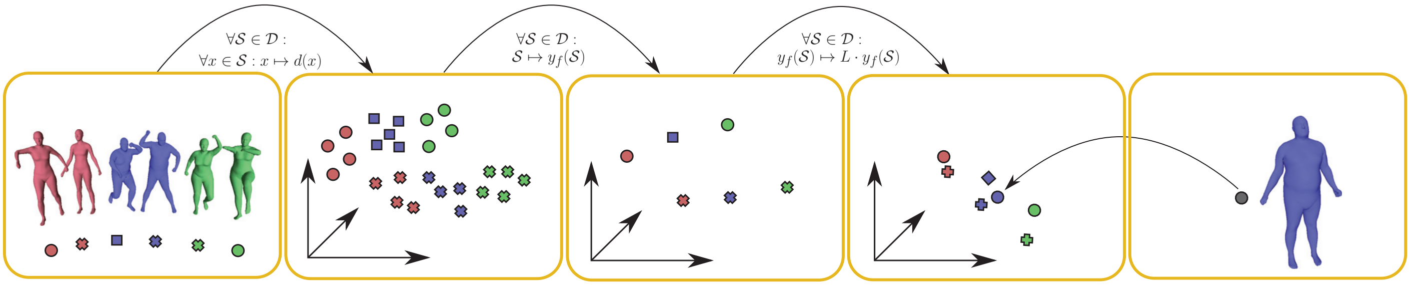

Collection of Labeled Shapes | Spectral Point Descriptors | Combined Spectral Descriptors | Large Margin Nearest Neighbor | Query via *k*-nearest neighbors

Objective: Optimal Shape Descriptors

Given: A set of non-rigid shapes with their corresponding class labels.

Goal: Efficient estimation of similarity between labeled and new unlabeled shapes.

- **Find a shape descriptor space where classification is easy.**
- **Efficiently compute a global shape descriptor for any new shape.**
- **Accurately determine the class label of the new shape.**

Contributions

- **An intuitive way to measure similarity between non-rigid shapes**
- **Small set of shapes needed for training, compared to other methods**
- **State of the art classification accuracy and runtime**

Shape Analysis Tools

Laplace-Beltrami Operator

- **Invariant under isometric deformations**
- **Defined as the negative divergence of the gradient:** $\Delta f := -\textbf{div}(\nabla f)$,
- **LBO Eigen-decomposition of a shape,** $\Delta\phi_i = \lambda_i\phi_i$ **with** $0=\lambda_0\le\lambda_1\le\lambda_2\le\ldots\le+\infty$.

Heat Kernel Signature (HKS) [1]

- **The amount of heat that remains at a point** x **after diffusion time** t **is** $k_t(x,x)=\sum_{k=0}^{K-1} e^{-\lambda_k t}\phi_k(x)^2$
- **Concatenating for different time stamps:** $HKS(x)=(k_{t_1}(x,x),k_{t_2}(x,x),\ldots,k_{t_T}(x,x))^T$

Scale Invariant Heat Kernel Signature (SI-HKS) [2]

- **Observation: The HKS of a scaled shape is only shifted in time.**
- **Taking the Fourier transform of the HKS, undoing the phase and sampling at** q **frequencies:**

$$siHKS(x)=(|H(\omega_1)|,\ldots,|H(\omega_q)|)^T$$

Wave Kernel Signature (WKS) [3]

- **The avg. probability over time to locate a quantum particle at point** x **is:**

$$p(x)=\lim_{T\to\infty}\frac{1}{T}\int_0^T|\psi_E(x,t)|^2=\sum_{k=0}^{\infty}\phi_k(x)^2 f_E(\lambda_k)^2$$

- **Different choices of** f_E **give us shape properties at different scales:**

$$WKS(E,x)=(p_{e_1}(x),\ldots,p_{e_q}(x))^T$$

Large Margin Nearest Neighbor [4]

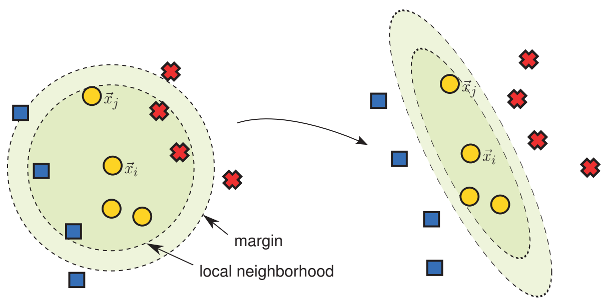

Euclidean Metric | Mahalanobis Metric

- **Finds a Mahalanobis distance that is optimal for *k*-nearest neighbors classification.**
- **Enforces maximum margin between intra-class and inter-class samples (as in SVMs).**
- **The learned descriptor space can be expressed as a linear transformation** L **of the inputs.**
- **A convex formulation of the objective function guarantees the optimal metric** $M=L^TL$**:**

$$\varepsilon(M)=(1-\mu)\sum_{i,j\rightsquigarrow i}D_M(\vec{x}_i,\vec{x}_j)+\mu\sum_{i,j\rightsquigarrow i}\sum_{l}(1-y_{il})[1+D_M(\vec{x}_i,\vec{x}_j)-D_M(\vec{x}_i,\vec{x}_l)]_+,\quad \mu\in[0,1]$$

Presented at the European Conference on Computer Vision, Amsterdam, 2016.

Combined Spectral Descriptors and LMNN

Observations

- **SI-HKS captures global shape properties, while WKS focuses on local shape features.**
- **A combination should describe each shape in a more unique way.**
- **Comparing point descriptors is intractable for retrieval.**

Combined Spectral Descriptors

Weighted Average: We compute a weighted average over all point descriptors $d(x)$ **computed from the points** x **of a given shape** $\mathcal{S}$**:**

$$y_f(\mathcal{S})=\sum_{x\in\mathcal{S}}w_x d(x)\quad\text{with}\quad w_x=\frac{a_x}{\sum_{y\in\mathcal{S}}a_y}\quad\text{and}\quad y_{CSD}(\mathcal{S})=\begin{pmatrix}y_{SIHKS}(\mathcal{S})\\ y_{WKS}(\mathcal{S})\end{pmatrix}$$

where a_x **is the area element associated with vertex** $x\in\mathcal{S}$.

Experimental Results

Space of Shapes Visualization

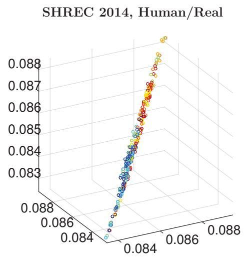

dimensions $1,2,3$ of the shape descriptors $y_f(\mathcal{S})$ before learning

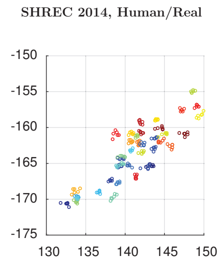

dimensions $1,2$ of the transformed shape descriptors $L\cdot y_f(\mathcal{S})$

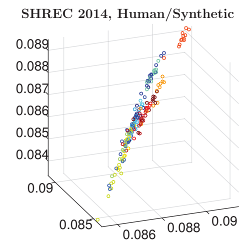

dimensions $1,2,3$ of the shape descriptors $y_f(\mathcal{S})$ before learning

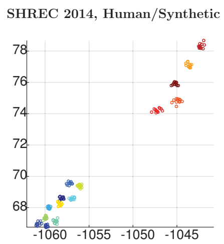

dimensions $1,2$ of the transformed shape descriptors $L\cdot y_f(\mathcal{S})$

MAP Comparison (%)

Method	Synthetic	Real
ShapeGoogle(VQ) [5]	81.3	51.4
Unsupervised DL [6]	84.2	52.3
Supervised DL [6]	95.4	79.1
RMVM [7]	96.3	79.5
CSD	82.67	50.75
CSD+LMNN	**99.67**	**97.92**

Code available at `https://github.com/tum-vision/csd_lmnn`

References

[1] 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. Wiley Online Library, 2009.

[2] Michael M Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1704–1711. IEEE, 2010.

[3] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *ICCV Workshops, 2011*, pages 1626–1633. IEEE, 2011.

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

[5] Alexander M Bronstein, Michael M Bronstein, Leonidas J Guibas, and Maks Ovsjanikov. Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics (TOG)*, 30(1):1, 2011.

[6] Roee Litman, Alex Bronstein, Michael Bronstein, and Umberto Castellani. Supervised learning of bag-of-features shape descriptors using sparse coding. In *Computer Graphics Forum*, volume 33, pages 127–136. Wiley Online Library, 2014.

[7] Andrea Gasparetto and Andrea Torsello. A statistical model of riemannian metric variation for deformable shape analysis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.