

SCREENED POISSON HYPER-FIELDS FOR SHAPE CODING

R. A. GULER [†], S. TARI [‡], AND G. UNAL ^{§¶}

Abstract. We present a novel perspective on shape characterization using the screened Poisson equation. We discuss that the effect of the screening parameter is a change of measure of the underlying metric space; also indicating a conditioned random walker biased by the choice of measure. A continuum of shape fields is created, by varying the screening parameter or equivalently the bias of the random walker. In addition to creating a regional encoding of the diffusion with a different bias, we further break down the influence of boundary interactions by considering a number of independent random walks, each emanating from a certain boundary point, and the superposition of which yields the screened Poisson field. Probing the screened Poisson equation from these two complementary perspectives leads to a high-dimensional hyper-field: a rich characterization of the shape that encodes global, local, interior and boundary interactions. To extract particular shape information as needed in a compact way from the hyper-field, we apply various decompositions either to unveil parts of a shape or parts of boundary or to create consistent mappings. The latter technique involves lower dimensional embeddings, which we call Screened Poisson Encoding Maps (SPEM). The expressive power of the SPEM is demonstrated via illustrative experiments as well as a quantitative shape retrieval experiment over a public benchmark database on which the SPEM method shows a high-ranking performance among the existing state-of-the-art shape retrieval methods.

Key words. Screened Poisson equation, Elliptic models for Distance Transforms, conditioned random walker, shape decomposition, Screened Poisson Encoding Maps (SPEM), non-negative sparse coding, non-rigid shape retrieval, level-set models.

AMS subject classifications. Second-order elliptic equations;

1. Introduction. Geometric information regarding the shape of objects is a significant component of visual information, which is one of the main sensory inputs utilized in our perception of the world. The question of how to best represent a shape mathematically for its use in artificial intelligence systems has been studied for many decades. Computational vision problems such as object recognition require a shape representation that should be primarily: well-descriptive of the object geometry; invariant with respect to a certain geometric transformation group for robustness; compact for efficient computation and storage. Many different approaches to represent an object's geometry have been proposed including mainly medial axes-based, boundary- or surface-based, and region- or volume-based shape representations. In this paper, we present a novel shape representation, where shape information is encoded inside the shape by exploiting internal distance relationships via the screened Poisson equation.

In the 90's, it was observed that shapes can be embedded as zeros of a function defined over the shape domain, opening the way to an active research area in implicit shape representations ([1], [2]). In this area of research, the signed distance transform was popularized heavily by the level-set framework [3] and its fast implementation [4]. The distance function is created via solution of the Eikonal equation $|\nabla u(x)| = 1, x \in \Omega$ subject to boundary condition $u|_{\partial\Omega} = 0$. The governing equation forces the absolute value of the gradient to be constant. Equipped with a suitable

[†]R.A. Guler is with the Faculty of Engineering and Natural Sciences, Sabanci University, Istanbul Turkey.

[‡]S. Tari is with the Computer Engineering Department, METU, Ankara, Turkey.

[§]G. Unal is with the Faculty of Engineering and Natural Sciences, Sabanci University, Istanbul Turkey.

[¶]Corresponding Author: gozdeunal@sabanciuniv.edu

^{||}This work was supported by TUBITAK Grant No: 112E320.

boundary condition, the solution $u(x)$ is interpreted as the shortest time needed to travel from the boundary to the point x . Signed Distance Transform (SDT) is formed by setting positive and negative of the distances exterior and interior to the shape or vice versa, facilitating regional encoding of shape domain and its exterior by minimal distances to the shape boundary. The shape is then represented as the zero level set of the signed distance transform (SDT). This representation of the shape, *i.e.* via embedding the shape boundary as the level set of SDT, became quite instrumental in developing approximate schemes for segmentation functionals and introducing shape knowledge in segmentation problems, *e.g.*, [5, 6, 7, 8, 9].

In the late 90's and the following decade, elliptic PDEs started to appear as alternative models for computing smooth distance fields. In [10], screened Poisson PDE is employed:

$$(1.1) \quad \begin{aligned} \Delta v - \frac{v}{\rho^2} &= 0 \\ v|_{\partial\Omega} &= 1, \end{aligned}$$

where $\frac{1}{\rho^2}$ is the screening parameter that controls the level of smoothing. The approximate distance field created by this PDE is smooth and differentiable, and has smooth level sets, in contrast to the level sets of the distance transform obtained from the Eikonal equation. With any given ρ value, the field's value of 1 at the shape boundary drop towards the interior of the shape. While a motivation in [10] was to create a shape scale space, demonstrated particularly for shape skeletons via the controlled smoothing parameter, in [11], the intuition of a random walker starting at an interior point and its mean hitting time required to reach the shape boundary led from its discrete interpretation to the continuous Poisson equation with zero Dirichlet boundary conditions on the shape boundary. Various measures based on the solution field were extracted and shape properties were used for classification of shapes as well as actions [12]. In [13], the authors also utilized the Poisson equation to derive a shape characteristic measure based on the variation over the streamlines of the solution field, and used it to differentiate between the shapes of anatomical structures for healthy and diseased populations. Recently, Poisson equation is revisited as a tool for robust skeletonization [14, 15]. In [16, 17], a connection between nonlinear Hamilton-Jacobi equations, for which the Eikonal equation is a special case, and the screened Poisson equation by taking $\rho \rightarrow 0$ is presented, along with an efficient approximate distance transform computation using FFT. The importance of the linearity of shape embedding space was brought into attention by the work of [18], [19] that represented contours as zero level set of a harmonic function in the solution of the Laplace PDE. The linearity property, which was also emphasized in [16], enables proper addition of shape fields, facilitating creation of shape template or atlas representations that stay within the original spaces of shapes. [20] solved heat flow with a fixed time parameter and used its normalized gradient field to obtain the closest scalar potential field with the same gradient. In [21], smooth distance fields are considered as L_p distance fields, where p is the control variable. A recent shape field related to the screened Poisson [22] is a fluctuating one consisting of both negative and positive values inside the shape by addition of a zero-mean constraint to the shape field. The zero-level set then partitioned the shape domain into two: one that corresponds to the central region, a coarse and compact shape, and one to the peripheral region, which included protrusions from a shape.

The screened Poisson PDE was employed for several other applications with a typically fixed screening parameter: for image processing applications as in image

filtering and sharpening of [23]; for mesh filtering applications as in anisotropic and interactive geometric filtering over meshes of [24]; and for surface reconstruction in [25]. [23] started from a variational perspective by writing out the gradient of an unknown function to be close to a given vector field as well as a term of data fidelity to a given function which "screens" the 2D Poisson equation. This was then Fourier transformed to show that the screened Poisson can be interpreted in frequency domain as a filtering operation for images, while it can be solved using an FFT or DCT. [24] extended [23] to meshes for localised editing by changing the Riemannian metric of the underlying space, proportional to surface curvature, as well as a multi-grid implementation of the equation. The effect of the fidelity value, i.e. the screening parameter, was also discussed to result in more dampening and amplification at low frequencies with smaller parameter values. [25] modified this method by putting positional constraints, i.e. the data fidelity, only over a set of input points rather than over the full domain. Adding a screening term to the Poisson surface reconstruction framework, the screening parameter was also adjusted to the resolution in a multi-grid implementation.

In a parallel line of research, from the heat equation perspective, the multi-scale property of the heat kernel led to development of shape signatures that take advantage of heat diffusion process on surfaces [26]. This line of work makes use of the spectral properties of the Laplace-Beltrami operator, which is the generalization of the Laplace operator from the Euclidean space to a Riemannian manifold. In [26], the heat kernel signature (HKS) at a point on the shape manifold is defined in terms of the weighted sum of the squares of the eigenfunctions at the point. The weights are given by the exponentials of the negated eigenvalues multiplied by the temporal variable t in heat flow. It was shown that under certain conditions (*i.e.*, if the eigenvalues of the operator are not repeated) the heat kernel signature is as informative as the family of heat kernel functions parameterized both in space and time. The HKS also relates to global point signatures [27], which are based on eigenfunctions normalized by square root of the corresponding eigenvalues, and to diffusion distance [28, 29, 30] between two points over the shape manifold, which is defined by the distances between the eigenfunctions at those two points. [31] constructs a scale-invariant HKS (SI-HKS) by logarithmically sampling the time-scale that translates into a time shift, which is then removed through taking a Fourier transform modulus to overcome the scale sensitivity of HKS. A volumetric extension of HKS was shown in [32].

Recently, Wave Kernel Signature [33] based on complex Schrodinger equation is presented as an alternative to HKS. The authors make the point that HKS employs a collection of low-pass filters parameterized by time variable, causing the suppression of high frequency shape information whereas the WKS captures both the high and the low frequency shape information.

Meanwhile, works such as Shape DNA [34] showed the utility of the eigenvalues of the Laplace operator, where the distances between shapes were expressed as the p -norm of the difference between the truncated eigenvalue sequences for the two shapes. In [35], a normalized shape DNA distance, called the weighted spectral distance is proposed.

Laplace-Beltrami eigenfunctions of surfaces proved to be extremely useful in applications of 3D shape matching and retrieval. In [36], it was shown that a bijective mapping between a given pair of shapes induces a transformation of a function of derived quantities between them. Furthermore, this transformation can be written as a linear map between selected basis functions over both surfaces, exemplified by the

Laplace-Beltrami eigenfunctions. [37] presented a method to perform shape matching in a reduced space in which the symmetries of shapes were identified and factored out. This was achieved within the functional maps framework of [36] where the functional linear map was decomposed into its symmetric subspace and its orthogonal subspace, and the former was utilised to carry out the shape matching between symmetric shapes. For joint analysis of multiple shapes, [38] presented a coupled construction of common Laplacian eigenfunctions using approximate joint diagonalizations.

In [39], a shape-aware interior-mesh distance was defined by propagating a distance measure defined on the mesh to the surface interior, while preserving distance properties. This was exemplified by the diffusion distance and mean-value coordinates selected as the barycentric coordinates. [40] later applied this idea to interpolating the Laplace-Beltrami eigenfunctions of the boundary into the interior volume by using barycentric coordinates. This way, a volumetric measure was constructed from the HKS, i.e. the interior HKS, and adopted to finding correspondences between volumes and shape retrieval.

1.1. Our Contribution. In this paper, we present a novel perspective on shape characterization using the screened Poisson equation. Both the Poisson and the screened Poisson equations found increased utility in various shape descriptors. As the screening parameter in (1.1) tends to ∞ , the screened Poisson equation approaches to the Poisson equation. The controlled smoothing provided by the screening parameter is advocated by some researchers and recent works [23, 15, 17, 22, 20, 25] rejuvenated the model.

Our work differs in several aspects. We consider multiple instances of the screened Poisson equation to decompose the sources of variability due to several factors including the boundary sources and the screening parameter, both of which are novel. We discuss that the effect of screening parameter is a change of measure of the underlying metric space, hence, fixing ρ^2 fixes the measure. Suitably sampling N values for the screening parameter and m points for the shape boundary $\partial\Omega$, we form a stack of $N \times m$ screened Poisson fields. We call this collection as a screened Poisson hyper field. This is not a scale-space in the usual sense but hides in it a two-dimensional scale space of shapes, coarsening in the direction of increasing ρ^2 and decreasing field values. We argue that the hyper field is a full characterization of all sorts of interactions between shape elements: local-global and boundary-interior. Then we discuss two low-dimensional embedding schemes, one to unveil parts and the other to produce consistent mappings, which we call Screened Poisson Encoding Maps (SPEM), for the purpose of shape matching and shape retrieval.

Encoding a change in the diffusion using the varying screening rates in the screened Poisson equation forms a remarkable parallelism with the class of methods in spectral shape analysis. We argue that a coverage of ρ^2 parameter space for (1.1) over the shape domain brings advantages over the coverage of the temporal parameter space for the heat kernel over the shape in terms of producing a direct volumetric shape representation. [32], extending the heat kernel signatures to volumes, noted that the boundary isometries of the HKS do not carry over to volume isometries, however, volumetric HKS can still faithfully model nearly isometric deformations, which are argued to be more suitable in modelling of natural articulations and deformations of solid objects. On the other hand, [40] propagated HKS on the surface towards the interior of the shape to be able to construct volumetric measures to benefit from nice properties of the HKS, including the surface isometry, the multi-scaleness and insensitivity to topological noise at small scales however at the

expense of its sensitivity to scale of the shape [32, 41]. Different from these earlier heat-kernel based approaches, here, we directly compute volumetric distances from the solution to the volumetric screened Poisson PDE, which enjoys properties such as multi-scaleness based on a varying screening-parameter that tunes smoothness of the level curves of the field, an adaptation to scale by an appropriate mapping, and a near isometry-invariance as demonstrated experimentally by the robustness of the proposed method under a 3D nonrigid shape retrieval application (§ 5.5).

As an alternative to the heat equation and its kernel, our work presents a different differential operator, a different kernel, and demonstrates the high-ranking performance of the SPEM to articulated pose and deformation in a publicly available large-scale benchmark data set: *SHREC'11 Shape Retrieval on Non-Rigid 3D Watertight Meshes* [42]. Our method, as shown in the presented 3D shape retrieval application, provides a robust and high-performance alternative to those methods based on shape's intrinsic surface properties. Furthermore, existence of fast solvers for the screened Poisson PDE, as realised by [23, 24, 17, 25] in other applications of image filtering and mesh processing, is another factor that makes it attractive to be adopted in a new shape representation idea as in this paper.

The organization of the paper is as follows. In § 2, we show separation of the sources of variability in the v -field, and present the construction of the new shape hyper-field. We expound properties of the new hyper-field and the SPEM in § 3 through a random walk interpretation, relation to geodesic distances, and a connection to spectral methods. In § 4 we present how decompositions on shape hyper fields via two alternative techniques produce consistent mappings and part decompositions. Finally in § 5, we present our experimental results followed by conclusions.

2. A new hyper-field. In this section, first the existing two-dimensional scale space parameterized by ρ and the values of v (§ 2.1) is explained. Then we describe the two dimensions of the new shape representation: the varying of the ρ (§ 2.2), and decomposition of the boundary sources (§ 2.3). The new hyper-field thus includes two dimensions of variability: (i) by variation of ρ , it covers the internal smoothing characteristics of v ; (ii) by variation of boundary sources, it covers interactions between individual boundary nodes versus all internal nodes. We note that the decomposition into those two dimensions do not create a true scale-space per se, however, creates a rich shape hyper-field representation from which descriptive volumetric shape encoding maps (SPEM) can be extracted.

2.1. A two-dimensional Scale Space. The information encoded in the resulting field v of Eq.(1.1), as a shape representation, is highly dependent on the value of ρ^2 . The influence of the parameter ρ^2 can be observed in Fig. 1, where different fields that arise using different ρ^2 values are presented for a cat shape. Smaller ρ^2 values lead to fields where distinct relations in the regions that are close to shape boundary (protrusions, indentations) are extracted, but are clueless about the central part of the shape and global interactions. In contrast, larger values of ρ^2 generate fields that are coarse in the regions close to the boundary, but able to capture global interactions within the shape. Unlike the level curves of the solution of the Eikonal Equation, the level curves of v (the solution of screened Poisson Equation) has smooth level sets, and as one moves along the gradient lines, the level curves gets smoother. As discussed in [43, 10]

$$(2.1) \quad v(\mathbf{x}) \approx \rho \left(1 + \frac{\rho}{2} \text{curv}(\mathbf{x}) \right) \frac{\partial v}{\partial n} + O(\rho^3)$$




Fig. 1: v fields for different values of ρ^2 .

where $\text{curv}(\mathbf{x})$ is the curvature of the level curve of v passing through the point \mathbf{x} at \mathbf{x} , and n is the direction of the normal. Thus, one can imagine a two-parameter family of level curves parameterized by v and ρ . Smoothing increases with a decrease in v and an increase in ρ^2 . This is a very interesting property. This explains how the linear screened Poisson mimics a non-linear reaction-diffusion. Though this observation was made in the early work of [10], the follow up work on screened Poisson typically focused on isolated treatment of the ρ^2 . Rangarajan [17, 16] took a very small value to approximate the Eikonal Equation, while Tari [15, 22] and Shah [44] used very large values.

We believe that isolated treatment is hindering full utilization of the controlled smoothing offered by the model. As we show in §4.2, once the entire scale space is utilized, both local and global interactions can be realized and a natural hierarchical central to peripheral decomposition of the shape domain is achieved without requiring the recent non-local term in [22].

2.2. Varying ρ^2 : Sweeping Internal Smoothing Characteristics. In a setting where the screening parameter is considered as an additional dimension to the spatial ones, it is clear that the $n + 1$ -Dimensional field calculated for a shape embedded in \mathbb{R}^n , where the parameter ρ^2 is swept from 0 to ∞ inherits all the information that is possible to be extracted using such a method about the shape. The collection of fields $\{v^\rho\}_{\rho^2}$ consists of a 1D family of functions that sweeps the ρ^2 dimension for each node on the lattice that the shape is described on. A field created using only one of these values would explain only limited portion of the variance. In order to capture this high dimensional information, we linearly sample ρ^2 values to N bins, and calculate v^ρ for each ρ_j^2 value for $j = 1, \dots, N$. Each v^ρ field as a single instance explain relatively little variation of the shape in comparison to the whole family.

We depict via an example that the v function is coding characteristics that extend beyond the distance to the nearest boundary point as well as curvature (Fig 2). We consider several nodes in a shape domain. They are marked with colored crosses. Each node has a different character: The blue one is central; the other four are closer to shape boundary, pink being the closest; and red is on a peripheral part (finger). The v versus ρ plot on the right depicts striking differences among $v(\cdot, \rho)$ profiles for these different shape nodes. For example, the two points colored red and pink respectively have closer profiles as they have comparable proximity to the shape boundary. However, the profiles are not nearly identical because the red node is residing in a thin part of the shape while on the contrary the pink one is not.

Fig. 3 demonstrates further coding characteristics of the v^ρ field. A set of 1D profiles ($v^\rho(\tilde{\mathbf{x}})$) for a set of locations $\tilde{\mathbf{x}}$ on a hand silhouette are depicted. Here, the point we emphasise is that the selected locations, $\tilde{\mathbf{x}} \in \Omega$ are equidistant to the shape boundary. Observe that the 1D curve describing the relation between v and ρ shows a quite different character for each point, which has the same Euclidean distance to the boundary, while the v -field is able to encode the diversity of the geometric shape




Fig. 2: Field value versus ρ at five selected nodes of distinct characters. v function is coding characteristics that extend beyond usual distances. A dense linear sampling is used between $\rho = 2$ and $\rho = 30$.

information among those points.




Fig. 3: Behavior of v^ρ in ρ dimension for sampled points on the domain that are equidistant to the boundary.

2.3. Fixing ρ^2 : Decomposition of Boundary Sources. The two-dimensional scale space is a continuous collection of simple closed curves parameterized by $[1, 0) \times \{1, 2, \dots, N\}$. For a fixed screening parameter, a one-dimensional scale space is formed by the collection of the level curves of the field v^ρ , which is a union of these level curves. This is not the only way to envision v^ρ . Thanks to the linearity of the equation, it is also possible to express v^ρ as a superposition of basis fields each of which is expressing the contribution due to a single “unit” of inhomogeneity.

In order to elaborate on the super-positioning aspect of the screened Poisson PDE for a fixed ρ^2 and better understand geometric properties induced by boundary interactions, we consider decomposing the sources of inhomogeneity in the boundary condition. Assuming that the shape boundary is given as a set of points $\partial\Omega = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_m\}$, we consider m independent PDEs:

$$(2.2) \quad \begin{aligned} \Delta v^{\mathbf{p}_i}(\mathbf{x}) - \frac{v^{\mathbf{p}_i}(\mathbf{x})}{\rho^2} &= 0 \\ v^{\mathbf{p}_i}(\mathbf{p})|_{\mathbf{p} \in \partial\Omega} &= \delta(\mathbf{p} - \mathbf{p}_i) \end{aligned}$$

where $v^{\mathbf{p}_i}$ denotes the solution when the only inhomogeneity is due to the point $\mathbf{p}_i \in \partial\Omega$. Thanks to the linearity of the equation, these “sub-fields” are the building

blocks that make up the field v described in (1.1):

$$(2.3) \quad v = \sum_{i=1}^m v^{p_i}$$

The super-positioning of the sources is demonstrated on a 1D example in Fig.4. The boundary condition on the third column is an addition of the two boundary conditions used in the first two columns. Hence the solutions in the third column are superpositions of the pair of solutions given on the respective row of the first two columns.

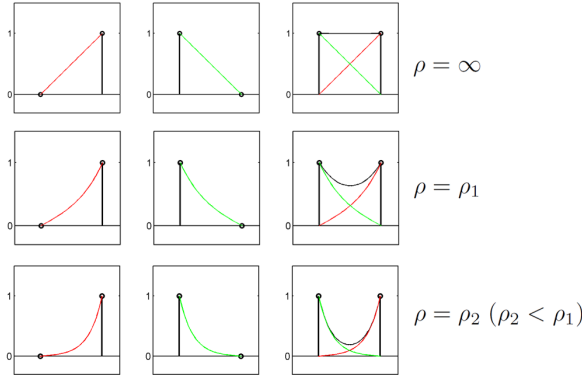


Fig. 4: Solutions of the screened Poisson equation for a 1D experiment using three different boundary conditions (columns) and three different ρ values (rows).

In Fig. 5, the logarithm of the field v^{p_i} obtained from a boundary point p_i on the hand shape is visualized on the left. It can be observed that the v field shows a sharp fall of its values over the fingers whereas a much less steep slope of fall is observed from the boundary points of the hand's side palm regions (e.g. close to the wrist). This different behaviour is expected. To analyse it on a simpler case, assume a spherical geometry with a source term $s(\mathbf{r})$ at the origin, and consider the Poisson equation: $\Delta v = s(\mathbf{r})$, the fundamental solution is: $g(\mathbf{r}) \propto \frac{1}{r}$, whereas for the screened Poisson equation: $\Delta v - \frac{v}{\rho^2} = s(\mathbf{r})$, the fundamental solution reads: $g(\mathbf{r}) \propto \frac{e^{-\frac{r}{\rho}}}{r}$ [45]. Hence with a nonzero source term the solution is given by:

$$(2.4) \quad v(\mathbf{r}) = \int_{\Omega} d\mathbf{r}' s(\mathbf{r}') \frac{1}{|\mathbf{r} - \mathbf{r}'|} e^{-\frac{|\mathbf{r} - \mathbf{r}'|}{\rho}}.$$

For a spherical symmetric case, the source is diffused to its surrounding points by a convolution with a kernel inversely proportional to the distance between the source and the given point for the standard Poisson case, whereas for the screened Poisson the convolving kernel is in addition weighted by a decaying exponential. Although for the arbitrary geometric configuration of our boundary conditions we cannot write an integral equation to solve for the result, we can observe the exponential decay effect in our v -field from a single source point to other points. With a union of all boundary sources, the effect is even more pronounced. Similarly in § 2.2, we changed the rate of decay by varying ρ to probe this property. We will further discuss the relation of the v -field to geodesic distances in § 3.2.

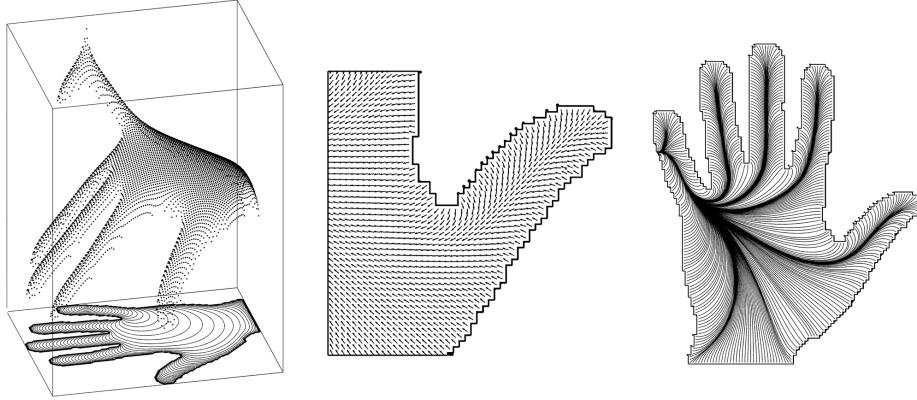


Fig. 5: Restricting the boundary inhomogeneity to a single point \mathbf{p}_i on the little finger. a) Iso-contours (bottom) and values of v -field using $\log(v^{\mathbf{p}_i})$ visualised as a point cloud; b) Normalized gradient, $\frac{\nabla_v \mathbf{p}_i}{|\nabla_v \mathbf{p}_i|}$ for the 'thumb'; c) Streamlines obtained by tracking along the normalized gradient directions.

2.4. Putting it altogether: The New Hyper-field. By considering a total of $N \times m$ screened Poisson equations, we form a stack of fields. This stack of fields hides separation of several sources of variability due to all kind of interactions: local, global, region and boundary. The schematic depiction is given in Fig. 6. Intuitively, this can be best explained as simultaneous decomposition layers.

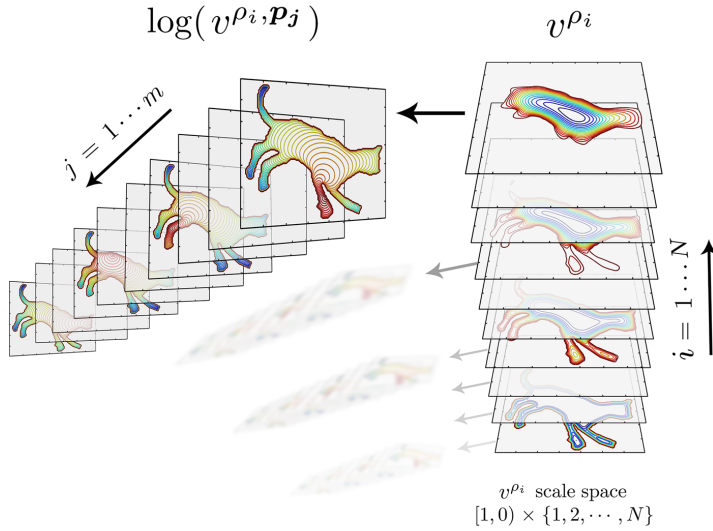


Fig. 6: Separating sources of variability in the shape hyper-field.

In the first decomposition layer, ρ^2 is varied to obtain a stack of fields $\{v^{\rho_i}\}_{i=1 \dots N}$. Each slice in the stack is an interpretation of the shape with a certain bias – choice of measure, and is a collection of shape boundaries embedded as level curves hence parameterizable by a continuous parameter $s \in (0, 1]$. This is the second layer of de-

composition. The stack of fields $\{v^{\rho^i}\}_{i=1\dots N}$ as parameterized by $(0, 1] \times \{1, \dots, N\}$ defines a 2D scale space of shapes, coarsening in the direction of increasing ρ^2 and decreasing s . At the final decomposition layer, the effect of inhomogeneity (note that the solution to the PDE in (1.1) is the trivial solution in the absence of inhomogeneities) is individuated by considering m -fields $v^{\rho^i} \mathbf{P}^j$ for $j = 1 \dots m$, which add up to $v^{\rho^i} = \sum_{j=1}^m v^{\rho^i} \mathbf{P}^j$. This last layer added by the boundary source sweep is built on top of the nonlinear scale-space of ρ and level curves of v , hence maintains a more complex structure.

The hyper field provides a rich characterization of the shape. We will present how to extract this information in a robust way in §4.

3. Screened Poisson: Properties.

3.1. Screened Poisson as a conditioned random walk. In this section, we expound the underlying stochastic interpretation of the v^ρ field in order to gain more intuition into its coding properties. Specifically, we are interested in understanding better the effects of 1) the change of ρ , and 2) the boundary interactions.

First, let us shift the inhomogeneity in the boundary condition in (1.1) to the right hand side as a source term, and then consider an inhomogeneous heat equation: $(\Delta + \frac{\partial}{\partial t})u(\mathbf{x}, t) = f(\mathbf{x})$. On one hand, the steady state solution as $t \rightarrow \infty$ is $1 - v^\rho$ for $\rho \rightarrow \infty$ (i.e., the solution of the Poisson Equation). On the other hand, the transient solution is

$$(3.1) \quad u(\mathbf{x}, t) = \int p_t(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) d\mu(\mathbf{y})$$

where μ is the Lebesgue measure and $p_t(\mathbf{x}, \mathbf{y})$ is the transition probability from point \mathbf{x} to \mathbf{y} in time t . The transition probability (also called heat kernel) is given by the Gaussian function:

$$(3.2) \quad p_t(\mathbf{x}, \mathbf{y}) = \frac{1}{(4\pi t)^{d/2}} \exp\left(-\frac{|\mathbf{x} - \mathbf{y}|^2}{4t}\right)$$

Now let μ be a measure on a Riemannian manifold \mathcal{M} . The inhomogeneous heat equation with the corresponding Laplace (-Beltrami) operator on the manifold is

$$(\Delta_\mu + \frac{\partial}{\partial t})u(\mathbf{x}, t) = f(\mathbf{x})$$

The transient solution is given by (3.1). Let us examine the effect of screening following Grigoryan [46]. We note that introducing screening to the Poisson equation corresponds to a change of measure. Let $\tilde{\mu}$ be the new measure, then $\Delta_{\tilde{\mu}}$ is related to $(\Delta_\mu - \frac{1}{\rho^2})$ by the Doob h-transform.

$$(3.3) \quad \Delta_{\tilde{\mu}} = \frac{1}{h} \circ (\Delta_\mu - \frac{1}{\rho^2}) \circ h \longrightarrow \Delta_{\tilde{\mu}} v = \frac{1}{h} (\Delta_\mu - \frac{1}{\rho^2})(vh)$$

To summarize, diffusion is a stochastic Markov process, indeed a Brownian motion with heat kernel as its transition probability. In the case of the diffusion governed by screened Poisson, the new transition kernel \tilde{p}_t that relates to the original transition kernel is the heat kernel on the Riemannian manifold with measure $d\tilde{\mu} = h^2 d\mu$ [46]. For a random walk on a network, when $p_t(\mathbf{x}, \mathbf{y})$ is induced by conductances c_{xy} , then \tilde{p}_t is induced by conductances $\tilde{c}_{xy} = h(x)h(y)c_{xy}$ [47] [48]. This means that the

conditioned random walk behaves like the unconditioned walk but is biased by an isotropic drift h .

The conditioned random walk with a certain ρ^2 -value affects a point in the shape domain with a certain bias, making it possible to probe multiple random walkers going through different conductances over the shape. We believe that this is how the continuum of fields encodes the shape characteristics both locally and globally with its varying screening rates or biases. This can also be interpreted as Brownian motions with different drift amounts, the zero drift corresponding to the unconditioned random walk, hence pure diffusion without any screening term of the standard Poisson equation.

For a fixed ρ^2 , the field $1 - v^\rho$ is a superposition of multiple random walks on a manifold with a measure $\tilde{\mu}$. Note that the transient solution (after the change of measure) for the time-dependent equation then would be given by

$$(3.4) \quad 1 - v^\rho(\mathbf{x}, t) = h_\rho(\mathbf{x}) \int \tilde{p}_t(\mathbf{x}, \mathbf{y}) f(\mathbf{y}) h_\rho(\mathbf{y}) d\mu(\mathbf{y})$$

At the steady-state, the transition kernel becomes only a function of distance independent of t . Thus, separating the boundary condition to a set of points, and solving the screened Poisson PDE for each single point as in Eq. (2.2), each field value $v^{\mathbf{p}_i}(\mathbf{x})$ (after a normalization) is interpreted as the probability that the biased random walker emanating from \mathbf{p}_i to arrive at the locations \mathbf{x} . We note that the intuition of the boundary condition on a random walker was mentioned by [20] for the heat flow, with the zero Dirichlet boundary condition implying absorption of heat that leads to random walker "falling off" the grid. With this interpretation, the way we set the point source on the boundary to unity while setting all other boundary points to zero implies that the probability of the walker falling off the grid differs substantially for different local geometric regions of the shape (see Fig. 3 for this effect).

3.2. Relation to Geodesic Distances. There is a strong link between the values of the $v^{\mathbf{p}_i}$ field and the geodesic distance from \mathbf{p}_i to another shape node, with the underlying Riemannian metric. A prominent aspect that forms this link is the gradient directions of $v^{\mathbf{p}_i}$, which are parallel to geodesics. The choice of boundary conditions configures the resulting gradient field. For instance, Dirichlet boundary conditions attract the flux to the medial locus. In Fig. 5 (middle and right figures), we show normalized gradient directions along with streamlines obtained by tracking points in the gradient direction. The link between the heat flow kernel (i.e. the $\rho \rightarrow \infty$ case) and geodesic distances was established by Varadhan: $-\sqrt{-4t \log(p_i)}$, where t corresponds to the amount of time that passes after heat diffusion starts [49]. Simply taking the logarithm of the v field leads to an encoding of the local relationships in a rather useful manner and preserves the gradient directions. The choice of logarithm stems from the exponential decay of the field (Eq. 2.4) also noted by [10], [16], and the logarithm of the field values become strictly negative, decreasing as the probabilities for the random walkers get less likely. We note that this is not an attempt to make the v -field values similar to Euclidean distances. Taking the square-root as in Varadhan's formula [49] also preserves the gradient directions but suppresses high rates of decay. This sort of treatment would compromise a very desirable property for part based analysis of shapes: at nodes that belong to articulated regions on the shape domain, as the probabilities for random walkers to go off the grid increases, the rate of decay increases drastically. This property was observed in Fig. 5 (on the left). Therefore, with the v -field, we are exploiting an exponential decay effect with a complementary

contribution from the shape boundary conditions, to construct a beneficial "geodesic distance" from the given shape geometry. Observe the effect of this complementary contribution in Fig. 3, where the points that have equal Euclidean distances to the boundary have v^ρ field values which encode a geodesic distance that both shows an exponential character and is more global in the sense that it is affected by the full shape boundary conditions.

3.3. Relation to Spectral Methods. The popularity of the heat-kernel-based methods in non-rigid shape matching is due to the usefulness of the heat kernel function in finding near-isometric correspondences between shapes. This is appealing because many expected deformations between shape surfaces, particularly the articulated motion, can be approximated by an isometric mapping. Because the isometry of a manifold preserves the heat kernel [46], heat kernel signature was shown to be isometrically-invariant in [26]. However, a volumetric isometric invariance was not sought for in the volumetric HKS of [32], and it was argued that the articulations and non-rigid deformations of solid objects do not follow a boundary isometry. Similarly, although we do not show an isometry property for our volumetric Screened Poisson Encoding Maps (SPEM), we discuss our approach against the heat-kernel-based approaches next. With μ as the Lebesgue measure, the heat kernel in (3.2) can be expressed as [46]:

$$(3.5) \quad p_t(\mathbf{x}, \mathbf{y}) = \sum_{k=1}^{\infty} e^{-\lambda_k t} \varphi_k(\mathbf{x}) \varphi_k(\mathbf{y})$$

where φ_k are the eigenvectors and λ are the eigenvalues of the Laplace operator: $\Delta_\mu \varphi + \lambda \varphi = 0$. Based on the heat kernel, Ovsjanikov [50] defined the heat kernel map $\Theta_q(\mathbf{x}) = p_t(\mathbf{q}, \mathbf{x})$, which measures the amount of heat transferred from a source point \mathbf{q} to other points \mathbf{x} over a given shape surface. The idea is to match the point from the target surface whose heat kernel map is closest to that of the given point in the reference surface. Hence, a correspondence between the two shapes can be established. On the other hand, by varying the t parameter, the heat kernel signature (HKS) [26] creates a 1-parameter family of functions from the diagonal of the heat kernel, also called the auto-diffusivity function: $p_{t_1}(\mathbf{x}, \mathbf{x}), \dots, p_{t_n}(\mathbf{x}, \mathbf{x})$.

The constructed 1-parameter family of functions based on time t in the heat kernel approaches is similar in spirit to our method. However, rather than the time variable t , we vary ρ variable in the screened Poisson operator. In the former, the temporal evolution of the heat operator is considered, hence the multi-scale heat diffusion characteristic in time is taken into account, whereas in our approach, the 1-parameter family of solution fields to the screened Poisson PDE with varying screening parameters provides the biased diffusion of the boundary sources, from the boundary towards the shape interior. Similar to the Heat Kernel Map [50], it would be possible to match shapes by sampling a set of source points \mathbf{q}_j inside the shape and using directly the 1-parameter family of the screened Poisson hyper-fields $\{v_{i=1 \dots N}^{\rho_i, \mathbf{q}_j}\}$ at points \mathbf{x} on the shape surface. Our work differs by the following: (i) in contrary to the heat kernel map approach, we put the sources on the boundary and diffuse those towards the inside of the shape with a different differential operator, i.e. the screened Poisson; (ii) instead of directly using the 1-parameter screened Poisson fields, we create a low-dimensional embedding of these functions over the ρ -dimension (§4.2). The embedding unveils the diffusion bias in projection maps which provide beneficial properties like scale adaptation, compactness and representation power, which are experimentally verified

(§5.4,5.5).

4. Extracting information from Hyper-fields.

4.1. Unveiling parts from the hyper field via sparse coding. We first focus on a single slice of the hyper-field (a fixed measure). This is a collection of m fields that forms a vector field and contains individuated boundary-internal node interactions of the shape. One may construct different useful measures from these interaction vectors. For instance, analysis of correlation between two internal nodes either inside the shape domain or between two boundary nodes, or between a boundary and an internal node are all possible using this collection. Even basic clustering methods such as k-means or Gaussian mixture models will lead to intuitive clusters of internal nodes. However, we chose to employ a specific matrix factorization technique (non negative sparse coding) to portray the decomposability of the global-local information to unveil the parts of shape.

In order to decompose the collection onto a set of components, we start with a normalized *log* field which has zero mean at each point:

$$\begin{aligned} \mathbf{V}^{\mathbf{p}_i}(\mathbf{x}) &= \log(v^{\mathbf{p}_i}(\mathbf{x})) - \frac{1}{m} \sum_{j=1}^m \log(v^{\mathbf{p}_j}(\mathbf{x})), \\ (4.1) \quad &= \log(v^{\mathbf{p}_i}(\mathbf{x})) - \log \sqrt[m]{v^{\mathbf{p}_1} v^{\mathbf{p}_2} \dots v^{\mathbf{p}_m}}(\mathbf{x}) \end{aligned}$$

Note that centering the *log*-field by its mean is equivalent to centering the field by its geometric mean:

$$(4.2) \quad \mathbf{V}(\mathbf{x}) = \log \left\{ \frac{v}{\sqrt[m]{v^{\mathbf{p}_1} v^{\mathbf{p}_2} \dots v^{\mathbf{p}_m}}} \right\}(\mathbf{x})$$

In order to apply the non-negative matrix factorization, the vector elements that are lower than the mean are replaced by zeros. Though this normalization procedure ignores a region within a certain proximity to the boundary node of interest, thanks to the centering of the data, remaining regions are encoded in a manner that allows distinction of prominent parts. The resulting non-negative vector field can now be analyzed as an additive combination of some bases, leading to a part-based representation. The non-negative measurements obtained by a normalization with the mean and median are depicted in Fig.7.




Fig. 7: Non-negative measurements: $y_j(i)$, where the same \mathbf{p}_i in Fig. 5 is used. *Left*: Normalization by *median*. *Right*: Normalization by *mean*.

Arranging the measurements $\mathbf{V}(\mathbf{x})$ into columns of a matrix $Y^{m \times |\Omega|}$ for each shape node \mathbf{x} , would allow the linear decomposition of the data as $Y \approx AS$, where

the matrix A is the mixing matrix with *basis vectors* as its columns. The rows of S contains the hidden components that encodes the contribution of each mixing vector while reconstructing the input vectors. When both factors A and S are forced to be non-negative, the decomposition corresponds to the method of non-negative matrix factorization (NMF) [51][52]. The non-negativity of the factors makes the representation additive as desired.

Many variants of NMF have been proposed since the pioneering work of Lee and Seung [52]. Due to its additive nature, NMF produces a sparse representation of the data, where the data is represented using inherent *active components*. Non-negative sparse coding (NNSC), introduced by Hoyer[53], forms an analogy between NMF and sparse coding [54]. NNSC provides control over the sparseness of the hidden components by adding a sparsity-inducing penalty to the objective function, which is very desirable feature for obtaining shape parts as active components that are described on the shape domain. By selecting k mixture elements, the objective function of the NNSC is formulated as:

$$(4.3) \quad \min_{A \in \mathbb{R}^{m \times k}, S \in \mathbb{R}^{k \times |\Omega|}} \sum_{i=1}^{|\Omega|} \left(\frac{1}{2} \|y_i - AS_i\|_2^2 + \lambda \|S_i\|_1 \right) \quad s.t. \quad A \geq 0, \forall i, S_i \geq 0,$$

where the first term forces minimization of reconstruction error and the second term forces the sparseness. λ controls the tradeoff between sparseness and accurate reconstruction of Y . Sparsity is enforced by using the L_1 norm, this formulation can also be considered as the constraint in the *Lasso* problem [55]. $\lambda = 0$ case is equivalent to NMF formulation (i.e. no additional sparsity). The problem is solved using the method of Mairal et al. [56][57], which outperforms method of Hoyer [53] in minimizing the objective function in batch mode.

The resulting shape decomposition as NNSC components for a hand shape is presented in Fig.8. The results are produced using the slice corresponding to $\rho \rightarrow \infty$. In the first experiment (Fig.8.a), the shape is decomposed into six components, with a larger λ value. The fingers and the central part of the hand are separated as expected. In the second experiment (Fig.8.b), twelve components are obtained with a relaxed sparseness constraint. Notice that the components that represent fingers, which are the most prominent parts, are preserved. Additional components represent the connection points of articulated parts to the central part of hand. Also, the central part of the hand is partitioned into three different parts. The fact that important parts are preserved even when the separation settings are relaxed illustrates the nature of the information preserved in the measurements and robustness of the representation.

4.2. Producing consistent mappings for shape correspondence: SPEM.

In the previous section, we have concentrated on a single slice in the hyper-field and demonstrated that sparse coding unveils parts by integrating local-global interactions. In this section, we focus on a complementary problem in shape analysis: defining real-valued functions on a shape domain that can be used for the purpose of matching or registration. In the other dimension of the hyper-field, the 1-parameter space that is spanned by varying ρ -values encodes the boundary-interior diffusion characteristics. Although it is possible to utilize this 1-parameter field directly for shape matching, we take one step further and we compactly code the variation in the ρ dimension to produce consistent mappings through a low-dimensional embedding. There is a vast number of dimensionality reduction approaches. We advocate use of principal component analysis (PCA) which produces consistent maps that exhibit adaptation to scale (see Fig. 9 and Fig. 10).

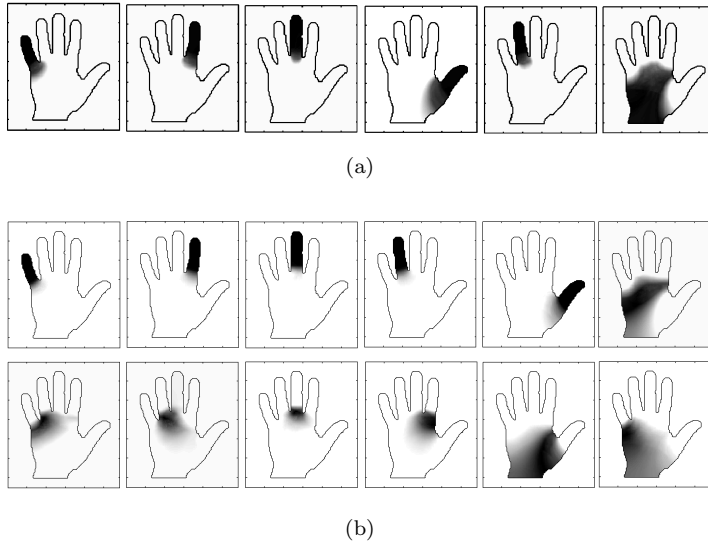


Fig. 8: a) NNSC components obtained using a large λ and $k = 5$; b) NNSC components obtained using a low λ and $k = 12$.

We find the linear PCA very intuitive as compared to some other popular decomposition methods. Orthogonality of the bases provides quite a consistent mapping across shapes. Independent component analysis based methods form inconsistent mappings. We have observed that non-linear methods such as locally linear embedding [58] or Diffusion Maps [29] over-learn the ρ -space, leading to less number of features and less consistency. Linear PCA also outperforms latent variable methods such as Probabilistic PCA solved by maximum likelihood estimation [59]. The data is created by a *linear* operator, and it is extremely *smooth*. We have observed consistency (among different poses of similar shapes) using a direct singular value decomposition even for the projections that explain variance as small as (10^{-14}) .

We now consider all ρ^2 -slices of the hyper-field, i.e. consider the $|\Omega| \times N$ matrix Y . Each column of Y is a v field for a certain choice of ρ^2 . The covariance matrix of Y is computed, and then decomposed to yield an orthogonal set of bases: the eigen maps $\Phi_n, n = 1, \dots, N$ of the hyper field. After the new bases are calculated, the hyper field can be projected to form N mappings, where each mapping \mathcal{P}_n is related to a measure of the variance explained by the n^{th} basis:

$$(4.4) \quad \mathcal{P}_n = Y\Phi_n.$$

The low-dimensional embedding facilitates a principled selection of a handful of projections maps, as we call them the SPEM (Screened Poisson Encoding Maps). We note that we observe some interesting properties such as almost perfect representation of the variability in just several bases (or projections). We relate this to the linearity and smoothness of the screened Poisson operator. Using these bases, we observe visual correspondence even in 2D shapes under a perspective transform.

4.2.1. Adaptation to Scale. The resulting eigenvectors Φ_n for a hand shape can be observed in Fig. 9. The eigenvectors adapt to global changes of the shape, leading to a robust representation. This is exemplified by scaling the hand shape.

Notice that the eigenvectors change because a specific characteristic that is detected for a larger distance corresponds to a larger ρ value. This adaptation does not mean that the field is scale invariant, because discretization in spatial and ρ domains would not allow direct invariance. However, the representation does not change abruptly as scale increases. In order to show this, we computed peak-SNR (PSNR) values (in dB) between the original hand shape (maximum value of distance transform is 20 pixels) and its scaled versions up to scale 4 (Fig. 9 scale changes are coded by colour on the bottom right). Note the slow monotonical change across scale for the projections, which provides coherence against shape scale changes.

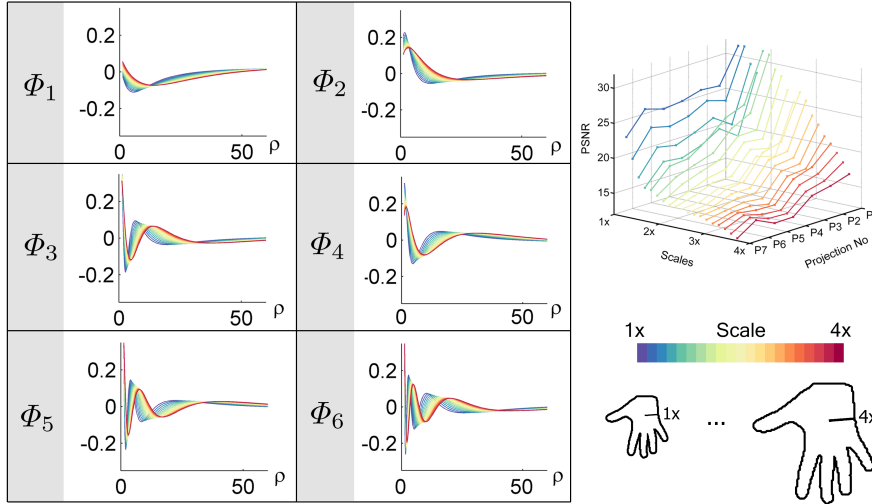


Fig. 9: For the hand shape: Left: calculated Φ_n colored according to corresponding scale ratios; Right: PSNR values for projections obtained using Φ_n across different shape scales show a slow monotonically changing behaviour, which provides a desired robustness to scale changes, color coded as shown on the bottom right.

Adaptation of the principal directions in ρ -space to *scale* is also presented in Fig. 10. Class of $\{n/4\}$ regular star polygons for $n = 9, \dots, 20$ are depicted, where all the vertices are lying on circles of a constant radius. As n increases, the shapes become more circular. This change of internal distance relationships affects the characteristics of the hyper-field in its ρ -dimension. The eigenvectors of the covariance matrices are altered in accordance, leading to robustness to scale changes. The six eigenvectors on the right are almost identical and they are calculated for the shapes that are scaled to have the same maximum distance to boundary. This property of the projections imply that the discriminability of the projections originate from local and global spatial relationships. The model offers a framework where globally similar shapes should have similar projections in locally similar regions, which makes it a promising tool for shape analysis along with robustness to global scale effects.

5. Results and Discussions. In this section, we demonstrate the expressive power and robustness of projections of the new hyper-field. After discussing computational issues, we first present qualitative results with sparse coding over boundary decomposition of the shape hyper field. Next, we show first a study on 2D shape classification that demonstrates the usefulness of the SPEM with a moment-feature based evaluation. Finally, we validate the new SPEM descriptor over the 3D *SHREC*

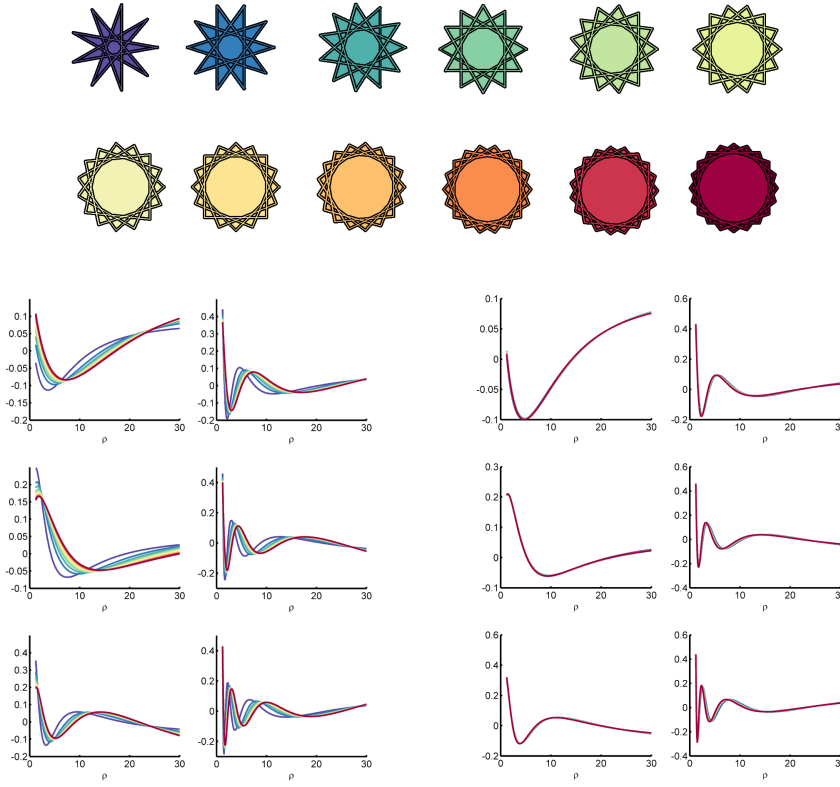


Fig. 10: *Top*: $\{n/4\}$ regular star polygons, for $n = 9, \dots, 20$ *Left*: First six eigenvectors Φ_1, \dots, Φ_6 for the shapes colored accordingly. *Right*: The first six eigenvectors Φ_1, \dots, Φ_6 for the shapes after re-scaling with respect to the maximum value of the shortest distance to the boundary.

benchmark data set [42].

5.1. Computational Aspects. Computation of each field v_ρ or $v_\rho^{\mathcal{P}}$ can be done in parallel in both approaches that we presented. Notice that we calculate projections \mathcal{P}_i using all the boundary nodes used as sources. Also, we fix ρ for the calculation of the fields for the sparse coding (NNSC) application. Calculating a field for each boundary node taken as a source for a 3D shape is not feasible, yet it is possible to apply a similar approach to calculation of fields over segmented regions on the boundary. This requires a fast initial partitioning of the boundary nodes with large granularity.

We implement the hyper-field as a sparse matrix vector multiplication on a NVIDIA Tesla K20c GPU. The screened Poisson operator is represented as a sparse matrix: $P_{|\Omega| \times |\Omega|} = (\Delta - \frac{1}{\rho^2} I) \chi_{(B.C.)}$, where Δ is the five-point discretization of the Laplace operator and $\chi_{(B.C.)} : |\Omega| \times |\Omega| \rightarrow \{0, 1\}$ is an indicator function for the edges that are allowed by the desired boundary condition. The gradient descent solution to the screened Poisson field is then obtained by iterating the following multiplication: $v^{n+1} = v^n (1 + \tau P)$, where τ is the artificial time step. In our implementation we use MATLAB Parallel Processing Toolbox and CUSP library [60], which is a generic

CUDA library for sparse linear algebra. The computation time of each screened Poisson field v_ρ or $v_\rho^{\mathbf{D}}$ for a shape of 250.000 voxels is 2 seconds. The total computation time of a hyper-field is directly proportional to the number of boundary segments for the approach in § 4.1 and number of samples from the ρ domain for SPEMs (§ 4.2). For 2D shapes, we concatenate sparse matrices of the operators $\Delta - \frac{1}{\rho^2}I$ and solve the fields simultaneously, which is not an option for 3D shapes due to GPU memory constraints. The calculation of the projections for each shape in 1000 shape database [15] takes approximately 3 seconds.

Each field is a solution to an elliptic linear PDE, which is a problem that occurs in various fields, and many fast alternative solvers exist. Adaptation of GPU’s is an ongoing study for more than a decade [61]. While sparse Cholesky decomposition and FFT based approaches work in subquadratic time [23, 20], multifrontal methods [62, 63] and multigrid methods [11, 64] can reach $O(N)$, which is the lower bound for the problem. Certainly, a more customized and efficient GPU implementation would lead to a faster computation, yet our implementation is satisfactory for the 2D and 3D experiments we present.

5.2. Boundary Decomposition Based on Regional Information. As described previously in §4.1, a natural application of the non-negative sparse decomposition of the shape hyper-field was partitioning of the shape domain into its “meaningful” components. The decomposition is applied to nodes on the shape domain based on their random walk distances to all the boundary nodes as in the demonstration on the hand shape in Fig. 8 in §. 4.1.

Here, we demonstrate another setting where the shape decomposition is achieved by minimizing the objective function in (4.3) in §.4.1 using the transpose of the measurement matrix Y without normalization (4.2). In this setting, in contrast to the previous setting §. 4.1, the boundary nodes are decomposed based on their random walk distances to all the nodes in the shape domain. Minimization of the reconstruction error in the objective function depends on boundary nodes and the regions that are associated with each node. That is, the boundary nodes that relate to similar regions are more likely to belong to the same boundary partition. Hence, the resulting decomposition of the boundary inherently depends on a regional partitioning of the shape. An example on a human figure is presented in Fig. 11, where decomposed parts and corresponding regions can be observed as active components and basis vectors (respectively) that are factorized from the hyper-field using NNSC.




Fig. 11: Decomposition of the human figure and associated regions. $k = 8$.

Introducing information about regional characteristics of a shape for decomposition of its boundaries leads to rather consistent results. We demonstrate this in Fig. 12 for three distinctively different cat poses. The structures after decomposition are very coherent. The sparse decomposition into eight boundary segments reveals the head, the front and rear parts of central body, the tail and four legs. In the third pose

only, an additional segment is included in the leg whose regional characteristics are altered due to the significant articulation and deformation, however the inconsistent segment can easily be detected and eliminated considering its low intensity.

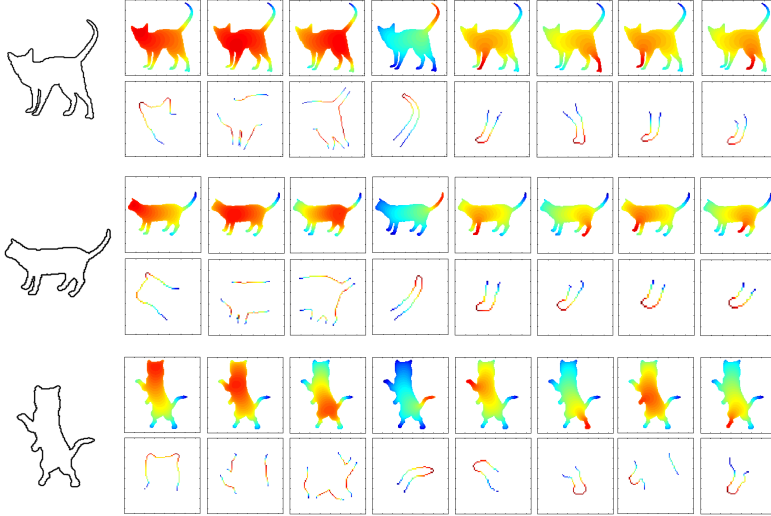


Fig. 12: Non-negative sparse decomposition over shape hyper-fields of three highly different cat poses partition shape boundary into: the head, the cat frontal body, the back body, its tail, and its legs in a consistent manner.

5.3. Orthogonal Projections Based On ρ^2 Sweep: SPEM. For the SPEM, we experimented with a set of shapes that are not necessarily related by isometry. In Fig. 13, the projections obtained using the first five principal components are presented for six different cat shapes. The first two projections, which explain most of the variance in the data, are much smoother compared to the remaining projections. In the first projection, it is observed that the nodes in the vicinity of the boundary attain highly positive values, hence, can be distinguished from the interior nodes, giving only a rough sense of central/peripheral separation. The second projection, on the other hand, exhibits a much stronger central/peripheral separation similar to [22]. The projections(SPEM) obtained using the third or higher eigenvectors encode more subtle details. For example, the third projection reveals ears, head, legs and tail of the respective cat. Notice that these explicitly expressed parts are intuitive and consistent across deformations of the cat shape.

Figure 14 demonstrates two things: (i) human figure, with different articulated motion as well as small local deformations shows the projections preserve their character across those nonrigid deformations; (ii) hand figure with occlusion, local deformation and noise effects show robustness of the projections against noise. The consistency, which can be observed among the projections over each row across the varying instances of the human and the hand shapes, is poised to provide the desired robustness in shape representation required for shape matching and recognition.

In Fig. 15, a 3D example of the SPEM is presented. The projections of the 4D hyper-field computed from a 3D horse form onto second to sixth projections are depicted. Since a 3D form conveys the exact geometry of a real world object as opposed to a 2D shape, which is a perspective projection, our projected fields are naturally




Fig. 13: Left-Right: First five projections(SPeM): $\mathcal{P}_{1,\dots,5}$ for 6 different poses of a cat shape, depicted in each row. Each column corresponds to a different projection mode. Hotter colors indicate positive and high values while colder colors indicate negative and low values. Consistency of projections across deformations of the cat shape is observed.

more consistent across arbitrary pose changes. In order to be able to visualize different sections distinguished by each projection, we applied a histogram based thresholding procedure. For each projection, one positive and one negative threshold is selected and the surfaces corresponding to the level-sets of those thresholds are visualized. Thresholds were selected at the first jump in the histogram for all the projections and the same threshold was used for the same projections of shapes under different poses. The same remark that was made about the smoothness of the projected fields in the 2D case holds for the 3D case as well. Although some of the thresholded parts can be detached, as in the blue neck part in the sixth projection, the consistency and the similarity of the 3D fields even after the thresholding are notable.

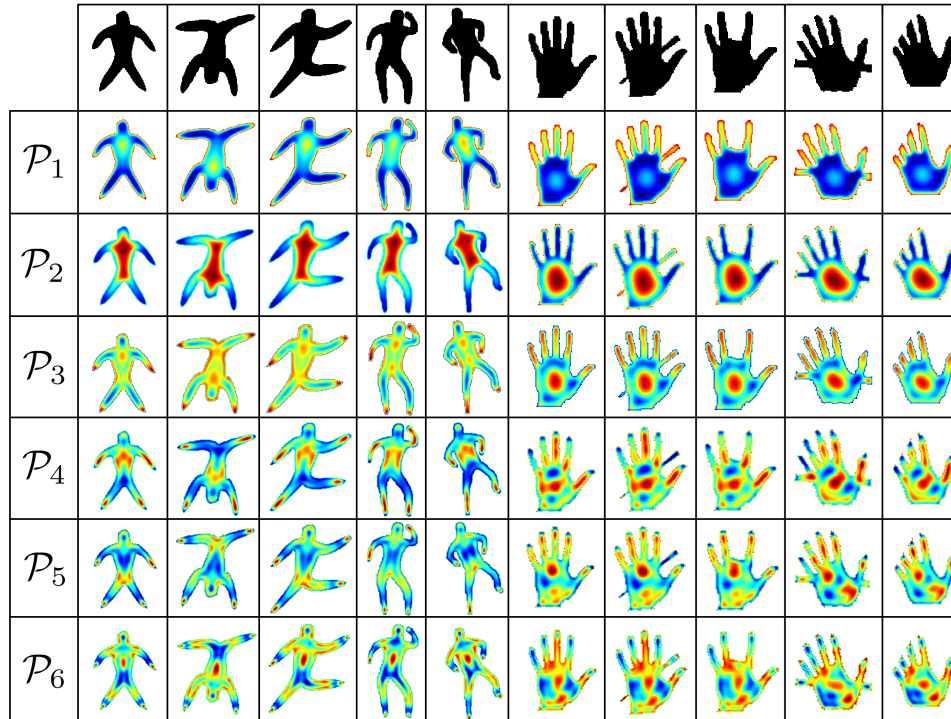


Fig. 14: First six projections(SPEM): $\mathcal{P}_{1,\dots,6}$ (on each row) for five different instances (on each column) of a human and a hand silhouette. Human figure displays articulated motion and local deformations. Hand figure displays different noise conditions: occluding a finger; shortening of fingers; protruding two new parts from the hand. Hotter colors indicate positive and high values while colder colors indicate negative and low values. Robustness of projections against occlusion, local deformation, and noise is observed.




Fig. 15: Top-Down: Second to sixth projections(SPEM): $\mathcal{P}_2, \dots, \mathcal{P}_6$ for three different poses of a 3D horse. Consistency of each projection across a row for different poses can be observed.

5.4. A Moment Based Evaluation of Consistency and Correspondence.

In order to quantitatively demonstrate the consistency of the projections, we conducted a classification experiment on the 2D "1000-shape" database, which is an extended version of [15]. The database consists of 50 classes, each containing 20 shapes varying significantly with severe deformations and articulations. From each class, 10 shapes are randomly selected as training data and remaining 10 is used for testing. In order to experiment with the classification performance using a moment-based representation, we extract a group of shapes with disjoint parts from each input single connected binary shape in the database.

The shapes in the database are scaled to a fixed maximum distance to boundary of 20 voxels. Next, the hyper-fields are created, and principal directions are calculated in the ρ dimension. To each projection obtained from a given shape, a basic k-means clustering is applied ($k = 3$) using the intensities of the projections. Thanks to the nature of the SPEM, the resulting cluster centers are very similar: One of the cluster means is very close to zero in terms of projection intensity value and the other two are from the shape nodes that have positive and negative intensity projections, characterizing positive and negative nodal domains. We use the mean of the corresponding cluster for both the negative and positive clusters to generate two new shape maps for each projection. This can be considered as a rough yet straightforward approach for detecting regions that behave similar in ρ^2 space, specifically in a certain principal direction of the hyper-field. We note that a common positive and a negative threshold value is utilized for all shapes in the database. In Fig.16, we exemplify the positive and negative shape clusters obtained by thresholding the first five projections on several cat shapes from the database.

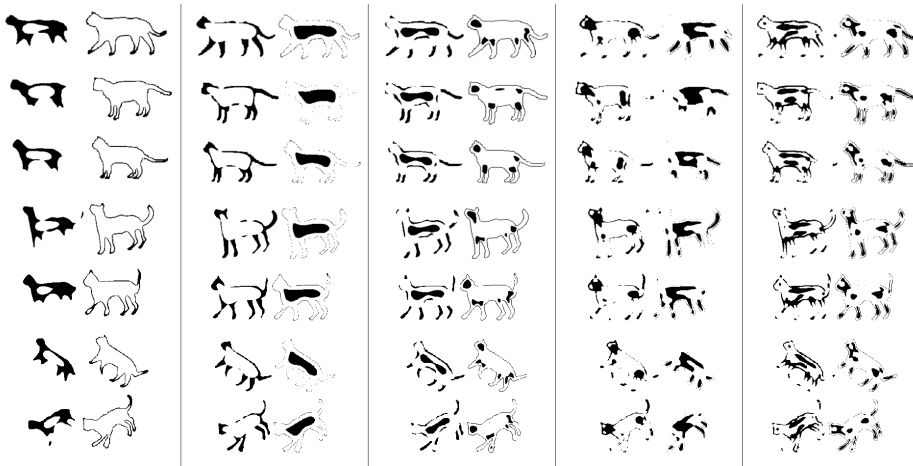


Fig. 16: Each row contains the negative-positive nodal domain clusters for corresponding to first 5 projections of 7 cat shapes.

As features, we computed Hu's seven invariant moments [65], which are invariant to translation, scale and orientation as features for the classification. The *weak* sense of similarity that these simple moments provide allows us to evaluate the correspondences more clearly. We train linear Support Vector Machines (SVM's) using moments of each generated shape both by stacking features in a cumulative manner and individually to each projection. The classification results for both experiments that are repeated 100 times, randomized over selection of 10 training and 10 test

shapes, can be observed at Fig.17.

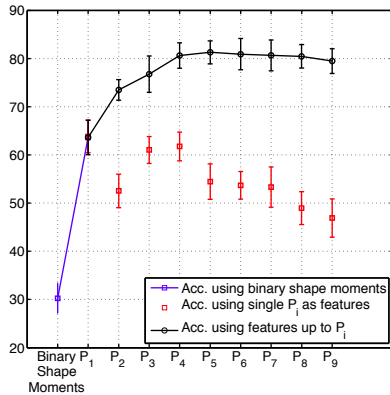


Fig. 17: SVM Classification accuracies using moment features of: binary shape moments (blue); individual thresholded projections (red); and cumulatively adding thresholded projections (black). Notice that the success rate jumps from %30 (blue) to %80 (blacks) when our approach is used.

From the experiments where the SVMs are trained using features from individual shape projections (red), it is clear that the moments for the newly generated shape maps are much more informative compared to only the input shape’s moments (blue). This alone shows that level curves of the SPEMs are consistent among shapes of same category and corresponding regions on shapes of the same category have similar projections. The moments obtained from even the ninth projections, which explain very little portion of the variance in the hyper-field, are almost twice as descriptive compared to the original shapes moments (see Fig.17). The monotonic behavior in classification performance obtained using the combinations of the moments (black) as features implies that new projections introduce new information (that even the moments can express), which is an observation that is greatly in accordance with orthogonality of the projections.

5.5. Non-Rigid Shape Retrieval Using Screened Poisson Encoding Maps.

To address the problem of retrieving similar shapes from a database given a query shape, we utilize the information extracted from the Screened Poisson Encoding Maps (SPEMs). We use a feature-based approach to obtain a compact global shape descriptor from SPEM using feature encoding methods.

An analogy between feature-based 3D shape retrieval and image retrieval was made in [66], where an image is treated as a collection of primitive elements, namely, local image descriptors as *visual words*. The analogy is formed by obtaining *geometric words* using multiscale diffusion heat kernels, which are represented by a geometric vocabulary using soft vector quantization. A similar feature-based approach is used in [40], where Interior Heat Kernel Signatures (iHKS) are used as geometric words with a similar representation proposed in [66]. Our retrieval approach is mainly similar to those in perspective, yet it differs in the way geometric words are obtained and the way the features are encoded.

As features, we use the SPEMs explained in § 4.2. Considering the nature of the problem, due to large variability of the shapes undergoing non-rigid deformations,

the features should be robust to bending and articulations, which cause topological changes in the volumetric representation. In Fig. 18, we present joint histograms of projections for several shapes that go through large pose changes. The shapes belong to the *SHREC'11* benchmark[42]. The histograms are obtained from the values of SPEMs: $\mathcal{P}_4(\mathbf{x})$ (horizontal axis) and $\mathcal{P}_3(\mathbf{x})$ (vertical axis) for all $\mathbf{x} \in \Omega$, the logarithm of the number of nodes in the bins are visualized. The choice of fourth and third projections is purely arbitrary, other projections also give coherent results.




Fig. 18: Joint histograms inside SPEMs: \mathcal{P}_4 vs \mathcal{P}_3 for corresponding shapes on the right. The histogram intensities are displayed using a logarithmic scale. The articulations have almost no effect on the joint histograms and there is large variation in histograms of shapes with different volumetric structures.

The histograms visualized in Fig.18 provide only a hint of what the feature space looks like, yet the distinctiveness of the volumetric information encoded is clearly revealed. Even for the shapes that go through large articulated motion and deformation, the representation remains unaltered. Also notice that the representation of woman and man shapes are more similar (yet still distinguishable) in comparison to the representation of other shapes that are less related.

In order to compactly represent a shape for the retrieval application, we use the feature encoding method: Vector of Locally Aggregated Descriptors (VLAD)[67]. VLAD characterizes the distribution of vectors with respect to the pre-computed centers, words that belong to a vocabulary. Unlike the *Bag of Features* or *soft vector quantization* approaches, where the distances of the features to centers are accumulated, the difference *vectors* from each feature to assigned center are aggregated.

For a node in the shape domain $\mathbf{x} \in \Omega$, the SPEM $\mathcal{S}(\mathbf{x})$ consists of the projections $\mathcal{P}(\mathbf{x})$: $\mathcal{S}(\mathbf{x}) = \{\mathcal{P}_1(\mathbf{x}), \dots, \mathcal{P}_d(\mathbf{x})\}$ where d is the number of projections used. A

codebook, $\mathcal{C} = \{\mathbf{c}_1, \dots, \mathbf{c}_k\}$, of k representative points $\mathbf{c}_i \in \mathbb{R}^d$ is acquired using k-means clustering algorithm. Each node in the shape domain is assigned to the nearest cluster center, as a hard vector quantization to obtain $\mathcal{NN} \in \mathbb{R}^d$:

$$(5.1) \quad \mathcal{NN}(\mathbf{x}) = \arg \min_{\mathbf{c}_i \in \mathcal{C}} \|\mathcal{S}(\mathbf{x}) - \mathbf{c}_i\|$$

The differences $\mathcal{S}(\mathbf{x}) - \mathbf{c}_i$ of the vectors assigned to each center \mathbf{c}_i are accumulated to obtain d dimensional *residual* sum vectors \mathbf{R}_i :

$$(5.2) \quad \mathbf{R}_i = \sum_{\mathbf{x} \in \Omega: \mathcal{NN}(\mathbf{x}) = \mathbf{c}_i} \mathcal{S}(\mathbf{x}) - \mathbf{c}_i$$

The aggregated residual vectors \mathbf{R}_i are then normalized so that they have unit L_2 norm. This is followed by a power normalization, which can be considered as a variance stabilizing transform:

$$(5.3) \quad \theta_i = \text{sign}(\mathbf{R}_i) \left(\frac{\mathbf{R}_i}{\|\mathbf{R}_i\|} \right)^\alpha$$

where the power operates element wise on the given vector. This normalization helps coping with the undesired burstiness of the encoded vectors. One of the sources of burstiness is the different number of nodes assigned to each center. In [68], several approximations are introduced that associate (5.3) to earlier work on variance stabilizing transforms on a compound Poisson distribution [69, 70]. Another source of variance in our case is the differences in the intensities of the projections. The SPEM projections that correspond to eigenvalues that express little variance have much lower intensities, hence little contribution in comparison to, say, \mathcal{P}_1 . The power transform enhances their contribution in the overall representation. The vectors θ are concatenated and again normalized to have unit L_2 norm as a vector of length $k \times d$ to form the final representation of a shape.

We conducted the retrieval experiment using the *SHREC'11 Shape Retrieval on Non-Rigid 3D Watertight Meshes* benchmark[42]. The database consists of 30 classes, each with 20 samples that inherit large intra-class variation in terms of articulated motion and non-rigid deformation. In our implementation, all of the shapes in the database were normalised to a constant scale of $\sim 250,000$ voxels. We used first six projections $d = 6$, and trained a vocabulary C of size $k = 32$ words, resulting into a 192 dimensional vector as a global descriptor of shape. The vocabulary is learned from 50000 randomly sampled nodes from each shape in the database. Standard deviations for the SPEM projections that belong to each projection index are also estimated using this training data. While evaluating, each SPEM projection is first L_2 normalized then multiplied by the standard deviation estimated for that projection index. The nodes that are next to the boundary are discarded while encoding, since they are less informative and might get affected from boundary discretization. The idea is to normalize the projection intensities, without changing the intensity range of the projections. We observed that such a normalization causes an increase in retrieval performance. For power normalization of VLAD, we use $\alpha = 0.25$ in Eq.(5.3). Correlation distance between the 192 dimensional global descriptor is used as a measure of dissimilarity between query shape and all shapes in the database.

The retrieval performance is demonstrated in Table 1 using five standard retrieval statistics: nearest neighbor (NN), first tier (FT), second tier (ST), E-measure (E), and discounted cumulative gain (DCG). For the details regarding the measures, we refer to [71]. Results for best four methods from the participants in the contest are presented. The results for these are taken from [72], where the methods are also briefly presented. Additional results in the table are taken from the work of Rustamov: iHKS [40] and Konukoglu’s spectral distance method: WESD [35]. Hybrid method results are obtained the way *SD-GDM + MeshSift* hybrid method is presented in [42], specifically, by applying a min-max algorithm to distances obtained using two methods and adding them up to obtain the final distance matrix. We merge our results with other methods in the proposed manner, using the distance matrices participated in SHREC contest, to demonstrate how combination of our work and existing work performs. We also present the precision-recall performance of our retrieval approach and the top four contestants in SHREC’11 track in Fig.19.

Table 1: SHREC’11 Retrieval on Non-Rigid 3D Watertight Meshes Database Results

		NN	FT	ST	E	DCG
Methods:	Our Method (SPEM)	99.8	97.4	98.6	73.3	99.3
	SD-GDM [73]	100.0	96.2	98.4	73.1	99.4
	iHKS[40]	99.5	92.2	95.5	71.0	98.0
	MDS-CM-BOF[74]	99.5	91.3	96.9	71.7	98.2
	ShapeDNA[34]	99.2	91.5	95.7	70.5	97.8
	WESD[35]	99.3	90.2	93.05	69.0	97.1
	MeshSIFT [75]	99.5	88.4	96.2	70.8	98.0
Hybrid	SPEM+SD-GDM	100.0	98.4	99.5	74.0	99.7
Methods:	SPEM+MDS-CM-BOF	100.0	97.5	99.1	73.7	99.6
	SPEM+MeshSift	99.7	97.6	98.9	73.5	99.5
	SD-GDM+MeshSift	100.0	97.2	99.0	73.6	99.6

Our method outperforms existing methods in terms of first-tier, second-tier and E-measure performance. It ranks the second for the NN and DCG measures with very close numbers. Only one of the nearest neighbor results out of 600 is retrieved falsely, which is a sample from the *ants* class. Generally, the method only fails when internal distance relations undergo large changes, which occasionally happen in ants class deformations. Among the methods compared, only iHKS uses volumetric information. Our method clearly outperforms iHKS with respect to all of the performance measures.

The results suggest that extracting volumetric information in a robust way can lead to enhancement in the non-rigid shape retrieval performance when compared to extracting information regarding only intrinsic surface properties. Moreover, as expected, combining volumetric information and surface information results in a significant boost in performance, as observed in resulting hybrid performance of SPEM in Table 1.

6. Conclusion. In this paper, we provided a novel distance hyper-field representation for shapes via screened Poisson PDE. An acknowledged advantage of the presented shape representation, which is valid for all distance-based shape representations, is that the shape model is free from the dimensions of the shape: the volumetric

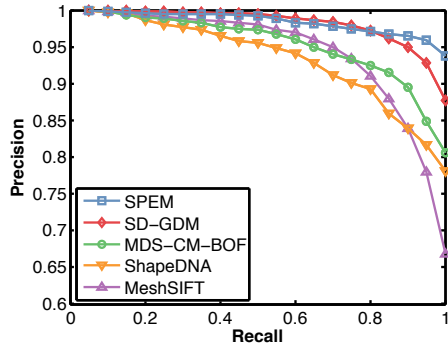


Fig. 19: Precision - Recall Performances in Shrec'11 Non-Rigid Database

description of the hyper-fields we propose is readily extendable to any shape embedding dimensions in \mathbb{R}^n for $n = 2, 3, 4, \dots$

For an $n - d$ shape with m boundary points, the hyper-field is a scale-space stack parameterized by the screening parameter, where each member of the stack v^ρ is a superposition of m fields, each of which is the solution of the PDE with a point source placed at a boundary point. The shape hyper-field captures all sorts of characteristic information within and on the boundary of the shape, therefore encodes both local-global and interior-boundary interactions.

The new representation is also tied to a collection of conditioned random walks, each walk emanating from a fixed boundary point and *walking* with a bias controlled by the screening parameter; and, the effect of the screening parameter is a change of measure.

Furthermore, we demonstrated extraction of shape information from the hyper-field, which is a rich high-dimensional representation, by two felicitous compact decompositions to exemplify: (i) natural shape partitions (NNSC); (ii) consistent shape maps through SPEM. The potential of extracting various shape descriptors from the introduced shape hyper-field was demonstrated over both a 2D "1000-shape" database [15], and a benchmark dataset *SHREC'11* [42]. The SPEM performance was evaluated by using the VLAD method [67] for volumetric feature encoding. The SPEM consistently ranked the first or the second in all measures, and ranked the first when a hybrid combination with top surface-based methods was computed. Another interesting property of the SPEM was its adaptation to scale, which was experimentally verified by its performance over the benchmark.

The shape hyper-field representation presented in this paper is constructed over the whole shape domain, therefore, it is certainly possible to either define new shape measures or adopt existing popular descriptors, even image descriptors, in order to apply the introduced hyper-field framework to shape matching, classification or partitioning problems.

To sum up, as we introduced the foundations of the new shape hyper-field, the SPEM volumetric features, and demonstrated its strength in a large-scale 3D nonrigid shape retrieval application, we believe that it will find ubiquitous use in various shape analysis applications.

REFERENCES

- [1] Sclaroff, S., Pentland, A.: Generalized implicit functions for computer graphics. *SIGGRAPH Comput. Graph.* **25**(4) (1991) 247–250
- [2] Savchenko, V.V., Pasko, A.A., Okunev, O.G., Kunii, T.L.: Function representation of solids reconstructed from scattered surface points and contours. In: *Computer Graphics Forum*. Volume 14., Wiley Online Library (1995) 181–188
- [3] Osher, S., Sethian, J.A.: Fronts propagating with curvature-dependent speed: algorithms based on Hamilton–Jacobi formulations. *J. Comput. Phys.* **79**(1) (1988) 12–49
- [4] Sethian, J.A.: *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*. Second edn. Volume 3 of Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, Cambridge (1999)
- [5] Leventon, M.E., Grimson, W.E.L., Faugeras, O.: Statistical shape influence in geodesic active contours. In: *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*. Volume 1., IEEE (2000) 316–323
- [6] Chan, T.F., Vese, L.A.: Active contours without edges. *Image Processing, IEEE Transactions on* **10**(2) (2001) 266–277
- [7] Tsai, A., Yezzi Jr, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W.E., Willsky, A.: A shape-based approach to the segmentation of medical imagery using level sets. *Medical Imaging, IEEE Transactions on* **22**(2) (2003) 137–154
- [8] Chen, Y., Tagare, H.D., Thiruvankadam, S., Huang, F., Wilson, D., Gopinath, K.S., Briggs, R.W., Geiser, E.A.: Using prior shapes in geometric active contours in a variational framework. *International Journal of Computer Vision* **50**(3) (2002) 315–328
- [9] Riklin-Raviv, T., Kiryati, N., Sochen, N.: Segmentation by level sets and symmetry. In: *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*. Volume 1., IEEE (2006) 1015–1022
- [10] Tari, Z., Shah, J., Pien, H.: Extraction of shape skeletons from grayscale images. *Computer Vision and Image Understanding* **66**(2) (1997) 133–146
- [11] Gorelick, L., Galun, M., Sharon, E., Basri, R., Brandt, A.: Shape representation and classification using the poisson equation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **28**(12) (2006) 1991–2005
- [12] Gorelick, L., Blank, M., Shechtman, E., Irani, M., Basri, R.: Actions as space-time shapes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **29**(12) (2007) 2247–2253
- [13] Haidar, H., Bouix, S., Levitt, J.J., McCarley, R.W., Shenton, M.E., Soul, J.S.: Characterizing the shape of anatomical structures with poisson’s equation. *Medical Imaging, IEEE Transactions on* **25**(10) (2006) 1249–1257
- [14] Aubert, G., Aujol, J.F.: Poisson skeleton revisited: A new mathematical perspective. *Journal of Mathematical Imaging and Vision* **48**(1) (2014) 149–159
- [15] Aslan, C., Erdem, A., Erdem, E., Tari, S.: Disconnected skeleton: shape at its absolute scale. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **30**(12) (2008) 2188–2203
- [16] Gurumoorthy, K.S., Rangarajan, A.: Distance transform gradient density estimation using the stationary phase approximation. *SIAM J. Math. Analysis* **44**(6) (2012) 4250–4273
- [17] Sethi, M., Rangarajan, A., Gurumoorthy, K.: The schrödinger distance transform (sdt) for point-sets and curves. In: *Computer Vision and Pattern Recognition, 2012 IEEE Conference on, IEEE (2012) 198–205*
- [18] Duci, A., Yezzi, A.J., Mitter, S.K., Soatto, S.: Shape representation via harmonic embedding. In: *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on, IEEE (2003) 656–662*
- [19] Duci, A., Yezzi, A., Soatto, S., Rocha, K.: Harmonic embeddings for linear shape analysis. *Journal of Mathematical Imaging and Vision* **25**(3) (2006) 341–352
- [20] Crane, K., Weischedel, C., Wardetzky, M.: Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics* **32**(5) (2013) 152
- [21] Belyaev, A.G., Fayolle, P.A., Pasko, A.A.: Signed Lp-distance fields. *Computer-Aided Design* **45**(2) (2013) 523–528
- [22] Tari, S., Genctav, M.: From a non-local ambrosio-tortorelli phase field to a randomized part hierarchy tree. *Journal of Mathematical Imaging and Vision* (2013) 1–18
- [23] Bhat, P., Curless, B., Cohen, M., Zitnick, C.: Fourier analysis of the 2d screened poisson equation for gradient domain problems. In: *Proc. 10th European Conf Computer Vision: ECCV. (2008) 114–128*
- [24] Chuang, M., Kazhdan, M.: Interactive and anisotropic geometry processing using the screened poisson equation. *ACM Trans Graphics (TOG)* **30**(4) (2011) 57–57

- [25] Kazhdan, M., Hoppe, H.: Screened poisson surface reconstruction. *ACM Trans Graphics (TOG)* **32**(3) (2013) 29–29
- [26] Sun, J., Ovsjanikov, M., Guibas, L.: A concise and provably informative multi-scale signature based on heat diffusion. In: *Computer Graphics Forum*. Volume 28., Wiley Online Library (2009) 1383–1392
- [27] Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: *Proceedings of the fifth Eurographics symposium on Geometry processing*, Eurographics Association (2007) 225–233
- [28] Lafon, S.S.: Diffusion maps and geometric harmonics. PhD thesis, Yale University (2004)
- [29] Coifman, R.R., Lafon, S.: Diffusion maps. *Applied and computational harmonic analysis* **21**(1) (2006) 5–30
- [30] Liu, Y.S., Li, Q., Zheng, G.Q., Ramani, K., Benjamin, W.: Using diffusion distances for flexible molecular shape comparison. *BMC Bioinformatics* **11** (2010) 480
- [31] Bronstein, M.M., Kokkinos, I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In: *Computer Vision and Pattern Recognition (CVPR)*, 2010 IEEE Conference on, IEEE (2010) 1704–1711
- [32] Raviv, D., Bronstein, M.M., Bronstein, A.M., Kimmel, R.: Volumetric heat kernel signatures. In: *Proceedings of the ACM workshop on 3D object retrieval*, ACM (2010) 39–44
- [33] Aubry, M., Schlickeewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: *Computer Vision Workshops*, 2011 IEEE International Conference on, IEEE (2011) 1626–1633
- [34] Reuter, M., Wolter, F.E., Peinecke, N.: Laplace-beltrami spectra as shape-dna of surfaces and solids. *Computer-Aided Design* **38**(4) (2006) 342–366
- [35] Konukoglu, E., Glocker, B., Criminisi, A., Pohl, K.: Wesd-weighted spectral distance for measuring shape dissimilarity. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **35**(9) (2013) 2284–2297
- [36] Ovsjanikov, M., M.B-Chen, Solomon, J., Butscher, A., Guibas, L.: Functional maps: A flexible representation of maps between shapes. *ACM Trans on Graphics (TOG)* **31**(4) (2013) 30–30
- [37] Ovsjanikov, M., Merigot, Q., Patraucean, V., Guibas, L.: Shape matching via quotient spaces. *Comp Graphics Forum* **32**(5) (2013) 1–11
- [38] Kovnatsky, A., Bronstein, M.M., Bronstein, A.M., Glashoff, K., Kimmel, R.: Coupled quasi-harmonic bases. *Computer Graphics Forum* **32**(2) (2013) 439–448
- [39] Rustamov, R., Lipman, Y., Funkhouser, T.: Interior distance using barycentric coordinates. *Comp Graphics Forum* **28**(5) (2009) 1279–1288
- [40] Rustamov, R.: Interpolated eigenfunctions for volumetric shape processing. *Vis Comput* **27**(11) (2011) 951–961
- [41] Litman, R., Bronstein, A., Bronstein, M.: Stable volumetric features in deformable shapes. *Computers & Graphics* **36**(3) (2012) 569–576
- [42] Lian, Z., Godil, A., Bustos, B., Daoudi, M., Hermans, J., Kawamura, S., Kurita, Y., Lavoué, G., Van Nguyen, H., Ohbuchi, R., et al.: Shrec’11 track: Shape retrieval on non-rigid 3d watertight meshes. *3DOR* **11** (2011) 79–88
- [43] Mumford, D., Shah, J.: Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on pure and applied mathematics* **42**(5) (1989) 577–685
- [44] Shah, J.: Gray skeletons and segmentation of shapes. *Computer Vision and Image Understanding* **99**(1) (2005) 96–109
- [45] Nishijima, K.: *Fundamental Particles*. W.A. Benjamin Inc., New York (1963)
- [46] Grigoryan, A.: Heat kernels on weighted manifolds and applications. *Cont. Math* **398** (2006) 93–191
- [47] O’Connell, N.: Conditioned random walks and the rsk correspondence. *Journal of Physics A: Mathematical and General* **36**(12) (2003) 3049
- [48] Doob, J.L.: *Classical potential theory and its probabilistic counterpart*. Volume 262. Springer-Verlag (1984)
- [49] Varadhan, S.S.: On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics* **20**(2) (1967) 431–455
- [50] Ovsjanikov, M., Merigot, Q., Memoli, F., Guibas, L.: One point isometric matching with the heat kernel. *Comp Graphics Forum* **29**(5) (2010) 1555–1564
- [51] Paatero, P., Tapper, U.: Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics* **5**(2) (1994) 111–126
- [52] Lee, D.D., Seung, H.S.: Learning the parts of objects by non-negative matrix factorization. *Nature* **401**(6755) (1999) 788–791

- [53] Hoyer, P.O.: Non-negative matrix factorization with sparseness constraints. *The Journal of Machine Learning Research* **5** (2004) 1457–1469
- [54] Olshausen, B.A., et al.: Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**(6583) (1996) 607–609
- [55] Tibshirani, R.: Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)* (1996) 267–288
- [56] Mairal, J., Bach, F., Ponce, J., Sapiro, G.: Online learning for matrix factorization and sparse coding. *The Journal of Machine Learning Research* **11** (2010) 19–60
- [57] Mairal, J., Bach, F., Ponce, J., Sapiro, G., Jenatton, R.: *Spams: Sparse modeling software* (2011)
- [58] Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**(5500) (2000) 2323–2326
- [59] Tipping, M.E., Bishop, C.M.: Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **61**(3) (1999) 611–622
- [60] Bell, N., Garland, M.: *Cusp: Generic parallel algorithms for sparse matrix and graph computations* (2012) Version 0.3.0.
- [61] Bolz, J., Farmer, I., Grinspun, E., Schröder, P.: Sparse matrix solvers on the gpu: conjugate gradients and multigrid. *ACM Transactions on Graphics (TOG)* **22**(3) (2003) 917–924
- [62] Schmitz, P.G., Ying, L.: A fast direct solver for elliptic problems on general meshes in 2d. *Journal of Computational Physics* **231**(4) (2012) 1314–1338
- [63] Xia, J., Chandrasekaran, S., Gu, M., Li, X.S.: Superfast multifrontal method for large structured linear systems of equations. *SIAM Journal on Matrix Analysis and Applications* **31**(3) (2009) 1382–1411
- [64] McAdams, A., Sifakis, E., Teran, J.: A parallel multigrid poisson solver for fluids simulation on large grids. In: *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association (2010) 65–74
- [65] Hu, M.K.: Visual pattern recognition by moment invariants. *Information Theory, IRE Transactions on* **8**(2) (1962) 179–187
- [66] Bronstein, A.M., Bronstein, M.M., Guibas, L.J., Ovsjanikov, M.: Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics (TOG)* **30**(1) (2011) 1
- [67] Jégou, H., Douze, M., Schmid, C., Pérez, P.: Aggregating local descriptors into a compact image representation. In: *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, IEEE (2010) 3304–3311
- [68] Jégou, H., Perronnin, F., Douze, M., Sánchez, J., Pérez, P., Schmid, C.: Aggregating local image descriptors into compact codes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **34**(9) (2012) 1704–1716
- [69] Bartlett, M.: The square root transformation in analysis of variance. *Supplement to the Journal of the Royal Statistical Society* (1936) 68–78
- [70] Freeman, M.F., Tukey, J.W.: Transformations related to the angular and the square root. *The Annals of Mathematical Statistics* (1950) 607–611
- [71] Shilane, P., Min, P., Kazhdan, M., Funkhouser, T.: The princeton shape benchmark. In: *Shape Modeling Applications, 2004. Proceedings, IEEE* (2004) 167–178
- [72] Lian, Z., Godil, A., Bustos, B., Daoudi, M., Hermans, J., Kawamura, S., Kurita, Y., Lavoué, G., Van Nguyen, H., Ohbuchi, R., et al.: A comparison of methods for non-rigid 3d shape retrieval. *Pattern Recognition* **46**(1) (2013) 449–461
- [73] Smeets, D., Fabry, T., Hermans, J., Vandermeulen, D., Suetens, P.: Isometric deformation modelling for object recognition. In: *Computer Analysis of Images and Patterns, Springer* (2009) 757–765
- [74] Lian, Z., Godil, A., Sun, X., Zhang, H.: Non-rigid 3d shape retrieval using multidimensional scaling and bag-of-features. In: *Image Processing (ICIP), 2010 17th IEEE International Conference on*, IEEE (2010) 3181–3184
- [75] Smeets, D., Keustermans, J., Vandermeulen, D., Suetens, P.: meshsift: Local surface features for 3d face recognition under expression variations and partial data. *Computer Vision and Image Understanding* **117**(2) (2013) 158–169