

Understanding and Exploiting Object Interaction Landscapes

SÖREN PIRK, Stanford University

VOJTECH KRS, KAIMO HU, SUREN DEEPAK RAJASEKARAN, and HAO KANG, Purdue University

YUSUKE YOSHIYASU, CNRS-AIST-JRL

BEDRICH BENES, Purdue University

LEONIDAS J. GUIBAS, Stanford University

Interactions play a key role in understanding objects and scenes for both, virtual and real world agents. We introduce a new general representation for proximal interactions among physical objects that is agnostic to the type of objects or interaction involved. The representation is based on tracking particles on one of the participating objects and then observing them with sensors appropriately placed in the interaction volume or on the interaction surfaces. We show how to factorize these interaction descriptors and project them into a particular participating object so as to obtain a new functional descriptor for that object, its *interaction landscape*, capturing its observed use in a spatio-temporal framework. Interaction landscapes are independent of the particular interaction and capture subtle dynamic effects in how objects move and behave when in functional use. Our method relates objects based on their function, establishes correspondences between shapes based on functional key points and regions, and retrieves peer and partner objects with respect to an interaction.

CCS Concepts: • **Computing methodologies** → **Shape analysis**; *Modeling methodologies*; Shape representations;

Additional Key Words and Phrases: Object functionality analysis, shape analysis, affordance analysis, object semantics, physical interactions, geometric modeling

ACM Reference format:

Sören Pirk, Vojtech Krs, Kaimo Hu, Suren Deepak Rajasekaran, Hao Kang, Yusuke Yoshiyasu, Bedrich Benes, and Leonidas J. Guibas. 2017. Understanding and Exploiting Object Interaction Landscapes. *ACM Trans. Graph.* 9, 4, Article 39 (March 2017), 15 pages. DOI: 0000001.0000001_2

1 INTRODUCTION

The goal of this work is to define a descriptor to explore interactions among rigid physical objects (animate or inanimate), such as a human sitting on a chair, grasping a cup, as well as water pouring into a glass, air flowing past a driving car, etc. Such interactions usually involve pairs or multiple objects in relative proximity or contact. Current representations of interactions are polarized at two extremes. On the one hand, there are very high level interaction descriptions like those in natural language ("John grasped his mug and poured coffee into it"); computer vision has expended a lot of effort on mapping images and videos to such descriptions (action recognition). On the other hand, there are the detailed simulations

The work was supported by the NSF grants CCF-1514305, IIS-1528025, EEC-1606396, NSG-1161480, CDS&E 1608762, EEC-1606396, ONR MURI N00014-13-1-0341, the Stanford AI Lab-Toyota Center for Artificial Intelligence Research, a Google Focused Research Award, the Max Planck Center for Visual Computing and Communications, the JSPS Strategic Young Researchers Visits Program for Acceleration Brain Circulations, and the National Science Foundation of China (61373071).

© 2017 ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, https://doi.org/0000001.0000001_2.

performed in computer graphics requiring full representation of the geometry of the environment, e.g., a CAD model of a mug, finite element models for deformable objects like human fingers grasping the mug, as well as material properties such as viscosity for the fluid (coffee) flowing into the mug.

In our work we aim for an *intermediate-level representation* capturing more information about the geometry and physics of the interaction than just a natural language description, yet abstracted away from all the details tied to the specific geometry and physics discretizations used in the simulation. Our motivation is to compare and understand the similarity between the actions of a human sitting on a chair or on a bed, while being able to distinguish the actions of a human hand grasping a mug from its handle, as opposed to from its rim.

We believe that the types of actions that an object participates in can provide insight into the semantic identity of the object. We want to understand in a more computational form the old notion that form and function are deeply inter-correlated. We study how to project actions and interactions onto each of the participating objects or media, so that an object's interactions become a new signature for the object in a way that captures the essential spatio-temporal aspects or patterns of the interaction *from the point of view of the object*. Our goal is, given an object, to summarize the interactions it has experienced into a mathematical descriptor, and to use such interaction landscapes to retrieve similar objects, to recover partner objects that fit a particular interaction (given a screw model, find the appropriate screwdriver), or to infer functionality of a new object.

Understanding shape functionality remains one of the key challenges in shape analysis and geometric modeling. Researchers have not only studied the perception of shapes [5, 49], but have also analyzed their structure, form, and similarity with a wide variety of computational models and representations [8, 32]. It has been recognized that shape modeling is a fundamental problem that must include a high-level understanding of the semantic meaning of a shape — one of the key observations of previous studies is that a shape is not purely described by its form but also by its function [21, 22]. Embedding a shape in its interaction space enables a more thorough understanding of its functionality, a key requirement in many application areas, e.g., shape correspondence and similarity, reconstruction and content creation.

Some of the previous approaches formulate the problem of understanding shapes by looking at human-centric interactions to learn shape affordances [25]. Moreover, capturing real humans allows us to learn the function of an environment and to infer possible interactions for unobserved scenes [42]. Others analyze shapes by




Fig. 1. Different interaction landscapes representing the interactions of a motion driver with a static object. We capture the motion trajectories (red) and encode their signatures into a descriptor that can be used to compare interactions. From left to right: a cloth simulation interacting with a support structure, a human model walking on a floor, a wind simulation interacting with a car, and a robotic hand grasping a cup.

defining functionality models or by using distinct labels for classifying interactions [26, 47]. More recently, it has also been recognized that spatial relations can infer a context-aware signature of object-object interactions [22] that even allow the analysis of shapes [21]. The diversity of these contributions illustrates the breadth of the interaction modeling area.

We introduce a descriptor that captures the spatio-temporal motions of interactions. We look at the trajectories of motion particles sampled from a partner object, as they move in the vicinity of the object of interest. We call this partner object the *motion driver*, as it generates these particle motions. In our coffee mug example, we look at particles on the human hand grasping the mug, or particles in the coffee fluid pouring into the mug. To observe the partner motion particle trajectories, we establish a set of *sensors* on or near the object of interest, each of which is tracking particle motions in a region of free space near it. The sensors then build a descriptor of the trajectories observed in the part of space each is responsible for — and altogether these local descriptors form our interaction signature.

We perform a number of experiments to demonstrate the relevance of our interaction descriptor for analyzing the functionality of shapes. In particular, we show that it is well-suited for shape correspondence and classification tasks. Additionally, we show that our approach can capture and encode multiple functions of an object, *i.e.*, a cup can be grasped at the top for carrying it, while it can also be grasped from the side for drinking. Fig. 1 shows examples of four interaction landscapes, each representing the interaction of a motion driver with a static object. We capture the motion trajectories (shown in red) and encode their signatures with a novel interaction encoding method. In summary, we claim the following contributions:

- We introduce a novel representation for proximal interactions between physical objects and a flexible framework that allows an exploration of the interaction dynamics driven by input data from various sources, *e.g.*, simulations, animations, motion capture, and RGB-D scans.
- We show how to factorize interactions into descriptors for each of the participating objects and how to compare such descriptors. The resulting interaction landscapes establish correspondences between shapes based on functionally-defined keypoints or regions.

- We evaluate our method against state-of-the-art descriptors such as the Light-field Descriptor (LFD), the Intersection Bisector Surface (IBS) and the Interaction Context (ICON), and show the relevance of our descriptor for applications, such as shape retrieval and saliency estimation.

2 RELATED WORK

Shape completion and segmentation, similarity assessment of shapes, and structure-aware modeling have been recognized as fundamental problems in a variety of application domains [32]. Traditionally, many of the previous methods in these areas rely on geometric descriptors to encode local features of shapes and forms. By employing functions on the geometry, these methods provide more discriminative attributes for relating shapes, often formulated with the goal of being robust or invariant against specific transformations [8]. Functions on shapes have also been the basis for the functional formulation of maps and correspondences between shapes [37]. More recently, efforts have shifted towards encoding the structural meaning of shape parts and high-level semantic information. Structure-aware modeling of shapes requires knowledge of symmetries [33, 48], similarities of parts [45], variability [54] and the object’s embedding in its context [22, 53].

Analyzing the uses and affordances of a model provides a deeper understanding of shapes. While early attempts explicitly model the functionality of shapes in the context of specific categories [47] or decompose models into sets of volumetric functional shape primitives [39], representing functionality allows the resolution of object-functionality correspondences, even for high variance data sets. Pechuk et al. [38] derive functionality labels by employing a multi-level hierarchy for the automated classification of parts. Laga et al. [26] encode the context of a model as structural relationship. Their approach models shapes as interconnected parts in order to derive pairwise-correspondences of 3D shapes and to learn the functionality of shape components through supervised classification. Zhu et al. [55] propose an indirect learning approach of poses through training a knowledge base.

Several approaches exist to explore interactions of human agents with shapes or even entire scenes. Bar-Aviv et al. [6] and Liu et al. [31] propose agent-centric systems that analyze the performance of human avatars for classification and interaction tasks. While Grabner et al. [18] only concentrate on analyzing the functionality

of chairs, Kim et al. [25] detect functional similarities by predicting human poses based on a trained affordance model. The human agent can also be observed from images [12] and videos [19]. Li et al. [30] and Song et al. [46] explicitly limit themselves to investigating visual grasp affordances from 2D measurements.

In SceneGrok, Savva et al. [42] observe the behavior of real humans interacting with an environment. They capture RGB-D data to infer action maps and train a classifier that allows the transport of the interaction knowledge to other, unseen environments. In a more recent work Savva et al. [43] show that it is possible to infer higher level semantic meaning from observed human-centric interactions. Knowledge about interactions can be employed as a means for reconstructing objects and scenes. Fisher et al. [16] used virtual agents to associate object arrangements with typical activities to derive more plausible scene setups. Tzionas et al. [50] introduced a 3D reconstruction pipeline based on how hands interact with a target object, which even works for featureless or symmetric shapes. Instead of focusing on the interaction of real humans with their environment, other works concentrate on encoding dynamic motions of virtual humanoid agents for retargeting and motion adaptation purposes [2, 20].

In contrast to the previous methods, we are interested in exploring interaction partners – complementary shapes that accomplish a function jointly with the given shape. This notion is similar to what Cain [10] proposed as *Motion Constraints*; however, our method focuses on capturing the dynamic nature of interactions of two or more shapes and not just static relationships. Moreover, we show that interactions can be factorized by analyzing their motion patterns in an appropriate vector space [52]. Related mathematical ideas based on “currents” for describing surfaces can also be found in Dürlemann [15].

Our approach is similar to the *Shape Flow* method proposed by Jiang and Martin [23], who detect and encode flow lines in videos and the *HON4D* descriptor introduced by Oreifej and Liu [35] that encodes interactions captured in RGB-D data. Unlike the previous specialized approaches, we propose a more general framework that allows us to encode a wide variety of different interactions in 3D setups. Additionally, we show that our signatures show similar characteristics to the recently proposed ICON descriptor [22] for static scenes, which makes them well suited for tasks such as estimating shape correspondences and assessing shape similarity.

Finally, our novel particle-based representation has multiple advantages over prior interaction encodings. It is independent of the specific geometric and physical model used to obtain the interaction simulation. Moreover, trajectory collections are nice mathematical objects on which a variety of analysis tools, such as vector fields, flow analysis, etc., can be applied.

3 OVERVIEW

The objective of our work is to encode and abstract the dynamic characteristics of interactions. More specifically, we focus on interactions of one object in motion, which we refer to as *motion driver*, with another static object (e.g., a chair or a cup) that is acting as the action receiver. The motion driver can be any source of motion; to explore the versatility of our approach, we have employed motions

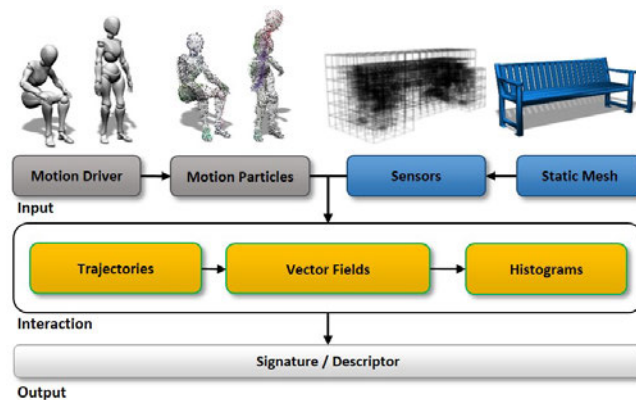


Fig. 2. Overview of the pipeline for computing interaction landscapes. Given an input mesh and a motion driver, we track the movement of motion particles attached to the moving object in local sensor regions. In each region, the sampled trajectories of the motion particles are converted into vector fields that are analyzed for unique patterns in the motion flow. By quantizing the vector space we derive a global signature that allows us to embed objects in their interaction space and to relate them based on their functionality.

from a variety of sources, such as simulators specialized for grasping or fluids, character animations, physics-based motions, and motion captured data.

The surface of the motion driver (see Fig. 2) is covered with *motion particles* that we use to track the motion. For particle-based simulations, such as Lagrangian fluids, we randomly select a subset of the actual simulation particles to track. Depending on the type of motion that causes the interaction, we provide the system with a number of parameters, such as the location and orientation of an emitter, the number of particles for a fluid flow, or the number and properties of transformations for animations and motion-capture data sets.

Our approach is based on the observation that the interaction of two shapes defines a unique pattern that characterizes their form and function. To encode the dynamics of an interaction we examine how a motion particle approaches an object. We sample the surface of the input shape to distribute sensors that track motion particles passing by, with each sensor measuring a region on or near the surface. These regions span an interaction space that we refer to as the *interaction landscape*. As it is not known where the interactions are likely to happen beforehand, we spatially subdivide the potential interaction space based on the input geometry. Every motion particle that enters one of the sensor regions creates a unique motion trajectory that characterizes the occurring interaction.

We transform the captured trajectories into local vector fields for each of the sensors. Operating on vector fields allows us to analyze the properties of the observed interaction using powerful mathematical tools. More specifically, we detect a number of attributes of the vector fields, such as the vector and tensor magnitude, dilation and vorticity, as well as particle velocity and orientation. We build histograms over these attributes to factorize the interaction and capture the response of different shapes to the same motion driver by computing their relative distance over a set of histograms.




Fig. 3. A human model sitting down on various chairs. Depending on the performed action, the motion trajectories, shown in red, can vary: compare (a) and (d). We capture the trajectories by uniformly distributing motion particles on the surface of the motion driver (human model). We record the direction and speed of the movement of each motion particle: (d), (e), and aggregate their trajectories into signatures.

4 INTERACTION LANDSCAPES

An interaction landscape is a 4-dimensional (three spatial and one temporal) interaction space of two shapes participating in an interaction. The input to our algorithm is a 2-dimensional surface mesh M of the observed static shape that can even be a triangle soup and a time variant motion driver D (the partner moving object or medium), both embedded in \mathbb{R}^3 . The motion driver encodes the source of movement that initiates the interaction; for our tests we used moving surfaces and simulations. We define an interaction landscape by specifying a set of sensors $S = \{s_1, \dots, s_n\}$ distributed in the interaction space of M and motion particles $A = \{a_1, \dots, a_m\}$ originating from D . The interaction space U is a fixed size bounding box with manually added extra space around the object, large enough to encompass all objects of a particular class. Depending on the type of motion, motion particles are either stochastically sampled on the surface (e.g., for animations), or directly advected by the simulation (e.g., for fluids). Moreover, we expect shapes of the same class to share a common orientation and require a canonical arrangement of the interaction landscape with respect to the observed shape.

4.1 Motion Drivers

Our method provides a more general means for exploring different types of interactions represented by the motions of a motion driver. To capture the variance of motions of shape parts relative to each other and against the observed shape, we compute a distribution of motion particles on the surface of the motion driver that covers the entire surface but with an emphasis on salient regions.

We ensure this distribution of samples on the surface of D by employing a bilateral farthest point sampling strategy. Given a motion driver D , we first initialize the particle set with a single position entry $P = \{p_1\}$, with p_1 being stochastically sampled on D . We iteratively add the farthest point s on D to P until a specified number of particles is distributed. To capture salient features of D , we use a bilateral distance measure, that enlarges distances in regions of the surface with finer details and thereby emphasizes the sampling in these regions [14]. We add new samples to P by

computing a density function ρ over all mesh vertices based on their *local feature sizes (lfs)* and define the distance measure $\bar{d}(s, p)$ as

$$\bar{d}(s, p) = d(s, p) \cdot \rho, \quad \rho(p) = 1/\text{lfs}(p)^2 \quad (1)$$

for each existing sample $p \in P$, where s is a particle sample candidate, $d(s, p)$ the Euclidean distance between s and p in \mathbb{R}^3 , and $\text{lfs}(p)$ the inverse density function that measures the saliency of p based on its distance to the medial axis of the mesh [3]. Vertices of the mesh are denoted as s and p , and the sample candidates s are selected from the remaining set of vertices of the mesh. For each p_i in P we compute the distance according to Eq. 1. This definition is not symmetric as we only consider the saliency of the points that already exist in P . The lfs is the distance to the medial axis of the mesh, which is insensitive to its resolution. To ensure plausible results for non-smooth shapes, we approximate the medial axis as discussed in Li et al. [29]. If the lfs value for a vertex is smaller than a threshold δ it is revised as δ , which ensures non-zero values for all vertices. Fig. 3 (e) and (f) illustrate the resulting farthest point sampling for an animated surface mesh.

For fluid flows and other Lagrangian simulations, positions for motion particles are directly derived from the simulation. In our implementation, fluid particles are randomly initiated. Hence, we can stochastically sample a subset of motion particles from the set of fluid particles. We track the selected motion particles until the life time of the corresponding fluid particle has ended and initiate a new particle at the fluid emitter. Other implementations of fluid simulations might require more refined sampling strategies to assure complete coverage of particle positions across the motion driver. For most of our tests we used a few hundred to a couple of thousand particles for both, moving surfaces and fluid flows.

4.2 Sensor Placement and Particle Tracking

To track the movement of motion particles we distribute a set of sensors S in the interaction space U . As we cannot assume any prior knowledge about where an interaction is going to happen (i.e., for simulations), we distribute sensors in the entire domain of U , by subdividing the space with an Octree. To emphasize the importance of motions in the vicinity of M , we produce a graded set of cubes



Fig. 4. We capture interactions by defining sensor regions in the domain around the observed shape (left). We compute samples on its surface (middle) and use them to produce a spatial subdivision of a graded set of cubes that get finer close to its surface. Location and size of each sensor correspond to the leaf nodes of an Octree that generates the sensor regions (right).

that gets finer closer to the shape surface. As we cannot rely on the tessellation of M to produce the levels of the Octree, we sample the surface with a 3D variant of *Poisson Disk Sampling*. We randomly select a triangle T of M and compute a sample candidate by

$$e_c = (1 - \sqrt{r_1})T_a + \sqrt{r_1}(1 - r_2)T_b + (\sqrt{r_1}r_2)T_c, \quad (2)$$

where T_a, T_b, T_c are the vertices of the triangle T in \mathbb{R}^3 and r_1, r_2 are random numbers in $[0, 1]$ [36]. We add e_c to the sample set E when $\forall e_i \in E : |e_i - e_c| > c$, where c is a user defined threshold that specifies the minimum distance between the samples. We then insert the samples into the Octree and subdivide each of its cells as long as they contain more than one sample. Thereby, threshold c is a control parameter for the depth of the Octree and hence implicitly defines the number of cells (Fig. 4). For most of our tests we used sample distances from 0.01 to 0.5, which roughly corresponds to 1% of the width of the interaction space U .

Sensors are placed at the cell centers of the Octree, each observing a region defined by the corresponding cell. This allows us to maintain cuboidal sensor regions, while the sensors are adaptively placed according to their distance to the input shape. Intuitively, motions close to the surface contribute the most in describing the interaction; however, we found that distant motions, can also be important for thoroughly describing an interaction. This can be seen in Fig. 3, (a-c): the head tilt of the model produces unique motion trajectories (visualized in red) that help to discriminate sitting interactions from other motion types. Moreover, the arrangement of shape parts of M affects how the motion driver interacts with the shape — some regions are more important than others. Employing an adaptive spatial subdivision allows us to place sensors with sensing regions of different sizes and thereby to capture particle motions with varying degrees of precision.

Each sensor s_i defines a cuboidal region r_{s_i} of size q_{s_i} , which tracks bypassing motion particles. We sample each particle's motion trajectory in r_{s_i} with a constant time step Δt and retain its position and velocity as it moves through the volume over time. After capturing the interaction, we uniformly discretize r_{s_i} into cells $c_{i,j}$ and average the positions and directions of all trajectory samples that fall into the same cell $c_{i,j}$ resulting in the vector $u_{c_{i,j}}$. For each sensor s_i we build a vector field v_i through trilinear interpolation of the discretized sensor region. The set of vector fields $V = \{v_1, \dots, v_n\}$ represents the observed interaction of two participating shapes — so we can now proceed to analyze the properties of the interaction by detecting patterns in sets of such vector fields.

4.3 Encoding Motion Flows

Each interaction is represented as a set of vector fields. To facilitate the comparison and clustering of interactions we define a distance measure between two sets of vector fields V_1 and V_2 , each representing one interaction of a motion driver and a static object. To compute this distance we compare locally captured motion flows. We compute the gradient for each vector field $v_i \in V_1, V_2$ to derive a set of its first-order attributes [52], which associate a real value to each cell $c_{i,j}$. In particular, we measure the magnitude of the attributes vorticity, dilatation, shear strain rate, and the gradient tensor field T (see Appendix A for the precise definitions). Additionally, we consider the magnitude of vector $u_{c_{i,j}}$ and the dot product of $u_{c_{i,j}}$ with the closest surface normal of the mesh as two vector attributes. Computing this set of attributes on v_i allows us to quantify the measured interactions I_1 and I_2 by computing a set of histograms.

Histograms. We compute a histogram for each attribute a of the vector field v_i , forming a set of histograms

$$\mathcal{H} = \{h_a, a \in A\}, \quad (3)$$

where h_a is a normalized histogram over the attribute a (Fig. 5, a). The histogram h_a is built by binning the corresponding attribute values of a vector field in the range $[0, 1]$. The contribution to the individual bins is weighted by an exponential scaling of the Euclidean distance of the relative position $p_{c_{i,j}}$ of cell $c_{i,j}$ to the closest point p_M on the surface of M . We compute the value of a given bin m_k by

$$m_k = \frac{1}{N} \sum_{j=1}^{n_c} t_j \exp\left(-\left(\frac{\|p_{c_{i,j}} - p_M\|}{r}\right)^2\right) \cdot \chi_E(t_j), \quad (4)$$

where n_c is the number of cells in the vector field, t_j the attribute value in the current cell, r the size of the sensor s_i and χ_E the indicator function that associates $\chi_E = 1$ if $t_j \in E$ and $\chi_E = 0$ otherwise, where $E = [\frac{k}{N}, \frac{k+1}{N}[$. Each bin is normalized by N , the number of all bins. The exponential weighting ensures that distant motions have less influence than those in the vicinity of M . Finally, we compute a set of global histograms

$$\overline{\mathcal{H}} = \{\overline{h}_a, a \in A\}, \quad (5)$$

where each histogram \overline{h}_a represents the average of all histograms h_a of the attribute a across all vector fields $v_i \in V$ (Fig. 5, b). To account for differently sized sensor regions we compute each bin \overline{m}_k of \overline{h}_a as

$$\overline{m}_k = \frac{\sum_{j=1}^n m_{j,k} r_j^3}{\sum_{j=1}^n r_j^3}, \quad (6)$$

where n is the number of histograms (and vector fields) of one attribute, $m_{j,k}$ the k -th bin of the j -th histogram h_a and q_j , the size of the corresponding sensor s . Additionally, we only consider sensors that tracked motion particles to reduce the complexity. For the sake of a clear presentation we describe the computation of the histograms as a two stage process, however, the global histograms can be directly computed from the vector fields.

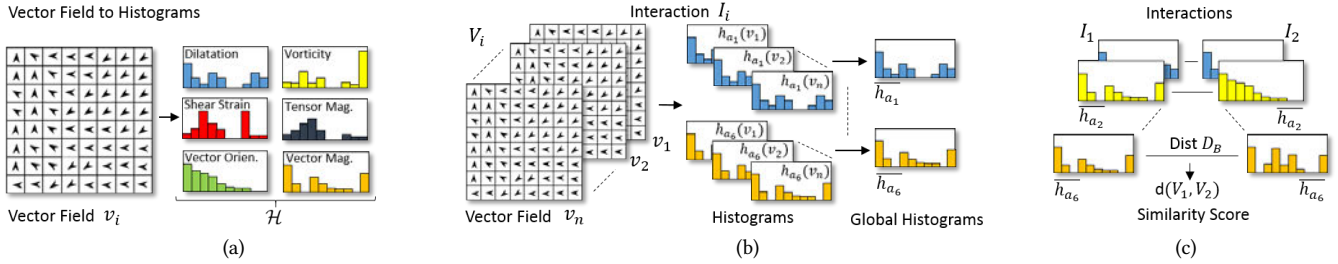


Fig. 5. A 2D illustration of our descriptor: each sensor region is converted into a vector field that allows us to derive a set of six attributes describing patterns in the captured motion flows, which we quantize by computing histograms for each attribute (a). We average the local attribute histograms over all vector fields to produce a global histogram for each attribute (b). Finally, we define a similarity score between two interactions by computing the sum of relative distances between two corresponding global histograms (c).

Interaction Descriptor. We define our interaction descriptor as a set of global histograms, encoding the motion patterns of an interaction of a motion driver and a static object. We compare interactions (Fig. 5, c) by computing the relative distance between two sets of histograms as the similarity score:

$$d(V_1, V_2) = \frac{\sum_{a \in A} w_a D_B(h_a, k_a)}{|A|}, \quad (7)$$

where w_a is a weight for attribute a , $|A|$ the number of attributes and D_B the Bhattacharyya distance between two global histograms h_a and k_a , associated with V_1 and V_2 respectively. Operating on the global histograms allows us to compare interactions, even if the number of used sensors for each of them differs. We use the Bhattacharyya distance, as it has proved effective at relating semantically similar interactions and discriminating ones that differ. Moreover, we found the weights w_a by experimentation; dilatation magnitude and vector orientation turn out to be the most important attributes, while shear strain rate and vector magnitude do not contribute much to interaction comparison.

Time variance. A time sensitive descriptor allows to differentiate interactions of varying speeds and time-steps, an important property for dynamic interactions. However, while exploiting this features is desirable in many scenarios, we aim at providing a means of controlling the importance of temporal information. For example, when grasping a cup, the speed of the grasp might differentiate motions, while it can be irrelevant for classifying the shape. Our approach can accommodate both modalities. When discretizing the sensor boxes into vector fields we compute a vector $u_{c_{i,j}}$ for each cell $c_{i,j}$ in the vector field v_i . This vector is the sum of all trajectory vectors located in $c_{i,j}$ normalized by the factor F :

$$u_{c_{i,j}} = \frac{\sum_{u_k \in c_{i,j}} u_k}{F}. \quad (8)$$

To differentiate similar interactions with different speeds, we choose $F = |c_{i,j}|$, effectively calculating the average vector, which retains the velocity of the motion flow. When we match similar interactions independent of their temporal resolution we choose $F = \|\sum_{u_k \in c_{i,j}} u_k\|$, which normalizes the vectors and thus only maintains information of the motion flow direction.

Alternative Approaches. Instead of averaging the captured motions we performed a number of experiments with more sophisticated mathematical methods for finding matching pairs of vector fields, e.g., through *Hungarian Optimization* and *Spectral Alignment* [28] and tested different distance measures such as Chi-Square or Earth-Mover-Distance [41], but found averaging the signatures performs better on our data sets. Moreover, we employed *Dynamic Time Warping (DTW)* [9] to temporally align the histograms before computing the signatures, however, we found that in most cases DTW did not improve much the quality of the clustering.

5 IMPLEMENTATION AND DATASETS

We implemented our method in C++ with OpenGL and GLSL. The system was deployed on a desktop computer with an Intel Xeon processor clocked at 3.7 GHz, with 32 GB RAM and an NVIDIA GTX 980 GPU. Results shown in the paper were generated with our own framework.

We focus on capturing time-variant interactions of one or multiple motion drivers and a single observed object. The motion trajectories are generated by tracking the motion driver (surface or particles) and by updating the corresponding motion particles efficiently through graphics hardware. Depending on the number of motion particles on the surface and the amount of tracked trajectories, this can be done at interactive rates. We exploit the interactivity of our system to provide the user with control of capturing the motion trajectories. We provide access to the control parameters, such as the number of motion particles, sampling speed of simulations and animations, and parameters for the fluid simulations, e.g., velocity, viscosity, and surface tension.

The initialization step is done during preprocessing and the most time demanding step is the farthest point sampling strategy of the motion driver that is currently implemented without any acceleration and has $O(n^2)$ complexity, where n is the number of mesh vertices. We compute the *local feature size (lfs)* as the approximated medial axis defined in Li et al. [29]. The most demanding part of the runtime module is the tracking of motion particle trajectories. We test if a motion particle is contained in a sensor volume and capture its motion trajectory, resulting in at worst $O(mn)$ complexity, where m and n is the number of sensors and particles, respectively.

We measured the sampling of motion drivers with 80–600 seconds (with 2k particles as the upper limit) and 2–20s for observed objects. Depending on the recording time, we capture around 100k–8M




Fig. 6. The simulation of a humanoid hand grasping small scale man-made objects: cup (a), phone (b), and bottle (c). The motion trajectories indicate a series of interactions performed on the corresponding object. Interactions of a fluid and a cup: the SPH fluid particles are poured into a cup (d), and their trajectories are captured (e). The trajectories of a stirring motion from a different interaction (f).




Fig. 7. Real world interactions: we used the SceneGrok [Savva et al. 2014] dataset which includes 45 RGB-D captures in 14 environments (a) to capture motion trajectories of humans (b) performing interactions that can be categorized into 8 classes. (c-e): motion capture data of a human interaction with physical objects can also be used in our framework.

Fig.	Motion Driver			Object		Sensors			Trajectories		Signature
	MP	MT	MY	OS	OT	S	OL	DS	T	TS	ST
3 (a)	2.0k	552s	AN	0.7k	0.40s	2.6k	8	6	17k	991k	0.54s
3 (b)	2.0k	552s	AN	0.3k	0.20s	1.6k	8	6	13k	896k	0.39s
16 (*)	2.0k	581s	AN	1.1k	0.43s	3.3k	8	6	14k	907k	0.65s
6 (e)	2.0k	-	FL	5.5k	2.23s	17.7k	7	5	28k	601k	2.82s
8 (*)	2.0k	-	FL	3.6k	1.76s	4.5k	8	10	10k	242k	1.74s
16 (†)	2.0k	-	FL	1.8k	0.65s	4.8k	6	6	20k	570k	0.65s
17 (*)	2.0k	-	FL	0.8k	0.21s	2.9k	8	10	22k	566k	0.57s
6 (a)	1.5k	161s	GR	0.9k	0.56s	3.2k	7	5	18k	201k	0.81s
8 (†)	1.6k	154s	GR	1.6k	0.75s	5.6k	8	5	21k	294k	1.63s
7 (a)	1.8k	80s	SG	1.8k	1.75s	6.6k	10	15	51k	639k	1.11s
7 (d)	2.0k	119s	MC	0.5k	0.18s	2.1k	9	8	5k	79k	0.34s

Table 1. Details on the processed interactions of some of the figures shown in the paper; MP=Motion Particles, MT=Motion Driver Sample Time, OS=Object Samples, OT=Object Sampling Time, S=Sensors, OL=Octree Level, DS=Domain Size, T=Trajectories, TS=Trajectory Samples, ST=Signature Time. The table is sorted by the type (MY) of motion driver: AN=Animation, FL=Fluid Simulation, GR=GraspIt! Simulation, SG=SceneGrok, MC=MotionCapture. We sampled motion trajectories with $\Delta t = 25ms$ and set the sample distance for the observed object to 0.1.

trajectory samples which roughly amounts to 10–80MB of motion data. Computing the signatures, which involves calculating vector fields and their corresponding histograms from the trajectories takes 0.30–3.0 seconds. Table 1 provides details on the complexity of some of the interactions shown in the paper.

5.1 Datasets

Shapes can interact in a variety of ways, which makes comparing and aggregating them a challenging task. To show the versatility of our system we employed data from various sources. Simulations and RGB-D data provide more meaningful signatures compared to precomputed data sets. This is mostly because rigged animations and motion capture data cannot be easily acquired, resulting in a sparse representation of an interaction when used for computing

the distances. Simulations can be used with arbitrary input meshes and can dynamically adapt. Moreover, we found that aggregating repeated interactions of the same two participating shapes improves the quality of the resulting signatures.

Simulations. We explored two different types of simulation. First, we explored interactions of a human hand grasping objects. We used the *GraspIt!* simulator [40], which is a publicly available tool for simulating grasping motions of human and robot hand designs. It accommodates arbitrary input meshes and detects the quality of grasping configurations by simulating collisions of hands and obstacles. We used *GraspIt!* to simulate the interaction of human and robotic hands with small scale man-made objects, e.g., cups, cellular phones, and books (Fig. 6, a-c). We do not simulate any advanced effects, such as force closure or precision grasps, but instead rely on *GraspIt!*'s support for detecting collisions and determining contact.

Second, we simulated fluids with Smoothed Particle Hydrodynamics (SPH) [34]. Unlike Eulerian approaches, particles are trackable in 3D space and thus easily allow detecting fluid-obstacle collisions. We randomly select fluid particles to capture motion trajectories of the fluid flow. We used SPH to explore objects interacting with fluids, such as a fluid poured in a cup, and wind fields, such as planes and cars. Fig. 6 (d-f) and 8 (middle row) show examples of fluid flows and the captured motion trajectories.

Animations. To account for full-body interactions of humans with large scale rigid objects such as furniture, we used animated humanoid meshes and motion captured data. We employed models from Adobe's Mixamo [1] model repository that provides a large number of humanoid models rigged to perform certain actions, e.g., jumping, sitting on a chair, or walking. We manually fitted these models to object meshes and captured their motion data when performing the predefined action. Fig. 11 shows three human models performing different motions.

Motion Capture and RGB-D. To show that our method works with real-world data, we employed the RGB-D scans from SceneGrok [42] and motion capture (MoCap) data from the KIT Whole-Body Human Motion Database [24], illustrated in Fig. 7. The SceneGrok dataset contains 45 RGB-D captures in 14 environments representing humans performing interactions that can be categorized into 8 different classes. We use the captured environment as static object and build a humanoid mesh that is updated based on the provided skeleton positions and used as the motion driver.

The MoCap dataset includes high-quality human whole-body motions that involve human-object interactions. In particular, we explored 10 sitting interactions each with 3-4 objects. To generate a surface model we fit the human body model (template mesh) [4] to MoCap data points. Specifically, we embedded a skeleton to the template using automatic rigging software [7] and scaled the body segments of the template to fit the body shape of the captured subject. For both data types we sample point locations from the produced surface mesh (Section 4.1) which is continuously updated to drive motion to capture interaction trajectories.

6 RESULTS, EVALUATION AND APPLICATIONS

We have introduced a novel descriptor for aggregating time-variant interactions of arbitrary shapes and motion drivers for 3-dimensional spaces. Unlike previous approaches we are not limited to static scenes and rigid surfaces but instead allow and exploit time variance for interactions. In this section we show results, provide an evaluation of our method and discuss its performance against other state-of-the-art shape and interaction descriptors. Moreover, we show that knowledge of interactions informs knowledge of the shape function that can be employed for shape classification and saliency estimation.

6.1 Results

Fig. 8 (top row) shows two *multidimensional scaling (MDS)* embeddings of simulated interactions. We see a plot of grasping interactions from *Graspit!* (left), objects in wind tunnels (middle), and a clustering of motion capture interactions (right). For both experiments our descriptor provides a meaningful aggregation of the interaction of shapes that allows us to relate and even retrieve them according to their function (Fig. 16). For the wind experiment, bicycles and motorbikes share some overlap due to their similarities in both their function and shape.

For most of our tests we randomly selected models from the Modelnet [51] and ShapeNet [11] model repositories; our database includes 157 objects of 20 classes. We tested our framework with an automated setup that loads scene files with the specification of motion driver and observed object. We run the interactions for around 60 seconds to capture enough trajectories; animations and *Graspit!* sets are tracked for their entire duration. To capture the interaction of cars, planes and motorbikes with wind fields we setup wind tunnel scenes with fixed particle emitters. For *Graspit!* we randomly chose the positioning of the hand relative to the observed shape for the simulation and aggregated the interaction into a data set readable by our framework.

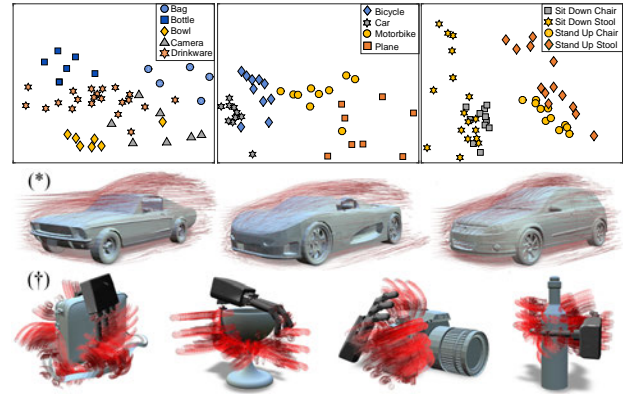


Fig. 8. Top: multidimensional scaling (MDS) plots of isolated clusters of categories. Left: interactions of simulated grasps from *Graspit!* and small scale objects. Middle: a plot of interactions of objects in wind tunnels. Right: motion captures of a human sitting on furniture. Our method generates discernable clusters, objects of overlapping classes share functional commonalities, also noticeable from their visual appearance. Middle + bottom: objects and their motion trajectories used for the MDS embeddings.




Fig. 9. Left: a subset of orientations of the motion driver around the up-axis of the observed shape. Right: a color-coded matrix that visualizes the distances between each of the configurations. We performed three experiments: rotation of the object around the up-axis (T_1), both of the other two axes (T_2) and the rotation of the motion driver (T_3).

Fig. 9 illustrates the behavior of our descriptor with respect to variations in the orientation of observed object and motion driver. On the left we show a subset of arrangements of a hand positioned around the up-axis of the cup. The color-coded distance matrix (right) visualizes the distances of all scenes we used in this experiment. We tested 58 (29 bottle, 29 mug) scenes representing different orientations of the observed object along the up-axis (T_1), both of the other axes (T_2), and of the motion driver (T_3). As we use global histograms to discriminate differences in the interactions, our descriptor is mostly agnostic to variations in the arrangements of shapes and involved motions.

Fig. 10 shows a result of ten scenes of one and two motion drivers interacting with a single static object. Our method allows to capture several interactions at the same time and to encode them as a meaningful signature. The image on the left shows the setup for this experiment. A grasping motion and a fluid are both interacting with a cup as the observed model. The MDS embedding on the




Fig. 10. Encoding two motion drivers. Left: the hand and the fluid are interacting with the observed cup at the same time. Right: a MDS embedding of multiple scenes showing the effect of using two motion drivers.

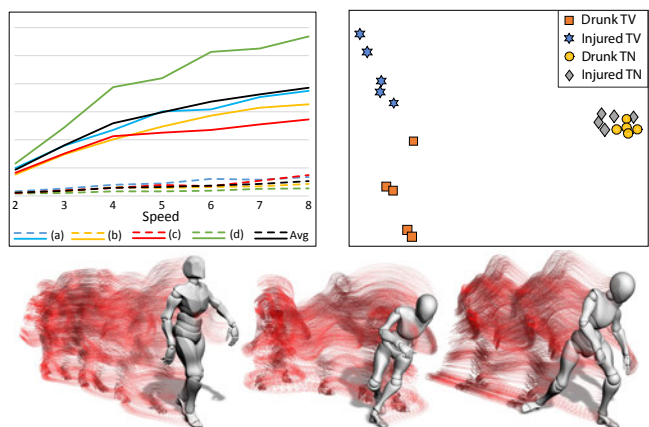


Fig. 11. Top left: relative distances of human-chair interactions (Fig. 3 a-d) showing the effect of normalizing (dashed lines) or emphasizing (solid lines) the velocity of an interaction. Taking the speed into consideration allows us to even differentiate subtle motions. The MDS embedding (top right) shows 20 animations of "drunk" and "injured" walks from the Mixamo model repository (bottom) where we computed the distances with (TV) and without (TN) time variance.

right shows that our method allows to produce a more principled clustering of the interactions when two motion drivers are used.

Finally, our descriptor enables to differentiate interactions based on their speed. Fig. 11 (top, left) shows the relative distances of human-chair interactions shown in Fig. 3 (a-d) with different speeds. As can be seen in the diagram normalizing or emphasizing the velocity in the computation yields different results. Taking the velocities into consideration acts like a switch that allows to further differentiate subtle variations in the motion flows. The MDS embedding (top right) shows 20 animations of "drunk" and "injured" walks from the Mixamo model repository (bottom) where we computed the distances with (TV) and without (TN) time variance (Section 4.3). Taking velocities into consideration yields finer clusters while normalizing the motions causes a stronger grouping.

6.2 Evaluation

Unlike previous methods, our descriptor uses time-varying data. Therefore, comparing it to existing techniques only has limited

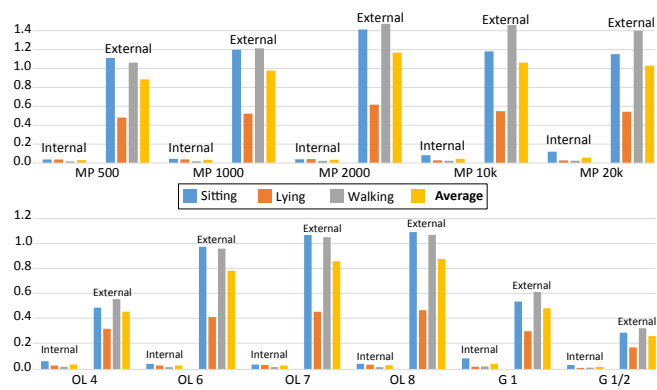


Fig. 12. Comparison of different settings for the number of motion particles and sensors in our framework. Top: the number of motion particles (MP); bottom: the number of sensors defined by the levels of an Octree (OL) and by a uniform grid (G). The diagrams show the relative distances of interactions within the same group (internal) and across different groups (external). We used five interactions per category, where we used different animations as motion driver for the interaction with chairs and beds (sitting, lying) or the floor (walking).

meaning. However, we measure its effectiveness for common shape classification and retrieval tasks. For these evaluations we employ precision and recall queries on our data set and compare the performance to existing descriptors on the same task, but with static scenes. Additionally, we discuss a number of experiments to show the effect of parameters, *e.g.*, the number of motion particles and sensors, specified in our framework.

Parameters and Method. For most of our tests we distributed between a few hundred and a couple of thousand motion particles on moving surfaces or advect them directly through the simulation. We tested different settings for the number of motion particles and sensors and found that using only a few particles generates sparse motion trajectories that do not provide enough data to extract meaningful signatures. On the other hand, placing many particles does not improve the results because of the discretization of sensor regions into vector fields. We computed vector fields of sizes 4^3 , 8^3 , and 16^3 and found that the method is somehow independent to the resolution in this range, however, size 8^3 performs best.

Fig. 12 shows the effect of different settings for the numbers of motion particles and sensors. We evaluated these changes by comparing the relative distances of three sets of animations (sitting, lying, walking) each containing five interactions. The descriptor is most effective when it differentiates interactions from the same group from those of different groups. Therefore, we measure the distance of interactions within the same group (internal) and across different groups (external). The best results were achieved using 2,000 motion particles and sensors generated by an Octree of level 7 or 8, which corresponds to about 3k active sensors for a domain of size 6 and an object sample distance of 0.1. Further increasing the number of the particles and sensors did not provide significantly better results. Additionally, we measured the effectiveness of using an Octree compared to a uniform grid with different cell sizes.




Fig. 13. Precision and recall diagram showing the performance of our descriptor compared to the LFD, IBS and ICON.

Comparison. To evaluate our method, we used the *Interaction Bisector Surface (IBS)* [53] and the more recently introduced *Interaction Context (ICON)* [22] descriptors, both of which describe interactions of groups of spatially arranged shapes by measuring the geometric properties of intersection surfaces between a center object and each of its surrounding shapes. ICON additionally provides *Interaction Regions (IR)* to facilitate functionality descriptors on these datasets that allow to classify shapes based on their function. However, both approaches only operate on static scenes of two or more participating shapes, which makes a direct comparison with our method difficult.

To compare IBS and ICON with our method we selected sets of animations and simulations in our dataset and converted them to a series of static scenes, by selecting snapshots of the motions. For each series, we selected pairs of motion driver and object, where the motion driver is closest to the surface of the observed shape and used this configuration as the input. We tested this approach on six categories: bags, bottles, and cups (grasping); chairs, benches, and beds (animations) which are composed of 65 time-variant scenes. Fig. 13 shows the precision and recall diagram. Our method allows to retrieve objects more reliably than other descriptors. Moreover, it returns objects that share the same functionality, even if their geometry differs.

We also used the light-field descriptor (LFD) [13] to evaluate our method as it provides a baseline for visual similarity based shape retrieval. It determines the similarity of two given shapes by their visual appearance. However, due to the semantic gap between geometry and function, shapes with the same functionality potentially show large geometric variations. Consequently, the performance of LFD is inferior compared to our method, with a precision under 50% for a recall of 60% (Fig. 13). Moreover, our method is capable of retrieving objects that have similar functionality, which LFD cannot reproduce.

Fig. 14 shows the evaluation of our method and ICON on the RGB-D scans from Savva et al. [42]. The dataset includes labeled interactions of a human in an environment captured from RGB-D scans. We use the provided motion information (captured skeleton) to build a stick figure mesh to sample and advect motion particles. The reconstructed environment mesh represents the observed object. We capture each type of interaction for the duration provided in the dataset. For ICON we used the figure mesh to compute the IBS and RS with the environment mesh.





Fig. 14. A set of interactions from SceneGrok showing the interaction of a human with an environment (top). Compared to ICON our method produces similar results with respect to distances across different classes of interactions (external) while it performs better in discerning interactions of the same class (internal).

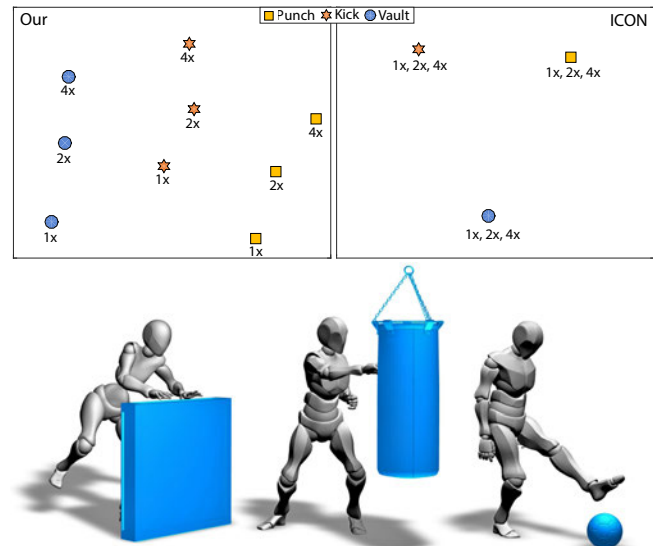


Fig. 15. Top row: multidimensional scaling (MDS) plots of our method (left) and ICON (right). While ICON can only differentiate the types of interactions, our method also allows to discriminate them according to their speed. Bottom row: we used three different interactions for this experiment at different speeds (1x, 2x, 4x).

Although our descriptor does not produce significantly different distances compared to ICON, it emphasizes the difference between interactions of the same groups (internal) and across different groups (external), with a ratio of 8.1 (ICON) to 11.2 (our) for the average distance. However, processing RGB-D scans poses a challenge in that they contain a lot of noise; the location of skeleton joints often jumps from frame to frame. To constrain these artifacts we only use sequences close to the interaction without considering transitions between them. This explains the similarity to ICON as many of the

interactions (e.g., *Read Book* or *Rest Feet*) do not show significant amounts of motion. In case reliable motion information is available (e.g., *Walk Floor*) our descriptor provides better results.

Fig. 15 illustrates the advantage of exploiting motions for discriminating interactions. We captured three animations of humanoid meshes interacting with static objects at different speeds (bottom row). As can be seen in the embedding space (top, left), our method allows to discern the type of interaction (punch, kick, vault) as well as variations in the speed at which they are performed.

Other descriptors do not consider speed as feature, which is illustrated in Fig. 15 (top, right). We used ICON for an evaluation on our datasets by taking static snapshots of the motion driver sampled at relative temporal distances of the animation. Each static snapshot is a similar object-motion arrangement, independent of the variations in speed. Hence, ICON can differentiate the type of interaction, but not the variations of the motion. The MDS plot on the right visualizes the interactions shown on the left. All static snapshots of the same interaction appear the same to ICON and are thus positioned at the same location in the embedding space.

It is important to note that our descriptor, although designed for motion comparison, provides comparable results to descriptors that are designed for static scenes.

6.3 Applications

To demonstrate the effectiveness of our interaction descriptor we illustrate its application for shape classification and saliency estimation. In particular, we perform tests by querying objects from a dataset, by estimating salient regions of shapes based on their interaction response, and by analyzing motions to predict the corresponding type of interaction. It is important to note that we expect shapes of the same class to share a common orientation and require a canonical arrangement of the interaction landscape with respect to the observed shape. Objects can store signatures of multiple interactions with various motion drivers.

Shape Retrieval. We quantify the performance based on the quality of shape retrievals from a dataset containing groups of labeled shapes (157 objects, 20 classes). Each object in the database stores a signature of the interaction with a motion driver that we manually assigned to each class. We select an unknown query object and perform an interaction with the same motion driver. Finally, we determine the precision and recall of shapes in the database by searching for the closest matches using the relative distance (Sec. 4.3 and Fig. 13) and by counting true and false positives until we recall all shapes from a given class. Fig. 16 shows the result for some of the objects used for this experiment. We query our dataset with an input shape (blue, right) to it we see the returned shapes in prioritized order (from left to right). Except for the failure case of the wineglass, our method returns semantically similar shapes (an object is a chair because you use it as a chair). Unlike the current approaches for describing static object relationships, capturing interactions as continuous motions implicitly reveals the affordance of the observed shapes in a more principled way. Although we rely on the specific arrangement of object and motion driver, this experiment emphasizes the potential of our method, as we are able to retrieve shapes that share similarities in their form and function.



Fig. 16. Visual comparison of models retrieved by our descriptor. The query object is shown in blue, the closest models according to our descriptor to the right; the retrieved objects share similar properties in descending order. The motion drivers used for this experiment were: sitting animation (chairs), fluid simulation (bicycles), grasping simulation (bottles, bags).




Fig. 17. A comparison of surface saliency estimations of our approach (top) and Lee et al. [2005] (bottom). Salient regions are highlighted as color overlay, red indicates a high response to the performed interaction, blue identifies uninvolved areas. The motion drivers for this experiment were: side grasp (flask), wind (car), top grasp (basket), animation (chair), wind (plane).

Saliency Estimation. Shapes are often synthesized without context and without knowledge about the semantic identity of their parts, making predictions on their functional use a difficult problem in structure-aware modeling. Previous methods proposed finding salient features of shapes through analyzing their geometric representation [17, 44], however, they only operate on shapes in isolation. Kim et al. [25] proposed to predict salient features by finding contact points of human agents interacting with shapes; their method explicitly focuses on human-centric saliency estimation.

We track interactions as continuous events in the interaction space of two shapes. We not only know where the motion driver is in direct contact with the observed shape, but also how it approaches its surface, or even how it initiates the interaction. We estimate salient regions by mapping the motion signature data to the geometry of a model. We find the sensor region for each shape vertex and its corresponding vector field. The saliency is measured for each vertex of the observed shape by computing the weighted sum of vector field attributes of the cell that contains the vertex. The weight for each attribute was found through experiments (Section 4.3). Additionally, we define a spherical region around the center of the current cell and include all adjacent cells with a distance smaller than the specified radius in the computation. We assign the normalized weighted sum to each vertex as the saliency measure. Fig. 17 (top) shows a

visualization of salient regions in direct comparison with the surface curvature based approach of Lee et al. [27] (bottom). Although we do not conclude that our saliency estimation provides better results, we would like to emphasize that an interaction-based estimation is a more general means that provides meaningful hints towards high level saliency detection. Compared to existing methods that solely rely on analyzing the geometry for finding salient regions, exploiting interaction information provides more details of the observed object to the cost of a more complex setup.

Shape Correspondence. Matching correspondences of semantically similar parts remains a challenging problem in many areas, such as shape completion, and segmentation, similarity assessment, or structure-aware modeling. To test our descriptor on this task, we find salient regions of shapes of the same category and determine the sum of their differences. Vertices of participating meshes are inserted into a uniform grid and the saliency value assigned to each vertex is used for the computation. Regions with the minimal sum of these differences can be defined as corresponding regions. This is similar to what Kim et al. [25] define for human-centric interactions. We can identify multiple interactions on the same object (e.g., grasping a cup from above or from the side). By finding the best match between these two kind of interaction signatures, our method allows us to detect corresponding regions of geometrically different shapes. Fig. 18 illustrates the results between a set of cups. The bottom row shows the result of an interaction transfer experiment. Here we capture a sitting interaction (animation) for the three chairs (a), (b), and (d). The results of (a) is transferred to the other chairs (c), (e).

Interaction Prediction. Finally, we show that exploiting time variations of interactions enables new types of applications that cannot be supported by static interaction descriptors. Our method can be used to predict the type of an interaction when only the first few seconds are known. To perform this experiment, we captured interactions and divided the motion trajectories into $N = 8$ time spans, for which we independently compute a signature. The k 'th signature of an interaction is computed based on the trajectories captured between the beginning t_{begin} and $t_{begin} + \frac{k}{N}d$, where d is the duration of the interaction. We found that computing the signatures from subsequent parts of the interaction is more reliable than comparing the individual time spans. The set of signatures for each interaction is stored in a database. We used 23 interactions of 3 categories (sitting down, walking, grasping) to produce a total of 184 signatures. To perform the prediction of an unseen interaction we progressively calculate the signature of the already available part of an interaction and compare it to the signatures in the database that cover the same number of time spans.

The performance of predicting the type of an interaction is visualized in Fig. 19. On the left we show a walking motion with color-coded trajectories. Each color represents a temporal sequence of the interaction. The PR diagram on the right shows how the performance gradually increases the more time spans become available. The graphs show the average precision and recall values for all interactions used in this experiment. With only one segment available we achieve a precision of 60% with a recall of 50%. The color in the PR diagram matches the color of the corresponding sequence up to where we compute the signature for the interaction.




Fig. 18. An example of matching correspondences with our descriptor. We can apply the same interaction to different shapes and find their corresponding regions if saliency is shared among the shapes. The top row shows three cups with color coded corresponding regions, handle (grasp from the side) and rim (grasp from top). The bottom row shows the result of an interaction transfer experiment. We capture a sitting interaction for the three chairs (a), (b), and (d). The captured interaction of (a) is transferred to the other chairs (c), (e).




Fig. 19. Interaction prediction: we divide interactions into equidistant time spans (left, start: blue, end: red) and store the signature for each part in a database. We predict an interaction by progressively comparing sequences of an unseen interaction to the previously encoded segments in the database. Right: a PR diagram that visualizes the quality of the interaction prediction. More sequences allow to more reliably predict an interaction. The color of the curves in the PR diagram matches the color of the sequence up to where we compare the interaction.

Current descriptors for interactions provide a means for analyzing and predicting the functionality of objects, but do not allow to predict interactions as dynamic motions. Unlike the previous approaches, our descriptor analyzes the motion flow and allows to reliably predict the type of interaction. Predicting interactions can be useful for the motion planning of autonomous agents and for objects in smart environments.

6.4 Discussion and Limitations

The main limitation of our approach is the absence of more meaningful interaction data. Predefined animations do not provide the same complexity in the motions found in the real-world or simulated data sets. Moreover, some of our current observations rely on small groups of predefined data resulting in a sparse representation of the captured interactions. It would be interesting to overcome this limitation and to test our method with more data from RGB-D video sequences of real-world interactions.

Unlike IBS [53] and ICON [22], we do not provide a hierarchical aggregation of interactions, but only consider the entire interaction landscape of the observed object, which could yield less meaningful signatures for complex interaction of many objects. We explicitly limit ourselves to a single object interacting with one or more motion drivers and would like to address more complex interaction concepts as future work. Moreover, our method is constrained by the need to place motion particles on moving surfaces. While this does not pose a problem for Lagrangian simulations, we rely on an expensive sampling of surfaces to find motion particle positions for solid objects. Moreover, only sampling a surface causes imprecision in capturing the motion trajectories and the corresponding interaction. However, capturing interactions through tracking complex geometric surfaces or even volumes does not yet seem feasible.

Instead of employing vector fields to quantize motions, potentially more refined representation exist. For example, directly defining and exploiting properties on the motion trajectories could be an interesting endeavor. However, to our knowledge, no method exists to readily explore such attributes of time-variant trajectories as geometric entities. Just finding correspondences or common parts within a single trajectory is a difficult problem that cannot be solved easily and at the required scale. Compared to more lightweight approaches that directly operate on sensor data, employing vector fields is less prone to artifacts and noise in the input data. The vector fields regularize the signal of the motion flow and can be efficiently analyzed. Our descriptor is based on six attributes and allows us to encode finer nuances of the motions compared to, *e.g.*, capturing the properties of normals of propagating surfaces as discussed in Oreifej and Liu [35].

7 CONCLUSION AND FUTURE WORK

We have presented a novel general-purpose descriptor that exploits time-variant properties of interactions observed in 3-dimensional spaces. We define interaction landscapes as the interaction space of two objects participating in an interaction. We distribute sensors in this interaction space and place motion particles on moving surfaces to track their movements. Our framework allows capturing interactions from a variety of sources, including simulations, animations, motion capture, and RGB-D scans. By encoding the resulting motion trajectories of the tracked particles, we can acquire information that supports a more complete understanding of shapes and their function. Moreover, we have shown that the signatures of observed interactions can be used to classify shapes and to find their salient regions.

Capturing the interaction of objects is essential in order to learn more about shapes and their function. Moreover, we believe that a general purpose descriptor provides a valuable means for exploring and classifying interactions with many potential avenues of future work. In this work we specifically limited ourselves to encoding static interaction landscapes of one motion driver and an observed object. However, in future work observing objects in dynamically changing interaction spaces seems like an interesting endeavor. Furthermore, the interactions of entire groups of entities, *e.g.*, as can be found in crowds or swarms, could be explored.

An interesting avenue for future work would be to explore our descriptor for other, non-proximal, interaction types as many interactions do not happen through object-object contacts, but remotely. This spans from physiological phenomena, such as capturing views or non-verbal human behavior to non-contact physical interactions, such as the transport of light or the effect of magnetism. Our method provides a general means for encoding interactions applicable to many domains. Finally, we would like to explore more elaborate methods to better encode finer nuances in the captured motion flows and to thereby enable our descriptor for more evolved applications in shape analysis, such as functional correspondence, shape synthesis, and symmetry detection.

ACKNOWLEDGMENTS

We would like to thank Mirela Ben-Chen for her insightful comments and suggestions on methods for 3D vector fields, Boris Thibert for discussing equations and proofreading the manuscript, and Torsten Hädrich for his implementation of the SPH-based fluid solver. The work was supported by the NSF grants CCF-1514305, IIS-1528025, EEC-1606396, NSG-1161480, CDS&E 1608762, EEC-1606396, ONR MURI N00014-13-1-0341, the Stanford AI Lab-Toyota Center for Artificial Intelligence Research, a Google Focused Research Award, the Max Planck Center for Visual Computing and Communications, the JSPS Strategic Young Researchers Visits Program for Acceleration Brain Circulations, and the National Science Foundation of China (61373071).

REFERENCES

- [1] Adobe. 2015. Mixamo. <https://www.mixamo.com/>. (2015).
- [2] Rami Ali Al-Asqhar, Taku Komura, and Myung Geol Choi. 2013. Relationship Descriptors for Interactive Motion Adaptation. In *Proc. of ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '13)*. ACM, 45–53. DOI : <https://doi.org/10.1145/2485895.2485905>
- [3] Nina Amenta, Marshall Bern, and Manolis Kamvyselis. 1998. A New Voronoi-based Surface Reconstruction Algorithm. In *Proc. of SIGGRAPH*. ACM, New York, NY, USA, 415–421. DOI : <https://doi.org/10.1145/280814.280947>
- [4] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. 2005. SCAPE: Shape Completion and Animation of People. In *Proc. of SIGGRAPH*. Los Angeles, California.
- [5] F.G. Ashby. 1992. *Multidimensional Models of Perception and Cognition*. L. Erlbaum. <https://books.google.com/books?id=eeQh4ESeOfsC>
- [6] Ezer Bar-Aviv and Ehud Rivlin. 2006. Functional 3D Object Classification Using Simulation of Embodied Agent.. In *BMVC*. British Machine Vision Association, 307–316.
- [7] Ilya Baran and Jovan Popović. 2007. Automatic Rigging and Animation of 3D Characters. *ACM Trans. Graph.* 26, 3 (jul 2007), 72:1–72:8.
- [8] S. Biasotti, A. Cerri, A. Bronstein, and M. Bronstein. 2015. Recent Trends, Applications, and Perspectives in 3D Shape Similarity Assessment. *Comp. Graph. Forum* (2015), n/a–n/a. DOI : <https://doi.org/10.1111/cgf.12734>
- [9] Alan W Black and Paul Taylor. 1997. Automatically Clustering Similar Units For Unit Selection In Speech Synthesis.. In *in Eurospeech97*. 601–604.
- [10] M. Caine. 1994. The design of shape interactions using motion constraints. In *Robotics and Automation, 1994. Proceedings., 1994 IEEE International Conference on*. 366–371 vol.1. DOI : <https://doi.org/10.1109/ROBOT.1994.351268>
- [11] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Sava, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. 2015. ShapeNet: An Information-Rich 3D Model Repository. *ArXiv e-prints* (Dec. 2015). [arXiv:cs.GR/1512.03012](https://arxiv.org/abs/1512.03012)
- [12] Yu-Wei Chao, Zhan Wang, Yugeng He, Jiaxuan Wang, and Jia Deng. 2015. HICO: A Benchmark for Recognizing Human-Object Interactions in Images. *ICCV* (2015).
- [13] Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. 2003. On Visual Similarity Based 3D Model Retrieval. *Comp. Graph. Forum* 22, 3 (2003), 223–232. DOI : <https://doi.org/10.1111/1467-8659.00669>
- [14] Jiating Chen, Xiaoyin Ge, Li-Yi Wei, Bin Wang, Yusu Wang, Huamin Wang, Yun Fei, Kang-Lai Qian, Jun-Hai Yong, and Wenping Wang. 2013. Bilateral Blue Noise

- Sampling. *ACM Trans. Graph.* 32, 6, Article 216 (Nov. 2013), 11 pages. DOI: <https://doi.org/10.1145/2508363.2508375>
- [15] Stanley Durrleman. 2010. *Statistical models of currents for measuring the variability of anatomical curves, surfaces and their evolution*. Ph.D. Dissertation. Université Nice - Sophia Antipolis, France.
- [16] Matthew Fisher, Manolis Savva, Yangyan Li, Pat Hanrahan, and Matthias Niessner. 2015. Activity-centric Scene Synthesis for Functional 3D Scene Modeling. *ACM Trans. Graph.* 34, 6, Article 179 (Oct. 2015), 13 pages. DOI: <https://doi.org/10.1145/2816795.2818057>
- [17] Ran Gal and Daniel Cohen-Or. 2006. Salient Geometric Features for Partial Shape Matching and Similarity. *ACM Trans. Graph.* 25, 1 (2006), 130–150. DOI: <https://doi.org/10.1145/1122501.1122507>
- [18] H. Grabner, J. Gall, and L. Van Gool. 2011. What makes a chair?. In *CVPR*. 1529–1536. DOI: <https://doi.org/10.1109/CVPR.2011.5995327>
- [19] Abhinav Gupta, Aniruddha Kembhavi, and Larry S. Davis. 2009. Observing Human-Object Interactions: Using Spatial and Functional Compatibility for Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 31, 10 (Oct. 2009), 1775–1789. DOI: <https://doi.org/10.1109/TPAMI.2009.83>
- [20] Edmond S. L. Ho, Taku Komura, and Chiew-Lan Tai. 2010. Spatial Relationship Preserving Character Motion Adaptation. *ACM Trans. Graph.* 29, 4, Article 33 (2010), 8 pages. DOI: <https://doi.org/10.1145/1778765.1778770>
- [21] Ruizhen Hu, Oliver van Kaick, Bojian Wu, Hui Huang, Ariel Shamir, and Hao Zhang. 2016. Learning How Objects Function via Co-analysis of Interactions. *ACM Trans. Graph.* 35, 4, Article 47 (July 2016), 13 pages. DOI: <https://doi.org/10.1145/2897824.2925870>
- [22] Ruizhen Hu, Chenyang Zhu, Oliver van Kaick, Ligang Liu, Ariel Shamir, and Hao Zhang. 2015. Interaction Context (ICON): Towards a Geometric Functionality Descriptor. *ACM Trans. Graph.* 34, 4, Article 83 (2015), 12 pages. DOI: <https://doi.org/10.1145/2766914>
- [23] Hao Jiang and David R. Martin. 2008. Finding Actions Using Shape Flows. In *Computer Vision – ECCV 2008*, David Forsyth, Philip Torr, and Andrew Zisserman (Eds.). Lecture Notes in Computer Science, Vol. 5303. Springer Berlin Heidelberg, 278–292. DOI: https://doi.org/10.1007/978-3-540-88688-4_21
- [24] Germany Karlsruhe Institute of Technology, Karlsruhe. 2015. KIT Whole Human Motion Database. <https://motion-database.humanoids.kit.edu/>. (2015).
- [25] Vladimir G. Kim, Siddhartha Chaudhuri, Leonidas Guibas, and Thomas Funkhouser. 2014. Shape2Pose: human-centric shape analysis. *ACM Trans. Graph.* 33, 4, Article 120 (2014), 12 pages. DOI: <https://doi.org/10.1145/2601097.2601117>
- [26] Hamid Laga, Michela Mortara, and Michela Spagnuolo. 2013. Geometry and Context for Semantic Correspondences and Functionality Recognition in Man-made 3D Shapes. *ACM Trans. Graph.* 32, 5, Article 150 (2013), 16 pages. DOI: <https://doi.org/10.1145/2516971.2516975>
- [27] Chang Ha Lee, Amitabh Varshney, and David W. Jacobs. 2005. Mesh Saliency. *ACM Trans. Graph.* 24, 3 (July 2005), 659–666. DOI: <https://doi.org/10.1145/1073204.1073244>
- [28] M. Leordeanu and M. Hebert. 2005. A spectral technique for correspondence problems using pairwise constraints. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, Vol. 2. 1482–1489 Vol. 2. DOI: <https://doi.org/10.1109/ICCV.2005.20>
- [29] Pan Li, Bin Wang, Feng Sun, Xiaohu Guo, Caiming Zhang, and Wenping Wang. 2015. Q-MAT: Computing Medial Axis Transform By Quadratic Error Minimization. *ACM Trans. Graph.* 35, 1, Article 8 (Dec. 2015), 16 pages. DOI: <https://doi.org/10.1145/2753755>
- [30] Y. Li, J. L. Fu, and N. S. Pollard. 2007. Data-Driven Grasp Synthesis Using Shape Matching and Task-Based Pruning. *IEEE Transactions on Visualization and Computer Graphics* 13, 4 (July 2007), 732–747. DOI: <https://doi.org/10.1109/TVCG.2007.1033>
- [31] Zhenbao Liu, Caili Xie, Shuhui Bu, Xiao Wang, Junwei Han, Hongwei Lin, and Hao Zhang. 2015. Indirect shape analysis for 3D shape retrieval. *Computers & Graphics* 46 (2015), 110–116. DOI: <https://doi.org/10.1016/j.cag.2014.09.038> Shape Modeling International 2014.
- [32] Niloy Mitra, Michael Wand, Hao (Richard) Zhang, Daniel Cohen-Or, Vladimir Kim, and Qi-Xing Huang. 2013. Structure-aware Shape Processing. In *SIGGRAPH Asia 2013 Courses (SA '13)*. ACM, New York, NY, USA, Article 1, 20 pages. DOI: <https://doi.org/10.1145/2542266.2542267>
- [33] Niloy J. Mitra, Mark Pauly, Michael Wand, and Duygu Ceylan. 2013. Symmetry in 3D Geometry: Extraction and Applications. *Comput. Graph. Forum* 32, 6 (2013), 1–23.
- [34] Joe J Monaghan. 1992. Smoothed particle hydrodynamics. *Annual review of astronomy and astrophysics* 30 (1992), 543–574.
- [35] O. Oreifej and Zicheng Liu. 2013. HON4D: Histogram of Oriented 4D Normals for Activity Recognition from Depth Sequences. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*. 716–723. DOI: <https://doi.org/10.1109/CVPR.2013.98>
- [36] Robert Osada, Thomas Funkhouser, Bernard Chazelle, and David Dobkin. 2002. Shape Distributions. *ACM Trans. Graph.* 21, 4 (Oct. 2002), 807–832. DOI: <https://doi.org/10.1145/571647.571648>
- [37] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. 2012. Functional Maps: A Flexible Representation of Maps Between Shapes. *ACM Trans. Graph.* 31, 4 (2012).
- [38] Michael Pechuk, Octavian Soldea, and Ehud Rivlin. 2008. Learning Function-based Object Classification from 3D Imagery. *Comput. Vis. Image Underst.* 110, 2 (May 2008), 173–191. DOI: <https://doi.org/10.1016/j.cviu.2007.06.002>
- [39] E. Rivlin, S.J. Dickinson, and Azriel Rosenfeld. 1994. Recognition by functional parts [function-based object recognition]. In *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR '94., 1994 IEEE Computer Society Conference on*. 267–274. DOI: <https://doi.org/10.1109/CVPR.1994.323839>
- [40] Columbia University USA Robotics Lab, Computer Science Laboratory. 2012. GraspIt! <http://www.cs.columbia.edu/~cmatei/graspit/>. (2012).
- [41] Y. Rubner, C. Tomasi, and L.J. Guibas. 1998. A metric for distributions with applications to image databases. In *Computer Vision, 1998. Sixth International Conference on*. 59–66. DOI: <https://doi.org/10.1109/ICCV.1998.710701>
- [42] Manolis Savva, Angel X. Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. 2014. SceneGrok: Inferring Action Maps in 3D Environments. *ACM Trans. Graph.* 33, 6, Article 212 (Nov. 2014), 10 pages. DOI: <https://doi.org/10.1145/2661229.2661230>
- [43] Manolis Savva, Angel X. Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. 2016. PiGraphs: Learning Interaction Snapshots from Observations. *ACM Trans. Graph.* 35, 4, Article 139 (July 2016), 12 pages. DOI: <https://doi.org/10.1145/2897824.2925867>
- [44] Philip Shilane and Thomas Funkhouser. 2007. Distinctive Regions of 3D Surfaces. *ACM Trans. Graph.* 26, 2, Article 7 (June 2007). DOI: <https://doi.org/10.1145/1243980.1243981>
- [45] Oana Sidi, Oliver van Kaick, Yanir Kleiman, Hao Zhang, and Daniel Cohen-Or. 2011. Unsupervised Co-segmentation of a Set of Shapes via Descriptor-space Spectral Clustering. *ACM Trans. Graph.* 30, 6, Article 126 (2011), 10 pages. DOI: <https://doi.org/10.1145/2070781.2024160>
- [46] Hyun Oh Song, M. Fritz, Chunhui Gu, and T. Darrell. 2011. Visual grasp affordances from appearance-based cues. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*. 998–1005. DOI: <https://doi.org/10.1109/ICCVW.2011.6130360>
- [47] Melanie Sutton, Louise Stark, and Kevin Bowyer. 1994. GRUFF-3: Generalizing the domain of a function-based recognition system. *Pattern Recognition* 27, 12 (1994), 1743 – 1766. DOI: [https://doi.org/10.1016/0031-3203\(94\)90091-4](https://doi.org/10.1016/0031-3203(94)90091-4)
- [48] Art Tevs, Qixing Huang, Michael Wand, Hans-Peter Seidel, and Leonidas Guibas. 2014. Relating Shapes via Geometric Symmetries and Regularities. *ACM Trans. Graph.* 33, 4, Article 119 (July 2014), 12 pages. DOI: <https://doi.org/10.1145/2601097.2601220>
- [49] Amos Tversky. 1977. Features of similarity. *Psychological Review* 84 (1977), 327–352.
- [50] Dimitrios Tzionas and Juergen Gall. 2015. 3D Object Reconstruction from Hand-Object Interactions. In *International Conference on Computer Vision (ICCV)*. <http://files.is.tue.mpg.de/dtzionas/In-Hand-Scanning>
- [51] Zhirong Wu, Shuran Song, Aditya Khosla, Xiaoou Tang, and Jianxiong Xiao. 2014. 3D ShapeNets for 2.5D Object Recognition and Next-Best-View Prediction. *CoRR abs/1406.5670* (2014). <http://arxiv.org/abs/1406.5670>
- [52] Liefei Xu, H. Quynh Dinh, Philippos Mordohai, and Thomas Ramesay. 2011. Detecting Patterns in Vector Fields. *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition* 6 (2011).
- [53] Xi Zhao, He Wang, and Taku Komura. 2014. Indexing 3D Scenes Using the Interaction Bisector Surface. *ACM Trans. Graph.* 33, 3, Article 22 (2014), 14 pages. DOI: <https://doi.org/10.1145/2574860>
- [54] Youyi Zheng, Daniel Cohen-Or, and Niloy J. Mitra. 2013. Smart Variations: Functional Substructures for Part Compatibility. *Computer Graphics Forum* 32, 2pt2 (2013), 195–204. DOI: <https://doi.org/10.1111/cgf.12039>
- [55] Yuke Zhu, Alireza Fathi, and Li Fei-Fei. 2014. Reasoning about Object Affordances in a Knowledge Base Representation. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Lecture Notes in Computer Science, Vol. 8690. Springer International Publishing, 408–424. DOI: https://doi.org/10.1007/978-3-319-10605-2_27

A FIRST-ORDER VECTOR FIELD ATTRIBUTES

We analyze the captured motion flows by computing a set of first-order attributes on a vector field. As shown in Xu et al. [52] the gradient of a 3D vector field v_i is an asymmetric 3D tensor field that can be decomposed into symmetric S and asymmetric A components. While the diagonal and off-diagonal entries of S represent dilatation and shearing respectively, the off-diagonal entries of A measure vorticity. Given the tensor $T(p)$ of a point $p \in v_i$, the decomposition is defined as

$$T(p) = \begin{pmatrix} T_{11} & T_{12} & T_{13} \\ T_{21} & T_{22} & T_{23} \\ T_{31} & T_{32} & T_{33} \end{pmatrix} = \begin{pmatrix} \frac{\partial V_x}{\partial x} & \frac{\partial V_x}{\partial y} & \frac{\partial V_x}{\partial z} \\ \frac{\partial V_y}{\partial x} & \frac{\partial V_y}{\partial y} & \frac{\partial V_y}{\partial z} \\ \frac{\partial V_z}{\partial x} & \frac{\partial V_z}{\partial y} & \frac{\partial V_z}{\partial z} \end{pmatrix} = A + S,$$

with S and A given by

$$S = \begin{pmatrix} \epsilon_1 & \frac{1}{2}\theta_3 & \frac{1}{2}\theta_2 \\ \frac{1}{2}\theta_3 & \epsilon_2 & \frac{1}{2}\theta_1 \\ \frac{1}{2}\theta_2 & \frac{1}{2}\theta_1 & \epsilon_3 \end{pmatrix}, \quad A = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix},$$

where ϵ_i and θ_i represent the diagonal and off-diagonal components of S :

$$\epsilon_1 = \frac{\partial V_x}{\partial x}, \quad \epsilon_2 = \frac{\partial V_y}{\partial y}, \quad \epsilon_3 = \frac{\partial V_z}{\partial z},$$

$$\theta_1 = \frac{\partial V_z}{\partial y} + \frac{\partial V_y}{\partial z}, \quad \theta_2 = \frac{\partial V_x}{\partial z} + \frac{\partial V_z}{\partial x}, \quad \theta_3 = \frac{\partial V_y}{\partial x} + \frac{\partial V_x}{\partial y},$$

and ω_i the off-diagonal entries of A :

$$\omega_1 = \frac{1}{2} \left(\frac{\partial V_z}{\partial y} - \frac{\partial V_y}{\partial z} \right), \quad \omega_2 = \frac{1}{2} \left(\frac{\partial V_x}{\partial z} - \frac{\partial V_z}{\partial x} \right), \quad \omega_3 = \frac{1}{2} \left(\frac{\partial V_y}{\partial x} - \frac{\partial V_x}{\partial y} \right).$$

Based on this formulation, the attributes tensor magnitude M_t , dilatation magnitude M_d , magnitude of shear strain rate M_s and vorticity magnitude $M_{\vec{\omega}}$ are given by

$$M_t = \sqrt{\frac{1}{2} \sum T_{i,j}^2}, \quad M_d = \sqrt{\sum \epsilon_i^2},$$

$$M_s = \sqrt{\sum \theta_i^2}, \quad M_{\vec{\omega}} = \sqrt{\sum \omega_i^2},$$

with $i, j = 1, 2, 3$. Additionally, we derive the attributes vector orientation O , and vector magnitude M directly from the vector field. In contrast to the approach of Xu et al. [52], we do not consider the attribute pressure P , as it cannot easily be acquired for all motion drivers in our framework.

Received September 2016; revised January 2017; final version March 2017; accepted March 2017