

SALIENCY OF SUBTLE ENTITIES WITHIN 3-D POINT CLOUDS

R. Arav,* S. Filin

Dept. of Mapping and Geoinformation, Technion - Israel Institute of Technology, 32000, Haifa, Israel -
(reumaa, filin)@technion.ac.il

Commission II, WG II/3

KEY WORDS: Saliency; distinctness; 3-D point-clouds; subtle entities; laser scans

ABSTRACT:

Visual saliency is defined by regions of the scene that stand out from their neighbors and attract immediate attention. In image processing, visual saliency is frequently used to focus local analysis of key features. Though their advantage is largely acknowledged, little research has been carried concerning 3-D data, and even less in relation to data acquired by laser scanners for mapping. In this paper, we propose a new saliency measure for laser scanned point-clouds, governed by the neurological concepts of center-surround and low-level features. Adjusted to large point sets, we propose a fast geometric descriptor, which quantifies the distance of a point from its surrounding. We show that the proposed model highlights not only salient details in watertight models, but also in airborne and terrestrially scanned scenes that may hold subtle entities embedded within the topography. The detection of such regions paves the way to a myriad of applications, such as feature and pattern extraction, registration, classification, viewpoint selection, point-cloud simplification, landmark detection, etc.

1. INTRODUCTION

Visual saliency is defined by regions of the scene that stand out from their neighbors and attract immediate attention. Their detection is considered a key attentional mechanism that facilitates learning and survival by enabling organisms to focus their limited perceptual and cognitive resources on the most pertinent subset of the available sensory data (Frintrop et al., 2010). In image processing, salient regions act as preliminary cues in various applications, such as shape matching, object recognition, similarity estimation, registration, down-sampling, and visualization (Achanta et al., 2009; Li et al., 2013; Yuan et al., 2018; Li et al., 2019). Within 3-D point-clouds, visual attention can be harnessed to reduce the problem of scene understanding into rapid series of less demanding computational procedures, aimed to support localized visual analysis problems, e.g., detection, simplification, registration and others. However, the varying resolution, occlusions and the absence of topological information, make the estimation of saliency in this domain a challenge.

A number of biologically plausible models have been developed to explain the cognitive visual process of humans and animals, with two main principles at its foundation. The first, a neuroscience based – follows the observation that neurons in the retina are sensitive to regions which locally stand out from their surroundings. The second, an influence principle based – emerges from the joint impact of both goals (top-down influences) that animals have, and stimuli (bottom-up influences) that affect them (Li et al., 2013; Yuan et al., 2018). In computer vision, most saliency detection methods utilize those two principles in order to detect regions standing out. Itti et al. (1998) were first to introduce the center-surround concept combined with low-level features (e.g., color, contrast, orientation, or size). To do so, multi-scale image features were combined into a single topographical saliency map. van de Weijer et al. (2006) focused on color distinctness by estimating the probab-

ility of a color feature vector composed of the RGB channels and their derivatives to occur. Achanta et al. (2008) extended the center-surround rational by measuring the distance between pixel values in sub-regions within the image and in a later work, presented a frequency-tuned approach, where similar distance measure was used between smoothed images (Achanta et al., 2009). Li et al. (2013) and Wu et al. (2013) argued that global contrast should also be considered when one region is similar to its surrounds but still distinct in the whole scene. In the three-dimensional realm, saliency has been treated mostly in 3-D polygonal meshes, and defined as regions that are perceptually important. As an example, Lee et al. (2005) extended the model proposed by Itti et al. (1998) to meshes, but instead of considering color, intensity, and orientation of regions, the authors utilized the surface mean curvature as the most important attribute. Wu et al. (2013) incorporated global rarity into the model by evaluating the saliency of a cluster of vertices in meshes.

Little research has been carried out concerning point-cloud data (Shtrom et al., 2013; Wang et al., 2015; Kobyshev et al., 2016; Guo et al., 2018; Ding et al., 2019). In general, point-clouds do not hold any topological information on which most mesh-based methods rely on in the evaluation of saliency, and they suffer from inherited data acquisition characteristics such as noise, occlusions and varying point densities across the data. Most approaches proposed to use the directional changes between a point and its neighbors, described by fast point feature histogram (FPFH; Rusu et al., 2009) as the main contributor to saliency. This is achieved by a hierarchical model that measures the similarity and dissimilarity by the distances between the points' FPFH (Shtrom et al., 2013; Tasse et al., 2015; Kobyshev et al., 2016; Ding et al., 2019). While Shtrom et al. (2013) and Kobyshev et al. (2016) proposed to measure the dissimilarity between each point in the cloud, Tasse et al. (2015) accelerated the process by considering dissimilarities between clusters of points. Only recently, Ding et al. (2019) combined the cluster- and point-level saliencies proposed in

* Corresponding author

Tasse et al. (2015) and Shtrom et al. (2013), respectively, to accommodate for the global rarity. This also improved the robustness of the approach to noise, and has led to better focused results. Wang et al. (2015) used a top-down influences (via goals) and relied heavily on the fact that within a mobile scan of a road and its surrounding inventories, salient features are geometrically different from the main road. Guo et al. (2018) defined the point descriptor based on principal component analysis (PCA). The descriptor was composed of sigma-sets, extracted from the covariance matrix of each point's normal and curvature. Nonetheless, with descriptor-based estimation the computation has exponential runtime and memory complexity in terms of the quantity of the low-level features. Notably, most methods were tested on watertight 3-D point-clouds, which were acquired by table scanners in a controlled environment (e.g., Tasse et al., 2015; Guo et al., 2018; Ding et al., 2019).

This paper studies the detection of salient regions in open environments, scanned by airborne and terrestrial laser scanners. Such regions, conspicuous within their surroundings, will facilitate focused detection of embedded entities within natural environs. There, entities may be small in size, subtle in appearance and may wear a variety of forms. In such cases, the object-to-background transition is usually smooth, while surface roughness and measurement noise may obscure their transition. Therefore, distinctness is neither expressed in the uniqueness of the normal, as in urban environments (Shtrom et al., 2013; Kobyshev et al., 2016) nor in their difference from the dominant plane (Wang et al., 2015), but rather in the variation of the surface itself. To identify distinctness, we propose a new saliency measure which is adapted to open scenes and is governed by surface geometry, while maintaining the neurological concepts of center-surround and low-level features. Adjusted to large 3-D point-clouds, we aspire for a simple geometric descriptor, which enables to estimate the distinctness of each point from its surrounding, at low computational overhead, while being attuned to surface variations of relatively smooth scenes. We evaluate the proposed method against state-of-the-art approaches applied on models that are frequently used in such studies, and on complex airborne and terrestrially scanned point-clouds. The new saliency can be further integrated as preliminary step for local analysis of key entities for object extraction, classification, registration, smart down-sampling, etc.

2. METHODOLOGY

We define salient surface elements using geometrical properties that relate to perceptually dominant features. These may be bends within the topography, subtle elements inlaid within it, or distinctive topographic features (e.g., peaks, pits, ridges, valleys or saddle points). As the object-to-background transition is reflected by surface geometry, we focus on means to quantify the variations within it. Operating on a set of points, we develop discrete approximation for surface parameters and elaborate on the definition of points' neighborhood, with an analysis of its ensuing impact.

2.1 Point neighborhood

We consider two neighborhood classes: *i*) k -nearest neighbors based (k-nn) – where the set is composed of a predefined k number of the nearest neighbors; and *ii*) a radius based (r-nn) – where *all* points within a specific radius (or window) are considered as neighbors. The first ensures that the same number

of points composes the neighborhood, but incorporates no spatial or distribution related consideration. The latter values same sized regions as the defining criterion and is more likely to ensure symmetric distribution of neighbors.

When point density is known, it is customary to use the k-nn approach as it is likely to yield a computationally manageable number of neighbors (e.g., Weinmann et al., 2017). When densities vary between sets, size considerations become the preferable choice among the two categories.

For the extraction itself, speedup of the performance, which is key for implementation aspects, promotes compromise in accuracy of the results (e.g., Arya et al., 1998). Our experiments show that the approach implemented in the fast library for approximate nearest neighbors (FLANN; Muja, Lowe, 2009) yields the best performance among the other alternatives, even though our data is of lower dimension (three) compared to what the FLANN was intended to, and is the approach adopted here.

2.2 Surface normal computation

It was noted that surface normals are a prime feature to characterize surface variations (Shtrom et al., 2013; Tasse et al., 2015; Guo et al., 2018). A variety of methods has been proposed for their estimation, including fitting an implicit surface to the neighborhood and estimating its gradient; finding the tangent plane itself by analyzing the points' distribution; or using principal component analysis (Rusu et al., 2009; Sirmacek et al., 2016). In essence, all approaches are variants of the principal component analysis (PCA) of the point distribution, and this is the approach adopted here. Given a point \mathbf{q} and its k -neighbors, the covariance matrix is computed by:

$$\mathbf{C} = \frac{1}{k} \sum_{i=1}^k [(\mathbf{p}_i - \mathbf{q}) \cdot (\mathbf{p}_i - \mathbf{q})^T]. \quad (1)$$

and its principal components are given by the eigenvalues and vectors:

$$\mathbf{C}\mathbf{v}_j = \lambda_j \mathbf{v}_j, \quad j \in \{1, 2, 3\} \quad (2)$$

where \mathbf{v} , and λ are the eigenvector and values, respectively. The covariance matrix \mathbf{C} is positive semi-definite ($0 \leq \lambda_3 \leq \lambda_2 \leq \lambda_1$) and \mathbf{v}_3 , corresponding to λ_3 , is the approximation of \mathbf{n}_q , the surface normal. We take advantage of the fact that a closed form solution exists to the third degree characteristic polynomial, and instead of applying a general computation of the roots, we compute the eigenvalues as follows (Kopp, 2008):

$$\begin{aligned} \lambda_1 &= \mu + 2\nu \cdot \cos \theta \\ \lambda_2 &= \mu + 2\nu \cdot \cos(\theta + 120^\circ) \\ \lambda_3 &= 3\mu - \lambda_1 - \lambda_2 \end{aligned} \quad (3)$$

where:

$$\begin{aligned} \mu &= \frac{1}{3} \text{tr}(\mathbf{C}) \\ \nu &= \sqrt{\frac{\text{tr}[(\mu \cdot \mathbf{I} - \mathbf{C})^2]}{6}} \\ \theta &= \frac{1}{3} \cos^{-1} \left[\frac{1}{2} \det(\mathbf{C} - \mu \mathbf{I}) \right] \end{aligned} \quad (4)$$

and \mathbf{I} is a 3×3 identity matrix. As our interest is only in \mathbf{v}_3 , we compute it directly as the cross-product between two rows of $\mathbf{C} - \lambda_3 \mathbf{I}$, further reducing the runtime. As ambiguity in the

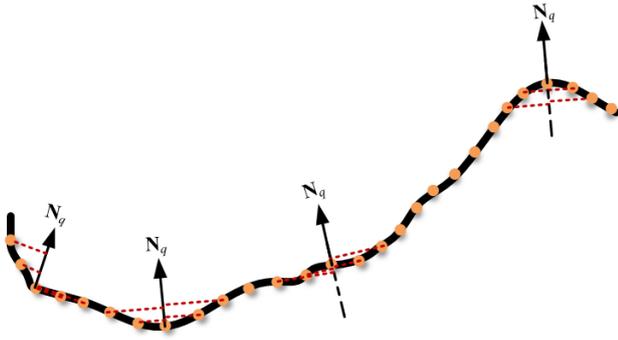


Figure 1. Convexity measure: the neighboring points are projected onto the center point normal direction \mathbf{n}_q .

normal sign can arise, we orient all normals toward a single viewpoint, v_p , such that:

$$\mathbf{n}_q = \begin{cases} \mathbf{n}_q, & \text{if } \mathbf{n}_q \cdot (v_p - \mathbf{q}) \geq 0 \\ -\mathbf{n}_q, & \text{otherwise.} \end{cases} \quad (5)$$

for all points.

2.3 Non-parametric surface curvature (convexity)

Surface normals may prove useful along building edges or other breaklines, elements that generate strong discontinuities. Aspiring for cues which are better attuned to entities that are embedded within the topography, we turn to surface curvature that measures variations in the normal direction.

Three different strategies for curvature estimation were evaluated and tested, including a numerical one, a parametric and a non-parametric, which we present here. Among the three, the non-parametric is the simplest, fastest and most importantly – provided the best characterization. The non-parametric approach quantifies the convexity/concavity of the surface in each point by examining the characteristics of the points' distribution around it. To do so, we sum the projections of neighboring points on \mathbf{n}_q , the normal direction to the center point (Fig. 1). Convexity is measured by:

$$C_q = \frac{\sum_{i=1}^k \mathbf{n}_q^T \cdot (\mathbf{p}_i - \mathbf{q})}{k} \quad (6)$$

where \mathbf{p}_i is the i -th neighbor of \mathbf{q} . In its bare form, effects of measurement noise and surface roughness would be documented within the convexity value, as each point contributes to the computation, leading to wrong estimations. Therefore, the projection must be dominant enough over the level of background noise and surface roughness. Assuming that the projections are part of a Normal distribution $N \sim (\varepsilon, \sigma_c)$, where ε is the surface roughness and σ_c the accuracy of the curvature estimation. In an attempt to quantify the effect of surface roughness and ranging accuracy, Baruch, Filin, 2011 derived an accuracy measure for the curvature based on the measurement ranging, m_0 :

$$\sigma_c = \pm m_0 \frac{\sqrt{6}}{d^2}. \quad (7)$$

and also the effect of surface roughness:

$$\varepsilon = \frac{2\Delta Z}{d^2} \quad (8)$$

where ΔZ the minimal detection level and d half the window size. Their focus was to test how surface characteristics and measurement noise affect curvature values. Using our approach, surface roughness is estimated directly from the variations in the surface itself, i.e., ΔZ , while the curvature accuracy is an agglomeration of the measurement and the normal accuracy. Nonetheless, as the curvature is evaluated along the vertical component of the normal, only the vertical accuracy should be considered. Therefore, the curvature accuracy is estimated in a much simpler manner, directly from the measurement accuracy, i.e., $m_c = m_0$. We establish an hypothesis test $H_0 : |C| \leq \varepsilon$ and $H_1 : |C| > \varepsilon$ as the alternative, to determine if a projection is noise, namely:

$$\left| \frac{C - \varepsilon}{\sigma_c} \right| \leq Z_{1-\frac{\alpha}{2}} \quad (\alpha = 5\%) \quad (9)$$

where $Z_{1-\alpha/2}$ the normalized Gaussian distribution and α the confidence level. Only projections that do not answer the criteria are used.

Non viable points – Estimation of the non-parametric values depends upon an approximately even distribution around the queried point. However, this estimation is compromised when the distribution is uneven, which mostly occur at the scan edges, or near large areas of occlusions. Such points are discarded or set with a zero value. Their detection can be readily accommodated by projecting the points to the plane orthogonal to the normal, using the projection matrix:

$$\mathbf{P}_q = \mathbf{I} - \mathbf{n}_q \mathbf{n}_q^T. \quad (10)$$

This facilitates the estimation of the center of gravity, and thereby its deviation from \mathbf{q} itself using simple mathematical operations that hardly affect the overhead.

2.4 Directional saliency

We expect a good model of saliency to highlight interesting changes within a laser scan of an open environment (be it natural or urban). Because the analyzed point-cloud is large, and the saliency itself is a preliminary stage to further processing, we seek the minimal information required to distinguish one area from another. Following Itti et al. (1998), we develop a center-surround operator, according to whom the distinctness of a point or a region is estimated by measuring the deviations of a wider *surround* from a narrower *center*.

We begin by defining zero-saliency where the surface orientation changes smoothly. There are two properties that control the saliency: the normal, which accommodates for the surface orientation, and the normal curvature. The center-surround operation for each property is defined by using a weighting function so that points within the immediate surrounding, as well as in distant areas, are given lower weights, while closer regions are given higher ones. Borrowing from the notion of a band-pass filtering, we construct a weighting function that suppresses the immediate and the distant surroundings, while encouraging a band at radius ρ (Fig. 2). For this we define a weighting function which is based on the Normal distribution whose center is at ρ , the distance which gets the highest weight, and the σ that defines the effective distance (i.e., the size of the surrounding):

$$\omega(x, y) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{\sqrt{x^2+y^2}-\rho}{\sigma} \right)^2}. \quad (11)$$

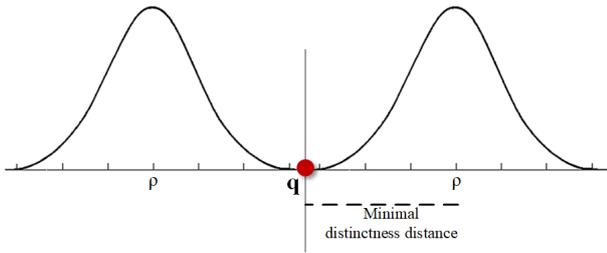


Figure 2. Weighting function for saliency detection: a center-surround filter that emphasizes regions at a distance ρ while suppressing closer and distant features, a 2-D view.

Using this function we estimate:

$$d\mathbf{n} = \iint_{\mathcal{S}} \|\mathbf{n}(x_0, y_0) - \mathbf{n}(x, y)\| \cdot \omega(x, y) dx dy \quad (12)$$

$$d\mathcal{C} = \iint_{\mathcal{S}} [\mathcal{C}(x_0, y_0) - \mathcal{C}(x, y)] \cdot \omega(x, y) dx dy \quad (13)$$

where (x_0, y_0) the planimetric coordinates of point \mathbf{q} . The proposed scheme follows the center-surround principal, by estimating the similarity of the point in reference to the normal and curvature to its surrounding.

Roughness and noise considerations – Rough texture areas and noise signals should not be marked, even though the surface changes both in curvature and normal. Therefore, within the effective distance, the change in either normal or curvature should be larger than an a priori ε , so that it satisfies:

$$d\tilde{\mathbf{n}} = \begin{cases} 0 & Var(d\mathbf{n}) > \varepsilon_{\mathbf{n}} \\ d\mathbf{n} & Var(d\mathbf{n}) \leq \varepsilon_{\mathbf{n}} \end{cases}, \quad (14)$$

$$d\tilde{\kappa} = \begin{cases} 0 & Var(d\kappa) > \varepsilon_{\kappa} \\ d\kappa & Var(d\kappa) \leq \varepsilon_{\kappa} \end{cases}.$$

Both evaluations are performed by establishing hypothesis tests. As the variance is being evaluated, the χ^2 -test is used.

Since the normal and curvature are two different, non-related, measures, we normalize them so that both elements have equal contribution to the saliency map. We considered different normalization functions \mathcal{N} , however, a simple linear one proved to be effective. The saliency measure is given by:

$$\zeta = \frac{1}{2} (\mathcal{N}(d\tilde{\mathbf{n}}) + \mathcal{N}(d\tilde{\kappa})) \quad (15)$$

Discrete implementation – When applying Eq. (12) to point-cloud data the saliency maps are estimated by using a discrete form:

$$d\mathbf{n}_i = \frac{1}{K} \sum_{j=1}^K \omega_{ij} \cdot (\mathbf{n}_i - \mathbf{n}_j) \quad (16)$$

$$d\kappa_{\mathbf{n}_j} = \frac{1}{K} \sum_{j=1}^K \omega_{ij} \cdot (\kappa_i - \kappa_j)$$

with K the neighborhood size. The weighting function, ω_{ij} , is approximated according to Eq. (11).

3. RESULTS AND ANALYSIS

Application of the proposed method is demonstrated on datasets representing a variety of scenarios and data spans. We compare the performance of our method to state-of-the-art approaches. Notably, in cases where the reference point-based saliency measure failed, due to the size of the data or to their complex nature, we compare our results to image-driven approaches. Evaluation is performed visually, as is the common practice in saliency related works, since the objective of visual saliency is to highlight conspicuous regions.

3.1 Watertight benchmark models

We demonstrate the application of the proposed saliency method on two of the common benchmark models. One is the Max Planck bust and the second is the Stanford Dragon. Three approaches are utilized here, Shtrom et al. (2013); Tasse et al. (2015) and Guo et al. (2018). We evaluate both quality and runtime. As these datasets are devoid of noise, and their point density is fixed, we use here a k-nn approach to define the neighborhoods.

Fig. (3) provides a comparative evaluation of the saliency methods applied to the benchmark models. It shows that our results concentrate indeed on the interesting regions, such as the facial area in the Max Planck bust, while FPFH-based approaches (Shtrom et al., 2013; Tasse et al., 2015) highlight wider areas. It also shows that in the dragon model our saliency produces less noise in the body than other approaches, and that it assigns high values at the edges of the model. Notably, runtime performance is better than state-of-the-art (Table 1), when considering also the programming language by which these measures were implemented.

	# points	Runtime (s)	Language
Shtrom et al. (2013)	437K	5400	Python
Tasse et al. (2015)	437K	250	C++
Guo et al. (2018)	437K	44	C++
Our method	437K	40	Python

Table 1. Runtime comparison of the proposed saliency estimation with state-of-the-art approaches on the dragon watertight model.

3.2 Airborne laser scan

Next, we demonstrate the application of the proposed method on an airborne laser scan of natural surrounding that features an alluvial fan ($31^{\circ}20'N$, $35^{\circ}25'E$; Fig. 4a). The data consists of gullies and collapse sinkholes (Abelson et al., 2003). The point-cloud spans an area of $480 \times 375 \text{ m}^2$ with point density of $\sim 8 \text{ points/m}^2$. Visually, we can identify two gullies and fifty-eight sinkholes as salient regions. The gullies are at width of 5 m and 9 m and depth of $\sim 2\text{-}6 \text{ m}$, dissecting the scan and forking towards the west; the sinkholes are at varying diameters and depths (4-20 m, 0.5-4.5 m, respectively) scattered along the scan, while some are formed within the gullies, or at their edge.

Comparing saliencies to point-based approaches proposed by Shtrom et al. (2013), Tasse et al. (2015) or Guo et al. (2018) was not possible here. The size of the data, which was too large, and its nature, that spans a wide region at lower resolution than of a terrestrial or table scanner, made them difficult to apply. To generate a comparative evaluation set, we examine the application of two image-based saliency approaches proposed in

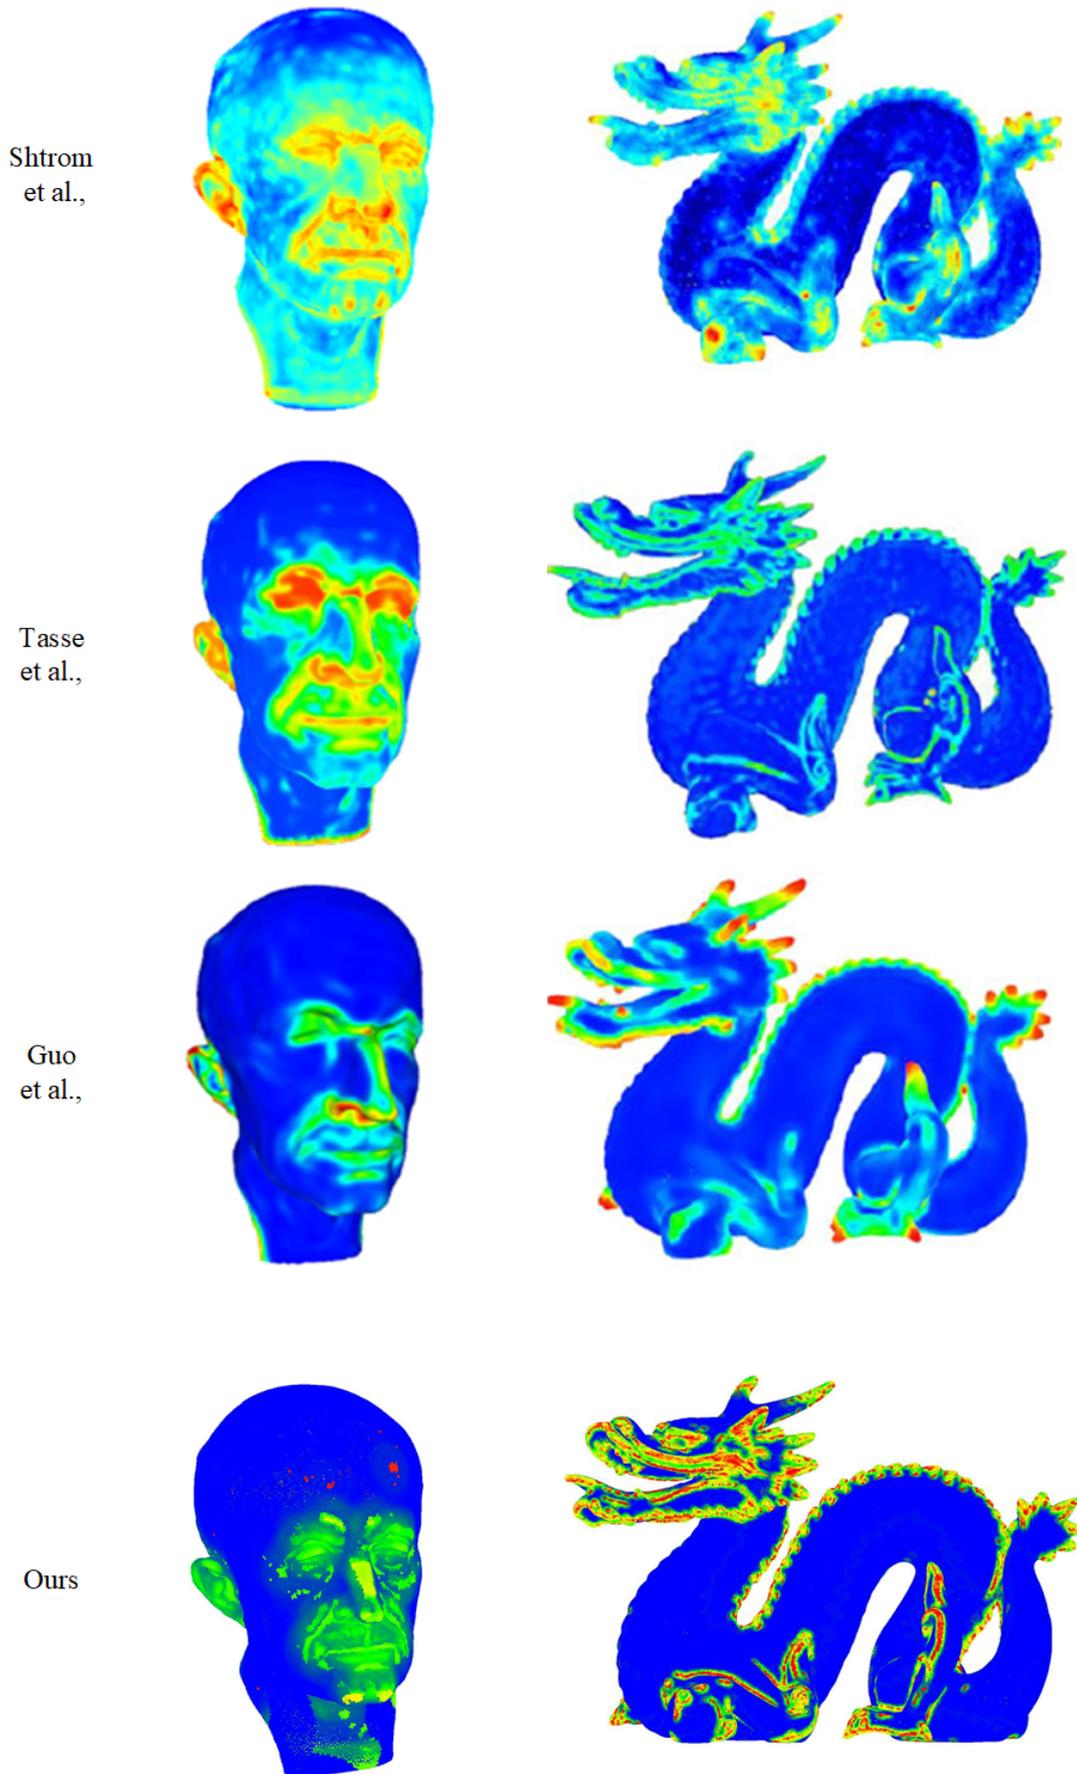


Figure 3. Comparison with point-based saliency estimations of two watertight benchmark models.

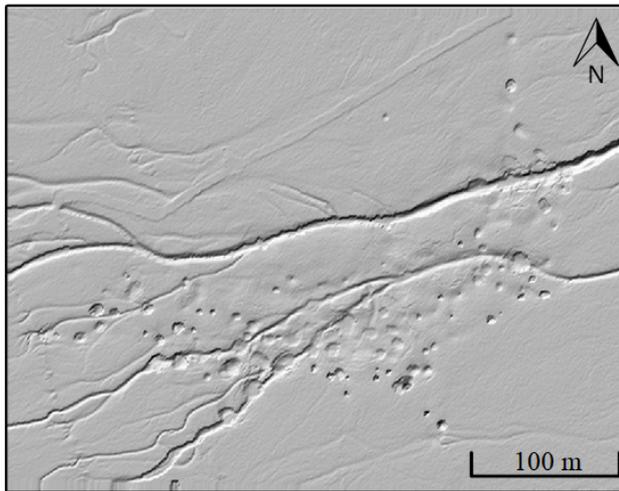


Figure 4. Airborne laser scan of an alluvial fan (31°20'N, 35°25'E).

Achanta et al. (2008) and Achanta et al. (2009). One is global in nature, while the second is local. The global approach (Achanta et al., 2009) is based on estimating the saliency by:

$$\zeta(x, y) = |\bar{I}_\mu - \hat{I}_\sigma| \quad (17)$$

where \bar{I}_μ the mean value of the elevations, before blurring; \hat{I}_σ is the smoothed version of I computed by difference of Gaussian band-pass filter on the elevation values with the ratio of σ_1/σ_2 of 1.6. The local approach (Achanta et al., 2008) computes the saliency by local contrast of sub-regions within the data, i.e.,

$$\zeta_{i,j} = D \left[\left(\frac{1}{N_1} \sum_{p=1}^{N_1} I_p \right), \left(\frac{1}{N_2} \sum_{q=1}^{N_2} I_q \right) \right] \quad (18)$$

where D is a Euclidean distance function; I_q, I_p are the elevation values; and N_1, N_2 are local neighborhoods.

Fig. (5a) shows the saliency results of the first approach (Achanta et al., 2009) with four levels of blurring, as proposed by the authors. One can see that the eastern part of the northern gully is more salient than others, while some of the gullies are completely ignored (yellow arrows, as an example). As for the sinkholes, shallow ones have been ignored completely (such as the upper turquoise arrows), and some have been identified as one unit of low-saliency region (lower turquoise arrow). This can be attributed to their depth, as the saliency here is estimated according to the features elevation with respect to the global one. Deeper entities appear more salient than others, attracting the detection towards them. Fig. (5b) shows saliency estimations using Achanta et al. (2008). Here the saliency is estimated locally, and therefore, most of the gullies and sinkholes are marked as salient. However, clusters of sinkholes are detected as a low saliency region – as locally they are not as salient (e.g., lowest turquoise arrow). Our directional saliency applies a local concept, only instead of changes in depth, it measures the variation in both normal direction and curvature. Fig. (5c) depicts the saliency map according to the convexity measure. Here, shallow entities are extracted, regardless of their neighboring features (e.g., yellow arrows), with banks of gullies emphasized. Shallow sinkholes are now highlighted, especially their rim, and the sinkhole cluster has now been mostly separated and the sinkholes are identified individually. In reference to

the non-parametric convexity measure, which plays a significant role in the saliency estimation, Fig. (6) shows the results on its application on this complex datasets. It clearly reveals all the entities, while keeping them separated one from the other. The evaluation shows that the proposed saliency highlights the various entities, especially at transition zones, facilitating a higher detection rate.

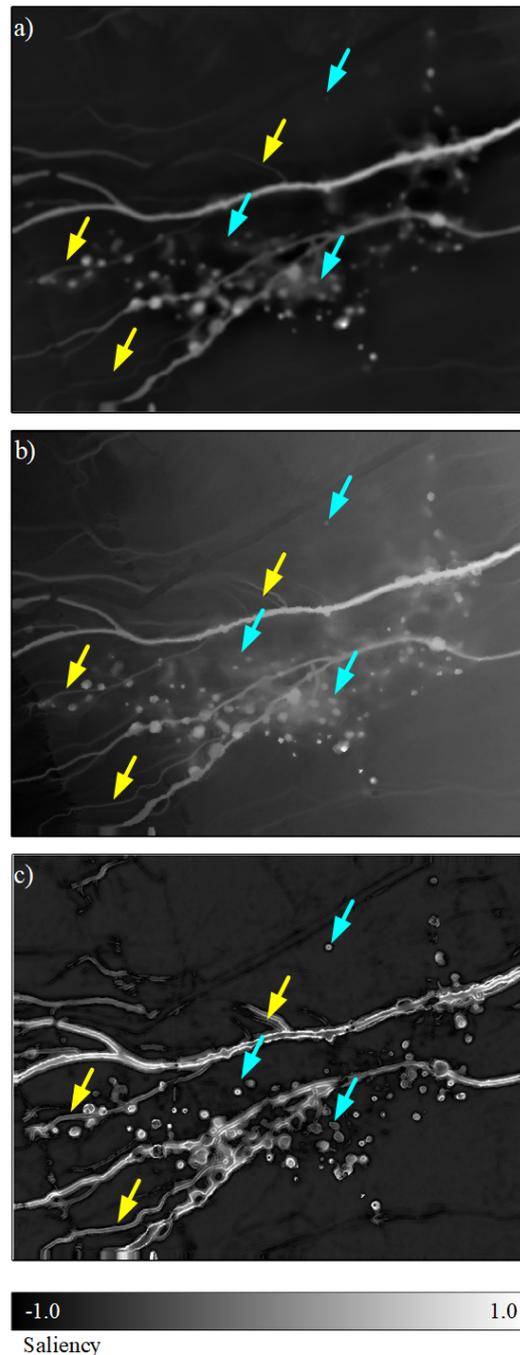


Figure 5. Saliency estimations of the alluvial fan: a) Global method (Achanta et al., 2008) with $\sigma = \frac{1}{16}$ at four levels of blurring; b) Local method (Achanta et al., 2009) with sub-region size of 5 m; c) proposed directional saliency, with $\sigma = 0.2 m$ (based on empiric tests) and $K = 8 m$. The arrows mark examples to shallow gullies (yellow) or sinkholes (turquoise) that are emphasized by our method.

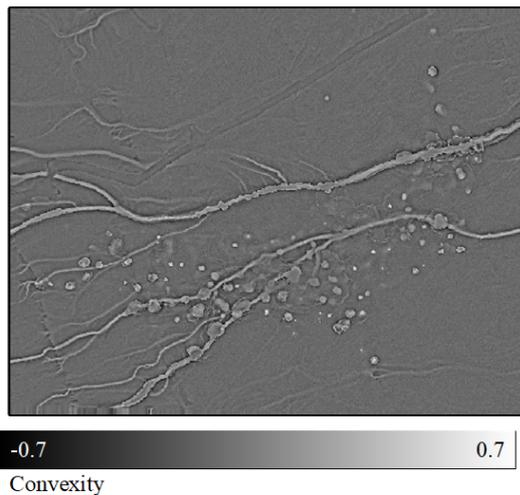


Figure 6. Convexity estimation at the alluvial fan, using a neighborhood of radius 4 m, in accordance to the approximate mean size of the entities.

3.3 Terrestrial laser scan

The last dataset is a terrestrial based scan of an archaeological site, The Leopard Temple, located in the 'Uvda Valley, Southern Israel (29°57'N, 34°58'E). The site, dated to 7500 BP (based on C¹⁴ evaluation), is considered to have been in use for ~4000 years (Avner, 2002). East of the main temple, a unique specimen of 16 animal-like figures, made of small stones affixed to the ground, and arranged along a 15 m stretch (Fig. 7). The figures were identified as leopards, due to their raised tails and the dark stones that symbolize their spots. Scans were acquired by the Leica c10 terrestrial laser scanner, with resolution of 0.1° (Fig. 7b). The smallest stones are 1-3 cm, with 7 cm space between them, while the larger ones are 10-30 cm, positioned less than 4 cm apart (Fig. 7b). Most stones are no higher than 2 cm above the ground, ranging up to 10 cm, with one reaching 20 cm above ground. Here we focus our discussion on a characteristic detail from the entire site, which features 3-10 cm long stones, and 1-5 cm high, and one large 30 × 20 × 15 cm stone (Fig. 8a). Notably, the terrain itself, though flat, is not smooth.

We compare our results to the approach proposed by Shtrom et al. (2013), as it sets the basis for most of the later works (Tasse et al., 2015; Kobyshev et al., 2016; Ding et al., 2019). The authors there propose to measure the global and local difference of the features based on the uniqueness of the normals in the scan. When applied, results emphasize the leopards' wider surrounding, in addition to ground related features (Fig. 8b). However, it does not emphasize details, only a wide region as salient.

Applying our method, we begin with the estimation of the non-parametric convexity measure. The neighborhood size is set to the minimal object size (3 cm). However, Fig. (8c) shows that at this size, the roughness of the surface is modeled as well. Removal of ground texture using Eq. (9) and the application of the saliency, which also considers normal changes, lead to a cleaner depiction of the leopard images (Fig. 8d). Note that ground related regions are suppressed and marked as "non-interesting". Minimal distinctness distance was set to 7 cm, in accordance with the stone spacing, which allows the distinction between small stones.

In sum, our approach, being center-surround based, has managed to localize on the leopard images, while not over- or undershooting the detection of details.

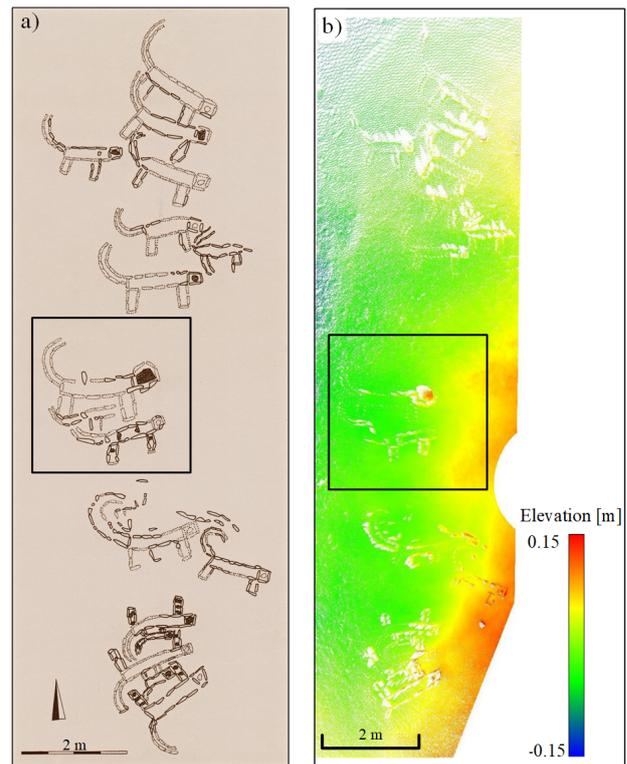


Figure 7. The Leopard Temple: a) site map (with permission of Israel Antiquities Authority[©]); b) scanned point-cloud

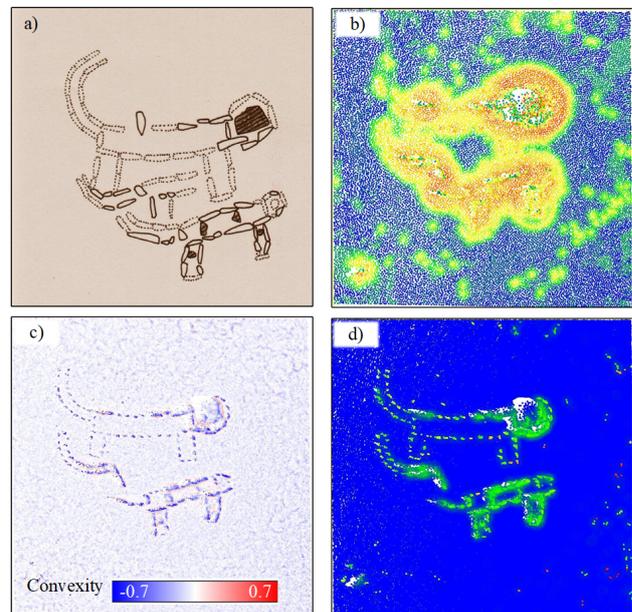


Figure 8. The Leopard Temple results: a) detail from Fig. (7); b) application of the saliency method proposed by Shtrom et al. (2013); c) proposed convexity estimation; d) our saliency estimation.

4. CONCLUSIONS AND FUTURE WORK

This paper introduced a new saliency measure for 3-D point-clouds that considers not only the direction of a local surface compared to the scene, but is also aware of the velocity and magnitude of the change. This is carried by characterizing the underlying surface, and quantifying its change with respect to its close neighborhood. Compared to state-of-the-art applica-

tions, our method runs faster, and yields comparable, if not better results. Within open environments, our approach facilitates the detection of salient features in both disturbed and smooth surfaces, where subtle entities are embedded within the topography, and where other approaches fail. The proposed saliency measure can be further applied to object extraction, point-cloud reduction, registration, classification, and enhanced visualization schemes, while improving the computational overhead, manipulation and analysis of these data.

ACKNOWLEDGMENTS

We would like to thank the Stanford Computer Graphics Laboratory for the Dragon benchmark model and to the Computer Graphics Group in the Max-Planck-Institut für Informatik for the Max-Planck bust.

References

- Abelson, M., Baer, G., Shtivelman, V., Wachs, D., Raz, E., Crouvi, O., Kurzon, I., Yechieli, Y., Meir, A., Gidon, B., Vladimir, S., Daniel, W., Eli, R., Onn, C., Ittai, K., Yoseph, Y., 2003. Collapse-sinkholes and radar interferometry reveal neotectonics concealed within the Dead Sea basin. *Geophysical Research Letters*, 30(10).
- Achanta, R., Estrada, F., Wils, P., Süstrunk, S., 2008. Salient region detection and segmentation. *Computer Vision Systems*, 66–75.
- Achanta, R., Hemami, S., Estrada, F., Süstrunk, S., 2009. Frequency-tuned salient region detection. *Computer vision and pattern recognition, 2009. cvpr 2009. ieee conference on*, IEEE, 1597–1604.
- Arya, S., Mount, D. M., Netanyahu, N. S., Silverman, R., Wu, A. Y., 1998. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *Journal of the ACM*, 45(6), 891–923.
- Avner, U., 2002. Studies in the Material and Spiritual Culture of the Negev and Sinai Populations, During the 6th-3rd Millennia B.C. PhD thesis, The Hebrew University, Jerusalem.
- Baruch, A., Filin, S., 2011. Detection of gullies in roughly textured terrain using airborne laser scanning data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(5), 564–578.
- Ding, X., Lin, W., Chen, Z., Zhang, X., 2019. Point Cloud Saliency Detection by Local and Global Feature Fusion. *IEEE Transactions on Image Processing*, 1–1.
- Frintrop, S., Rome, E., Christensen, H. I., 2010. Computational visual attention systems and their cognitive foundations. *ACM Transactions on Applied Perception*, 7(1), 1–39.
- Guo, Y., Wang, F., Xin, J., 2018. Point-wise saliency detection on 3D point clouds via covariance descriptors. *Visual Computer*, 34(10), 1325–1338.
- Itti, L., Koch, C., Niebur, E., 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1254–1259.
- Kobyshev, N., Riemenschneider, H., Bódis-Szomorú, A., Gool, L. V., 2016. 3D Saliency for Finding Landmark Buildings. *2016 Fourth International Conference on 3D Vision (3DV)*, 267–275.
- Kopp, J., 2008. Efficient numerical diagonalization of Hermitian 3×3 matrices. *International Journal of Modern Physics C*, 19(03), 523–548.
- Lee, C. H., Varshney, A., Jacobs, D. W., 2005. Mesh saliency. *ACM Transactions on Graphics (TOG)*, 24, ACM, 659–666.
- Li, T., Song, H., Zhang, K., Liu, Q., Lian, W., 2019. Low-rank weighted co-saliency detection via efficient manifold ranking. *Multimedia Tools and Applications*.
- Li, X., Lu, H., Zhang, L., Ruan, X., Yang, M. H., 2013. Saliency detection via dense and sparse reconstruction. *Proceedings of the IEEE International Conference on Computer Vision*, 2976–2983.
- Muja, M., Lowe, D., 2009. FLANN – Fast Library for Approximate Nearest Neighbors, user manual.
- Rusu, R. B., Blodow, N., Beetz, M., 2009. Fast Point Feature Histograms (FPFH) for 3D registration. *2009 IEEE International Conference on Robotics and Automation*, IEEE.
- Shtrom, E., Leifman, G., Tal, A., 2013. Saliency Detection in Large Point Sets. *Computer Vision (ICCV), 2013 IEEE International Conference on*, 3591–3598.
- Sirmacek, B., Shen, Y., Lindenbergh, R., Zlatanova, S., Diakite, A., 2016. Comparison of ZEB1 and Leica c10 indoor laser scanning point clouds. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 3, 143–149.
- Tasse, P. F., Kosinka, J., Dodgson, N., 2015. Cluster-based point set saliency. *Proceedings of the IEEE international conference on computer vision*, 163–171.
- van de Weijer, J., Gevers, T., Bagdanov, A., 2006. Boosting color saliency in image feature detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(1), 150–156.
- Wang, Z., Zhang, L., Fang, T., Mathiopoulos, P. T., Tong, X., Qu, H., Xiao, Z., Li, F., Chen, D., 2015. A multiscale and hierarchical feature extraction method for terrestrial laser scanning point cloud classification. *Geoscience and Remote Sensing, IEEE Transactions on*, 53(5), 2409–2425.
- Weinmann, M., Hinz, S., Weinmann, M., 2017. A hybrid semantic point cloud classification-segmentation framework based on geometric features and semantic rules. *PFG—Journal of Photogrammetry, Remote Sensing and Geoinformation Science*, 85(3), 183–194.
- Wu, J., Shen, X., Zhu, W., Liu, L., 2013. Mesh saliency with global rarity. *Graphical Models*, 75(5), 255–264.
- Yuan, X., Yue, J., Zhang, Y., 2018. RGB-D Saliency Detection: Dataset and Algorithm for Robot Vision. *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 1028–1033.