Selective Range Data Acquisition Driven by Neural-Gas Networks

Ana-Maria Cretu, *Student Member, IEEE*, Pierre Payeur, *Member, IEEE*, and Emil M. Petriu, *Fellow, IEEE*

***Abstract*—The collection of the rich flow of information provided by the current generation of fast vision sensing systems brings new challenges in the selection of only relevant features out of the avalanche of data generated by those sensors. This paper discusses some aspects of intelligent sensing for advanced robotic applications, with the main objective of designing innovative approaches for automatic selection of regions of observation for fixed and mobile sensors to collect only relevant measurements without human guidance. The proposed neural-gas-network solution selects regions of interest for further sampling from a cloud of sparsely collected 3-D measurements. The technique automatically determines bounded areas where sensing is required at a higher resolution to accurately map 3-D surfaces. Therefore, it provides significant benefits over brute-force strategies as scanning time is reduced and the size of the data set is kept manageable. Experimental evaluation of this technology is presented for 3-D surface measurement and modeling.**

***Index Terms*—Feature detection, neural gas, neural networks, selective sensing, surface modeling, 3-D vision.**

I. Introduction

THE CURRENT generation of 3-D vision data acquisition devices (e.g., laser scanners) offers high measurement speed, which often results into a large amount of data. Due to the general lack of knowledge on appropriate accuracy levels for correct description of shape and geometry, the data acquisition process is often long and complex. Reducing the complexity and size of these data sets is one of the key techniques required to operate subsequent applications at a reasonable computational cost. To tackle this issue, the most widely exploited trend in the literature implies the postprocessing of large data sets obtained by acquisition devices. Frequently, the proposed algorithms rely on the user input for providing parameters such as the desired density of sampling, the regularity of sampling, and the minimum distance between samples. This is a difficult task as the user is not always aware of the appropriate level of accuracy required for a model to be further processed, and the adjustment of such parameters can lead to a lengthy trial-and-error procedure. A significant advance is expected from the implementation of automated selective procedures to determine regions of interest and collect only relevant measurements that are contributing to the modeling applications.

The main objective of this research is the design of an innovative approach to achieve automatic selection of regions of observation for vision sensors to collect only relevant measurements without human guidance. The relevant regions of interest are extracted from 3-D point clouds during the acquisition procedure to prevent an avalanche of data and the related excessive processing load. Starting from an initial fast sparse scan of an object, a neural-gas network is used to selectively identify areas of interest for additional scanning to improve the accuracy of the model. The final model is a multiresolution 3-D model with a higher resolution in areas rich in features.

The paper, which is an extension of our previous work [1], is structured as follows. We start by showing the state of the art in postprocessing and sampling of large data sets in the form of point clouds in Section II. We detail our proposed solution for selective sensing in Section III. Section IV presents a comparison with classical sampling solutions. It also shows experimental results for data sampling using vision sensors, including the quantitative evaluation of errors and training times. Finally, we present future research directions and draw the conclusions.

II. Literature Review

In general, there are three sampling policies proposed in the literature: uniform sampling, random sampling, and stratified sampling. In uniform (regular or grid) sampling, samples are spread such that the probability of a surface point to be sampled is equal for all surface points. The method is popular because it can easily be implemented and ensures complete coverage of the surface within the sensor's field of view. However, the cost is high, since to achieve adequate sampling density over those regions requiring the highest resolution, the sampling density must be uniformly high everywhere. In random sampling, each point of the object has an equal chance of being selected, but only a lower number of points are collected. As the percentage of sampled points increases, the cost gets higher to eventually reach the one of uniform sampling. The risk here is that samples randomly collected can miss important features.

Another type of sampling is stratified sampling. This technique involves the collection of evenly spaced samples by subdividing the domain into nonoverlapping clusters and by sampling independently from each partition. Such a method ensures that an adequate sampling is applied to all partitions.

Manuscript received June 30, 2008; revised February 7, 2009. First published April 24, 2009; current version published July 17, 2009. This work was supported in part by the Communications and Information Technology Ontario Centre of Excellence and in part by the Natural Sciences and Engineering Research Council of Canada. The Associate Editor coordinating the review process for this paper was Dr. Dario Petri.

The authors are with the Sensing and Modeling Research Laboratory, School of Information Technology and Engineering, University of Ottawa, Ottawa, ON K1N 6N5, Canada (e-mail: acretu@site.uottawa.ca; ppayeur@site.uottawa.ca; petriu@site.uottawa.ca).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIM.2009.2015643

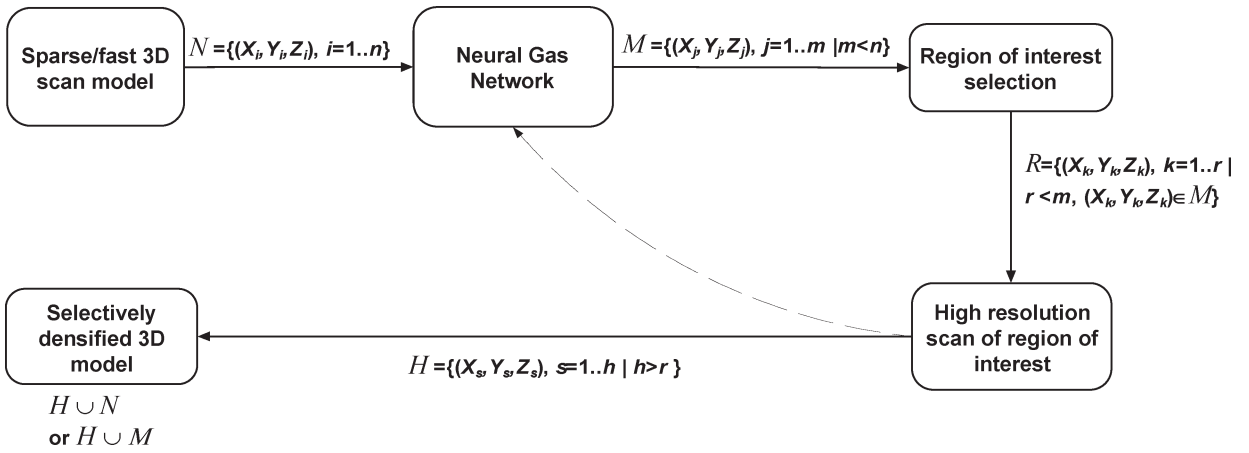

Fig. 1. Proposed framework for selective sensing.

It is often exploited in the context of postprocessing of large point clouds or meshes [2]–[8], where a subdivision of models into grid cells occurs. In such cases, sample points that fall into the same cell are replaced by a common representative.

The 3-D model is first voxelized with an octree, and one sample is output for each voxel. The common representative for a voxel can be selected according to a probability that decays as the distance of the sample to the center of the voxel increases [2]. The representative sample can be chosen to be the measured point that is closest to the average of points that fall into the same voxel [3] or the point whose normal is closest to the average of the points in the same voxel [4]. Alternatively, principal component analysis can be performed for each cell of the octree to efficiently decimate data [5].

Surface-based clustering can be employed as well instead of volumetric voxelization. With this technique, clusters are built by collecting neighboring samples while taking into account the local sampling density [6]. Points are incrementally added to a cluster until a maximum size and/or a maximum allowed variation is reached.

Another category of solutions for the decimation of large data sets is based on boundary segmentation [7], [8]. Meshes are segmented into the boundary of the original domain and the interior samples. Each part is separately simplified.

All the above methods are not meant to be incorporated in the actual sampling procedure but rather to postprocess collected data. An approach to integrate the sampling procedure into the measurement process is proposed by Pai *et al.* [9], [10]. They consider a known mesh of the object under study, as well as a set of parameters such as the maximum force exerted on the object, the maximum probing depth, and the number of steps for the deformation measurement. During probing, an algorithm generates the next position and orientation for the probe based on the specifications and the mesh of the object under test. It simultaneously performs proximity checks and verifies the expected contact location of the probe with the mesh based on line intersection. However, the procedure is not selective and therefore can result in the collection of data for all the points over the mesh.

The work detailed in this paper proposes an innovative automated selective data acquisition algorithm to guide online the probing of only relevant points for the objects under study. The relevant regions of interest are extracted from 3-D point clouds during the acquisition procedure to prevent an avalanche of data and the related excessive processing load. "Online" in this context refers to the use of the algorithm in each successive step of the measurement procedure, as opposed to postprocessing the data collected by a sensor as performed in [2]–[8]. The proposed algorithm has the purpose of increasing the relevance and speed of probing by online progressive refinements. The algorithm is not adaptive like the one proposed in [9] and [10] in that it does not adjust the probing procedure as a result of the interaction with the object. It is rather selective, in that it picks out the main geometrical features that characterize the object under study. An iterative use of the algorithm results in multiresolution object models, with a higher resolution in the areas rich in geometrical features.

III. Proposed Framework

Meant to be incorporated directly in the sampling procedure, the proposed automated selective scanning framework uses a self-organizing neural network. The latter automatically selects regions of interest for further refinement from a cloud of 3-D sparsely collected measurements. The framework is depicted in Fig. 1.

Starting from an initial low-resolution scan of n 3-D points of an object, denoted N in Fig. 1, a neural-gas network is employed to model the corresponding point cloud. The resulting neural-gas map contains m nodes, denoted M, that approximate the point cloud. The regions that require additional sampling to ensure an accurate model are detected by finding higher density areas in the neural-gas map. This is done by first applying a Delaunay tessellation to the neural-gas output map. Subsequently, all the edges of triangles that are larger than a set threshold are removed from the tessellation. The threshold is automatically computed based on the length of vertices for every triangle in the initial tessellation. The subset of nodes R extracted from the neural-gas map drives the rescanning over the regions of interest to obtain an extra set of points H at a higher resolution. A multiresolution model is then built by augmenting the initial sparse model with the points collected at a higher resolution over the regions of interest $(H \cup N)$. If slightly less accurate but more compact models are desired, the initial 3-D points can

be approximated by their corresponding nodes in the neural-gas map. In this case, the selectively densified model can be obtained by augmenting the initial neural-gas map with the higher resolution data samples $(H \cup M)$.

A. Neural-Gas Network

The use of a self-organizing architecture is justified by its ability to quantize the given input space into clusters of points with similar properties. As mentioned in the literature review, clustering is an efficient way to compress data. In our previous research, we proved that a neural-gas network is able to cluster both geometric and elastic properties of the objects embedded in a modeled point cloud [11]. The neural-gas network is selected instead of other self-organizing architectures due to its capability to capture fine details, unlike other architectures that tend to smooth them, such as the Kohonen self-organizing map [12].

The neural-gas algorithm can formally be described as follows [13], [14]. The algorithm starts by initializing the set S of network nodes to contain m units c_i with corresponding reference vectors $w_{c_i} \in \Re^p$. Each unit c has an associated p-dimensional reference vector that indicates its position in the input space. The reference vectors are randomly chosen according to a probability density function or from a finite set.

In each training step, an input vector x is presented to the network, and the winning neuron that best matches the input is identified using the equation

$$s(x) = \arg\min_{c \in S} \|x - w_c\| \tag{1}$$

where $\|.\|$ denotes the Euclidean vector norm. The neurons to be adapted in the learning procedure are selected according to their rank. The rank of neurons is the rank they have in an ordered list of distances between their weights and the input vector. Each time a new input vector x is presented to the network, a neighborhood ranking index list is built $(j_0, \ldots, j_{N-1})$; here, w_{j_0} is the reference vector closest to x, w_{j_1} is the reference vector second closest to x, and w_{j_k} is the reference vector such that k vectors w_i exist with

$$\|x - w_i\| \leq \|x - w_{j_k}\|. \tag{2}$$

In each training step, the best matching neuron $s(x)$ at time t is computed using the minimum Euclidean distance criterion (1). All neurons are then ordered according to (2). Based on this rank, a certain number of units are adapted.

The neurons' weights are updated according to the following rule:

$$w_j(t+1) = w_j(t) + \alpha(t) h_\lambda\left(k_j(x, w_j)\right)\left[x(t) - w_j(t)\right] \tag{3}$$

where $\alpha(t) \in [0, 1]$ describes the overall extent of the modification, and h_λ is one for $k_j(x, w_j) = 0$ and decays to zero for higher values according to

$$h_\lambda\left(k_j(x, w_j)\right) = \exp\left(-\frac{k_j(x, w_j)}{\lambda(t)}\right) \tag{4}$$

where $k_j(x, w_j)$ is a function that represents the ranking of each reference vector w_j [14]. If j is the closest to input x, then $k = 0$, for the second closest, $k = 1$, and so on. The learning rate $\alpha(t)$ and the function $\lambda(t)$ are both time dependent. These parameters are slowly decreased during the learning process to ensure that the algorithm converges. Usually, the following time dependences are used [13], [14]:

$$\alpha(t) = \alpha_0 \left(\frac{\alpha_T}{\alpha_0}\right)^{t/T} \tag{5}$$

$$\lambda(t) = \lambda_0 \left(\frac{\lambda_T}{\lambda_0}\right)^{t/T} \tag{6}$$

where the constants α_0 and λ_0 are the initial values for $\alpha(t)$ and $\lambda(t)$, α_T and λ_T are the final values, t is the time step, and T is the training length. For the time-dependent variables, some initial and final values have to be chosen. The algorithm continues to generate random input signals x while $t < T$.

In the context of this paper, the neural gas is employed to model point clouds collected during a sparse scan of an object via an active range finder. The network starts with the points N in this point cloud (as defined in Fig. 1) in the form of (X_i, Y_i, Z_i), $i = 1, \ldots, n$, with n being the size of the point cloud, and an initial configuration of unconnected nodes. During adaptation, the latter moves over the data space, and the model asymptotically contracts toward the points in the input space, respecting their density. Therefore, the nodes progressively take the shape of the objects encoded in the point cloud. The final map obtained by the neural gas after the adaptation (learning) procedure is denoted M, as shown in Fig. 1.

The training is stopped early to avoid that the nodes become uniformly distributed instead of capturing details. The point clouds collected are rasterlike models, and their density is uniform. Therefore, the neural gas, which inherently tends to respect the density in the point cloud, would tend to build uniformly dense maps after a long training time, as opposed to keeping nodes in the regions rich in features. This justifies the shortened training period and ensures that the density of the selected probing points will be higher in the regions with more pronounced variations in the geometric shape.

B. Regions of Interest Selection

A simple technique is used to detect higher density regions, which are the ones of interest in the neural-gas map. The principle is illustrated in Fig. 2, where a typical configuration of the nodes M in the neural-gas map is shown after the adaptation procedure [Fig. 2(a)]. A Delaunay triangulation is first applied over the output map to connect the nodes of the neural-gas map. The resulting tessellation is depicted in Fig. 2(b). It can be observed that two areas of high density of points are identified in the present example, both related to the existence of 3-D features (due to the modeling properties of neural gas). These two areas are shown shaded in Fig. 2(b). Such areas of high density of points are represented by small triangles (short edges) in the tessellation. Next, the triangulation is successively traversed, and the length of vertices between every pair of

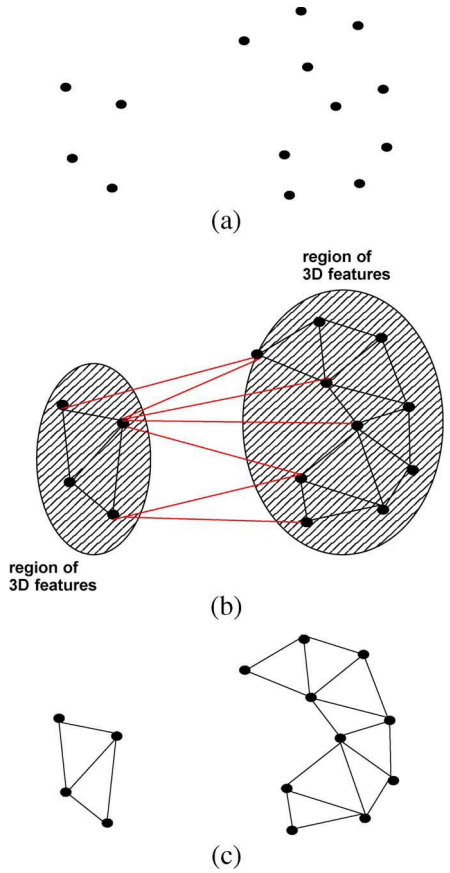

Fig. 2. Delaunay triangulation and triangle-removal procedure. (a) Example of a possible configuration of nodes in the neural-gas map after adaptation. (b) Delaunay triangulation applied on the nodes with long edges between the two regions of 3-D features marked in red. (c) Configuration after the removal of the long edges from the Delaunay triangulation.

points for every triangle is computed. The mean value of all these lengths is estimated, and a threshold is set equal to this value. All the vertices longer than the threshold value are then removed from the model. The removal of the edges longer than the threshold ensures the identification of close points and, therefore, dense areas. The long edges, marked with red in Fig. 2(b), are removed from the tessellation. The result of the removal procedure is shown in Fig. 2(c) and represents the set of points R belonging to the regions of interest. The remaining triangles and the associated points identify those regions where the density of nodes is higher and that require additional sampling. Supplementary data are collected only over these regions to further collect the set of points H that will augment the model in the higher resolution regions. A selectively densified multiresolution model is constructed either by augmenting the initial sparse low-resolution point cloud with the higher resolution data samples $(H \cup N)$ or by augmenting the initial neural-gas map with the higher resolution data samples $(H \cup M)$.

IV. Experimental Results

A. Test Objects

The proposed method for range imaging with selective sampling is experimentally evaluated on three objects with different sorts of features: a toy triceratops, a foam armchair, and a mock-up car door, as depicted in Fig. 3. The areas of the head, the neck, and the horns correspond to the regions of interest for the triceratops model. For the armchair, the regions of interest are the edges. For the door model, the regions of interest are found around the doorknob and the door-opening gap. Moreover, the 3-D measurements on these three objects come from different ranging sensing systems. This offers a supplementary dimension in the evaluation of the capability of the neural-gas solution to cope with features of different nature and complexity and with data obtained with various sensing mechanisms.

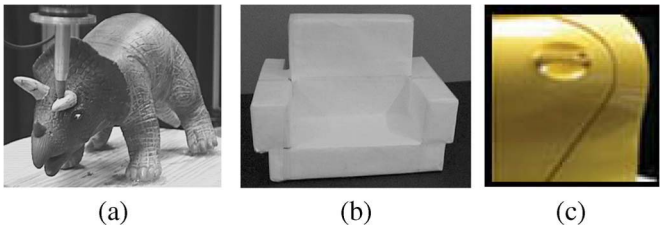

Fig. 3. Objects used for testing. (a) Toy triceratops. (b) Foam armchair. (c) Mock-up car door.

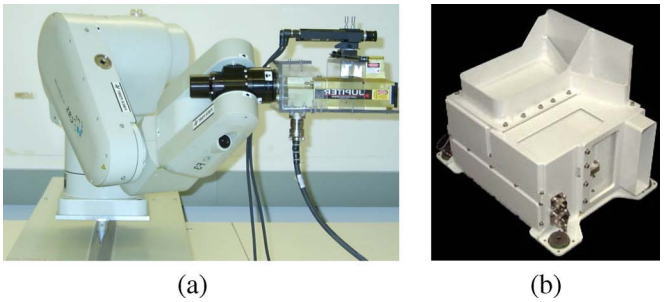

Fig. 4. Measurement equipment. (a) Jupiter laser scanner. (b) Neptec's full-image LMS range sensor.

To create the point cloud for the toy triceratops in Fig. 3(a), an initial coarse mesh is built on a cloud of points collected using a stereo–vision system. As the data were made available to us [10], no further details are provided here on the sensing equipment and the data collection procedure. This coarse mesh is then interpolated on a high-resolution mesh. The model of the armchair depicted in Fig. 3(b) is obtained using an automated Jupiter laser scanner mounted on a robot arm [15], as shown in Fig. 4(a). The high-resolution point cloud shown in Fig. 3(c) is collected with the Neptec Design Group Inc.'s Laser Metrology System (LMS) range sensor, depicted in Fig. 4(b).

Starting from an initial fast sparse scan of each object or a subsampled point cloud, a neural-gas network is employed to model the data, in the form of (X_i, Y_i, Z_i) coordinates, for each of the above point clouds. The data set is normalized prior to the neural-gas mapping such that its variance is unity, as required by the learning procedure.

Testing is performed for several sizes of the initial fast sparse scan to identify how the resolution of this initial scan influences the modeling results and what would be the smallest scan size that allows for the modeling of fine features in the modeled objects. The training is performed using Matlab code on a machine based on a Pentium IV processor working at 1.3 GHz and a RAM of 512 MB.

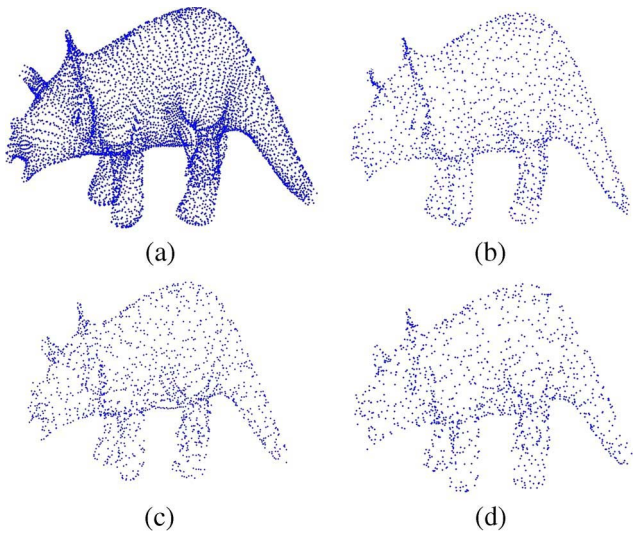

Fig. 5. Comparison of the neural-gas map with classical sampling techniques. (a) Initial sparse point cloud (6113 points). (b) Neural-gas map (1485 points). (c) Uniform sampling result (1528 points). (d) Random sampling result (1528 points).

B. Comparison With Classical Sampling Solutions

To validate the proposed framework, a comparison is initially performed between the neural-gas map and the results obtained with the classical uniform and random sampling procedures, respectively. To perform the comparison, the same sparse scan of an object is used as a starting point, and a neural-gas adaptation, a uniform sampling, and a random sampling algorithm are concurrently applied on this scan. To ensure a common basis for comparison, the same number of points is imposed at the output of each algorithm.

Fig. 5 shows the results of the comparison for the toy triceratops depicted in Fig. 3(a). For this example, the initial sparse point cloud of the triceratops contains 6113 points and is depicted in Fig. 5(a). The output is forced to be approximately 1500 points. Fig. 5(b) shows the neural-gas map of 1485 points (equivalent to a map of 33×45) obtained after the adaptation over the initial sparse scan. A set of 1528 uniformly sampled points from the same initial scan is depicted in Fig. 5(c), while an equal set of 1528 randomly sampled points is shown in Fig. 5(d). It can be observed by comparing Fig. 5(b) with Fig. 5(c) and (d) that due to its modeling properties, the neural-gas map better characterizes the fine features of the object under study when the same number of sample points is collected. The neural-gas map provides a better contour of the object, and the area around the neck and horns is much better defined than in the case of both uniform and random sampling.

The same experiments are repeated, but in this case, the number of points obtained at the output is not fixed. Tests are performed with the same initial scan for neural gas, uniform sampling, and random sampling, until reasonably good modeling results are obtained for all sampling techniques. The comparison is made on the minimum number of nodes in the neural-gas map or of samples that is required, respectively, for each method to capture the features of the object under study. It leads to the conclusion that about 1400 points in the neural-gas map allow for the visual identification of all the features in the model. Approximately double the number of points, i.e., over 3000 points, is required for the uniform and random sampling techniques to capture the same level of fine features in the triceratops point cloud. This demonstrates that the neural gas can provide more compact models while preserving fine details. This represents a clear advantage in the context of selective sampling.

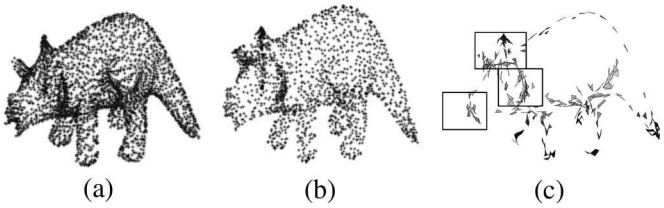

Fig. 6. (a) Initial scan of the toy triceratops. (b) Neural-gas model for a map size of 40×45. (c) Detected regions of interest for further sampling.

Moreover, the neural-gas map constitutes a good basis to identify the bounded regions of interest in the scan to guide the sensors for additional scanning. The designed edge-removal procedure described in Section III-B allows for an automated selection of bounded areas that contain geometrical features. Neither the uniform sampling procedure nor the random one possesses such characteristics. All these aspects confirm that the neural-gas solution is appropriate in the context of selective sampling and that it performs better than classical sampling algorithms.

C. Experimental Results for Test Objects

The point-cloud model of the toy triceratops is representative of objects with a roughly uniform distribution of sampling points and multiple small features. An initial sparse point cloud of 3065 points of the triceratops, representing roughly 25% of the full resolution scan of 12 226 points available, is shown in Fig. 6(a). The trained neural-gas network having as input this initial sparse point cloud and a map size of 40×45 (about 58% of the initial point cloud) represents a compressed model of the data set, as depicted in Fig. 6(b). In spite of some noise, the network is able to identify all the areas that require additional scanning, e.g., the areas around the head, the neck, and the horns of the triceratops. These areas, shown framed in rectangles in Fig. 6(c), can be scanned at a higher resolution to obtain the selectively densified model of the triceratops.

The armchair is representative of an object without many features but with sharp depth transitions. It is composed of rasterlike distributed sampling points. The results for the neural-gas modeling and the steps performed for the selection of regions of interest are presented in Fig. 7. Fig. 7(a) represents 3845 points in the initial scan that corresponds to roughly 25% of the full resolution scan of 15 382 points available. The data are used as provided by the Jupiter laser scanner, without any filtering procedure. This point cloud is provided to a neural-gas network with a map size of 45×40 (about 46% from the size of the initial point cloud). The obtained output map represents a compressed model for the data set in which the weight vectors consist of the 3-D coordinates of the object's points and is shown in Fig. 7(b). The artifacts in the neural-gas map are due

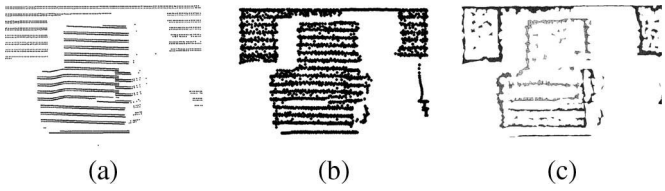

Fig. 7. (a) Initial scan of an armchair. (b) Neural-gas model for a map size of 45×40. (c) Detected regions of interest for further sampling.

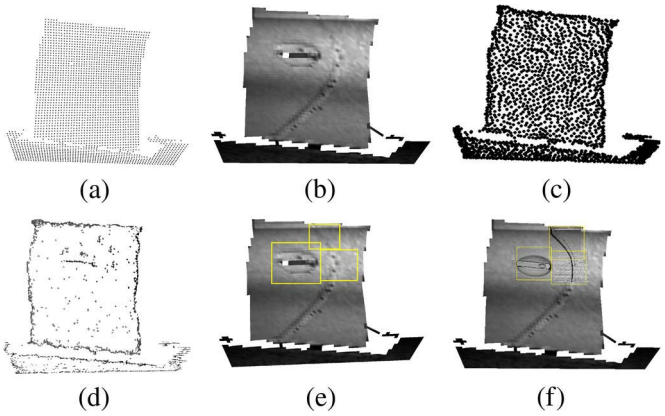

Fig. 8. Modeling steps for the mock-up car door. (a) Initial point cloud. (b) Rendered mesh model. (c) Neural-gas model. (d) Higher density areas in the neural-gas model. (e) Identified regions of interest. (f) Selectively densified model.

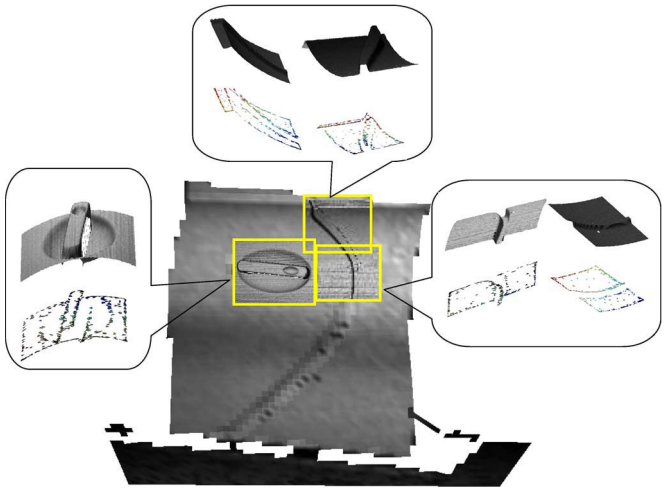

Fig. 9. Model of the car door showing selected areas that require additional sampling from different viewpoints and the regions of interest identified in a neural-gas map for each of them.

to the fact that the training is stopped early, as explained in Section III-A. Since there are no other significant details in the model apart from the edges and the back plane, these transition areas, shown in Fig. 7(c), get selected as regions of interest for additional sampling.

The more complex case of the car door is shown in Figs. 8 and 9. This object presents fine details as the regions representing the door opening and the doorknob are very small in comparison with the whole model and so is the number of points representing them. A subsampling is initially performed to obtain a very low resolution point cloud of 4096 points, as shown in Fig. 8(a). This represents 0.39% of the high-resolution scan of 1 048 576 points available. The data provided by the Neptec's full-image LMS range sensor are less noisy, due to its higher resolution when compared to the Jupiter laser scanner. Fig. 8(b) shows the rendered mesh model embedded in the initial point cloud. The normalized point cloud is provided as an input to a neural-gas network with a map size of 40×45 (43% of the number of points in the initial very low resolution sparse scan), and the resulting neural-gas map is depicted in Fig. 8(c). The regions of high density in the neural-gas output map, shown in Fig. 8(d), are identified by building the Delaunay triangulation and removing large edges from its triangles, as described in Section III-B. The regions of interest that require additional sampling to ensure an accurate model, excluding the contours of the door, are superimposed on the rendered model and depicted in Fig. 8(e). The resulting selectively densified multiresolution model including the extra samples over the regions of interest is presented in Fig. 8(f). It contains 111 596 points, which represents a reduction of about 90% in the number of points when compared to the full high-resolution scan.

The same procedure can iteratively be repeated for each of the regions of interest detected in the previous step. The extra points collected over each selected region can be provided as input to a neural-gas network to further detect fine details that are worth to be scanned at higher resolutions. Fig. 9 presents the detailed areas of interest detected in the previous step from different viewpoints. For each of the regions, it also shows areas of interest identified in the neural-gas map. Each higher resolution scanned region contained about 5000 samples, and a map size of about 30% of this number of points was selected. In all the presented cases, the regions of interest for additional scanning are correctly identified. The number of steps in which the procedure is repeated depends not only on the accuracy needed for the application but also on the type of scanner employed (e.g., its level of granularity and the maximum allowed resolution).

D. Selection of the Neural-Gas-Map Dimension

To fully benefit from the proposed method, research has been done to identify ways to select an appropriate neural-gas-map size for the objects under study that ensures the identification of regions of interest. In particular, the investigation looked into the correlation between the neural-gas-map size and the size of the initial point cloud. For this purpose, tests were run for all the objects under study for different sizes of the initial sparse scan and different neural-gas-map sizes. These tests revealed that a neural-gas network, regardless of its neural-gas-map size cannot capture enough features in an initial point cloud with a size that is smaller than 2000 points, as illustrated in the first row of Fig. 10. The second, third, fourth, and fifth rows of Fig. 10 show some good modeling results obtained for 2500–3000 points, 4000–5000 points, 15 000–16 000 points, and over 16 000 points as the initial point-cloud size, respectively. The "good" modeling results are identified as those compressed models with the highest compression rate that offer at the same time an optimal balance between the quality of the model (low relative error) and the time required for training. The relative error is computed as an

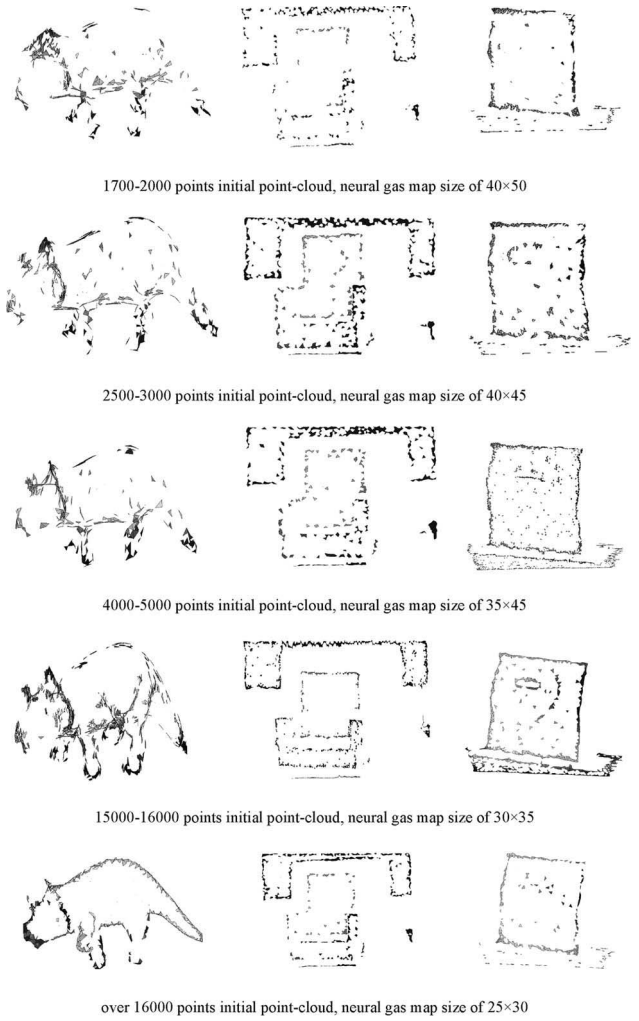

Fig. 10. Examples of regions of interest detected for different sizes of the initial point cloud and different neural-gas-map sizes.

average Euclidean distance between each data vector and its winning neuron and therefore shows how close the modeled data are to the initial point cloud.

It can be observed that starting from 2500 points in the initial sparse scan, the network is able to capture the edges of the armchair but only barely captures the triceratops and door features. To correctly capture the area around the horns of the triceratops and the fine details of the door model, at least 4000 points are required. The same observations were made for other objects, leading to the conclusion that the initial point cloud should contain a minimum of 4000 points for objects exhibiting tiny features.

The same tests revealed that a larger neural-gas-map size better models a lower resolution initial point cloud, while a smaller map size is enough for a higher resolution initial point cloud. As shown in Fig. 10, a neural-gas-map size of 40×50 is required to successfully capture some of the features of the object under study for the initial scan with 1700–2000 points, while a map size of 25×30 is sufficient for initial point clouds of over 16 000 points.

Overall, the visual quality improves with the increase in the size of the initial point cloud. This fact is also proven by a decrease in the error when the map size increases, as shown in Fig. 11. Figs. 11 and 12 depict the evolution of the error and training time for different sizes of the initial point cloud and of the neural-gas map in the cases of the triceratops and the armchair. The error, even if it remains low overall, tends to decrease with the increased size of the neural-gas map and reaches lower values for a higher resolution initial point cloud. Therefore, a higher resolution of the initial sparse scan increases the accuracy of the sampling procedure for larger map sizes. However, the training time increases with an increase in the map size. Furthermore, the training time increases with an increase in the resolution of the initial point cloud, as can be seen in Fig. 12. A compromise must then be identified on the sizes of the initial point cloud and the neural-gas map.

On the other hand, one would expect that even higher resolution initial scans would give better results in terms of identification of regions of interest. However, this is not necessarily true, as shown in the fifth row of Fig. 10. For initial point-cloud sizes larger than 16 000 points, for the triceratops model, the contour is better highlighted, but the model loses the fine details around the legs and the belly. In the model of the door, the additional points make the door gap almost invisible, since the relative number of points representing it is significantly reduced in comparison with the size of the point cloud.

Table I summarizes the correlation that exists between the number of points in the initial point cloud and the neural-gas-map size that provides reasonably good modeling results for the corresponding initial scan. The table shows that according to the number of points in the initial scan n, one can identify a map size, as a percentage of this number of points, that ensures good modeling results. The larger the initial scan is (within the upper limit of about 16 000 points), the lower the relative map size required to capture the details. An increase in the percentage within the shown range brings better results but in a longer time. In spite of the fact that the ranges proposed for point-cloud sizes and the neural-gas-map sizes can slightly vary with the characteristics of the objects under study (the size of objects and the number and size of features), Table I can be employed as a guideline to choose an adequate neural-gas-map size according to the required quality of the model (accuracy) and training time.

V. Conclusion

All the examples considered during this investigation demonstrate the ability of a neural-gas map to capture fine details in a point cloud sparsely collected on objects independently of the sensing technology and the distribution of points. By finding the high-density areas in the equivalent neural-gas map model rather than in the original point cloud, the proposed selective sampling procedure can be used to guide a vision sensor to collect only measurements in those regions that are dense in 3-D features. These regions are of interest for the improvement of accuracy of the obtained models, as well as for saving a large amount of less relevant data in the scans. The observed correlation between the number of points in the initial point cloud and the size of neural-gas map also provides the user with specific guidelines for choosing an adequate neural-gas map size according to the needs of the application. This greatly

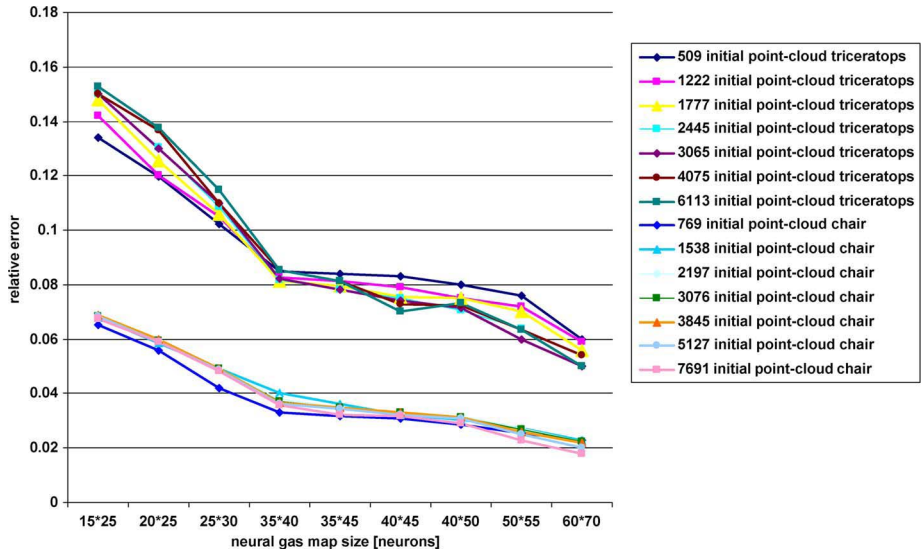

Fig. 11. Evolution of the error with the neural-gas-map size for different initial point-cloud sizes of the triceratops and the armchair.

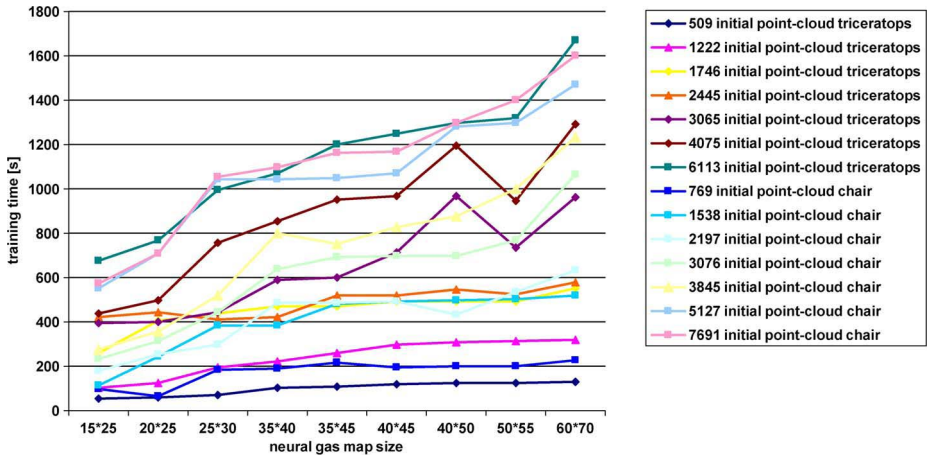

Fig. 12. Evolution of training time with the neural-gas-map size for different initial point-cloud sizes of the triceratops and armchair.

TABLE I
APPROXIMATE NEURAL-GAS-MAP SIZES RECOMMENDED FOR DIFFERENT SIZES OF INITIAL POINT CLOUDS

Size of initial scan point-cloud (number of points, n, size of N)	*Approximate neural gas map size to provide reasonably good results (number of nodes, m, size of M)*
2000-3000	(60%-100%) *n*
3000-4000	(50%-90%) *n*
4000-5000	(30%-70%) *n*
5000-6000	(20%-60%) *n*
6000-7000	(15%-50%) *n*
7000-16000	(10%-40%) *n*

contributes to automate the scanning process. Further research on the topic is directed toward means to eliminate the constraints of a fixed map size imposed by the neural-gas solution.

Under a wider perspective, the same procedure can be extended to sample elastic behavior by tactile probing under the assumption that changes in geometry can be correlated with changes in the elastic behavior of a given object. The fact that similar neural-network architectures can be simultaneously used for selective tactile and selective vision sensing opens the door to the development of multiresolution composite geometric and elastic models based on the proposed approach.

ACKNOWLEDGMENT

The authors would like to thank Dr. C. English from Neptec Design Group, Inc. and Dr. J. Lang for providing access to the high-resolution range data of the door and the triceratops models.

REFERENCES

[1] A.-M. Cretu, P. Payeur, and E. M. Petriu, "Selective vision sensing with neural gas networks," in *Proc. IEEE Int. Instrum. Meas. Technol. Conf.*, Victoria, BC, Canada, May 2008, pp. 478–483.
[2] D. Nehab and P. Shilane, "Stratified point sampling of 3D models," in *Proc. Eurographics Symp. Point-Based Graph.*, Zurich, Switzerland, Jun. 2004, pp. 49–56.
[3] R. Dillmann, S. Vogt, and A. Zilker, "Data reduction for optical 3D inspection in automotive applications," in *Proc. IEEE Int. Conf. Multisensor Fusion Integr. Intell. Syst.*, Taipei, Taiwan, Aug. 1999, pp. 159–164.
[4] K. H. Lee, H. Woo, and T. Suk, "Point data reduction using 3D grids," *Int. J. Adv. Manuf. Technol.*, vol. 18, no. 3, pp. 201–210, Aug. 2001.

[5] A. Kalaiah and A. Varshney, "Statistical point geometry," in *Proc. Eurographics Symp. Geom. Process.*, Aachen, Germany, Jun. 2003, pp. 107–115.
[6] M. Pauly, M. Gross, and L. P. Kobbelt, "Efficient simplification of point-sampled surfaces," in *Proc. IEEE Conf. Vis.*, Boston, MA, Nov. 2002, pp. 163–170.
[7] D. Uesu, L. Bavoil, S. Fleishman, J. Shepherd, and C. T. Silva, "Simplification of unstructured tetrahedral meshes by point sampling," in *Proc. IEEE Int. Workshop Volume Graphics*, E. Gröller and I. Fujishiro, Eds., Jun. 2005, pp. 157–238.
[8] H. Song and H.-Y. Feng, "A point cloud simplification algorithm for mechanical part inspection," in *Information Technology for Balanced Manufacturing Systems*, vol. 220, W. Shen, Ed. Berlin, Germany: Springer-Verlag, 2006, pp. 461–468.
[9] D. K. Pai, K. van den Doel, D. L. James, J. Lang, J. E. Lloyd, J. L. Richmond, and S. H. Yau, "Scanning physical interaction behavior of 3D objects," in *Proc. ACM Siggraph*, Aug. 2001, pp. 87–96.
[10] J. Lang, D. K. Pai, and R. J. Woodham, "Acquisition of elastic models for interactive simulation," *Int. J. Robot. Res.*, vol. 21, no. 8, pp. 713–733, 2002.
[11] A.-M. Cretu, E. M. Petriu, and P. Payeur, "Neural network mapping and clustering of elastic behavior from tactile and range imaging for virtualized reality applications," *IEEE Trans. Instrum. Meas.*, vol. 57, no. 9, pp. 1918–1928, Sep. 2008.
[12] A.-M. Cretu and E. M. Petriu, "Neural network-based adaptive sampling of 3D object surface elastic properties," *IEEE Trans. Instrum. Meas.*, vol. 55, no. 2, pp. 483–492, Apr. 2006.
[13] T. M. Martinetz, S. G. Berkovich, and K. J. Schulten, "Neural-gas network for vector quantization and its application to time-series prediction," *IEEE Trans. Neural Netw.*, vol. 4, no. 4, pp. 558–568, Jul. 1993.
[14] B. Fritzke, *Some Competitive Learning Methods*, 1997. [Online]. Available: http://www.neuroinformatik.ruhr-uni-bochum.de/ini/VDM/research/gsn/JavaPaper/
[15] P. Curtis, C. S. Yang, and P. Payeur, "An integrated robotic multi-modal range sensing system," in *Proc. IEEE Int. Instrum. Meas. Technol. Conf.*, Ottawa, ON, Canada, May 2005, pp. 1991–1996.

Ana-Maria Cretu (S'04) received the Master's degree from the University of Ottawa, Ottawa, ON, Canada, where she is currently working toward the Ph.D. degree with the Sensing and Modeling Research Laboratory, School of Information Technology and Engineering.

She is the author of more than 25 technical papers. She served as a Reviewer for three journals. Her research interests include neural networks, tactile sensing, 3-D object modeling, and multisensor data fusion.

Dr. Cretu is a Student Member of the IEEE Instrumentation and Measurement Society and the IEEE Computational Intelligence Society. She served as Technical Committee Member for five international conferences.

Pierre Payeur (S'90–M'98) received the Ph.D. degree in electrical engineering from the Université Laval, Quebec, QC, Canada, in 1999.

Since 1998, he has been with the School of Information Technology and Engineering, University of Ottawa, Ottawa, ON, Canada, where he is currently an Associate Professor and the Director of the Sensing and Modeling for Automation and Robotic Intelligence Research Group. He is also a founding member of the Vision, Imaging, Video and Autonomous Systems Research Laboratory and a member of the Sensing and Modeling Research Laboratory. He regularly serves as an external consultant for companies in industrial automation, sensing, and robotic applications. He is the author or a coauthor more than 70 technical papers and has supervised 40 researchers. He also regularly serves as a Reviewer for several journals and transactions in the fields of instrumentation and measurement, computer vision, and robotics. His research interests include machine vision, motion capture, probabilistic 3-D modeling, range data processing, tactile sensing, robot guidance, and teleoperation.

Dr. Payeur is a member of the IEEE Robotics and Automation Society, the IEEE Instrumentation and Measurement Society, and the Ordre des Ingénieurs du Québec. He served as the General Chair for the IEEE International Workshop on Robotics and Sensors Environments (ROSE) in 2007 and 2008, as well as the Technical Program Chair and a Committee Member on more than 35 international conferences on instrumentation and measurement, computer vision, and robotics.

Emil M. Petriu (M'86–SM'88–F'01) received the Dipl.Eng. and Dr.Eng. degrees from the Polytechnic Institute of Timisoara, Timisoara, Romania.

Since 1985, he has been with the University of Ottawa, Ottawa, ON, Canada, where he is currently a Professor and the University Research Chair with the Sensing and Modeling Research Laboratory, School of Information Technology and Engineering. He is the author or a coauthor of more than 300 technical papers and two books. He is the Editor of two books. He is the holder of two patents. His research interests include robot sensing and perception, interactive virtual environments, human–computer symbiosis, soft computing, and digital integrated circuit testing.

He is a Fellow of the Canadian Academy of Engineering and the Engineering Institute of Canada. He is a corecipient of the 2003 IEEE's Donald G. Fink Prize Paper Award and a recipient of the 2003 IEEE Instrumentation and Measurement Society Award.