Real-time Image Processing manuscript No.
(will be inserted by the editor)

Sergio Orts-Escolano · Vicente Morell · Jose Garcia-Rodriguez · Miguel Cazorla · Robert B. Fisher

Real-time 3D semi-local surface patch extraction using GPGPU

Application to 3D object recognition

Received: - / Revised: -

Abstract Feature vectors can be anything from simple surface normals to more complex feature descriptors. Feature extraction is important in order to solve various computer vision problems: e.g. registration, object recognition and scene understanding. Most of these techniques cannot be computed online due to their complexity and the context where they are applied. Therefore computing these features in real-time for many points in the scene is impossible. In this work a hardware-based implementation of 3D feature extraction and 3D object recognition is proposed in order to accelerate these methods and therefore the entire pipeline of RGBD based computer vision systems where such features are typically used. The use of a GPU as a General Purpose processor (GPGPU) can achieve considerable speed-ups compared with a CPU implementation. In this work advantageous results are obtained using the GPU to accelerate the computation of a 3D descriptor based on the calculation of 3D semi-local surface patches of partial views. This allows descriptor computation at several points of a scene in real-time. Benefits of the accelerated descriptor have been demonstrated in object recognition tasks. Source code will be made publicly available as contribution to the Open Source Point Cloud Library (PCL).

Keywords Real-time · GPGPU · RGBD-data · 3D local shape descriptor · Object recognition.

Sergio Orts-Escolano · Vicente Morell · Jose Garcia-Rodriguez · Miguel Cazorla
Computer Technology Department
University of Alicante
E-mail: sorts@dtic.ua.es

Robert B. Fisher
Institute of Perception, Action and Behaviour
University of Edinburgh
E-mail: rbf@inf.ed.ac.uk

1 Introduction

In recent years, the number of jobs concerned with 3D data processing has increased considerably due to the emergence of cheap 3D sensors capable of providing a real time data stream. The Kinect device[1] and 3D laser scanners are examples of these devices. Besides providing 3D information, these devices can also provide color information of the observed scene. The availability of real-time 3D streams has provided a key resource for solving many challenging problems. However, the computing power of the CPU is still an obvious bottleneck when running systems in real time. Additionally, as these systems are becoming more complex, huge computational resources are demanded, especially when hard real-time constraints are required.

With the advent of the GPU as a General Purpose Graphic Processing Unit (GPGPU) many methods related to 3D data processing have been implemented on GPUs in order to accelerate them. Examples of this can be found in the calculation of feature descriptors and 3D keypoint extraction on the GPU. In [5], the GPU performs curvature estimation of 3D meshes in real time. A parallel implementation using the GPU and the CUDA language from NVIDIA considerably accelerates the Histograms of Oriented Gradient (HOG) computation [20]. In [8], a Point Feature Histograms (PFH) GPU implementation is proposed allowing its computation in real-time on large point clouds. Despite the mentioned methods above, which demonstrate the feasibility of the GPU for 3D feature extraction, there are still comparatively few methods implemented with respect to all those currently prevalent in the state of the art. Furthermore, more complex descriptors have not been implemented yet on the GPU. One example of this is a descriptor based on the extraction of 3D surface patches representing the underlying model. It can also be noted that the integration of these methods in complete systems, that

[1] Kinect for XBox 360: http://www.xbox.com/kinect Microsoft

require real-time constraints, is still very low. Kinect Fusion [10] has been one of the first works where the GPU has been used as the main core processor, allowing the reconstruction of 3D scenes in real time.

This work goes a step further in the application of massively parallel processor architectures, such as the GPU, to 3D data processing tasks and their integration to complex 3D vision systems with real-time constraints. 3D data has been obtained mainly as RGB-Depth maps provided by the Kinect sensor. Advantageous results have been obtained using the GPU for accelerating the extraction of a feature descriptor based on the 3D calculation of semi-local surfaces patches. This descriptor is computationally expensive to compute on the CPU and also requires some preprocessing steps for its computation: normals estimation, surface triangulation and keypoint detection. Since these preprocessing steps are common for a large number of applications their acceleration and integration in the pipeline of a GPU architecture becomes essential to make progress in 3D data real-time processing.

Other motivations for this work include the existing gap of 3D object recognition solutions based on models that support real-time constraints; until now most of the proposed works that supported real-time constraints are view-based. For example in [1], a local feature descriptor for RGB-D images is proposed. This descriptor combines color and depth information into one representation. However, 3D information possibilities are still negligible, it depends mostly on textures and illumination of the specific scene. In [15], a combined descriptor formed by 3D geometrical shape and texture information is used to identify objects and its pose in real-time. The proposed system is accelerated using GPU achieving real-time processing. However, 3D information is only used to extend shape information and considerably relies on texture information, making it sensitive to scene illumination conditions. Model-based approaches are less sensitive to illumination, shadows, and occlusions of the scene allowing more robust object recognition systems improving also its pose estimation. Until now, works that make use of 3D local shape descriptors like SPLASH [23], Spin Images [11], Spherical Spin Image [21], surface patch representation [3], 3D tensor-based representation [16], 3D SURF [14], etc. have rarely been processed in real-time is also useful and have not been integrated in complex computer vision systems. These 3D model-based descriptors are the most popular techniques applied to free-form object recognition but in this work we are going to focus on 3D tensor-based representation [17] as it has demonstrated good performance in free-form object recognition even under significant background clutter and it outperforms other state-of-the-art descriptors as Spin Images [17]. 3D tensor-based computation requires some preprocessing steps that are common to other descriptors therefore being able to process those in real-time. The proposed gpu-based implementation may be extended to other 3D local shape descriptors.

The paper is structured as follows: in Section 2 the proposed descriptor to be accelerated is presented. In the following Subsections 2.1 and 2.2 the GPU implementation of pre-processing steps are detailed. In Section 2.3 a GPU-based tensor extraction implementation is described. Next, in Section 3, performance results are shown for different steps and hardware configurations. Finally, Section 4 presents a real application in object recognition tasks with time constraints to validate our implementation, followed by our main conclusions and future work.

2 GPU-based tensor extraction algorithm

The feature descriptor proposed for implementation on the GPU is based on the descriptor introduced in [16]. This descriptor is based on the calculation of semi-local surface patches obtained by a range camera. It has been used successfully for different applications like global registration and 3D object recognition, where other descriptors like Spin Images [12] or Geometric Histograms [7] obtained worse results. Additionally, in [17] it is demonstrated how this descriptor can be successfully applied in object recognition problems with high level of occlusion. The main problem of this descriptor is its high computational cost, which is prohibitive for running in a conventional CPU under real-time constraints.

This descriptor extracts a semi-local model of the scene, computing semi-local features that assist the local object recognition even under conditions of occlusion. This feature is referenced in the following sections as a tensor. A tensor is defined by the surface of the model that is intersected by each voxel on a centered grid. This set of values define a third order tensor.

To compute the required descriptor some preprocessing steps are necessary. These steps have also been implemented on the GPU in order to accelerate the entire pipeline. Therefore it has been necessary to implement on the GPU the following processes: depth map and color map transformation to a coloured point cloud, noise removal, normal estimation and surface reconstruction. A general system overview for semi-local surface patch extraction is shown in Figure 1. Moreover, in Figure 2 it can be seen the different steps required prior to the extraction of the tensor and their 3D visualization after each step. As we can appreciate, most steps are computed on the GPU taking advantage of parallel computing power of the GPU and avoiding transfers between CPU and GPU after each step. Pseudo-code of the entire GPU-based tensor extraction algorithm is presented in Algorithm 1. Moreover, pseudo-code snippets for all GPU-based preprocessing steps are presented in next sections.

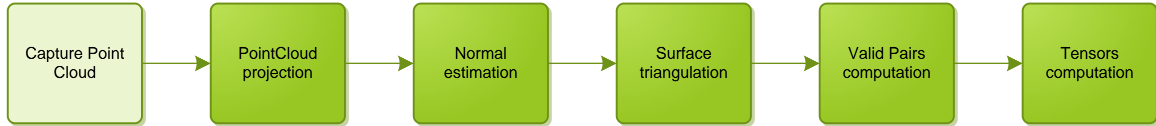

Fig. 1 General system overview. Steps coloured in dark are computed on the GPU.

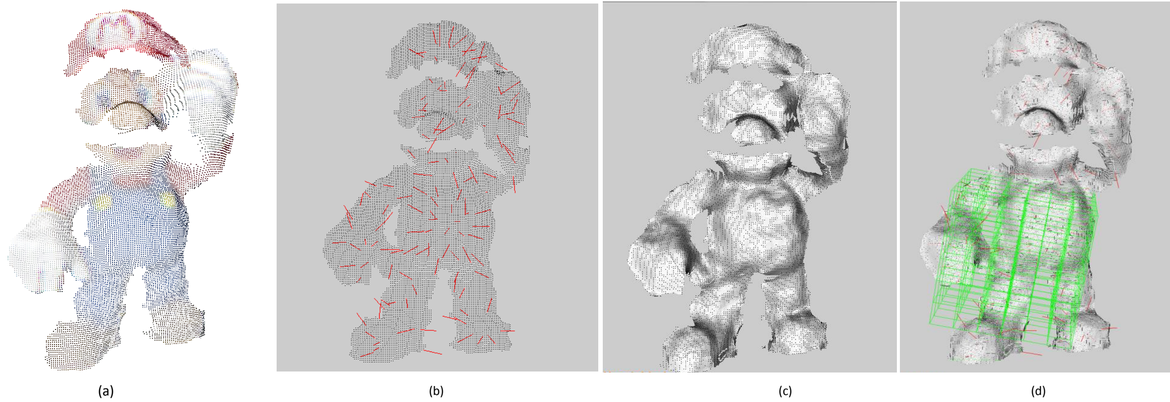

Fig. 2 (a) Point cloud obtained after transforming depth and color maps provided by the Kinect sensor. (b) Normal estimation. (c) Surface reconstruction. (d) Feature descriptor extraction: 3D tensor computed over a partial view

input : A depth map M_d of size 640×480
output: A set of 3D tensors $T = \{t_0, t_1, t_2, ..., t_N\}$ that describe the input data

1 *Depth map is transferred to the GPU memory*;
2 cudaMemCopyHostToDevice(d_M_d, M_d);
3 $d_M_{filtered} \leftarrow$ gpuBilateralFiltering(d_M_d);
4 $d_P_{xyz} \leftarrow$ gpuPointCloudProjection($d_M_{filtered}$);
5 $d_N_{xyz} \leftarrow$ gpuNormalEstimation(d_P_{xyz});
6 $d_Tri \leftarrow$ gpuSurfaceTriangulation(d_P_{xyz}, d_N_{xyz});
7 $d_V_{pairs} \leftarrow$ gpuValidPairsComp(d_P_{xyz}, d_N_{xyz});
8 cudaMemCopyDeviceToHost(V_{pairs}, d_V_{pairs});
9 **for** $i \leftarrow 0$ **to** $|V_{pairs}|$ **do**
10   *tensors are computed in parallel at thread level and tensor level*;
11   AsyncGpuTensorComp(v_i,d_P_{xyz},d_N_{xyz},d_Tri, d_T);
12 **end**
13 cudaMemCopyDeviceToHost(T, d_T);

Algorithm 1: Pseudo-code of the GPU-based 3D tensor extraction algorithm. $d_$ prefix means that variable is allocated in the GPU memory.

2.1 RGB-D processing on the GPU

In this work we focus on the processing of 3D information provided by the Kinect sensor. This processing is performed on the GPU with the aim of achieving a real-time implementation.

The overall goal is the implementation of systems that offer interaction with the user. Kinect sensor provides a RGB map with color information M_c and a disparity map M_d. The first step to carry out on the GPU with the aim of accelerating future steps is the projection of the depth and color information in a three-dimensional space, where the depth and colour information is aligned, allowing the production of a coloured point cloud that represents the scene, Figure 3.

The relationship between a disparity map provided by the Kinect sensor and a normalized disparity map is given by $d = 1/8 \cdot (d_{off} - kd)$, where d is the normalized disparity, kd is the disparity provided by the Kinect and d_{off} is a particular offset of a Kinect sensor. Calibration values can be obtained in the calibration step [13]. In this way the relationship between depth and a disparity map is given by the following equation:

$$z = \frac{b \cdot f}{1/8 \cdot (d_{off} - kd)} \tag{1}$$

where b is the baseline between the infrared camera and the RGB camera (in meters), and f is the focal distance of the cameras (in pixels). Once the depth map M_d is

obtained calculating the depth z for all points, the projection of each point in 3D space is given by:

$$
\begin{aligned}
p_x &= z \cdot (x - x_c) \cdot 1/f_x \\
p_y &= z \cdot (y - y_c) \cdot 1/f_y \\
p_z &= z
\end{aligned} \quad (2)
$$

where $p \in R^3$, x and y are the row and the column of the projected pixel, x_c and y_c are the distances (in pixels) to the map centre and f_x and f_y are the focal distances of the Kinect sensor obtained during the calibration [13].

This transformation can be computed independently for each pixel of the map, so it fits perfectly on massively parallel architectures such as the GPU, accelerating processing time related to the CPU implementation. As this transformation is often followed by other processing steps, it is not necessary to copy data back to the CPU memory and we therefore avoid the latency caused by these transfers by storing the projected 3D points on the GPU memory. Pseudo-code of the kernel executed onto the GPU is shown in Algorithm 2. In section 3.1 we show the acceleration factor and time of execution obtained by the GPU implementation. All these methods are developed in C++. GPU programming is done using the CUDA language created by NVIDIA [18]. Finally, 3D data management (data structures) and their visualization is done using the PCL[2] library.

input : A depth map M_d of size 640×480
output: Projected point cloud P_{xyz} into 3D space

```
1  __global__ void;
2  gpuPointCloudProjectionKernel( Md );
3  {
4      This kernel is executed creating one thread for each
       pixel in parallel;
5      int u = threadIdx.x + blockIdx.x * blockDim.x;
6      int v = threadIdx.y + blockIdx.y * blockDim.y;
7      float z = Md [v][u] / 1000.f;
8      float px = z * (u - cx) * fx_inv;
9      float py = z * (v - cy) * fy_inv;
10     float pz = z;
11 }
```

Algorithm 2: Pseudo-code of the GPU-based point cloud projection algorithm.

2.1.1 Noise removal: Bilateral filtering

In structured light imaging a predefined light pattern is projected onto an object and simultaneously observed by a camera. The appearance of the light pattern in a certain region of the camera image varies with the camera-object distance. This effect is used to generate a distance image of the acquired scene. The predefined light patterns can be e.g. gray codes, sine waves, or speckle patterns. Speckle patterns are used in popular structured light (infrared) cameras like the Microsoft Kinect. This method of obtaining 3D information from the scene presents problems when the surfaces have a high level of specularity (reflection of the incident light) making it impossible for the sensor to obtain depth information about some surfaces [27]. The same problem occurs in the case of objects that are very far away from the sensor. Therefore, if we want to extract coherent information of the observed surfaces it is necessary to minimize this observation error. In previous works, simple filters such as the mean or the median have been used as they correct the error and run in real-time. As the computing power of the GPU can be applied in this step, this allows the application of more complex filters that are able to reduce the depth map error without removing important information, such as edge information. An example of these filters, the Bilateral filter [24], is able to remove noise of the image whilst preserving edge information. This filter was used originally in color and grey scale images to reduce the noise while keeping edge information, but we can also use it to reduce the noise on depth maps obtained from 3D sensors like the Kinect. A Bilateral filter is a combination of a domain kernel, which gives priority to pixels that are close to the target pixel in the image plane, with a range kernel, which gives priority to the pixels which have similar values as the target pixel. This filter is often useful when it is necessary to preserve edge information because of the range kernel advantages. The new value of a filtered pixel is given by:

$$P_f = \frac{1}{K_p} \sum_{q \in \omega} V_q f(||p - q||) g(||V_p - V_q||) \quad (3)$$

where K_p is a normalization factor, ω is the neighbourhood of the target pixel, V_q is the target pixel value and P_f is the filtered value of pixel p. This equation also contains the domain kernel and the range kernel: $f(||p-q||)$, $g(||V_p - V_q||)$. Often, f and g are Gaussian functions with standard deviation σ_s and σ_r.

Section 2.1.2 shows how the estimation of the normal vectors is improved after applying bilateral filtering, providing more stable normal vectors by removing original noise presented on the depth map. Most 3D features extracted from the scene are based on the curvature of the geometry, which is calculated using information from normal vectors at each point in the scene, therefore obtaining more stable normal vectors leads to more accurate scene knowledge.

The calculation of filtered values at each pixel of the image can be calculated independently and therefore is well suited for parallel architectures like the GPU. In [2] and [26] GPU implementations able to run in real time were proposed. The runtime is considerably improved, allowing filtering in real time depth maps generated by the sensor. In Section 3 GPU and CPU runtimes and

[2] The Point Cloud Library (or PCL) is a large scale, open project [22] for 2D/3D image and point cloud processing.

Fig. 3 Left: Depth map. Center: RGB Map. Right: Projected Point cloud.

speed-ups for our implementation on different graphics boards are presented.

2.1.2 Normal estimation

Estimation of normal vectors on a geometric surface has been widely used in many application areas such as computer graphics: generating realistic illumination of the surfaces and in computer vision as geometric features of the observed environment: keypoints with high curvature (corner or edge points).

Given a geometric surface it is possible to estimate the direction of the normal vector at a point obtaining the outward facing vector of the surface. However, we still have a point cloud without information about the surfaces that compose it, therefore we approach the normal vector estimation of a point efficiently by using its neighbourhood to calculate the normal vector. Here we focus on the estimation of the plane that best fits the neighbourhood of points using least squares. In this way the point normal vector is calculated as the plane normal vector.

The search for this plane is reduced to the calculation of eigenvalues and eigenvectors, Principal Component Analysis (PCA) of the covariance matrix created using the neighbourhood of the point on which we want to know its normal vector. The orientation of the normal vector is easily calculated because we know the point of view of the scene, in this case the Kinect position, so that all normal vectors must be facing consistently toward the point of view satisfying the following equation:

$$\mathbf{n_i} \cdot (v_p - p_i) > 0 \tag{4}$$

where $\mathbf{n_i}$ is the calculated normal vector, v_p is the point of view, and p_i is the target point. In cases where this constraint is not satisfied, it is necessary to reverse the sign of the calculated normal vector.

Once we have the organized point cloud stored in the memory of the GPU, the normal estimation process using PCA can be performed efficiently on the GPU. The normal vector calculation is performed on the GPU independently at each point of the scene, considerably accelerating the runtime. Pseudo-code of the GPU-based normal estimation algorithm is shown in Algorithm 3. Moreover, thanks to the previous noise removal using bilateral filtering, normal vectors obtained are much more stable than normal vectors computed directly from the original depth map that does not take into account the borders and corner points of the scene. In Figure 4 we can see this effect.

input : A projected point cloud d_P_{xyz}
output: Point cloud of normals d_N_{xyz}

1. __global__ void;
2. gpuNormalEstimationKernel(P_{xyz}, k);
3. {
4. *This kernel is executed creating one thread for each point in parallel*;
5. int u = threadIdx.x + blockIdx.x * blockDim.x;
6. int v = threadIdx.y + blockIdx.y * blockDim.y;
7. *Compute Covariance matrix centered at point p using k neighbours*;
8. d_N_{xyz}[u][v] = compCovarianceMat(u,v,k,N);
9. d_N_{xyz}[u][v] = checkOrientation();
10. }

Algorithm 3: Pseudo-code of the GPU-based normal estimation algorithm

2.2 Surface triangulation on the GPU

In [9] an efficient method to triangulate organised point clouds is presented. In this section we present an accelerated and robust implementation of this method. In the original work a triangulation method for 3D points, obtained from range cameras or structured light, is proposed. Using sensors such as the Kinect, 3D points can be accessed using their matrix organization using x for the row and y for the column. In this way, a 3D point $p_{x,y}$ can be accessed using a 2D indexing system. Using this representation it is possible to obtain the scene surface from the point cloud captured by the sensor. The method assumes that the viewpoint is known and in this way it is possible to calculate the angle formed by the viewpoint vector v_p and the target point $p_{x,y}$ and the vector formed by the target point $p_{x,y}$ and its neighbour points $p_{x+1,y}$ or $p_{x,y+1}$. If points fall into a common

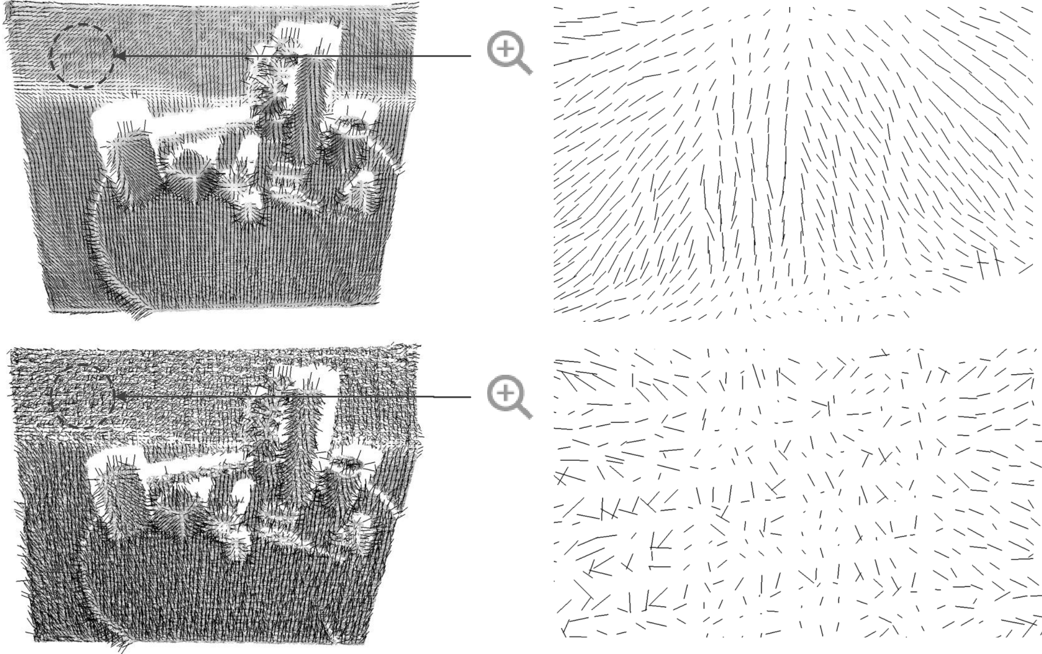

Fig. 4 Bottom row. Normal estimation using the original map. Top row. Normal estimation using the filtered map (bilateral filtering). It can be observed that the normal estimation is improved resulting in more stable normal directions. This effect can be visually observed on plane surfaces where the normals estimated using a noisy map are much less stable than normals computed over a filtered map.

line of sight with the viewpoint from where the measurements are taken, one of the underlying surfaces occludes the other. If all checks are passed the triangle is added to the mesh, otherwise a hole arises in the final reconstructed mesh. Moreover, if the sensor cannot acquire a valid depth measurement for a certain pixel that triangle is also rejected creating a hole. This Figure 5 visually shows the proposed condition for point triangulation.

Our proposed method is more robust than the original method, as the normal information at each point of the scene is used as an additional condition for meshing the point cloud. As the triangulation of the points can be done independently, the algorithm has been ported to the GPU, where each GPU thread tests the point we are targeting to form a triangle with its neighbourhood. As a result, we obtain a vector with all the triangles. Pseudocode of the GPU-based surface triangulation algorithm is shown in Algorithm 4.

Invalid triangles are created on points that do not satisfy the proposed constraint to keep the organization of the point cloud. Finally, the condition to create an edge between two points is formulated as follows:

$$edge_{valid} = (|v_{v_p,p} \cdot v_{p,q}| \leq \cos \epsilon_{\theta_{pov}}) \wedge$$
$$(||p-q||^2 \leq T_d) \wedge$$
$$(|n_p \cdot n_q| \leq \cos \epsilon_{\theta_n}) \tag{5}$$

where $\epsilon_{\theta_{pov}}$ is the angle existing between two points and the point of view establishing whether or not these points are occluded. This angle value is established based on the visual analysis showed in the Figure 5 and also based on results provided in [9]. The maximum distance between two points is T_d. This distance is obtained in real-time based on point cloud resolution. For that, the average distance between the targeted point and its neighbourhood k is calculated. Next, based on the mean of these average distances and standard deviation, threshold T_d is given by: $T_d = \overline{d_k} + \sigma_d$ where $\overline{d_k}$ is the mean distance and σ_d is the standard deviation. Finally, ϵ_{θ_n} is the established threshold for the maximum angle between two normal vectors. This is calculated in the same way as T_d, obtaining an angle threshold.

The proposed method allows us to obtain fast approximate meshing of the input point cloud. This fast meshing method is therefore used in Section 2.3 for computing the proposed 3D semi-local surface patch descriptor in real-time, as it requires a surface representation of the scene. The proposed accelerated meshing method takes advantage of the knowledge about the point of view position and also takes advantage of having already calcu-

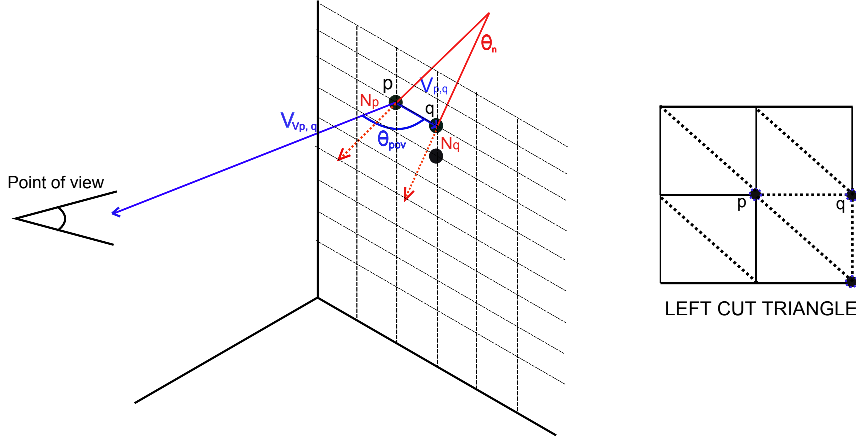

Fig. 5 Left: Point triangulation condition. In this image can be observed how the condition for creating an edge establishes that the angle θ_{pov} formed by vectors $v_{v_p,p}$ and $v_{p,q}$ must be within an established threshold $\epsilon_{\theta_{pov}}$. This threshold assures that points are not occluded among themselves. The Euclidean distance between p and q also must be smaller than an established threshold T_d, dynamically calculated according to mesh resolution and its standard deviation. Right: Triangles are established by left cut checking constraints between points.

lated normal vectors on GPU memory for every point of the scene. The GPU implementation achieves run times considerably lower compared to the CPU. The GPU implementation achieves processing frame rates close to 30 fps for 640 by 480 depth maps while the CPU implementation achieves a frame rate close to 6 fps. Figure 6 shows a point cloud mesh obtained using the proposed method.

input : A projected point cloud d_P_{xyz}
input : A point cloud of normals d_N_{xyz}
output: List of triangles d_Tri

1 __global__ void;
2 gpuTriangulationKernel(d_P_{xyz}, d_N_{xyz});
3 {
4 *This kernel is executed creating one thread for each point in parallel*;
5 int u = threadIdx.x + blockIdx.x * blockDim.x;
6 int v = threadIdx.y + blockIdx.y * blockDim.y;
7 *check constraints with neighbour points*;
8 if (isValidTriangle (i, index_down, index_right));
9 addTriangle (d_Tri);
10 if (isValidTriangle (index_right, index_down, index_down_right)) ;
11 addTriangle (d_Tri);
12 }

Algorithm 4: Pseudo-code of the GPU-based surface triangulation algorithm.

2.3 Tensor computation on the GPU

Once point cloud normal information and surface triangulation are obtained, pairs of points along with their normals are selected to define local 3D coordinate bases for tensor computation. To avoid the C_2^n combinatorial explosion of the points, a distance constraint is used on their pairing. This distance constraint allows the pairing between only those points that are within a previously specified distance. The distance constraint also ensures that the vertices that are paired are far enough apart so that the calculation of the coordinate bases is not sensitive to noise but close enough to maximize their chances of being inside the same surface. The maximum and minimum distances between points are based on point cloud resolution, being $d_{min} = pcl_{res} * 5$ and $d_{max} = pcl_{res} * 14$. pcl_{res} is calculated for each point cloud captured by the Kinect sensor in real time, allowing the movement of the sensor. In addition to this distance constraint, an angle constraint θ_d is defined between valid pairs of points, so that points with approximately equal normals are not paired (since their cross product will result in zero). This mutual angle must be higher than 15 degrees allowing the use of the mean value of these normals as an axis for the coordinate bases. Each point is paired with only its three closest neighbours, limiting the number of possible pairs to $3n$ per view. In practice, due to the constraints this number is lower than $3n$.

Pair point calculation is accelerated using as many threads as points in the point cloud, in this way each GPU thread checks its corresponding point pair with its neighbours. Moreover, the matrix organization of the point cloud is used for improving this search. In this way, each thread of the GPU performs the search of valid pairs only in a defined window around the targeted point. The size of this window is based on the maximum distance constraint d_{max} and the point cloud resolution pcl_{res}: $windows_{radius} = d_{max}/pcl_{res}$ giving the radius of the windows in pixels. In Section 3.1 a runtime comparison is presented, the CPU implementation applies the same technique for search acceleration.

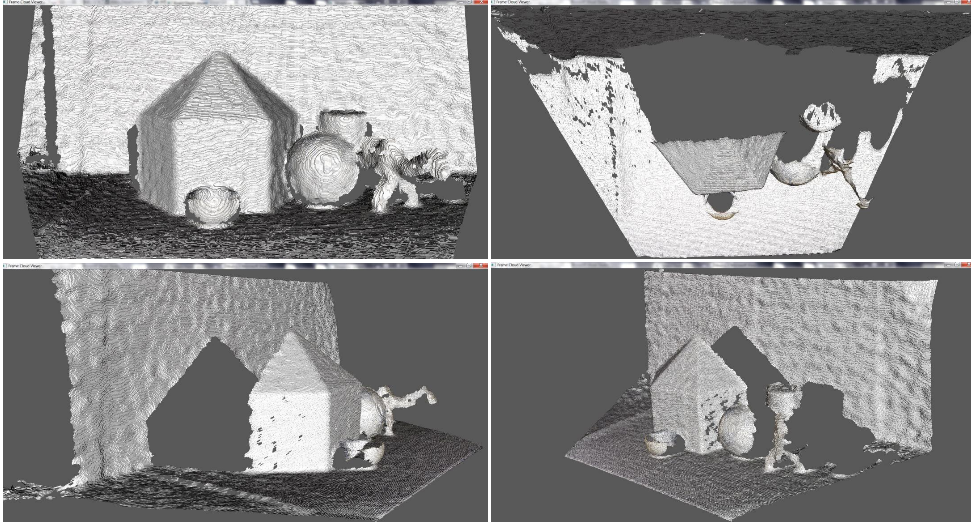

Fig. 6 Point cloud meshing using the proposed method. Note that some holes and gaps still exist in the approximate surface reconstruction due to the noisy information obtained from the Kinect sensor.

Once a valid list of point pairs is obtained, a local 3D basis is defined for each valid pair in the following manner: the center of the line joining the two vertices defines the origin of the new 3D basis. The average of the two normals defines the z-axis. The cross product of the two normals defines the x-axis and finally the cross product of the z-axis with the x-axis defines the y-axis. This 3D basis is used to define a 3D grid centered at its origin. This step is also computed in parallel on the GPU for each valid pair of points.

For the grid computation, which will define the feature descriptor, it is necessary to define two more parameters. The first one is the number of voxels that compose the grid n_{voxels} and the size of each of these voxels $voxel_{size}$. Modifying the number of voxels and so the size of the grid, causes the obtained descriptor to contain local, semi-local or global information of the scene. In the experiments done in [17] it is demonstrated how for the object recognition task, a size of $10 \times 10 \times 10$ grid allows the extraction of a descriptor with semi-local information of the object allowing identification even under a high level of occlusion. The size of the voxel $voxel_{size}$ is defined dynamically according to the point cloud resolution. Once the grid is defined, the surface area of the mesh surface intersecting each voxel of the grid is stored in a third order tensor. This tensor is a local surface descriptor which corresponds to a semi-local representation of the object where the pair of points are lying. Sutherland Hodgman's polygon clipping algorithm [4] is used for calculating area intersections between polygons and voxels. In this way an entry is made at the corresponding element position in the tensor. Since more than one triangulated facet can intersect a single voxel, the calculated area of intersection is added to the area already present in that voxel as a result of intersection with another triangulated facet. To avoid checking all triangles that compose the scene, a growing approach is used, which starts by checking the triangles that lie in the pair of points selected and growing along its neighbourhood until all the checked triangles are not intersected with the corresponding voxel. This approach is used in both CPU and GPU versions, allowing a fair runtime comparison. Finally extracted tensors are compressed by squeezing out the zero elements and retaining the non-zero values and their index positions in the tensor. These compressed tensors together with their respective coordinate basis and the mutual angle between their normals are called a tensor representation of the view.

The computation of each tensor is considerably accelerated using the GPU because there is no dependency between the calculation of the intersected area in each voxel of the grid. Therefore, $Dim_x \times Dim_y \times Dim_z$ threads are executed on the GPU organized as a three dimensional grid. Each thread calculates the intersected area between the mesh and its corresponding voxel, storing the calculated area in the position accessed by its indexes. See Figure 7. Due to the 3D index organization that the CUDA framework provides, the calculation of corresponding indexes is greatly accelerated. Sutherland Hodgman's polygon clipping algorithm is also executed by each thread in parallel. Pseudo-code of the GPU-based 3D tensor computation algorithm is shown in Algorithm 5. Additionally, there is also no dependency between the computation of different tensors, thereby the computation of different tensors is overlapped occupy-

ing all the available resources on the GPU. Performance results are shown in Section 3.1.

```
input : A projected point cloud d_P_xyz
input : A valid pair of points d_N_xyz
input : List of triangles d_Tri
output: 3D tensor t_i
1  __global__ void;
2  gpuTensorCompKernel( v_i,d_P_xyz,d_N_xyz,d_Tri );
3  {
4      This kernel is executed creating one thread for each
       bin of the grid in parallel;
5      int x = threadIdx.x + blockIdx.x * blockDim.x ;
6      int y = threadIdx.y + blockIdx.y * blockDim.y ;
7      int z = threadIdx.z + blockIdx.z * blockDim.z ;
8      binLimits = computeBinLimits(cloud,tri) ;
9      d_Neigh_Tri = compIndexNeighTriangles() ;
10     calculate area that clip with the corresponding bin;
11     for i <- 0 to |d_Neigh_Tri| do
12         area += clipTriangle(cloud,tri,binLimits) ;
13     end
14     t_i[x][y][z] = area ;
15 }
```

Algorithm 5: Pseudo-code of the GPU-based 3D tensor computation algorithm

All extracted tensors from a partial view of an object are stored with their coordinate basis allowing the use of this information for grouping all tensors with similar angle between their normals. In this way an efficient matching is possible for different applications such as partial view registration and object recognition. This collection of tensors is stored during a training phase creating a hash table for efficient retrieval during test phase.

3 Experimental results

GPU versions of the proposed method described in this document has been tested on a desktop machine with an Intel Core i3 540 3.07Ghz and different CUDA capable devices. GPU implementations were first developed on a laptop machine equipped with an Intel Core i5 3210M 2.5 Ghz and a CUDA compatible GPU. Table 1 shows different models that have been used and their main features. We used different models ranging from the integrated GPU on a laptop to a more advanced model, demonstrating that the GPU implementations can be executed on different GPUs and that they can obtain good execution times on different graphic boards with different number of cores.

GPUs are ideally suited to executing data-parallel algorithms. Data-parallel algorithms execute identical units of work (programs) over large sets of data. The algorithms can be parallelized for efficiency when the work units are independent and are able to run on small divisions on the data. One critical aspect of designing parallel algorithms is identifying the units of work and determining how they will interact via communication and synchronization. A second critical aspect is analyzing the data access patterns of the programs and ensuring data locality to the processing units. It is also necessary to consider the program execution pipeline in order to avoid unnecessary data transfers,

These three critical aspects have been satisfied by our GPU implementations because every step was decomposed as an independent execution unit. This is possible since there are no dependencies during their computation. Moreover, all the computed data on the GPU is not transferred back to the CPU until the entire pipeline is completed, thus avoiding expensive memory transfers. Finally, threads access memory using data-patterns in order to ensure locality to the processing units. Most steps implemented on the GPU use a 2D map of threads accessing memory in a coalesced way. For the tensor computation a 3D grid of threads is used.

3.1 Performance

The performance obtained by the GPU implementation allow us to compute the proposed methods under real-time constraints. In Table 2 we can see the different steps that have been accelerated using the GPU and their different runtime and the speed-ups achieved for the different graphics boards. The obtained acceleration is relative to a CPU implementation of the proposed method. In general the best performance was obtained with the graphics board with the largest number of CUDA cores (GTX480) and the largest memory bandwidth.

These results demonstrate how the proposed methods are suitable for massively parallel architectures such as the GPU, where each thread processes one of the points of the scene. Another interesting aspect of the results shown in Table 2 is that GPU implementations allows us to compute operations that are prohibitively slow on the CPU in real-time such as normal estimation, keypoint detection or surface triangulation. Moreover, in table 2 it is shown how the entire computation of 200 tensors in the GPU is performed in less than 0.5 seconds for the faster device achieving a 93x performance boost related to the CPU implementation and allowing the computation of the descriptor at different points of a scene in real-time.

Another remarkable aspect of the performance obtained for the overall system is that tensor computation is not only parallelized at thread level, it is also parallelized at task level computing simultaneously different tensors. As tensor computation is not dependent, it can be parallelized using different CUDA streams on the GPU. This technology allows executing in parallel as many kernels as possible in different queues and therefore allows to exploit available resources on the GPU [25]. We decided to exploit the possibility of launching multiple kernels concurrently using CUDA streams, overlapping

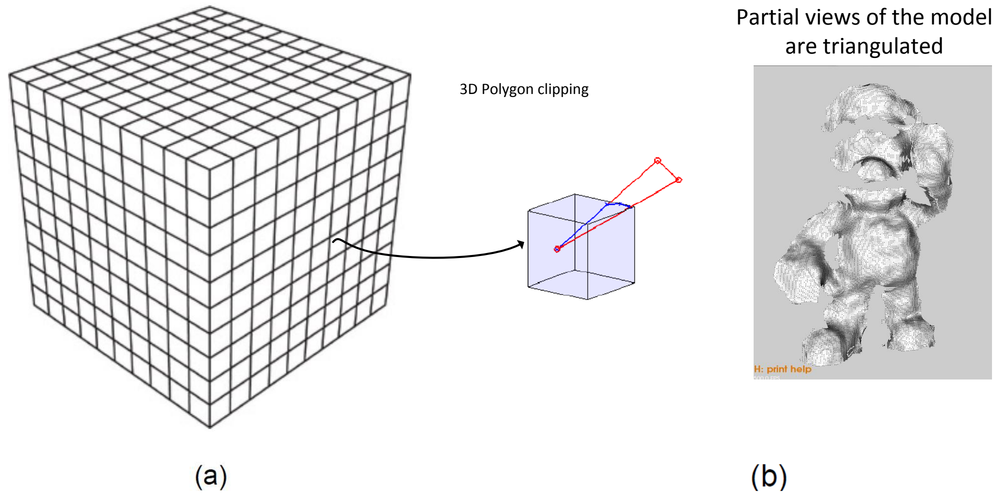

Fig. 7 (a) Launching $Dim_x \times Dim_y \times Dim_z$ threads in parallel where each GPU thread represent a voxel of the grid. (b) Each thread with indexes i, j, k calculates the area of intersection between the mesh and its corresponding voxel using Sutherland Hodgman's polygon clipping algorithm. Taking advantage of thread indexes, the calculated area is stored in a flattened vector.

Device Model	CUDA cores	Global Mem	Bandwidth Mem
Quadro 2k	192	1 GB	41.6 GB/s
GeForce GTX 480	480	1.5 GB	177.4 GB/s
GeForce GT630M	96	1 GB	32 GB/s

Table 1 CUDA capable devices used in experiments

Step	GT630M	GTX480	Q2k	CPU	GT630M	GT480	Q2k
Bilateral filtering of depth map	11ms	**5ms**	8ms	1008 ms	91.63x	**201.6x**	126x
Point cloud projection	2ms	**1ms**	1ms	50ms	25x	**50x**	50x
Normal estimation	9ms	**1ms**	8ms	190ms	21.11x	**190x**	23.75x
Compute surface triangulation	5ms	**2ms**	4ms	121ms	24.25x	**40.3x**	30.25x
Compute cloud resolution	7ms	**4ms**	6ms	330ms	47.14x	**82.5x**	55x
Compute valid pairs	71ms	**9ms**	35ms	4479ms	63x	**497x**	127.97x
Compute third order tensor	6ms	3ms	4ms	130 ms	31.6x	**43.33x**	32.5x
Total GPU time for extracting 200 tensors	854 ms	**490ms**	687ms	45887ms	53.72x	**93.64x**	66.79x

Table 2 Runtime comparison and speed-up obtained for proposed methods using different graphics boards. The fastest run times were achieved by the graphics board NVIDIA GTX480. Runtimes are averaged over 50 runs.

the paradigm of task parallelism to that of data parallelism. In order to analyse and confirm stream parallel execution we profiled the algorithm using the NVIDIA Visual Profiler [19], which allows to visually appreciate how stream computation is performed along the time and also multiprocessors occupancy on the GPU. In Figure 8 it can be seen algorithm computation timeline using and not using streams to overlap computations. Runtime execution and multiprocessors occupancy is greatly improved thanks to concurrent kernel execution using streams. Runtime is improved by a speed-up factor of 5x overlapping tensors computation with many kernels enqueued in different streams and launched concurrently.

We fixed the number of streams to 16 after testing different number of streams. In Figure 9 it can be seen how the runtime is improved as the number of streams is increased obtaining maximum performance and occupancy on the GPU using values larger than 4. Indeed, in Figure 8 it is shown how the maximum number of tensors that are calculated simultaneously is 8 without taking into consideration the maximum number of streams specified. This occurs due to the occupancy of all the available resources by the kernels running concurrently.

Finally, in Figure 10 an experiment computing different number of tensors is performed and the speed-up compared to the CPU version is presented. From Figure 10 we can conclude that the speed-up obtained by the GPU version is increased as the number of tensors is also increased achieving a larger speed-up factor. Computing

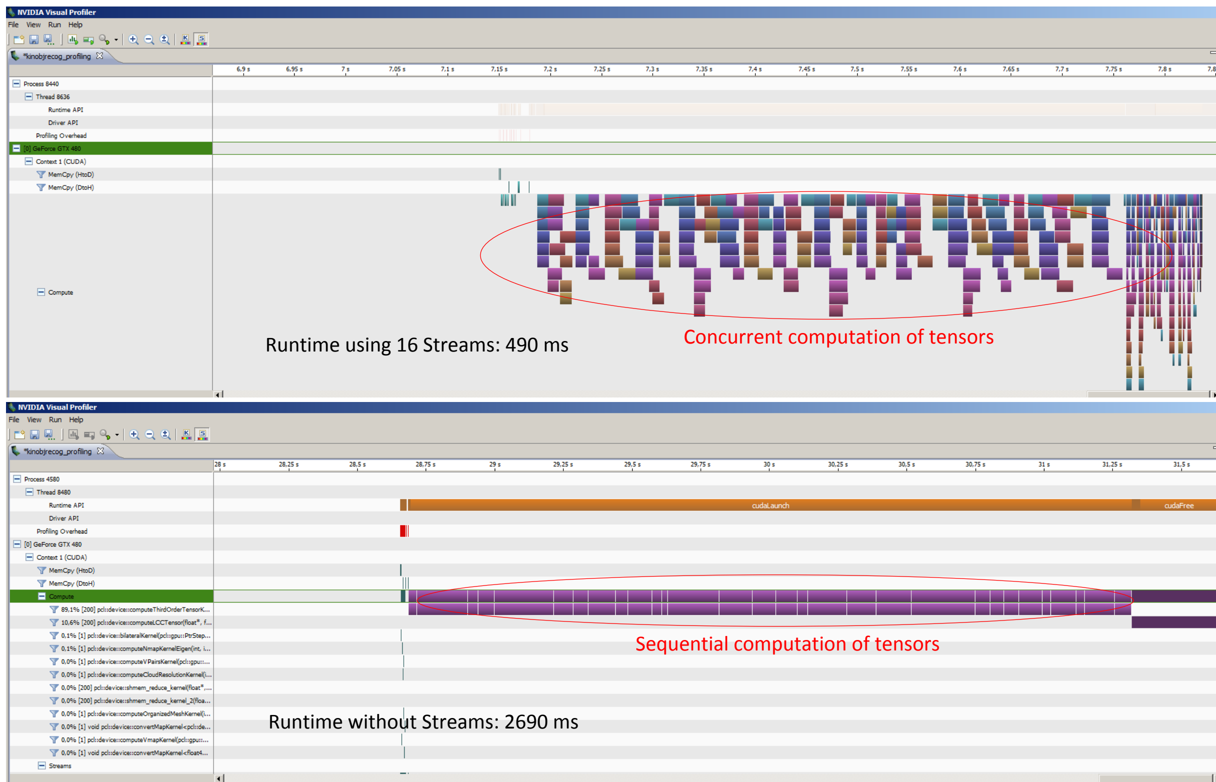

Fig. 8 Profiling computation of tensors using streams (top) and without streams (bottom). On the top of the figure (using streams) it can be seen how up to 6 kernels run simultaneously occupying all available resources on the GPU.

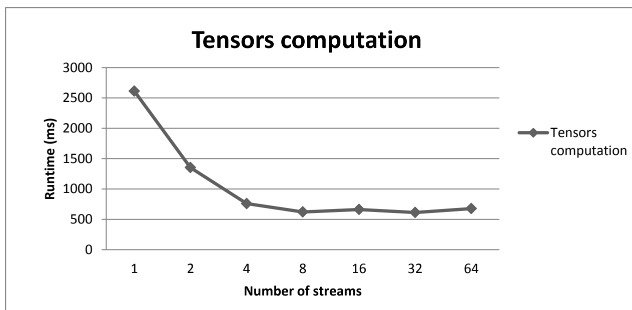

Fig. 9 Tensors computation runtime using different number of streams. Number of tensors is fixed to 200 and the device used is the NVIDIA GTX 480.

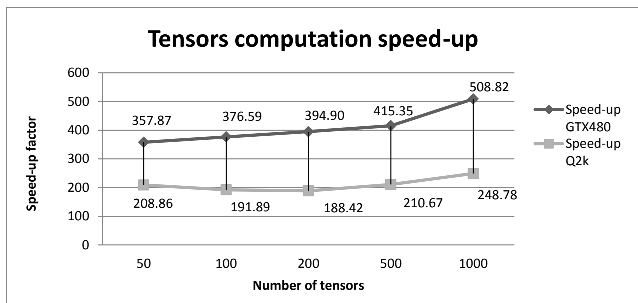

Fig. 10 Speed-up achieved compared to sequential CPU version for the computation of a different number of tensors.

times obtained using the CPU version are prohibitively for time-constrained applications.

4 Robot Vision: 3D object recognition

In this section, we show an application where the use of the accelerated semi-local surface feature extraction process allows to detect and recognize objects under cluttered conditions in real-time. The main goal of this application is the recognition of objects under real-time constraints in order to integrate the proposed algorithm in mobile robotics. Our method is designed to use only depth information because of the need for robots to work under bad or no illumination conditions.

To validate the proposal we tested the proposed feature in a similar application as it was done in the original work [17] where the semi-local surface feature is successfully used to recognize objects in cluttered scenes. For our experiments we have captured data from a Kinect sensor and tested the accelerated feature with some clut-

tered scenes. To do that, first a small library of models is constructed offline, storing all extracted tensors in an efficient way using a hash table. After, online object recognition is performed using cluttered scenes. Although the accelerated feature is tested using 3D data obtained from the Kinect sensor, this method is developed for managing 3D point sets collected by any kind of sensor and could be extended to other datasets.

We created a toy dataset to validate our proposal since the main goal of this work is to achieve real-time performance and integrate 3D data processing onto the GPU. Further analysis on recognition rates and feature parameters are already presented in the original work [17]. In addition, a deeper analysis on parametrization will be carried out in future works as this topic is out of the scope of the current work.

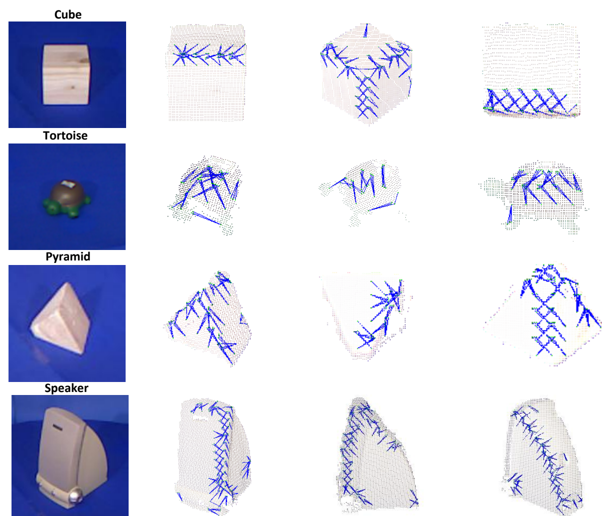

Fig. 11 Model library consisted of 4 real models. Each object consists of several partial views. For every partial view of a model, tensors are computed (blue lines) describing the model by extracting 3D surface patches (tensors).

4.1 Offline learning

To recognize objects using our real-time tensor extraction algorithm, first a model library is built extracting tensors from different views of free-form objects. Each partial view is represented with tensors and they are stored in an efficient way for being used after in an online recognition phase. In Figure 11 some partial views of the models used to build the library are presented. Moreover, tensors extracted for some of the views are showed in Figure 11. For each of these views the process explained in Section 2.1 is computed, obtaining as a result a set of tensors that describe each partial view.

Since multi-view correspondence using linear matching methods algorithms would be unaffordable, a hash map is introduced using the angle θ_d of the tensors as a key for the retrieval. The hash map is quantized in bins of θ_b degrees obtaining a good balance of tensors per bin and boosting the query performance. For this application the hash map is quantized into bins of 5 degrees.

In contrast to the original implementation of this kind of hash map presented in [17] and to integrate the matching process onto the GPU pipeline, the tensor library is stored in the GPU memory performing the tensors matching in parallel on the GPU and considerably accelerating its performance.

4.2 Online recognition

Once the model library is built and loaded in the GPU memory, the application is ready to start recognizing objects from a scene captured in real-time. Therefore, the input to our application is a point cloud of a scene. The point cloud is processed following the pipeline presented in Section 2.1. Once tensors from the scene are computed these are matched against the model library previously stored on the GPU memory. The matching process, as it was introduced in previous section, is performed in parallel onto the calculated entry using the angle as a key for the hash map. For the calculated entry are launched as many threads as tensors are stored in that bounded bin, computing in parallel a correlation coefficient originally presented in [17]. The correlation coefficient measures the similarity between the scene tensors and possible candidates stored in the model library. The correlation coefficient in the overlapped area between two tensors is calculated as follows:

$$C_c = \frac{n_q \sum_{i=1}^{n_q} p_i q_i - \sum_{i=1}^{n_q} p_i \sum_{i=1}^{n_q} q_i}{\sqrt{n_q \sum_{i=1}^{n_q} p_i{}^2 - \left(\sum_{i=1}^{n_q} p_i\right)^2} \sqrt{n_q \sum_{i=1}^{n_q} q_i{}^2 - \left(\sum_{i=1}^{n_q} q_i\right)^2}} \quad (6)$$

where p_i and $q_i (i = 1...n_q)$ are the respective elements of the model tensor T_m and scene tensor T_s in their region of overlap. Matchings whose $C_c < t_c$ are discarded ($t_c = 0.5$ based on results presented in [17]). The remain tensors are considered as possible correspondences.

Once all correlation coefficient has been calculated in parallel for a scene tensor T_s and considered as true possible correspondences, minimum correlation coefficient value is found and considered as the true correspondence. The reduction operation [6] to obtain the minimum value in parallel is also performed on the GPU pipeline using classical divide and conquer approach to find the minimum value.

4.3 Performance and results

In this section, some experiments related to the performance of the parallel matching performed on the GPU are presented, comparing performances obtained by the GPU and CPU versions.

Model library size (tensors)	2000	4000	16000	64000
Runtime CPU (ms)	215	398	1589	6414
Runtime GTX480 (ms)	75	140	534	2130
Speed-up GTX480	2.87x	2.84x	2.98x	3.01x

Table 3 Runtime comparison and speed-up obtained for matching process. As the size of the model library is increased the speed-up achieved is slightly larger. Runtimes are averaged over 50 runs.

In Table 3, it is shown how the matching process: correlation factor computing step and the reduction of the minimum value is computed faster on the GPU. The speed-up achieved is close to 3x but the most important is that the computing matching process on the GPU also avoid transferring data back to the CPU side after computing tensors on the GPU, step which obtains an important acceleration factor compared to the CPU as it is shown in Section 3.1. In this experiment, matching process is tested using different sizes of the model library, ranging from 4 to 64 objects. The number of objects was simulated copying the real model library comprised of 4 objects. It is considered that every object is described extracting tensors from 6 partial views obtaining an average of 1000 tensors per model.

Finally, in Figure 12 a scene computed using the proposed method is visually presented. Tensors are calculated randomly over the scene and matched tensors are labelled with the closest model in the library. Multiple labels are shown in Figure 12 as all tensors present in the scene are evaluated and matched against the library model. Voting strategies within clusters may be performed in order to further accelerate the object recognition process.

The number of tensors evaluated over the scene is 200 as experiments have demonstrated that evaluating over 200 tensors in most of cases achieves the recognition of all objects in the scene. A similar study was made in the original work [17]. Some wrong labelling appears in Figure 12 (Top) for the speaker as the partial view of the scene does not have enough geometric information to find a correspondence in the database. However, in the Scene 2 (Bottom) as the partial view contains more geometric information of the speaker, it is correctly recognized. For other objects as the tortoise, cube and pyramid, as similar views of the objects are present in the database and partial view of the scene has enough geometric information, the algorithm does correctly find tensors that match the model stored in the library. Total computation for the GPU took around 800ms using the NVIDIA GTX480, managing 3D object recognition problem in real-time and therefore enabling its integration in mobile robotics.

5 Conclusions

The highlights of this paper are as follows:

- Our primary concern is the integration of 3D data processing algorithms in complex computer vision systems. Experiments have demonstrated that GPGPU paradigm allows to considerably accelerate algorithms regard to CPU implementations and to run these in real-time.
- Within the 3D data algorithms used in the proposed pipeline, some progress have been made towards a faster and more robust point cloud triangulation algorithm, obtaining a GPU implementation that runs at 30 fps.
- Advantageous results are obtained in the use of the GPU to accelerate the computation of a 3D descriptor based on the calculation of 3D semi-local surface patches of partial views, thus allowing descriptor computation at several points of a scene in real-time.
- Matching process have also been accelerated onto the GPU, taking advantage of the GPU pipeline and achieving a speed-up factor of 3x regard the CPU implementation.
- We have implemented a prototype of the proposed pipeline and it has been tested with a real application obtaining satisfactory results in terms of accuracy and performance. We show that implemented prototype took around 800 ms with a GPU implementation and performing 3D object recognition of the entire scene.

Further work will include adding other processing steps to the GPU pipeline: hypothesis verification using ICP techniques on the GPU and using multi-GPU computation to improve performance and to manage computation of tensors and their matching on different devices.

Acknowledgements This work was partially funded by the Valencian Government BEFPI/2012/056, and by the European Network of Excellence on High Performance and Embedded Architecture and Compilation (HiPEAC). Experiments were made possible with a generous donation of hardware from NVDIA.

References

1. Manuel Blum, Jost Tobias Springenberg, Jan Wülfing, and Martin Riedmiller. A learned feature descriptor for object recognition in rgb-d data. In *ICRA*, pages 1298–1303, 2012.
2. Derek Chan, Hylke Buisman, Christian Theobalt, and Sebastian Thrun. A Noise-Aware Filter for Real-Time Depth Upsampling. In *Workshop on Multi-*

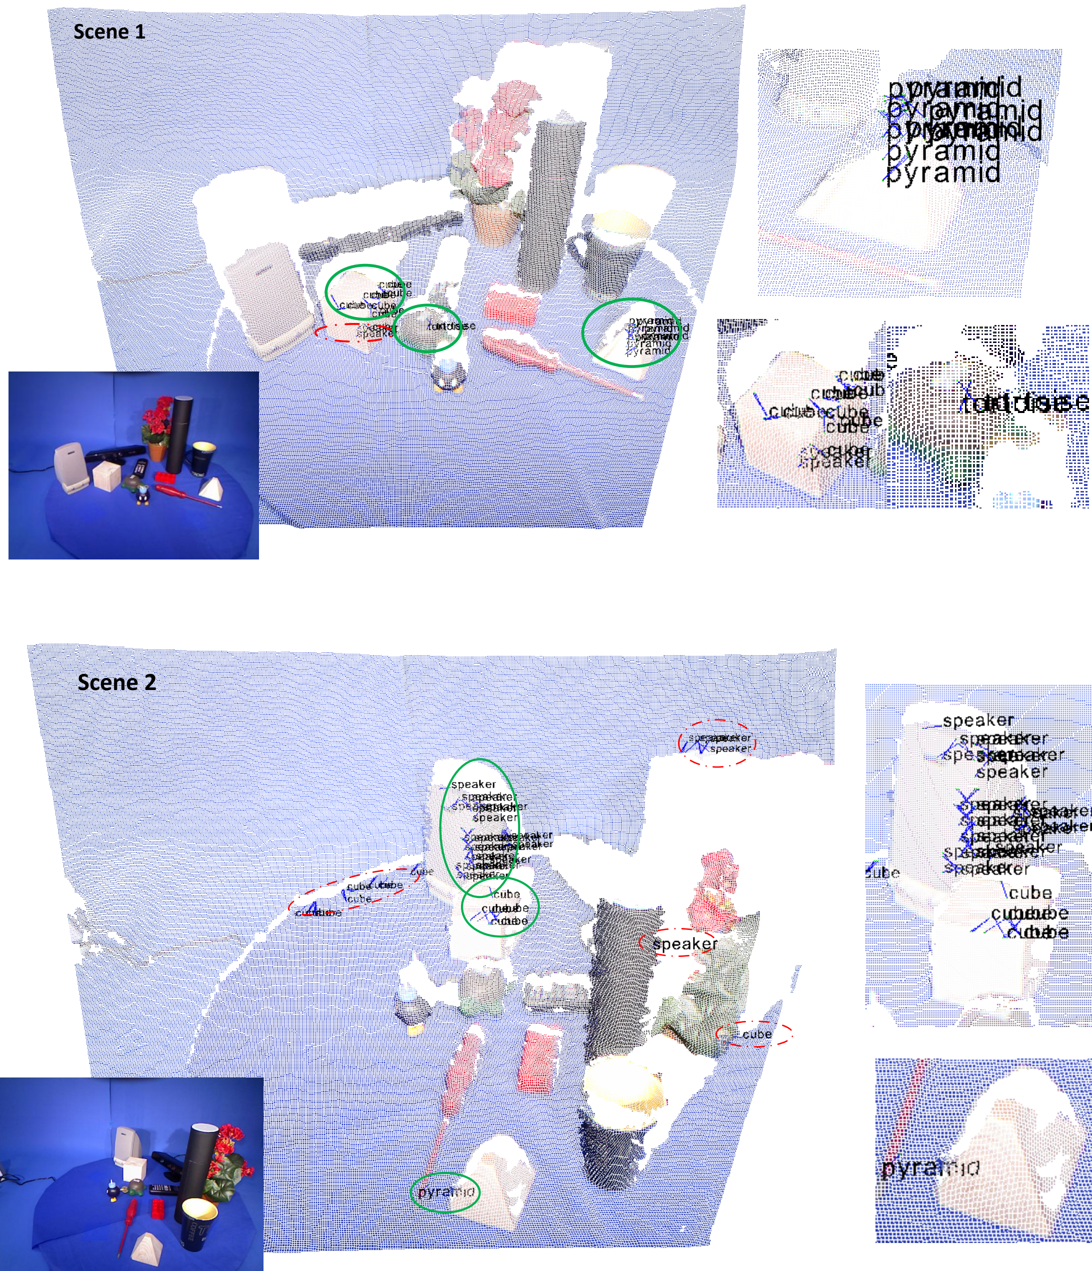

Fig. 12 3D object recognition in cluttered scenes. Different partial views of two scenes are showed. Multiple labels are shown as all computed tensors are evaluated and matched against the library model.

camera and Multi-modal Sensor Fusion Algorithms and Applications - M2SFA2 2008, Marseille, France, 2008. Andrea Cavallaro and Hamid Aghajan.

3. Hui Chen and Bir Bhanu. 3d free-form object recognition in range images using local surface patches. *Pattern Recogn. Lett.*, 28(10):1252–1262, July 2007.
4. James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer graphics: principles and practice (2nd ed.).* Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1990.
5. Wesley Griffin, Yu Wang, David Berrios, and Marc Olano. Real-time gpu surface curvature estimation on deforming meshes and volumetric data sets. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1603–1613, October 2012.

6. Mark Harris. Optimizing parallel reduction in cuda. *NVIDIA Dev. Technology*, 2008.
7. Gnter Hetzel, Bastian Leibe, Paul Levi, and Bernt Schiele. 3d object recognition from range images using local feature histograms. In *CVPR (2)*, pages 394–399. IEEE Computer Society, 2001.
8. M. Himmelsbach, T. Luettel, and H.-J. Wuensche. Real-time object classification in 3d point clouds using point feature histograms. In *Proceedings of the 2009 IEEE/RSJ international conference on Intelligent robots and systems*, IROS'09, pages 994–1000, Piscataway, NJ, USA, 2009. IEEE Press.
9. Dirk Holz and Sven Behnke. Fast range image segmentation and smoothing using approximate surface reconstruction and region growing. In *Proceedings of the 12th International Conference on Intelligent Autonomous Systems (IAS)*, Jeju Island, Korea, June 2012.
10. Shahram Izadi, Richard A. Newcombe, David Kim, Otmar Hilliges, David Molyneaux, Steve Hodges, Pushmeet Kohli, Jamie Shotton, Andrew J. Davison, and Andrew W. Fitzgibbon. Kinectfusion: real-time dynamic 3d surface reconstruction and interaction. In *SIGGRAPH Talks*, pages 23:1–23:1, 2011.
11. Andrew E. Johnson and Martial Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):433–449, 1999.
12. Andrew E. Johnson and Martial Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(5):433–449, May 1999.
13. Kourosh Khoshelham and Sander Oude Elberink. Accuracy and resolution of kinect depth data for indoor mapping applications. *Sensors*, 12(2):1437–1454, 2012.
14. Jan Knopp, Mukta Prasad, Geert Willems, Radu Timofte, and Luc Van Gool. Hough transform and 3d surf for robust three dimensional classification. In *Proceedings of the 11th European conference on Computer vision: Part VI*, ECCV'10, pages 589–602, Berlin, Heidelberg, 2010. Springer-Verlag.
15. Wonwoo Lee, Nohyoung Park, and W Woo. Depth-assisted real-time 3D object detection for augmented reality. *ICAT'11*, 2:126–132, 2011.
16. A. S. Mian, M. Bennamoun, and R. A. Owens. A novel representation and feature matching algorithm for automatic pairwise registration of range images. *Int. J. Comput. Vision*, 66(1):19–40, January 2006.
17. Ajmal S. Mian, Mohammed Bennamoun, and Robyn Owens. Three-dimensional model-based object recognition and segmentation in cluttered scenes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(10):1584–1601, October 2006.
18. NVIDIA. *NVIDIA CUDA Programming Guide 5.0.* 2012.
19. NVIDIA. Nvidia visual profiler, 2012.
20. Victor Prisacariu and Ian Reid. fasthog - a real-time gpu implementation of hog. Technical Report 2310/09, Department of Engineering Science, Oxford University, 2009.
21. S. Ruiz-Correa, L.G. Shapiro, and M. Melia. A new signature-based method for efficient 3-d object recognition. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages I–769–I–776 vol.1, 2001.
22. Radu Bogdan Rusu and Steve Cousins. 3D is here: Point Cloud Library (PCL). In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011.
23. F. Stein and G. Medioni. Structural indexing: efficient 3-d object recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 14(2):125–145, 1992.
24. C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proceedings of the Sixth International Conference on Computer Vision*, ICCV '98, pages 839–846, Washington, DC, USA, 1998. IEEE Computer Society.
25. Lingyuan Wang, Miaoqing Huang, and Tarek A. El-Ghazawi. Exploiting concurrent kernel execution on graphic processing units. In Waleed W. Smari and John P. McIntire, editors, *HPCS*, pages 24–32. IEEE, 2011.
26. Jakob Wasza, Sebastian Bauer, and Joachim Hornegger. Real-time preprocessing for dense 3-d range imaging on the gpu: Defect interpolation, bilateral temporal averaging and guided filtering. In *ICCV Workshops*, pages 1221–1227, 2011.
27. Zhengyou Zhang. Microsoft kinect sensor and its effect. *MultiMedia, IEEE*, 19(2):4 –10, feb. 2012.